ICWall: a Calibrated Stereo Tiled Display from Commodity Components

T. v.d. Schaaf*
Faculty of Sciences, Vrije Universiteit Amsterdam
M. Koutek†
EWI Faculty, Delft University of Technology
D.M. Germans‡
Faculty of Sciences, Vrije Universiteit Amsterdam
H.E. Bal§
Faculty of Sciences, Vrije Universiteit Amsterdam

Abstract

Recent developments in the fields of parallel rendering and high resolution tiled displays have made it possible to apply these technologies to build large and scalable stereo displays for use in Virtual Reality applications. This paper presents the implementation of a high-resolution stereo tiled display (2x8 tiles), built from low-cost commodity components. Among the problems that arise when building such a system, the most challenging is multi-projector alignment and calibration. We describe our method of aligning the left and right-eye projectors using an automatic approach rather than the time consuming alignment of the projectors by hand. We compare two implementations of this method: a single-pass and a two-pass rendering method to adjust projector images for alignment of the tiles. We demonstrate such a stereo-calibrated tiled display in action and we present recommendations for using this system to overcome remaining issues.

CR Categories: I 4.9 [Image Processing and Computer Vision]: Applications; I 3.2 [Graphics Systems]: Distributed/network Graphics; I 3.7 [Three-Dimensional Graphics and Realism]: Virtual Reality

Keywords: Tiled displays, VR system architecture, stereo graphics, parallel rendering

1 Introduction

In recent years we observe increasing popularity of tiled displays of various dimensions and resolutions. Tiled displays provide a scalable and high resolution display by tiling together projector outputs on a large screen. Often, due to the excellent price-performance ratio of desktop PCs with modern graphics cards, [Funkhouser and Li 2000] the projectors are driven by rendering clusters interconnected via a fast network (FastEthernet [Boden et al. 1995], Myrinet or Infiniband).

Existing high-end Virtual Reality systems such as the CAVE and ImmersaDesk [Cruz-Neira et al. 1993; Pape et al. 1999] offer an excellent way to perform immersive visualization of scientific data. The 3D interaction devices available for these systems offer a natural way of interacting with the data. Unfortunately, these system use a single projector per screen, which results in clearly visible pixels, even when using high resolution projectors. A second problem is the scalability in terms of the amount of spectators. Only a limited amount of spectators can be part of the Virtual Reality experience. These issues can be solved by using Tiled Displays with the capability of stereoscopic rendering. Because Tiled Displays have a scalable amount of projectors, very high resolution images can be generated. Also, they offer the possibility to create bigger screens, suitable for rooms that can hold a large audience. Currently, very few stereo Tiled Displays have been realized.

In order to realize a stereo tiled display, two main issues need to be addressed. First, geometric alignment and photometric calibration of left and right (eye) tiles have to be separately executed. Then both left and right projected images have to be registered onto the screen. Second, the graphics pipeline has to deliver stereo views on the virtual scene providing proper depth vision. Here, the frustum, inter-ocular distance and the fusion point of the viewer are important. Unlike existing VR systems (like the CAVE [Cruz-Neira et al. 1993]), a stereo tiled display has multiple viewers that are not tracked. Because of this, the image can not be calculated correctly for every viewer and an average solution has to be considered.

In this paper we will present our work on realizing the ICWall stereo tiled display at the Vrije Universiteit of Amsterdam. The first demonstrations of our stereo prototype took place in January 2003. This display is set up in a class room for use in the university education programme as well as by scientific researchers with VR applications. The ICWall consists of 16 projectors, controlled by 9 PCs to generate a stereo image of roughly 2x4096x1524 pixels.

The contributions of this paper are:

- We discuss hardware setup considerations in order to realize a stereo tiled display.
- We present a sub-pixel-accurate and deterministic method to geometrically calibrate the display for both eyes.
- We describe how sub-frusta for each projector are obtained in our implementation: either by using direct homographic matrix transformation or through the two-pass rendering.
- We present several example applications for the ICWall.

The rest of this paper is structured as follows. Section 2 discusses a related work in the area of stereo display and Tiled Display Calibration. Section 3 explains the hardware setup of our stereo Tiled Display. Section 4 describes how geometric, photometric and viewpoint calibration are performed to achieve a convincing stereo image. Section 5 discusses several applications using the display. Finally, we conclude with Section 6.

*e-mail: Tom@few.vu.nl
†e-mail: M.Koutek@ewi.tudelft.nl
‡e-mail: Desmond@germansmedia.nl
§e-mail: Bal@cs.vu.nl

2 Related Work

2.1 Stereo Displays

Our focus in this paper is on stereo graphics within the realm of Virtual Reality. Until recently, this field was dominated by high end systems such as the CAVE and ImmersaDesk [Cruz-Neira et al. 1993; Pape et al. 1999]. Recently, commodity component setups based on desktop PCs support stereo at lower costs [Leigh et al. 2001; Pape et al. 2002].

The research group at Fraunhofer-IGD developed the HEyeWall, a stereo tiled display, based on Infitec stereo separation [Kresse et al. 2003]. From their work it becomes clear that maintaining a large stereo tiled displays with manually aligned projectors is very time consuming. Our work is complementary to [Kresse et al. 2003], as we present a working stereo tiled display based on back-projection and linear polarization using automatic calibration in software. We will present basic guidelines to make acceptable stereo quality without the need to manually adjust projectors.

Very recently, growing market interest has attracted commercial companies to the field of stereo tiled displays. Commercial tiled display setups usually offer high quality seamless stereo images by using modern high-end projectors, special optics, image warping units, semi-automated calibration schemes and high quality screens. One example of a commercial display is the I-wall by Barco, an Infitec-based stereo tiled display (2x2 tiles). The main drawback of these systems is their high price. Our approach was to realize a stereo tiled display using cost-effective off-the-shelf components.

2.2 Tiled Display Calibration

Important issues for tiled display realization are the geometric and photometric calibration of each projector output in order to produce a large and seamless image with equal luminance and chrominance over the entire surface. Recent developments in geometric calibration have shown that a single-pass deterministic homographic calibration method is most accurate [Raskar 2000]. This method puts very little constraints to the physical positioning and orientation of each projector. More specific, as long as each projector is aimed at the same surface and overlaps with the neighboring projectors, a seamless display can be constructed. In more recent work [Raskar et al. 2003], Raskar et al. describe techniques that enable clusters of self-configuring projectors to automatically form a seamless image image.

Homography-based calibration in the case of a flat display surfaces was introduced by M. Steele[Steele et al. 2002], performing continuous monitoring of projected images on a flat surface. Work by K. Li at Princeton explored how an inexpensive uncalibrated camera observing a flat wall can be used to bring a display wall with multiple projectors in geometric alignment[Chen et al. 2000].

Roughly, the method works as follows. After projecting a known stimulus on each tile, a computer connected to a CCD camera determines the orientation of each projector by fitting a homographic transformation to stimulus features that are recognized from a test image captured by the camera. From this transformation, the configuration of the entire screen can be inferred. Our implementation of this method results in sub-pixel accurate geometric calibration of the screen.

Notable in the field of photometric calibration is work by Majumder [Majumder and Stevens 2002]. Here they give a thorough analysis of various issues in photometric calibration and give a phenomenological method to equalize luminance (the most important component) of the picture, using what they call LAMs (Luminance Attenuation Maps). This mapping is based on a per-pixel analysis and does not attempt to functionally model the phenomena that cause the luminance variations. Continuing on this research is the PixelFlex2 [Raij et al. 2003] project, a display system that is capable of automatically calibrating casually aligned projectors.

3 Setup

This section will explain the hardware setup of our stereo Tiled Display: the ICWall. The display was setup as a large screen in a lecture room of the faculty, to allow a large number of spectators (more than 60). We choose to use passive stereo on a back-projection screen. Our argument is that active stereo with LCD shutter glasses, with one projector for each tile, requires more expensive projectors with a high refresh rate. The cheaper alternative for active stereo, using two cheaper projectors for each tile in combination with LCD or mechanical shutters, would of course require a tight synchronization of all the shutters, which is a technical problem in itself.

Three option for passive stereo were considered: linear polarization, circular polarization, and Infitec. At that time, the most cost effective choice was to use linear polarization. Using circular polarization on a back-projected system leads to very bad stereo ghosting effects and unpleasant light intensity variations over single tiles. The Infitec color correction that is used in current Infitec displays was not yet available during the design and construction phases of the display. Without this correction, disturbing color differences between the left and right channels results in less convincing stereo images. Even with color correction, color discrepancies between left and right-eye images remain a problem. Unfortunately, linear polarization is not without problems either. When using back projection, the screen must maintain the polarization angle of the light, while also emitting light as diffuse as possible. Screens that combine these properties are very expensive.

We selected an affordable back-projection screen from Stewart-Film for its polarization properties. Light passes through and is distributed with the polarization angle left unchanged. Unfortunately, this comes at the expense of the Lambertian characteristics of the screen. Light is not scattered equally in all directions, which complicates luminance calibration.

Figure 1: The frame, projectors and driving PCs at the ICWall site.

To avoid the internal image polarization of commodity LCD-projectors, we use DLP-projectors (Philips UGO), two for each tile. Each set of projectors is connected to a PC with a dualhead NVidia GeForce 4 graphics card. The PCs as well as one extra host machine are interconnected via a Myrinet network, allowing high throughput and low network latency. That is needed while we are using AURAII, a distributed scene graph API, and the individual tiles are rendered in parallel on the rendering slaves. Certainly, the scene graph changes from the master node has to be updated on the slaves as fast as possible. Figure 1 shows our projection setup.

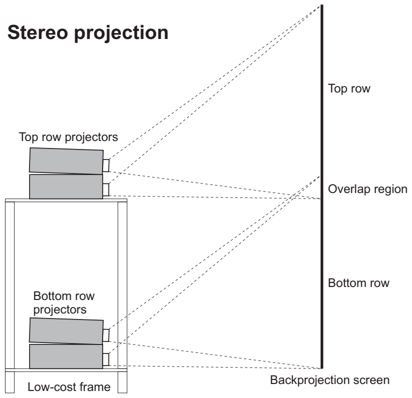

Figure 2: Each projector pair is stacked with the top projector slightly tilted.

We choose a projector configuration that allows easy and low-cost maintenance of the projectors. For each tile, both projectors are mounted so that they can shift and rotate with approximately 1 degree of freedom. The top projector is slightly tilted to approximately project onto the same area as the bottom one (Figure 2). In our setup, projectors need not be aligned very precisely: making sure that all projection areas have an overlap with the adjacent projection areas is sufficient. This overlap area is required to allow the geometric calibration process to generate a smooth transition from one tile to the other.

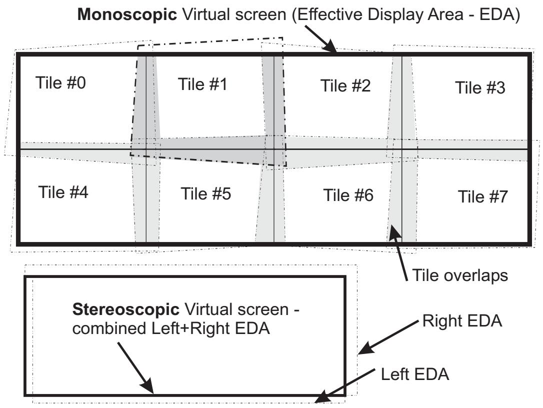

Figure 3: Tile overlaps and construction of a virtual screen.

Figure 3 shows an example of a rough manual alignment. The different projection areas of 8 overlapping projectors (for one eye) will be combined into a single virtual screen by the calibration software. In stereo mode, the left and right-eye virtual screens are again combined into a single virtual screen.

4 Calibration

This section describes the geometric, photometric and viewpoint calibration applied to the display. Geometric calibration achieves a correspondence between coordinates of each individual projector and coordinates on the screen, photometric calibration achieves a correspondence between intended image luminance and chrominance (sent to the projector), and actual luminance and chrominance. Viewpoint calibration achieves a correspondence between the viewer in front of the screen and the utilized 3D frustum in the virtual world that is being displayed on the screen. We will look in detail at all three calibration processes and show how they become important in producing a convincing stereo image.

4.1 Geometric Calibration

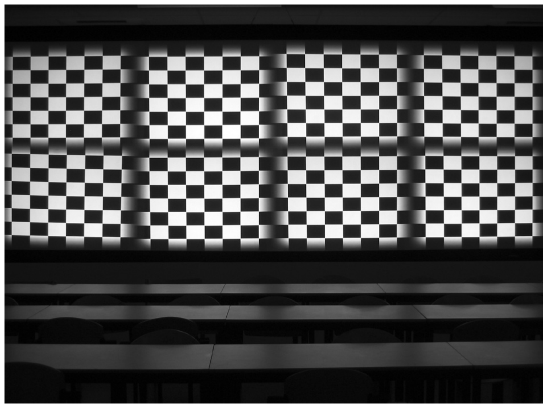

Figure 4: The resulting capture of the CCD camera for one eye. From these patterns, the homographic perspective plane transformation of each projector is retrieved.

To characterize the position and orientation of the individual projectors, we detect features of a synthetic stimulus generated by each projector. In our case we use a checkerboard pattern of known size (Figure 4). The checkerboard pattern on the display is captured using a digital camera. In order to compensate for the lens distortion of this camera, the image of the checkerboard pattern is adjusted, which is a simple procedure once the distortion of the lens has been measured.

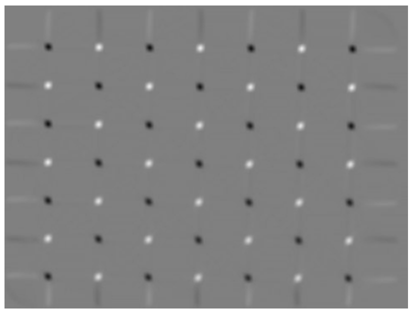

Figure 5: Typical convolution result of matched-filter convolution on 8x8 checkerboard stimulus. The white and black dots are extreme responses, indicating full match and full anti-match (a checkerboard crossing where white and black are swapped).

The features that are being looked for, are the crossings in the checkerboard. These crossings can be detected by performing a convolution of the captured image with a *matched filter* that has an extreme response where the filter matches the crossing [Vuylsteke and Oosterlinck 1990]. Figure 5 shows a graphical representation of the response for a single projector. The white and black dots indicate full match and full anti-match (a checkerboard crossing where white and black are swapped). Additionally, [Vos 1998] has shown that one can obtain sub-pixel accurate localization of the checkerboard crossings by calculating the center-of-mass of the responses around each extremum.

Via this technique, we arrive at a list of landmarks for each tile/eye combination. We know that the list should represent the stimulus under a homographic transformation:

$$\begin{pmatrix} x_{out} \\ y_{out} \\ w_{out} \end{pmatrix} = \begin{pmatrix} a & b & c \\ d & e & f \\ g & h & 1 \end{pmatrix} \star \begin{pmatrix} x_{in} \\ y_{in} \\ 1 \end{pmatrix}$$

$$\text{or:} \quad x_{out} = \frac{ax_{in}+by_{in}+c}{gx_{in}+hy_{in}+1}$$

$$y_{out} = \frac{dx_{in}+ey_{in}+f}{gx_{in}+hy_{in}+1}$$

Here, (x_{in}, y_{in}) are the input coordinates (the known crossing locations in normalized projector space) and (x_{out}, y_{out}) are the output coordinates (the recognized crossings in CCD camera space). a,b,c,d,e,f,g and h are parameters of the transformation. We can now fit this transformation to the list of landmarks using a non-linear least squares fitting algorithm.

To achieve more accuracy, the transformation can be enhanced by simultaneous modeling of the projector lens distortions. Here, the transformation is augmented as follows:

$$r(x_{out}, y_{out}) = \sqrt{(x_{out} - x_c)^2 + (y_{out} - y_c)^2}$$

$$x_{ldm}(x_{out}, y_{out}) = x_{out} + \alpha_2 r(x_{out}, y_{out})^2$$

$$y_{ldm}(x_{out}, y_{out}) = y_{out} + \beta_2 r(x_{out}, y_{out})^2$$

Where r is the distance of the point (x_{out}, y_{out}) with a parametric center (x_c, y_c), (x_{ldm}, y_{ldm}) is the new output point and α_2 and β_2 are distortion parameters. To obtain the distortion parameters, we take a generic polynomial in r(xout,yout) with parameters alpha0, alpha1, alpha2, etc. Note that in order for the parameter estimation to work, we need independent parameters. Therefore, we need to leave out the first two parameters. The constant term of this polynomial denotes a translation, which is already coded in the homographic transformation. The first-order term denotes a scaling, which is also coded in the homographic transformation. Because third- and higher-order terms have negligible influence, they are dismissed as well. This leaves only the second-order term, both for x and for y. Using this in our implementation, we arrive at sub-millimeter accuracy on the display surface.

With a transformation for each tile/eye combination, we can now setup quadrilaterals in screen space that represent the projection area of each tile/eye combination. From this information, we configure the contents of the display. First the largest rectangle is found that fits on the screen for both eyes, this rectangle is called the *virtual screen*. When this rectangle is found, it is intersected with each quadrilateral. the resulting polygons represent the *projection area polygon* of each tile/eye combination. Intersections between adjacent projection area polygons give the *overlap polygons*.

The homographic transformations can be applied to the projector images in two ways, which will be described in Section 4.4.

4.2 Photometric Calibration

Photometric calibration comes down to equalizing the luminance and the chrominance response of each tile. For effective stereo output, it is important that images for both eyes are similar in luminance and chrominance. [Majumder and Stevens 2002] shows that chrominance-differences are much less noticeable, so we concentrate on equalizing luminance only. Commodity DLP projectors add white to the image in order to boost the contrast, which complicates modeling of the chrominance considerably.

We measure luminance with the same CCD camera as used for geometric calibration. What needs to be said here is that the mapping between resulting camera output values and actual luminance is unknown. In order to calibrate this, one could follow [Debevec and Malik 1997]. However, we choose to model the entire process from projector input to camera output, including all calibration issues related to projectors or cameras.

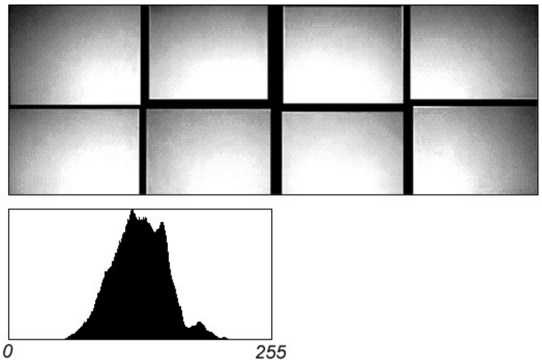

Figure 6: Luminance response due to non-Lambertian characteristics of the back-projection screen. The hot spots indicate a direct line of sight with the heart of the projectors. Note that for the histogram, the overlap areas are not taken into account.

A large contribution to luminance differences is the fact that the back-projection screen is not a perfect Lambertian source. Light that passes through and is diffused along a large angle is much less bright than light that passes through and is not bent at all (Figure 6).

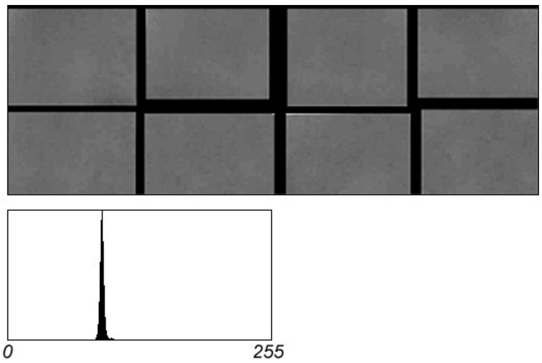

Figure 7: Luminance response after applying the modeled correction. Note that the spread in the accompanying histogram is now minimal. Note again that for the histogram, the black overlap areas are not taken into account.

For each tile/eye-combination, the luminance difference can be modeled by a second-order polynomial in the distance to the center of a hot spot. To estimate the parameters, all projectors are set to output a specific level of gray on all pixels. A capture is made of the outputs of the projectors and the model is fitted to the projector response using a non-linear least squares fitting method. Figure 7 shows the luminance response for the display after this correction.

Because the unwanted luminance effect is viewer-dependent, a perfect compensation does not exist and we have to choose a preferred location for the viewer. To allow scalability in the audience, one viewer for which the image is compensated is not enough, and so we would have to blend luminance compensations of a whole range of viewers.

To utilize the luminance correction and fuse the overlap areas smoothly, an *alpha mask* is constructed for each projector. This mask is multiplied with the projector output to locally attenuate the signal (Figure 8). For this, we transform the projection area and overlap polygons back to projector coordinates, using the inverse of the geometric transformation. The projection area polygon functions as a main mask. In order to have two adjacent tiles blend towards each other (Figure 9), the overlap polygons are used to produce gradual fall-offs towards the sides of the projection. Finally, the mask is attenuated with the luminance model.

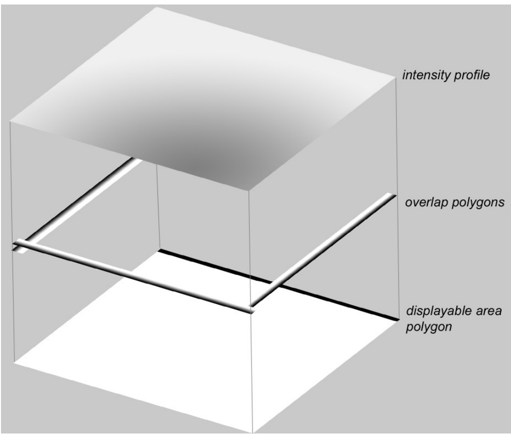

Figure 8: The alpha mask is built out of the projection area polygon, the various overlap polygons that produce a gradient to the sides and a luminance profile.

It is important to realize that this method only allows us to model linear effects in the luminance transfer from projector to camera. This is sufficient for a convincing stereo image, despite higher order effects that are still noticeable in the output (gamma effects and a non-linear camera response).

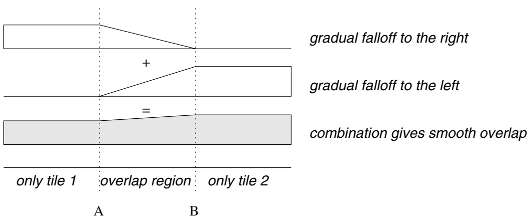

Figure 9: The gradual falloff of the overlap polygons for two adjacent tiles add together to form a smooth overlap.

The effect of our geometric and photometric alignment procedure is shown on Figure 10. The left picture shows a closeup of an overlap area between two projectors. The right picture shows the same area using our alignment procedure.

4.3 Viewpoint Calibration

Correct separation of left and right eye images is not enough to create convincing stereo. Each of the two images must be rendered using the proper perspective for each eye. To obtain an optimal depth image, we must take into account several factors (Figure 11): the size of the screen, the the positions of left and right eye with respect to the screen, and the fusion point of the eyes (the point the user is looking at).

The 3D viewing frustum governs how 3D world coordinates are transformed to the 2D surface of the screen. Like in traditional

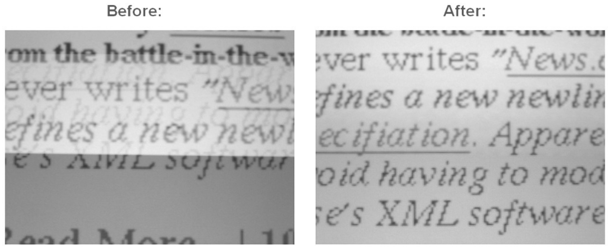

Figure 10: Successful alignment and correct edge blending

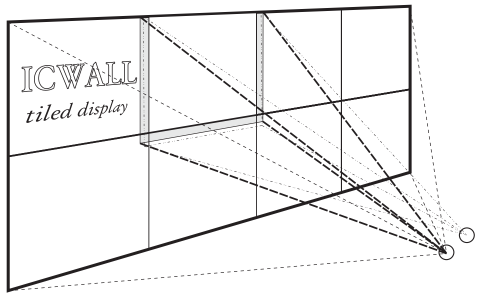

Figure 11: Stereo tiled projection.

VR systems, we require off-axis frustum projection for convincing stereo images. Because our audience is not tracked, we require a static viewer that represents an average of the audience. Figure 12 shows left and right eye frusta when the viewer is placed 5 meters in front of a 6 meter wide screen. In case of a display with multiple projectors, we need to calculate a sub-frustum for each projector incorporating any possible overlap with other projectors, using the information acquired by our geometric alignment procedure as described in 4.1. There are two methods for doing this, which are described in 4.4.

The inter-ocular distance is the distance between the left and right eye. For audiences to view the screen for prolonged periods of time, it is important that this distance is calibrated on several test-viewers. If the distance is chosen to be too wide or too narrow, viewers will feel uncomfortable when viewing the screen.

The fusion point is the exact location that both eyes are looking at. Choosing the fusion point too close to the viewer will give an uncomfortable cross-eye impression, whereas a far fusion point tends to leave the viewer with no reference. For any (traditional) display, viewers will naturally place the fusion point at the surface of the display. We choose the same for our display and place the fusion point at the center of the display surface.

4.4 Implementation

We have implemented the described Geometric, photometric and viewpoint calibration in the parallel rendering software (Aura/VIRPI [Germans et al. 2001]) that drives the ICWall tiled display.

There are two methods for obtaining the proper sub-frustum for a given projector, given the homography matrix acquired by the geometric calibration, and the viewing frustum acquired during the viewpoint calibration.

The most straightforward and fastest method to calculate the sub-frustum is by multiplying the projection matrix of the application with the 3D extension of the 2D homography matrix.

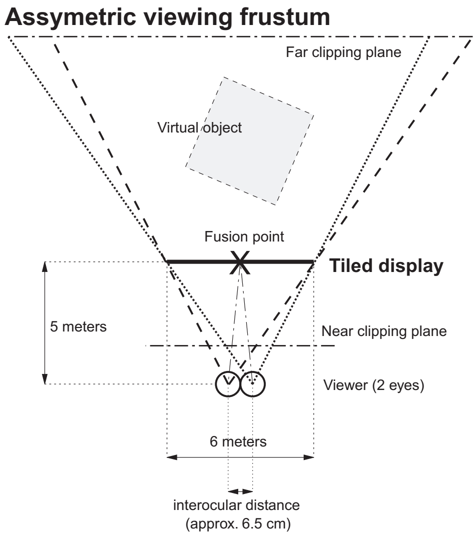

Figure 12: The viewing frustum and the physical properties of the display and the viewer need to coincide.

$$M_{3D-homography} = \begin{pmatrix} a & b & 0 & c \\ d & e & 0 & f \\ 0 & 0 & 1 & 0 \\ g & h & 0 & 1 \end{pmatrix}$$

Unfortunately, this homography matrix is not a clean affine transformation matrix and has undesirable side-effects. For example, it moves near and far planes of the frustum, which requires compensation. Even when applying several adjustments to the homography matrix and near and far values, these planes will appear warped and not precisely at their specified locations. On top of that, some implementations of OpenGL can no longer draw display lists, as they are culled by optimizations based on near and far values.

The main problem with this approach is that we try to apply a 2D homographic transformation matrix to a 3D projection. A more correct way of solving this problem consists of two steps: first we draw the 3D scene (Figure 13) and then apply the transformation to the resulting 2D image (Figure 14).

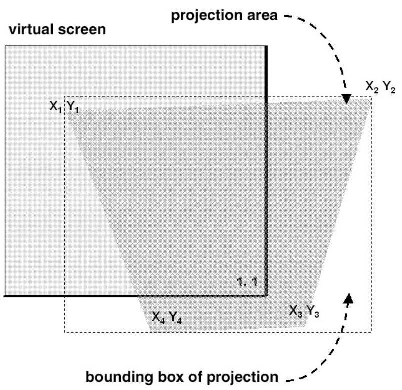

Figure 13: Two-pass geometric calibration, first part.

For the first step, we need to find the coordinates of the projection area of the projector in virtual screen coordinates. Because our homographic matrix is defined in the virtual screen coordinate system, acquiring the coordinates from this matrix is trivial. Because graphics hardware supports only rectangular textures, we draw that part of the scene that exists within the bounding box of the projection area into texture memory.

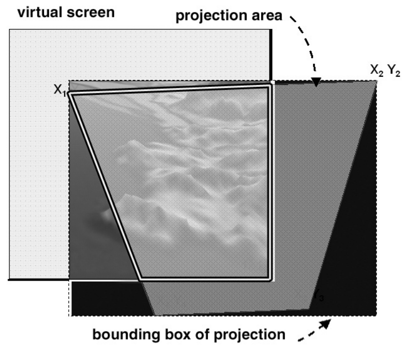

Figure 14: Two-pass geometric calibration, second part.

For the second step, all we need to do is draw the texture on a quadrilateral filling the entire framebuffer and adjust texture coordinates in such a way that only that part of the texture that was within the projection area is visible. Constructing the matrix that calculates these coordinates requires some tricks to get the right values for the homogeneous texture coordinates and is beyond the scope of this document.

The first implementation requires only a very small, cost per frame independent of the scene being rendered (one matrix multiplication). The second implementation has a larger cost (but also independent of the scene), because of the two-pass rendering. Although the second implementation is less efficient, the resulting image is guaranteed to be correct. Fortunately, rendering directly into texture memory makes sure that the image remains on the graphics hardware and keeps overhead relatively small. Measured impact of this method is only significant on applications with framerates that are well beyond the refresh rate of current display devices.

5 Results and Applications

To evaluate the validity of our calibration methods and our stereo setup, we implemented several tests that could be used to perform two types of informal measurements. We present here an informal geometric linearity test and a viewpoint calibration test.

5.1 Stereo Projection Tests on the Tiled Display

First, we measured the accuracy of our alignment by displaying a set of horizontal and vertical lines that, based on calibration information, should be placed with intervals of exactly 10 cm. Physical measurement of the real distances between lines on the screen (Figure 15) gives an indication of the error within the calibration system. If we, for example, do not compensate for camera lens distortions, we find big errors on the edges of the screen, where the lens distortion is worst. When using our full calibration procedure we find only very small errors, varying from about 1*mm* for the central areas of the ICWall to 5*mm* closer to the borders.

Further, from this experiment it became clear that our method delivered very good geometric continuity over the tiles, see Figures 15-19. The quality of the alignment procedure (8 tiles and 16 projectors involved)is show in Figures 10 and 16. For the measurement purposes during the stereo projections we set the interocular distance (IOD) to zero.

In the center of the screen, the middle of the bright white cross is drawn by all 8 projectors overlapping that area. Practically, there is no visible mis-alignment, which is a very good result.

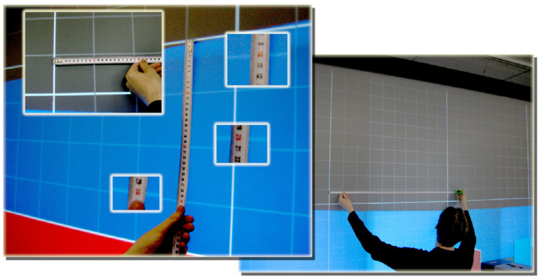

Figure 15: ICWall grid linearity measurements.

To test the viewport calibration and depth perception, we performed an augmented reality test where the virtual scene would extend the length of two real-world boxes. The boxes where placed on known positions on a table in front of the screen, but not physically touching the screen. The stereo camera viewpoints were calibrated for a known viewing position in the room. We used meters as the metrics for the virtual scene.

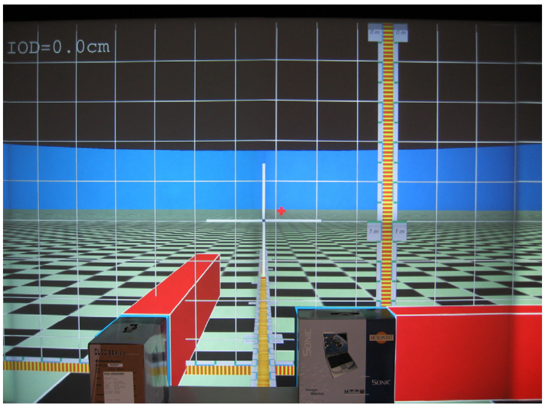

Figure 16: ICWall stereo viewpoint calibration and visual scene registration test: front view

The stereo images where captured using a camera positioned in the known viewing position (Figure 16, 17). These images show that our stereo tiled display together with AURA graphics library, can provide calibrated stereo views on the virtual scenes from various view angles. The geometric continuity over the tiles is preserved not only in the screen plane with the grid, but also in the depth, see the checkerboard pattern in the ground plane. The registration test, where we tried to match physical with virtual object, shows that the real and virtual perspective nicely correspond. There is a little mis-registration, which was probably caused by the placement of the camera-on-tripod. Because IC Wall display is not supported by the head tracking, is it for us a sufficient result that our we can convincingly present depth information in the stereo mode. During immersive applications and presentations we a set of predefined viewpoints in the classroom (i.e. sitting in the first row, or in the middle of the class, or standing 2m from the screen).

5.2 Stereo Applications

In this section we will show a few examples of applications where the stereo tiled display has impact. One of the applications is a

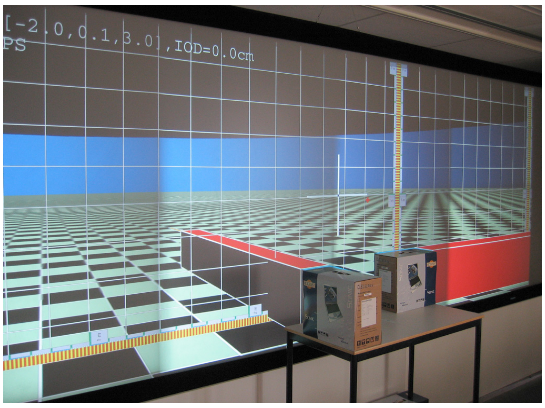

Figure 17: ICWall stereo viewpoint calibration and visual scene registration test: left side view

molecular visualization. We use it to visualize large molecular structures from molecular dynamics simulations.

Figure 18 shows this application displaying bacteriophage Alpha3 (consisting of more than 300.000 atoms), rendered using the rendering cluster at 3 frames per second in stereo. Each atom is represented by a (60 vertex) lit sphere without any level-of-detail optimizations. For rendering we are using AuraII scene graph library.

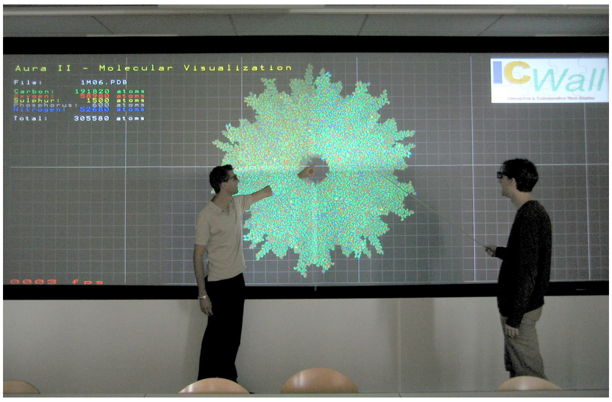

Figure 18: Molecular visualization on the ICWall.

Figure 19 shows a stereo application displaying isosurface extraction from a (standard benchmark) medical data set. The isosurface is generated by the Visualization Toolkit (VTK), which is coupled to the AuraII. Note that Figures 19 and 18 have been partially illustrated to correct for the missing visual depth information in the ICWall photos.

6 Conclusions and Discussion

This paper demonstrates how a high-resolution stereo tiled display can be realized using low-cost commodity components. Our setup uses the cost effective approach for image separation: passive stereo using linear polarization. We achieve with this a sufficient stereo separation, but the users should not tilt to much their heads. The amount of the crosstalk can be diminished by avoiding high-contrast between edges in the scene and the background color. Currently Infitec passive stereo is a good alternative, but both technologies have their drawbacks and advantages. Despite the differences in color and intensity of each projector and the non-Lambertian characteristics of our screen, we show that it is possible

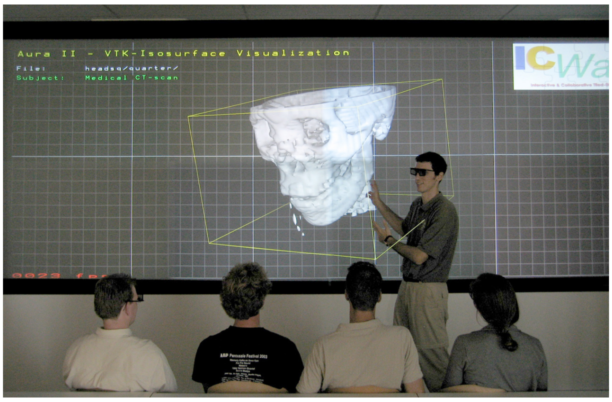

Figure 19: VTK isosurface visualization on the ICWall.

to generate a convincing stereo image. The best view on the stereo projection is from the central axis in the middle of the room. At stronger viewing angles, the hot-spots become more visible, because of the non-Lambertian characteristics of the screen.

Because manually aligning a large number of projectors for a stereo tiled display is time consuming, we present an automatic calibration procedure. We apply a method based on matched-filter convolution, to obtain sub-pixel accuracy in our geometric alignment. This precise calibration enables us to properly align left and right-eye images. By performing photometric calibration we can compensate for the difference in luminance over the surface of the screen for a particular viewing position. Finally, calculation of the correct viewing frustum by taking into account the viewer position, the inter-ocular distance and the fusion point, makes the stereo image comfortable to watch.

Several informal measurements confirm that our calibration is accurate, although more thorough measurements are certainly required. Finally, we discussed several applications illustrating the use of this system in practice.

Currently the ICWall at the Vrije Universiteit is one of just few places where scientists can use such a large tiled display in stereo mode for visualization of their data and immersive VR presentations. It is not an easy task to setup such a system and keep it ”alive”. But it is definitely possible and also affordable.

7 Acknowledgements

This work was carried out in the context of the Virtual Laboratory for e-Science project (www.vl-e.nl). This project is supported by a BSIK grant from the Dutch Ministry of Education, Culture and Science (OC&W) and is part of the ICT innovation program of the Ministry of Economic Affairs (EZ).

References

BODEN, N. J., COHEN, D., FELDERMAN, R. E., KULAWIK, A. E., SEITZ, C. L., SEIZOVIC, J. N., AND SU, W.-K. K. 1995. Myrinet — a gigabit-per-second local-area-network. *IEEE Micro 15*, 1 (Feb.), 29–36.

CHEN, Y., CLARK, D. W., FINKELSTEIN, A., HOUSEL, T. C., AND LI, K. 2000. Automatic alignment of high-resolution multi-projector displays using an uncalibrated camera. In *IEEE Visualization 2000*, 125–130.

CRUZ-NEIRA, C., SANDIN, D., AND DEFANTI, T. 1993. Surround-screen Projection-based Virtual Reality: The Design and Implementation of the CAVE. In *Computer Graphics (SIGGRAPH '93 Proceedings)*, J. T. Kajiya, Ed., vol. 27, 135–142.

DEBEVEC, P., AND MALIK, J. 1997. Recovering High Dynamic Range Radiance Maps from Photographs. In *ACM SIGGRAPH'97 Conference Proceedings*, 369–378.

FUNKHOUSER, T., AND LI, K. 2000. Guest editors' introduction: Large-format displays. *IEEE Computer Graphics and Applications 20*, 4 (July), 20–21.

GERMANS, D., SPOELDER, H., RENAMBOT, L., AND BAL, H. 2001. VIRPI: A High-Level Toolkit for Interactive Scientific Visualization in Virtual Reality. *Immersive Projection Technology and Virtual Environments 2001* (may), 109–120.

KRESSE, W., REINERS, D., AND KNPFLE, C. C. 2003. Color consistency for digital multi-projector stereo display systems: the heyewall and the digital cave. In *EGVE '03: Proceedings of the workshop on Virtual environments 2003*, ACM Press, New York, NY, USA, 271–279.

LEIGH, J., DAWE, G., TALANDIS, J., HE, E., VENKATARAMAN, S., SANDIN, D., AND DEFANTI, T. 2001. AGAVE: Access Grid Augmented Virtual Environment. In *Proceedings of the AccessGrid Retreat*.

MAJUMDER, A., AND STEVENS, R. 2002. LAM: Luminance Attenuation Map for Photometric Uniformity in Projection Based Displays. In *ACM Virtual Reality and Software Technology*, 3–11.

PAPE, D., ANSTEY, J., BOGUCKI, M., DAWE, G., DEFANTI, T., JOHNSON, A., AND SANDIN, D. 1999. The ImmersaDesk3 - Experiences with a Flag Panel Display for Virtual Reality. In *Proceedings of the 3rd International Immersive Projection Technology Workshop*, 107–112.

PAPE, D., ANSTEY, J., AND DAWE, G. 2002. A Low-Cost Projection Based Virtual Reality Display. In *The Engineering Reality of Virtual Reality 2002 SPIE Electronic Imaging: Science and Technology 2002*, 483–491.

RAIJ, A., GILL, G., MAJUMDER, A., TOWLES, H., AND FUCHS, H. 2003. Pixelflex2: A comprehensive, automatic, casually-aligned multi-projector display. In *IEEE International Workshop on Projector-Camera Systems (PROCAMS-2003), Nice, France*.

RASKAR, R., VAN BAAR, J., BEARDSLEY, P., WILLWACHER, T., RAO, S., AND FORLINES, C. 2003. ilamps: Geometrically aware and self-configuring projectors. *ACM Transactions on Graphics (TOG) 22*, 3 (July), 809–818.

RASKAR, R. 2000. Immersive Planar Display using Roughly Aligned Projectors. In *Proceedings of the 2000 IEEE Conference on Virtual Reality (VR-00)*, IEEE, 109–116.

STEELE, R. M., WEBB, S., AND JAYNES, C. 2002. Monitoring and correction of geometric distortion in projected displays. *Journal of WSCG 10*, 2, 429–434.

VOS, F. 1998. *A System for Measuring, Modelling and Visualizing Corneal Shapes Based on Pseudo Random Encoding*. PhD thesis, Vrije Universiteit Amsterdam.

VUYLSTEKE, P., AND OOSTERLINCK, A. 1990. Range image acquisition with a single binary-encoded light pattern. *IEEE Trans. Pattern and Anal. Machine Intell. 12*, 2, 148–164.