Large Scale Projector Alignment for the Wide Area Virtual Environment

Nagaraj Raparthi
Eric Acosta
Benjamin Wagley
Alan Liu

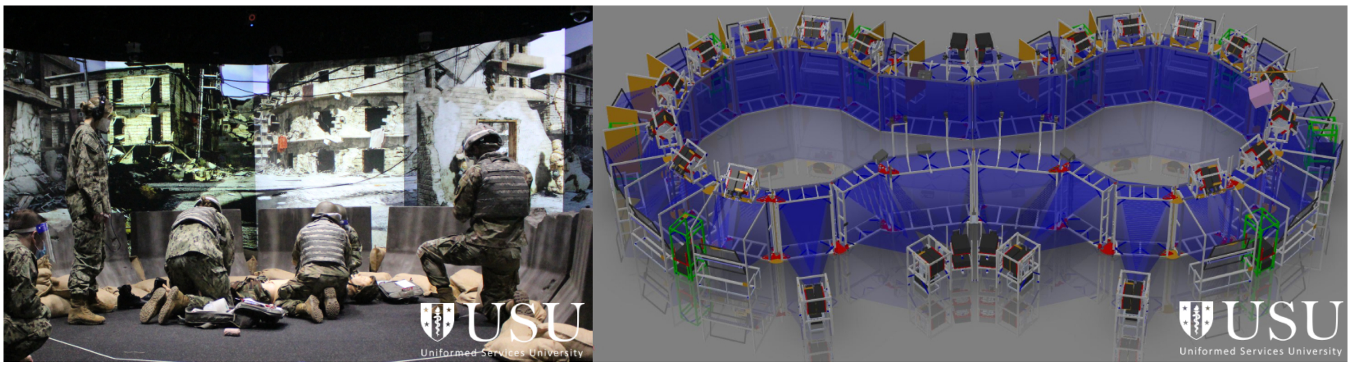

Figure 1. Medical training and projector setup inside the Wide Area Virtual Environment.

Abstract

The Wide Area Virtual Environment (WAVE) is an 8,000 sq. ft. virtual reality facility for medical team-based training. A high resolution tiled stereoscopic display, involving an array of 48 projectors, creates an immersive environment. Projector alignment is critical to the training experience. However, frequent projector image drifting makes it difficult to maintain alignment. The need for a permanently installed solution, and the unique shape of the tiled display imposes difficulties in employing other methods. This paper describes a scalable image warping method developed to automate projector alignment. The multi-threaded and distributed approach makes very frequent alignments possible via a tablet and inexpensive webcams.

ACM Reference Format:
Nagaraj Raparthi, Eric Acosta, Benjamin Wagley, and Alan Liu. 2021. Large Scale Projector Alignment for the Wide Area Virtual Environment . In *SIGGRAPH Asia 2021 (SA '21 Posters), December 04-13, 2021.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3415264.3425474

SA '21 Posters, December 04-13, 2021, Virtual Event, Japan

ACM ISBN 978-1-4503-8113-0/20/11.
https://doi.org/10.1145/3415264.3425474

1 Introduction

The Wide Area Virtual Environment (WAVE) presents a unique setting for medical team learning within an immersive virtual environment [Liu et al. 2018]. To engage students in the immersive stereoscopic experience, the shape of the WAVE was designed as two near 360 degree field-of-view circular pods connected by a corridor. Figure 1 illustrates the structure of the WAVE. The system is comprised of 24 tiled screens with modular displays. Each module includes a screen, and a stereo pair of projectors and image generators; resulting in a total of 48 projectors and image generators.

Misaligned stereoscopic displays can result in eye discomfort, and in some cases disorientation among the users. Unfortunately, the presented images from the WAVE's projectors drift daily resulting in misalignments like shown in Figure 2. Manual realignment and warping of the projected images is not feasible since it takes a qualified technician a full day to finish aligning all the projectors inside the WAVE. This paper presents a solution developed to automate this process which makes it possible for a single person to align and warp all tiled displays within 35-50 minutes.

2 Related Work

The use of multiple projectors for higher resolution tiled displays is well documented [Chen et al. 2000; Majumder et al. 2000]. Over the years, various computer vision based techniques have been developed to align multi-projector displays [Brown et al. 2005]. Most of these techniques implement warping algorithms which digitally distort images to

Figure 2. Daily drifts cause mis-alignments between screens.

align the projected images [Raskar et al. 1999; Surati 2005]. Typically, inexpensive and commonly available webcams are placed in front of the projected display to capture images of the mis-aligned projectors and perform computer vision algorithms. Using this technique Raskar et al. are able to accommodate various displays such as corner walls and curved screens. SimWall is an example of a tiled large scale stereoscopic system that supports both geometry and photo metric calibration for projector alignment [Xie et al. 2007]. In some cases camera calibration can also be avoided [Chen et al. 2000].

Many of the existing methods require the cameras to capture the entire display in order to align the projected images. The shape of the WAVE's structure makes this impossible. It is also difficult to place cameras in the training space due to props covering the front of the screens. To overcome these limitations, an image registration method was developed to calculate the parameters required to warp the projected images into alignment.

3 Method

As the daily projector image drifting and misalignment presented a major challenge, the goal was to develop a method to align the WAVE's stereoscopic tiled display prior to starting the training sessions. This required a permanently installed, fast and accurate alignment solution that a single person can run with ease.

This section describes the method used to align the 48 projectors in the WAVE. The image registration algorithm developed to automate projector alignment is first described. Next, the distributed implementation of the method used to parallelize the alignment is discussed.

3.1 Alignment Algorithm

The projectors have an integrated warp map that can be used to remap the pixel locations of their displayed image. Alignment of the WAVE's tiled display is achieved by generating and uploading warp maps for each projector. Computing a warp map requires establishing a warp grid to control how the image is manipulated. Offsetting the warp grid points defines how the image's sub-regions within the grid are distorted. The projector provides a manual method of manipulating a warp grid to align its display. Automating this process involves 1) displaying a warp grid, 2) capturing an image of the screen with a webcam, 3) computing the warp grid point offsets (via a computer vision algorithm) required to align the projector and 4) computing and uploading a warp map for the projector.

A custom application is run on each image generator to display a warp grid. A 5x5 warp grid was sufficient for accurate alignment.

The shape of the WAVE prevents a single camera from capturing the entire tiled display. One Logitech 4K webcam was installed per screen. Due to space constraints behind the screens, each camera can only capture an image of its screen, as shown in Figure 3. This makes it possible to align stereo projector pairs for each screen with traditional methods. However, this does not align the projectors between neighboring screens. To overcome this limitation, a ground truth reference was established for each screen after one alignment of the tiled display. This reference is only captured once, but makes it possible for each screen's projector pair to be warped independently to match its reference, but still be aligned with the other 46 projectors.

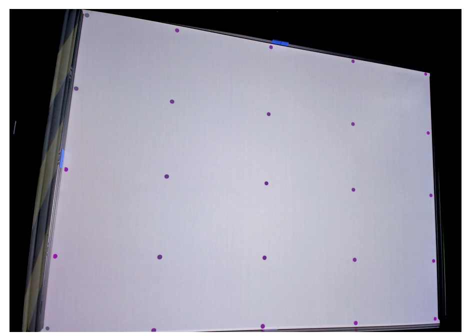

Figure 3. Image captured from the webcam with the warp grid points displayed.

The coordinates of the grid points in the webcam's image serve as the ground truth fiducials for the image registration alignment algorithm. The image generator of a screen connected to a webcam captures an image, and utilizes OpenCV's Hough Circle detection algorithm to identify the coordinates of the grid points in image space. Image distortions resulted in significant error in the coordinates using only the default algorithm. The circle detection algorithm was augmented by scanning a bounding box around the detected circle to identify the centroid. This lead to consistently accurate coordinate identification even for distorted circles.

Subsequent warping and alignment works by displaying the same (now misaligned in the display) 25 grid points and computing the offsets required to align the points to the ground truth of each screen. A simple, but effective, point walking algorithm was developed to compute the offsets.

Image-space vectors between the points' coordinates in the current webcam's image and the ground truth coordinates are used to guide the points' movement. The offsets cannot be used directly due to camera image distortions, camera and screen misalignment and a difference in camera and screen resolution. To reduce overshooting the target location, each point is walked a portion of it's offset vector. Moving a point 60% along the offset vector worked really well. The point walking algorithm iteratively moves the points until they are all within 1 pixel of their targets in image space. Once the walking algorithm completes, the offsets for the warp points are determined from the initially displayed grid points and their final locations.

The computed warp grid offsets are used to compute a warp map by solving the Laplace equation in 2D with the finite difference method: $\Delta^2 f = \frac{\partial^2 x}{\partial x^2} + \frac{\partial^2 y}{\partial y^2} = 0$ subject to Dirichlet boundary conditions. An elementary Jacobi Iteration scheme is first used to fix the boundary values and then choose initial guesses for the interior points. Then, in multiple iterations ϕ is calculated at all of the interior 5X5 grid points and the 16 warp regions formed, according to the formula below. The iterations are completed once a minimum error threshold is met at which point the equations would be solved with the most accuracy.

$$\phi_{i,j}^{m+1} = \phi_{i,j}^{m} + \frac{1}{4}[(\phi_{i+1,j}^{m} + \phi_{i-1,j}^{m}) + (\phi_{i,j+1}^{m} + \phi_{i,j-1}^{m}) - 4\phi_{i,j}^{m}]$$

The projector alignment algorithm requires all 25 warp grid points to be visible on the screen. As the projected image overshoots the physical screen, some warp grid points may not be initially displayed on the screen. A center offset calculation step was introduced to run prior to computing the offsets for all 25 points. The center offset performs the same offset calculation and warp map generation steps as described above, but only using the center grid point. After the center offset warp map is uploaded, all 25 points are guaranteed to be displayable on the screen. In the event that a boundary grid point is still not on the screen it is moved incrementally towards the grid's center until it becomes visible. Since the center offset serves as a "global offset," it is simply added to the offsets calculated for all 25 points.

3.2 Implementation

Due to the modular nature of each screen, it was possible to implement a distributed projector alignment system. Figure 3 provides an overview for one screen. Each screen module includes a screen, a stereo pair of projectors and image generators (PCs), and a Logitech 4K web camera.

An HTTP based Image Server application was developed to display the grid points, capture webcam images and perform the computer vision algorithms for offset calculations. This application runs on each image generator.

An HTTP server application, Warp Server, was developed to function as the conductor for the alignment system which runs on a centralized computer server. It communicates with the distributed Image Servers throughout the alignment process. The Warp Server parses the information received from the Image Servers to generate the warp maps. It is also responsible for communicating with all the projectors to shutter them when needed and handles warp map uploading. A user can access the Warp Server's web interface on any device with a browser. The web interface provides functionalities such as projector de-warping, ground truth calculation, center offset calculation, warp grid offset calculation, warp map generation, preview and commit. A tablet is used to access this web-page and wirelessly align the WAVE.

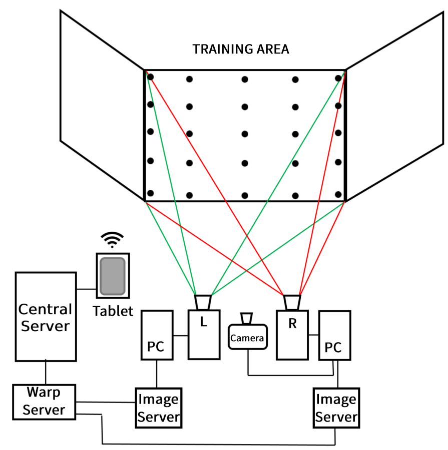

Figure 4. Architecture setup of one tiled screen module.

To begin the alignment process, a user first selects the projectors and de-warps them. A de-warp is achieved by uploading a warp map with zero-offsets for the warp grid. Next, a center offset step is performed. This is followed by an offset calculation step for the 5x5 warp grid points. Since the screens are independent of each other, all offset calculation steps for the WAVE's screens can be performed in parallel. When an offset (or center only offset) calculation is requested, the Warp Server shutters the right projector of a screen and sends a request to the left eye's Image Server to compute the offset calculation. Once the Warp Server receives the left eye's offsets from the Image Server it shutters the left projector, unshutters the right projector, and requests an offset calculation from the right eye's Image Server. Since only one webcam is used per screen, the Warp Server also facilitates the communication required for the left eye's Image Server to display its grid points, while the right eye's Image Server captures images and computes the offsets. The offset calculations are complete once the Warp Server receives offsets from all Image Servers. The final steps involves computing the warp maps, uploading the warp maps for preview and

Table 1. Time taken to perform alignment

Alignment Step	Time(mins)
Ground Truth Calculation	< 1
Center Offset Calculation	10 - 12
Grid Offsets Calculation	10 - 12
Warp Map Generation	< 1
Warp Map Upload and Commit	6 - 7

then committing the warp maps to the projectors' memory. All of these final steps are multi-threaded and performed in parallel. To verify the alignment the user can display various test patterns on all the screens.

4 Results and Conclusion

The described method is used to align the WAVE's projectors for the 4 corridor and 20 pod screens, oriented in landscape and portrait mode respectively. The projectors within the same screen and neighboring screens can be aligned within 0-2 screen pixels (or 0-1/8 inch). Figure 5 illustrates. The biggest challenge was due to some webcams needing to be installed close to the floor to remain out of the projector's light path. This resulted in fewer pixels, and greater lens distortion towards the top of the screens, as shown in Figure 3. The worst area is located around the screens' top corners. For those cases, the majority of the screen is aligned within 0-2 screen pixels. However, the top corners for those worst cases had an error up to 6 screen pixels (or 3/8 inch), as shown in Figure 6. Running the offset calculation a second time to refine the alignment corrected those areas within 0-2 screen pixels, as also shown in Figure 6.

The Digital Projection HIGH-lite Laser II projector is used in the WAVE. The image generators and central server are equipped with an Intel Core i7 processor with 64 GB of RAM running Windows 10. On average, the WAVE can be aligned within 35 minutes for one alignment pass, and 50 mins if a second refinement pass is required. The distributed and multi-threaded implementation contributed significantly to the alignment system speed. A breakdown of each step is provided in Table 1.

The described system is used to align the WAVE located at the Val G. Hemming Simulation Center. Reducing the one day manual alignment time, down to under an hour, makes it possible to optimize the projector alignment prior to a training session. Utilizing 4k webcams made it possible to compute sub-pixel accurate offsets. The described implementation relied on the projector's warp map capability to align the projectors. However, the method can still work for other warping methods such as NVIDIA's Warp and Blend, or using a custom distortion mesh within a rendering engine by formatting the warp grid/map results appropriately.

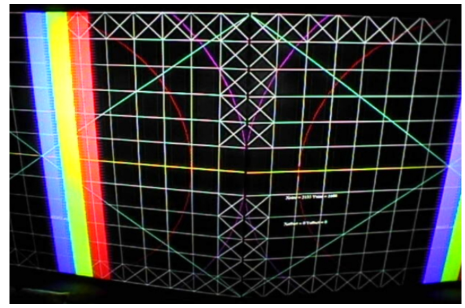

Figure 5. Aligned 4 projectors of neighboring screens displaying the same pattern.

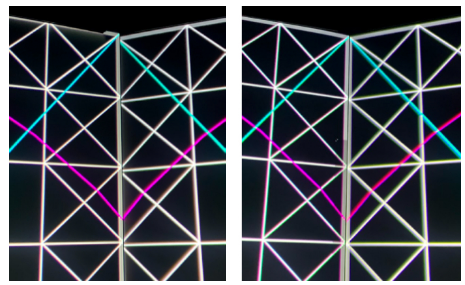

Figure 6. Seam between two worst case top screen corners after one alignment (left) and corrected after a second alignment (right).

References

M. Brown, A. Majumder, and R. Yang. 2005. Camera-based calibration techniques for seamless multiprojector displays. *IEEE Transactions on Visualization and Computer Graphics* 11, 2 (2005), 193–206. https://doi.org/10.1109/TVCG.2005.27

Yuqun Chen, Douglas W. Clark, Adam Finkelstein, Timothy C. Housel, and Kai Li. 2000. Automatic Alignment of High-Resolution Multi-Projector Display Using an Un-Calibrated Camera. In *Proceedings of the Conference on Visualization '00* (Salt Lake City, Utah, USA) *(VIS '00)*. IEEE Computer Society Press, Washington, DC, USA, 125–130.

Alan Liu, Eric Acosta, Jamie Cope, Valerie Henry, Fernando Reyes, Joseph Bradascio, and Wesley Meek. 2018. The Wide Area Virtual Environment: A New Paradigm for Medical Team Training. In *Augmented Cognition: Users and Contexts*, Dylan D. Schmorrow and Cali M. Fidopiastis (Eds.). Springer International Publishing, Cham, 293–304.

Aditi Majumder, Zhu He, Herman Towles, and Greg Welch. 2000. Achieving Color Uniformity Across Multi-Projector Displays. (08 2000). https://doi.org/10.1109/VISUAL.2000.885684

Ramesh Raskar, Michael Brown, Ruigang Yang, Wei-Chao Chen, Greg Welch, Herman Towles, Brent Seales, and Henry Fuchs. 1999. Multi-Projector Displays Using Camera-Based Registration. *IEEE* (08 1999). https://doi.org/10.1109/VISUAL.1999.809883

Rajeev Surati. 2005. Scalable Self-Calibration Display Technology for Seamless Large-Scale Displays. (08 2005).

Li-jun Xie, Yao Zheng, Ting-jun Yang, Wenxuan Gao, and Ning-he Pan. 2007. SimWall: A practical user-friendly stereo tiled display wall system. *Journal of Zhejiang University Science A* 8 (04 2007), 596–604. https://doi.org/10.1631/jzus.2007.A0596