

Seamless Multi-Projector Display on Curved Screens

Jeroen van Baar¹, Thomas Willwacher¹, Srinivas Rao² and Ramesh Raskar¹

¹ Mitsubishi Electric Research Labs (MERL), Cambridge, MA USA

² University of Maryland, College Park, MD USA

Abstract

We describe a new technique to display seamless images using overlapping projectors on curved surfaces. Our method addresses issues such as automatic registration, smooth intensity blending and efficient rendering. Previous techniques for automatically registered seamless displays have focused mainly on planar displays. Techniques for curved screens currently involve cumbersome manual alignment to make the installation conform to the intended design.

Keywords:

Calibration, Projectors, Seamless Displays, Curved Surfaces, Rendering, Virtual Reality.

Categories and Subject Descriptors: I.3.3, I.3.7 [Computer Graphics]: Picture/Image Generation, Virtual Reality; I.4.1, I.4.5[Image Processing and Computer Vision]: Digitization and Image Capture, Reconstruction

1. Introduction

Large seamless display using overlapping projectors is an emerging technology for constructing high-resolution semi-immersive visualization environments capable of presenting high-resolution images from scientific simulation, large format images in entertainment and surround environment for instruction. In this way, they complement other multi-projector technologies such as the CAVE [3], Blue-C [11,6] and well-defined tiled planar [5] or curved displays [6].

A major drawback of those systems is the reliance on precise display geometry. Projectors and display surfaces must be in exact, pre-defined alignment to produce seamless imagery. This restriction requires special purpose mounting and display infrastructure and a great deal of time and effort to deploy and maintain. This also highly constrains the physical space where the display system can be deployed, usually forcing the space to become dedicated to housing the system as a permanent fixture.

Current packaging systems for tiled displays also make it somewhat difficult to build compact or wrap around displays that will ultimately be needed for immersive display systems. We believe the prototype system described in this paper is a contribution in that direction. In particular, we have designed a system for deploying modestly sized (4-6 projector), but highly flexible curved display composed entirely of commodity-based components. To accomplish

this, we have incorporated robust camera-based geometric registration and corrective warping. This self-calibrating display removes all but the simplest tasks, reducing traditional restrictions on projector alignment and display surface shape. We provide a practical, relatively inexpensive multi-projector display system suitable for use by groups such as museums, libraries, schools, small business, and exhibitors.

New automatic approaches for creating seamless displays for curved screens, either parametric or non-parametric, are currently lacking in the literature. We exploit new image transfer functions recently published in the computer vision field to improve on the previous calibration and rendering procedures. We note that a purely parametric approach is not feasible given that a well-defined curved display screen is difficult to build.

2. Previous Work

In commercially available planar displays, alignment is typically performed manually. However, many research groups have exploited camera-based approaches to automate this process [14, 12, 8, 10, 4]. In addition, recent software developments ease the use of PC clusters equipped with graphics cards to power immersive projection environments where multiple video projectors form a high resolution and large surface display [1]. For curved screens, commercial systems allow projection of a navigation grid

pattern. Then all the overlapping projectors are manually aligned with the grid [13].

In [9] a non-parametric approach based on discrete sampling of the display area using a collection of cameras was proposed. It involves computing Euclidean three-dimensional parameters of the projectors, i.e. internal and external calibration of the projective devices, in addition to determining the 3D geometric shape by zipping meshes. However, this approach is too difficult to implement in a practical setup.

In this paper, we propose several new algorithms. The well-known planar projective transform (or planar homography) is traditionally used to geometrically register and pre-warp images for planar displays. However, curved screens present new challenges. First, a parametric approach such as pre-warping homography can be used only in limited cases when the screens are curved. Second, due to the display configuration, a single camera may not be able to view all the projections. Finally, the intensity blending functions [7] for seamless transition in the overlap area requires an accurate geometric relationship between projectors and cameras.




Figure 1 Multiple overlapping projectors augmented with cameras (red discs). (a) projection onto a convex curved surface. (b) projection onto a concave curved surface.

3. Prototype System

Our current prototype system can accurately register overlapping projections on curved surfaces that adhere to the quadric equations. The system consists of projector-camera

pairs. The cameras are depicted as red circles in Figure 1. The intrinsic parameters for the camera and Euclidean rigid transformations for the projector relative to the camera in the projector-camera pair have been determined in an off-line process. This projector-camera pair calibration or simply calibration, is not strictly necessary. However, it greatly improves the registration accuracy of multiple overlapping projectors.

Figure 1 shows two possible setups of multiple projector-camera pairs, with the projections overlapping on a curved surface. The setups in the figure illustrates that for a specific camera some projections are entirely visible, some projections are only partially visible and some projections are entirely obscured. Currently our prototype system can only register projections that are entirely or partially visible. For example projectors 1 and 2 in Figure 1a, and projectors 1, 2, and 3 in Figure 1b.

The setup depicted in Figure 1a is an example of the projectors projecting on a convex curved surface. With respect to the curved surface, the projectors are projecting on the outside. In Figure 1b the projectors are projecting on the inside, or the concave side of the curved surface. In section 6 we show results of our technique on both a convex and concave curved surface.

4. Geometric Registration

To geometrically register the projectors we choose one of the projector-camera pairs as our *base-pair*, indexed as k . All registration information is calculated relative to this pair, similar as for planar display surfaces [8, 9, 14]. The outline of our registration technique is as follows:

1. Project structured light pattern with projector P_k .
2. All cameras record image of current projection.
3. Determine 3D reconstructed points of display surface for all cameras, which correspond with the structured light features.
4. Determine extrinsic parameters from 3D reconstructed points and detected image features for projectors relative to camera C_k .
5. Fit quadric to the 3D points in all camera image spaces for which the structured light projection was visible.
6. Take quadric fit as initial estimate for quadric transfer between C_k and refine re-projection errors with non-linear minimization.
7. Apply quadric transfer parameters to pre-warp the input image into camera image spaces.

Given a calibrated projector-camera pair, we can reconstruct the 3D points for the structured light features using triangulation. The reconstructed points will model the shape of the underlying curved surface. The projected structured light pattern onto the surface is (partially) detected by the cameras of neighboring projector-camera pairs. From the corresponding 2D image points and 3D surface points, we can determine the Euclidean rigid transformations of the neighboring cameras with respect to the projector currently projecting the structured light pattern.

Because our curved surfaces adhere to the quadric equations, given the Euclidean rigid transformations we can fit quadrics to the reconstructed 3D points. We perform a fit for all cameras for which the structured light projection was visible. We use a simplified form of the quadric transfer, the mapping between two perspective views of a quadric. The quadric fit to the 3D reconstructed points is used as an initial estimation of the quadric transfer parameters. This initial estimate and finding the pose of a camera from known 3D points on a quadric are not very accurate. The latter is due to the lack of depth variation across the quadric surface. Therefore a non-linear refinement is used to minimize the re-projection error for the quadric transfer parameters.

4.1. Intensity Normalization

In order to create seamless displays the perceived increased brightness in overlapping areas has to be corrected. This correction is sometimes denoted as *intensity normalization* [9]. The goal of this normalization is to adjust the brightness in overlapping regions such that it is the same as the (perceived) brightness in non-overlap regions. One possible way to achieve such a brightness adjustment is by calculating some weight function that can be used as an *alpha* channel during rendering [7, 9]. The accurate registration we achieve using our technique meets the requirement for calculating such a weight function to correctly adjust the brightness.

5. Rendering

For rendering, we treat the quadric transfer as a homography via the so-called polar plane plus a per-pixel shift. We pre-warp the input image into each camera's image space. Note that warping an image using a quadric transfer is different than rendering quadric surfaces. Current PC class graphics hardware provides support for user-defined vertex and pixel shaders. We have implemented the quadric transfer using a simple vertex shader program.

We currently have two simple programs for rendering. One renders a triangle and the other one renders a textured rectangle. In the case of the triangle we compute the pro-

jection of each vertex of a densely tessellated triangle using the quadric transfer. Similar for the case of the rectangle, but first we map the input image as a texture onto a densely tessellated rectangular mesh. The mesh is then animated using simple rigid transformations.

6. Results

We have implemented a system to demonstrate our techniques on an eLumens VisionStation hemispherical dome surface. The hemispherical surface, depicted in Figure 2 has a diameter of 1.5 meters. First we show the results of registration of three projections on the concave side, i.e. inside or front, of the dome surface. Second, we show results of registration of two projections on the convex side, i.e. outside or rear, of the dome surface. The casual placement of the dome surface results in a small deviation from its spherical form. However we will demonstrate that our technique still produces accurate results.



Figure 2 The eLumens VisionStation hemispherical dome surface. The dome surface has a diameter of 1.5m.

Figure 3 shows the structured light pattern projections of each of the three projectors on the dome surface. The projections show how each projector is oriented with respect to the dome surface. The relatively large rotation of the second projection shows that our method works for arbitrarily oriented projectors. Figure 4 shows the accurate registration of the three projections. The left image in Figure 4 shows each projector simultaneously rendering a triangle, and all three triangles are accurately registered. The right image demonstrates that three triangles are rendered. The three projections are separated using an offset surface. Figure 5 shows some examples of an animated textured object. All three projections remain accurately registered during animation of the textured object.




Figure 3 Structured light projections of three projectors on the concave side of the dome surface.




Figure 4 A triangle rendered simultaneously with three projectors. (left) The triangles are accurately registered on the dome surface. (right) The three projections are separated using an offset surface.



Figure 5 Four images from a sequence of accurately registered animated textured rectangles. Each of the three projectors projects a rectangle with a different texture. Projector P1 projects the rectangle textured with the “original” image. P2 projects the rectangle textured with the inverse of the original image. Similar for P3 but the inverse is augmented with red circles. The top left image shows the result with P3 entirely obscured. The top right image shows the result with P1 entirely obscured. The two bottom row images show results with P1 partially obscured.



Figure 6 Two overlapping projectors on the convex, i.e. rear, side of the VisionStation dome surface.

We have also applied our technique for registering projections on the convex side of the dome surface. This dem-

onstrates the generality of our technique. Projecting on the convex side would be particularly useful if the dome was made of transparent material. The top image in Figure 6 shows a setup using two overlapping projectors. The middle and bottom image show the projected structured light patterns. Note that the projected patterns are more distorted due to the convex curved surface. For this reason we limit our demonstration to two overlapping projectors on the convex side of the dome surface.




Figure 7 A triangle rendered simultaneously with two projectors. (left) The triangles are accurately registered on the convex side of the dome surface. (right) The two projections are separated using an offset surface.

The left image of Figure 7 shows the rendering of a triangle for the registered projectors on the convex side of the dome surface. The right image shows again that the registration is only achieved on the dome surface. Finally, Figure 8 shows some examples of an animated textured object on the convex surface.

7. Conclusions and Future Work

We have presented a technique to display seamless images using overlapping projectors on quadric curved surfaces. We believe the proposed approach is ideal for building curved screens, such as cylindrical, dome shaped surfaces and their variants for various applications. Multi-projector displays on curved screens are widely used in planetariums, flight simulators, large visualization centers and for large format curved panoramic movies. Our automatic registration eliminates tedious setup and maintenance for building curved screens systems, and hence reduces cost.




Figure 8 Four images from a sequence of accurately registered animated textured rectangles. Each of the two projectors projects the rectangle with a different texture. Projector P1 projects the rectangle textured with the “original” image. P2 projects the rectangle textured with the inverse of the original image. The top left image shows both projections in full overlap. The top right image shows the object partially outside the projection of P1. The bottom left image shows a close-up of this case. The texture is suddenly cut off along the projection boundary, while the projections remain geometrically aligned. The bottom right image shows the object entirely outside the projection of P1. Only the inverse textured rectangle projected with P2 is now visible.

Hardware support for rendering is achieved by implementing the quadric transfer as a vertex shader. For moderately high complex virtual scenes this would allow a dedicated cluster of PCs to render at interactive frame rates. This is similar to the performance achieved for planar displays.

We are currently working on extending our technique to work for general curved surfaces. Our current prototype can accurately register projections on surfaces that adhere to the quadric equations. The problems with such surfaces are that (a) they are usually not readily available and have to be custom manufactured, and (b) deviations from a quadric surface may result in noticeable registration errors.

We also would like to extend our technique for the case where some projections are entirely obscured from a certain camera’s viewpoint. The desired geometric information will have to be determined by cascading the geometric information between some number of in-between, mutually visible projector-camera pairs. This would allow easy and flexible setup of wrap-around displays to enhance the immersive experience.

References

1. Bues, M.; Blach, R.; Stegmaier, S.; Häfner, U.; Hoffmann, H.; Haselberger, F. “Towards a scalable high performance application platform for immersive virtual environments,” Eurographics Workshop on Virtual Environments, 2001.
2. Yuqun Chen, Douglas W. Clark, Adam Finkelstein, Timothy Housel, and Kai Li, “Methods For Avoiding Seams On High-Resolution Multi-Projector Displays Using An Un-calibrated Camera,” Technical Report TR-618-00, Department of Computer Science, Princeton University, April 2000.
3. Cruz-Neira, C., Sandin, D.J., DeFanti, T.A., Kenyon, R., and Hart, J.C. “The CAVE, Audio Visual Experience Automatic Virtual Environment.” Communications of the ACM, June 1992, pp. 64-72.
4. FuturesLab, ActiveMural, Argonne National Labs.
5. Jupiter Inc, <http://www.jupiter.com>.
6. Kunz Andreas, Spagno Christian, “Modified Shutter Glasses for Projection and Picture Acquisition in Virtual Environments,” Proceedings of the IEEE Virtual Reality 2001, Yokohama, Japan, March 2001, pp.281-282.
7. Lyon, Paul. “Edge-blending Multiple Projection Displays On A Dome Surface To Form Continuous Wide Angle Fields-of-View,” pp. 203-209. Proceedings of 7th I/ITEC, 1985.
8. Princeton University, Scalable Display Wall.
9. R. Raskar, M. S. Brown, R. Yang, W.-C. Chen, G. Welch and H. Towles. “Multi-Projector Displays Using Camera-Based Registration.” In Proceedings of IEEE Visualization 1999, October 1999.
10. Raskar R., van Baar, J., Chai, J. A Low Cost Projector Mosaic with Fast Registration. In Proceedings of Asian Conference on Computer Vision, 2002
11. Staadt Oliver, Gross Markus, Kunz Andreas, Meier Markus, Gross Markus, “The Blue-C: Integrating Real Humans into a Networked Immersive Environment”, ACM Collaborative Virtual Environments 2000.
12. Stanford Interactive Workspaces Project, Information Wall.
13. Trimension Inc, <http://www.trimension-inc.com/>.
14. Ruigang Yang, David Gotz, Justin Hensley, Herman Towles, and Michael S. Brown. “PixelFlex: A Reconfigurable Multi-Projector Display System,” IEEE Visualization 2001. San Diego, CA.