Table-Top Spatially-Augmented Reality: Bringing Physical Models to Life with Projected Imagery

Ramesh Raskar, Greg Welch, Wei-Chao Chen
University of North Carolina at Chapel Hill
{raskar, welch, ciao}@cs.unc.edu
http://www.cs.unc.edu/~raskar/Tabletop/

Abstract

Despite the availability of high-quality graphics systems, architects and designers still build scaled physical models of buildings and products. These physical models have many advantages, however they are typically static in structure and surface characteristics. They are inherently lifeless. In contrast, high-quality graphics systems are tremendously flexible, allowing viewers to see alternative structures, facades, textures, cut-away views, and even dynamic effects such as changing lighting, moving automobiles, people, etc.

We introduce a combination of these approaches that builds on our previously-published projector-based Spatially-Augmented Reality techniques. The basic idea is to aim multiple ceiling-mounted light projectors inward to graphically augment table-top scaled physical models of buildings or products. This approach promises to provide very compelling hybrid visualizations that afford the benefits of both traditional physical models, and modern computer graphics, effectively "bringing to life" table-top physical models.

1. Introduction

In [Raskar98c] we introduced the general notion of Spatially Augmented Reality (SAR), where physical objects are augmented with images that are integrated directly in the user's environment, not simply in their visual field. For example, images can be projected onto real objects using light projectors, or embedded directly in the environment with flat panel displays. For the purpose of this paper we concentrate on the former, in particular for the specific case where multiple ceiling-mounted projectors are aimed inward so that they illuminate and can augment table-top scaled physical models of buildings or other objects.

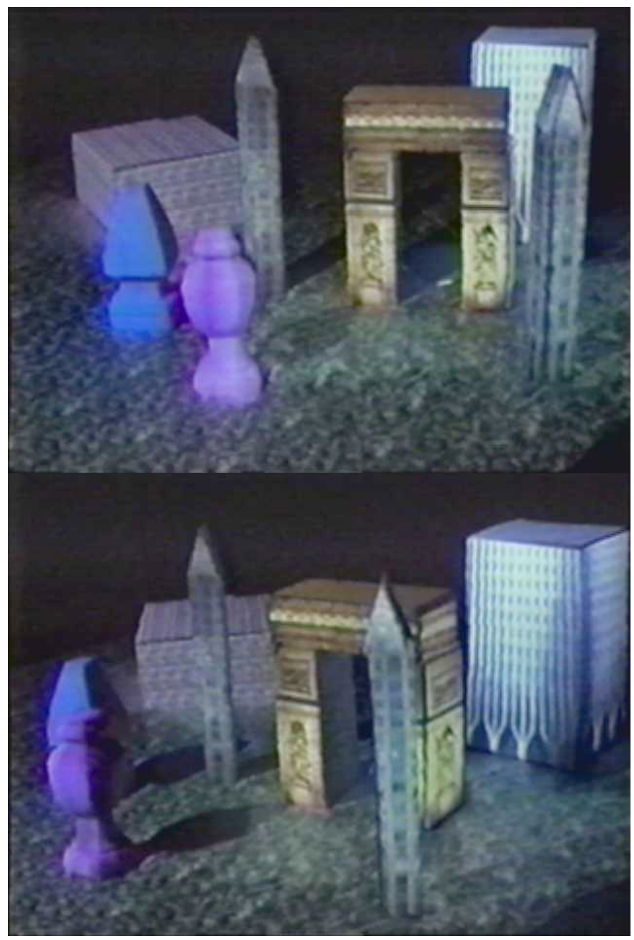

Figure 1. Two different views of simple physical models augmented with projected imagery. (The underlying physical models are shown in

This setup promises very compelling hybrid visualizations that afford benefits heretofore exclusively afforded by either physical or graphics models. Like traditional physical models, the augmented physical model could be viewed in 3D from any position around the table, by multiple people, without head-tracked stereo glasses. However as is typical with modern computer graphics, one could easily depict alternative surface attributes, changing lighting conditions, dynamic objects, and other helpful 2D information. If one was willing to wear tracked stereo glasses, one could make virtual modifications to the physical model, adding or removing components, depicting internal structure, etc. In either case, multiple inward-pointing projectors can be used to afford very high resolution and highly-saturated imagery.

Figure 2. The underlying physical models from Figure 1. The physical objects are wood, brick, and cardboard.

When more than one inward-pointing projector is used to create virtual imagery, two central problems need to be solved:

- We need to calibrate the display environment and achieve static registration between the projectors and objects in the working volume. If 3D virtual imagery is desired, the corresponding head-tracking system also needs to be registered.
- We also need to solve the more difficult problem of generating seamless images by achieving geometric registration between overlapping projections. Building on our previous work [Raskar99a] we propose to use static video cameras and standard active computer vision techniques to compute the necessary 3D representation of the projector parameters and the surfaces of the real objects. The problem of achieving seamless imagery with multiple projectors has been explored for simple configurations by [Humphreys99] [Panoram] [Raskar98d] [Raskar99a] [Trimensions]. However, in this case, due to the presence of concave objects or a collection of disjoint objects, the regions of overlap between the two projectors are not necessarily contiguous.

1.1 Applications

The hybrid physical/graphics model-based SAR approach described in this paper has certain restrictions when compared to pure physical or graphics model approaches. However, it offers an interesting new method to realizing compelling high-fidelity illusions of virtual objects and surface characteristics coexisting with the real world. Two example applications are augmented visualization of table-top architectural models of one or more buildings, and augmented visualization of bench-top parts and procedures for assembly line workers or repair technicians.

In the first example, an architect could provide clients with a compelling form of walk-around scaled model of the real (proposed) buildings or complex. In a simple demonstration in [UnderKoffler99], a single projector is used to illuminate blocks on a table-top for urban planning. With SAR, at a minimum, assuming the surfaces of the physical model are diffuse white, the approach could be used to "paint" different colors and textures onto the surfaces of the physical model. (The textures could also convey some notion of 3D surface perturbations by using bump mapping for example.) In addition she could show the clients the building as it would appear under varying lighting conditions, including night time with building lights on, daylight with the sun in varying positions, and both over varying seasons. Finally, she could show the clients parts of the internal structure of the building, including pipes, electrical wiring, etc.

In the second example, an assembly line worker could be guided through the assembly process via spatially-augmented information. Head-mounted display AR has been used for this application at the Boeing Corporation [Curtis98]. Using the techniques in this paper we believe one could achieve the same effects without the need of a head-mounted display, using inward-pointing projectors to render instructional text or images on a white work surface.

1.2 Hybrid Model Visualization

In purely virtual environments (VE), one renders graphics models of real objects, usually together with computer generated virtual objects. In contrast, the basic notion of Augmented Reality (AR) is to enhance physical objects with computer generated virtual objects. In the case of immersive (HMD-based) VE, the user sees physical and virtual objects at the same limited spatial and temporal resolution and fidelity. One advantage of projector-based Spatially-Augmented Reality [Raskar98c], like optical-see-through HMD-based AR, is that the spatial and temporal fidelity of the physical object is preserved and only the additional data is rendered at limited resolution. In contrast, with video-see-through HMD-based AR, images of virtual objects are rendered and superimposed with video images of the physical objects, so again the user sees physical and virtual objects at the same limited resolution and fidelity.

Here we are interested in Spatially-Augmented Reality in the specific case where the physical object being augmented by projected imagery is itself a model of interest—in fact, a *physical model* that matches the basic structure of the *graphics model*, the representation used by the computer. In the most basic example, the visualization of a building (for example) makes use of both a physical model and a graphics model of the building. The physical model has the proper structure or shape, but no color or texture. The graphics model minimally includes the structure (identical to the physical model), the colors, the textures, and any other surface attributes. In addition, the graphics model might contain some purely virtual components for which there is no physical counterpart. In effect, the user is viewing a hybrid physical and graphics model, getting advantages from both.

1.3 Projector Configurations

Previously, multiple overlapping projectors have been used primarily to create large panoramic displays. The user typically stands in front of the displayed images (e.g. InfoMural [Humphreys99], InfinityWall [Czernuszenko 97]) or inside the large field-of-view display environment (e.g. Cave [Cruz-Neira93], Office of the Future [Raskar98a] [Raskar99a]). We call this an *inside-looking-out* projection system. In most cases, one aligns the projectors so that the neighboring projections overlap side-by-side. The region on the display surfaces simultaneously illuminated by two or more projectors is usually a (well-defined) single contiguous area. Further, the corresponding projector pixel coordinates change monotonically. This is similar to the monotonic ordering of corresponding pixels in stereo camera pairs.

Here we envision a table surrounded and illuminated by a collection of ceiling-mounted projectors, where users can visualize and possibly interact from anywhere around the table. We call this an *outside-looking-in* projection system. One can imagine using the projectors to render onto a simple display surface such as a sphere or a cube, creating a crystal-ball type visualization system. Another setup would be looking into a concave hemispherical bowl illuminated to render high-resolution 2D or head-tracked 3D imagery that you can walk around.

In this paper we are more interested in visualization system where one can change 2D attributes such as color or texture, and possibly 3D attributes, of known three-dimensional physical models that themselves form the display surface. We have previously demonstrated [Raskar98a] how to render perspectively correct images on smooth but non-planar display surfaces. In this case, due to the presence of concave objects, or a collection of disjoint objects, the regions of overlap between two or more projectors are not necessarily contiguous, and corresponding pixels do not maintain monotonic ordering. This is a major difference and creates new challenges when a seamless image of the virtual object is to be rendered in the overlap region. In this paper, we discuss the motivation for such a system and suggest an approach for calibration and rendering for such a setup.

2. Usefulness

At one extreme, if a detailed physical model of an object is available, the model is clearly going to be higher resolution, more responsive, easier on the eyes, essentially better than almost anything Virtual Reality (VR) has to offer—for a static model. At the other extreme, clearly pure VR has the advantage in that you can show the user "anything," static or dynamic, without the need for a physical model. We believe that this hybrid Spatially-Augmented Reality approach can offer some of the advantages of each of the two situations, when a physical model is either readily available or obtainable. We believe that the combination has significant potential. Even simple static demonstrations are extremely compelling, bright, clear, and easy to look at. (Please see the video available at project webpage.)

In general, assuming you want to augment a physical object with 2D or 3D graphical information, you have a several alternatives [Milgram94]. For example, you could use a video or optical see-through head-mounted display.

In fact, one major advantage of Spatially Augmented Reality achieved using light projectors is that the user does not need to wear a head-mounted display. (In [Bryson97] and [Raskar98c] the various advantages of spatially immersive displays over head-mounted displays for VR and AR have been noted.) In video see-through AR, or pure VR for that matter, the physical and virtual objects are both rendered at a limited pixel resolution and frame rate i.e. limited spatial and temporal resolution. In the hybrid SAR approach however, the spatial resolution depends only on the display parameters of the projector such as its frame buffer resolution, field of view and distance from the illuminated object. The spatial and temporal resolution of static scene is independent of the viewer location or movement. Thus, using a fixed set of projectors much higher resolution imagery, text or fine detail can be realized.

If only surface attributes of real objects are to be changed, then the calibration, authoring and rendering are much easier. In this case, the rendering is viewer independent, no stereo display (projection) is necessary and multiple people around the real object can simultaneously see the augmentation. Even if the virtual objects are not strictly surface attributes, but are near the real surfaces on which they are displayed, the eye-accommodation is easier. Most of these advantages are shared by all spatially-augmented reality setups.

To be fair, such a hybrid approach has some disadvantages. The approach cannot in general be said to be better than pure physical or graphics models, but better than each in certain respects under certain circumstances, and worse in others. For example, you must have or be able to obtain (using our methods for example) a graphics model of the physical model. Also, one of the advantages of video see-through AR is that virtual imagery can easily occlude the images of real objects. In projector-based SAR, if the surfaces of the physical model are not pure white, one might not be able to completely occlude portions of the physical model, should that be necessary.

3. Methods

We have developed a simple interactive approach to modifying the surface characteristics of multiple table-top physical models. The approach essentially involves manually adjusting projected image texture coordinates to visually align with the physical models. While not sophisticated, we have shown the results to many people, and the overwhelming consensus is that these simple results are extremely compelling. (Please see the video available at project webpage.)

More significantly, building on our previous work we have developed a comprehensive automatic approach for modifying the surface characteristics of the physical model, and adding 3D virtual objects. While we are still working on demonstrating this full approach, we have demonstrated individual portions, and hope to have a full demonstration soon.

The full approach for augmenting physical models involves first determining the relationships between various components in the environment and their parameters. These components include video cameras, light projectors, physical model and the head-tracking system. We refer to this as the *calibration phase*. Next, the user might need to interactively associate parts of the graphics model with the corresponding parts of the physical model, or they might want to alter parts of the graphics model. We refer to this as *authoring*. Finally, during run time we use advanced rendering techniques to augment the physical model with perspectively correct virtual objects for the head-tracked user.

3.1 Calibration

We propose to use multiple ceiling mounted inward-looking static video cameras to capture geometric information about the physical model. The video cameras can themselves be calibrated by observing a common calibration pattern such as a cube with carefully pasted checkerboards on each of its visible side [Tsai86][Faugeras93]. After the intrinsic and extrinsic parameters of the cameras are computed, the calibration pattern can be removed. By projecting active structured light with projectors, calibrated stereo camera pairs can be used to compute the depth in the scene. The primitives in the structured light could be a dense set of binary encoded dots projected by each projector. By stitching together the depth values computed by each stereo camera pair, one can create a 3D surface representation of the entire physical model. Since multiple projectors will be used, it is necessary to create a unique and continuous geometric representation of the physical model so that we can display overlapping images without visible seams. The extracted physical model can be stored as a polygonal model, the *graphics model*. During depth extraction, we can also determine the correspondences between 2D pixel coordinates of a given projector and the 3D locations illuminated by those pixels. If corresponding pixels for six of more 3D surface points are known, one can calibrate the projector and find the projection parameters of that light projector. Finally, if a head-tracking system is used, the transformation between the tracker's coordinate system and the working volume's coordinate system can be computed by taking readings of the tracker sensor at multiple positions and corresponding positions of the sensor computed by triangulation with calibrated stereo camera pairs.

When more than one projector illuminates a part of the physical model, we need to ensure that the projected images are geometrically aligned. This is analogous to creating photo mosaics by stitching images together as in [Szeliski96][Sawheney97]. We need to compute correspondences between multiple projector pixels. Each camera observes which pixel of different projectors illuminated the same surface point on physical model. The set of projector pixels in correspondence can be indirectly calculated from these observations [Raskar98d] [Raskar99a].

3.2 Authoring

One of the important tasks in achieving compelling augmented reality is to create association between the physical objects and the graphics primitives that will enhance those objects when projected. Examples of graphics primitive are lines, text, texture mapped polygons or even complete 3D (virtual) objects. For example: which texture image should be used for the face of a building model? What color distribution will look better for a physical model? A user interface is critical in creating the graphics primitives with different shape, color and texture. A similar user interface is required for positioning and aligning the graphics primitive so that it is correctly projected on the desired part of the physical model.

3.3 Rendering

If one wants to change only the surface attributes of the physical model such as color or texture, then it may not be necessary to completely compute the 3D graphics models of the physical model or the projection parameters of the light projectors. For example, if the user wants to change color of one face of a building on a tabletop architectural model then s/he only needs to find the set of pixels from one or more projectors that illuminates that face of the building. Those pixels can be determined interactively without explicit 3D representation. The pixels can be colored or applied pre-warped textures to change the appearance of the face of the building. On the other hand, if the 3D graphics model of the building and projector parameters are known, then one can easily pre-compute the set of projector pixels that illuminate the face of the building. When only surface attributes (for diffuse surfaces) are changed, the rendering can be assumed to be view-independent and no head tracking is necessary.

When virtual objects are to be rendered in head-tracked 3D, one can use the two-pass rendering method described in [Raskar98a]. With this method virtual objects can be made to appear perspectively correct even when the underlying surfaces of the physical model is not planar. In the first pass, the desired image of the virtual object for the user is computed and stored as a texture map. In the second pass, the texture is effectively projected from the user's viewpoint onto a polygonal graphics model of the physical model. The polygonal graphics model, with the desired image texture mapped onto it, is then rendered from the projector's viewpoint. This is achieved in real-time using projective textures [Segal92]. As described in [Raskar99a], usually a third pass of rendering is necessary to ensure that the overlapping images projected from multiple projector are geometrically aligned.

When multiple projectors overlap, the luminance in the overlap region may be much greater than that in regions illuminated by only one projector. Thus in addition to geometric alignment between projected images, it is also necessary to achieve intensity normalization. The problem of generating seamless images using multiple projectors has been explored for large wide-field-of-view displays [Panoram] [Trimensions] [Raskar99a] [Raskar99b], as well as two-dimensional arrays of flat projections [Humphreys99] [Czernuszenko97]. In such cases, the overlap region is typically a (well-defined) contiguous region on display surface as well as in each projectors frame buffer. The intensity of projector pixels is weighted using feathering (also known as intensity roll-off or soft-edge) techniques so that the overlapping images blend to create a single seamless image. In case of multiple projectors looking inwards, if we have a single convex physical object illuminated by a rectangular projected image, the overlap region for any two projector is also contiguous. However, typically the physical model is made up of non-convex objects or a collection of disjoint objects resulting in overlap regions that are fragmented in each projector's frame buffer. In [Raskar99a] we described and demonstrated an image blending technique to achieve geometric alignment and intensity normalization to create seamless images from multiple projectors. The image blending technique can be used even if the single contiguous overlap region is not rectangular or the illuminated surface is not flat. When the overlap region is not contiguous, however, one first needs to identify the pixels in each projector's frame buffer that illuminate surface(s) also illuminated by at least one other projector. Using a simple region-growing algorithm in each projector's frame buffer it should be possible to identify the different islands of overlapping regions. The image blending technique described in [Raskar99a] can then be used for each of these islands.

4. Registration Issues

In augmented reality, preserving the illusion that virtual and real objects coexist requires proper alignment and registration of virtual objects to real objects [Azuma94][State96][Neumann96]. Traditional AR methods use body-centric coordinate system to render synthetic objects, and SAR methods use a fixed world coordinate system to render them. However, in both cases, the static and dynamic registration errors are caused by a number of factors such as system delay, optical distortion and tracker measurement error, and are difficult to address with existing technology. The tracking requirements for registration in SAR are similar to spatially-immersive display (SID-VR) systems because real and virtual objects lie in the same fixed world-coordinate system. Thus, static calibration errors can play an important role in registration. They include correct estimate of transformations between display devices, tracker and world coordinate system. In the case of HMD-based AR, such errors result in the apparent translation of virtual objects with respect to the real objects. As noted in [Cruz-Neira93] and [Raskar98c], in SAR such errors lead to fundamentally different types of artifacts. For example, when the additional imagery is simply modifying the surface attributes, the rendered imagery is viewer independent and remains registered with static real objects. If 3D virtual objects are displayed on part of the physical model with which they are expected to be registered, then as described in [Cruz-Neira93][Raskar98c], the dynamic errors results in *shear* of virtual objects instead of translation. Finally, if floating 3D objects are to be displayed, the mis-registration is similar to HMD-based AR. This will also be the case if interaction with virtual objects involves movement with respect to the real objects.

5. Conclusion

In this paper we have presented the idea of augmenting physical models by surrounding them with light projectors and displaying seamless images on the surfaces of those objects. This method appears to be especially effective when the surface attributes of the real objects need to be augmented, for example surface color or texture. Multiple users can stand around and view the modified surface attributes without stereoscopic projection, glasses or HMD. We have described how the setup can be used to augment physical models by displaying perspectively-correct 3D virtual objects.

The hybrid visualization method can augment physical models with white diffuse surfaces by blending images from multiple projectors. However, currently this technique appears to be somewhat limited to visualization and not suited for complicated interaction with the virtual objects. One also needs to address the issue of aliasing if physical models with sharp edges are illuminated by limited resolution images. Shadows can also create a problem.

We look forward to refining our ideas and the related algorithms, developing better authoring programs, and to pursuing some of the many applications we have in mind. In our current setup, we have simply painted different colors and textures on top of the physical model. However, we plan to construct a complete setup with head-tracking to display 3D virtual objects in the next few months. In the end we believe the approach promises very compelling hybrid visualizations that afford the benefits of both traditional physical models, and modern computer graphics, effectively "bringing to life" table-top physical models.

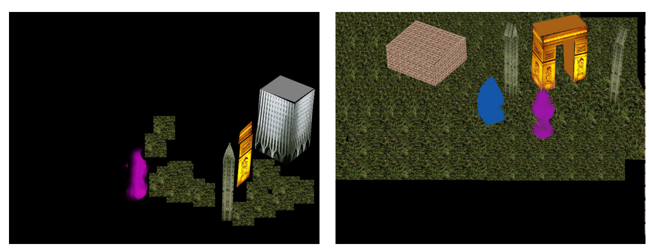

Figure 3. The images in projector framebuffers for the scene in Figure 1.

6. Experiment

We have demonstrated a simple table-top physical model illuminated by two video projectors and augmented by painting different textures and colors. The video of the demonstration is available at the project website. As shown in Figure 2, the scene is made up of white-colored wooden objects, cardboard boxes and bricks. The textures and colors are interactively painted using Adobe Photoshop. For facades, it is sufficient to specify four points to achieve the necessary pre-warping of textures.

Figure 1 shows the augmented scene with two projectors. The video also shows how colors can be interactively changed (in this case spray-painted). Then we show contribution of each projector. When we turn on the room lights, one can see the simplicity of the physical model (also shown in Figure 2). The images in each of the two projector framebuffers are shown in Figure 3. Although this experiment uses only two projectors and a simple physical model, complex architectural models when illuminated with multiple projectors and viewed with head tracking will be more pleasing to look at than with the pure VR or HMD-based AR displays. Please see the project web site http://www.cs.unc.edu/~raskar/Tabletop for more media.

Acknowledgements

This research is supported by the National Science Foundation agreement ASC-8920219: “Science and Technology Center for Computer Graphics and Scientific Visualization”, Intel Corporation, and the “National Tele-Immersion Initiative” sponsored by Advanced Networks & Services, Inc. We would like to thank members of the Office of the Future group at UNC Chapel Hill and in particular Prof. Henry Fuchs and Herman Towles for useful discussions and support.

7. References

[Azuma94] Azuma, R., Bishop, G. Improving Static and Dynamic Registration in an Optical See-through HMD. Proceedings of SIGGRAPH 94 (Orlando, Florida, July 24-29, 1994). In *Computer Graphics* Proceedings, Annual Conference Series, 1994, ACM SIGGRAPH, pp. 197-204.

[Bryson97] Bryson, Steve, David Zeltzer, Mark T. Bolas, Bertrand de La Chapelle, and David Bennett. The Future of Virtual Reality: Head Mounted Displays Versus Spatially Immersive Displays, SIGGRAPH 97 Conference Proceedings, Annual Conference Series, ACM SIGGRAPH, Addison-Wesley, pp. 485-486, August 1997.

[Cruz-Neira93] Carolina Cruz-Neira, Daniel J. Sandin, and Thomas A. DeFanti. 1993. Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE, SIGGRAPH 93 Conference Proceedings, Annual Conference Series, ACM SIGGRAPH, Addison Wesley.

[Curtis98] Dan Curtis, David Mizell, Peter Gruenbaum, and Adam Janin, “Several Devils in the Details: Making an AR App Work in the Airplane Factory”, First IEEE Workshop on Augmented Reality (IWAR'98), November 1, 1998, San Francisco, CA.

[Czernuszenko 97] Marek Czernuszenko, Dave Pape, Daniel Sandin, Tom DeFanti, Gregory L. Dawe, Maxine D. Brown, “The ImmersaDesk and Infinity Wall Projection-Based Virtual Reality Displays”, Computer Graphics, May 1997.

[Faugeras93] O. Faugeras. Three-Dimensional Computer Vision: A Geometric Viewpoint. MIT Press, Cambridge, Massachusetts, 1993.

[Humphreys99] Humphreys, Greg and Pat Hanrahan. “A Distributed Graphics System for Large Tiled Displays” to be published in Proceedings of IEEE Visualization 99, San Fransisco, CA, October 24-29, 1999.

[Milgram94] P Milgram, H Takemura, A Utsumi and F Kishino. Augmented Reality: A class of displays on the reality-virtuality continuum. *SPIE Vol. 2351-34, Telemanipulator and Telepresence Technologies,* 1994.

[Neumann96] U. Neumann and Y. Cho, “A Self-Tracking Augmented Reality System”, ACM International Symposium on Virtual Reality and Applications, ISBN: 0-89791-825-8, pp. 109-115, July 1996

[Panorama] Panoram Technology. http://www.panoramtech.com

[Raskar98a] Ramesh Raskar, Greg Welch, Matt Cutts, Adam Lake, Lev Stesin, and Henry Fuchs. 1998.The Office of the Future: A Unified Approach to Image-Based Modeling and Spatially Immersive Displays, SIGGRAPH 98 Conference Proceedings, Annual Conference Series, Addison-Wesley, July 1998.

[Raskar98b] Raskar, Ramesh, Matt Cutts, Greg Welch, Wolfgang Stuerzlinger.Efficient Image Generation for Multiprojector and Multisurface Displays, Rendering Techniques ‘98, Drettakis, G., Max, N. (eds.), Proceedings of the Eurographics Workshop in Vienna, Austria, June 29-July1, 1998

[Raskar98c] Raskar, Ramesh, Greg Welch, Henry Fuchs. 1998. "Spatially Augmented Reality," First IEEE Workshop on Augmented Reality (IWAR'98), November 1, 1998, San Francisco, CA.

[Raskar98d] Raskar, Ramesh , Greg Welch, Henry Fuchs. 1998.Seamless Projection Overlaps Using Image Warping and Intensity Blending, Fourth International Conference on Virtual Systems and Multimedia, Gifu, Japan. November 1998.

[Raskar99a] Raskar, Ramesh, Michael S. Brown, Ruigang Yang, Wei-Chao Chen, Greg Welch, Herman Towles, Brent Seales, Henry Fuchs. 1999. "Multi-Projector Displays Using Camera-Based Registration," to be published in Proceedings of IEEE Visualization 99, San Fransisco, CA, October 24-29, 1999.

[Raskar99b] R. Raskar and M. Brown. Panoramic imagery using multiple projectors on planar surfaces. Technical Report TR-99-016, University of North Carolina, Chapel Hill, 1999.

[Sawhney97] H.S. Sawhney and R. Kumar. True multi-image alignment and its applications to mosaicing and lens distortion correction. In IEEE Comp. Soc. Conference on Computer Vision and Pattern Recognition (CVPR'97), 1997.

[Segal92] Mark Segal, Carl Korobkin, Rolf van Widenfelt, Jim Foran, and Paul E.Haeberli. 1992. Fast Shadows and Lighting Effects using Texture Mapping, SIGGRAPH 92 Conference Proceedings, Annual Conference Series, Addison Wesley, volume 26, pp. 249-252, July 1992.

[State96] State, A., Hirota, G., Chen, D.T., Garrett, W.F., Livingston, M.A. superior Augmented Reality Registration by Integrating Landmark Tracking and Magnetic Tracking. Proceedings of SIGGRAPH '96 (New Orleans, LA, August 4-9, 1996). In *Computer Graphics* Proceedings, Annual Conference Series, 1996, ACM SIGGRAPH.

[Szeliski96] R. Szeliski. Video mosaics for virtual environments.IEEE Computer Graphics and Applications, 16(2):22-30, March 1996.

[Trimensions] Trimensions. http://www.trimensions-inc.com/

[Tsai86] Tsai , Roger Y. An Efficient and Accurate Camera Calibration Technique for 3D Machine Vision, Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, Miami Beach, FL, pp. 364-374, 1986.

[UnderKoffler99] Underkoffler, J., and Ishii, H., Urp: A Luminous-Tangible Workbench for Planning and Design, in Proceedings of Conference on Human Factors in Computing Systems (CHI '99), Pittsburgh, Pennsylvania USA, May 15-20, 1999), ACM Press, pp. 386-393.