Spatially Augmented Reality

Ramesh Raskar, Greg Welch, Henry Fuchs
{raskar, welch, fuchs}@cs.unc.edu , (919)-962-1761
Department of Computer Science, University of North Carolina at Chapel Hill,
Chapel Hill, NC 27599, U.S.A.

Abstract

To create an effective illusion of virtual objects coexisting with the real world, see-through HMD-based Augmented Reality techniques supplement the user's *view* with images of virtual objects. We introduce here a new paradigm, *Spatially Augmented Reality* (SAR), where virtual objects are rendered directly within or on the user's *physical space*.

A key benefit of SAR is that the user does not need to wear a head-mounted display. Instead, with the use of spatial displays, wide field of view and possibly high-resolution images of virtual objects can be integrated directly into the environment. For example, the virtual objects can be realized by using digital light projectors to "paint" 2D/3D imagery onto real surfaces, or by using built-in flat panel displays.

In this paper we present the rendering method used in our implementation and discuss the fundamentally different visible artifacts that arise as a result of errors in tracker measurements. Finally, we speculate about how SAR techniques might be combined with see-through AR to provide an even more compelling AR experience.

1.   Introduction

In *Spatially Augmented Reality (SAR)*, the user's physical environment is augmented with images that are integrated directly in the user's environment, not simply in their visual field. For example, the images could be projected onto real objects using digital light projectors, or embedded directly in the environment with flat panel displays. For the purpose of this paper we will concentrate on the former. While the approach has certain restrictions, it offers an interesting new method to realizing compelling illusions of virtual objects coexisting with the real world. The images could appear in 2D, aligned on a flat display surface, or they could be 3D and floating above a planar surface, or even 3D and floating above an irregular surface.

For 2D non-head-tracked SAR, the images representing virtual objects do not continuously change with user motion. For example, in the Luminous Room [UnderKoffler97], the user's environment is enhanced with synthetic images projected on flat surfaces. However the user can be head-tracked and the images update dynamically to create an illusion that virtual objects are registered to real objects. Shuttered glasses can be used to facilitate stereo imagery, further enhancing the 3D effect of the virtual imagery. For the purpose of this paper we will focus on technologies for head-tracked SAR where virtual objects are rendered on irregularly shaped real objects. While not appropriate for every application, in this method the user does not need to wear a head-mounted display. Instead, with the use of spatial displays, wide field of view and possibly high-resolution images of virtual objects can be integrated directly into the environment.

It is our work on the Office of the Future [Raskar98a] that led to the realization and implementation of the SAR paradigm as described and analyzed in this paper. Specifically, when exploring the use of irregular (non-planar) surfaces for spatially immersive displays, we realized that the registration problem becomes somewhat unusual one of having to register 2D imagery with 3D physical geometry. This is similar to conventional AR applications, except for one crucial point: the 2D imagery exists on a virtual image plane that is attached to the (fixed) physical displays surface instead of

and ARC's dome shaped displays [Benette98] are other examples of SID. However, in OOTF the display surfaces are not limited to the designated flat walls (or parameterized surfaces) and could be everyday surfaces.

3. Methods

To create an illusion that virtual objects are registered to real objects for a moving user, we need to know the position of the user, projection parameters of the display devices, the shape of the surfaces of real objects in the physical environment and be able to render virtual objects on those surfaces. Here, we will describe a method for each of these in a unified projector-based SAR system.

3.1. Display surface shape extraction

The 3D surface shape extraction can be achieved using a calibrated [Tsai86] projector-camera pair where structured light patterns are projected and observed by the camera. [Raskar98c] describes a near-real time method to capture the 3D shape of the display surface and [Raskar98a] describes a unified approach to capture and display on irregular (non-planar) surfaces.

3.2. Rendering and viewing method

Projecting images on irregular surfaces so that they appear correct to a static user have been described in [Dorsey91] [Max91] [Jarvis97] [Raskar98d]. In [Raskar98a], a real time technique to generate such images for a moving head-tracked user was introduced. Here, we will describe how it can be used for SAR even when not all the viewing parameters are known. Let $V = V_i * V_e$ represent the user's eye perspective projection matrix, where, V_i is the projection matrix and V_e is the transformation matrix (subscripts i and e are used for intrinsic and extrinsic parameters, respectively). $E = E_i * E_e$ is an intermediate projection matrix, which shares the center of projection (COP), with the user's eye and P represents the projector's perspective projection matrix. Display surface model is **D** and **G** is the graphics model we wish to render. We will use the notation $I(r,g,b,z) = M*[G]$ to indicate that image I is generated from 3D colored model G using perspective projection matrix **M** (using any rendering method). Note that $M^{-1}*I(r,g,b)$ represents a set of 'colored' rays and $M^{-1}*I(r,g,b,z)$ represents colored surface interpolating 3D points due to termination of the rays at the corresponding depth. We want to present an image $I = V*[G]$ to the user with *only* the COP for **V** known. (Note that in general it is difficult to correctly estimate the COP for **V**, due to the obvious difficulty in co-locating the tracking sensor at the human eye COP. However, as we point out below, erroneously estimating the COP results in a fundamentally different visual artifact than head-mounted AR, an artifact that would be less noticeable for typical applications.)

Rendering

Step I :
(a) Compute the desired image color by rendering **G**:

$$I1(r,g,b) = E*[G]$$

(b) Update only the depth buffer (without overwriting color) by rendering the display surface model **D**.
$$I1(z) = E*[D]$$

As noted in the previous section, the system needs only the location of user's eye and not the orientation to render perspectively correct images of virtual objects on real surfaces. Hence, change in only the orientation of the user does not change the way real objects are 'painted'. Similarly, an error in the measurement of orientation in the tracker will not result in misregistration error. An error in measurement of position of the user will result in shearing of virtual objects that are projected on the real surfaces with which they are expected to be registered. The shear visible at any point on the virtual object is proportional to the distance between the virtual point from real display surface on which the object is projected. Therefore, a virtual object not in 'contact' with a real surface but fixed in world will exhibit shear as well as swimming when there is position measurement error.

4.   Advantages of SAR

A key benefit of SAR is that the user does not need to wear a head-mounted display. In [Bryson97] various advantages of spatially immersive displays over head-mounted displays have been noted. SAR shares similar benefits. In SAR, large field-of-view images can be generated with greater amount of integration of virtual objects with real world and also to improve sense of immersion if necessary. Projector-based SAR allows possibly higher resolution and bright images of virtual objects, text or fine details. Since virtual objects are typically rendered near their real-world location, eye accommodation is easier.

5.   Problems with SAR

The most crucial problem with projector-based SAR is its dependence on display surface properties. A light colored diffuse object with smooth geometry is ideal. It is practically impossible to render vivid images on highly specular, low reflectance or dark surfaces. The ambient lighting can also affect the contrast of the images. This limits application of SAR to controlled lighting environments with restrictions on type of objects with which virtual objects will be registered. For front-projector-based SAR, shadows of the user can create problems. This can be partially overcome using multiple projectors. SAR also allows only one active head-tracked user at any instant in the environment because the images are created in the physical environment rather than in individual user space. Time multiplexed shuttered glasses can be used to add more users that are active and head-tracked.

6.   Future Work

We have built a proof-of-concept system and demonstrated the SAR ideas in our laboratory, however we are anxious to test the approach with a real application. For instance, something like one of the many applications listed in the introduction, or an application that involves a more complicated physical display surface e.g. room-sized terrain visualization.

While the majority of our efforts to date have been focused solely on projector-based SAR, a hybrid environment could be built with, for example, digital light projectors and a see-through HMD. While this would require the user to wear an HMD, we believe that such a hybrid has the potential to offer the best of both worlds, combining all of the advantages of conventional see-through AR and this new SAR paradigm. As we have experience building both conventional AR systems and projection-based SAR systems, we look forward to building and experimenting with such a hybrid environment.

[Milgram94b] 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.

[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," to appear in SIGGRAPH 98 Conference Proceedings, Annual Conference Series, Addison-Wesley, July 1998.

[Raskar98b] Ramesh Raskar, Matt Cutts, Greg Welch, and Wolfgang Stürzlinger. 1998. "Efficient Image Generation for Multiprojector and Multisurface Displays," UNC Computer Science Technical Report TR98-016, University of North Carolina at Chapel Hill, March 1998.

[Raskar98b2] Raskar, Ramesh, Matt Cutts, Greg Welch, Wolfgang Stüerzlinger. "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, Henry Fuchs, Greg Welch, Adam Lake, Matt Cutts. "3D Talking Heads: Image Based Modeling at Interactive Rate using Structured Light Projection," University of North Carolina at Chapel Hill, Dept of Computer Science, Technical Report TR98-017, 1998

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

[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.

[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.

[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.

[UnderKoffler97] John Underkoffler. "A View From the *Luminous Room*," Springer-Verlag London Ltd., Personal Technologies (1997) 1:49-59.