Projected Imagery in Your "Office of the Future"

Greg Welch, Henry Fuchs, Ramesh Raskar, and Herman Towles
University of North Carolina at Chapel Hill

Michael S. Brown
University of Kentucky

In July 1998 we presented a long-term vision for a project we call the "Office of the Future" at ACM Siggraph 98.[1] Our dream, depicted in Figure 1, is that some day your office will have ceiling-mounted digital cameras and projectors that work together to support high-resolution projected imagery, human-computer interaction, and dynamic image-based modeling. In this article we discuss the display aspects, including motivations, challenges, techniques, and the future of projected imagery in your office.

Someday, high-resolution projected imagery will surround you in your office. The walls, your desk, and even the floor will serve as your computer desktop.

The magic of so many pixels

Most of us will put up with a small display if we have to carry it around. For example, information appliances such as the Palm Pilot have apparently struck such a nice balance of usefulness and size that people put up with the limited display size. But in your office, where you spend a great deal of time and work on "big" projects, small displays can be frustrating. In the Twentieth Anniversary Edition of his 1975 book *The Mythical Man-Month*, Fred Brooks observed

> The so-called "desktop metaphor" of today's workstation is instead an "airplane-seat" metaphor. Anyone who has shuffled a lap full of papers while seated in coach between two portly passengers will recognize the difference—one can see only a very few things at once. The true desktop provides overview of and random access to a score of pages. Moreover, when fits of creativity run strong, more than one programmer or writer has been known to abandon the desktop for the more spacious floor.[2]

If you have flown recently, you can appreciate this characterization, especially if you have tried to use a laptop computer on your airplane seat-tray "desktop." Yet even in the office, where we can have a reasonably sized CRT or LCD monitor, we are often pressed for real estate both on the screen and on the physical desktop. The solution for most people now is to use multiple monitors. However, even if the operating system will let you logically join the respective desktop regions, the physical displays are clearly disjoint. Furthermore, multiple monitors take up useful physical desktop space.

Projector-based display systems, on the other hand, offer the attractive combination of dense pixels over a wide area. Examples include the University of Illinois at Chicago room-sized CAVE[3] and InfinityWall

1 A conceptual sketch of the "Office of the Future" vision. (Sketch by Andrei State, UNC Chapel Hill.)

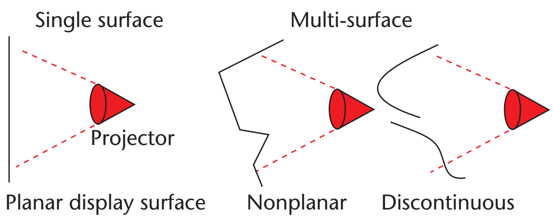

2 Conventional use of projectors (left) versus some situations that might arise from use in the office (right): surface(s) might be discontinuous or nonplanar, and the projection will be oblique in many cases.

display systems; Princeton University's 8 × 18-foot planar, scalable, display wall system[4]; the DataWall at the Massachusetts Institute of Technology; the PowerWall at the University of Minnesota; and commercial systems such as those by Trimension. The combination of physical scale and high-resolution imagery supports a natural means of interacting with the data: you can move up close to see more detail, or step back and look at the big picture. Your feet do the zooming—something we are accustomed to in the real world. Multiple people can enjoy that same capability at once, allowing natural interaction with local colleagues.

In your office, on every surface

Projector-based display systems are so impressive that we want one in our office for every-day work. Unfortunately the rear-projection display wall paradigm doesn't fit well within the office. Most systems are built in a large, dedicated laboratory space. A system for your office should minimize intrusive infrastructure, making the best use of existing space and structure. We envision using ceiling-mounted projectors to render imagery onto as many existing surfaces in the office as possible, turning every-day surfaces such as walls, desktops, and even floors, into display surfaces. Beyond visualization we want to use the projected imagery to interact with our computers. We want our office to embody visions such as those being explored by the Tangible Media Group at MIT, where cameras and projectors serve as "I/O bulbs," facilitating image-based human-computer interaction.[5,6]

As you look around your office, you might wonder how you would display imagery on the various surfaces. You will get some help from new technologies, but because we cannot overcome the laws of physics, you'll also need to redecorate. You will likely designate certain broad surfaces (walls or desktops) as your primary display surfaces. You'll cover these surfaces with material that controls the direction, amplitude, and phase of the projected imagery. To support occasional projection onto secondary surfaces, you'll choose materials and colors that minimize photometric problems as much as possible. For example, if you're willing to install light-colored carpet, imagery could "spill" onto the floor, letting you see traces of a document as you drag it from one display area to the real trash can, or to another (disjoint) area on the opposite side of the room. Some day you might even have the freedom to "abandon the [computer] desktop for the more spacious floor."[2]

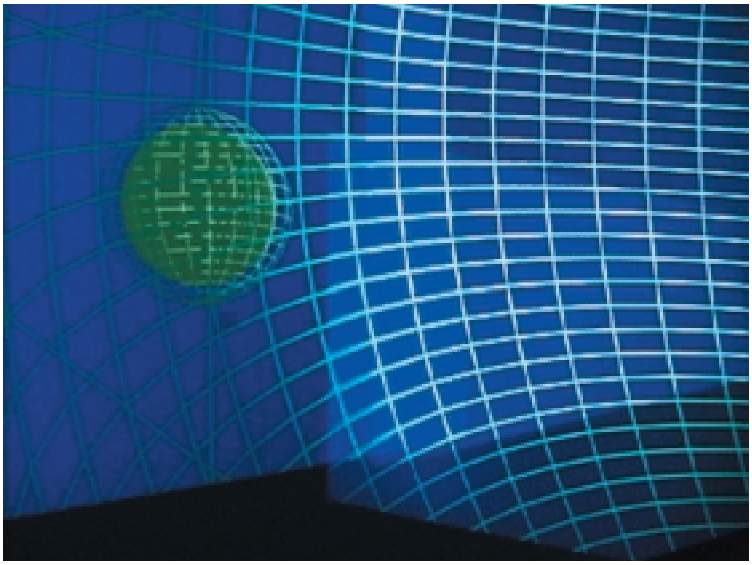

3 This image shows the need for intensity and color blending—the overlap from three projectors results in brighter regions in the overlap area (center of the image).

Fundamental challenges

Our vision and the corresponding sketch in Figure 1 are by definition futuristic. There remain many challenges to realizing even the projector-based display aspects.

Geometric issues

One of the most significant challenges is the unusual geometric relationship likely to exist between the projectors and display surfaces. (See Figure 2.) We sometimes talk about putting projectors "anywhere" and projecting onto "anything." We don't mean that literally, but even a reasonable arrangement of projectors around the office can result in unusual projection geometry. Thus the projected imagery must be prewarped so that it appears correct from wherever you sit. If the image content is strictly 2D—your Web browser, e-mail, and word processing—it's probably acceptable to make the imagery look geometrically correct from a single, fixed viewpoint. For head-tracked 3D imagery, it must be done continuously as you move.

The sampling and computational demands for perceptual geometric correctness are nontrivial. Consider that adults with unimpaired vision can resolve approximately 60 lines per degree of visual arc.[7] At one meter of distance—across your desk in the Office of the Future—this corresponds to approximately 3,500 × 3,500 pixels per square meter, with 0.3 millimeters of geometric tolerance. This means that, ideally, imagery would be projected at this resolution and geometrically corrected to this tolerance.

Image intensity and color

While some of the geometric issues may seem unique to our project, intensity and color problems affect virtually all projector-based display systems. The problems exist within the imagery of individual projectors, across regions where different projectors overlap (Figure 3), and in places where part of one display surface reflects

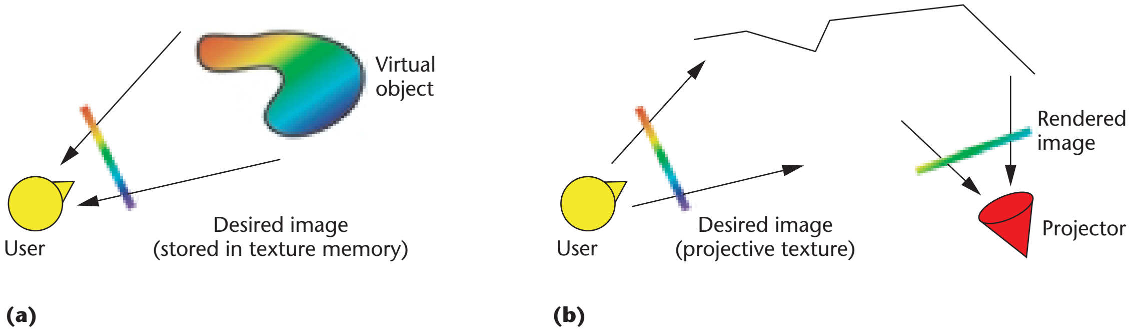

4 The two steps to our two-pass approach to rendering on nonplanar surfaces such as those shown in Figure 2. (a) Compute texture. (b) Project texture and rendering from projector's viewpoint.

light onto another (an inter-reflection). Humans are quite sensitive to such variations: we can perceive as little as a 2-percent variation in brightness[8] and a 2-nm variation in wavelength.[9] Similar challenges arise with monitors, printers, scanners, and film recorders.

For projected imagery there are essentially three relevant factors: the nature of the light source, the surface properties, and human perception. Light factors (projected and ambient) include the visibility of a surface with respect to the light, spatial resolution, angle of incidence with a surface, distance, and polarization. Surface factors include texture, gain, reflectance, persistence, polarization, and orientation. Human perception of light varies with wavelength, intensity, viewer distance to the surface, and image resolution. All of these issues are interrelated, making the challenges even more difficult.

The control of light in general may be one of the most difficult problems we face. Interreflections pose a problem for all projector-based display systems, but more so for us, given the relatively small volume of the typical office, where conditions for reflections abound. If only we had a means to remove light. As our colleague Gary Bishop once said, "We need some new physics!"

Projector technology

Five years ago a typical 3-lens CRT-based projector weighed almost 200 pounds, had less than one million pixels and 500 lumens of brightness, and cost $25,000 to $40,000. Today projectors can weigh less than five pounds and have 1,000+ lumens of brightness for under $7,000. Thanks to growing competition, it appears that the situation will only get better. New, reflective micro displays based on standard CMOS technology offer the hope of smaller, smarter, and less expensive digital projectors—perhaps as little as $100 some day.

However, we continue to face fundamental projector technology challenges for multiprojector systems in general and our Office of the Future in particular. While brightness and efficiency are important, dynamic range of intensity and color become even more important as we attempt to address the intensity and color challenges. In addition, we have surveyed many projectors and continue to see problems with inter- and intra-projector intensity and color uniformity. The characteristics of the optics can pose distortion, convergence, and focus problems. (A common practice is to adjust projector zoom to the point of minimal distortion, to translate the projector until the image is the proper size, and then to focus.)

Range of focus or depth of field poses a particular challenge to projecting onto multiple nonplanar surfaces. For now we address the problem by careful arrangement of the projectors. In the future we look forward to projectors with more coherent light sources; for example, see http://www.colorvision-lasers.com/. Finally, in addition to a digital interface for more precise geometry and color, flexible high-bandwidth controls would support time-division multiplexing for multiple stereo views and imperceptible structure-light for dynamic image-based modeling.[1]

On 24-25 February, the US Department of Energy organized an Advanced Display Workshop at the San Diego Supercomputer Center. Researchers from universities, national laboratories, and private companies met to discuss requirements for next-generation display devices, with a specific focus on light projectors. The group has stated plans to publish a white paper with prioritized display needs.

Some practical approaches

Despite all these problems, we're working on practical methods aimed at realizing our Office of the Future.

Multiprojector, multisurface rendering

In 1998 we introduced an efficient two-pass rendering approach to address the geometric problems of Figure 2.[10] In the first pass we render the desired image—an image that reflects what users should see given their current viewpoint and the graphics model (Figure 4a). We use OpenGL projective textures to effectively "project" that desired image onto a 3D model of the display surface, then render the result from the viewpoint of the projector (Figure 4b). We send this final (prewarped) image to the projector so that when users look at the display surface, they see an undistorted version of the desired image. This two-step process involves additional work, but the amount of work is constant, related only to the complexity of the display surface. Also, in many graphics engines the texture stack is hardware accelerated.[11]

This two-pass approach depends on the complete 3D calibration of the display surface and the projector. We first use a known physical calibration pattern—a textured wooden cube 70 cm on a side—to estimate the internal and external parameters for a pair of cameras. We then remove the cube and estimate the display surface geometry as follows. We illuminate a projector pixel *m* to create a feature on the display surface, observe the feature in each camera, and compute the correspond-

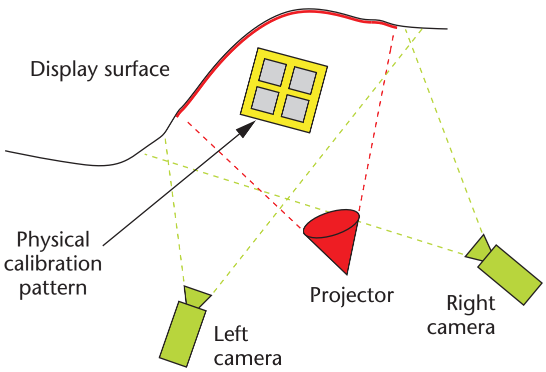

5 Display surface map and projector calibration setup.

ing 3D point *M* with respect to the camera. We repeat this process for each projector pixel to get a 3D display surface map *D*. Using the points from this map, and the corresponding projector pixel coordinates, we estimate the internal and external projector parameters.

This works well for a single projector, as in Figure 5, but new problems arise with multiple projectors. In particular, recall that the above method results in a distinct depth map in the space of each projector. Subtle errors in the geometry estimation may go unnoticed with a single projector, but where two projectors overlap, the differences in the estimated geometries can become noticeable.

To address this problem, we use a geometric optimization technique to unify the depth maps. This has the effect of moving each underlying 3D point, which corresponds to some projector pixel, to a new 3D location. This results in a new depth map for each projector, which optimally transitions to the neighboring maps. However a new point may no longer project back to the original projector pixel, but to a slightly different pixel. To address this image shift, we apply a final post-rendering image warp to move the pixel back to its proper location. This warp is static, so it only needs to be calculated once and can be implemented as a final texture mapping. In a separate article[12] we present the details of the entire calibration process, which for a two-projector setup can be completed in about 15 minutes.

Image intensity and color

A common approach to intensity blending in rectangular overlap regions is to create a smooth intensity roll-off near the edges. However, for the arbitrary overlaps that result from casual projector placement, the attenuation function must be more complex. For each projector pixel of each projector we compute a normalized weight in the range $0 \le \alpha \le 1$ that tends to zero approaching the edge of the frame buffer in the overlap region. We then load each alpha mask into the alpha-blending hardware of our graphics engine. The details for this and the preceding geometric blending appear in a separate publication.[12] Figure 6 shows the three alpha masks used to blend the imagery in Figure 3, and Figure 7 (next page) shows the final blended imagery, as well as some additional results.

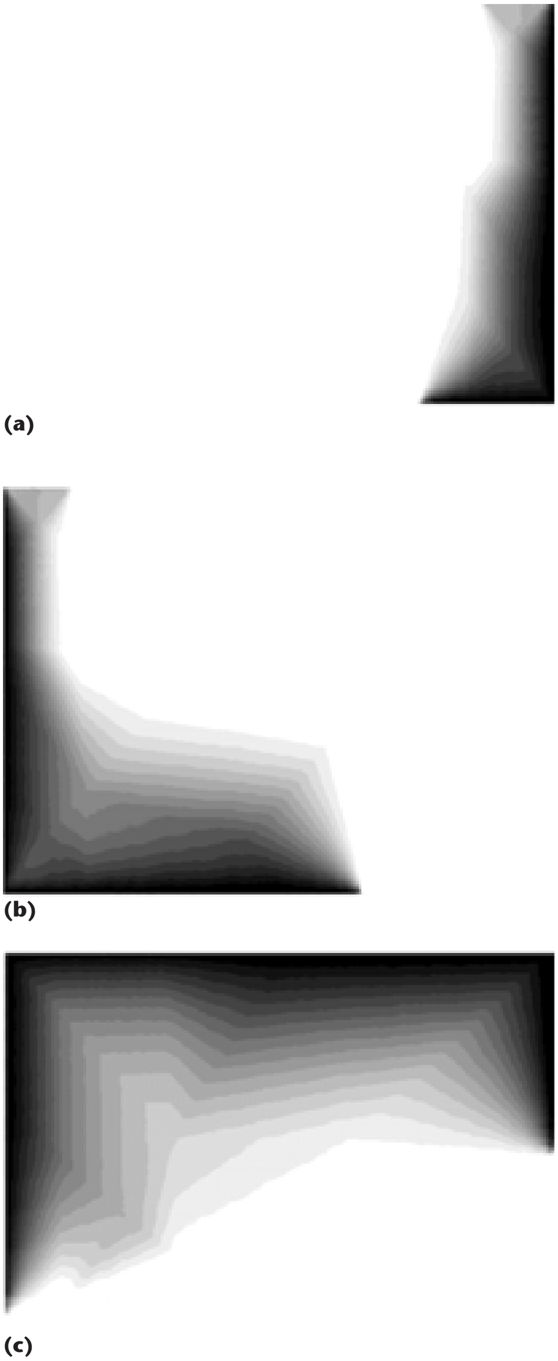

6 The three alpha masks used to blend the imagery in Figure 3. For this figure we gave the masks black borders to improve the visibility. (a) Mask for the left-most projector in Figure 3. (b) Mask for the upper right projector in Figure 3. (c) Mask for the lower right projector in Figure 3.

Beyond blending in the overlap regions, we're working on methods to normalize the image intensity and color gamut over multiple projectors. We have used a spectroradiometer (Photoresearch PR-705) to accurately characterize the radiometric response of various projectors among identical models and across different manufacturers. Using this information and fundamental principles of photometry, we're developing methods to map individual projectors to a common perceptually linear response. Currently, for real-time operation we are limited to three-channel (global) correction per projector. This work will be presented at (and published with) IEEE Visualization 2000 in October.[13]

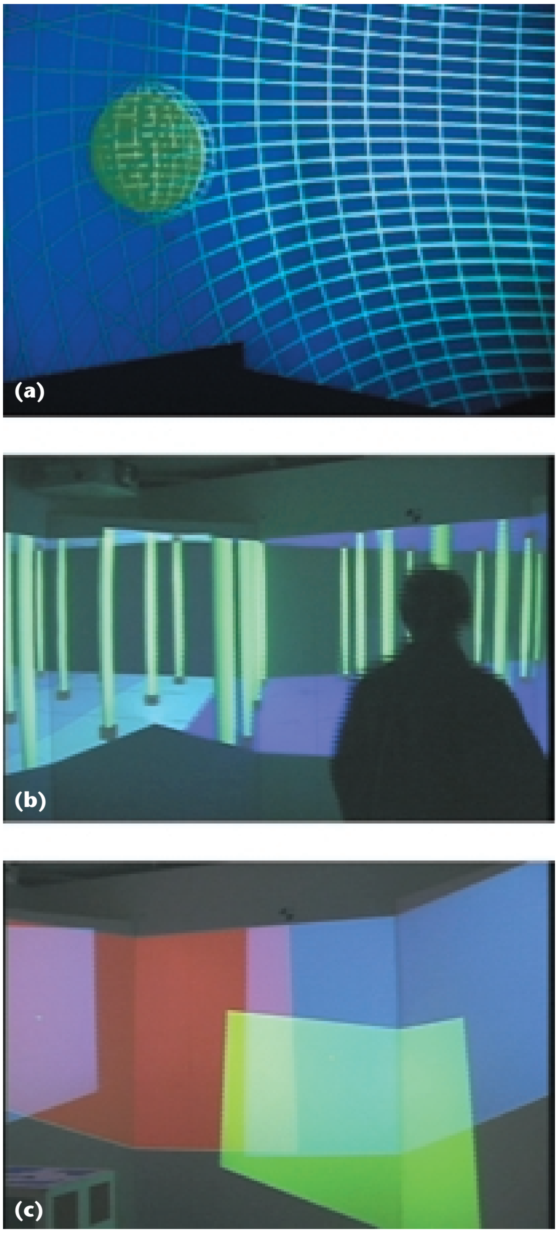

7 (a) Image showing the final blended imagery corresponding to the unblended imagery in Figure 3. (b) A panoramic image formed with four projectors. (A head-tracked user stands in the foreground.) (c) We replaced the blended projector images with solid colors to show the contribution from each projector.

Continuous autocalibration

A common problem with multiprojector systems is maintaining geometric and photometric calibration. Device characteristics continually change with temperature and supply voltage variations, the aging of projector bulbs, and physical perturbations. Recently we developed a stochastic autocalibration method that uses cameras and whatever imagery the user is projecting to refine a display surface estimate while the system is in use. At every frame we use the projector frame-buffer contents and the current display-surface estimate to predict the location of an image feature in the camera's image plane. We then search nearby in the actual camera image for the corresponding feature and correct the display surface estimate using the difference. We repeat this on a single pixel per frame basis, visiting the pixels repeatedly in a pseudorandom fashion. The longer the system is used, the better the estimate of the display surface. We're working on extending this to the external parameters of the projectors and cameras, and eventually photometric parameters.

Pieces of the puzzle

The results of our work toward the Office of the Future aren't likely to be directly useful in isolation. We aren't attacking some related problems that need solving to realize a practical and useful system. Happily, others are working on those problems, and it's reasonable to expect the work to converge at some point. For example, in addition to multiprocessing and load balancing for 3D graphics,[4] a software framework that encapsulates device-specific details (tracking, stereo, and renderings) would let developers concentrate on applications. The CAVE Library solves this nicely for certain 3D display systems. Similarly, researchers at Princeton have developed a Displaywall Toolkit,[4] and researchers at Stanford are working on interfaces as part of their Interactive Workspaces project.[14]

Today's desktop window systems (X, MacOS, and Windows) deal nicely with simple multiple-channel displays, organizing the desktop as non-overlapping, rectangular plots of planar pixel turf. Of course, significant challenges remain in designing a window server that makes transparent the geometric and photometric rendering complexities of multiprojector and multisurface rendering. For example, questions remain in how to address how the logical desktop should relate to disjoint physical areas of display around the office.

If you're interested in our vision of the Office of the Future and want to do something to help us—and yourself—go out and buy some projectors! Set them up in your office, and you too will never want to go back to the office of yesterday. ■

Acknowledgments

We acknowledge Gary Bishop for his ideas and practical suggestions. We thank Brent Seales from the University of Kentucky for his collaboration during his sabbatical here at UNC last year and for bringing along one of his top students, Michael Brown (coauthor). We also acknowledge our other team members at UNC, listed at http://www.cs.unc.edu/Research/stc/office/index.html. Finally, we thank Andrei State for his visionary sketch (Figure 1).

This work is supported by NSF Cooperative Agreement No. ASC-8920219: "Science and Technology Center for Computer Graphics and Scientific Visualization"; the "National Tele-Immersion Initiative" sponsored by Advanced Network and Services, Inc.; DARPA grant No. N66001-97-1-8919: "Three Applications of Digital Light Projection for Tiled Display and 3D Scene Capture"; and "Automatically Reconfigurable Arrays of Projectors for Wide-Area Display," Lawrence Livermore National Laboratory High Performance Computing Group, Contract B504967. We also thank Intel Corporation for its generous equipment donations as part of Technology for Education 2000.

References

1. R. Raskar et al., "The Office of the Future: A Unified Approach to Image-Based Modeling and Spatially Immersive Displays," *Computer Graphics,* Ann. Conf. on Computer Graphics and Interactive Techniques, M.F. Cohen, ed., ACM Press, New York, 1998, pp. 179-188.
2. F.P. Brooks, Jr., *The Mythical Man-Month: Essays on Software Engineering*, Twentieth Anniversary Edition, Addison-Wesley, Reading, Mass., 1995, pp. 194-195.
3. C. Cruz-Neira, D.J. Sandin, and T.A. DeFanti, "Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE," *Computer Graphics*, Vol. 27, Ann. Conf. on Computer Graphics and Interactive Techniques, Siggraph Conf. Proc., ACM Press, New York, Addison-Wesley, Reading, Mass., 1993, pp. 135-142.
4. R. Samanta et al., "Load Balancing for Multi-Projector Rendering Systems," *Proc. 1999 Eurographics/Siggraph Workshop on Graphics Hardware*, ACM Press, New York, Addison-Wesley, Reading, Mass., Aug. 1999, pp. 107-116. (See also http://www.cs.princeton.edu/omnimedia/index.html, cited 20 Oct. 1999.)
5. J. Underkoffler, B. Ullmer, and H. Ishii, "Emancipated Pixels: Real-World Graphics In The Luminous Room," *Computer Graphics*, Ann. Conf. on Computer Graphics and Interactive Techniques, A. Rockwood, ed., Siggraph Conf. Proc., ACM Press, New York, Addison-Wesley, Reading, Mass., 1999, pp. 385-392.
6. J. Underkoffler, "A View from the Luminous Room," *Personal Technologies*, Vol. 1, No. 2, June 1997, pp. 49-59.
7. R.A. Weale, *Focus on Vision*, Harvard University Press, Cambridge, Mass., 1982, p. 15.
8. J. Laycock and R. A. Chorley, "The Electro-optical Display/Visual System Interface: Human Factors Considerations," Royal Aircraft Establ., Farnborough, UK AGARD-AG-255, Oct. 1980.
9. R.E. Bedford and G. Wyszecki, "Wavelength Discrimination for Point Sources," *J. Optical Society of America*, No. 48, 1958, pp. 129-135.
10. R. Raskar et al., "Efficient Image Generation for Multiprojector and Multisurface Displays," *Proc. 9th Eurographics Rendering Workshop*, G. Drettakis and N. Max, eds., Springer-Verlag, Vienna, 1998, pp. 139-144.
11. M. Segal, "Fast Shadows and Lighting Effects Using Texture Mapping," *Computer Graphics*, Ann. Conf. on Computer Graphics and Interactive Techniques, Siggraph Conf. Proc., ACM Press, New York, Addison-Wesley, Reading, Mass., 1992, pp. 249-252.
12. R. Raskar, "Multi-Projector Displays Using Camera-Based Registration," *Proc. Conf. on IEEE Visualization 99: Celebrating 10 Years*, ACM Press, New York, 1999, pp. 161-168.
13. A. Majumder et al., "Color Matching of Projectors for Large Tiled Displays," to appear in *Proc. IEEE Visualization 2000,* IEEE Computer Society, Los Alamitos, Calif.
14. G. Humphreys and P. Hanrahan, "A Distributed Graphics System for Large Tiled Displays," *Proc. IEEE Visualization 99*, IEEE Computer Society, Los Alamitos, Calif., 1999.

***Greg Welch** holds the position of Research Assistant Professor in the Computer Science Department at the University of North Carolina at Chapel Hill. He received a BS in electrical technology from Purdue University, and an MS and PhD in computer science from the University of North Carolina at Chapel Hill. He is a member of the IEEE Computer Society and the Association for Computing Machinery (ACM).*

***Henry Fuchs** is Federico Gil Professor of Computer Science. Fuchs started the Pixel-Planes/PixelFlow high-performance image generation team and has led teams that developed displays ranging from augmented reality head-mounted stereo displays to true 3D vari-focal mirror displays. His awards include the 1992 ACM Siggraph Computer Graphics Achievement Award, the 1992 NCGA Academic Award, and the 1997 MMVR Satava Award. Fuchs was elected to the National Academy of Engineering and the American Academy of Arts and Sciences in 1997.*

***Ramesh Raskar** is a PhD candidate in computer science at the University of North Carolina at Chapel Hill. He received a BS in electronics and telecommunication from the Government College of Engineering, Pune, India, and an MS in computer science from the University of Iowa. His research interests include graphics and computer vision.*

***Michael S. Brown** is a PhD candidate in computer science at the University of Kentucky, Lexington. He received a BS in computer science from the University of Kentucky and spent a year and a half at the University of North Carolina at Chapel Hill as a visiting PhD student. Brown's research interests include graphics, computer vision, and digital libraries.*

***Herman Towles** holds the position of Senior Research Associate in the Computer Science Department at the University of North Carolina at Chapel Hill. He received a BS and MS in electrical engineering from the University of Memphis, Tennessee.*

Readers may contact the authors at the Office of the Future Project, University of North Carolina at Chapel Hill, Dept. of Computer Science, Chapel Hill, NC 27599-3175, e-mail {welch, fuchs, raskar, towles}@cs.unc.edu. Contact Brown at the University of Kentucky, Dept. of Computer Science, Lexington, KY 40506-0108, e-mail mbrown@dcs.uky.edu.