Omnidirectional Vision and Tele-Presence

Henry Fuchs

Recent years have witnessed dramatic increases in the number of research groups & countries around the world interested in the notion of Tele-Immersion, or Tele-Presence, creating for the participant the convincing illusion of being in a physical place different from the his/her location. Among the most significant components of this "sense of presence" is the ability to look around the remote environment, as naturally as if one were there. Although the notion of omnidirectional vision within Tele-Presence is in many ways beyond the capability of even today's most advanced systems, it remains a key goal of research projects.

Conceptual sketch of the Office of the Future, by artist Andrei State.

UNC, in collaboration with research groups at Brown University, the University of Pennsylvania, Advanced Network & Systems, and the University of Utah, have been working for most of the last decade on issues relating to omnidirectional vision within Tele-Immersive presence. While the general capability is still beyond reach, we have experimented with, identified, and pursued intermediate goals to explore different aspects of different, overlapping subsets of this capability, in order to determine which are most effective, and to single out directions of future research, and future systems.

A major choice within the enormous space of possibilities is the limiting number of participants and sites involved in a Tele-Immersive system. Although the though of a limitless number of sites and participants per site is tempting (some researchers are pursuing "multi-player" environments in which thousands of people can participate, as in the conferences & cafes of the current time), we have been aiming at a system in which the number of users and sites is quite few; a system wherein the locations are ubiquitous office environments or small meeting rooms, where we believe the vast majority of useful work takes place both today & in the foreseeable future.

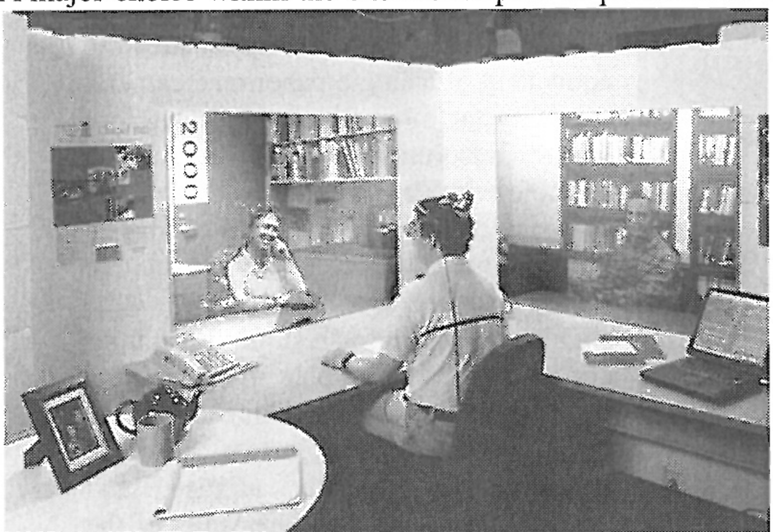

Graduate research assistant Ruigang Yang sits in the current prototype system of the Office of the future, facing two remote collaborators at different physical locations in the country.

Within the design space of our current system, we have explored the design and prototype implementation of systems varying from single individuals in offices, to small group (say, six) in small meeting rooms, and scaling the number of sites from two to four. We have performed experiments with collaborators under the umbrella of two organizations: initially, the NSF Science & Technology Center (STC) for Computer Graphics & Scientific Visualization, a 5-site compendium in which the emphasis is on Telecollaboration, and our closest colleagues (in the application of collaborative design & fabrication of mechanical components & assemblies) are Brown University and the University of Utah; and, for past three years, the USA National Tele-Immersion Initiative (NTII), funded by Advanced Network & Services (ANS) and directed by Chief Scientist Jaron Lanier. (Our closest collaborators in the latter are the University of Pennsylvania, our fellow STC site, Brown University, and ANS itself.

Within the restricted design space of few sites and users, we have explored several different options. The first useful one was a 360-degree panoramic vision scenario, in which an individual in a private office wishing to participate in a meeting in a different location is able to do so. For this experiment, in which our principle research colleagues were initially the University of Utah group within the STC, we constructed a camera cluster of 10 – 12 cameras, which represented the remote participant's field-of-vision in the meeting room, and thus "replaced" that person at the discussion table. The camera cluster (following design of Vic Nalwa at Bell Labs, the details of which can be seen at http://www.fullview.com) possessed the property that the center of projection of all cameras was at the same physical location (to within manufacturing and calibration tolerances), and therefore represented the equivalent of a single panoramic camera, one with far more pixels (and therefore higher resolution) than would be possible with the same NTSC resolution of a single camera. The rates of images of the cameras, since data from each one was streamed parallel, were simultaneously refreshed at 30 frames per second. The display in the remote participant's office was achieved via five ceiling-mounted projectors whose imagery was casually aligned, and generally covered the bare walls of the office. Calibration was achieved before the meeting, by moving the camera cluster into the office, and positioning it precisely at the location where the individual's chair


The camera cluster, as seen from various positions in the meeting room.

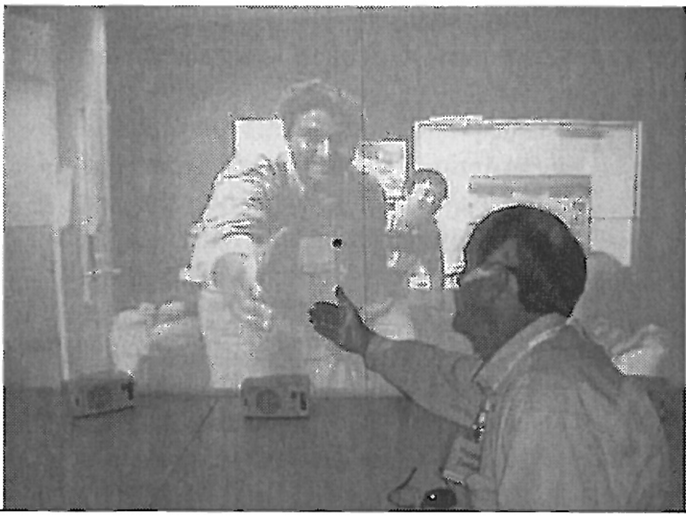

OOTF project led engineer Herman Towles, in the office environment, shakes hands with research assistant Ramesh Raskar, in the meeting room.

would be, and then by projecting individual pixels of each of the projectors and cameras, observing which camera- and projector-pixel locations would coincide, and creating an association between these matching camera- and projector-pixels. Later, during the course of the meeting, these associations were used to determine, for each camera pixel, which one or more projector pixels it mapped onto. The resulting prototype system gave a convincing illusion of omnidirectional visual presence, albeit for non-stereo monocular vision. As such, several users, including author, preferred to use the system wearing spectacles in which the non-dominant eye was darkened, so as to prevent any conflicting illusion provided by stereo views of the real office and walls, even though the images projected onto them was monocular.

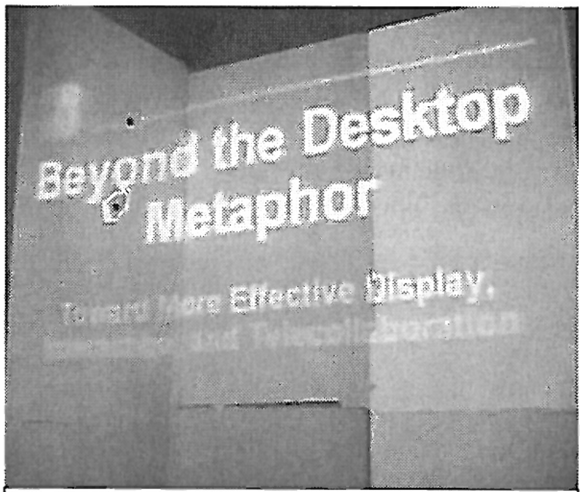

The view of a presentation in the meeting room, as seen from the office environment. Note that the differences in geometries of the rooms lead to the distortion correction of the text on the meeting room's display wall; it is seen perspectively correct from the calibration position.

This and similar arrangements, with two-dimensional panoramic images from the camera cluster and an arbitrary number of displaying projectors, have advantage that they are easy to set up, to calibrate, and appear to work with an arbitrary number of both cameras and projectors. However, on the list of disadvantages are that differing geometries of the office end and the meeting room can be distracting when displayed monocularly, when the distance to the display surface is remarkably different from the distance to the corresponding object being displayed. A further disadvantage is the increasing distortion that is inevitable when the viewer moves away from calibration position in the office environment. Extensibility (making the system truly bi-directional) is a further significant problem that has yet to be solved.

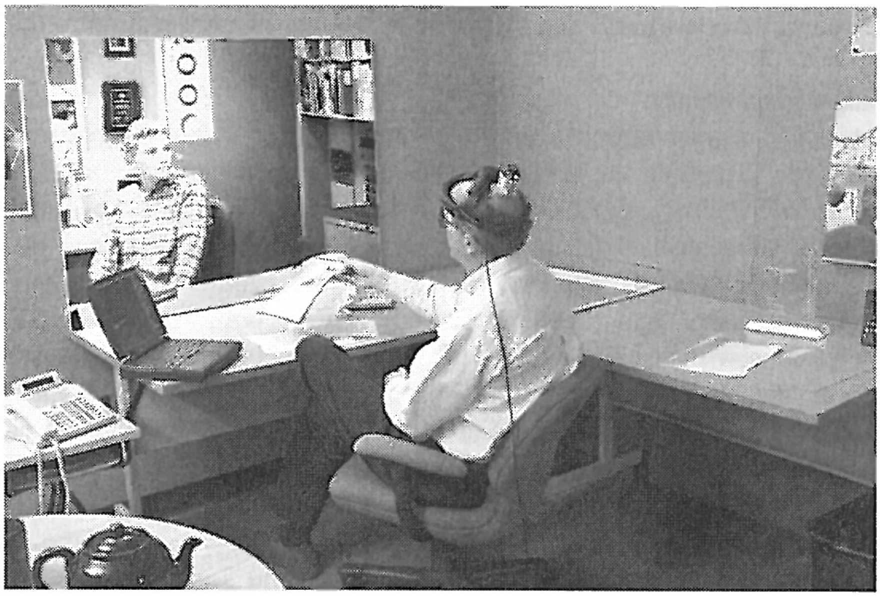

For all these reason we have focused in the last several yrs on three-dimensional acquisition and display in the meeting room and office environments, together with our closest collaborators at the University of Pennsylvania, as well as others within NTII. Due to the extremely demanding problem of extracting precise robust depth values for each camera pixel, our current operating prototypes emphasize a "through the window" paradigm (see photo, above right), rather than the 360-degree omnidirectional visualization. This is achieved with seven cameras, all aimed at the single participant at each of the remote "office" sites; three-dimensional surfaces are extracted from the images obtained from these cameras at 1 –2 frames per second, and are merged with pre-scanned three-dimensional models of office environments, which were attained using the UNC-developed Laser Range Finder & camera system of Prof Lars Nyland. Display for this data is achieved with polarized light seen through passive stereo glasses, while tracking is done using the UNC-developed Hiball tracker, which is composed of a cluster of optical sensors, housed in a golf ball-size device worn on the user's head, that observe selected infrared illuminated LEDs.

A combination of the above-described technologies is what we look forward to in order to increase performance of true three-dimensional acquisition and display that will give omnidirectional head-tracked three-dimensional views, affording the best of panoramic two-dimensional views and the three-dimensional "through-the-window" views that are currently being improved upon. Such displays in one's office environment will not only enable a dramatically increased sense of presence and facilitate a richer collaboration among distant participants, but also enable a new class of tools and capabilities for local collaboration within the office, and for individual human-computer interaction, and work done on one's PC.

In his 1965 IFIP paper, Ivan Sutherland related that the ultimate display for a computer would serve as a "window" to the world inside it. With tele-immersion, that window can now engulf the user, and allow him/her to experience the real world, either inside the computer or out, or some combination of real and synthetic, whatever is most suiting to the task at hand.

Acknowledgements:
Most of this work resulted from collaborations first within the 5-site NSF Science & Technology Center for Graphics & Visualization, whose members are at Brown University, California Institute of Technology, Cornell University, the University of Utah, and UNC Chapel Hill. Especially close collaborators have been the teams at Brown, lead by Andries van Dam, and the team at Utah, lead by Richard Riesenfeld and Elanie Cohen. Most of the work relating to computer vision has been a collaboration with the GRASP Lab at University of Pennsylvania,with Ruzena Bajcsy (the lab 's director) and with Kostas Daniilidis. Much of the work of the past three years has been performed as part of the National Tele- Immersion Initiative, funded by Advanced Network & Services, and led by its chief scientist, Jaron Lanier. Much of the UNC work has been led by Greg Welch and by Herman Towles (the UNC project 's lead engineer). Other current members of the UNC team include Wei-Chao Chen, Ruigang Yang, Sang-Uok Kum, Andrew Nashel, Scott Larsen, Greg Coombe, David Gotz, Justin Hensley, Aditi Majumder, Jim Mahaney, Kurtis Keller, and John Thomas. Finally, we thank Caroline Green for pulling together the text and photos for this extended abstract on very short notice.

Bibliography

Chen, Wei-Chao, Herman Towles, Lars Nyland, Greg Welch, and Henry Fuchs. Toward a Compelling Sensation of Telepresence: Demonstrating a portal to a distant (static) office. Proceedings of IEEE Visualization 2000,Salt Lake City, UT.

Nyland, Lars, David McAllister, Voicu Popescu, Chris McCue, Anselmo Lastra, Paul Rademacher, Manuel Oliveira, Gary Bishop, Gopi Meenakshisundaram, Matt Cutts, and Henry Fuchs. The Impact of Dense Range Data on Computer Graphics.Proceedings of Multi-View Modeling and Analysis Workshop (MVIEW99), (Part of CVPR99), (Fort Collins, CO), June 23-26, 1999.

Raskar ,Ramesh, Greg Welch, Matt Cutts, Adam Lake, Lev Stesin and Henry Fuchs. The Office of the Future: A Unified Approach to Image-Based Modeling and Spatially Immersive Displays. Proceedings of ACM SIGGRAPH 1998, Orlando FL.

Welch, Greg, Henry Fuchs, Ramesh Raskar, Michael Brown, and Herman Towles. (2000). "Projected Imagery In Your Office in the Future," IEEE Computer Graphics and Applications, July/August 2000: 62-67.

Welch, Greg, Gary Bishop, Leandra Vicci, Stephen Brumback, Kurtis Keller, D'nardo Colucci. 1999. "The HiBall Tracker: High-Performance Wide-Area Tracking for Virtual and Augmented Environments," Proceedngs of the ACM Symposium on Virtual Reality Software and Technology 1999 (VRST 99), University College London, December 20-22, 1999.