Scape: Supporting Stereoscopic Collaboration in Augmented and Projective Environments

Hong Hua
University of Hawaii at Mãnoa

Leonard D. Brown and Chunyu Gao
University of Illinois at Urbana-Champaign

Our collaborative augmented reality system for multiple users simultaneously creates outside-in workbench views and inside-out life-size walk-through views, bridging the virtual and augmented reality paradigms.

With the trend toward scientific and economic globalization, users and researchers alike have a great interest in tools and infrastructures that let multiple users at local and remote sites effectively perform collaborative tasks in an intuitive manner. Ideally, these users should behave in the same way as with face-to-face collaboration.

There already exists a large body of research efforts in the areas of computer-supported collaborative work and multiuser collaboration and telecollaboration infrastructures to facilitate collaborative interfaces.[1,2] Recent efforts have attempted to develop tools and infrastructures to support local and remote collaboration in 3D virtual and augmented environments.[3,4]

Here we introduce a collaborative infrastructure—stereoscopic collaboration in augmented and projective environments (Scape)—shown in Figure 1a. Our goals for the system include enabling users to

- view the task from their individual perspectives,
- have equal and natural access to the task,
- perceive the presence of local and distant members,
- communicate by both verbal and nonverbal means, and
- dynamically switch focus between the shared workspace and interpersonal communication space.

The system provides a shared space in which multiple users (a theoretically arbitrary number) can concurrently observe and equally interact with a 3D synthetic environment, viewing the space from their individual perspectives, while face-to-face cooperation is preserved. The system simultaneously creates an outside-in bird's eye view (Figure 1b) and an inside-out walk-through view (Figure 1c) of a 3D data set through the workbench and the immersive room representations. Scape seamlessly merges the paradigm of virtual reality with that of augmented reality in a single system and provides multimodality interface devices and unique widget interfaces to facilitate collaborative work in 3D environments.

Scape's core display components are multiple custom-designed prototypes of head-mounted projective displays (HMPD). Scape's ubiquitous projective environment mainly consists of a 3×5 foot workbench, a $12 \times 12 \times 9$ foot four-walled mural display surrounding the workbench, and selected object surfaces in the room. The multi-modality interface devices include a Hiball 3000 optical tracking system by 3rdTech (http://www.3rdtech.com), a Flock of Birds (FOB) magnetic tracking system by Ascension Technology (http://www.ascension-tech.com), a wireless DataGlove by 5DT (http://www.5dt.com), and a set of custom-designed widgets to facilitate interaction and collaboration tasks in 3D environments. A custom-designed cross-platform application-programming interface (API) allows high and medium-level controls over the Scape workspace to enable augmented interaction and collaboration.

To more accurately reflect a user's real experience with the system, we captured the photographs included in this article—except those conceptual simulations and illustrations and hardware implementations—through the eye position in the display.

Background and related work

The two main areas related to our work are 3D col-

Published by the IEEE Computer Society

laborative interfaces and infrastructures, and the HMPD technology.

3D collaborative interfaces and infrastructures

Several different approaches for facilitating 3D collaborative work exist. An attractive and yet expensive solution is projection-based spatially immersive displays. Examples of these include CAVE-like systems[5] or the responsive workbench,[6] which let users view stereoscopic images with LCD-shutter glasses. In these displays, users can see each other and therefore preserve face-to-face communication. Unfortunately, these systems render images from a single user's viewpoint, and therefore the stereo images are perspective correct only for the tracked user. The other nontracked users will experience both point-of-view and motion distortion. This limitation makes it impossible to provide all users with a visually consistent environment.

Several efforts attempt to overcome this limitation. Agrawala proposed the two-user responsive workbench in which two people simultaneously view individual stereoscopic image pairs from their own viewpoints by using four frame buffers.[7] This approach sequentially renders two pairs of stereoscopic images at one-quarter the display frame rate. The system cuts the display frame rate in half—compared to the single viewer approach—for each user, causing a noticeable flicker (for example, 30 Hz for each eye with ordinary display hardware having a 120-Hz maximum frame rate). Kitamura[4] proposed an alternative solution, IllusionHole, in which three or more people simultaneously observe individual image pairs from independent viewpoints without the sacrifice of frame rate. The IllusionHole display consists of a normal bench display and a display mask, making each user's drawing areas invisible to others. However, this system has a limited maximum number of users, a limited movement space for each user, and a small viewing area for each user.

Augmented, or mixed, reality interfaces superimpose graphics and audio onto the real world, combining the advantages of virtual environments and enhanced reality. Thus, they facilitate the development of collaborative interfaces while supporting seamless interaction with the real world, reducing the functional and cognitive seams.[8]

Some efforts have pursued the collaborative augmented-reality metaphor. For example, Rekimoto used tracked handheld LCD displays in a multiuser environment and miniature cameras attached to the LCD panels, enabling virtual object compositing on video images of the real world.[9] Billinghurst[8] and Szalavari[10] proposed using see-through HMDs with head and body tracking in a collaborative interface, allowing multiple local or remote users to work in the real and virtual worlds simultaneously. Bimber and colleagues proposed the concept of the virtual showcase, which lets two to four tracked users interact with the virtual content of the showcase while maintaining the augmentation of the virtual contents with real artifacts.[3]

HMPD technology

Augmented reality applications predominantly use optical or video see-through HMDs. The HMPD[11] is an emerging technology that lies on the boundary of conventional HMDs and CAVE-like projective displays.[5,6] Unlike a conventional, optical see-through HMD, an HMPD (illustrated in Figure 2a, next page) replaces eyepiece-type optics with projective lenses. An image on the miniature display—located beyond the focal point of the lens, rather than between the lens and focal point as in a conventional HMD—is projected through the lens and thus forms a magnified real image. A beam splitter reflects the real image outward into the object space, rather than directly toward the eye space. Unlike a conventional projector, an HMPD replaces a diffusing screen with a retro-reflective screen. The retro-reflec-

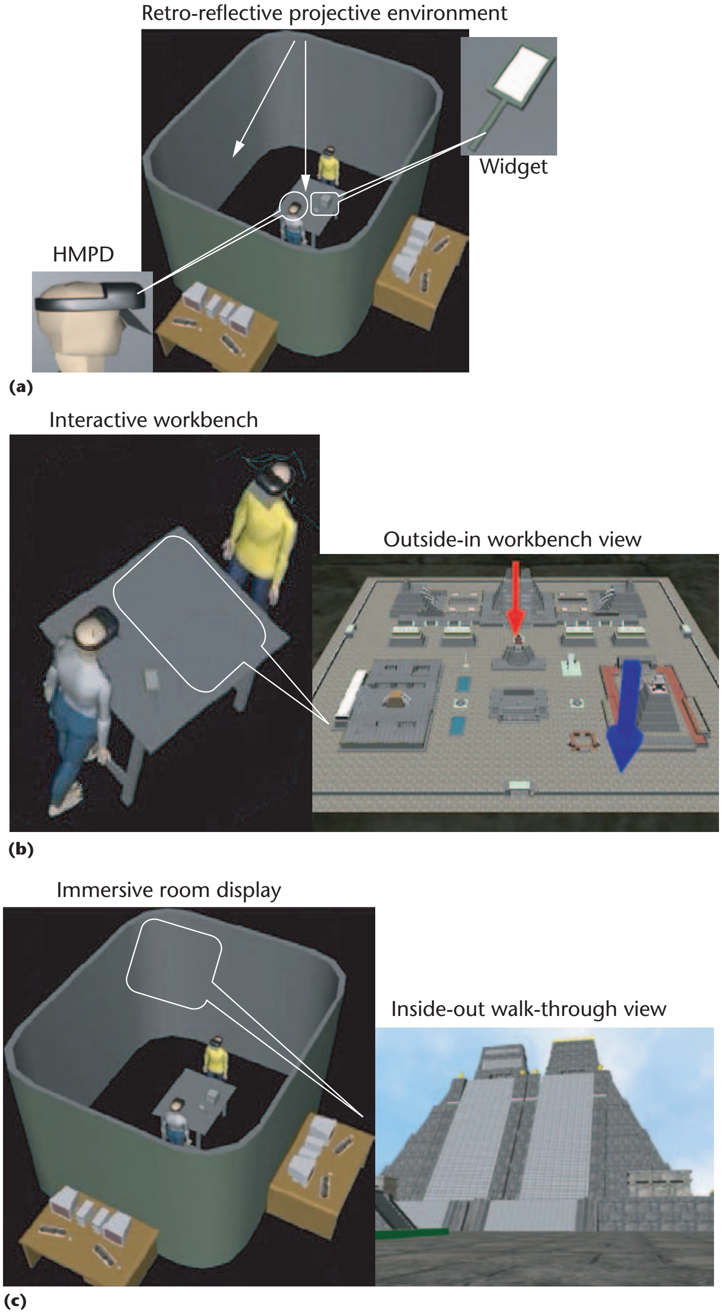

1 Scape (a) concept and components, (b) simulation of outside-in bench view, and (c) simulation of inside-out walkthrough view.

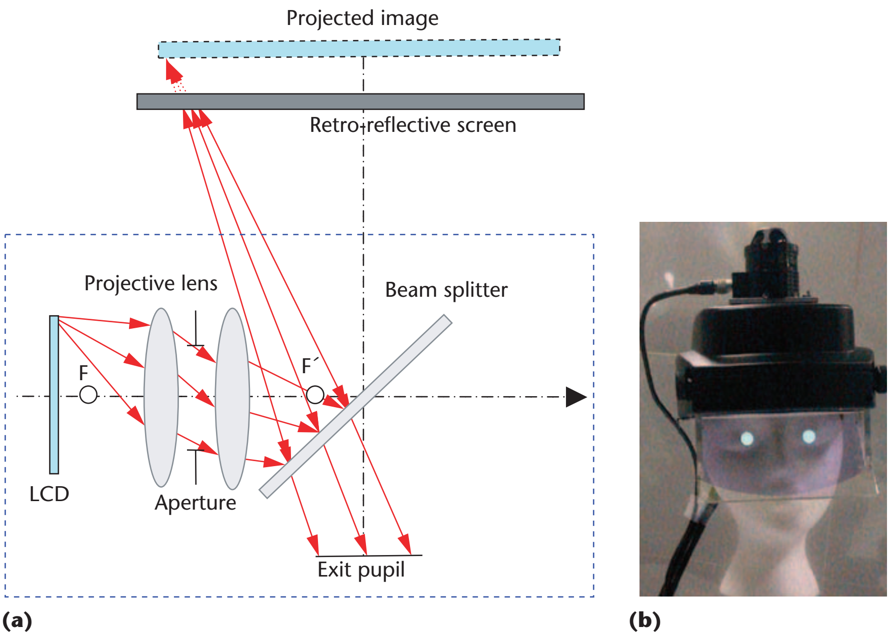

2 Head-mounted projective display: (a) conceptual illustration and (b) compact prototype.

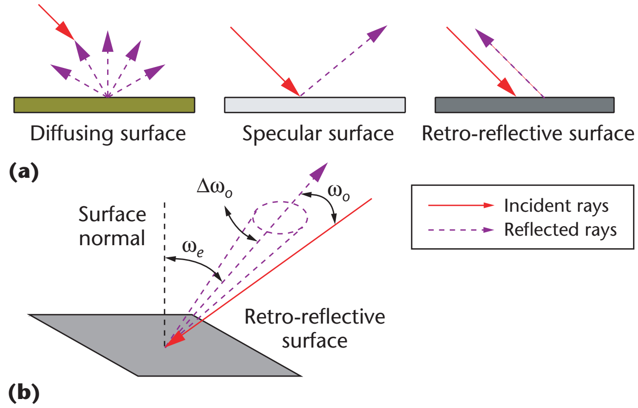

3 (a) Difference in retro-reflection from diffusing and specular reflections, and (b) Characterization of retro-reflective material.

tive screen reflects back on itself a ray hitting the surface at any angle; exactly in the opposite direction (see Figure 3a). Therefore, the projected rays are retro-reflected back to the exit pupil, where the eye is positioned to observe the projected image.

Researchers have demonstrated that the HMPD concept yields 3D visualization capabilities with a large field of view (FOV), lightweight and low distortion optics, and correct occlusion of virtual objects by real objects.[12-14] Studies have recognized it as an alternative solution to a wide-range of augmented applications.[15,16] In collaboration with colleagues at the University of Central Florida and the Michigan State University, we have developed a compact prototype, shown in Figure 2b, using a custom-designed lens.[14] The prototype achieves 52 degrees FOV and weighs 750 grams. Our work on the Scape collaborative infrastructure relies on this prototype.

Scape concept

Applying retro-reflective surfaces in the workspace and integrating multiple, head-tracked HMPDs and interface devices can extend single-user HMPD technology to a multiuser mode. The HMPD technology supports collaborative applications by

- enhancings the real world with 3D computer-generated information,
- providing the capability to create an arbitrary number of individual viewpoints with nondistorted perspectives and without cross talk among users, and
- retaining natural face-to-face communication among local participants.

Shared workspace for 3D collaborative environments

One example of shared workspaces is a retro-reflective workbench equipped with multiple head-tracked HMPDs and multimodality interaction devices (see Figure 1b). We refer to this configuration as the *interactive workbench*, through which multiple participants, wearing head-tracked HMPDs, view and manipulate a 3D data set superimposed on its physical counterpart placed on the bench. We might deliberately coat the physical objects with retro-reflective film. Each user has an individualized perspective of the synthetic environment and retains the privilege of equal access. If registration is appropriately achieved, when two users point to the same part of the data set, their fingers shall touch, which does not happen in CAVE systems. The number of users is unlimited, depending on availability of display and computing resources. Additionally, displaying multiple independent views offers the intriguing possibility of presenting different aspects or levels of detail of a shared environment in each user's view.

Screen shape options include many possibilities, such as clothing or other nonregular shaped surfaces. For example, we made a cylindrical display (see Figure 4) in which we installed the retro-reflective cylinder on a revolving platform. An angular sensor measures its rotation. Together with multiple head-tracked HMPDs and interface devices, this platform provides a shared medium for users to equally interact with a visualized subject intuitively. The cylindrical screen's shape and location don't distort the perceived shape or location of visualized objects.

Our workbench and cylindrical display configurations provide outside-in semi-immersive perspectives of a 3D data set, as do the responsive workbench[7] and Illusion-Hole;[4] although users can only explore the data set from an exocentric perspective. Using the HMPD technology, it's also possible to create a CAVE-like room-sized space with inside-out immersive perspectives, such as a walk-through. Therefore, Scape combines the functionality of an interactive workbench with that of a room-sized immersive display to create outside-in and inside-out perspectives simultaneously (see Figure 1a).[17]

A microscene is registered with the physical workbench and objects placed on the bench, while the surrounding virtual environment can be considerably larger than the physical room. A user can rely on the workbench environment to conduct augmentation tasks while immersed in a life-size virtual environment visualized through the room. Switching between miniature and immersive views occurs seamlessly by maintaining a unique world-coordinate system for all users and a continuous viewing system for each user between the two different scales of visualization. Moreover, the miniature visualization on the workbench can also represent a different level of detail from that of the walk-through.

Design considerations and limitations

The Scape implementation is mainly affected by the characteristics of retro-reflective materials. In HMPDs, retro-reflection indicates that user perception of image shape and location is independent of the retro-reflective screen shape and location (see Figure 2a). Practically, however, a retro-reflective material can only work well for limited angles. We use three angles in describing the performance of retro-reflective materials: entrance angle ω_e, observation angle ω_o, and cone angle $\Delta\omega_o$ (see Figure 3b). The entrance angle specifies the angular range in which a material remains highly retro-reflective. The observation angle gives the angle difference between an incident ray and the reflected ray. The cone angle specifies a reflected beam's angular width. Wide entrance angle, zero observation angle, and zero cone angle are preferred in an HMPD. Imperfect reflective properties have a direct or indirect impact on imaging characteristics and quality, and thus affect certain aspects of the Scape design.

Screen shape. In contrast to a diffusing screen, we can tailor a truly retro-reflective screen into arbitrary shapes without causing image blurring and introducing distortions to virtual images. Practically, retro-reflection is only dominant within ±40-degree entrance angles for available samples, which sets up a constraint on screen shape. As incident angles increase, the drop of reflectivity results in gradual vignetting effects on image brightness. Therefore, for a given visual field, we predict that a concave shape can improve image brightness of marginal fields, but a convex shape will worsen the brightness. When designing a multiwall room display, round corners—rather than squared corners as in CAVE systems—minimize the gradual drop in image luminance because of the decrease of reflectivity when incident angle θ increases (see Figure 5).

Screen position and room size. We can apply a truly retro-reflective screen to any possible location in the physical space without causing image blurring or degrading image quality. A screen with a nonzero observation angle leads to slightly blurred images if the screen is far away from the focused image plane, which sets up constraints on room size. In this case, it is desirable to

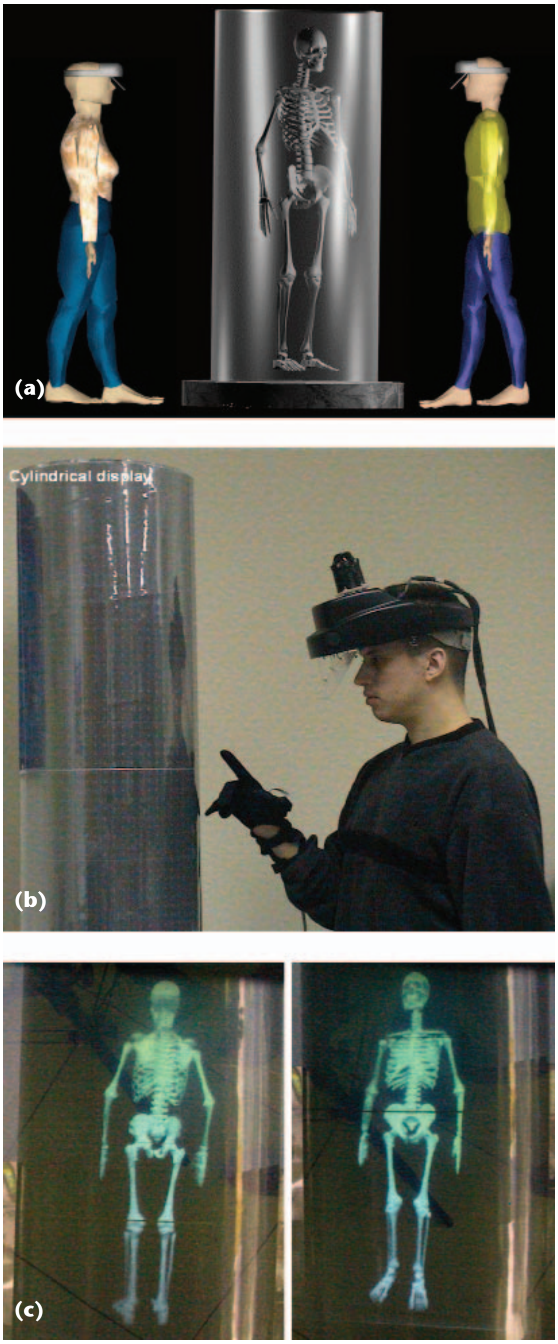

4 Cylindrical display for collaborative visualization: (a) conceptual illustration, (b) prototype implementation, and (c) two individual views through the HMPDs.

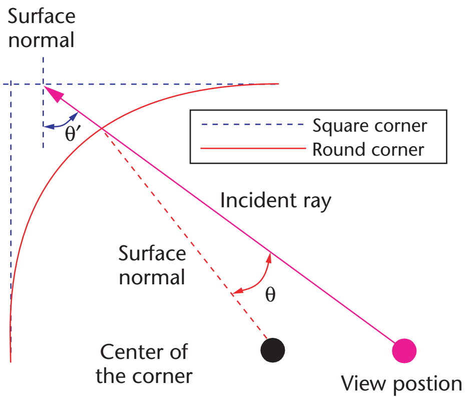

5 Screen shape design: round versus square corners.

have a zero-observation angle. The few samples we tested have up to 3 degrees of observation angle and the best performance material available is less than 0.5 degrees.

Cross talk. For a truly retro-reflective screen, the stereo pair of images projected for the left and right eyes, or different channels of images projected for individual users, are naturally separated. Practically, a nonzero cone angle can cause cross talk if a screen is too far away from a user. The cone angle of the reflected beam is as narrow as 0.4 degree for the sample we used, thus cross talk between the left and right eyes can possibly occur when the user is more than 9 meters away from a screen. However, for such a distance no cross talk is present if two users stand side by side. Therefore, with this sample, 9 meters is the upper bound of a walk-through display. In this case, it's possible to generate as many unique perspectives as needed for each user, without introducing cross talk from any other participants. The performance of different material technologies varies.

Functional augmentation. The optical see-through capability in HMPD allows augmenting the real world with computer-generated imagery, while preserving the direct awareness of the user's surrounding environment. This can occur without tradeoffs in visual quality, as would occur with video see-through HMDs. The challenge, however, is to achieve geometric registration (size, depth, and so on) and physical attribute matching (resolution, reflectance, and so on) of virtual objects with respect to their physical counterparts. Furthermore, users can only perceive virtual objects when they look toward surfaces coated with the retro-reflective materials. Thus, an HMPD can intrinsically provide correct occlusion cues of computer-generated virtual objects by real objects, and users can naturally switch their focus of interest between the real and virtual workspaces. This does not happen in conventional optical see-through HMDs. The challenge is, however, that retro-reflective materials must be deliberately applied to physical surfaces for a particular application. Otherwise, virtual objects will erroneously disappear when a virtual object is intentionally floating between a real object and the user. This challenge might impose limitations on the scope of applications.

FOV. Wide FOV is desirable for wearable displays. The feasible FOV of an HMPD is constrained by that of the projective optics and is limited by the maximum entrance angle of retro-reflective materials. In fact, to avoid significant degradation of luminance in peripheral visual fields, the maximum entrance angle of the screen used in a system sets up an HMPD's upper-bound FOV. For example, for the tested samples, the HMPD's FOV should be equal to or less than 80 degrees, if we assume a flat retro-reflective screen.

Environmental lighting. The lack of brightness is a common problem in LCD-based optical see-through HMDs. This problem is aggravated in HMPDs because light passes through the beam splitter multiple times, leading to the loss of at least 75 percent of the light. Thus, we must dim environmental lighting to a minimal level, which might make reading difficult. The lack of image luminance also limits viewing distance, and therefore the size of the room.

Portability and other issues. Our current prototype is not yet as light and compact as LCD shutter glasses used in CAVE systems. This limitation certainly causes mobility concerns in a collaborative workspace. On the other hand, an HMPD is more compact and lighter than conventional HMDs due to projective optics' characteristics. For example, the proprietary lens design for our prototype is about 6 grams per eye, which leaves open the possibility of designing more compact prototypes. The necessity of a retro-reflective screen imposes a limitation on face-to-face interaction with wall displays but not with workbench displays, similar to the shadow problem in a CAVE environment. We might partially or fully solve this limitation in the near future. In the meantime, retro-reflective screens can be portable; minimal requirements for wall alignment during the installation process significantly reduce the system's cost and construction time.

Scape implementation

With the previous design considerations in mind, we describe a preliminary implementation of the Scape hardware and software.

Hardware

The Scape display environment currently consists of a workbench and a four-walled immersive cage made from retro-reflective film, multiple head-tracked HMPDs, multimodality interface devices, computing facilities, and networking tools.

In our implementation, the workbench is a tabletop providing a 3 × 5 feet surface area and coated with removable retro-reflective film from 3M. We've placed this in the middle of a room, about 3 feet above the floor. The four-walled immersive cage surrounds the workbench setup. Figure 6a specifies its shape, which is composed of four 6-foot-wide flat walls and four arch corners with 3-foot radii. The walls are 9 feet high. The round corners minimize the gradual drop in image luminance, as explained previously.

We assembled the walls and corners before coating them with retro-reflective film. We designed one corner to act as a hinged door. The enclosure allows full control of lighting. Coating the floor and ceiling with the film would enable a six-wall cage. We use a Hiball3000 sensor by 3rdTech for head tracking purposes, therefore we installed our ceiling with a 12 × 12 foot array of LED strips. Because of the minimal requirements on wall alignment and the low cost of the film, we built the reflective cage for much less than the cost of building a CAVE. Figure 6b shows Scape's physical setup.

An SGI Octane workstation with a R12000 processor and a four-channel VGA output option drives the HMPDs. Each two-channel set outputs a pair of stereoscopic images to an HMPD. A Hiball3000 optical tracker or an Ascension FOB magnetic tracker detects each user's head position and orientation. The system generates the stereoscopic image pairs without distortion for each user and according to their individual viewpoints. Currently, only one tracked HMPD is regularly available for experiments, but we could easily expand the system to support an arbitrary number of users when hardware resources are available. We simulate a second viewpoint to perform preliminary experiments, and we are designing and building a second prototype.

Interface devices and widgets

Scape employs a set of generic interface devices to manipulate and interact with virtual environments. Besides the typical devices, such as mouse and keyboard, an Ascension FOB magnetic tracker tracks moving objects such as hand or interface widgets. A 5DT DataGlove manipulates 3D virtual objects on the bench and navigates the walk-through immersive environment in the room. We also developed a set of unique augmented devices or widgets to facilitate interaction and collaborative interfacing.

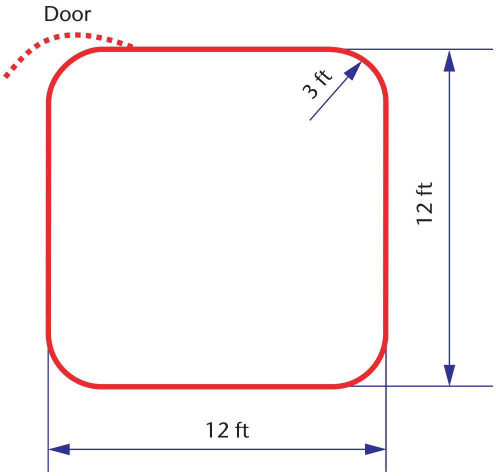

(a)

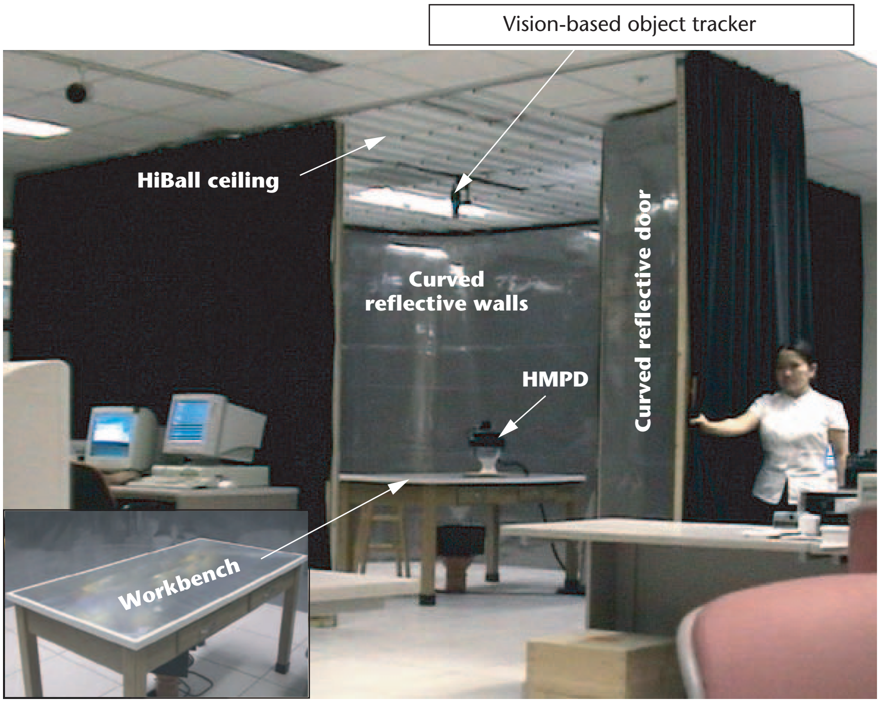

(b)

6 **Scape implementation: (a) shape and size specification of the room and (b) experimental setup.**

Vision-based object tracker. To support augmentation and interaction, a vision-based 2D object tracking method recognizes and tracks physical objects placed on the workbench. An infrared camera with infrared lamps mounted on the ceiling continuously captures the image of the objects placed on the bench (see Figure 7a, next page). In a dimmed lighting condition, the low-level coaxial infrared illumination makes the image of the retro-reflective background extremely bright and that of the diffusing objects dark. Thus, we apply segmentation algorithms to group and recognize the objects and determine their 2D position and orientation.

Under different application contexts, this tracking method with minor modifications can track multiple physical objects in augmented environments, recognize simple hand gestures to interact with virtual environments without special attachments or hand markers, and use widgets to facilitate cooperation among multiple users. For example, we configured the method to recognize and locate multiple number-coded ID checkers in our Aztec city application (see Figure 7b), which we discuss later, and reconfigured it to locate an arbitrary number of game stones placed on the bench in our augmented GO game (see Figure 7c).[13]

Magnifier metaphor. Given that the workbench presents a context of a miniature visualization of a 3D data set at a low detail, we developed a magnifier widget to let a user examine detailed views of the virtual data on the workbench. The magnifier is a handheld device coated with retro-reflective film and has a motion tracker attached (see Figure 8a). A virtual magnifier camera is associated with a macroscene at a higher detail level than the bench view. While moving the magnifier above the bench, a user perceives a magnified view superimposed on the bench view corresponding to the

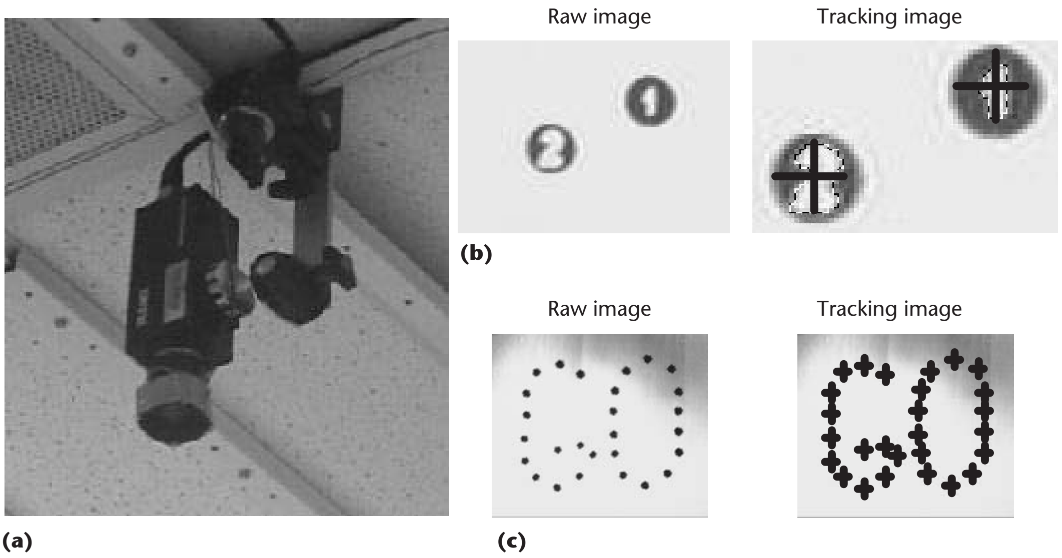

7 **Vision-based object tracker: (a) experimental setup, (b) tracking physical IDs in Aztec Explorer, and (c) tracking physical stones in augmented GO.**

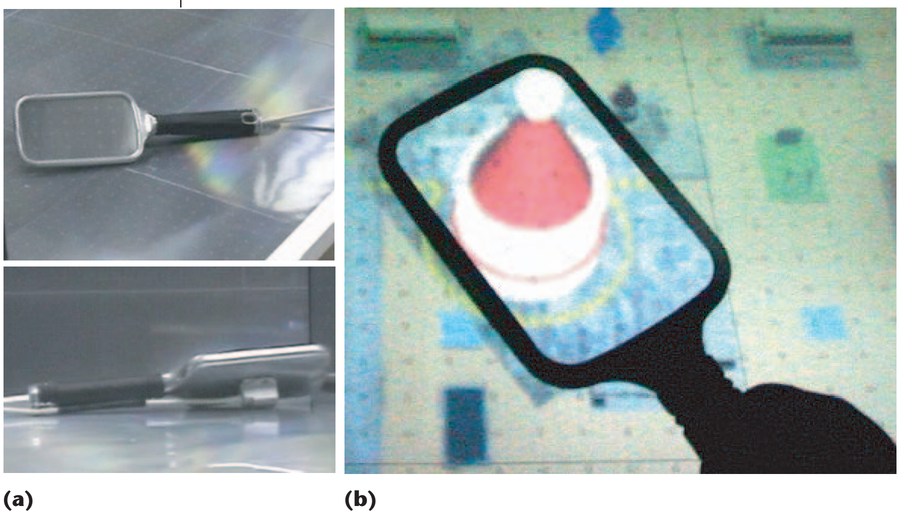

8 **Magnifier widget (a) implementation and (b) at work.**

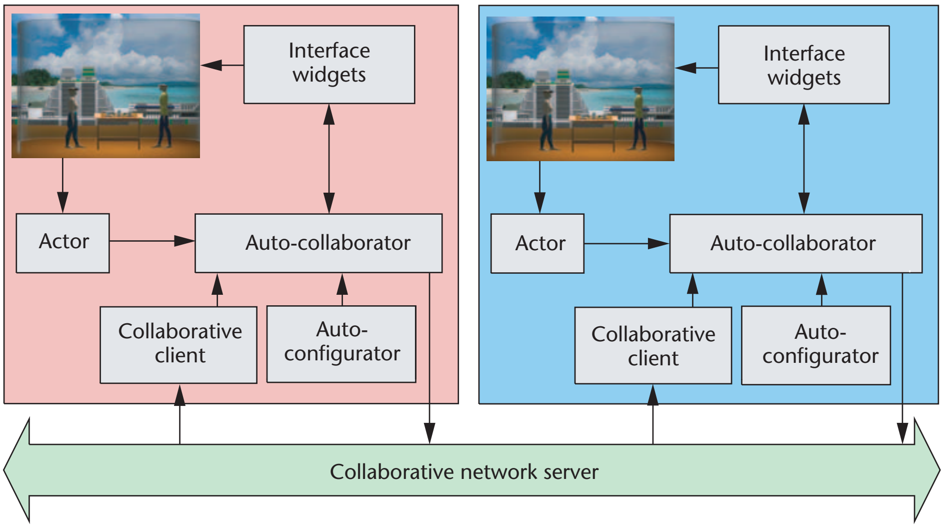

9 **Scape software architecture.**

spot of interest as captured by the magnifier's virtual camera (see Figure 8b). Thus, the magnifier metaphor naturally creates a through-a-window visualization at a medium level of detail that lies between the immersive macroscene and semi-immersive microscene.

Software architecture for augmented collaboration

The Scape toolkit is a cross-platform, modular, and extensible framework allowing high- and medium-level control over the Scape workspace (see Figure 9). Among its medium-level constructs are those typical of virtual environments, including a scene graph generator, stereo rendering constructs, and interface device drivers. The toolkit employs a client–server architecture for networking. We built this architecture on the Cavernsoft G2 API (see http://www.openchannelsoftware.org). A generic collaborative server class routes or multicasts arbitrary data over a TCP/IP network, and we customized the collaborative client modules for task-specific communication. Because Scape lets multiple users command several virtual and augmented interaction modes simultaneously, we also require robust controls for state management (we focus on the actor and auto-collaborator controls).

User management via actor. We associate each user with an actor object encapsulating all the real and virtual components of that user. The actor maintains a user's viewpoint for multiple scales of visualization, interface devices such as head and limb trackers and gesture gloves, transformations, and public and private data. Each actor optionally possesses a graphical avatar, which can convey presence to offsite collaborators in networked applications or assist interaction with the environment's other virtual components.

For reasons of security, ethics, and/or convenience, we don't presume the symmetric access of all users to all data. Hence we limit the accessibility of certain data and devices by constraining their ownership. Implicitly, by allowing multiple ownerships of the same private data, a subset of actors can confer on that private data as a group, independently of the larger actor community. In this methodology, user-defined behaviors help maintain and update the states of certain scene graph components within specific actors. Scape then loads the private data onto the scene graph for exclusive rendering for particular owners' views; it remains unloaded otherwise. In the case of augmented widgets, actors not owning a widget will see no virtual component when they manipulate the widget's physical device. For them, the widget is essentially turned off. The ownership requirement also suggests that a widget might identify and interface intelligently with each actor, restoring unique saved states or preferences from previous encounters.

Automation via auto-collaborator. We are implementing an auto-collaborator class that encapsulates all components of a multiuser Scape application (see Figure 9). By invoking the system's prepackaged controls, a developer can build simple, augmented virtual environments with minimal coding. The auto-collaborator provides default support for three modes of collaboration:

- interactive local collaboration in which multiple users at the same physical site have equal access to the augmented simulation;
- passive remote collaboration in which local users alter the state of the augmented simulation dynamically, while remote users only perceive the simulation's public contents and can't influence its state; and
- interactive remote collaboration in which both local and remote users symmetrically interact with the augmented simulation.

Thus far, we've only implemented the first two modes; we will develop the third mode in future work.

A nontrivial calibration procedure precisely aligns the coordinate systems of the various components so that the synthetic visualization is consistent and continuous for all users. An auto-configurator model reads each device's specific calibration data from a text initialization file. Due to space limitations, we omit here the details of our calibration methods.[13,18]

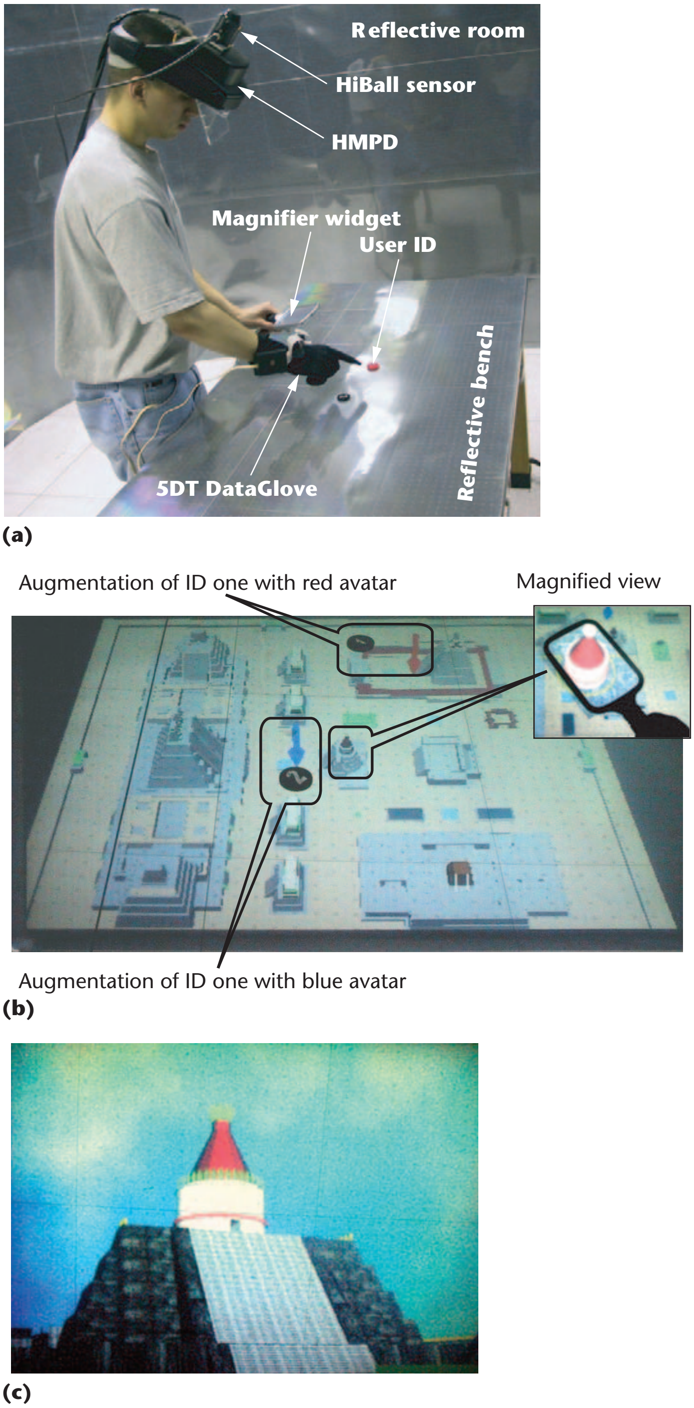

10 **Aztec explorer's architecture (a) experimental setup, (b) augmented workbench view, and (c) immersive walk-through view for user one.**

Aztec explorer

An application example demonstrates some of Scape's characteristics as well as its API framework and aspects of the interface and cooperation features that we implemented. Figure 10 illustrates the Aztec explorer's architecture and sample views.

We obtained the 3D model, an Aztec city, from 3D Cafe (see http://www.3dcafe.com), and we enhanced it with texture mapping used for walk-through and magnifier views. Our 3D scene graph mainly consists of a microscale scene at a low level of detail (which users can examine via the workbench) and a textured macroscale scene at a high level of detail (which users can walk-through in the room). We've rendered two individual perspectives for two head-tracked users. Users can either discuss the Aztec city planning with other participants through the workbench view or explore its architectural style via the walk-through. The application can support unlimited users if resources are available.

In the workbench-only collaboration mode, users share exactly the same microscene but from individual perspectives, so collaboration takes place in an intuitive face-to-face manner. Users point to a spot of interest with their hands to direct the group's focus. Participants share the magnifier widget, allowing them to closely examine the magnified view of a particular temple (see Figure 10b). The workbench supports augmentation tasks. The microscene must properly register with the physical bench and other physical components placed on the bench. The head trackers provide absolute measurements of the users' physical positions relative to the workbench, maintaining correct registration and perspectives when they walk around the bench.

Each user receives a unique physical ID, such as an encoded checker piece, which we used in our experiments (see Figure 10a). The vision-tracking interface simultaneously recognizes multiple IDs and determines their 2D locations relative to the microscene. Each user's ID location on the bench corresponds to a unique location in the macroscene. This acts as a transport mechanism to facilitate rapid navigation in a sufficiently large walk-through scene. By manipulating their physical IDs on the workbench, users can instantly move their virtual position to the site of interest in the corresponding macroscene. For example, user one starts the tour from the north end of the city, while user two starts from the center of the city (as shown in Figure 10b).

Each user has a virtual avatar (for example, a color-coded arrow) in the microscene and all participants can view it in the bench view. Each avatar represents the current macroscene location of its associated user and updates accordingly as the user walks through the scene. This provides the user and other participants with an awareness of each other's location. The virtual avatars are registered properly with the physical checkers. When users transport to a location with their checker ID, their arrow avatar initially registers with the physical checker (for example, the blue avatar for user two registers with checker two in Figure 10b). While navigating through the macroscene, the users' workbench avatars update accordingly and don't necessarily maintain superposition with their checkers until the next ID manipulation (for example, the red avatar for user one doesn't superimpose with checker one in Figure 10b). The bench view is like a shared map to explore the large city. When a user refers to the map view, both the avatars and physical IDs help to quickly identify the other users' locations or this user's previous locations.

In the walk-through mode, the macroscene is a fully immersive walk-through environment, which is experienced as significantly larger than the physical room (see Figure 10c). Scape takes measurements of the head-tracker relative to its corresponding ID position to navigate the macroenvironment. Users can also move their viewpoint by making forward or backward gestures with the DataGlove or relocating their IDs on the bench—rather than physically walking forward or backward—which overcomes the constraints of physical room size. The virtual user avatar is always updated accordingly when the head-tracker or glove is updated. This navigation mechanism lets users rapidly shift their focus of interest by moving the physical IDs on the bench or by slowly walking through the site with head-tracker and glove.

Future work

In light of the great interest in 3D collaborative interfaces and infrastructures, we developed a multiuser collaborative infrastructure. In the future, besides further efforts on the miniaturization and mobility of the HMPD helmets, we will develop the remote collaboration component. This will help integrate the visual and audio acquisition facilities and evaluate the system as a tool for stereoscopic collaborative applications with our collaborative laboratories over high-speed networks. Meanwhile, we will work on the interface design to make local or remote collaboration experiences more comfortable. We plan to develop other physical widgets to facilitate a more intuitive human–computer interface and collaboration in 3D environments. We can also integrate some interface techniques developed for ubiquitous computing into the system with the goal of offering an intelligent space for 3D visualization, computing, and collaboration. ■

Acknowledgments

We thank Jannick Rolland of the ODALab at the University of Central Florida, Narendra Ahuja of the Computer Vision and Robotics Lab at the University of Illinois at Urbana-Champaign, and Frank Biocca of the MIND Lab at Michigan State University for their stimulating discussions and collaboration. We especially acknowledge 3M for generously supplying the film. This article is based on work supported by National Science Foundation Grant IIS ITR 0083037 and 0301874.

References

1. W. Buxton, A. Sellen, and M. Sheasby, "Interfaces for Multiparty Videoconferencing," *Video Mediated Communication*, K. Finn, A. Sellen, and S. Wilber, eds., Erlbaum, 1997, pp. 385-400.
2. H. Ishii, M. Kobayashi, and J. Grudin, "Integration of Inter-Personal Space and Shared Workspace: Clearboard Design and Experiments," *Proc. Conf. Computer-Supported Cooperative Work* (CSCW), ACM Press, 1992, pp. 33-42.
3. O.Bimber et al., "The Virtual Showcase," *IEEE Computer Graphics and Applications*, vol. 21, no. 6, 2001, pp. 48-55.
4. Y. Kitamura et al, "Interactive Stereoscopic Display for Three or More Users," *Proc. Siggraph*, ACM Press, 2001, pp. 231-240.
5. C. Cruz-Neira, D.J. Sandin, and T.A. DeFanti, "Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE," *Proc. Siggraph*, ACM Press, 1993, pp. 135-142.
6. W. Kröger, and B. Fröhlich, "The Responsive Workbench,"

IEEE Computer Graphics and Applications, vol. 14, no. 3, 1994, pp. 12-15.
7. M. Agrawala et al., “The Two-User Responsive Workbench: Support for Collaboration through Individual Views of a Shared Space,” *Proc. Siggraph*, ACM Press, 1997, pp. 327-332.
8. M. Billinghurst and H. Kato, “Collaborative Mixed Reality,” *Mixed Reality: Merging Real and Virtual Worlds*, Y. Ohata and H. Tamura, eds., Ohmsha and Springer-Verlag, 1999, pp. 261-284.
9. J. Rekimoto, “Transvision: A Hand-Held Augmented Reality System for Collaborative Design,” *Proc. Virtual Systems and Multimedia* (VSMM), Int’l Soc. Virtual Systems and Multimedia, 1996, pp. 18-20.
10. Z. Szalavari et al., “Studierstube: An Environment for Collaboration in Augmented Reality,” *Virtual Reality*, vol. 3, no. 1, 1998, pp. 37-48.
11. R. Kijima and T. Ojika, “Transition between Virtual Environment and Workstation Environment with Projective Head-Mounted Display,” *Proc. IEEE Virtual Reality* (VR), IEEE CS Press, 1997, pp. 130-137.
12. H. Hua et al., “Engineering of Head-Mounted Projective Displays,” *Applied Optics*, vol. 39, no. 22, 2000, pp. 3814-3824.
13. H. Hua et al., “A Testbed for Precise Registration, Natural Occlusion and Interaction in an Augmented Environment Using a Head-Mounted Projective Display,” *Proc. IEEE Virtual Reality*, IEEE CS, 2002, pp. 81-89.
14. H. Hua, Y. Ha, and J. Rolland, “Design of an Ultralight and Compact Projection Lens,” *Applied Optics*, vol. 42, no. 1, 2003, pp. 97-107.
15. N. Kawakami et al., “Object-Oriented Displays: A New Type of Display Systems—From Immersive Display to Object-Oriented Displays,” *Proc. IEEE Int’l Conf. Systems, Man, and Cybernetics* (SMC), vol. 5, IEEE Soc. Systems, Man, Cybernetics, 1999, pp. 1066-1069.
16. J. Parsons and J.P. Rolland, “A Non-Intrusive Display Technique for Providing Real-Time Data Within a Surgeon’s Critical Area of Interest,” *Proc. Medicine Meets Virtual Reality*, 1998, J.D. Westwood et al., eds., IOS Press and Ohmsha Ltd., pp. 246-251.
17. H. Hua et al., “A New Collaborative Infrastructure: Scape,” *Proc. IEEE Virtual Reality*, IEEE CS Press, 2003, pp. 171-179.
18. C. Gao, H. Hua, and N. Ahuja, “Easy Calibration of a Head-Mounted Projective Display for Augmented Reality Systems,” *Proc. IEEE Virtual Reality*, IEEE CS Press, 2003, pp. 53-60.

Hong Hua *is an assistant professor in the Department of Information and Computer Sciences at the University of Hawaii at Mãnoa. Her research interests include stereoscopic displays, virtual and augmented reality, human–computer interaction, human and computer vision, and optical engineering. Hua has a PhD from the Beijing Institute of Technology in opto-electronic engineering.*

Leonard D. Brown *is a graduate student in the Department of Computer Science, University of Illinois at Urbana-Champaign. His research interests include virtual and augmented reality, human–computer interaction, and computer graphics. Brown has a BS from West Virginia University in computer science.*

Chunyu Gao *is a PhD student in the Department of Electrical and Computer Engineering, and the Beckman Institute for Advanced Science and Technology, University of Illinois at Urbana-Champaign. His research interests include computer vision, virtual and augmented reality, camera and imaging, and stereoscopic displays. Gao has a BSE from the Beijing Institute of Technology in opto-electronic Engineering.*

Readers may contact Hong Hua at Post 317, Dept. of Information and Computer Sciences, Univ. of Hawaii at Mãnoa, 1680 East-West Rd., Honolulu, HI 96822; hhua@hawaii.edu.