TECHNICAL OPINION

Eric W. Tatham

Getting the Best of Both Real and Virtual Worlds

A simple but effective way to facilitate visual occlusion using a see-through display.

In contrast to virtual reality systems, generally designed to immerse the user as fully as possible within a synthetic environment, computer-augmented reality supplements real-world stimuli with computer-generated elements. Visually, this is achieved by electronic or optical superimposition of computer graphics with a user's view of the real world. The potential applications are wide ranging, but there are a number of hurdles to overcome if such systems are to reach fruition.

The greatest challenges relate to maintenance of accurate spatial and temporal registration of real and virtual entities when objects are moved or the point of view changes. However, even if these problems are resolved, inherent limitations associated with current "see-through" displays remain when it comes to producing convincing integration of real and virtual elements. Most problematic, they produce a graphic overlay that is transparent and easily washed-out in regions where the background is

bright and, hence, are incapable of simulating the occlusion effects essential for appropriate depth perception. Proposed here is a simple but effective way of facilitating visual occlusion and improving the degree of color control in see-through augmented reality displays.

While most people are familiar with the concept of combining computer-generated graphics with real-world imagery (exemplified, for example, by films such as "Jurassic Park"), achieving similar effects in a head-mounted display over a real-time view of our actual environment is a more difficult proposition. The principle of such computer augmentation of reality has its roots in the early head-mounted display devised by Sutherland [8] and the subsequent head-up displays designed for military pilots. It is only recently that more wide-spread interest and potential has developed.

Visual augmentation offers many of the advantages of a synthetic world while retaining the obvious value of interaction with reality. Getting the best of both
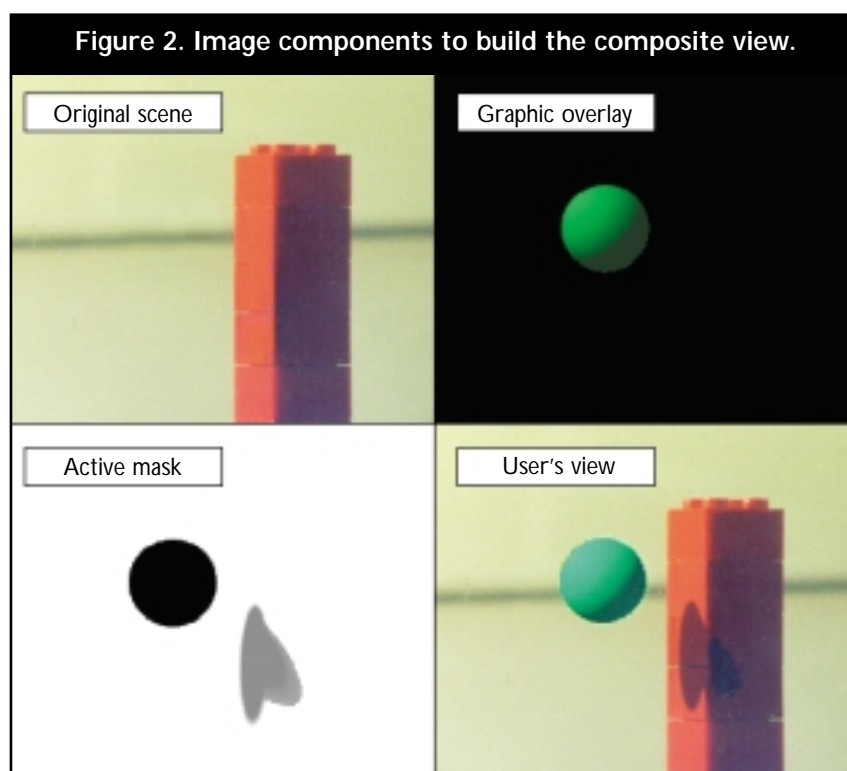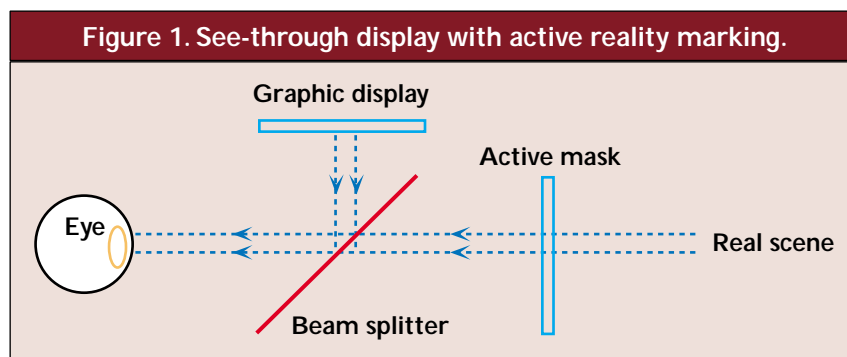
worlds in this way offers many new and exciting possibilities, including assistance in manufacturing and maintenance [4, 5], medical imaging [3], robotic teleoperation [7], and design visualization [1]. However, significant difficulties are inherent, particularly with respect to applications for which the real and virtual elements need to be convincingly integrated. For example, accurate spatial and temporal registration of computer graphics with a real scene remains crucial to the success of most applications, especially those providing manufacturing or surgical guidance. Whereas small tracking inaccuracies might not be noticeable in an immersive VR system, even very small angular errors in detecting the orientation of an augmented reality headset can result in a large displacement in the registration of the graphics.

Ultimately, more accurate methods of position and orientation tracking are required, as well as effective methods of tracking over larger distances. Further registration problems can occur due to

latency between changes in the real scene and the corresponding computer graphic update, as the graphics almost inevitably lag behind. These problems provide the focus for much of the current research in this field [2], but seamless visual integration of real and virtual worlds also depends on effective simulation of other visual factors.

It is of fundamental importance in conveying a convincing depth perception that, where appropriate, real objects appear in front of the virtual, occluding parts that cannot be seen, and virtual objects that are suitably interposed before the real backgrounds. In addition, achieving an accurate composite also depends on realistic simulation of shadows, color-bleeding, and other illumination effects arising from the interaction between real and virtual elements [6].

Even when registration problems are resolved and real-time depth mapping becomes readily available, effective visualization of an augmented composite will not be possible with current see-through displays. These are generally designed so computer-generated images are reflected over the real-world view using partly silvered mirrors angled appropriately in front of each eye. Consequently, the resultant overlay is inherently transparent, ghost-like, and unconvincing. Occlusions cannot be properly represented and the localized adjustment of real-world color quality necessary for simulating changes in illumination is impossible. Although sufficient control is available in augmented reality systems based on a video composite, see-through arrangements are necessary in cases where the synthetic

Figure 1. See-through display with active reality marking.

Figure 2. Image components to build the composite view.

elements are to be overlaid on an unpixelated view of reality.

In order to retain the direct-view advantages of a see-through, augmented reality display while also providing occlusion and facility for other desired visual interactions, a modified display arrangement is proposed. The basic alteration to existing displays is the addition of an active mask (shown in Figure 1). The computer-generated component is viewed by reflection in a partly silvered mirror. But now the real world is viewed, not only through

the mirror, but also via a transparent active panel placed along the viewer's line of sight.

The computer display and transparent panel image are both spatially registered and temporally synchronized, with the transparent element acting as an active mask, selectively reducing the intensity of light reaching the viewer's eye. On the other hand, the reflected image selectively increases the light reaching the eye. This arrangement allows for significant flexibility and control, making it possible to both reduce and increase the intensity

of light reaching the viewer from any region of a scene. Using the transparent display element to create an opaque mask and the reflected element to display the superimposed graphic object allows virtual entities to visually occlude a real background. The active mask also can be used to generate areas of neutral density that reduce light received from selected portions of the actual world and enable the simulation of virtual shadows within a real scene.

For demonstration purposes, a short animation has been produced in which a virtual sphere appears to orbit a real LEGO pillar (Figure 2). The original scene is shown together with the reflected graphic sphere overlay, its corresponding mask, and composite view seen by the user. Besides producing occlusion and shadow effects, using an active mask capable of filtering color, permits selective filtering of areas of real-world

color. Thus, employing such a mask facilitates production of any desired color-bleeding or illumination effects.

Although the images in Figure 2 were produced using a LCD panel as the active mask, this is far from ideal. Such panels introduce attenuation and distortion. A promising alternative arrangement dispenses entirely with the need for an active transparent panel, as well as obviating the requirement for a partly silvered mirror, through the use of a spatial light modulator such as the Digital Micromirror Device (DMD) developed by Texas Instruments. This device comprises an array superstructure of micromechanical aluminum mirrors each associated with a memory bit. Each mirror is an approximately 16μm square and can be tilted electrostatically depending on the state of the underlying memory cell. An appropriate arrangement of optical elements is needed to project the real

and virtual images onto the DMD surface and then onward to the user's eye. Individual mirrors in the DMD can be directed to project portions of a real or virtual image while rapid switching allows the real and virtual to be mixed in varying proportions.

Whatever the hardware used for implementation, the principle of incorporating active reality masking overcomes significant inherent limitations of current see-through displays and provides a degree of control that helps enable computer-augmented reality to fulfill its promise as a highly versatile tool for the future. ▣

Eric W. Tatham (E.W.Tatham@ open.ac.uk) is a Computer Science lectturer at The Open University, U.K.

References

1. Ahlers, K., et al. Distributed augmented reality for collaborative design applications. In *Proceedings of Eurographics '95* (Maastricht, Netherlands, Aug. 1995), 3–14.
2. Azuma, R. Tracking requirements for augmented reality. *Commun. ACM 36*, 7 (Jul. 1993), 50–51.
3. Bajura, M., Fuchs, H., and Ohbuchi, R. Merging virtual objects with the real world: Seeing ultrasound imagery within the patient. *Comput. Graph. 26*, 2 (1992), 203–210.
4. Caudell, T. and Mizell, D. Augmented reality: An application of heads-up display technology to manual manufacturing processes. In *Proceedings of Hawaii Int. Conf. System Sciences* (Jan. 1992), 659–669.
5. Feiner, S., MacIntyre, B., and Seligmann, D. Knowledge-based augmented reality. *Commun ACM 36*, 7, (Jul. 1993), 52–62.
6. Fournier, A., Gunawan, A., and Romanzin, C. Common illumination between real and computer generated scenes. In *Proceedings of Graphics Interface '93.* (Toronto, May 1993), 254–262.
7. Milgram, P., Drascic, D., Grodski, J., Restogi, A., Zhai, S., and Zhou, C. Merging Real and Virtual Worlds. In *Proceedings of IMAGINA '95* (Monte Carlo, Feb. 1995), 218–230.
8. Sutherland, I. A head-mounted three Dimensional Display. In *Proceedings of the Fall Joint Computer Conference* (Washington DC, 1968), 757–764.