Editor: Theresa-Marie Rhyne

Information at a Glance

Blaine Bell, Steven Feiner, and Tobias Höllerer

Columbia University

Personal digital assistants (PDAs) combined with wireless networking are beginning to let mobile users access the same information they once used only from the office or at home. However, researchers have long realized that it's important not to treat mobile computing simply as a scaled-down version of desktop computing. Rather, mobility has profound consequences for the way in which we use computers.[1] For example, desktop users might concentrate on a computer display for long periods, but mobile users are often mobile because they want to interact with the world around them. Consequently, a common situation when using a PDA to find information about our immediate environment is to continuously switch between looking around and looking at the PDA display.

What if we could instead visualize and interact with information directly in the context of our surroundings? Our research group is exploring how augmented reality[2] (AR) could someday make this possible. AR integrates a complementary virtual world with the physical world—for example, by using head-tracked see-through head-worn displays to overlay graphics on what we see. Instead of looking back and forth between the real world and a PDA, we look directly at the real world and the virtual information overlaid on it. At the heart of this approach is *context-aware computing*,[3] computing systems that are sensitive to the context in which they operate, ranging from human relationships to physical location. For example, information might be tied to specific locations within a global, Earth-centered, coordinate system.[4]

How can we design effective mobile AR user interfaces? We've been trying to answer this question in part by developing the experimental AR research prototypes we describe here. In AR, as in work on information visualization using desktop technologies, the amount of information available can far exceed what a system can legibly display at a given time, necessitating information filtering. Julier et al.[5] have developed information filtering techniques for AR that depend on the user's goals, object importance, and proximity. In the descriptions that follow, we assume that a system can accomplish information filtering of this sort and that our system is displaying everything it should.

1 Statues labeled by our AR system.

Information displays in AR

We're especially interested in how an AR system could provide information at a glance to aid a mobile user exploring an unfamiliar environment. Such a user might be a tourist visiting a foreign city or a native trying to find a building from its address. For example, Figure 1 shows the display of one of our prototypes in the rotunda of Columbia's Low Library. Here, the user sees the identities of the four statues on the balcony as overlaid labels. (The system automatically created this image by video mixing the virtual information in real time with the output of a video camera whose position and orientation it tracked.) With a conventional guidebook, the user would have needed to find a textual description or picture of the statues and correlate it with the statues—a time-consuming task. In contrast, using our prototype, the user can view this information simply by glancing at the statues.

View management

The labels in Figure 1 are automatically positioned and sized so that they neither overlap any of the statues in view, nor overlap each other. To eliminate ambiguity, the system connects each label to its statue with an arrow. We use the term *view management*[6] to refer to maintaining visual constraints of this sort on the projections of virtual objects on the view plane. These constraints can involve both physical and virtual objects, as in this case, in which the system constrained virtual

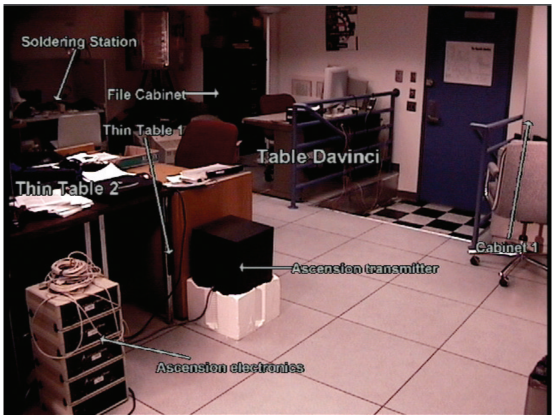

2 Laboratory labeled by our AR system.

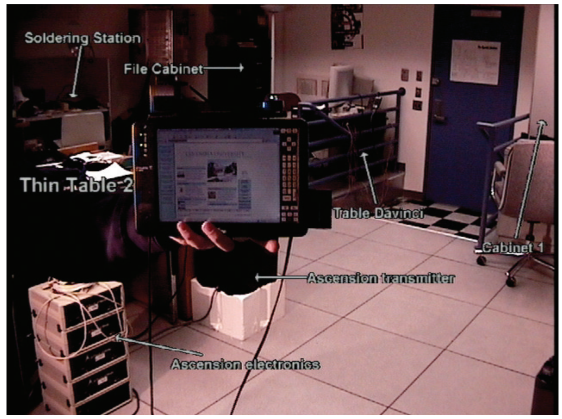

3 Tracked handheld display takes priority over labels.

labels relative to the physical objects that they overlay.

Our approach to view management begins with a model of the physical environment and its contents; a specification of the virtual annotations; and information about the properties, relationships, and relative priorities of all the virtual and real objects. The system represents each object with the visible portion of its projection on the view plane the user sees. Initially, only the real, physical objects (over which the system has no control) are visible prior to the system laying out virtual objects. In general, we let a virtual object obscure another only if the obscuring object is assigned a higher priority than the object obscured. In this example, the four statues and their labels all have the same priority, so none of the labels overlap the statues or each other. Furthermore, each label is related to its statue, a relationship that is realized visually by drawing arrows from the labels to their statues.

Our system maintains these constraints in real time for the user's current position and orientation, relying on an algorithm that approximates each object's projection by its upright rectangular extent.[6] To avoid visual discontinuities between frames as the user and objects move relative to each other, the system uses various techniques to take into account how information was laid out in previous frames when calculating the current frame's layout. For example, a label's position in the previous frame will take precedence over a position the system calculated independently for the current frame if it determines the new placement will result in too large a jump.

Figure 2 shows how our system annotates a set of selected objects in our lab. In this case, we defined all labels to be of equal priority to each other but of higher priority than the objects that they label. Consequently, the system automatically lays out the labels so that each attempts to occupy the visible silhouette of the object with which it's associated (with the overlap representing the labeling relationship). If a label can't be legibly placed with a position and size that obeys this constraint, then the system tries to place it outside but close to its object's visible portion, without overlapping the other labels and objects, and with an arrow from the label to the desired object.

4 Pop-up window provides additional information about selected object.

Figure 3 shows how we can accommodate a six-degree-of-freedom-tracked handheld computer that we don't want the labels to obscure. We gave the tracked handheld computer a higher priority than any of the objects in Figure 2 and no associated label. As a result, the handheld computer obscures much of "Table Davinci," resulting in its label changing from a large internal label to a smaller external one; the "File Cabinet" and "Ascension Transmitter" labels and arrows move up and down, respectively, out of the handheld computer's way; and the "Thin Table 1" label disappears.

Asking for more information

Users can also ask our system for more information. Figure 4 shows what happens when a user requests more information about the statue of Sophocles. The system creates a pop-up window that displays a brief textual description of the Athenian playwright and a close-up view of his statue. The pop-up window replaces the label, with the same priority and relationship to the statue. Therefore, it also avoids overlapping the statues and the other labels.

We've used a similar approach in an outdoor prototype. Our testbed mobile AR system[7] (see Figure 5, next page) consists of an external frame backpack on which we mounted a laptop computer with a 3D

5 Outdoor mobile AR system presents overlaid graphics on a tracked see-through head-worn display.

6 Restaurant information appears in a pop-up window.

7 Overlaid graphics recreate a building's furniture from a century ago.

graphics accelerator and a centimeter-level-accurate real-time-kinematic Global Positioning System and Global Navigation Satellite System (GLONASS) receiver for position tracking. The user wears a color, stereo, see-through head-worn display tracked by an inertial orientation tracker and holds a handheld display. In the figure, the user is standing on Broadway in New York City across the street from The West End restaurant. The backpack computer's screen displays the graphics overlaid by the see-through head-worn display.

Figure 6 shows a view through the head-worn display of Tom's Restaurant—the eatery immortalized by Suzanne Vega and Jerry Seinfeld—and a related pop-up window. (Unlike the other figures in this article, we photographed this figure directly through an optical see-through head-worn display rather than through a video see-through display.) In this case, the pop-up window has a brief description; a picture of the restaurant's interior; and links to the restaurant's menu, Web page, and customer reviews. The pop-up window is automatically positioned so it doesn't obscure the restaurant's projection.

Traveling back in time

AR systems have the potential to show users information about the past as well as the present.[7] Figure 7 shows a view of the Low Library rotunda after the user has asked to view its appearance in 1900. Here we use a model that recreates the rotunda's furnishings 100 years ago, when it served as the reading room of Columbia's library. In this image, virtual bookcases and desks overlay the user's view of the room. (Although we captured this image through a video see-through display, it doesn't take advantage of the ability to obscure existing infrastructure in the physical world. Therefore, a collection of contemporary furniture set up for an ongoing conference is also visible.)

Conclusions and future work

This article shows some examples of how an AR system can annotate a user's environment to visualize related information. Unlike desktop visualization environments in which the system can exercise an arbitrary amount of control over what the user sees, in head-tracked AR, the user determines the viewing specification and the physical environment determines which objects are visible. Thus, we've developed an approach in which our system interactively redesigns overlaid graphics to convey the desired information, taking into account what has already been presented to avoid visual discontinuity.

We're currently exploring several research directions. In the examples we describe here, we explicitly set object constraints and priorities in advance, which can be rigid and tedious. In contrast, we're developing a rule-based component that uses a knowledge base to determine constraints and priorities interactively, based on the higher-level relationships among objects, and on models of the user and the situation. To determine the effectiveness of this approach, we're designing a series of task-oriented usability studies to compare user performance under different versions of the overlaid user interface and to guide the design of our knowledge base. ■

Acknowledgments

This research is partially funded by ONR Contracts N00014-99-1-0249, N00014-99-1-0394, and N00014-99-1-0683; NSF Grant IIS-00-82961; and gifts from Intel, Microsoft, and Mitsubishi. Oliver Zeller created the model of Low Library.

References

1. B. Rhodes, "The Wearable Remembrance Agent: A System for Augmented Memory," *Personal Technologies*, vol. 1, no. 4, Dec. 1997, pp. 218-224.
2. R. Azuma et al., "Recent Advances in Augmented Reality," *IEEE Computer Graphics and Applications*, vol. 21, no. 6, Nov./Dec. 2001, pp. 34-47.
3. T. Moran and P. Dourish, eds., special issue on Context-Aware Computing, *Human–Computer Interaction*, vol. 16, nos. 2-4, Dec. 2001, pp. 87-95.
4. J.C. Spohrer, "Information in Places," *IBM Systems J.*, vol. 38, no. 4, 1999, pp. 602-628.
5. S. Julier et al., "Information Filtering for Mobile Augmented Reality," *Proc. IEEE and ACM Int'l Symp. Augmented Reality* (ISAR 00), IEEE CS Press, Los Alamitos, Calif., 2000, pp. 3-11.
6. B. Bell, S. Feiner, and T. Höllerer, "View Management for Virtual and Augmented Reality," *Proc. ACM Symp. User Interface Software and Technology* (UIST 2001), ACM, New York, 2001. pp. 101-110.
7. T. Höllerer, S. Feiner, and J. Pavlik, "Situated Documentaries: Embedding Multimedia Presentations in the Real World," *Proc. Third Int'l Symp. Wearable Computers* (ISWC 99), IEEE CS Press, Los Alamitos, Calif., 1999, pp. 79-86.

Readers may contact Blaine Bell at Columbia Univ., 450 Computer Science Building, 1214 Amsterdam Ave., Mailcode 0401, New York, NY 10027-7003, email blaine@cs.columbia.edu.

Readers may contact department editor Theresa-Marie Rhyne by email at tmrhyne@ncsu.edu.