4. Gaver, W., Moran, T., MacLean, A., Lövstrand, L., Dourish, P., Carter, K. and Buxton, W. Realizing a video environment: EuroPARC's RAVE system. In *Proceedings of CHI'92* (Monterey, Calif. May 3–May 7, 1992), ACM, New York, 1992. pp. 27–35.
5. Hansen, W.J. and Haas, C. Reading and writing with computers: A framework for explaining differences in performance. *Commun. ACM 31*, 9 (Sept. 1988).
6. Ishii, H. and Miyake, N. Toward an open shared workspace: Computer and video fusion approach of Team-WorkStation. *Commun. ACM 34*, 12 (Dec. 1991).
7. Johnson, J., Roberts, T., Verplank, W., Smith, D., Irby, C., Beard, M., Mackey, K. The Xerox Star: A retrospective. *IEEE Comput. 22*, 9 (Sept. 1989).
8. Johnson, W., Rao, R., Hellinek, H., Klotz, L., Card, S. Bridging the paper and electronic worlds: Paper as a user interface. In *Proceedings of INTER-CHI'93* (Amsterdam 24–19 Apr. 1993).
9. Krueger, M. *Artificial Reality II*. Addison-Wesley, 1991.
10. Luff, P., Heath, C.C. and Greatbatch, D. Tasks-in-interaction: Paper and screen based documentation in collaborative activity. In *Proceedings of CSCW'92*.
11. Newman, W. and Wellner, P. A desk supporting computer-based interaction with paper documents. In *Proceedings of CHI'92* (Monterey, Calif. May 3–7, 1992) ACM, New York, 1992.
12. Olson, J.S., Olson, G.M., Mack, L.A., Wellner, P. Concurrent editing: the group's interface. In *Human-Computer Interaction—INTERACT'90* (Cambridge, England), Elsevier, Amsterdam 1990. pp. 835–840.
13. Rheingold, H. *Virtual Reality*. Secker and Warburg, 1991.
14. Seybold, A. The DOC.IT *Andrew Seybold's Outlook on Professional Computing 11*, 2 (Sept. 1992).
15. Tang, J., Minneman, S. VideoDraw: A Video interface for collaborative drawing. *ACM Trans. Inf. Syst. 9*, 2 (Apr. 1991).
16. Wellner, P. The DigitalDesk calculator: Tactile manipulation on a desk top display. In *Proceedings of the ACM Symposium on User Interface Software and Technology* (UIST '91) (Nov. 11–13, Hilton Head, S.C. 1991).
17. Wellner, P. Tactile manipulation on the DigitalDesk. (video) in *CHI '92 Special Video Program, ACM SIGGRAPH Video Review* 79.
18. Wellner P. Self calibration on the DigitalDesk. EuroPARC Tech. Rep. EPC-93-109, 1993.
19. Wellner P. Adaptive thresholding on the DigitalDesk. EuroPARC Tech. Rep. EPC-93-110, 1993.
20. Wellner, P. and Freeman, S. The DoubleDigitaldesk: Shared editing of paper documents EuroPARC Tech. Rep. EPC-93-108, 1993.

CR Categories and Subject Descriptors: C.3 [**Computer Systems Organization**]: Special-purpose and application-based systems; H.1.2 [**Information Systems**]: Models and Principles—*User/Machine systems*; H.4.1 [**Information Systems**]: Information Systems Applications—*Office Automation*; J.0 [**Computer Applications**]: General

General Terms: Design, Human Factors

Additional Key Words and Phrases: DigitalDesk

About the Author:

PIERRE WELLNER is a final-year Ph.D. candidate at the University of Cambridge Computer Lab and a research scientist at Rank Xerox EuroPARC. Current research interests include human-computer interaction, computer-augmented environments, and novel input/output devices. **Author's Present Address:** Rank Xerox PARC, Cambridge EuroPARC, 61 Regent Street, Cambridge, United Kingdom, CB2 1AB; email: wellner@europarc.xerox.com

Acknowledgments

On February 24, 1993, we held a small workshop on "Augmented Reality and Ubiquitous Computing" at MIT. Many of the authors in this issue participated, as well as other researchers who have done substantial work in the field. We discussed issues we had in common and the defining characteristics of this field. These discussions helped shape our thinking, brought together related new ideas from different areas, and confirmed to us the importance of this approach to the future of human-computer interaction.

We wish to acknowledge the participants: Hal Abelson, Fred Martin, Nick Papadakis, Bill Stasior from MIT; Ron Azuma from University of North Carolina-Chapel Hill; Thomas Baudel and Michel Beaudouin-Lafon from LRI—University of Paris-Sud; Bill Buxton, Mike Dixon, Marvin Theimer, Mark Weiser from Xerox PARC; Steven Feiner from Columbia University; George Fitzmaurice, Paul Milgram from University of Toronto; Craig Hubley from Hubley & Assoc.; Myron Krueger from Artificial Reality; Murray Mazer from DEC—CRL; Peter Robinson from University of Cambridge; Zary Segall from Carnegie Mellon University; and Meg Withgott from Interval Research Corp.—*Pierre Wellner, Wendy Mackay, and Rich Gold*

WENDY MACKAY, GILLES VELAY, KATHY CARTER, CHAOYING MA, AND DANIELE PAGANI

Augmenting Reality: Adding Computational Dimensions to Paper

Despite claims about the "paperless office," we find that today's technology, rather than replacing paper, has increased our use of it. We use tiny Post-it notes and large pads of flipchart paper. We fill in preprinted forms and buy products with barcode labels. We annotate our calendars and telephone directories. We continue to read books, newspapers and memos. Chances are that you are reading these words on paper.

Since we are committed to using paper, how can we best augment it? Computers provide a number of possible dimensions, such as computational, semantic, syntactic, graphic, and temporal. Thus, columns of numbers become amenable to "what-if" speculations on an electronic spreadsheet; or they can appear as dynamic weather patterns or rotating molecular models when fed into a scientific visualization program. Words on a page can be translated into French, checked for spelling or analyzed for writing style. Two-dimensional hand-drawn sketches can be edited, replicated or made more precise; or they can be rendered into 3D and rotated or projected onto other objects. A paper cartoon can be animated, or a still image can become a movie.

We are using the DigitalDesk (see the accompanying article "Interacting with Paper on the DigitalDesk") to explore these dimensions. Figure A shows the Digital Drawing Board, which lets designers sketch as they normally do and then modify their drawings with the aid of the computer. Figure B shows Mosaic, which lets users work with paper storyboards to create computer-controlled video productions.

Digital Drawing Board

The goal of this project is to let designers work as they normally do in the studio and still take advantage of the computational power of CAD systems. Most designers surround themselves with sketches, pinning images to the walls and spreading them around the drafting table. They use a variety of tools and materials, which they place within easy reach and use in a free flow of activity

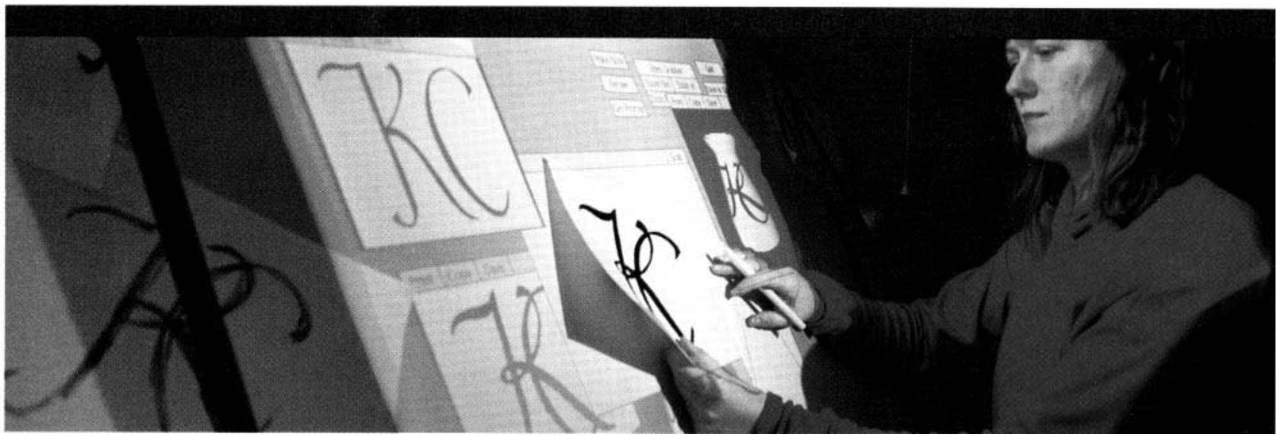

Figure A. Digital Drawing Board

when sketching. They walk around, stand back or look closely at the sketches as the ideas unfold, maintaining a very active physical interaction with the design space.

The Digital Drawing Board consists of an A1 (635mm × 889mm) digitizer and an LCD data projector, which projects computer-generated images onto the digitizer. The designer can use the board as both a conventional drawing board and as a computer. The digitizing stylus lets the designer use gestures to organize electronic objects and view them over a large space.

A video camera mounted beside the projector and connected to a frame grabber supports the transfer of images from paper into the computational environment. Various computations can then be applied to the digitized images. A prototype application has been built that digitizes sketches and uses them as texture maps for rendering the surfaces on solids of revolution. The photograph shows a designer sketching monograms that are then rendered onto the surface of a vase by the computer. The designer can try a wide variety of different sketches and explore how alternative designs appear in 3D. Ultimately, the images could be projected onto the 3D-object itself. Other applications of the Digital Drawing Board can provide support to architects, map-makers, electrical engineers, industrial designers and others who need to sketch ideas rapidly in a large space and then take advantage of the computer to enhance or compute aspects of their drawings.

Mosaic

Mosaic addresses the problem of extending paper along the temporal dimension, providing support for managing video and other time-based data. The initial Mosaic application is designed to support video producers, who use a combination of video-editing equipment and paper storyboards to manage the development of their videos. A storyboard consists of a set of elements, each containing a sketch or image of the "best frame" of the video, associated text, and notes that describe the action. Storyboards provide an efficient way to sketch action sequences and develop "what-if" scenarios; they are flexible, portable, and easily annotated. The designer can lay out a large number of storyboard elements and view multiple alternatives at the same time. However, there is no easy way to associate the offline storyboard with the video that has been created on-line. Therefore, video producers must move to a completely different system to view and edit the moving video.

Figure B. Mosaic

Mosaic provides a new interface that combines the benefits of paper storyboards with computer-controlled video. The Mosaic DigitalDesk consists of an LCD projector and video camera positioned over a standard desk with a small (11.5- × 9cm) LCD monitor mounted in the surface of the desk. The video camera captures handwritten annotations from the paper storyboard, commands such as "print" or "play" and glyphs or numbers that identify the particular storyboard element. The LCD projector and monitor work together to provide a high-resolution video image with associated text and handwritten annotations. Video is stored on a write-once videodisc for high access speed and image quality. (Mosaic could also work with videotape or with digitized video.)

The video producer can lay out the elements of a storyboard in any order on the desk and then identify and play any sequence of clips. The user can view the storyboard elements in different sequences by placing them under the camera. Mosaic recognizes and remembers which elements should be included and in what order and can then play them on the desk, print them, or save them as a new storyboard. Future versions of Mosaic will use a second camera to allow users to simply point at the storyboard elements, rather than physically move them under the camera. This lets the user move smoothly between the spatial and temporal views of the video.

Mosaic also allows users to issue commands directly from the desk. For example, the user can place the word "New" under the camera, link a new set of storyboards together, and then use "Play" to view it or "Save" to store it. If the user wants to capture notes scribbled on paper storyboard elements, it is possible to grab the image of the note and incorporate it into the computer-based version of the storyboard. Users can also search for storyboard elements with a particular characteristic, such as those containing a particular phrase, and use that information to automatically generate new storyboards. Future applications will extend these capabilities using the EVA system, which allows researchers to analyze multiple streams of multimedia data. We are interested in developing a user interface that allows flexible exploration of temporal data of any kind. **C**