






INTERACTING WITH PAPER ON THE DIGITAL DESK

Pierre Wellner

We interact with documents in two separate worlds: the electronic world of the workstation, and the physical world of the desk. Interaction styles in these two worlds do not resemble each other, functions available are different, and integration between the two is limited. Each world has advantages and constraints that lead us to choose one or the other for particular tasks (see Figure 1). Unfortunately, choosing to interact with a document in one world means forgoing the advantages of the other. A great challenge to office system designers is to provide the best of both types of documents.

The ways we physically interact with electronic documents are limited compared to our interactions with paper, pencils, erasers, rulers, and other traditional tools on the desk. When interacting with objects in the physical world, we take advantage of natural skills developed over our lifetimes. We use our fingers, arms, 3D vision, ears and kinesthetic memory to manipulate multiple objects simultaneously, and we hardly think about how we do this because the skills are embedded so deeply in our minds and bodies. Although electronic documents provide valuable functions such as spell checking, numerical calculations, and keyword searching, people must sacrifice highly developed tactile skills with paper in order to access these functions.

The classic approach to solving this problem is represented by the "desktop metaphor," developed in the 1970s [7]. By making electronic workstations analogous to the physical desk, users can take advantage of their knowledge of the physical world, the computer becomes more familiar, and less learning is required. This approach has remained successful, and electronic documents are continuing to acquire more properties of physical documents (high-resolution color, portability, pen-based interaction, for example).

Where a property in the physical world has an advantage over the computer, human-computer interaction (HCI) researchers tend to try to find a way to enhance the computer with that desired property. Many items, in addition to desktops and paper documents are put into the computer. This tendency, taken to an extreme, is virtual reality (VR), where the user abandons the real world to be completely surrounded by the computer [13].

real world can be overcome in VR, and physical tools can be made obsolete by more flexible, virtual alternatives.

Physical tools can be hard to replace, however. At one time it seemed that paper might become obsolete, for example, and visionaries predicted the "paperless office" would dominate within a few years. But the trouble is that people like paper. It is easier to read than a screen [5], it is cheap, universally accepted, tactile, and portable. According to some studies, paperwork in the office has increased by a factor of six since 1970, and is now growing at 20% annually [14]. Like electronic documents, paper has properties that people just cannot seem to give up, making it resilient in the face of computer-based alternatives [10].

Consequently, we have two desks: one for "paper pushing" and the other for "pixel pushing." Although activities on the two desks are often related, the two are quite isolated from each other. Printers and scanners provide a way for documents to move back and forth between desktops, but this conversion process is inconvenient. Wang's Freestyle system [3], for example, was a "paperless office" system that attempted to coexist with partially paper-based processes. A key factor necessary for adoption of that system was to minimize the printing and scanning re-

quired, because too many of these tasks would cause the process to revert entirely back to paper.

Trade-offs between electronic and paper documents can make the choice of medium difficult, but imagine if we did not have to choose, and we had a space where documents could be both paper and electronic at the same time. Instead of putting the user in the virtual world of the computer, we could do the opposite: add the computer to the real world of the user and create a Computer Augmented Environment for paper (see guest editors' introduction.). Instead of replacing paper with computers, we could enhance paper with computation.

The Xerox PaperWorks product [8] takes a step in this direction with its fax-based paper user interface (UI) to a storage and retrieval system. With this system, ordinary paper forms are enhanced to control a PC through a fax machine. These paper documents gain some properties of electronic documents, but fax machines are slow compared to computer screens. Response time is limited by the delay it takes to scan and print a page, and this limits the range of interaction techniques possible.

Another approach to enhancing paper documents with computation is to create the opposite of the desktop metaphor. Instead of making the workstation more like a desk, we can make the desk more like a workstation. This is the aim of the DigitalDesk. On this desk, papers gain electronic properties, and electronic objects gain physical properties. Rather than shifting more functions from the desk to the workstation, it

shifts them from the workstation back onto the desk.

The DigitalDesk

The DigitalDesk is a real physical desk on which you can stack your papers, lay out your favorite pencils and markers, and leave your coffee cup, but it is enhanced to provide some characteristics of an electronic workstation. A computer display is projected onto the desk, and video cameras pointed down at the desk feed an image-processing system that can sense what the user is doing (see Figures 2 and 3). No desktop metaphor is needed because it is *literally* a desktop.

The DigitalDesk has the following three important characteristics:

- it projects electronic images down onto the desk and onto paper documents,
- it responds to interaction with pens or bare fingers (hence *DigitalDesk*), and
- it can read paper documents placed on the desk.

Each of these characteristics is further discussed in the section on implementation issues, but first, to illustrate how the desk can be used, the following section describes some example applications.


Example Applications

The DigitalDesk provides a computer-augmented environment in which paper gains electronic properties that allow it to overcome some of its physical limitations. There are many possible ways to enhance ordinary paper documents on this desk. The following subsections describe three working prototype applications: a calculator, PaperPaint, and the DoubleDigitalDesk. Other examples (such as a French-to-English translation system) are described elsewhere [11, 17] and in the accompanying sidebar by Mackay et al.

Calculator

The calculator is a simple and familiar application that can benefit from the DigitalDesk. People using calculators often enter numbers that are printed on a piece of paper lying on the desk, and they must copy the numbers manually into the calculator

Figure 1. Comparison of advantages of electronic and paper documents






Figure 2. Schematic diagram of a DigitalDesk




Figure 3. The first DigitalDesk prototype was cobbled together from a 1120-by-780 "ScratchPad" display, an overhead projector, some spare video cameras, a cooling fan, polarizing filter, cardboard, and a lot of tape. Versions based on commercial computer projection units have since been built.

in order to perform arithmetic on them. Transcribing these numbers can constitute a large proportion of the keystrokes when using a calculator, and a large proportion of the errors.

The DigitalDesk Calculator (previously described in [16]) addresses this problem by providing another means of entering numbers. It allows people to place ordinary paper documents on the desk and simply point at a printed number to enter it into the calculator. In the working prototype, users can point with a pen or bare finger, and a rectangle is projected in front of the finger to indicate which number is selected. When the user taps the desk, the system reads this number with a camera, recognizes the digits, and treats them as though they had been typed into the calculator (see Figure 4).

This example application shows how a paper document can become more like an electronic document by allowing selection of numbers for calculation. Of course, it is also possible to scan a paper receipt or annual report through an optical character recognition (OCR) program to convert the entire paper document to electronic form, but this is a less interactive, more batch-oriented process, and the resulting electronic document has lost all physical properties of the original. The purpose of the DigitalDesk is not to convert paper into electronic documents; it is to support rapid and direct computer-based interaction with selected regions of paper documents.

In this prototype, numbers are entered into a projected calculator tool, but it would also be possible to use a physical calculator that was connected to the DigitalDesk by wire or infrared. Another possibility is to leave out the calculator, and project results back onto the paper. A "paper spreadsheet" could be designed on which pen-based interaction techniques would operate in the same way on both ink and projected numbers (See [17] for a video envisionment of how this and other possible applications might work).

PaperPaint

Although "select and paste" is now a standard feature of electronic docu-

ments, it is awkward to perform the same operation with real paper, requiring a photocopier, scissors, and some adhesive. The DigitalDesk, however, makes it possible to select and paste paper documents in the same way that we select and paste electronic documents. A simple paint program has been implemented (PaperPaint) in which a sketch on paper can be electronically selected by sweeping out an area of the paper with a stylus; the projector displays a rectangle on the paper to indicate what is selected. When the stylus is raised, the system snaps a picture, and the projected rectangle is replaced by a thresholded electronic copy of the area. This copy can then be moved about and copied to other parts of the paper. Sliding this electronic copy over the drawing to place it somewhere else is very similar to sliding a paper copy (see Figure 5).

This application allows users to construct a mixed paper and electronic drawing. Currently, the system cannot move the projected image to follow the movement of paper, so users must keep the paper motionless. A combined scanner/printer (not yet implemented) would be necessary to make the projected marks permanent and allow users to take the merged document away. Even without such a printer, the system could be made to recognize papers whenever they are placed on the desk (using a system such as Marcel [11]) and fill in the appropriate electronic parts.

User testing of PaperPaint revealed another (unexpected) way of using this tool, which is also very powerful. Instead of constructing a mixed paper and projected drawing, some users noticed they could construct a purely projected drawing from selected portions of their paper sketches. They can sketch a figure on paper, move it to the desired location in the projected drawing, then select it so that it remains "pasted down" in that location after moving the paper away. The effect is similar to that of dry-transfer lettering or rubber stamping, but from any piece of paper onto projected electronic drawings. This encourages the use of hand-sketched or printed templates of *paper*-based "clip art" that can be

naturally placed and rotated into position with the fingertips. This interaction technique is quite different from the standard "select and paste" found on most workstations and takes advantage of unique qualities of the DigitalDesk: using both hands for manipulating and pointing as well as the superimposition of paper and electronic objects.

DoubleDigitalDesk

When working together, people often use documents and frequently must simultaneously see and modify these documents. Two people on separate continents cannot normally write on, point at, or otherwise manipulate the same paper document, but this is another constraint of physical paper than can be addressed by the DigitalDesk.

Shared editing of documents has been the focus of a number of research projects (see [6, 12, 15] for examples). Most of this work has concentrated on screen-based documents, but the DoubleDigitalDesk makes it possible to share real paper documents. It allows users in two separate locations to "share" their physical desks, both seeing, editing and writing on each other's paper documents. For a full description of this project, including discussion of the use of multiple pointing devices, see [20].

In this application, each DigitalDesk continuously grabs images from its local desk and projects scaled, thresholded images from the remote desk. The result is that both users see what is on both desks. When a paper document is placed on desk A, its image is projected onto desk B and vice versa. The projections are digitally scaled and positioned to provide the same view to each participant, and the contrast is adjusted so that feedback (the image of your own desk transmitted back from the remote desk) is minimized. Each user can draw with a real pen on both paper and electronic documents, and the other user will see these marks appear in the corresponding places. Hand motions are also transmitted, enabling users to see each other point to certain areas on the documents (see Figure 6). Since the partner's hands block

Figure 4. (Calculator) This photograph was taken just after the user selected the number 4834 on a piece of paper and the number was recognized and put into the calculator as though it had been typed.





Figure 5. (PaperPaint) This user has selected the sketch of the window and made two copies of it. Now he is about to move a copy of the flower that he drew.




Figure 6. (DoubleDigitalDesk) The local user is drawing on paper with red ink, while the remote user's paper and hand can be seen as having just finished drawing an 'O'.




Figure 7. White areas indicate the difference between two frames in which the finger has moved

Figure 8. Example of the poor-quality image result of simple thresholding

the view of what is underneath them, just as with an ordinary desk, this must be dealt with through social protocols and speech. Not pictured in the figure is an audio link through telephones or speakerphones. Another useful addition is a face-to-face audio-video link through, for example, the EuroPARC RAVE system [4].

This system could scale up to support three or more shared desks by "or-ing" all the remote thresholded images together, or by displaying each desk's image in a different color. Of course, as more people share the space, participants will need to be more careful not to interfere with each other's documents and writing.

Implementation Issues

The following sections discuss some of the challenges involved in implementing a DigitalDesk along with strategies for addressing these challenges. These issues are discussed under three subsections: interaction on the desk, reading paper documents, and self-calibration.

Interaction on the Desk

One aim of the DigitalDesk is to go beyond so-called "direct manipulation" with a mouse (which in fact is not direct at all) and to explore the possibilities of "tactile interaction" with real and electronic objects using the fingers. The DigitalDesk merges paper and electronic documents not only by adding electronic properties to paper documents, but also by adding physical properties to electronic documents. Since we use pens and fingers to interact directly with paper on our desks, we should be able to interact with electronic documents on the DigitalDesk in the same way.

One way of interacting with electronic objects with bare fingers is through video-based finger tracking. For some applications, obscuration of fingers by other fingers or parts of the body can be a problem, but with desk work this does not seem to be a significant difficulty because the hands have a limited range of motion and they mostly remain in a 2D plane. Pointing things out to the computer is much like pointing them out to another person. Therefore it is easy for users to learn not to cover

the object being pointed to. A bare finger is too thick, however, to indicate small objects such as a single letter, so the user must also be able to point precisely with a pen or other thin object.

A wide range of interaction techniques are possible using video-based finger tracking, as demonstrated, for example, by Krueger [9]. His system and others, rely on the hands being viewed against a plain background in order to make them easier to distinguish. Unfortunately, this is impossible on a DigitalDesk because there are various paper documents, pens and pencils present in addition to the user's hands. It is difficult to distinguish these objects from fingers and pointers by shape only. A more effective strategy is to look for motion, assuming that most objects seen on the desk do not move, other than the user's hands and the objects they are holding. An effective way to pick out moving objects from the background is to capture successive frames and to examine the image produced by subtracting successive values of each pixel in the two frames. The result, when applied to images of a moving hand, for example, is shown in Figure 7. This is a better image to start from than the original cluttered image, but further processing is required to remove noise and to locate the precise position of the fingertips. In the future, more sophisticated techniques will be necessary to track multiple fingers and to recognize gestures.

Determining when the user taps on the desk is difficult when only processing images from the overhead camera. One solution to this problem is to detect finger taps in the same way a person might: by listening. A microphone is attached to the bottom of the desk and the system monitors the amplitude of the signal to determine when the user taps on the desk. This technique works well, but sometimes it confuses other taps, bumps on the desk, or hand claps with a finger tap. The finger-following and tap-monitoring tasks must be carefully synchronized, because the raised finger does not appear in the same place as the lowered finger and it moves very fast. If the image-processing system lags just a little,





Figure 9. Four-point warping

Figure 10. Four ways to sweep out a selection rectangle

then it reports the finger in the wrong place at the time of the tap. Another way to detect tapping is to use a touch screen. Unlike the microphone, it can provide dragging information as well as extra location data. A problem with desk-based touch screens (named the "Midas Effect" by Steve Freeman) is that users tend to rest their hands on them and everything touched can be interpreted as input.

Projected Display. Having discussed issues on the input side of desktop interaction, we turn now to output: the projected display. Projection from above provides similar capabilities as a large flat display screen and faces the same size versus resolution trade-off, but it has the key advantage that computer-generated images can be superimposed onto paper documents. This is necessary for creating merged paper and elec-

tronic documents, and for providing feedback when making selections on paper. There are some problems, however, with overhead projection. One potential problem is shadows; it is not possible, for example, to lean down to look at a projected image too closely. In practice, however, shadows are hardly noticed when the projector is mounted above a horizontal desk (see the following section entitled "User Experiences"), but special measures must be taken to avoid shadow problems on a nearly vertical surface such as the Digital Drawing Board [2].

The brightness of the room can also be an issue with projection. The projectors used in these experiments work quite well with normal fluorescent lights, but a bright desk lamp or direct sunlight can make the display unreadable, limiting the desk's usability in some settings. An additional problem with projection is that not all surfaces make good screens. The projection area should be white in order to see images most clearly, and some executives may be reluctant to hide their polished mahogany desk under a piece of paper or other screen.

Reading Paper Documents

For the DigitalDesk to read selected portions of paper documents the following steps are necessary: image capture, thresholding, and (in the case of text) character recognition.

Image Capture. Document images are captured through an overhead video camera, but a difficulty with standard video cameras is their low resolution compared to scanners. One way to solve this problem, used by the Marcel document recognition system [11], is to prescan documents at high resolution and then use the low-resolution camera image to look up the corresponding scanned image. Prescanning is inconvenient for many interactive applications, however, so the prototypes described in this article instead use two cameras, one of which is zoomed in close to the desk to obtain a high-resolution image (about 200 spots per inch). Since this means that only a portion of the desk is used for capturing document images in a high-resolution format, a window is pro-

jected onto the desk to indicate the active area to the user. More cameras could easily be added to cover the entire desk, but this has not yet been necessary, because the applications tried so far only use small parts of a document at a time, and sliding a piece of paper into the camera's window is so easy. In the long run, high-definition television and advances in low-cost, integrated digital cameras will make this approach more economical.

Thresholding. The image produced from a video camera and frame grabber is grayscale (typically eight-bits-per-pixel), even when it represents a white sheet of paper with black ink. This grayscale image must be thresholded, or converted to a one-bit-per-pixel black-and-white image before it can be used for character recognition or any of the other example applications previously described.

Simple global thresholding is not adequate for obtaining an image suitable for character recognition. In normal office lighting, the range of brightness on different parts of the desk varies greatly, so a global threshold creates large patches of black-and-white with indistinguishable text (see Figure 8).

In order to make a good one-bit-per-pixel image of a black-and-white document, the system must use an adaptive thresholding algorithm which varies the threshold value across the image according to its background value at each pixel. Some adaptive thresholding algorithms produce very good results but require more than one pass through the image and are too slow to support user interaction. It is possible to get nearly as good a result in a single pass, however, by calculating the threshold value at each point from an estimate of the background illumination based on a moving average of local pixel intensities [19]. This method is fast and can also be combined with a scaling operation if necessary.

Finally, when dealing with text, the thresholded image is skew-corrected and recognized by an OCR server (in this case, Xerox Imaging System's ScanWorkX). If the resolution is high enough relative to the

text size, then it returns the associated ASCII string. Because this process is not guaranteed to be accurate, it is important to provide both quick feedback (as in the calculator example), and a simple way for the user to correct unrecognized characters.

Self Calibration

To support interaction on the desk, projected feedback, and selective grabbing of images through the camera, the DigitalDesk requires a two-step calibration process. First, absolute positions on the digitizing tablet must be mapped to positions on the display in order to provide feedback. Second, positions on the display must be mapped to corresponding positions in the frame grabber in order to support grabbing of selected areas on the desk. If finger tracking is used for pointing, however, then only the second step is necessary.

Unfortunately, many factors conspire to make this calibration difficult. The projected display is not a perfect rectangle (there are optical distortions such as "keystoning"), the camera and/or tablet may be rotated relative to the projected display, and it may be necessary for the camera to view the projected display from an angle. Also, vibrations caused by air conditioners or slamming doors cause movements which disrupt the calibration, as do any adjustments to the equipment.

Obtaining the data to calibrate the pointing device to the display is relatively straightforward: a series of points are displayed and the user is prompted to touch them with a pointer. Obtaining data for calibrating the camera to the display is not as simple. One obvious way to do this is to adjust the size and shape of a projected rectangle on the desk while looking at its image on a video monitor. The rectangle can be aligned with the edges of the monitor, producing rough calibration data, but this assumes there is no keystoning, no rotation, and no underscanning by the video monitor. In general, none of these assumptions is correct, so even after performing this tedious manual procedure, the result is poor.

A better approach is to project an image that can be located by the image-processing system, allowing

the system to self-calibrate without any assistance from the user [18]. The current system projects a thick "plus" sign, and uses image morphology routines by Bloomberg [1] to pinpoint the center of the mark in frame grabber coordinate space.

Two calibration points are not enough, however, to calculate an accurate mapping; therefore the system currently uses a four-point calibration system to compensate for rotation and keystoning. To calculate the mapping from four points it uses the following equations (see Figure 9). With four point pairs, the two sets of four simultaneous linear equations can be quickly solved by Gaussian Elimination to find the values of c_{1-8} . Then, a fifth plus mark is projected and its location is checked to make sure it is close enough to the position predicted by the mapping. In practice, the result is accurate to within one or two display pixels, allowing the user to select areas on the desk and rely on the displayed feedback to precisely indicate what will appear in the grabbed image.

User Experiences

Although no formal experiments have been conducted with the DigitalDesk, a number of people have performed tasks on it, and their reactions were noted. All subjects had previously used traditional workstations or PCs, and they all said they found the desk generally comfortable and natural to work with. One unexpected result of the tests was that no one was bothered by the projection system and the shadows cast by hands and other objects. In fact, after using the desk for almost 15 minutes, one user asked if it was made of glass and looked under it to see where the display was coming from! It may be that people are so used to the shadows cast by overhead lights and desk lamps that these types of shadows are hardly noticed. At the end of each session, subjects were asked how they felt about using the desk compared to a traditional workstation. Their overall reaction was that it was quite similar, but they commented specifically that they had "more space," it was "more healthy than a screen," "easier on the eyes," and "more manual."

Handedness

Unlike a traditional workstation, UIs on the DigitalDesk must take account of handedness. If feedback is projected to the lower left of the pointer, for example, then a right-handed person has no trouble seeing it, but a left-handed person does have trouble because it is projected on the hand. Of course, handedness affects the use of any pen-based system, but with a projected display, shadows strengthen the effect. Not only is feedback affected, but so is the general layout of applications. The French-to-English translation application, for example, inadvertently assumed users were right-handed, with paper documents on the left, and projected definitions on the right. Left-handed subjects were inconvenienced by this set-up because it required them to reach their arm further than right-handed subjects, and at the same time, their arms hid the paper they were reading. These handedness issues could be addressed by using the overhead camera to automatically detect with which hand the user is pointing, and this information could be used by applications. A pop-up menu, for example, would be projected to the left of the pointer for a right-handed person, and to the right of the pointer for a left-handed person.

Obscuring Selections

Subjects also noticed a difference between selecting pixels with a workstation and selecting marks on paper with a DigitalDesk. On a workstation, we can obscure something with the pointer as we select it and the system still knows what is underneath. When pointing at paper with the DigitalDesk, however, the system must be able to see the paper, and this means that fingers and other pointing devices must be out of the way. It may be possible to address this issue by storing previously snapped (or scanned) images; the Marcel system, for example, does not have this problem. A solution may be unnecessary, however, because people do not seem to have much difficulty learning how to interact with the system in a way that keeps selections visible. When sweeping out a rectangle in the PaperPaint application, for ex-

ample, there are four ways of doing this as illustrated in Figure 10. If right-handed people use method ①, or if left-handed people use method ②, they end up with a picture of their pen in the selections. They usually have no difficulty switching to one of the other three methods instead, and they do not seem to repeat the mistake (except sometimes just for the fun of selecting their own finger or hand to watch it get copied). In general, the system cannot see the selection unless the user can see it too, and that is easy to learn.

Selection feedback can also play an important role in preventing users from obscuring their selections. If a fixed selection rectangle is centered about the pointer, to use an extreme example, then it is impossible to get the pointer out of the way. If the selection rectangle floats slightly ahead of the pointer, however, it is easy to avoid placing the pointer inside. If the system had handedness detection, it would even be possible to adapt the selection feedback, preventing users from sweeping out a rectangle the wrong way.

Conclusions

Instead of making the electronic workstation more like the physical desk, the DigitalDesk does the opposite: it makes the desk more like the workstation, and it supports computer-based interaction with paper documents. Experience so far with this desk is encouraging. It enables people to use paper documents to perform useful tasks that are more awkward to do in other ways: tasks such as copying a long number into a calculator, translating a foreign word, replicating part of a sketch, or remote shared editing of paper documents. The interaction style supported by this desk is more tactile than "direct manipulation" with a mouse, and it seems to have a wide variety of potential applications, well beyond those described in this article. Some issues to be addressed when implementing a digital desk include camera and projector resolution, finger following, good adaptive thresholding, and calibration, but there exist ways to address each of these issues, as illustrated by the

working prototypes described in this article.

This work can be seen as a step toward better integration of paper documents into the electronic world of personal workstations, making paper-based information more accessible to computers. The motivation for this work, however, is just the opposite. The goal is not to enhance computers by giving them better access to paper; it is to enhance paper by giving it better access to computers. There is a difference between integrating the world into computers and integrating computers into the world. The difference lies in our perspective: Do we think of ourselves as working primarily in the computer but with access to physical world functionality, or do we think of ourselves as working primarily in the physical world but with access to computer functionality? Much of the research in human-computer interaction seems to emphasize the former perspective, yet many useful ideas can be gained from the latter. Instead of making *us* work in the computer's world, let us make *it* work in our world.

Acknowledgments

Many have contributed, providing inspiration, help getting equipment to work, and insightful comments on this project and this article. The following people deserve special mention: Stu Card, Kathy Carter, Ian Daniel, Mike Flynn, Steve Freeman, Bill Gaver, Austin Henderson, Mik Lamming, Allan MacLean, Linda Malgeri, Mike Molloy, Tom Moran, William Newman, Peter Robinson, Z Smith and Mark Weiser. Photographs are by Len Dance. 

References

1. Bloomberg, D. and Maragos, P. Image algebra and morphological image processing. In *SPIE Conference Proceedings* (San Diego, Calif., July 1990).
2. Carter, K. Computer aided design: Back to the drawing board. In *Proceedings of Creativity and Cognition*, Loughborough, Apr. 1993.
3. Francik, E., Rudman, S., Cooper, D. and Levine, S. Putting innovation to work: Adoption strategies for multimedia communication systems. *Commun. ACM.* 34, 12 (Dec. 1991).

4. Gaver, W., Moran, T., MacLean, A., Löfstrand, 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 INTERCHI'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

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© ACM 0002-0782/93/0700-086 \$1.50

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 Stasiar 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