

Escritoire: A Personal Projected Display

Mark Ashdown and Peter Robinson
University of Cambridge, UK

We created a system called *Escritoire* that uses two overlapping projectors to create a projected display for a personal computer. A large low-resolution region fills an entire desk while a high-resolution region accommodates the user's focus of attention. The system works in real time and can be used by one person at a desk or by remote participants to create a shared visual space.

The face of personal computing with its desktop metaphor has remained largely unchanged since the release of the Xerox Star in 1979. We created a system called the *Escritoire* to explore the issues in making and using a display that is much more like a real desk—the every day workspace that people use. Our system differs from the conventional interface in a number of ways: it provides much more space, supports the affordances that make paper so useful, uses input techniques different from those of the standard interface, and supports remote collaboration. Please see the sidebar “Conventional Interfaces” (on p. 36) to compare our system with older existing methods.

Escritoire

Escritoire uses the overlapping displays from multiple digital projectors to make a horizontal display, which is as large as a traditional desk but still has high resolution in the region where it's needed, close to the user (see Figure 1). Thanks to reductions in the prices of digital projectors, increases in the power of commodity video cards, as well as the ability to drive the system from a single desktop computer, it's financially and practically feasible as an alternative to the conventional computer interface.

Large vertical screens are often used for presentations or visualization, but *Escritoire's* display is closer to horizontal and thus has different affordances. It acts more like an architect's drafting board, allowing the user to sit comfortably and peer over documents as if they were papers on a desk. The items displayed on the desk are sheets of virtual paper, and documents from

standard application programs can be printed to this virtual form in the same way as they would be printed to physical paper. Two pens provide bimanual input over the entire desk area, and a client-server architecture means that you can link multiple desks to allow remote collaboration in a task space.

Projected desk display

For the desk display we created a horizontal desk surface with a display and interaction area of 36×48 inches. The size and cost of projectors has been falling for some time but the number of pixels they provide has not risen as rapidly. Projectors are now available for around \$2,000 that output 1,400 lumens of light, weigh around one kilogram, and have a resolution of $1,024 \times 768$ pixels. The mass market for such devices to be used in meeting rooms and lecture theaters allows projectors at this resolution to be manufactured cheaply, but moving to higher resolutions quickly becomes prohibitively expensive. To create a display that fills a desk but also has high enough resolution near the user to render a life-sized document legible, we combined two projectors to create what we call a *foveal display*, which has a large, low-resolution periphery that fills the desk and a small, high-resolution fovea for performing detailed work.

Because the light from the fovea projector spreads over a smaller area than that from the periphery projector, its image appears brighter. This hasn't been a problem, and we even consider it an advantage, as we describe in the “User trials” section. Figure 2 shows how the projectors are arranged. Baudisch et al.¹ made a vertical screen for displaying a conventional user interface that also uses this multiresolution technique, but they used an LCD panel for the high-resolution area and a projector for the surrounding region.

Various groups have made multiprojector display walls for scientific visualization in which the room behind a back-projected display houses an array of projectors, a cluster of rendering and computation nodes, and a high-speed network. Our system is the opposite of this in terms of price and size. Rather than being an expensive installation for groups of people to book time on for limited periods, it's a personal projected display, suitable as the interface to a personal workstation. A vertical wall display is useful for presentations or lectures where the user can make small changes, then step back to survey the

results and let onlookers view the changes. However, sitting at a desk display lets users perform detailed work as they would on a normal desk with physical sheets of paper.

Projector calibration

When a projector is aligned so that it projects orthogonally onto a flat surface, the resulting image is rectangular. In oblique projection, where the projector is not aligned to the surface, the image becomes distorted. Requiring orthogonal projection would severely restrict the placement of the projectors and would necessitate precise mechanical adjustment, so we correct the projected image's distortion by warping the image before it's displayed. If we assume a pinhole model for the projector that's the opposite of a pinhole camera, then the image will be distorted by a projective transformation of the following form:

$$\begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} = \mathbf{H} \begin{bmatrix} x \\ y \\ w \end{bmatrix}$$

where \mathbf{H} is a 3×3 matrix called a *planar homography* and the vectors are homogeneous points on the 2D surfaces of the desk and projector frame buffer. We can calculate this homography using a closed-form, least-squares solution from four or more point correspondences. We obtain the homography by projecting a grid of nine known points from each projector and selecting them with a pen input device, then apply the inverse transform to the projected image to correct the distortion. The calibration procedure for both projectors takes about one minute. To avoid projecting twice onto any part of the desk we transform the foveal region into the periphery's frame buffer then black out the resulting quadrilateral.

PCs now have powerful 3D video cards as a standard that can perform projective transformations on images. The image is prepared in the texture memory of the video card then warped by texture mapping it onto a quadrilateral. We can manipulate the x , y , and homogeneous w values for the four vertices to adjust the warping and set them by using the appropriate homography. The two main application programming interfaces for exploiting 3D video hardware in this way are DirectX or OpenGL.² Commodity video cards can perform projective transformations quickly: The cards we tested can warp a $1024 \times$



Figure 1. *Escritoire* is a desk-sized display where users can manipulate documents and images.

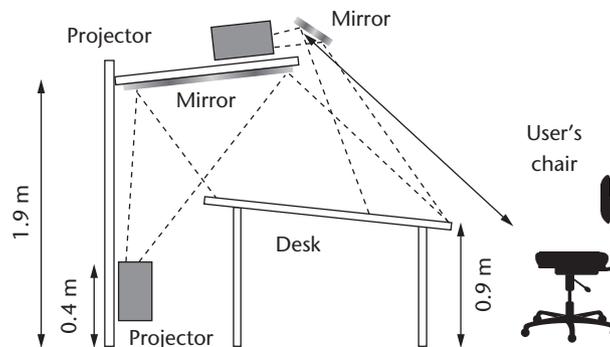


Figure 2. *Dual-projector arrangement of the foveal display.*

768 image in 0.3 milliseconds. Updating the texture is an issue because the hardware is designed to draw large numbers of polygons with static textures, rather than to update the textures rapidly, so we have optimized our system to minimize the amount of updating that's necessary. We achieve 30 frames per second for the two-projector display driven by a single dual-head video card. The pinhole assumption is often a poor one for cameras, especially cheap webcams, but we found that it works well for our projectors.³

Interface

The DigitalDesk,⁴ and subsequent projects, addressed the prevalence of paper by augmenting paper with projected graphics, but this has various problems, particularly for remote collaboration. In that case the paper only exists at one of

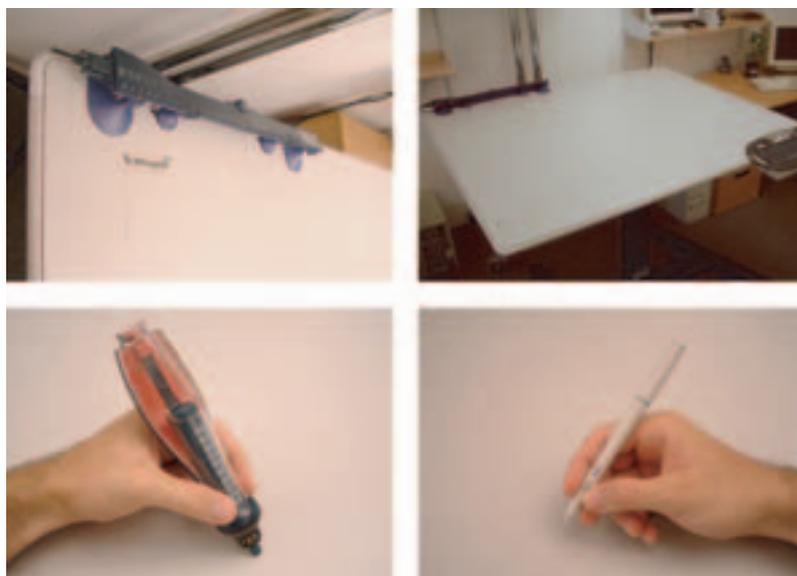


Figure 3. We combined (bottom left) an ultrasonic pen and (top left) its receiver bar with (top right) a large digitizer and (bottom right) its pen. The receiver bar attaches horizontally to the digitizer's back edge.

manipulated (Figure 1). After the normal features of real paper—which might be called literal functionality—have been provided, magical functionality can be added, such as sorting, searching, instant transmission, and new interface techniques that present an animated GUI.

Bimanual input

Our system required two-handed input over the entire desk. Commercial devices that allow simultaneous bimanual input over a desk-sized surface aren't available, so we used a large digitizer with a pen for the desk and combined it with an ultrasonically tracked whiteboard pen (Figure 3). One tracking system is based on electromagnetism and the other on ultrasound so they don't interfere. The digitizer pen is

the collaborating sites, forcing an asymmetry in the interaction of the participants. We have opted instead for a fairly literal emulation of a real desk, where sheets of virtual paper are arranged and

thinner and more accurate so we assigned it to the user's dominant hand, and gave the chunky and less accurate ultrasonic pen to the nondominant hand. Table 1 (on p. 38) summarizes the dif-

Conventional Interfaces

We used four main themes when designing the *Escritoire*: large displays using new technology (as opposed to normal monitors), affordances of physical paper that cause it to be as popular as ever, alternative interaction techniques to the keyboard and mouse, and remote collaboration of a different style to the standard videoconference.

Display size

With a computer monitor, the conventional display space is 9×12 inches. If you imagined this space as an actual desk, working with multiple documents on this desk would be annoying because there's only room for one sheet on top at a time, but this lack of space is common in GUIs where only one window can be usefully displayed. A conventional display fills about 20 to 40 degrees at the center of the visual cone and is meant to be read without rotating the neck. The much larger space available on a traditional desk's top exercises the viewer's peripheral vision and permits a different style of work than that which is possible with the desktop metaphor.

For decades, computers have almost exclusively used screens with diagonal measurements of 14 to 21 inches, but now there's much interest in small-screen, mobile devices. However, in fixed locations such as offices, full advantage can be taken of devices at the opposite end of the scale: very large displays. The use of multiple monitors is becoming popular and techniques are being developed to allow the existing desktop metaphor to be exploited on such systems, but the graphical workspace requires partitioning because of technical constraints; clearly a continuous display would be more desirable.

Work on display devices such as the large vertical screen of the Xerox Liveboard has required many tacit assumptions of the traditional GUI to be revised. Also, large displays have extra issues to address, for instance when a large amount of space is available, the arrangement of items becomes more impor-

tant. As Miller¹ notes, the location of information in the world is important. The fact that space is unimportant to modern information processing systems shouldn't mean that spatial arrangement is ignored; rather, it should mean that we leave it fully to the user's disposition.

Affordances of paper

The affordances of an object are the perceived and actual properties that determine how we use it. The rise of the personal computer brought predictions of the paperless office, but that dream hasn't been forthcoming because paper has affordances that electronic systems haven't surpassed. Paper is more than just a convenient way to display sequential pages of text. Papers on a desk are cognitive artifacts. They remind us of accessible information and pending tasks. They're knowledge in the world that complements knowledge in the head.

Peripheral vision and kinesthetic sense are important for someone using a real desk, particularly for early design work where a designer will place tools, sketches, and other material around a central work area within which the design is formed. A recent study² showed that kinesthetic cues aid spatial memory. Participants were better able to remember the location of objects on the screen when they used a touch screen rather than a mouse, because they were helped by the memory of their hands' positions.

Input techniques

A new interface based on the affordances of paper requires new interaction techniques. Also, a display that's much larger than a standard monitor is qualitatively as well as quantitatively different. Simply displaying a conventional GUI on a large screen does not work.³ For instance, menu bars become difficult to use when they're a long way from the user's center of attention, and text displays become cumbersome when the head must be rotated to see the whole display.

The mouse is a good input device for a conventional monitor where it can provide reasonable accuracy and still allow the entire display to be accessed with a single hand movement. A large display requires something different like direct input with a pen, which is a natural device and has the advantage of kinesthetic cues. On a large display control and feedback should be centered on the area the user is focusing on, which can be approximated with the pen location.

Bimanual input (using two hands) has been shown to have manual and cognitive benefits, but continuous input from two hands isn't used in the conventional interface. A computer interface that uses bimanual input should respect the difference between the dominant and nondominant hands that occurs in other tasks. For instance, when writing, the nondominant hand holds and periodically repositions the paper so that the dominant hand can stay in a limited area while it moves the pen. A study on the use of electronic and paper documents for a sum-

marization task⁴ found that the navigation of paper documents is fast and automatic unlike the electronic case. The advantages of paper include the use of two hands to overlap navigation with other activities, anticipatory page turning, the ability to lay out paper in space, and reading and writing spaces that can be accessed concurrently and independently. Fixity of information relative to physical pages is also important for navigation because it lets the reader acquire incidental knowledge of the location of information.

Collaboration

Buxton⁵ uses the term *person space* for the type of videoconferencing where participants see each other's faces via cameras and *task space* for when they share a virtual space where they can both interact. The intuitive appeal of conventional video communication that creates a person space prompted forecasts of its wide-scale adoption—the Picturephone from Bell Labs, for example, was introduced publicly in 1964 and at the time was predicted to replace the existing voice-only telephone by the early 1970s. However, except for limited use in business settings, it has not become a substitute for face-to-face meetings. A major advantage of a desk that holds virtual rather than real items is that the desk surface and its contents can be shared with people at remote sites to allow collaboration in a task space. Krueger's Videodesk⁶ was a prescient example of this type of remote collaboration. We believe that collaboration in a task space, such as that provided by two linked Escriitoare desks, will be more useful than a conventional videoconference for many tasks, and studies have demonstrated activities for which allowing participants to share a task domain is more useful than letting them see each other's faces.⁷

References

1. G.A. Miller, "Psychology and Information," *American Documentation*, vol. 19, no. 3, 1968, pp. 286-289.
2. D.S. Tan et al, "Kinesthetic Cues Aid Spatial Memory," *Proc. Computer-Human Interaction: Extended Abstracts on Human Factors in Computing Systems*, ACM Press, 2002, pp. 806-807.
3. K. Swaminathan and S. Sato, "Interaction Design for Large Displays," *Interactions*, vol. 4, no. 1, 1997, pp. 15-24.
4. K. O'Hara and A. Sellen, "A Comparison of Reading Paper and Online Documents," *Proc. Special Interest Group on Computer-Human Interaction Conf. Human Factors in Computing Systems*, ACM Press, 1997, pp. 335-342.
5. W. Buxton, "Telepresence: Integrating Shared Tasks and Personal Spaces," *Proc. Conf. Graphics Interface*, Morgan Kaufmann, 1992, pp. 123-129.
6. M.W. Krueger, "Environmental Technology: Making the Real World Virtual," *Comm. ACM*, vol. 36, no. 7, 1993, pp. 36-37.
7. A.H. Anderson et al., "Video Data and Video Links in Mediated Communication: What Do Users Value?" *Int'l J. Human-Computer Studies*, vol. 52, no. 1, 2000, pp. 165-187.

Table 1. Differences between pens for the user's dominant and nondominant hands.

	Dominant	Nondominant
Resolution	High	Low
Cost	High	Low
Buttons	3	1
Grip of pen	Sleek	Chunky



Figure 4. Web browser being accessed through Virtual Network Computing and placed on the desk with other documents.



Figure 5. Piles of virtual paper. The pile splits open to reveal its contents as the pen moves over it.

ferences. We've allowed the nondominant hand to perform coarse tasks such as positioning sheets on the desk display, and gave fine tasks such as writing to the dominant hand. The nondominant hand is just as good at selecting targets as the dominant hand when the distances and targets are large.⁵

Virtual paper

We implemented sheets of virtual paper of several types on the *Escritoire*. Users can place PDF documents on the desk, annotate them with the digitizer pen, and save the annotations in the file so that they can be viewed later with a standard PDF viewer. Instead of printing to physical paper from an application like a word processor or spreadsheet, users can print to PDF and then read the document on the desk display.

Bitmapped images like JPEGs can be placed on the desk and annotated, and the changes are saved for later viewing. We also created a Virtual Networking Computing client (VNC; <http://www.realvnc.com>) that users can place on the desk display. This lets application programs on virtually any computer be viewed on the desk along with the other items, and controlled with the pen. Figure 4 shows a Web browser, accessed via VNC, being used on the desk.

To save space on the desk and group documents we have added the notion of piles to the interface. The time required to maintain a filing system and the cognitive difficulty of creating appropriate categories for information mean that people often create vaguely classified piles on their desks. The physical arrangement of the documents then reminds the person of tasks to be performed—recognition is easier than recall—and means that the information is easily accessible. Apple's Pile Metaphor⁶ allows icons representing documents to be placed in piles to form queries to an information retrieval system. Our piling system lets sheets of virtual paper be dragged into and out of piles, and to new positions in a pile. As the user moves the digitizer pen over a pile it is animated and splits open to allow browsing (see Figure 5).

Remote collaboration

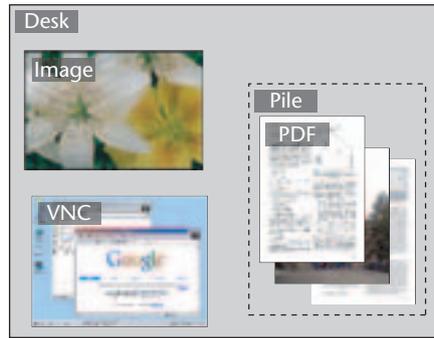
The image warping described previously, and the handling of the input devices, requires low-level, performance-dependent code.

The device-independent code and data that implement the sheets of virtual paper are different, so we split the software into a client and a

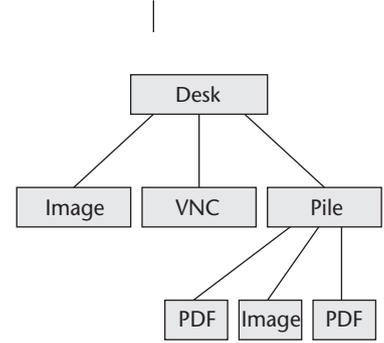
server that communicate with each other over a standard Internet connection. The client program runs on a computer that has the input and output devices of the Escritoire display (the projectors and pens). Several clients can be connected to the same server to allow participants to collaborate around a common set of information. It is a “what you see is what I see” (WYSIWIS) system. Multiple desk displays running the client software can connect to a single server that stores all the system state and makes a shared visual task space available to the collaborators. We use Webcams, microphones, and standard videoconferencing software to augment that task space with a person space based on a video and audio channel, so the participants can speak to each other as they interact and can look up to see each other on monitor screens.

Table 2. The different characteristics for Escritoire’s client and server programs.

Characteristics	Client	Server
Control flow	Sequential	Event driven
Programming language	C++	Java
System dependence	Dependent	Independent
State storage	Stateless	Stateful



(a)



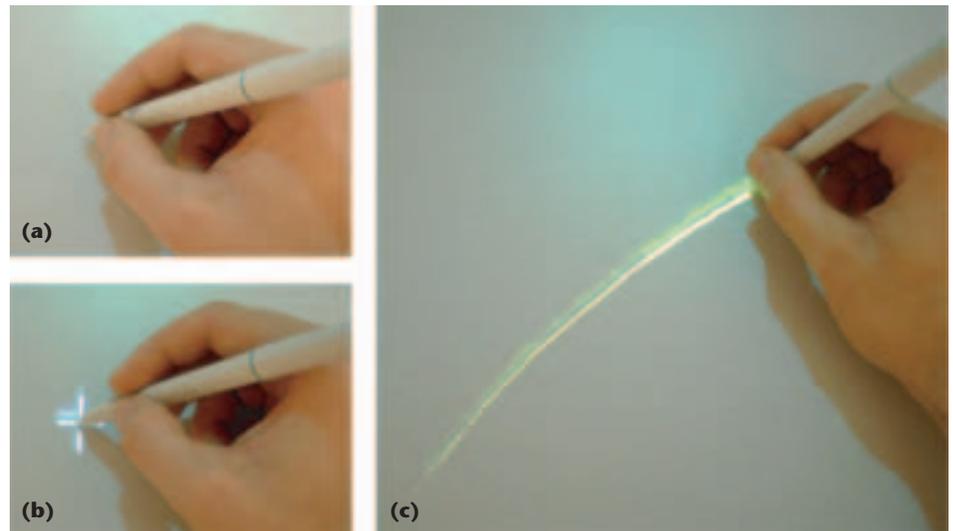
(b)

Figure 6. (a) Documents and other items can be arranged on Escritoire’s desk display. (b) These, including the underlying desk surface, are represented by objects stored in the server.

Client–server system

A clear division of labor exists between client and server, and Table 2 summarizes the differences. The client is a hardware-dependent program that handles the pen input devices and performs the real-time graphics warping to compensate for oblique projection. The control flow of this program is simple—it loops around handling input events and updating the display as quickly as possible to obtain maximum performance from the hardware. All of the system’s state is stored on the server, which is a hardware-independent Java program. The server responds to events from the clients connected to it, and sends them updates when a sheet on the desk changes its appearance or location.

To the server, the sheets are objects that might contain complex code to respond to events or do asynchronous processing and updates, but the client simply deals with bitmaps that visually represent the sheets of virtual paper. This model, which resembles that of the X Window System, helps keep a clean divide between client and server, and allows a single-user system to host both client and server on one computer or a multiuser system to have clients distributed across the Internet. Figure 6



(a)

(b)

(c)

Figure 7. We implemented three options for cursors: (a) no cursor, (b) crosshairs that show the current position of the pen, and (c) fading traces that display an animated history of locations.

shows an example arrangement of items on a desk display and the tree of objects that the server would use to represent it.

Cursors

We implemented three cursor options for the pens: no cursor, crosshair, and trace (see Figure 7). The cursors are duplicated on the desks of all participants in a conference, with the crosshair

No extra training was needed to get participants to work collaboratively with our system.

simply showing the pen's position, and the trace displaying an animated history of its previous locations over the last second or so. We found a fading trace of 0.7 second to be effective. Traces improve gestural communication for both creator and viewer,⁷ especially in the presence of network jitter. The shared surface isn't just a medium for drawing and making other permanent changes—much of the interaction will consist of ephemeral gestures between participants that complement their conversation. This is especially true in a domain that's visual rather than textual and where the participants don't have precise words for the items they're showing to each other.

User trials

We obtained feedback from visitors to our computer laboratory who used *Escritoire*. We also conducted two user trials to gain experience and qualitative results from our system: one with single users and one with pairs collaborating between sites.

We conducted single-user tests with seven employees at Thales Research & Technology who funded this work. They were introduced to the desk display and pens and shown how to perform some basic tasks, then they were asked to repeat two sets of tasks. First they circled spelling mistakes in textual documents, then they placed images in piles based on their content. The test took around 45 minutes for each participant, which included answering some general questions about the interface.

We conducted collaborative tests between remote sites about 100 miles apart connected via the Internet. Initially we showed each of the six participants 30 sheets on the desk, each containing information and a photograph of a house. We then paired up the participants to collaborate remotely. The pairs had to find the best

house from groups of ten, and did this three times, once for each of the cursor options in Figure 7. The audio, video, and desk messages were carried on a 256 Kbps digital subscriber line link. Afterwards, the participants rated on a scale of one to five whether they found the audio, video, and desk channels useful. We also asked them whether any problems were caused by the difference in resolution between periphery and fovea, the difference in brightness between periphery and fovea, or the latency of the remote interaction. It took each participant around two hours to complete this trial.

Single-user results

Participants only needed a few minutes of training to use the system, even though they had never used it before. They could easily use the large display area because they could survey the whole desk surface at a glance to see which documents were available and could quickly reach out and drag one to the fovea. In the single-user case, participants overwhelmingly preferred to have no cursors following the pens—unlike a relative pointing device like a mouse, cursors are not needed for a direct pointing device like the pen, so they just get in the way.

Various issues regarding the pens were mentioned. The pen buttons are difficult to press in particular combinations, so we adjusted the interface to make those combinations unnecessary. The angle at which the pen is held can affect the reported position, although this effect is not so pronounced for the more accurate digitizer pen. Also, unlike the mouse, the pen can be lifted from the surface then put down in a different location, so we added lift events to inform the server when the pen is moved away from the surface.

The use of these events, which are not present in a conventional GUI, mean that the pile browsing feature shown in Figure 5, for example, is turned off when the pen is lifted from the surface causing the pile to revert to its normal form. Front projection is generally problematic for vertical screens because a presenter can obscure the image by walking in front of the projector, but because the image on the *Escritoire*'s display is projected obliquely from the back of the desk the user can lean considerably forward without obscuring the display. Some obscuring does occur around the hands of the user but no participants complained about it. We believe this is because people are accustomed to items being lit from above so they move their hands automati-

cally when they shadow something.

Errors in pen registration due to parallax differences are common on back-projected displays and tablet PCs that have thick glass screens, but an advantage of front projected displays is that they don't suffer from this problem. We used two different digitizers for the two systems we assembled. These had different surface characteristics with one of them deviating significantly from the ideal white diffuse surface—these devices aren't designed to be illuminated by projectors, so their optical properties should be determined in advance.

Collaborative results

No extra training was needed to get participants to work collaboratively with our system—we simply showed them the videoconference screen and told them to start working with the person at the other site. In contrast to the single-user case, participants much preferred the trace option from Figure 7. The pen traces let participants explicitly gesture to each other, but they're also useful because they allow one participant to view, or simply remain aware of, the actions of the other at all times.

All participants strongly agreed with the statements that the audio and desk channels were useful for the task, but responses to the video channel were much weaker (see Table 3). We believe that for most tasks a task space such as that provided by the shared desk surface will be important, and for many it will be more important than the person space provided by the conventional video channel of a videoconference. The small latency of the interaction and difference in resolution between fovea and periphery weren't a significant problem. The difference in brightness between the two regions was considered helpful by two participants, because it delineates the regions and emphasizes the high-resolution area.

Some issues warrant further investigation. The precise actions assigned to the dominant and nondominant hands could be refined with more trials. Two participants said that they would like to browse through a pile with the nondominant hand while making notes with the dominant hand. This wasn't possible because pile browsing and writing were only available to the dominant hand. Also, private workspaces would be useful, where a user can keep material that shouldn't be seen by remote participants.

Table 3. Responses from the six participants of the collaborative tests.

Test Statements	Stongly Agree	Agree	Neither	Disagree	Strongly Disagree
Audio useful	6	0	0	0	0
Video useful	0	2	2	0	2
Desk useful	6	0	0	0	0

Future work

In the user trials involving remote collaboration, a start-up period of a few minutes was required for each task to download the bitmap data for the many sheets on the desk from server to client. We have since added lossless compression, which works well for the PDF documents we use (which are mostly text). The bitmap data for them is reduced to around 2 percent of the documents' original size, thus reducing the start-up time to a few seconds. We're also adding functionality to make it easy for participants in a collaborative session to drag PDF documents and images from their laptops or desktop machines onto the desk in one simple movement. A participant in a collaborative session will run a small program on their laptop that will accept files and transmit them to the server, which will then add them to the shared surface. This will also provide a private workspace in which the participant can keep items until they're needed.

Graphics cards now commonly have 256 Mbytes of memory, which at 72 dots per inch and 16 bits per pixel is enough to store more than 200 A4 sheets. The sheets on the desk could be stored in separate textures, and then effects like rotating and zooming could be implemented easily with virtually no extra computational cost. Luminance and chrominance matching for multiprojector displays are active areas of research. Luminance correction is undesirable for the *Escritoire's* display because the brightness of the fovea would have to be greatly reduced to match it with the periphery. This would waste most of the power of the projector, but chrominance matching could be exploited.

We experimented with extra projectors to create large displays on the walls of an office to increase the amount of space available to arrange documents. A device we call a wand controls desk and wall displays from a distance by pointing at items, thus documents can be stored on the wall display until they are needed, as if they were on a bookshelf. We used a magnetic tracker for our wand, but it is disrupted by metal objects so an ultrasonic tracker or laser pointer might be

a better solution.

The Everywhere Displays project⁸ is exploring the use of a steerable projector to create movable displays controlled by specialized computer vision. The wall displays complement the Escritoire's desk display by providing a private space for users sharing a desk surface. The wall displays also extend the concept of a hierarchy of displays where the larger ones are further from the user and more coarsely rendered and controlled.

We're continuing this project by working with several partners, each of whom will have at least one instance of the Escritoire hardware. The Escritoire technology will be applied to applications in crisis management and collaborative design. We hope to reduce the amount of travel necessary between the partners by collaborating via the networked desks. Having more than two sites will allow us to link three or more desks together, which will prompt new insights and developments. **MM**

Acknowledgments

Thales Research & Technology UK funded this work and helped with the construction of the hardware and the user trials. We gratefully acknowledge their involvement.

References

1. P. Baudisch, N. Good, and P. Stewart, "Focus Plus Context Screens: Combining Display Technology with Visualization Techniques," *Proc. 14th Ann. ACM Symp. User Interface Software and Technology (UIST)*, ACM Press, 2001, pp. 31-40.
2. M. Ashdown and P. Robinson, "Experiences Implementing and Using Personal Projected Displays," *Proc. IEEE Int'l Workshop on Projector-Camera Systems*, 2003; <http://lear.inrialpes.fr/people/triggs/events/iccv03/cdrom/index.php>.
3. M. Ashdown, *Personal Projected Displays*, PhD dissertation, Computer Laboratory, Univ. Cambridge, 2004.
4. P.D. Wellner, *Interacting with Paper on the DigitalDesk*, PhD dissertation, Computer Laboratory, Univ. of Cambridge, 1994.
5. P. Kabbash, I.S. MacKenzie, and W. Buxton, "Human Performance Using Computer Input Devices in the Preferred and Non-Preferred Hands," *Proc. ACM Conf. Human Factors in Computing (InterCHI)*, ACM Press, 1993, pp. 474-481.
6. R. Mander, G. Salomon, and Y.Y. Wong, "A 'Pile' Metaphor for Supporting Casual Organization of Information," *Proc. Special Interest Group on Computer-Human Interaction Conf. Human Factors in Computing Systems*, ACM Press, 1992, pp. 627-634.
7. C. Gutwin, "Traces: Visualizing the Immediate Past to Support Group Interaction," *Proc. Graphics Interface*, 2002, AK Peters, pp. 43-50.
8. G. Pingali et al., "Steerable Interfaces for Pervasive Computing Spaces," *Proc. First IEEE Conf. Pervasive Computing and Comm.*, IEEE CS Press, 2003, pp. 315-322.



Mark Ashdown currently holds a postdoctoral fellowship at the University of Tokyo Institute of Industrial Science in Japan. His work focuses on large projected interfaces and calibration for

such displays, including the use of computer vision, input devices, graphics, and human-computer interaction. He recently completed his PhD in Computer Science at the University of Cambridge Computer Laboratory.



Peter Robinson is the Deputy Head at the University of Cambridge Computer Laboratory in England, where he previously received his undergraduate degree in mathematics and PhD

in computer science. His research explores new technologies and techniques for human-computer interaction, with an emphasis on the use of video projectors and cameras in the user interface.

Readers may contact Mark Ashdown at mark@ashdown.name.

For further information on this or any other computing topic, please visit our Digital Library at <http://www.computer.org/publications/dlib>.