Experience with the Escritoire: A Personal Projected Display

Mark Ashdown, Peter Robinson
University of Cambridge, UK

***Abstract*— We have created a system called the *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. We use commodity graphics cards to warp the projected images in real time to compensate for the unconstrained positions of the projectors. To allow documents to be shared in a visual space between remote participants we have split the software into a hardware-dependent client that handles the input and output, and a platform-independent server that controls the desk contents. We have conducted two sets of user studies, one with single users one with pairs at separate sites.**

***Index Terms*— projector, foveal display, desk display, bimanual input, remote collaboration.**

I. The Conventional Display

The face of personal computing with its desktop metaphor has remained largely unchanged since the release of the Xerox Star in 1979. We have 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 everyday workspace that people use. Our system is different from the conventional interface in a number of ways: it provides much more space, it has been designed to support the affordances that make paper so useful, it uses input techniques different from those of the standard interface, and it supports remote collaboration.

A. Space

Figure 1 depicts an imaginary 9 by 12 inch desk. Working with multiple documents on this desk would be annoying because there is only room for one sheet to be on top at a time, but this lack of space is common in graphical user interfaces where only one window can be usefully displayed at a time. A conventional display fills about 20-40 degrees 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 real desktop exercises the viewer's peripheral vision and permits a different style of work to that which is possible on the desktop metaphor.

For decades computers have almost exclusively used screens with diagonal measurements of 14 to 21 inches but there is now much interest in small-screen, mobile devices. 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 partitioning of the graphical

Fig. 1. The 9 by 12 inch desk. It is too small to work on, yet this lack of space is common in graphical user interfaces.

workspace is due to technical constraints and a continuous display will 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 graphical user interface 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 important. George A. Miller noted that the location of information in the world is important [1]. The fact that space is unimportant to modern information processing systems should not mean that spatial arrangement is ignored, but rather that it is placed wholly at the disposition of the user.

B. Affordances of Paper

The affordances of an object are the perceived and actual properties that determine how it is used. The rise of the personal computer brought predictions of the paperless office, but that dream has not been forthcoming because paper has affordances that have not been surpassed by electronic systems. Paper is more than just a convenient way to display sequential pages of text. Papers on a desk are cognitive artifacts. They remind their owner of information that is available and of pending tasks. They are knowledge in the world that complements knowledge in the head. Peripheral vision and kinaesthetic sense are important for someone using a real desk, and 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. In a recent study [2] kinesthetic cues were shown to aid spatial memory. Participants were better able to remember the location of objects on the screen when they used a touchscreen rather than a mouse, because they were helped by the memory of the positions of their hands.

C. Input Techniques

A new interface based on the affordances of paper requires new interaction techniques. Also, a display that is much larger than a standard monitor is qualitatively as well as quantitatively different. Simply displaying a conventional graphical user interface on a very large screen does not work [3]. For instance, menu bars become difficult to use when they are 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 do 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 mentioned above. On a large display control and feedback should be centered on the area to which the user is attending, which can be approximated by the pen location.

Bimanual input—using two hands—has been shown to have manual and cognitive benefits, but continuous input from two hands is not used in the conventional interface. A computer interface that does use bimanual input should respect the difference between the dominant and non-dominant hands that occurs in other tasks. For instance, when writing the non-dominant 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 summarization task [4] 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 allows the reader to acquire incidental knowledge of the location of information.

D. Collaboration

Buxton [5] uses the term *person space* for the type of videoconferencing where participants see each other's faces via cameras, and *task space* for the situation where they share a virtual space with which 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 was introduced publicly in 1964 and at the time was predicted to replace the existing voice-only telephone by the early 1970s—but, 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 VIDEODESK [6] 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 Escritoire 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].

E. The Escritoire

We have developed a system called the Escritoire that uses the overlapping displays from multiple digital projectors to make a horizontal display that is as large as a traditional desk, but which still has high resolution in the region where it is needed, close to the user. We call this combination of projectors a *foveal display* (Figure 2), although unlike a head-mounted version the difference between regions is intended not to be imperceptible but to focus the user's attention. Thanks to reductions in the prices of digital projectors and increases in the power of commodity video cards, and to the fact that the system can be driven from a single desktop computer, it is financially and practically feasible as an alternative to the conventional computer interface.

Large vertical screens are often used for presentations or visualization, but the Escritoire's display is closer to horizontal and thus has different affordances. It acts more like an architects 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 area of the desk, and a client-server architecture means that multiple desks can be linked to allow remote collaboration in a task space.

Several projects, starting with the DigitalDesk [8], have explored the idea of a horizontal projected display, although they have tended to be relatively small, single-projector systems. Another paper on the Escritoire contains a review [9]. Recently the falling size and cost of projectors has prompted much work on multi-projector displays, new calibration methods, and novel applications. The Procams workshop on projector-camera systems (`http://www.procams.org/`) is a good example of work in this area.

Next we will describe the Escritoire's foveal display and how it is calibrated. After that, the interface with its bimanual input and virtual paper, the client-server architecture that allows remote collaboration, and user trials with single users and collaborating pairs. We will finish with directions for future work and conclusions.

II. Projected Desk Display

We have created a horizontal desk surface with a display and interaction area of 36×48 inches (approximately A0 size). The size and cost of projectors has been falling for some time

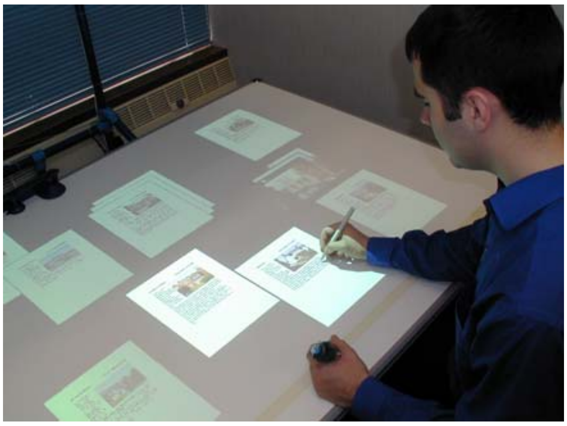

Fig. 2. The Escritoire is a desk-sized display on which items like documents and images can be manipulated.

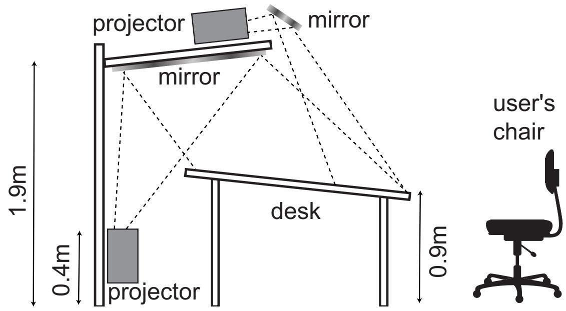

Fig. 3. The dual-projector arrangement of the foveal display.

but the number of pixels that they provide has not risen so rapidly. Projectors are now available for around $2000 that output 1400 lumens of light, weigh around one kilogram, and have a resolution of 1024 by 768 pixels. The mass market for such devices to be used in meeting rooms and lecture theatres 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 have 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* in which detailed work can be performed. Because the light from the fovea projector is spread over a smaller area than that from the periphery projector its image appears brighter, but we have found this not to be a problem, and it can even be considered an advantage, as described in the User Trials section below. Figure 3 shows how the projectors are arranged. Baudisch *et al.* [10] have made a vertical screen for displaying a conventional user interface that also uses this multi-resolution technique, but they use an LCD panel for the high-resolution area and a projector for the surrounding region.

Various groups have made multi-projector 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, but 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 is 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, but sitting at a desk display a user can perform detailed work as one would on a normal desk with physical sheets of paper.

A. Projector Calibration

When a projector is aligned so that it projects orthogonally onto a flat surface the resulting image is rectangular, but 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 have chosen to correct distortion of the projected image by warping the image before it is displayed. If we assume a pin-hole model for the projector, that is the opposite of a pin-hole 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 in the 2D surfaces of the desk and the projector framebuffer. This homography can be calculated 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 framebuffer of the periphery then black out the resulting quadrilateral.

PCs now have powerful 3D video cards as standard which 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. The x, y, and homogeneous w values for the four vertices can be manipulated to adjust the warping and are set using the appropriate homography. Video cards can be exploited through DirectX or OpenGL [11]. Commodity video cards can perform projective transformations very quickly: the cards we have tested can warp a 1024×768 image in under 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 is necessary. We achieve 30 frames per second for the two-projector display driven by a single dual-head video card. The pin-hole assumption is often a poor one for cameras, especially cheap webcams, but we have found that it works well for the projectors we have used [9].

III. Interface

The DigitalDesk [8], 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 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 manipulated (Figure 2). 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 graphical user interface.

A study of the reading of online and paper documents [4] describes various advantages of paper: support for annotation while reading, quick navigation, and flexibility of spatial layout. It suggests the use of a larger screen and a large virtual workspace with overview and detail renderings, but the lack of spatial continuum between the focus of attention and the periphery in that case is given as a disadvantage of the method over the continuous display of a large desk.

A. Bimanual Input

One requirement of our system was two-handed input over the entire desk. Commercial devices that allow simultaneous bimanual input over a desk-sized surface are not available, so we used a large digitizer for the desk and combined it with an ultrasonically-tracked whiteboard pen (Figure 4). One tracking system is based on electromagnetism and the other on ultrasound so they do not interfere with each other. The digitizer pen is thinner and more accurate so we have assigned it to the user's dominant hand, and given the chunky and less accurate ultrasonic pen to the non-dominant hand. The differences are summarized in Table I. We have allowed the non-dominant hand to perform coarse tasks such as positioning sheets on the desk display, and have given fine tasks such as writing to the dominant hand—the non-dominant hand has been shown to be just as good at selecting targets as the dominant hand when the distances and targets are large [12].

TABLE I
DIFFERENCES BETWEEN THE PENS FOR THE USER'S DOMINANT AND NON-DOMINANT HANDS.

	Dominant	Non-dominant
resolution	high	low
cost	high	low
buttons	3	1
grip of pen	sleek	chunky

B. Virtual Paper

We have implemented several types for the sheets of virtual paper on the Escritoire. PDF documents can be placed on the desk and annotated with the pen, and the annotations are saved 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 one

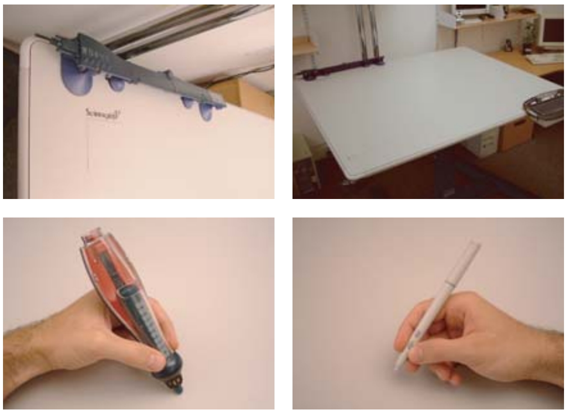

Fig. 4. We have combined an ultrasonic pen and its receiver bar (left) with a large digitizer (right). The receiver bar is attached horizontally to the back edge of the digitizer. The non-dominant hand performs coarse tasks with the ultrasonic pen, while the dominant hand performs more detailed work with the more accurate digitizer pen.

Fig. 5. A web browser is being accessed through VNC, and is placed on the desk with other documents.

prints to PDF then reads 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 have also created a VNC (`www.realvnc.com`) client that can be placed on the desk display. This allows application programs on virtually any computer to be viewed on the desk along with the other items, and controlled with the pen. Figure 5 shows a web browser, accessed via VNC, being used on the desk.

To save space on the desk and allow documents to be grouped 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 causes the person to be reminded of tasks to be performed—recognition is easier than recall—and means that the information is easily accessible. Apple's Pile Metaphor [13] allowed icons representing documents to be placed in piles to form queries to an information retrieval system. Our piling system allows sheets of virtual paper to be dragged into and out of

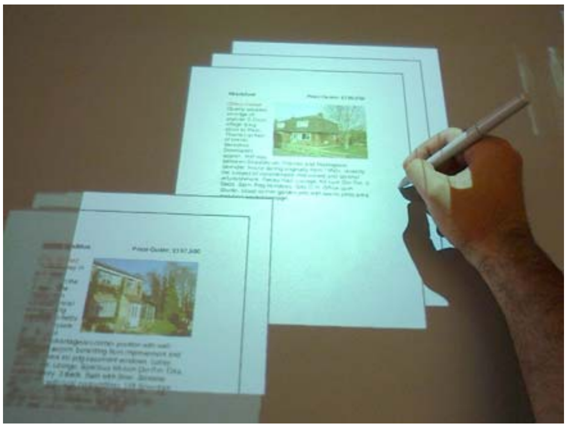

Fig. 6. Piles of virtual paper: the pile splits open to reveal it's contents as the pen is moved over it.

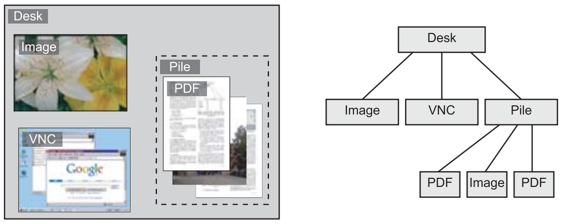

Fig. 7. Documents and other items can be arranged on the desk display of the Escritoire (left). These, including the underlying desk surface, are represented by objects stored in a tree on the server (right).

piles, and to new positions in a pile. As the digitizer pen is moved over a pile it is animated and splits open to allow browsing (Figure 6).

IV. Remote Collaboration

The image warping described above, 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 very different, so we have split the software into two parts, 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—and several clients can be connected to the same server to allow participants to collaborate around a common set of information. It is a WYSIWIS (What You See Is What I See) system. Multiple desk displays running the client software can connect to a single server that stores all the state of the system and makes a shared visual task space available to the collaborators. We have used 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.

A. Client-Server System

There is a clear division of labor between client and server, and the differences between them are summarized in Table II. 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 state of the system 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 may contain complex code to respond to events or do asynchronous processing and updates, but the client simply deals with bitmaps that are the visual representations of the sheets of virtual paper. This model, which is similar to that of the X Window System, helps keep a clean divide bewteen client and server, and allows a single-user system to host both client and server on one computer or a multi-user system to have clients distributed across the Internet.

Figure 7 shows an example arrangement of items on a desk display, and the tree of objects that the server would use to represent it.

TABLE II
The different characteristics of the Escritoire's client and server programs.

	Client	Server
control flow	sequential	event driven
programming language	C++	Java
system dependence	dependent	independent
state storage	stateless	stateful

B. Cursors

We have implemented three cursor options for the pens: no cursor, cross hair, and trace (Figure 8). The cursors are duplicated on the desks of all participants in a conference, with the cross hair simply showing the position of the pen, and the trace displaying an animated history of it's previous locations over the last second or so—we have found a fading trace of 0.7 seconds to be effective. Traces have been shown to improve gestural communcation for both creator and viewer [14], especially in the presence of network jitter. The shared surface is not 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 is very visual rather than textual, and where the participants do not have precise words for the items they are showing to each other.

V. User Trials

We have obtained feedback from visitors to the Computer Laboratory who have used the Escritoire. We have also conducted two user trials to gain experience and qualitative results from our system: one with single users, and one with pairs collabrating between sites.

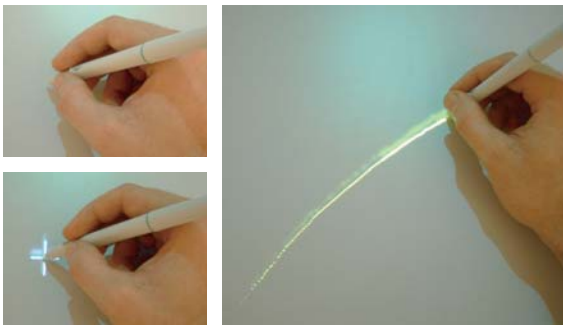

Fig. 8. We have implemented three options for cursors: no cursor, cross hairs that show the current position of the pen, and fading traces that display an animated history of locations.

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 trail 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 each of the six participants was shown 30 sheets on the desk, each containing information and a photograph of a house, then they were paired up and collaborated remotely. The pairs were given the task of finding the best house from groups of ten, and did this three times, once for each of the cursor options in Figure 8. The audio, video, and desk messages were carried on a 256kbps DSL link. Afterwards, the participants were asked to rate, on a scale of one to five, whether they found the audio, video, and desk channels useful, and they were asked 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. Each participant was occupied for around 2 hours with this trial. The full results of the trials [9] have been summarized below.

A. *Single-User Results*

Participants only needed a few minutes of training in order to use the system, even though they had never used it before. They could easily make use of 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 highlighted. 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 graphical user interface, mean that, for instance, the pile browsing feature shown in Figure 6 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 forward a considerable way 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 automatically when they are shadowing 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 do not 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 are not designed to be illuminated by projectors so their optical properties should be determined in advance.

B. *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 8. The pen traces allow participants to explicitly gesture to each other, but they are 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 (Table III). We believe that for most tasks a task space such as that provided by the shared desk surface will be very 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 were not 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 non-dominant hands could be refined with more trials—two participants said that they would have liked to have browsed through a pile with the non-dominant hand while making notes with the dominant hand. That was not 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 should not be seen by the remote participants.

TABLE III
RESPONSES FROM THE SIX PARTICIPANTS OF THE COLLABORATIVE TESTS.

	strongly 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

VI. 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 very well for the PDF documents we have used which are mostly text. The bitmap data for them is reduced to around 2% of its original size, thus reducing the start-up time to a few seconds. We are 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 are needed.

Graphics cards now commonly have 256 Megabytes of memory, which at 72 dots per inch and 16 bits per pixel is enough to store over 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 multi-projector 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, which would waste most of the power of the projector, but chrominance matching could be exploited.

We have 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* is used to control 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 have used a magnetic tracker for our wand but it is disrupted by metal objects so an ultrasonic tracker or laser pointer may be a better solution. The Everywhere Displays project [15] 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 that are sharing a desk surface, and extend the concept of a hierarchy of displays where the larger ones are further from the user and more coarsely rendered and controlled.

We will be 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 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.

VII. CONCLUSION

We have described a system called the Escritoire that has a desk-sized horizontal display with bimanual input that allows a different style of work to that possible on the conventional computer interface with it's small vertical screen, keyboard, and mouse. Two projectors overlap to create a foveal display that fills the desk but also has high resolution where it is needed. This personal projected display is feasible as the output device for a personal computer—the computer, video card, projectors, and pen devices are all standard components and collectively cost under US $10,000. Also, the cost of the most expensive part, the projectors, is falling. The whole front-projection system requires little extra floor space over that of the desk, fits below a normal office ceiling, and can be used under normal lighting.

Current graphics hardware can deliver the performance necessary to warp the imagery in real time, providing at least 30 frames per second for two projectors connected to a single graphics card. We have used 1024×768 projectors to get a resolution of approximately 62 dots per inch in the fovea, but projector resolutions are increasing and 1600×1200 projectors should become affordable in the next few years. These would give the fovea of our display a resolution similar to that of current LCD monitors.

Tests participants were able to use the system after only a few minutes of practice—it is easy to forget how long it takes to be able to confidently use a keyboard and mouse but the timescale will be in days or weeks rather than minutes. They found it easy to perform actions with the non-dominant hand that do not require high accuracy, such as moving a sheet. Collaborating users appreciated the ability to gesture to each other that the pen traces gave them, and they found the task space of the desk more useful than the person space of the video conference for the task we gave them.

The Escritoire is constructed from standard components and exploits users' existing manual skills to form a personal projected display for performing everyday tasks for which people traditionally use their desks. In fixed locations, such as offices, where space and mobility are not limiting factors, large-format interfaces will become popular and the Escritoire is an example of one with a lower price and requiring less space than existing multi-projector display walls, and which has different affordances due to it's horizontal configuration.

ACKNOWLEDGMENT

Thales Research & Technology UK funded this work and helped with the construction of the hardware and the user trials. Their involvement is gratefully acknowledged.

REFERENCES

[1] G. A. Miller, "Psychology and Information," *American Documentation*, vol. 19, no. 3, pp. 286–289, 1968.

[2] D. S. Tan, R. Pausch, J. K. Stefanucci, and D. R. Proffitt, "Kinesthetic Cues Aid Spatial Memory," in *Extended Abstracts of CHI 2002*, 2002, pp. 806–807.

[3] K. Swaminathan and S. Sato, "Interaction Design for Large Displays," *Interactions*, vol. 4, no. 1, pp. 15–24, 1997.

[4] K. O'Hara and A. Sellen, "A Comparison of Reading Paper and On-line Documents," in *Proceedings of CHI 97*, 1997, pp. 335–342.

[5] W. Buxton, "Telepresence: Integrating Shared Task and Person Spaces," in *Proceedings of Graphics Interface '92*, 1992, pp. 123–129.

[6] M. W. Krueger, "Environmental Technology: Making the Real World Virtual," *Communications of the ACM*, vol. 36, no. 7, pp. 36–37, 1993.

[7] A. H. Anderson, L. Smallwood, R. MacDonald, J. Mullin, and A. Fleming, "Video Data and Video Links in Mediated Communication: What Do Users Value?" *International Journal of Human-Computer Studies*, vol. 52, no. 1, pp. 165–187, 2000.

[8] P. D. Wellner, "Interacting with Paper on the DigitalDesk," Ph.D. dissertation, University of Cambridge Computer Laboratory, 1994.

[9] M. Ashdown, "Personal Projected Displays," Ph.D. dissertation, University of Cambridge Computer Laboratory, 2004.

[10] P. Baudisch, N. Good, and P. Stewart, "Focus Plus Context Screens: Combining Display Technology with Visualization Techniques," in *Proceedings of UIST 2001*, 2001, pp. 31–40.

[11] M. Ashdown and P. Robinson, "Experiences Implementing and Using Personal Projected Displays," in *IEEE International Workshop on Projector-Camera Systems (Procams 2003)*, 2003.

[12] P. Kabbash, I. S. MacKenzie, and W. Buxton, "Human Performance Using Compuer Input Devices in the Preferred and Non-Preferred Hands," in *Proceedings of InterCHI 93*, 1993, pp. 474–481.

[13] R. Mander, G. Salomon, and Y. Y. Wong, "A 'Pile' Metaphor for Supporting Casual Organization of Information," in *Proceedings of CHI 92*, 1992, pp. 627–634.

[14] C. Gutwin, "Traces: Visualizing the Immediate Past to Support Group Interaction," in *Proceedings of Graphics Interface 2002*, 2002, pp. 43–50.

[15] G. Pingali, C. Pinhanez, A. Levas, R. Kjeldsen, M. Podlaseck, H. Chen, and N. Sukaviriya, "Steerable Interfaces for Pervasive Computing Spaces," in *Proceedings of the First IEEE Conference on Pervasive Computing and Communications (PerCom 2003)*, 2003, pp. 315–322.

Mark Ashdown has recently finished a PhD in Computer Science at the University of Cambridge Computer Laboratory, where he previously obtained the BA in Computer Science. His PhD work has focussed on large projected interfaces, calibration for such displays including the use of computer vision, input devices, graphics, and human-computer interaction. He is currently working at the University of Tokyo Institute of Industrial Science.

Peter Robinson is Reader in Computer Technology and Deputy Head of Department at the University of Cambridge Computer Laboratory in England. His research interests are in the general area of applied computer science. The main focus for this is human-computer interaction, where he has been leading work for some years on the use of video and paper as part of the user interface. He also works on electronic design automation and, in particular, on support for self-timed circuits.