

# Experiences Implementing and Using Personal Projected Displays

Mark Ashdown  
mark@ashdown.name

Peter Robinson  
pr@cl.cam.ac.uk

University of Cambridge Computer Lab, J. J. Thomson Avenue, Cambridge CB3 0FD, U.K.  
<http://www.cl.cam.ac.uk/Research/Rainbow/>

## Abstract

We have created a personal projected display called the *Escritoire*. Two projectors overlap to create a horizontal desk display, and two pens allow two-handed input. We describe the issues we encountered implementing this system and the results of experiments with single users working on their own and with pairs of users collaborating using separate desks. Cheap portable projectors form a display that fills the desk and also has a high-resolution region in front of the user. The display works under normal office lighting and front projection has not been a problem. Commodity 3D video cards can easily perform the necessary image warping, using planar homographies, to compensate for rough projector positioning. We describe various implementation issues with DirectX and OpenGL. Individual users were quickly able to use the system, and the colour and intensity mismatch between projectors was not a problem. For remote collaboration participants found an audio channel and shared desk surface to be much more useful than a video channel.

## 1. Introduction

The face of personal computing has remained largely unchanged since it was defined by the Xerox Star on its release in 1979. It consists of a desktop computer with a monitor for output, a mouse and a keyboard for input, and a graphical interface that uses windows, icons, menus, and a pointer. There are various problems with this model of the human-computer interface, and various enhancements are possible that are not supported by the conventional desktop metaphor. One problem is lack of space. The falling prices and increasing quality of display devices, particularly digital projectors, make large personal displays possible. These large displays require new styles of interaction because simply putting the desktop metaphor on much larger screen does not work [13]. Another problem is the




Figure 1. The *Escritoire* is a personal projected display. (a) Two overlapping projectors create a desk display on which documents can be arranged. (b) Items can be moved from the low-resolution to the high-resolution region to view them in detail.

prevalent use of paper which the personal computer was supposed to eliminate. People still use paper because it has affordances that have not been superseded. An addition that would improve the interface is continuous input with two hands, which has been proved to have manual and cognitive benefits [9]. Another addition is the ability for users to collaborate in real time. Seeing each others faces, as in a standard video conference, is not particularly useful in most tasks—users should be allowed to share the documents and data on which they are working, so those materials can ground their conversation and form a basis for reviewing information and suggesting changes.

Predictions of the paperless office have not been forthcoming, and people now use more paper than they ever did. The DigitalDesk [15], and other systems after it, attempted to address this issue and some of the problems above by taking the view that peoples' desks will always be covered in paper and by augmenting sheets of real paper with projected

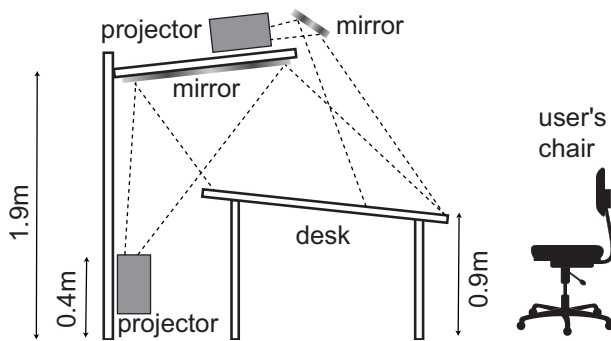


Figure 2. A shelf above the desk holds a projector that creates a small high-resolution fovea in front of the user. The projector behind the desk creates a large low-resolution periphery.

graphics to give them some computational properties. With the work described here we have taken the opposite route—simulating the properties of paper in a wholly computerized interface. Documents are increasingly becoming available in electronic form, but before paper is replaced as a medium for reading and annotating documents the infrastructure for storing and delivering electronic documents must be complemented by a suitable user interface. The hardware of tablet computers will be useful for mobile users, but for fixed settings such as homes or offices the full versatility of a desk-sized display can be exploited. An interface based on virtual paper rather than augmented paper has many advantages, such as the computer’s complete knowledge and control of the contents of the desk, and support for distributed collaboration in which the users can interact with a document in the same way irrespective of where it is physically stored.

The Escritoire [2] is a personal projected display that allows documents to be perused and manipulated like physical documents on a real desk (Figure 1). This paper describes lessons that have been learned during the implementation and use of this system that will be applicable to other projected displays and to projector-camera systems, although we do not use computer vision to calibrate the display of the Escritoire. The system has a large horizontal display where one projector fills the entire area while another provides high resolution in the region just in front of the user. The projected graphics are warped by commodity 3D video hardware to compensate for the distorting effects of oblique projection and imprecise projector positioning. The projection surface is a large digitizing tablet with a cordless pen, and this is combined with an ultrasonic pen to allow two-handed input. The sheets of virtual paper on the desk mimic physical paper, and can be arranged and annotated with the two pens. Interaction is possible between two users with desk displays in different locations because the

system is split into client and server programs and multiple clients can connect to a single server. Two instances of the Escritoire hardware have been assembled, at the Computer Laboratory in Cambridge, and at Thales Research and Technology who funded this work. The system uses TCP sockets to allow users to collaborate over any IP network. The details of the protocol are outside the scope of this paper, as are the precise features of the user interface. This paper covers the physical arrangement of the hardware, the use of video hardware to warp and align the projected graphics, an experiment on how individual users found the system, and finally an experiment on using the system for distributed collaboration.

## 2. Projected Display

Projectors have been tiled to make multi-projector display walls [8] but these are expensive installations intended as visualization facilities to be shared among groups of people. Typically a high speed network links a cluster of compute nodes that perform a simulation, and the results are sent to a cluster of rendering nodes that send their video output to the projectors. In contrast, the Escritoire is a projected display system for a single user that is intended to replace a normal workstation—it is a personal projected display [1]. A single PC with a dual-head graphics card drives two portable projectors that are mounted above and below the desk which is a large digitizer. The complete hardware arrangement costs less than 10,000 euros.

Because the Escritoire has a horizontal display it lends itself to a different type of work than a wall display. On a wall the user can make small changes then step back to view the whole surface or let onlookers view the changes, as would occur during a presentation. Sitting at the desk display the user can perform detailed work as he would on a normal desk with physical sheets of paper. The cost and complexity of the desk display must be balanced with the need for high resolution to perform detailed work. This is achieved by using one projector to fill the A0-sized desk with a low-resolution *periphery*, then using a second projector to create an A3-sized *fovea*—a region in front of the user where high resolution is needed most. Baudisch *et al.* [3] have also explored the idea of combining a high resolution focal region with a lower resolution periphery, by combining an LCD screen with a digital projector. However, they use a vertical surface that displays a conventional interface with a keyboard and mouse for input, and they do no graphics warping which means that calibration is purely mechanical and manual, orthogonal projection must be used, and the user must be positioned to avoid obscuring the optical path for the peripheral projection. Figure 2 shows the arrangement of the Escritoire’s projectors. One is placed on a shelf above the desk to get it as close to the desk surface as pos-

sible. The other is hidden behind the desk, and its image is reflected in a mirror on the bottom of the shelf.

For the fovea, image quality is paramount so a front-silvered mirror is used to reflect its image down onto the desk. The quality of the periphery is not so important so a standard, cheaper, back-silvered mirror is used. The fovea is brighter because its light is spread over a smaller area. It has higher resolution in pixels per millimetre for the same reason. Since the two regions have different resolutions the sheets of virtual paper must be prepared at different resolutions. To avoid artifacts in the fovea these sheets are stored at the fovea resolution so scaling is only necessary when preparing the image to display in the periphery. The differences between the two display regions are summarized below.

	Fovea	Periphery
display size	small	large
light intensity	high	low
display resolution	high	low
mirror	front-silvered	back-silvered
graphics scaling	no	yes

Two-handed input is achieved by combining two devices that do not interfere with each other: a large electromagnetic digitizer, and an ultrasonically tracked pen that was designed to be attached to a whiteboard. By analogy with the differences between fovea and periphery for the visual output, the pen input devices have been chosen to give different qualities to the user's two hands. The digitizer pen in the dominant hand is for precise tasks like writing. It has a button in the nib that senses when it is pushed down onto the surface, and two on the shaft that are pressed with the fingers. The ultrasonic pen in the non-dominant hand has lower accuracy and only one button, and is used to move items around on the desk. The differences between the two pens are summarized below.

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

The second instance of the *Escritoire* hardware, which was constructed at Thales Research and Technology, uses two Sanyo PLC-XW20A projectors which supply 1100 lumens, have 150 watt lamps, and draw 250 watts in total. The official minimum throw distance for these projectors is 1.43 metres which is too far, even at minimum zoom, to get an image small enough for an A3-sized fovea. However, we have found that a distance of around 0.8 metres is possible—projectors can generally focus on surfaces significantly closer than the distances for which they have been

designed. For the second instance of the *Escritoire* hardware, rather than reflecting the top projector's image down onto the desk using a front-silvered mirror we mounted the projector at a sharp angle so it points directly onto the desk. Although this violates the instructions that state that it should not be placed at an angle of more than 10 degrees, we have not noticed any problem in practice.

Projectors generate lots of heat. The room where the Sanyo projectors were used, which is 2.5×5 metres, was warmed considerably by them, requiring the windows to be opened after about an hour. For long-term use, ducts could be used to channel heat away from rooms whose air conditioning cannot cope [4]. The projected desk display can be used under normal office lighting, although, as with conventional monitors, direct sunlight should be avoided. The use of front projection rather than rear projection has not been a problem. Projected imagery does appear on a user's hands when they are moved over the desk surface but this has never been significant. The use of oblique projection from behind the desk allows the user to lean over to get a better look without occluding the projected image (Figure 3), thus avoiding a common drawback with front-projection.

### 3. Graphics

Precise mechanical positioning and periodic adjustment of the projectors and display surfaces would be prohibitive for a general purpose, personal display. Instead the graphics are warped in real time using commodity 3D video hardware to compensate for the distorting effects of oblique projection. The user selects some projected targets in the fovea and periphery using the digitizer pen which has an accuracy of ±0.25mm, and these are used to calculate a planar homography from each projector to the co-ordinate space of the digitizer. Sukthankar *et al.* [12] have used a computer vision technique to obtain the necessary homography au-

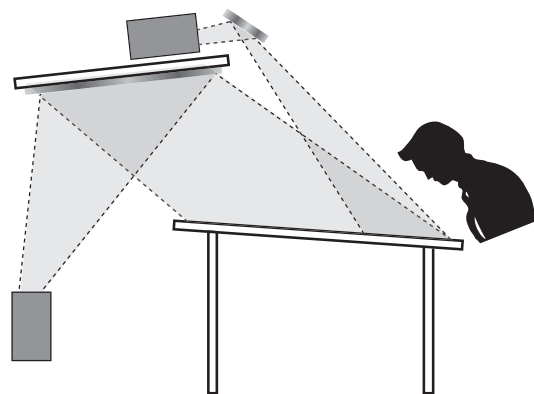


Figure 3. Oblique projection allows the user to lean forward without obscuring the projected image.




Figure 4. The 2D co-ordinate spaces of the projectors, desk, and pen input devices must be considered. Calibration involves computing planar homographies between these spaces that will allow a point in one to be transformed to a point in any of the others.

tomatically, but in our case manual intervention cannot be avoided altogether because it is necessary to calibrate the ultrasonic pen. At least four points are necessary for each projector—we use nine and calculate a least-squares solution. Calibration takes about one minute and only needs to be repeated when the hardware is moved. Various co-ordinate spaces must be considered: the two images to be displayed which are prepared in textures on the video card, the framebuffers corresponding to the two projectors, the physical surface of the desk, and the co-ordinate spaces of the two pen input devices. The two-dimensional projective spaces are linked together in a tree (Figure 4) where the transitions are accomplished via planar homographies that are measured during calibration. This is related to the method of using a homography tree to calibrate a multi-projector display wall [5].

The warp applied to the projected image is a projective mapping. This corresponds to a pin-hole model for the projector, and experimentally we have found it is a good match for the distortion experienced by the projected image, and is better than the bivariate quadratic mapping that was used for the DigitalDesk [15].

To avoid projecting on the fovea region twice it is blanked in the periphery framebuffer by transforming the fovea texture to the periphery framebuffer using the tree in Figure 4, then filling that area with a black quadrilateral.

The functionality of modern video cards is usually accessed through DirectX [6] or OpenGL [16]. Using 3D hardware to warp 2D textures requires some care. The large memories on modern video cards are ample for storing images in textures, but there are still limits on texture sizes because manufacturers expect the programmer to use many small textures rather than one or two big ones. Current limits are typically 512, 1024, or 2048 pixels square, and dimensions must be powers of two. We have used 1024×1024 textures for 1024×768 projectors. The warped quadrilat-

eral can be drawn as two triangles. To avoid triangles being omitted during back-face culling the triangle vertices should be specified in the correct order, or back-face culling should be disabled.




Figure 5. (a) The rectangular image should be warped using a projective mapping. (b) Incorrect homogeneous values for the vertices will result in an obviously incorrect warp. (c) Out-of-range texture co-ordinates will result in the left side of the texture wrapping onto the right side.

In the days before modern video hardware it would have been necessary to divide the quadrilateral into an array of smaller pieces and apply an affine warp to each one. This is what researchers on the Tele-Graffiti project did recently [14], but it would have been better to warp the image using perspective texture mapping and homogeneous co-ordinates as described here. When the  $x$  and  $y$  screen co-ordinates of a triangle vertex are calculated, the homogeneous  $w$  co-ordinate should also be calculated and used because although it does not affect the position of the vertex on the screen it does affect the texture mapping (Figure 5(b)). Negative  $w$  values may cause a problem because they would not normally occur in the rendering of a 3D scene. If they actually occur, the homography matrix can be multiplied by -1 to avoid them without changing the mapping that the homography represents. Each vertex on the screen also has  $u$  and  $v$  texture co-ordinates that specify which part of the texture is used. They should be in the range 0.0 to 0.999—because textures are often tiled the co-ordinates are taken modulo 1, so a value of 1.0 refers to the same position on the texture as 0.0. Figure 5(c) shows the left edge of a texture wrapping onto the right edge. The scale factor between texture and screen should be in the range 0.5 to 2.0 because otherwise bilinear interpolation does not produce a good result—either pixels in the texture will be omitted from the calculations (Figure 6(a)), or the resolution of the projector will be wasted because the texture has insufficient resolution to utilize it (Figure 6(b)).

Direct3D, the part of DirectX that exploits 3D video hardware, offers immediate-mode rendering for simple sets of primitives and retained-mode rendering for managing complex scenes. Since the geometric model of our scene is extremely simple, immediate mode is best. We have used flat shading to avoid unwanted lighting effects, enabled texture perspective to ensure that the  $w$  co-ordinates are used




Figure 6. The scale factor from texture to screen should be between 0.5 and 2.0. (a) If the texture has too many pixels, some of them will be omitted. (b) If the texture has too few pixels, the resolution of the projector will not be fully exploited.

as described above, and have drawn the warped quadrilaterals by making triangle fans from the sets of four vertices. A clipper object for each display handles the case when a triangle protrudes over the edge of the screen.

OpenGL is like the immediate-mode part of Direct3D. It projects simple primitives from 3D to a 2D screen, then renders them. If one imagines the texture to be in a plane where  $z = 0$ , the  $3 \times 3$  planar homography matrix to map from texture to screen can be placed in a  $4 \times 4$  projection matrix which can be pushed onto OpenGL's matrix stack. The rectangular texture can be drawn as a quadrilateral using the `GL_QUADS` primitive.

Drawing even large polygons with large textures is fast with current hardware. To draw a  $1024 \times 1024$  texture so that it fills most of a  $1024 \times 768$  screen and therefore has about 750,000 pixels takes 0.15 to 0.45ms on the Matrox Millennium G400 cards that we have used. The hardware is designed to draw polygons quickly and texture map them with static textures—updating the textures is not so fast. It takes 4-10ms to update an entire  $1024 \times 1024$  texture on the AGP cards we have used, and about 3 times that for the PCI cards. When moving a  $640 \times 480$  image across the desk our original system based on a 900 Mhz PC with 256 MB of RAM can update the two-projector display at 30 frames per second. That system uses separate AGP and PCI cards for the two projectors due to previous limitations in Direct3D, but this is no longer necessary. To minimize texture update times, it has been necessary to use update regions to ensure that only the parts of a texture that will change are replaced.

#### 4. Single User

The desk surface of the *Escritoire* allows sheets of virtual paper generated from PDF documents and images to be arranged, put in piles, and annotated. The annotations are written to the original files so they can be reviewed off-line. Figure 7 shows a user annotating some images. We have implemented three options for pen cursors that are shown

in Figure 8. We have used the two *Escritoire* systems that we have assembled at two different sites for demonstrations, and issues about the user interface have subsequently arisen. For instance, people cannot drag with the digitizer pen without engaging the nib button, so we modified the code to cause only one button press to be reported at a time—the button that was pressed first is used and later presses are discarded. We also explicitly tested the interface in two experiments: one with single users described in this section, and the one with two users interacting over a network described in the next section.




Figure 7. Some images have been placed on the desk, and the user is annotating one of them.




Figure 8. We have implemented three cursor options for the pens: (a) no cursor, (b) cross hairs, or (c) traces.

Seven people who had not used the *Escritoire* before performed some tasks with the system, then they were asked some questions about it. Each participant was first shown how to use the two pens to move and annotate items on the desk. The participant then performed two tasks. The first task was to highlight the spelling mistakes in a set of four textual documents, where each document was a single page and contained one mistake (Figure 9). This task was performed three times, once for each type of cursor. The order in which the cursor options were used was varied among the participants.

The second task was to put a set of 15 images of animals,

buildings, and flowers from a stock photography library into three piles depending on their content (Figure 10). We say that the pile in the figure has a south-east direction because of the direction along which the images are arranged from front to back. There are three other directions, so the task was performed four times, once for each direction. The piling feature is similar to the *pile metaphor* developed by Apple [11]




Figure 9. The first task for the individual users was to (a) inspect some pages of text one at a time and (b) highlight the spelling mistake in each one.

After finishing the tasks, participants were asked which type of cursor they preferred, and were asked for comments about the cursor types. They were asked which pile direction they preferred and for comments about the piles. Finally they were asked for comments about the two pens, and about the system generally. The whole experiment took 30 to 40 minutes for each participant.

All participants took only a few minutes of instruction to be able to use the desk display with two pens. One person was split between preference for no cursor and preference for the cross cursor, which made the totals: none 5.5, cross 1.5, trace 0. Responses did not indicate that one pile direction was much better than the others, with 4 participants expressing no preference. A pile splits open automatically when the pen moves over it to allow the user to browse through its contents. This caused some confusion because items moved around spontaneously. The digitizer pen is effective up to 25 mm from the surface, so the a pile could be affected even if the user was just passing his hand over the desk without intending to cause any reaction. We modified the software to use *proximity events* from the digitizer to cause a pile to return to its default state when the pen is moved away from the desk.

The pens work best when they are held perpendicular to the surface, although this is unnatural for most people. As the pen is tilted the reported location changes, although this is not a problem if the person who is using the desk is the one who calibrated it. It may be a problem if over time the user varies the angle at which the pen is held. To get events from the ultrasonic pen we used a program that comes with the hardware to make the pen emulate the mouse, and processed the mouse events that were subsequently generated. A problem that arose with this approach is that the program




Figure 10. The second task for the individual users was to put images into piles based on their content. We have experimented with four directions for arranging the items in a pile, including the direction shown here.

interpreted the pen dwelling in a small area for more than about 0.5 seconds as a mouse click rather than the start of a drag. Some users instinctively pressed the pen on an item and then held it there while they decided where to move the item, but they had to be instructed to move the item immediately after touching it. Initially the projected display fell outside the active area of the digitizer and users were confused when the pen stopped working as they moved an item outside that area. Upon noticing this problem we recalibrated the system to confine the display to the active area.

Users tended to lean on the desk so that it moved, causing the projected graphics to become misaligned. This experiment was performed on a CalComp DrawingBoard IV digitizer which has a white surface that is good for front projection but which has rather feeble fixings. The other digitizer we have used is a Summagrid V which is much more sturdy and would have avoided this problem.

The differences in colour and intensity between fovea and periphery are not a problem with the *Escritoire's* desk display. None of the users complained about them. This is unlike multi-projector display walls where luminance and chrominance matching are important to hide the seams between projectors [10].

## 5. Multiple Users

Two *Escritoire* desks can be linked over the Internet to allow real-time collaboration on documents. We have used this to augment the audio and video channels of a conventional video conference. Rather than simply seeing each other's faces, the participants share textual and graphical information which gives them a basis for conversation and allows them to point and annotate while discussing the material.

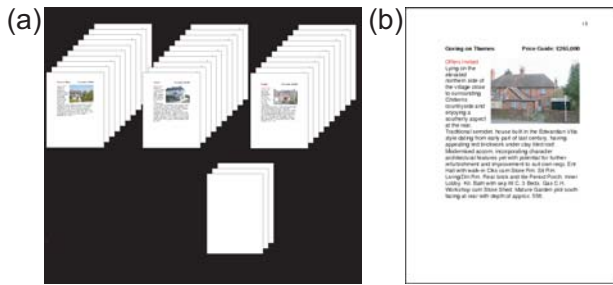


Figure 11. Pairs of users were given the task of collaborating remotely using two linked Escritoire desks to pick the best house from groups of houses. (a) Initially each user was shown the 30 houses and was given some blank sheets on which to make notes. (b) Each page contained a textual description and a photograph of one house.

We performed an experiment with three pairs of participants. In each case one person was using an Escritoire system in Cambridge, while another was using one at Thales in Reading about 100 miles away. The participants were told that the aim of the exercise was to find the best value house out of a group of houses. Initially they worked separately. A set of 30 sheets, each containing details and a picture of one house, was shown to each participant, who was also given some blank sheets of virtual paper on which to make notes (Figure 11). The two desks were then connected and 10 of the original 30 houses were placed on the shared surface. The pair of participants was set the task of working together to find the best three houses in order. This task was repeated three times, and each time a different one of the cursor types from Figure 8 was used. One of the desks was in the same room as the server and was connected to it via a 100 Mbps LAN, while the other was connected to it via the Internet which was available to the server computer through an ADSL modem offering 256 kbps. Each set of 10 houses took 20-30 minutes to discuss so each group took around 1.5 hours for the collaborative part of the experiment. Afterwards participants were asked which cursor type they preferred, and were asked for comments on cursors. They indicated their agreement on a five-point scale to the statements ‘the audio channel was useful for the task’, ‘the video channel was useful for the task’, ‘the desk interaction was useful for the task’, and ‘the amount of desk space available aided the task’. They were asked whether the difference in resolution and the difference in brightness between fovea and periphery were a problem, and whether the latency of the desk interaction was a problem. Finally they were asked for general comments.

All participants were quickly able to use the desk for remote collaboration without any training. One person was split between preference for the cross cursor and the trace which made the totals: none 0, cross 0.5, trace 5.5. The

direct nature of pointing at items with a pen means that in the single user case it is obvious where the pen is pointing so the cursor was just a hindrance, but in this collaborative case one user can use the cross cursor to indicate locations to the other user, and can use the trace to make gestures to the other user. The trace also gives continuous feedback that makes it easier to follow what the other user is doing. Gutwin [7] has made various recommendations for implementing traces like these and has also shown that they improve communication between remote collaborators. Three participants said the difference in resolution between fovea and periphery was not a problem, and three said it was because text was unreadable outside the fovea and items had to be moved to the fovea to use them. It would be nice to have high resolution over the whole desk surface but we have made a trade-off between the cost, size, and complexity of the system and the number of pixels that it can display. All participants said the difference in brightness was not a problem. One of them said he did not notice it, and another said it was useful because it was an obvious cue to the edge of the high resolution region. The system took a few minutes to transfer all of the pixel data for the sheets of virtual paper across the ADSL connection, but after this initial burst the connection had ample capacity to carry the Escritoire’s message data together with the video and audio streams. Using the system through a server 100 miles away feels the same as using a server in the same room so latency was not a problem, but there were a few occasions when the system hung for a few seconds, presumably because of surges in network traffic.

## 6. Conclusion

The creation and testing of the Escritoire have shown that a projected desk display can provide a large workspace for manipulating documents, and for collaborating on them with remote parties.

Two projectors can be connected to a commodity PC and arranged in a normal office to provide the image sizes required for the foveal display. Projective warping can compensate for rough projector alignment. Use of front-projection was not a problem. The difference in colour between the two projectors was not a problem, and the difference intensity, if anything, was helpful to delineate the display regions. Higher resolution projectors would improve the interface but current models are designed for displaying video or making presentations—the high end of the mass market provides XGA (1024×768) resolution. UXGA (1600×1200) projectors would be better but they are very expensive. Remote controlled focus and zoom would make physical configuration easier but would increase the price of the system.

Commodity 3D video cards can easily perform the warp-

ing required for oblique projection of even large images. Care must be taken to avoid unnecessary time being spent updating the texture. If the computer rendering the images is geographically separated from the one that stores them, transmission should be optimized and large updates should be avoided when human interaction is occurring.

It only took a few minutes for people to learn to use the Escritoire's desk display with two-handed pen input. This is impressive when one considers how long it takes to learn to use a keyboard and mouse. The ultrasonic pen is easy to grab and its operation is simple: it is just pressed down onto the surface, moved, then lifted. Users found it easy to use with the non-dominant hand. People could easily make use of the large display area because they could survey the whole thing at a glance to see which documents were available, and could quickly reach out and drag one to the fovea. Pen traces helped remotely collaborating users point and gesture to each other.

Useful person-centred and task-centred collaboration using video, audio, and desk interaction could be conducted over a standard ADSL connection. Two improvements that could improve the experience are minimization of the pixel data sent from server to client, possibly by using appropriate compression, and some kind of quality of service guarantee across the network.

There is currently much interest in miniature interfaces for mobile devices and PDAs fuelled by advances in technologies such as small screens, batteries, and ubiquitous networks. We believe that in fixed locations such as offices and homes, where space and mobility are not limiting factors, large-format interfaces will become popular. The prevalence of multi-monitor systems indicates users' appetite for screen space. Projectors can be combined to create large displays, and the Escritoire is an example of one with a lower price and size than existing multi-projector display walls for visualization or presentations. The desktop metaphor does not translate well to large displays, but the ease of direct manipulation with pens makes it a good way to control a desk-sized display. The Escritoire uses commodity components and exploits the existing manual skills that people have to form a personal projected display for performing the everyday tasks for which people use their desks.

## Acknowledgments

We gratefully acknowledge the involvement of Thales Research & Technology, part of the Thales Group. They funded this work, and the single and multiple user experiments were conducted at their site using Thales employees.

## References

- [1] Mark Ashdown. *Personal Projected Displays*. PhD thesis, University of Cambridge Computer Laboratory. Will be available early in 2004.
- [2] Mark Ashdown and Peter Robinson. The Escritoire: A Personal Projected Display. In *11th International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision (WSCG 03)*, pages 33–40, 2003.
- [3] Patrick Baudisch, Nathaniel Good, and Paul Stewart. Focus Plus Context Screens: Combining Display Technology with Visualization Techniques. In *Proceedings of UIST 01*, pages 31–40, 2001.
- [4] Gary Bishop and Greg Welch. Working in the Office of "Real Soon Now". *IEEE Computer Graphics & Applications*, 20(4):76–78, 2000. Available at <<http://www.cs.unc.edu/~welch/oorsn.html>>.
- [5] Han Chen, Rahul Sukthankar, Grant Wallace, and Kai Li. Scalable Alignment of Large-Format Multi-Projector Displays Using Camera Homography Trees. In *Proceedings of IEEE Visualization 2002*, pages 339–346, 2002.
- [6] Microsoft DirectX. <<http://www.microsoft.com/windows/directx>> (accessed 27 June 2003).
- [7] Carl Gutwin. Traces: Visualizing the Immediate Past to Support Group Interaction. In *Proceedings of Graphics Interface 2002*, pages 43–50, 2002.
- [8] Greg Humphreys and Pat Hanrahan. A Distributed Graphics System for Large Tiled Displays. In *Proceedings of IEEE Visualization '99*, pages 215–224, 1999.
- [9] Andrea Leganchuk, Shumin Zhai, and William Buxton. Manual and Cognitive Benefits of Two-Handed Input: An Experimental Study. *Transactions on Human Computer Interaction*, 5(4):326–359, 1998.
- [10] Aditi Majumder and Rick Stevens. Color Non-Uniformity in Projection Based Displays: Analysis and Solutions, 2003. to appear in *IEEE Transactions on Visualization and Computer Graphics*.
- [11] Richard Mander, Gitta Salomon, and Yin Yin Wong. A 'Pile' Metaphor for Supporting Casual Organization of Information. In *Proceedings of CHI 92*, pages 627–634, 1992.
- [12] Rahul Sukthankar, Robert G. Stockton, and Matthew D. Mullin. Smarter Presentations: Exploiting Homography in Camera-Projector Systems. In *IEEE Proceedings of ICCV 2001*, pages 247–253, 2001.
- [13] Kishore Swaminathan and Steve Sato. Interaction Design for Large Displays. *Interactions*, 4(1):15–24, 1997.
- [14] Naoya Takao, Jianbo Shi, Simon Baker, Iain Matthews, and Bart Nabbe. Tele-Graffiti: A Pen and Paper-Based Remote Sketching System. In *8th International Conference on Computer Vision (ICCV) 2001*, page 750, 2001.
- [15] Pierre D. Wellner. Interacting with Paper on the DigitalDesk. *Communications of the ACM*, 36(7):87–97, 1993.
- [16] Mason Woo, Jackie Neider, and Tom Davis. *OpenGL Programming Guide: The Official Guide to Learning OpenGL, Version 1.2*. Addison Wesley, 1999. ISBN 0201604582.