

Number 538



**UNIVERSITY OF  
CAMBRIDGE**

Computer Laboratory

## The Escritoire: A personal projected display for interacting with documents

Mark Ashdown, Peter Robinson

June 2002

15 JJ Thomson Avenue  
Cambridge CB3 0FD  
United Kingdom  
phone +44 1223 763500  
<http://www.cl.cam.ac.uk/>

© 2002 Mark Ashdown, Peter Robinson

Technical reports published by the University of Cambridge  
Computer Laboratory are freely available via the Internet:

*<http://www.cl.cam.ac.uk/TechReports/>*

Series editor: Markus Kuhn

ISSN 1476-2986

# The Escritoire: A personal projected display for interacting with documents

Mark Ashdown and Peter Robinson  
University of Cambridge Computer Laboratory  
J J Thomson Avenue, Cambridge CB3 0FD, U.K.  
{msda2,pr}@cl.cam.ac.uk

## ABSTRACT

The *Escritoire* is a horizontal desk interface that uses two projectors to create a *foveal display*. Items such as images, documents, and the interactive displays of other conventional computers, can be manipulated on the desk using pens in both hands. The *periphery* covers the desk, providing ample space for laying out the objects relevant to a task, allowing them to be identified at a glance and exploiting human spatial memory for rapid retrieval. The *fovea* is a high resolution focal area that can be used to view any item in detail. The projected images are continuously warped with commodity graphics hardware before display, to reverse the effects of misaligned projectors and ensure registration between fovea and periphery. The software is divided into a hardware-specific client driving the display, and a platform-independent server imposing control.

**Keywords:** H.5.2, virtual paper, augmented reality, foveal display, focus plus context, bimanual input.

## MOTIVATION

The advent of the personal computer with a graphical display brought predictions of the paperless office, but this dream has not been forthcoming. Although the point at which it is converted from electronic to hardcopy has, in many cases, become later in a document's lifetime, paper still has affordances that have not been superseded by the abilities of computer applications. Highly structured tasks such as searching a library are now accomplished through a computer rather than with a cabinet of cards, but less structured work such as writing an academic paper that draws on ideas from various sources, generally still uses paper documents [28].

The hardware and software of e-books – batteries, memory, LCD displays, wireless networking, standard data formats, user interfaces – has progressed to the point where they seem a viable alternative to paper books, but their adoption now depends



**Figure 1:** The 9×12 inch desk; too small to work on, yet this lack of space is common in graphical user interfaces.

on other factors – distribution, encryption, availability of compelling content, copyright protection, billing and collection of royalties [10]. Similarly, standards such as Adobe's PDF allow documents to be exchanged in a standard format, but before they can supplant paper as a medium for documents the affordances of paper must be addressed. Once virtual paper can do what physical paper can its other capabilities will become very useful, such as its extension from single user interaction to distributed collaboration.

For decades the computer has almost exclusively presented a visual interface on a screen with a 14 to 21 inch diagonal, and received input via a keyboard and mouse. Weiser envisaged ubiquitous computing accessed through tabs, pads and boards – display and interaction devices of different sizes, posing varied challenges in implementation and interface design [32]. The archetypal personal computer has a pad-sized display. There has recently been much interest in the tab-sized, small screen technology of the many mobile devices being developed to exploit the economical and near ubiquitous access now available to digital networks [31]. This style of interface is useful for nomadic users, but when at home or in the office the size and versatility of a board-sized display can be fully utilized.


The rising quality and falling cost of LCD and plasma screens and video projectors using DLP technology [35], mean that large displays will form the interface of choice for personal workspaces, but scaling up the interface will produce a change, not only in quantity, but also in quality [30].

Figure 1 depicts the imaginary ‘9×12 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. An analogy can be drawn with the memory hierarchy of a computer where the inability to store a program’s working set in the appropriate memory level results in ‘thrashing’. In the case of the desk, the inability to simultaneously view the documents that are required to make useful progress results in the continual overhead of swapping a different document to the top of the pile. This hinderance to interleaved access qualitatively restricts the user’s working practices. No one would have a desk this small, yet current incarnations of the ‘desktop metaphor’ (more aptly called the ‘office analogy’ [19]) only allow one application window to be viewed adequately at any time.

This paper describes the *Escritoire*; a desk on which projected sheets of *virtual paper* can be read and manipulated. It may remove the point in many documents’ lifetimes at which transfer to hardcopy is necessary. A cordless pen in each hand is used for input, two overlapping projectors shining onto a desk form a *foveal display* that provides a large visual context while supplying a high detail focus, and a client-server design supports distributed collaboration. Images and documents can be arranged and annotated on the desk, and the new interface encompasses the desktop metaphor, allowing existing applications to be used along side the virtual paper, combining the standard pad-sized computer display with a board-sized desk surface. The two-handed input allows the user to interleave actions and helps to avoid decomposing a task to the unnaturally low level often imposed by traditional one-handed techniques.

## RELATED WORK

The interface of the *Escritoire* builds on three main areas of previous research; visualization techniques whose goal it is to make more information available on a traditional pad-sized display, large displays that actually have more space, and desk interfaces that provide the more natural horizontal surface that humans use for non-computerized tasks.



**Figure 2:** An unmodified program is projected onto the wall and controlled at a distance by pointing a ‘wand’ at it.

## Visualization Techniques

Fisheye views [27] have been used to provide displays of focus plus context by sacrificing the original geometric relationships between graphical objects. They allow large vector diagrams such as graphs to be visualized on a standard display, but are less suitable for bitmapped images such as photographs because of the severe distortions that are introduced. The Document Lens displays a large document in its entirety on the screen, while allowing the user to interactively move their focus to any part of the document which is then viewed in full detail [24]. Pad++ takes a different approach, creating an infinite 2-dimensional information space containing graphical objects [4]. At any point in time the user is viewing a particular portion of the information space, at a particular distance that determines the size and resolution of the objects. The project addresses the practical problems of achieving high frame rates and scalability, and rather than attempting to impose metaphors on the interface it applies ‘interface physics’ – uniform sets of simple rules that allow the user to navigate quickly and effortlessly in the information space. The University of Maryland has continued the project with Jazz [12], a Java API that allows applications to be written for, and reside in, this zoomable information space. Data Mountain [23] replaces the standard web favourites menu with a group of miniature graphical versions of the web pages that are positioned on a 3-dimensional plane. The system takes advantage of human spatial memory, and experiments have shown that it reduces reaction time and error counts.

## Large Displays

Early work on new interaction paradigms using

large displays was undertaken in the MIT Media Room. Dataland [7] allowed a user to arrange items such as maps, letters, calendars, and photographs on the wall-sized display of the room, while a monitor at the user's side displayed a small portion in high detail like a magnifying glass. A related project, the World of Windows [6] presented an ensemble of simultaneous video streams on the wall. Such a deluge of information would ordinarily be overwhelming, but the user's gaze direction was used to identify the stream on which they were focused, then the volume of the other streams was reduced, and following a period of prolonged attention the video image was enlarged. Preliminary work on the *Escritoire* involved devising a new wall-sized display controlled using a 6 degree-of-freedom tracker [2]. Standard applications are projected on the wall and controlled at a distance using a device like a wand (Figure 2).

Recent work has used projected displays to present peripheral information about standard application programs. Kimura [15] displays representations of the user's activities on the peripheral display, reminding them of background tasks that are not currently visible on their monitor and allowing them to quickly switch to them. Baudisch has combined an LCD monitor and video projector to make a 'focus plus context screen' [3] – a very large monitor that has high resolution in the focal region, just like a standard monitor, plus a large contextual area that supports the current task. However, image warping has not been addressed, calibration is purely manual, and the combination of the displays simply involves placing a black window over the large display in the position occupied by the small display. Sukthankar et al. [29] have developed a method for efficiently calibrating a projected display to a non-orthogonal surface for use when making presentations. A similar method is used by the *Escritoire* to prewarp projected images at over 30 frames per second. The Office of the Future project at the University of North Carolina [20] envisages an office adorned with a multitude of inexpensive cameras and projectors that are used to infer the geometry and reflective properties of all surfaces. The projectors then use the surfaces to display interactive 3-dimensional objects in the room.

### Desk Interfaces

Earlier work in the Computer Laboratory on the DigitalDesk [33, 34] combined real paper and projected digital information on a horizontal desk surface. Further work investigated animated paper documents [25, 26] in which a normal book, when opened on the desk, would be recognized using computer vision techniques and augmented with

extra properties, such as an applet that could rate the user's performance at sketching a graph in a mathematics textbook. Later systems [1, 14] pursue the same goals, recognizing sheets of paper from printed bar codes and allowing interaction with projected graphics. The approach with the *Escritoire* is not to include paper in the interface, with the problems of computer vision and one-way transfer of information from the physical to digital domains that this entails, but to cater for the affordances of physical paper that are not available in the 'desktop metaphor'.

### FOVEAL DISPLAY

The display of the *Escritoire* comprises two digital projectors that are combined to create what we call a *foveal display*. The arrangement is shown in Figure 3. The user sits in front of a desk, with a projector on the floor reflecting its image in a mirror on the bottom of a shelf. This creates the *periphery* – a large, low-resolution display that covers the desk. The image from a second projector on top of the shelf is reflected down onto the area directly in front of the user. This creates the *fovea* – a small, high-resolution display that accommodates the user's focus of attention. The desk is actually a 36×48 inch digitizing tablet with a cordless pen.

The design objective for the fovea is high image quality, while for the periphery it is ease of construction and speed of operation. The fovea uses a 1024×768 pixel projector, a front-silvered mirror to avoid multiple reflections, and as small a projection angle as possible to avoid focusing problems. The periphery uses a 640×480 projector and a




Figure 3: The double projector arrangement of the foveal display.

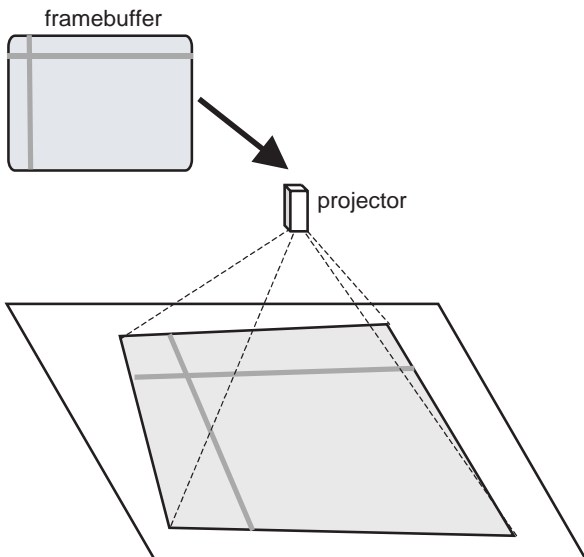
back-silvered mirror.

When the fovea and periphery images are projected they are distorted because the projectors are not aligned to the desk. Precise and regular mechanical positioning has been used in the past to alleviate this problem in large multi-projector displays, but this would be prohibitive for a general home or office computer. The *Escritoire* allows the projectors and mirrors to be positioned roughly, then compensates for the distortion using commodity graphics hardware.

### Co-ordinate Spaces

Graphical objects displayed on the *Escritoire* are positioned in *desk space* – a 2-dimensional plane with device-independent units of metres. Three other co-ordinate spaces must be considered: events from the pen input device occur in *device space*, and the image to be displayed is created in a rectangle in *texture space* then transformed to a quadrilateral in *framebuffer space* before being projected onto the desk.

The distortion undergone by each image as it is projected onto the desk is assumed to be the reverse of a pin-hole camera – a projective transform. The 2D image in the framebuffer is distorted as it is projected, producing a 2D image on the desk (Figure 4). The image to be displayed is therefore prepared in a rectangular texture on the graphics card, then transformed to produce a quadrilateral in the framebuffer which will appear as a rectangle again when projected on the desk. The transformation is calibrated using a set of points sampled from the digitizer onto which the image is projected.



**Figure 4:** Each image is distorted as it is projected. The distortion is reversed with a projective transform.

Figure 5 shows the four coordinate spaces that are considered and the necessary transformations between them. The transforms are all projective so they can be represented as  $3 \times 3$  matrices applied to homogeneous points in the co-ordinate spaces.  $P$  is predetermined by the relative positions of the projector and desk, and projector properties such as the zoom setting. The projective mapping  $F$  can be written as,

$$F = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & 1 \end{bmatrix}$$

Where the eight constants  $a-h$  are calculated by requiring the user to select four displayed points with the pen, thus providing four corresponding pairs of framebuffer  $(x, y)$  and device  $(u, v)$  points. The points are substituted into the equation below.

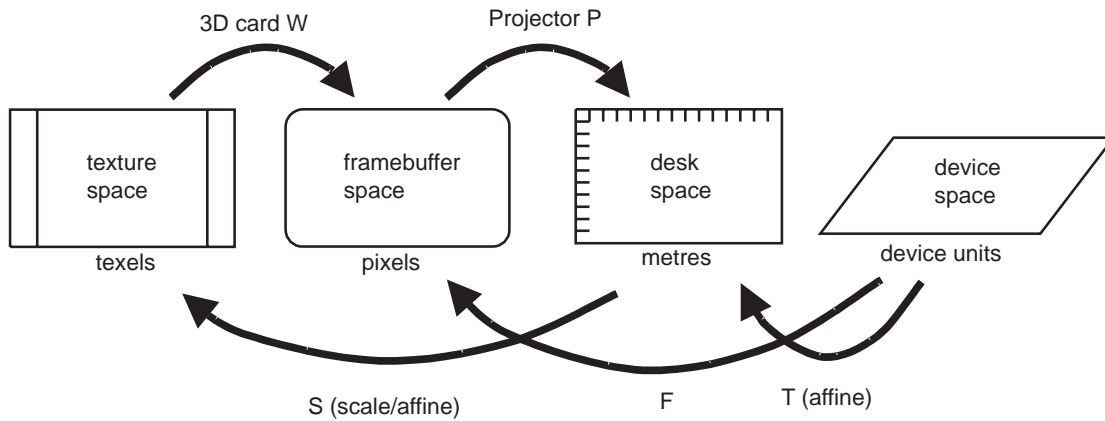
$$\begin{bmatrix} u_0 & v_0 & 1 & 0 & 0 & 0 & -u_0x_0 & -v_0x_0 \\ u_1 & v_1 & 1 & 0 & 0 & 0 & -u_1x_1 & -v_1x_1 \\ u_2 & v_2 & 1 & 0 & 0 & 0 & -u_2x_2 & -v_2x_2 \\ u_3 & v_3 & 1 & 0 & 0 & 0 & -u_3x_3 & -v_3x_3 \\ 0 & 0 & 0 & u_0 & v_0 & 1 & -u_0y_0 & -v_0y_0 \\ 0 & 0 & 0 & u_1 & v_1 & 1 & -u_1y_1 & -v_1y_1 \\ 0 & 0 & 0 & u_2 & v_2 & 1 & -u_2y_2 & -v_2y_2 \\ 0 & 0 & 0 & u_3 & v_3 & 1 & -u_3y_3 & -v_3y_3 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \\ d \\ e \\ f \\ g \\ h \end{bmatrix} = \begin{bmatrix} x_0 \\ x_1 \\ x_2 \\ x_3 \\ y_0 \\ y_1 \\ y_2 \\ y_3 \end{bmatrix}$$

Gaussian elimination is used to calculate  $a-h$ . Calibrating a projective transform like this is explained by Heckbert [11]. A similar method has been used by Sukthankar [29] using a camera to detect the projected calibration points. The *Escritoire* could use a camera in a similar way, but the user would still be required to select a number of projected points to calibrate the digitizer so there would be no saving in set up time.


Once transform  $F$  has been determined, the user chooses the position of the rectangular display on the desk using the keyboard to scale, translate and rotate it. Transform  $T$  is updated to reflect the changes.  $S$  is determined by the size on the desk of the projected display and the known resolution of the projector. The fovea and periphery each have their own  $S$  and  $F$  transforms corresponding to their separate textures and framebuffers, but  $T$  is associated with the pen input device and independent of the projectors. The desk-space locations and dimensions of graphical objects are device-independent.

### Image Warping


Once the four transforms have been calculated, transformations between any two co-ordinate spaces can be performed. Events from the pen are transformed to desk-space locations. When the user updates the position of the display on the desk



**Figure 5:** The four co-ordinate spaces and the perspective transforms between homogeneous points in the spaces. The units of measurement are shown under each space.



**Figure 6:** The two overlapping projected displays are warped due to oblique projection



**Figure 7:** The fovea and periphery regions appear square on the desk surface, having been warped using a hardware-implemented projective transform.

during calibration, the locations in the framebuffer to which the corners of the image in texture space must be transformed are calculated. The  $x$ ,  $y$  and homogeneous  $w$  values are used to create two triangles that form the quadrilateral in the framebuffer, which is texture mapped with the image to be displayed and passed to immediate-mode rendering routines. Commodity graphics cards can perform the required perspective transform at an ample speed – the system achieves 32 frames per

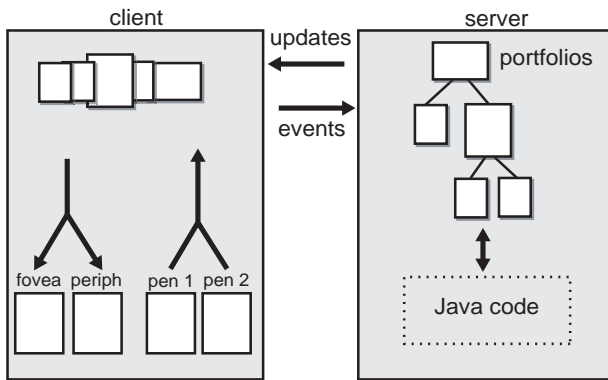
second when moving an A4 sheet (300,000 16-bit pixels) across the desk on a 900MHz PC with AGP Matrix Millennium G400™ and PCI NVIDIA Riva TNT2™ graphics cards. Projection by the periphery projector over the fovea display is avoided by blanking that region in the periphery framebuffer. The locations in the fovea framebuffer of the corners of its warped quadrilateral are transformed to desk-space, then from there to the periphery framebuffer. This allows a quadrilateral composed of two black triangles to cut out the foveal region.

### Pen Input

Digitizing tablets that support multiple pointing devices at different times are available, but we have found no large-format device that supports multiple pens simultaneously. To allow a second pen to be used on the desk, a Mimio™ [18] pen is combined with the digitizing tablet. The Mimio tracks its pen with ultrasound while the tablet uses electromagnetism, so the two do not interfere with each other. Similar to the distinction between fidelity and practicality of the fovea and periphery, the digitizing tablet provides a high resolution ( $\pm 0.25\text{mm}$ ) three-button pen for the dominant hand (the right hand of a right-handed person and conversely), while the ultrasonic pen for the non-dominant hand is a much lower resolution device with only one button that simply detects location of the pen when it is pressed to the desk surface. The Mimio receiver, normally mounted on the side of a whiteboard, is positioned at the top of the digitizer to avoid obstruction of the pen's ultrasonic emissions by either hand.

### CLIENT-SERVER ARCHITECTURE

The software of the Escritoire is split into two parts – the client and the server. The client is



**Figure 8:** Client-server architecture of the Escritoire. The server maintains a tree of portfolios, while the client displays their graphical representations and relays events back to their server implementations.

a hardware-dependent program written for Windows 2000 using Direct3D, designed to achieve the graphics performance necessary for high frame rates. The server is a platform-independent Java program that handles the control of, and responses from, the objects on the desk. The server stores all of the state, allowing the client to be restarted at any time.

The virtual paper is split into *tiles* and *portfolios*. The tiles reside on the client displaying their contents on the foveal display and accepting input events from the pens. They have a z-order which determines which ones appear on top of others. The portfolios on the server form the programmatic side of each sheet of virtual paper, accepting events and updating their appearance. They are held in a tree structure from which the tile order is derived by traversal, as shown in Figure 8. The table below summarizes the differences between the client and server software.

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

The protocol between client and server has been kept simple to keep the two as independent as possible. The client sends four types of message to the server; a **hello** message initiates a session and triggers an update to all tiles, a **ready** message indicates that the client has finished updating the display and is ready for more updates, **pen** events are specified in desk co-ordinates (metres) so a device-independent, and **keyboard** events signal key events. The server sends six types of message; **create tile** causes the client to make a new tile of the specified size, **move tile** causes a tile to be moved to a specified desk-space location, **update tile** is accompanied by pixel data to modify a

portion of an existing tile, **destroy tile** removes a tile from the list, **order tiles** specifies a new z-order for the tiles, and a **burst terminator** ends a series of events from the server.

### Client Pull versus Server Push

Should updates to the tiles be pushed out by the server when they are available, or requested by the client when it is ready? Message traffic between client and server is characterized by long periods of inactivity when the user is thinking, punctuated by periods of dense activity. Server push will be best for the quiet periods to avoid the client polling unnecessarily. Client pull, however, will allow traffic in dense periods to be synchronized to the client's display updates. Assuming graphics processing and other client functions are the bottleneck restricting frame rates, this will allow successive updates to the same area of a tile to be coalesced at the server, thus gracefully degrading update frequency without overwhelming the client.

The Escritoire uses a combination of client pull and server push to get the benefits of both techniques. The client sends a **ready** message to the server and consequently receives a burst of  $n$  messages followed by a **burst terminator**. It processes the messages, then if  $n > 0$  it polls the server again, but if  $n = 0$  it waits for more messages without polling. The server sends all pending messages in response to a **ready** message, but if there were no messages to send apart from the **burst terminator** it enters server push mode, whereby another burst will be started as soon as a message is available to send.

## VIRTUAL PAPER

Each sheet of virtual paper has an underlying portfolio on the server. The Java code is derived from a general Portfolio class and is like that of a Java applet – the server combines all incoming events into a single queue from which they are passed to the portfolio via a callback mechanism, and the portfolio updates its graphical appearance then makes a simple call to inform the server of the modified region so the appropriate tile can be updated. The four types of portfolio that can currently be created on the server – desk, jotter, PDF and VNC – are described below.

### Desk

The Desk object is the root of the portfolio tree. It contains other portfolios that are positioned by events from the pen in the non-dominant hand. Other events that are not handled by the desk object are passed on to its child portfolios.



## Jotter

The jotter portfolio allows an image in a standard format to be edited on the desk. The user positions the picture using the non-dominant hand, providing a frame of reference in which the dominant hand can work. This is a fast and easy way to move a picture to the fovea for viewing or editing and then replace it, as predicted in experiments by Buxton et al. [13].

## PDF


Another type of portfolio makes Adobe PDF documents available on the Escritoire. There is an example in Figure 11, being dragged by the left hand. Perusing a word-processed document or slide show on the desk simply requires it to be converted to a PDF file by the same method that would normally be used to print it to paper. Multiple documents can be viewed simultaneously, and the fovea, which is slightly larger than a sheet of A3 paper, allows documents to be placed side by side for comparison. The annotation strokes are added to the PDF file which can then be viewed offline using a viewer such as Adobe's Acrobat Reader™.

## VNC


Virtual Network Computing (VNC) [22] is a remote display system that translates the graphical interface of a computer into a standard protocol that is transmitted across an IP network. It allows, for example, the display of an X server to be viewed and controlled, using keyboard and mouse, from a Windows machine, or the desktop of a Mac to be controlled from a web browser. A portfolio has been written that supports the VNC protocol, thus a computer using a conventional desktop can be controlled from the Escritoire. Figure 11 shows an X Windows display lying on the desk along with the other items. Events from the pen in the dominant hand are translated into mouse events. Displays from window systems have proved to be good candidates for the warping process employed by the Escritoire because they are optimized for bitmapped representation.

## DISCUSSION

The Escritoire provides a large interactive surface on which to arrange images, documents and conventional application programs. The design has been motivated by the natural working practices of humans as evidenced by their physical workspaces, rather than the technological constraints that restrict current GUIs to small screens, with associated interface conventions and their drawbacks. Overlapping projectors and image warping at full



**Figure 9:** The projection is parallel to the axis through the middle, causing the shadow to fall on the outside of each hand.



**Figure 10:** Dragging a document into the fovea.

frame rate are used to create a foveal display that has a high resolution fovea in which the user performs detailed work, and a large periphery that provides much more space that is available on a conventional screen. This new interface not only caters for the direct focus of the eyes and hands, but fills the vision, allows data to be placed as far as the arms can reach, and reproduces the kinesthetic sense one gets from working with physical media such as paper.

As with the DigitalDesk, front projection has not been a problem. This may be partly because the projectors are positioned in front of the user so the shadows fall on the opposite side of the pen tips from the information being viewed (Figure 9), and partly because people are accustomed to light from above being shadowed and to information being occluded by the hands, so they automatically move their hands away when they obstruct the view. Work on the Office of the Future project has addressed intensity blending between projected images [21], but the visible boundary and difference



**Figure 11:** A set of slides is on the top left in the periphery, an image is on the left, the left hand is dragging a document, and the right hand is operating an X Windows program.

in image intensity, far from being a problem here, are a useful cue to the extent of the fovea and a further facet to the distinction between it and the periphery – the fovea, in addition to being superior in other ways, has a more intense image due to its light being spread over a smaller area.

Currently a single texture with a global warp is used for each projected image. The quadrilateral in the framebuffer could be split into parts creating local mappings requiring more calibration points between digitizer and projector, although this has not proved necessary in practice. Local mappings for the calibration of the ultrasonic pen are being investigated. With graphics cards now commonly containing 128Mb of memory, a texture bitmap on the card could be used for each sheet on the desk. At 72dpi with 16 bits per pixel over 100 A4 sheets could be stored in video memory. The warp that compensates for projector alignment to be combined with an individual warp for each sheet, enabling rotations and scalings to be used as inexpensive interactive effects.

The *Escritoire* presents a document-centred view of work being undertaken. Traditional graphical components such as menu bars are not present because these are unsuitable for very large dis-

plays [30]. The user has direct access to all of the documents on the desk, and can immediately bring any one to the fovea (Figure 10). The action of grabbing a document when it must be viewed in more detail comes instinctively. As with paper, modifications to the documents are not saved at the command of the user – their state and location when the client program is started is the same as it was the last time the *Escritoire* was used.

Computers were originally controlled via command line interfaces before graphical user interfaces became the mode of choice. GUIs still support CLIs by providing terminal windows, and similarly the desk interface of the *Escritoire* provides support for the GUIs of multiple machines, treating them simply as documents to be manipulated.

People use desk organization to remind themselves to do things, and to loosely categorize information by placing it in groups and piles that complement more formal, long-term filing systems [16]. The *Escritoire* supports this behaviour, but this time the messy desk contains ‘virtual mess’ that disappears at the touch of a button.

## FUTURE WORK

The ‘pile’ metaphor [17] developed at Apple Computer was intended to support informal categorization and browsing. The notion of piles will be added to the Escritoire, so any items on the desk can be grouped into piles, which can be browsed and broken up with simple gestures from the dominant hand. Unlike that previous work, the piles will not contain icons representing the documents, but the editable documents themselves.

The client-server architecture allows multiple clients to connect to a single server to share a workspace. This is currently possible using one client that has the hardware described and another that is a normal application program. The latter simulates the desk interface as well as is possible with a single mouse and a small screen. A second version of the hardware will be assembled and the pair will be used to augment a video conference. Such telecommunication can create *person space* and *task space* [8] – a standard video conference provides a person space but video has been found not to provide a significant advantage for many tasks. An Escritoire at each end of a video conference will provide a task space – an artificial shared space that is the domain in which the participants will work. This is expected to provide a much more useful facility for remote parties wishing to discuss and mark up documents and presentations – common tasks for employees of a multinational company. Once pen events are received by the server and queued for processing by the portfolios there is no distinction between the pens of different users. This should create interesting possibilities for bimanual techniques executed with one hand from one user, and one hand from another. Freeman [5,9] has considered such fine-grained interaction between direct manipulation tasks of multiple users.

The earlier work on projecting a GUI onto the wall and controlling it from a distance with a wand [2], will be combined with the desk interface to create three surfaces of increasing size and decreasing resolution. The user will point at items on the desk, then move them up onto the wall as if putting them on a book shelf. John von Neumann wrote in 1946 ‘We are . . . forced to recognize the possibility of constructing a hierarchy of memories, each of which has greater capacity than the preceding but which is less quickly accessible.’ This concept will be applied to the user interface to construct a hierarchy of interactive spaces, each of which has greater capacity than the preceding but which is less precisely represented and controlled.

## ACKNOWLEDGMENTS

This work was funded by Thales Research & Technology UK (<http://www.rtl.co.uk/>) whose support is gratefully acknowledged.

## REFERENCES

- [1] ARAI, T., MACHII, K., AND KUZUNUKI, S. Retrieving Electronic Documents with Real-World Objects on InteractiveDESK. In *Proceedings of UIST 95* (1995), pp. 37–38.
- [2] ASHDOWN, M., AND ROBINSON, P. The Writing’s on the Wall: Large, Remotely Controlled Displays. In *Proceedings of the First European Conference on Computer-Supported Collaborative Learning (Euro-CSCL) 2001* (2001), pp. 83–88.
- [3] BAUDISCH, P., GOOD, N., AND STEWART, P. Focus Plus Context Screens: Combining Display Technology with Visualization Techniques. In *Proceedings of UIST 01* (2001), pp. 31–40.
- [4] BEDERSON, B. B., AND D., H. J. Pad++: A Zooming Graphics Interface for Exploring Alternative Interface Physics. In *Proceedings of UIST 94* (1994), pp. 17–26.
- [5] BIER, E. A., AND FREEMAN, S. MMM: A User Interface Architecture for Shared Editors on a Single Screen. In *Proceedings of UIST 91* (1991), pp. 79–86.
- [6] BOLT, R. A. Gaze-Orchestrated Dynamic Windows. *Computer Graphics* 15, 3 (1981), 109–119. Proceedings of SIGGRAPH 81.
- [7] BOLT, R. A. *The Human Interface*. Lifetime Learning Publications, 1984.
- [8] BUXTON, W. Telepresence: Integrating Shared Task and Person Spaces. In *Proceedings of Graphics Interface ’92* (1992), pp. 123–129.
- [9] FREEMAN, S. M. G. *An Architecture for Distributed User Interfaces*. PhD thesis, University of Cambridge Computer Laboratory, 1993.
- [10] HARRISON, B. L. E-Books and the Future of Reading. *IEEE Computer Graphics & Applications* 2000, May/June (2000), 32–39.
- [11] HECKBERT, P. S. Fundamentals of Texture Mapping and Image Warping. Master’s thesis, University of California, Berkeley, 1989. See pages 17–21 for projective mappings between homogeneous co-ordinate spaces.
- [12] Jazz Project. <http://www.cs.umd.edu/hcil/jazz/>.

- [13] KABBASH, P., MACKENZIE, S. I., AND BUXTON, W. Human Performance Using Computer Input Devices in the Preferred and Non-preferred Hands. In *Proceedings of InterCHI 93* (1993), pp. 474–481.
- [14] KOBAYASHI, M., AND KOIKE, H. Enhanced-Desk: Integrating Paper Documents and Digital Documents. In *IEEE Proceedings of Asia Pacific Computer Human Interaction (APCHI '98)* (1998).
- [15] MACINTYRE, B., MYNATT, E. D., ET AL. Support for Multitasking and Background Awareness Using Interactive Peripheral Displays. In *Proceedings of UIST 01* (2001), pp. 41–50.
- [16] MALONE, T. W. How do People Organize Their Desks? Implications for the Design of Office Information Systems. *ACM Transactions on Office Information Systems 1*, 1 (1983), 99–112. Now called ACM Trans. on Inform. Sys.
- [17] MANDER, R., SALOMON, G., AND WONG, Y. A 'Pile' Metaphor for Supporting Casual Organization of Information. In *Proceedings of CHI 92* (1992), pp. 627–634.
- [18] Mimio™ from Virtual Ink. <http://www.mimio.com/>.
- [19] PEMBERTON, S. Views and Feelings: Metaphorically Speaking. *SIGCHI Bulletin* 27, 4 (1995), 96.
- [20] RASKAR, R., WELCH, G., CUTTS, M., LAKE, M., STESIN, L., AND FUCHS, H. The Office of the Future: A Unified Approach to Image-Based Modelling and Spatially Immersive Displays. In *Proceedings of SIGGRAPH 98* (1998), pp. 179–188.
- [21] RASKAR, R., WELCH, G., AND FUCHS, H. Seamless Projection Overlaps using Image Warping and Intensity Blending. In *Proceedings of Virtual System and Multimedia (VSMM) '98* (1998).
- [22] RICHARDSON, T., STAFFORD-FRASER, Q., WOOD, K. R., AND HOPPER, A. Virtual Network Computing. *IEEE Internet Computing* 2, 1 (1998), 33–38.
- [23] ROBERTSON, G., CZERWINSKI, M., AND LARSON, K. Data Mountain: Using Spatial Memory for Document Management. In *Proceedings of UIST 02* (1998).
- [24] ROBERTSON, G. G., AND MACKINLAY, J. D. The Document Lens. In *Proceedings of UIST 93* (1993), pp. 101–108.
- [25] ROBINSON, P. Digital Manuscripts and Electronic Publishing. *Editio 1999*, Autumn (1999), 337–346. ISBN 3-484-29513-9.
- [26] ROBINSON, P., SHEPPARD, D., WATTS, R., HARDING, R., AND LAY, S. A Framework for Interacting with Paper. In *Proceedings of Eurographics '97* (1997).
- [27] SARKAR, M., SNIBBE, S. S., TVERSKY, O. J., AND REISS, S. P. Stretching the Rubber Sheet: A Metaphor for Large Layouts on Small Screens. In *Proceedings of UIST 93* (1993), pp. 81–91.
- [28] SELLEN, A., AND HARPER, R. Paper as an Analytic Resource for the Design of New Technologies. In *Proceedings of CHI 97* (1997), pp. 319–326.
- [29] SUKTHANKAR, R., STOCKTON, R. G., AND MULLIN, M. D. Smarter Presentations: Exploiting Homography in Camera-Projector Systems. In *IEEE Proceedings of ICCV 2001* (2001).
- [30] SWAMINATHAN, K., AND SATO, S. Interaction Design for Large Displays. *Interactions* 4, 1 (1997), 15–24.
- [31] TREVOR, J., HILBERT, D. M., SCHILIT, B. N., AND KOH, T. K. From Desktop to Phonetop: A UI for Web Interaction on Very Small Devices. In *Proceedings of UIST 02* (2001), pp. 121–130.
- [32] WEISER, M. The Computer for the 21<sup>st</sup> Century. *Scientific American* 1991, Sept (1991), 94–104.
- [33] WELLNER, P. Interacting with Paper on the DigitalDesk. *Communications of the ACM* 36, 7 (1993), 87–97.
- [34] WELLNER, P. D. *Interacting with Paper on the DigitalDesk*. PhD thesis, University of Cambridge Computer Laboratory, 1994.
- [35] YODER, L. A. The State of the Art in Projection Display: An Introduction to the Digital Light Processing Technology. Available at [http://www.dlp.com/dlp/resources/white\\_papers.asp](http://www.dlp.com/dlp/resources/white_papers.asp).