

The Escritoire: A Personal Projected Display

Mark Ashdown & 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 surface that uses two overlapping digital projectors to create a *foveal display* – a large interactive surface that has a high-resolution area to accommodate the user's focus of attention. In contrast to high cost, multi-projector display walls for group presentations or scientific visualization, it is an augmentation of the personal computer, designed for common tasks such as document viewing and annotation. The system is calibrated using planar homographies, and then warps the projected images in real time to compensate for the unconstrained positions of the projectors. To allow documents to be shared by multiple clients, the software is divided into a hardware-specific client driving the display, and a platform-independent server controlling the desk contents.

Keywords

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

1. INTRODUCTION

The advent of the personal computer with a graphical display brought predictions of the paperless office, but this dream has not been forthcoming: paper still has affordances that have not been superseded by the abilities of computer systems. Highly structured tasks such as searching a library are now accomplished using 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 [Sel97].

Besides the natural affordances of a paper, the space available on a real desktop permits a different style of work to that possible on the 'desktop metaphor,' (which would be more aptly called the 'office analogy' [Pem96]). 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. This is analogous to the memory hierarchy of computer where the inability to store a program's working set in the appropriate memory level results in



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

'thrashing'. In the case of the desk, thrashing occurs as documents are repeatedly brought to the top to be read consecutively rather than concurrently. Visualization techniques lessen the problem by providing a 'window' onto a larger virtual workspace, but we believe that ideally computer interfaces would provide the physical space to which humans are accustomed in other forms of work.

The rising quality and falling cost of digital projectors mean that large displays will form the interface of choice for personal workspaces in the future, but the desktop metaphor does not perform well when simply transferred to a much larger screen [Swa97]. For instance, menu bars become difficult to use when they are a long way from the user's centre of attention, and conventional text displays become

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Journal of WSCG, Vol.11, No.1, ISSN 1213-6972
WSCG'2003, February 3-7, 2003, Plzen, Czech Republic.
Copyright UNION Agency – Science Press

cumbersome when the head must be rotated to see the whole display.

This paper presents the *Escritoire*; a desk on which sheets of *virtual paper* can be read and manipulated. It is the size of a standard, non-computerized desk (about 0.9 by 1.2 metres). Large vertical screens are typically used for presentations or visualization, but the *Escritoire*'s horizontal display acts more like an architect's drafting board, allowing the user to sit comfortably and peer over the documents as if they were papers on a desk. Two digital projectors form a *foveal display* – one projector covers the desk, creating a large visual context, while the other provides a high-resolution inset for detailed work. The virtual paper is more like real paper than conventional windows, and can be manipulated by both hands using cordless pens. Images and documents can be arranged and annotated on the desk, and a client-server design supports distributed collaboration. The new interface encompasses the desktop metaphor by allowing existing application programs to be used alongside the virtual paper.

2. RELATED WORK

The interface of the *Escritoire* builds on three main areas of previous research – visualization techniques that aim to make more information available on a conventional display, personal interfaces that present a large display area, and displays that combine multiple digital projectors.

Visualization Techniques

Focus plus context techniques [Leun94] distort a graphical representation of a data set to provide a detailed focal area while still accommodating the entire set on the screen. An example is the Document Lens [Rob93] which displays an entire multi-page document while allowing the user to simultaneously view any part of it at a legible resolution. Pad++ and its successor Jazz [Jazz] take a different approach, creating an infinite 2-dimensional space containing graphical objects, where the user alternates between viewing large areas at low resolution, and zooming in to view parts in detail. Data Mountain [Rob98] replaces the favourites menu on a web browser with a set of miniature graphical versions of the pages. Experiments showed that spatial memory reduced subjects' reaction times and error rates – this result is exploited by the *Escritoire* which also displays a large group of documents from which a desired one can be quickly identified then viewed in detail.

Large Personal Interfaces

Preliminary work on the *Escritoire* involved devising a new wall-sized display controlled using a 6 degree-of-freedom tracker [Ash01]. 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. Kimura [Mac01] displays representations of a user's tasks on a projected 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' [Bau01] – a very large monitor that has high resolution in the focal region, plus a large contextual area. This system superficially seems similar to the *Escritoire* because it combines a large context with a high-resolution focus, but it is purely an enlarged version of the conventional interface. Conventional programs are controlled using the standard peripherals – the keyboard and mouse – and no projector calibration is performed. Sukthankar et al. [Suk01] have developed a method for efficiently calibrating a projected display to a non-orthogonal surface for use when making presentations, and a similar method is used by the *Escritoire* to pre-warp projected images.

Earlier work at Cambridge on the DigitalDesk [Wel93] combined real paper and projected digital information on a horizontal desk surface. Further work investigated animated paper documents [Rob97] in which a normal book, when opened on the desk, would be recognized using computer vision techniques and augmented with interactive graphics. Later systems have pursued the same goals [for example Kob98] but the aim of the *Escritoire* is not to provide an experience half way between physical and virtual like these systems, but rather to capture many of the features that make using paper attractive, while also supporting the attractive features of virtual documents, such as the computer's complete control over the workspace, and remote collaboration




Figure 2. An unmodified XWindows program is projected onto the wall and controlled at a distance using a 'wand'.

between multiple users.

Multi-Projector Displays

Many groups have worked on combining multiple digital projectors to form a display wall for presentations or scientific visualization. One example is the Stanford Interactive Mural [Gui01] that uses 12 projectors driven by a cluster of 32 Linux PCs connected via a 1GB/s network. WireGL, and its successor Chromium [Hum02] allow OpenGL rendering to be distributed to multiple networked computers. The high price of multi-projector display walls with many compute and render nodes means that they are usually shared facilities for research laboratories, but the *Escritoire* is designed to be a personal system – an augmentation of the personal computer. The Office of the Future project [Ras98a] looks ahead to a time when an office can be augmented with a multitude of inexpensive projectors and cameras to provide an immersive display for a head-tracked user. In contrast the *Escritoire* uses a single standard PC with two graphics cards to drive its projectors and receives pen input. The main expense is that of the projectors, whose cost and size is falling as the technology improves and they become more widespread.

2. FOVEAL DISPLAY

The display of the *Escritoire* comprises two digital projectors that are combined to create what we call a *foveal display*. Figure 3 depicts the hardware arrangement. A projector on the floor reflects its image in a mirror on the bottom of a shelf to create 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 to create the *fovea* – a small, high-resolution display that accommodates the




Figure 3. The two-projector arrangement of the foveal display. The user sits on the right.

user's focus of attention. The desk is actually a 0.9 by 1.2 metre digitizer 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 high refresh rate. The fovea uses a 1024 by 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 by 480 projector and a back-silvered mirror.

When the fovea and periphery images are projected they are distorted because the projectors are not aligned to the desk. Precise mechanical calibration would be prohibitive in a general home or office system so the *Escritoire* compensates for the distortion using commodity graphics hardware.

Co-ordinate Spaces

Four 2-dimensional co-ordinate spaces are employed by the *Escritoire*: graphical objects are arranged in *desk space* which has device-independent units of metres, events from the pen occur in *pen space*, and the image to be displayed is formed in *texture space* then transformed to a quadrilateral in *framebuffer space* before being projected onto the display surface. Figure 4 depicts the co-ordinate spaces and transformations between them. The distortion \mathbf{H}_{fp} undergone by an image as it is projected onto the surface is assumed to be the reverse of a pin-hole camera – a projective transformation, also known as a planar homography. The 2D image in the framebuffer becomes a distorted 2D image on the surface so 3D graphics hardware is used to undo the distortion with \mathbf{H}_{tf} to get a rectangular image on the surface. Because a point can be converted from one space to any other via a projective transformation (possibly by inverting or composing the transformations in the figure) points can be represented by homogeneous vectors of the form (x,y,w) , transformations are achieved via multiplication by 3 by 3 matrices, and the warping of the graphics from texture to framebuffer can be performed by a commodity video card.

Transformation \mathbf{H}_{pf} from pen to framebuffer is calculated by requiring the user to select a number of projected targets on the surface. It has the form,

$$\mathbf{H}_{pf} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix},$$

where the nine constants a to i can be determined from four or more correspondences between pen (x,y) and framebuffer (u,v) points (the scale of the matrix is unimportant so there are only eight degrees




Figure 4. The four co-ordinate spaces and the projective transformations between homogeneous points in those spaces. The units of measurement are shown under each space. H_{fp} is the warp undergone by the projected image which depends on the characteristics and location of the projector, and H_{ff} is the warp performed by the video card to correct that distortion.

of freedom). Hartley and Zisserman [Har00] explain how this projective transformation can be calibrated using the following equation,

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

The matrix, which we will call \mathbf{A} , is created by stacking the 4, or more, point correspondences. The values a to i are obtained as the eigenvector corresponding to the smallest eigenvalue of $\mathbf{A}^T\mathbf{A}$. Planar homographies have been used in this way by Sukthankar [Suk01] using a camera to detect projected calibration points. Adding a camera to the Escritoire would give no time saving for the user because selection of a number of projected targets

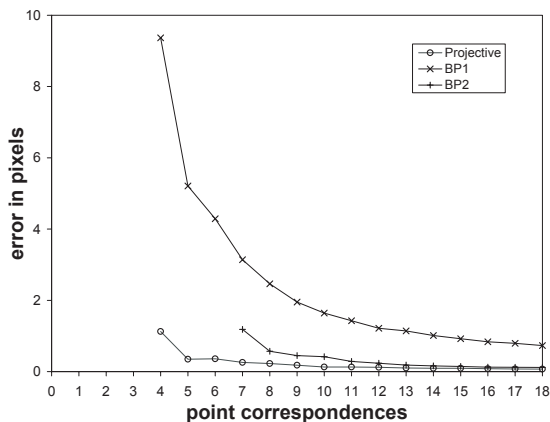


Figure 5. Average error for transformations to correct the distortion in the projected image.

would still be required to calibrate the digitizer.

Once transformation H_{pf} has been calculated, the user chooses the position of the rectangular display on the surface using the keyboard to scale, translate and rotate it. Transformation H_{pd} , a similarity transformation, is updated to reflect the changes. H_{dt} is a scaling matrix determined by the size of the rectangular display and the known resolution of the projector. The fovea and periphery each have their own H_{pf} and H_{dt} transformations corresponding to their separate textures and framebuffers, but H_{pd} is associated with the pen input device and independent of the projectors.

Image Warping

Once the four transformations have been calculated, points can be transformed between any two co-ordinate spaces. Events from the pen are transformed to desk-space locations. When the user calibrates the display, the vertices of the quadrilateral in the framebuffer to which the texture must be warped are calculated. The x , y and homogeneous w values are used to create a triangle fan in the framebuffer which is texture mapped with the image and passed to an 'immediate mode' rendering routine. Commodity 3D video cards can easily perform the perspective transformation at the required 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™ video cards. The time taken to warp the texture is negligible – performance depends on the ability to *update* the texture rapidly. Projection by the periphery projector over the fovea display is avoided by transforming the quadrilateral in the fovea framebuffer to the common desk-space, then transforming it to the periphery framebuffer, where a black quadrilateral is drawn over its contents.

We evaluated the use of the projective transformation to compensate for distortion of the projected image by displaying a 16 by 10 grid of targets on a 640 by 480 projected display, selecting them with the digitizer pen, then calibrating transformations with a subset of the points and calculating the average error in the mapping for the remaining points. We used a projective transformation as described above, and also bivariate polynomials which have the form,

$$u = c_1 + c_2x + c_3y + c_4xy + c_5x^2 + c_6y^2,$$

$$v = c_7 + c_8x + c_9y + c_{10}xy + c_{11}x^2 + c_{12}y^2,$$

where the constants c_1 to c_{12} are obtained from point correspondences in a similar manner as for the projective case. Figure 5 shows the results. BP1 is the form of the bivariate polynomial used on the DigitalDesk [Wel93] where $c_5 = c_6 = c_{11} = c_{12} = 0$. This was used to compensate for minor distortion rather than for the general projector positioning used with the Escritoire. BP2 is the full version with twelve degrees of freedom. The projective transformation performs best, and achieves an average error of just 0.35 pixels with five point correspondences, which validates the choice of this transformation for the Escritoire's display.

Pen Input

Koike et al. [Koi01] have used an infrared camera and computer vision techniques to allow a person to interact with projected graphics using their hands. An interesting new technology for multiple users is DiamondTouch [Die02] which allows a number of users to interact simultaneously with a large, touch-sensitive surface. We have avoided using cameras in favour of dedicated input devices, which are robust, and more accurate for tasks like writing. Digitizer tablets are available that support multiple pointing devices at different times, but we have found no large-format device that supports multiple pens simultaneously, so we have added a Mimio™ pen to




Figure 6. The 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 implementations on the server.

the digitizer. The two do not interfere because the Mimio is based on ultrasound while the digitizer uses electromagnetism. Similar to the distinction between the foveal and peripheral displays, the digitizing tablet provides a high resolution ($\pm 0.25\text{mm}$) three-button pen for the user's dominant hand, while the ultrasonic pen for the non-dominant hand has lower resolution and no buttons. The Mimio receiver, normally mounted on the side of a whiteboard, is positioned along the top of the digitizer to avoid obstruction of the ultrasound paths by the user.

2. CLIENT AND SERVER

The software of the Escritoire consists of a client and a server. This is to allow a pair of desks with their client software to connect to a single server to provide a shared *task space* that will augment the *person space* of a video conference. The client is 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. We run the client and server on separate PCs which communicate via TCP/IP over a standard 100MB/s Ethernet network.

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 6. The table below summarizes the differences between the client and server software.

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

Protocol

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, a **ready** message indicates that the client has updated the display and is ready for more messages, **pen** messages signal the pen state and location independent of device resolution, and **keyboard** messages signal key events. The server sends six types of message; **create tile** causes the client to make a new tile of a specified




Figure 7. Virtual paper on the desk display: a set of slides (top left, in the periphery), an image (bottom left), a document being dragged by the left hand (top right), and an XWindows program being operated by the right hand (bottom right).

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

A design decision for the Escritoire was whether updates to the tiles should 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, but client pull will allow traffic in dense periods to be synchronized to the client's display updates – assuming graphics processing on the client is the bottleneck, this will allow successive updates to the same area of a tile to be coalesced at the server, thus gracefully degrading the frame rate.

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

receives, in reply, 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.

3. 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 incoming events into a single queue, then passes them to the portfolio via a callback mechanism. The portfolio updates its graphical appearance then informs the server of the modified region so the appropriate tile can be updated. The four types of portfolio 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. Most events are passed on to its children.

Jotter

The jotter allows an image in a standard format to be edited on the desk. The user positions it 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 work on the image as predicted in experiments by Buxton et al. [Kab93].

PDF

Another type of portfolio makes Adobe PDF documents available. An example in Figure 7 is being dragged by the left hand. A word-processed document or slide show is output to PDF from any standard application program, which is analogous to printing it to paper. Multiple documents can be arranged on the desk, and the fovea allows two documents to be placed side by side for comparison. Annotations are added in a vector format to the PDF file which can then be viewed offline using software such as Adobe's Acrobat Reader™.

VNC

Virtual Network Computing (VNC) [Ric98] is a remote display system that translates a conventional graphical user interface into a protocol that allows, for instance, the display of an X server to be viewed and controlled from a Windows machine. The VNC portfolio allows a computer with a conventional interface to be controlled from the Escritoire. Figure 7 shows an XWindows display lying on the desk with 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 because they are optimized for bitmapped displays.

4. DISCUSSION

The Escritoire's large foveal display fills user's vision, allows data to be placed as far away as the arms can reach, and reproduces the kinaesthetic sense one gets from working with physical media such as paper.

As was found 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 areas around the pen tips are not shadowed (Figure 8), and partly because people are accustomed to their hands shadowing light from above so they automatically move them if the information underneath is occluded.

Work on the Office of the Future project has addressed intensity blending between projected images [Ras98b], but rather than being a problem here, the visible edge between fovea and periphery is




Figure 8. The projector is above and in front of the hands, so the information near the pen tips is not shadowed.

a useful cue to the boundary of the fovea, and the higher intensity of the fovea image due its light being spread over a smaller area complements the increase in resolution.

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 could then be combined with a similarity transformation for each sheet, enabling rotations and scalings with virtually no extra computational cost.

The Escritoire presents a document-centred view of work being undertaken. The user has direct access to all of the documents on the desk, and can immediately bring any one to the fovea (Figure 9). The action of grabbing a document to view it in more detail comes instinctively. As with paper, modifications to the documents are not saved at distinct times at the command of the user, but instead their state and location is saved automatically and preserved between sessions.

People use desk organization to remind themselves to




Figure 9. Dragging a document into the fovea.

of tasks, and to loosely categorize documents into piles that complement more formal, long-term filing systems [Mal83]. The *Ecritoire* supports this behaviour, but this time the messy desk contains 'virtual mess' that disappears at the touch of a button. The arrangement of documents could be enhanced by allowing them to be placed at any angle. This feature for conventional windows has been developed by Beaudoin-Lafon [Bea01].

5. FUTURE WORK

The 'pile' metaphor [Man92] was developed at Apple Computer to support informal categorization and browsing. Piles that are modified with gestures from the dominant hand are being added to the *Ecritoire*, but unlike the previous work the piles do not contain document icons, but rather the editable documents themselves.

A second version of the hardware will be assembled and the pair will be used to augment a video conference. This is expected to prove useful for remote parties wishing to discuss and mark up documents and presentations – common tasks for employees of a multinational company. There will be no distinction at the server between pen events from different users, which should create interesting possibilities for co-operation in bimanual tasks. Freeman [Bie91] has considered such fine-grained interaction between multiple users.

The earlier work to create a projected interface on the wall [Ash01] is being combined with the desk, adding a third surface with greater size and lower 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 extended to the user interface to construct a hierarchy of interactive surfaces, each of which has greater capacity than the preceding but which is less precisely rendered and controlled.

6. ACKNOWLEDGMENTS

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

7. REFERENCES

- [Ash01] Ashdown, M. and Robinson, P. The Writing's on the Wall: Large, Remotely Controlled Displays, proc. European Conference on Computer-Supported Collaborative Learning (Euro-CSCL) '01, pp. 83-88, 2001.
- [Bau01] Baudisch, P. et al. Focus Plus Context Screens: Combining Display Technology with Visualization Techniques, proc. UIST '01, pp. 31-40, 2001.
- [Bea01] Beaudouin-Lafon, M. Novel Interaction Techniques for Overlapping Windows, proc. UIST '01, pp. 153-154, 2001.
- [Bie91] Bier, E. A. and Freeman, S. MMM: A User Interface Architecture for Shared Editors on a Single Screen, proc. UIST '91, pp. 79-86, 1991.
- [Die02] Dietz, P.H. and Leigh, D.L. DiamondTouch: A Multi-User Touch Technology, proc. UIST '02, pp. 219-226, 2002.
- [Gui01] Guimbretiere, F. et al., S. and Winograd, T. Fluid Interaction with High-Resolution Wall-Size Displays, proc. UIST '01, pp. 21-30, 2001.
- [Har00] Hartley, R. and Zisserman A. Multiple View Geometry in Computer Vision, Cambridge University Press, 2000.
- [Hum02] Humphreys, G. Chromium: A Stream Processing Framework for Interactive Rendering on Clusters, proc. Siggraph '02, pp. 693-702, 2002.
- [Jazz] Jazz Project, <http://www.cs.umd.edu/hcil/jazz/>
- [Kab93] Kabbash, P. MacKenzie, S.I. and Buxton, W. Human Performance Using Computer Input Devices in the Preferred and Non-preferred Hands, Proc. InterCHI '93, pp. 474-481, 1993.
- [Kob98] Kobayashi, M. and Koike, H. EnhancedDesk: Integrating Paper Documents and Digital Documents, proc. Asia Pacific Computer Human Interaction (APCHI) '98.
- [Koi01] Koike, H., Sato, Y. Kobayashi, Y. Integrating paper and digital information on EnhancedDesk: a method for real-time finger tracking on augmented desk system, ACM Trans. on CHI, 8 (4), pp. 307-322
- [Leun94] Leung, Y. K. and Apperley, M. D. A Review and Taxonomy of Distortion-Oriented Presentation Techniques, ACM Transactions on CHI, 1 (2), pp. 126-160
- [Mac01] MacIntyre, B. Mynatt, E.D. et al. Support for Multitasking and Background Awareness Using Interactive Peripheral Displays, proc. UIST '01, pp. 41-50, 2001.
- [Mal83] Malone, T. W. How do People Organize Their Desks? Implications for the Design of Office Information Systems, ACM Trans. Office Info. Systems 1, No. 1, pp. 99-112, 1983.
- [Man92] Mander, R. et al. A 'Pile' Metaphor for Supporting Casual Organization of Information, proc. CHI '92, pp. 627-634, 1992.
- [Pem96] Pemberton, S. Views and Feelings: Metaphorically Speaking, SIGCHI Bulletin 27, No. 4, p. 96, 1996.
- [Ras98a] Raskar, R. et al. The Office of the Future: A Unified Approach to Image-Based Modelling and Spatially Immersive Displays, proc. of Siggraph '98, pp. 179-188, 1998.
- [Ras98b] Raskar, R. Welch, G. and Fuchs, H. Seamless Projection Overlaps using Image Warping and Intensity Blending, proc. Virtual System and Multimedia (VSMM) '98, 1998.
- [Ric98] Richardson, T. et al. Virtual Network Computing, IEEE Internet Computing 2, No. 1, pp. 33-38, 1998.
- [Rob93] Robertson, G. G. and Mackinlay, J. D. The Document Lens, proc. UIST '93, pp. 101-108, 1993.
- [Rob97] Robinson, P. et al. A Framework for Interacting with Paper, proc. Eurographics '97, 16, No. 3, 1997.
- [Rob98] Robertson, G. Czerwinski, M. and Larson, K. Data Mountain: Using Spatial Memory for Document Management, proc. of UIST '02, pp. 153-162, 2002.
- [Sel97] Sellen, A. and Harper, R. Paper as an Analytic Resource for the Design of New Technologies, proc. CHI '97, pp. 319-326, 1997.
- [Suk01] Sukthankar, R. Stockton, R.G. and Mullin, M.D. Smarter Presentations: Exploiting Homography in Camera-Projector Systems, proc. ICCV '01, pp247-253, 2001.
- [Swa97] Swaminathan, K. and Sato, S. Interaction Design for Large Displays, Interactions 4, No. 1, pp. 15-24, 1997.
- [Wel93] Wellner, P.D. Interacting with Paper on the DigitalDesk, Communications of the ACM 36, No. 7, pp. 87-97, 1993.