

Large-Format Computer Interfaces

Ph.D. First Year Report and Research Proposal

Mark S. D. Ashdown

University of Cambridge Computer Laboratory,
J J Thomson Avenue, Cambridge CB3 0FD, UK

www.mark.ashdown.name

July 2000

Contents

- 1 Background
 - 1.1 Introduction
 - 1.2 Space
 - 1.3 Display
 - 1.4 Bimanual Input
 - 1.5 Collaboration
 - 2 Related Developments
 - 2.1 Multiple Virtual Displays
 - 2.2 Visualisation Techniques
 - 2.3 Large Displays
 - 2.4 Tangible User Interfaces
 - 2.5 Virtual Environments
 - 2.6 Virtual Desktops
 - 2.7 Bimanual Input
 - 2.8 Collaboration
 - 3 Work Done
 - 3.1 Pen Input
 - 3.2 Projected Display
 - 3.3 VNC
 - 3.4 Hardware
 - 4 Research Proposal
 - 4.1 Objectives
 - 4.2 Foveal Display
 - 4.3 Virtual Paper
 - 4.4 Bimanual Input
 - 4.5 Collaboration
 - 4.6 User Study
 - 4.7 Timetable
- References

1 Background

1.1 Introduction

The advent of the personal computer with a graphical display [11,12] brought predictions of the paperless office, but this dream has not been forthcoming. Paper has affordances that have not been supplanted by the abilities of computer applications, and arguably computers now create more paper than they have replaced. Highly structured work involving precise data now utilises computers that provide large storage capacity and rapid access – a library is now searched using a computer rather than a cabinet of cards. Less structured work, such as writing a paper that draws on ideas from various sources, still generally makes use of paper documents [1]. Documents that are created by users or obtained from digital libraries start off in digital form, but are usually printed before being used. These physical documents and the offices in which they reside provide methods of working not duplicated in computerised information environments, like the ability to avoid the overhead of categorising information that is going to be used soon, by leaving it on the side of one's desk

The hardware and software of e-books – batteries, memory, LCD displays, networking, standard formats, display systems, user interfaces – has progressed to the point where they seem a viable alternative to paper books, but their adoption now depends on other factors – distribution, encryption, availability of compelling content, copyright protection, billing and royalties [69]. Similarly, standards such as PostScript and 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 abilities will become very useful, such as its ability to be used simultaneously by people in different locations at the same time.

The purpose of this research is to explore ways in which the wider affordances of paper can be replicated by a computer interface: use of space in a work environment, high resolution when viewed closely, natural interaction using two hands, and document-centred collaborative activity.

1.2 Space

For decades, the computer has almost exclusively presented a visual interface consisting of a screen with a diagonal measuring 14 to 21 inches and a keyboard and mouse for input. Mark Weiser's vision of Ubiquitous Computing [10] includes tabs, pads and boards – display and interaction devices of different sizes posing different challenges of implementation and man-machine interface design. The archetypal personal computer has a pad-sized display. There has recently been much investigation into tab-sized, small-screen technology applicable to the many mobile devices now being developed to exploit the economical and near ubiquitous access now available to wireless digital networks. Research in this area has included text entry using pens on PDAs, and systems for scanning prose and lists of textual items, for example. When one is at home or in the office, in a stable environment where the computer is used on a regular basis, the desire is for the largest and most versatile display surface possible. As the technology becomes cheaper and better supported, board-sized displays will become the interface of choice for many users as suitability for the task, rather than technological and economic constraints, becomes the criterion for selecting a display system.

Interaction with programs on the typical graphical user interface found today is accomplished via 'windows', but monitor screens often allow only one window to be used at a time because the others are partly or wholly obscured. By analogy to the use of real paper on a real desk, this is equivalent to having a 9 by 12 inch desk, just big enough to hold a single stack of papers (Figure 1). Writing a paper based on a number of documents involves alternately moving the documents to the top of the stack to read them, then moving a writing pad to the top of the stack to write one or two sentences. By analogy with virtual memory used in operating systems, the absence of simultaneous access to the current working set results in




Figure 1 The 9 by 12 inch desk

thrashing. The overhead involved in swapping limits the level to which tasks can be interleaved, qualitatively restricting the user's work practices.

O'Hara and Sellen [2] report on a study in which participants were asked to read and summarise a document. It compares the strategies employed by those using paper with the strategies employed by those reading and composing textual information on a computer using a commonly used word processing application. The critical differences centre around the major advantages paper offers in supporting annotation while reading, quick navigation, and flexibility of spatial layout. The use of two hands and the tangibility afforded by paper allowed quick movement through paper documents and anticipatory page turning, allowing reading, writing and navigation to be interleaved and overlapped. By contrast, in the on-line condition subjects did not make the frequent back and forth movements between making notes and reading. Actions lacked continuous feedback, and spatial constraints of the interface interfered with quick, flexible movement. When using paper, subjects laid out documents in space to gain an overall sense of structure, facilitate cross-referencing, and interleave reading and writing. The spatial arrangement was dynamic with pages being moved in and out of the centre of attention. Subjects using the computer found that when a document was viewed at a readable resolution, less than one page would fit on the screen. This meant that much of the necessary contextual information for developing a sense of text and location lay beyond the window boundaries.

Bolt [7] notes that 'Psychological research supports the notion that people learn readily, and without effort, the whereabouts of things in space' and highlights three points about information held by computers: (i) Word processors are superb instruments for handling information but they don't let you see the overall pattern of organisation (ii) Information can be managed spatially—that is, it can be assigned to certain fixed positions in space to help us recall its location. (iii) Computers, unlike libraries, have a way of hiding the whereabouts of information. It was noted in 1968 [9] that compared to traditional methods using paper and books, space is unimportant to modern information processing systems, but rather than erasing all spatial aspects they could be placed wholly at the disposition of the user. In order to provide a spatial frame of reference for information, an information system might build in or simulate such a frame.

1.3 Display

The limited size of CRT monitors has motivated projects to create large displays which have various benefits: scientific visualisation is possible on the PowerWall [41], a six by eight foot, eight megapixel display; immersive environments such as the CAVE [42] can be created; and new interaction techniques can be tested on displays like the eight-projector Interactive Mural of the iRoom at Stanford [43].




Figure 2 $16 \mu\text{m}^2$ mirrors on a Digital Micromirror Device chip

Texas Instruments' Digital Light Processing (DLP) technology uses, as the basis for a projection system, a Digital Micromirror Device, which is an array of mirrors built over conventional CMOS SRAM address circuitry (Figure 2)[26]. Currently up to 2.3 million of the $16 \mu\text{m}^2$ mirrors can be produced on a chip, each being independently able to tilt by ± 10 degrees when switched on or off. Such projection systems are now replacing film projectors in cinemas. DLP, LCD and plasma screens provide the final link in the digital video infrastructure. Large display surfaces using these technologies will become commonplace in the home and office as CRTs are phased out, enabling new applications and providing new challenges for interface design.

Level-of-detail degradation in the periphery of head-mounted displays has been explored [21] as a means of raising frame rates and reducing latencies, while causing minimal impact on usability. This idea could be applied to large projected displays. Wellner describes the DigitalDeskLamp as a future direction from work on

the DigitalDesk [3], which could be pointed at surfaces in a room to reveal electronic objects. The balance between resolution and display area could be varied by moving the lamp. The Office of the Future project described below aims to use multiple projectors and cameras to display images on physical objects in the office. A flexible setup for projectors is envisioned whereby multiple projectors can be used to project into the same space in order to achieve higher resolution, or increase light intensity, although by the authors admission the setup is envisioned for the distant future and requires many input/output devices.

1.4 Bimanual Input

As noted above, computers are almost exclusively controlled by keyboard and mouse, with control being switched between the two for different actions. These input devices have proved to be very versatile, but research into alternatives suggests that improvements are possible, especially when two hands are used together [5].

The graphical objects that form the user interface of a computer are assigned two roles by Buxton [28]: they convey information about their function or about the subject of a task, and they act as controls that are manipulated by the user. From the display perspective the GUI is space multiplexed. However, virtually all interaction is undertaken using a single input device, the mouse for example, hence from the control perspective the GUI is time multiplexed. Figure 3(a) illustrates acquisition times involved in performing alternating tasks A and B in a time multiplexed GUI, which are necessary because graphical display objects must be alternately selected. Figure 3(b) illustrates the gains that are possible by assigning a distinct input device to each of the user's two hands. Each device can return to a "home" position after performing an action eliminating task-switching times, and tasks may be overlapped. Buxton proposes that all direct manipulation actions can be categorised as the combination of a foreground and a background task,

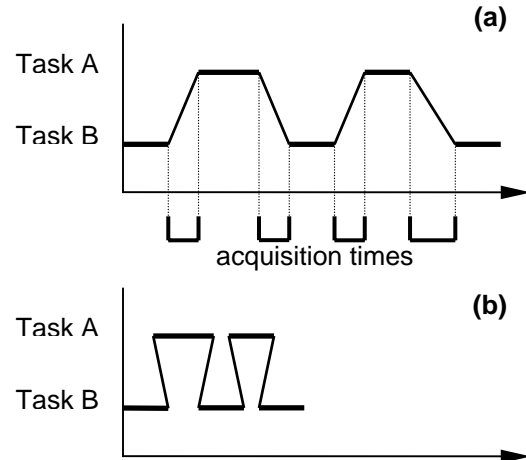


Figure 3 From [28] (a) acquisition times for alternating tasks (b) overlapping foreground and background tasks

performed by the dominant and non-dominant hands respectively. Experimental data suggests that performance increases can be gained by splitting the sub-tasks of compound continuous tasks between the two hands [29]. Even novice users have the requisite manual skills, and significant improvements can be made over common one-handed techniques. Bimanual manipulation is argued [40] to have two types of advantages: manual benefits from increased time-motion efficiency due to twice as many degrees of freedom being available to the user, and cognitive benefits as a result of reducing the load of composing and visualising the task at an unnaturally low level imposed by traditional unimanual techniques.

1.5 Collaboration

The discipline of *Computer-Supported Cooperative Work* (CSCW) arose in the late 1980s when the near ubiquity of networked computers in organisations prompted researchers and developers to start thinking about systems that let people work together [60]. The increasing power of personal computers, and advances such as the new display devices described above, will enable the now commonly used *groupware* systems such as email, to be joined by systems providing real-time interaction in media spaces.




Figure 4 The Relationship between Telecommunication and Artificial Reality [50]

Technology that facilitates *telepresence* – that which establishes a sense of shared presence or shared space among geographically separated members of a group – can create *person space* and *task space* [49]. Person space is the collective sense of copresence between group participants that is created by, for example, videoconferencing. Task space is generated by copresence in the domain of the task being undertaken, where each participant can make indications, annotations and changes to a copy of the data,

which are then made apparent to the other members of the group. Figure 4 illustrates the difference Krueger noted [51] between telecommunication and *artificial reality*, where with the latter the communication between two places creates a new place consisting of the information that is shared. A person space may be some amalgam of places A and B in Figure 4(a), while a task space may be a focus for the work of both parties as in Figure 4(b). This document describes a system that will provide a task space for participants.

CSCW can be undertaken by a face-to-face group or by a group distributed over many locations. Furthermore the activity may be real-time interaction or asynchronous, non-real-time interaction [53]. Figure 5 illustrates the space and time categories.

		Time	
		Same	Different
Place	Same	face-to-face interaction (presentation support)	asynchronous interaction (bulletin board)
	Different	synchronous distributed interaction (videophone)	asynchronous distributed interaction (email)

Figure 5 CSCW space-time matrix with example technologies

2 Related Developments

2.1 Multiple Virtual Displays

Rooms [6] is a system developed to address issues of space and spatial structure. The large area and resolution of a desk or dining table on which paper is arranged is compared to a conventional computer display, and the disparity used to motivate development of a system that aims to overcome some limitations of the conventional graphical user interface. Many workspaces, called rooms, can be created for different tasks, with links between them, represented as doors, forming a web that may be traversed. One room may be viewed at a time. A window may appear in a number of rooms with a different size and location in each. Windows can be moved between rooms and one room can be contained within another. An overview screen shows the rooms available, the links between them, and a miniature depiction of the contents of each.

Microsoft's Task Gallery [13,14] aims to provide support for task management and document comparison. 3D rendering hardware is used to generate a navigable spatial metaphor, which is reinforced by animation. The walls of the gallery are adorned with tasks (Figure 6) like those in Rooms, which have background images that provide a visual cue to differentiate them. A task has a set of application windows whose output is texture-mapped onto planar surfaces that are arranged on the background. The user can obtain a balance between context and focus by choosing a view between that of the whole room with all its tasks, and that of just one of the tasks. The user maintains a *selected set* of the windows currently in use – when anything is adjusted the windows in the set are scaled using the 3D hardware so that they all fit onto the screen without obscuring each other. The user's virtual left hand contains a palette that displays frequently used programs and favourite documents, which can easily be added to the current task. This palette is a Data Mountain [15] – a user interface for document management designed to take advantage of human spatial memory, where user arranges documents on a 3D surface using a 2D interaction technique. Documents are represented by images that remind the user of their contents, and a minimum distance between them is maintained to prevent loss of documents by occlusion. User tests showed the Data Mountain speeds up retrieval of web pages when compared to a menu of 'favourites'.



Figure 6 The Task Gallery from [13]. Tasks are positioned on the walls like paintings, with the current task filling the back wall.

2.2 Visualisation Techniques

Pad++ [19] was developed to make it easier and more intuitive to find specific information in large databases. Information is situated on an infinite 2D plain around which the user navigates by moving horizontally and vertically to find different objects, and zooming in and out to adjust the level of detail. A hypertext document, for instance, is represented as a page of text with the target pages of the links being smaller and positioned to the right. The pages linked from those documents are displayed in the same way relative to their ancestors. The directory tree of a file system is represented in a similar way. A sense of spatial context is maintained when moving between points by continuously zooming and panning whenever moving. When centring on a distant point, the view zooms out until both the old and new foci can be accommodated, then zooms in on the new focus. The paper defines semantic zooming, a technique whereby an object has various representations appropriate for the levels of detail at which it may be viewed

In contrast to the Pad++ approach, ‘fisheye views’ and ‘distortion-oriented presentation’ [23] are designed to provide focus plus context by ensuring that the entire data set is constantly visible. This is achieved by distorting the view of the graph, map or text document to allow a detailed view of the current focus and a course view of the rest of the data to be accommodated in a single display (Figure 7). This idea has been used in a multi-user text editor [25] that displays the whole of the current file in a small font, except for the user’s focus around which the font size is progressively enlarged. The activity of other users is signalled on the display with colour, and optionally some level of zoom applied to their focal positions in the text. The Document Lens [24] treats a document as a series of pages arranged in a 2D array. The user grabs a rectangular lens and pulls it around to focus on the desired area at the desired magnification. The presentation outside the lens is stretched to provide a continuous display of the global context. The stretching is implemented efficiently using affine transformations, allowing text documents to be viewed as a whole with an interactive visualisation.



Figure 7 Interactive Atlas from [23] presenting Scottish landmarks. Objects representing points of interest are embedded into the transformed coordinate space.

2.3 Large Displays

Early work on new interaction paradigms using large displays was undertaken at the Media Room at MIT’s Media Lab, which housed a number of projects concerned with spatially arranging and manipulating information. Dataland [7] provides the user with a large 2D graphical space in which to arrange data items such as maps, books, letters, calendars, photographs, motion video pictures and programs like the calculator used in the project. A low-resolution view of the space is displayed on a monitor at the user’s side, with a moveable “you-are-here” marker indicating the portion that is displayed in high detail on the back-projected, wall-sized display of the Media Room. The large screen is like a magnifying glass over Dataland. The marker is moved using joystick, the touch-sensitive surface of the small screen, or voice commands. A related project, the World of Windows [22], presents an ensemble of simultaneous video streams on the large screen. The multiplicity of such a display would ordinarily make coping with it untenable, but the intent of the research is to introduce order and control through a gaze-orchestrated interface. The system presupposes an information-gathering network that determines what information should be presented, and that sizes and groups the windows on the screen according to their information content. An eye-tracking mechanism calculates the user’s gaze direction. This is used to control the volumes of the soundtracks of the video streams either by increasing the volume of the selected stream, scaling the volumes by their spatial distance from the selected stream, or scaling the volumes by their temporal distance from the selected stream so the less recently selected streams are quieter. Gaze direction can also be used to identify the stream at which the user is looking and zoom in on it, following a time-out period or a joystick action. Applications are proposed for eye tracking such as non-uniform rendering of images so that the centre of the user’s attention is drawn in higher detail than the rest of the picture. “Put-That-There” [20] explored the combination of voice and gesture in manipulating simple items on the wall of the Media Room. Shapes of various forms and colours can be called into existence, moved about, replicated, have their attributes altered, and ordered to vanish. Pronouns and deictic references can be used, thus if the phrase “Put that there” is uttered, voice recognition software initiates sampling of the pointing direction of the tracker on the user’s finger when the words “that” and “there” are recognised. The selected object is then moved to the specified location.

Many of the early attempts to build computer support for meeting rooms included whiteboard-sized displays to provide a shared focus of attention, but these were controlled at a distance by keyboard and

mouse. The LiveBoard [55] was created at Xerox PARC to provide a directly interactive, stylus-based, large-area display. It is a back-projected display, approximately 4 x 3 feet, with a cordless pen that emits an optical beam that is located by a camera behind the screen, allowing the pen to be used at a distance. The pen could be used to control X Windows programs via a special driver. The most popular programs turned out to be the whiteboard and the slide show, which allowed annotation of slides. A survey of users revealed that above all else, people would like to have a display of a higher resolution – the resolution of the version of LiveBoard in the paper was 1120 x 780. Tivoli is a whiteboard application for the LiveBoard designed to support informal workgroup meetings. It is targeted at meetings where participants work closely together on a problem or idea, rather than at presentations, for example. The necessary core functionality was identified and implemented in Tivoli – simple pen and gesture scribbling and editing, multiple pages, saving and retrieving, printing, and importing images – while areas of extended functionality – remote collaboration, meeting management tools, integration with other ubiquitous devices – were anticipated as further research. Interaction with Tivoli prompted reconsideration of many tacit assumptions of traditional GUIs. A pen allows natural gestural input, does not require a cursor, and generates direct position data unlike the relative movements of a mouse that can be scaled to alter sensitivity. A very large interactive surface necessitates careful sizing and positioning of buttons, menus and pop-up messages, and places emphasis on keeping user control at the pen through techniques like gestural commands. The LiveBoard has three pens to support multiple users. One consequence is that tools on the display are shared, so for example, an icon cannot be highlighted to show the current pen line thickness since it can be different for different pens: a distinction must be made between *pen state* and *system state*. Work on Tivoli has continued, with emphasis on the structure and manipulation of object-oriented data.

Flatland [16] uses a digital projector and a wall-mounted board that can detect strokes with a stylus. It behaves like a conventional whiteboard for much of the time; simply displaying information until a pen is picked up and the user starts writing. The board can be split into segments by drawing lines, and information can be added to a segment by drawing inside it. The active segment grows while the inactive ones shrink to create more space without necessitating the erasure of information, and segments can be given behaviours so that the user's stylus strokes are interpreted, for instance, as parts of a map. On such a display surface, interaction techniques for creating and organising information must be considered carefully [36,37].

Rekimoto has developed a multiple device approach for supporting whiteboard-based interactions [30] because the large surface of such a display makes traditional GUI design ineffective. A PDA provides




Figure 8 The Pick-and-Drop operation [30,32]. The user taps the pen on an object in the PDA, then taps again on the large screen to place it

each user with a private workspace that complements the public space on the projected display of the whiteboard. The two linked displays are in Weiser's board and tab categories respectively. In a drawing application the PDA screen provides a tool palette and a temporary work buffer. The stylus used to control the PDA also works with the whiteboard so, using the pick-and-drop technique (Figure 8) [32], the user can select an object stored on the PDA then place it on the whiteboard. This technique can be used by a lecturer to create a whiteboard-style presentation using prepared digital information. Text entry facilities on the PDA can allow text to be written then copied to the large screen, rather than drawn as free-form strokes, and application data such as configuration information and help files can be browsed

without presenting them to the audience. An interaction technique called scoop-and-spread [33] uses gesture commands from a tracked PDA to scoop up images from a desk, or spread them out from the PDA

Rekimoto has also produced a computer augmented environment that uses cameras to identify and register the location and orientation of objects on a desk using printed visual markers – 2D *glyphs*

similar to bar codes [31]. Projectors produce images on the desk and wall (Figure 9). A technique called hyperdragging allows an object to be dragged off the edge of a notebook computer onto the desk or wall. Data can be associated with objects or transferred to other computers by dropping it on the desired destination, which requires only the readily available knowledge about the destination's physical location. The combination of two different displays – a high-resolution small display on the portable computer and a low-resolution large display on the table – represents the user's focal and peripheral information space.



Figure 9 Images are projected onto the desk and wall [31]

Work on The Office of the Future at the University of North Carolina [44] is aimed at building up a unified set of computer vision and computer graphics techniques to enable augmented reality applications. Real-time computer vision techniques will extract depth and reflectance information from visible surfaces in the office, which will be used to project images on the surfaces, render images of the surfaces, or interpret changes in the surfaces. Projectors will be used to create *spatially immersive displays* – personal 3D visualisation systems for a head-tracked user (Figure 10). The approach is designed to allow irregular display surfaces, arbitrarily positioned, overlapping light projectors, and a moving viewer. Calibration of the display involves geometric registration and intensity normalisation [45]. A camera observes the projected location of pixels from a number of projectors. These are used to create a mapping from projected pixels to coordinates in the image viewed by the camera, which are used to warp the projected images so they look correct to the viewer. The extent of projection for each projector is determined by turning on all pixels in the projector and thresholding the image from the camera. The extents can be intersected to determine the projector overlaps, which are used to define intensity blending functions that ensure uniform intensity across the projected display even when more than one projector contributes to a projected pixel. A panoramic display using these techniques has been created [46] as a “proof of concept”. High-resolution, wide field-of-view displays are envisioned that are easy to set-up, use and maintain.



Figure 10 A sketch of the Office of the Future [44]. Multiple cameras infer the geometry and reflectance of the visible surfaces and multiple projectors project images onto them.

2.4 Tangible User Interfaces

Tangible user interfaces augment real-world objects with computational features, in order to provide a natural, tactile experience. The BrightBoard is a whiteboard augmented with computational abilities [34,35]. A computerised interface is not placed between the user and the data, but instead the user draws on

the whiteboard in the usual manner and a camera observes the information. Images on the board can be saved, printed, faxed and emailed by making particular marks. Movement detected by the camera triggers an image to be thresholded, then marks are recognised and their configuration is added to a Prolog program that identifies commands that have been specified and executes appropriate script programs. Audio feedback is given to the user.

The Tangible Bits project [17] presents a vision of interactive surfaces coupling bits, atoms and ambient media. A main focus of this project is the metaDESK [18], which aims to physically instantiate many of the metaphorical devices the GUI has popularised. An application called Tangible Geospace causes the placement of a phicon (physical icon) of a building on the metaDESK, to be accompanied by the display on the back-projected surface of a map of the corresponding area around the building (Figure 11). Rotation and translation of the phicon causes corresponding transformation of the map. The map can be scaled by using two phicons simultaneously. Infrared light from emitters under the projection surface is reflected from the phicons and processed by a computer vision system used to track them. The ‘passive lens’ is a transparent surface incorporating a tethered magnetic tracker, which, when positioned on the desk, provides an alternative view of the section of the map covered by it. The ‘active lens’ is an arm-mounted flat-panel display using similar tracking technology that offers a 3D view of the buildings on the map when the user looks ‘through’ it.




Figure 11 A phicon in Tangible Geospace [18]

2.5 Virtual Environments

The alternative to the *augmented reality* and *tangible UI* approach of projects such as the BrightBoard, DigitalDesk and metaDESK is to immerse the user(s) in an artificial world – a *virtual environment*. Virtual reality headsets make the real world inaccessible so systems, such as CAVE [42], have been created that use large back-projected screens to provide stereoscopic views for a head-tracked user. They allow the user, for example, to see his body in the environment. The Two-User Responsive Workbench (Figure 12) uses frame interleaving to display a stereo image pair to each of two users who can collaborate face-to-face. The independent views provide interface possibilities: different representations of an object can be shown to the users to reflect their differing roles in the collaborative task, the objects can be rendered at different levels of detail depending on a users focus of attention, and users can have private workspaces.



Figure 12 Viewing a shared virtual environment on the Responsive Workbench [57]

2.6 Virtual Desktops

The DigitalDesk [3] is a system for interacting with projected images on a desk using a digitising tablet and stylus or the fingers. Information on real pieces of paper can be augmented with projected graphics. The

information can be transferred to digital documents via a camera using optical character recognition and other computer vision techniques. A number of prototype applications were developed for the system including a calculator operated by the fingers that allows numbers to be copied from a real piece of paper using OCR; and a paint program used to combine pictures on paper and computer-generated images.

A recent study [64] sought to measure the effects of the size and resolution of a digital desk on the performance of sketching and sorting tasks, when compared to a conventional monitor and a tablet with stylus. Participants did not object to the relatively low resolution of the desk. The whole of each display was used, and several participants found the surface of the desk to be too large and the tablet to be too small, suggesting that interfaces designed for large work surfaces should focus on putting primary information near the user and less important information in the periphery.

Interfaces like the desktop metaphor (more accurately called the *office analogy* [67]) are used to display and manipulate information, but they are not based on an understanding of how people actually use their desks, and organise their personal information environments. A study [66] of the way professional and clerical office workers organise the information in their desks and offices has a number of implications for designing 'natural' and convenient computer-based information systems. A very important function of desk organisation is to *remind* the user of things to do, not just to help the user *find* desired information. The location of documents relative to surfaces and to each other can convey attributes such as importance, urgency and subject matter

The *pile metaphor* [68] could be used to arrange documents on a desk display. Visual representations of the documents are arranged using direct manipulation and put in piles (Figure 13) that complement the folder filing system, which is used for more formal archiving. A pile can be sorted by various criteria, e.g. access time. Piles are browsed in different ways: the *edge* method involves looking at the sides of the pile for clues about the items within, the *restack* method involves removing items from the top and dealing with them, in the *hinge* method the pile is hinged open to reveal an item at a time, and finally the pile can be *spread out* to peruse its contents in parallel. The metaphor adds computational services to a system commonly used in managing physical documents.




Figure 13 [68] A pile created by the user hence dishevelled in appearance (left), and a pile created by the system using a set of rules (right)

2.7 Bimanual Input

An important feature of physical information media is the way they can be used with two hands. Toolglass and Magic Lenses [38] are user interface tools for use with bimanual input. A Toolglass widget is coarsely positioned by the non-dominant hand setting up a context in which the dominant hand can act, by clicking through the semi-transparent tool to alter the object below. A Magic Lens is a viewing filter that presents an alternative view of the object over which it is placed. Figure 14 shows a palette of outline colour widgets that can be used to change the line colour of a graphical object, over a magnification lens that increases the size of the graphical object below making the lines easier to select through the two tools. Such transparent tools need no permanent screen space so can be adapted to a wide range of display sizes; because tools are brought to the work area by the non-dominant hand the user's attention can remain focused on the work area; and because operations and views are spatially defined the user can work on a small area without changing the global context. Tablets, two-hands and transparency have been combined in a graphics application [39]. The design goals of the system are to maximise the amount of screen used for artwork, to




Figure 14 An outline colour palette widget over a magnifying lens [38]

avoid forcing the user to divert their visual attention from the artwork, and to increase the degrees of manipulation and comfort of input. Two pucks on two tablets, one for each hand, are used to control a Toolglass and cursor. Button presses and the location and orientation of each puck are used to control the application. This allows novel techniques for creating, and editing graphical objects, and rotating, panning and zooming the artwork.

In the system described above that links PDAs with a whiteboard (Figure 8), a tracking device was attached to the PDA that allowed it to be used like a Toolglass widget. When the PDA is held near an item on the whiteboard and an option is selected on its screen, an attribute of the item is changed to reflect the selection. The PDA provides a click-through interface like the Toolglass.

2.8 Collaboration

The VideoDraw system [58] was constructed to give two users a natural shared space that allowed drawing with pens and gesturing with the hands. Each user had a monitor onto which they drew directly with a marker pen, and a camera that continuously sent an image of the monitor screen, pen marks and hands to the other user's monitor. A sense of spatial awareness was evoked in the users. An attempt to combine the shared task space of VideoDraw with a shared person space was made with ClearBoard [59], which used the metaphor of a glass drawing surface between two people. ClearBoard-1 used video hardware and half-mirrors to provide each user with a surface to write on, through which he could see the other user and the marks he had made. The image of a partner and his drawing was reversed so both users had a common drawing orientation. Eye contact between users and gaze awareness indicating focus of attention created an enhanced feeling of copresence. ClearBoard-2 used a drawing program controlled by digitiser pens so that drawings could be saved, loaded and imported from other programs, and one user could alter the marks of the other which is not possible when real pens are used. ClearBoard-2, in particular, was designed as a bridge between two technology streams: groupware and video conferencing.

The DigitalDesk described above was augmented to create the DoubleDigitalDesk, which combines the views of two physical desks so users can collaborate remotely on a task involving paper documents [3]. Each DigitalDesk samples an image of the surface of its desk, possibly including the hand of the user, and sends a scaled thresholded version to the other machine, which displays it in addition to any graphics it would do normally. Freeman implemented multiple distributed input devices for the DoubleDigitalDesk and also the Multi-Device Multi-User Multi Editor (MMM)[70], a multi-user text editor with menus and a hierarchy of windows that can be used simultaneously by multiple users, each with his own mouse. Each user has a home area that is used to gain control of a cursor. Interference among users is reduced because fine-grained sharing allows users to simultaneously manipulate a single window or string of text, with confidence that the actions of one will not hinder those of another.

Meeting support is a major focus for groupware technology. Documents play an essential role in most meetings: people bring prepared documents to meetings to discuss with a group, and make hand-written notes that may form the basis of new documents that are used in subsequent meetings. DOLPHIN [27] is a system that supports such formal and informal information and the transformations between the types, using a hypermedia document model of nodes and links, with content in the form of *scribbles*, text and images. It was also designed to facilitate interactions of the various categories in Figure 5: face-to-face and distributed collaboration, with concurrent and asynchronous activity. Large interactive displays – Xerox LiveBoards [55] – provide public workspaces, while networked computers provide private workspaces for local participants and/or access for remote participants. All the participants may be located in a single room, or some may join in via the computer on their desk. A user may work on a document outside the meeting then produce it from a private workspace at the appropriate time in the meeting. Various pen gestures are used to structure the information objects on the LiveBoard.

The Collaborage system at Xerox PARC [52] uses vision technology to sense the information on a wall. The board on which items are arranged is a persistent background information source. A number of




Figure 15 VideoDraw: Two cross-connected television and camera pairs [58]

applications have been developed to test the capabilities of the system, such as the Project Task Wall (Figure 16) where cards representing categories are placed along the top of the board, tasks are placed in columns beneath them, and people are associated with tasks. The spatial arrangement of tasks is interpreted and images of the writing on the cards are captured. The task information is made available via a web interface. A pan-tilt camera grabs images of parts of the board that have changed, and these are combined to form an image mosaic. Cards contain glyphs, similar to those used by Rekimoto, which are located using a template-matching algorithm, then identified by zooming the camera in. The board has the advantages of a physical wall, presents a tangible user interface, and acts as an input device to the digital domain with just simple audio feedback like the BrightBoard. A physical board has a physical location and generally will be used by one person at a time, so Collaborage would fit into the same-place different-time category of Figure 5.



Figure 16 Category, task and person cards on the Project Task Wall [52]



Figure 17 Insight Lab [65]: Using walls as display spaces for information (left), and scanning a note to view associated digital data (right).

The Insight Lab developed at Anderson Consulting again presents a tangible user interface via the use of paper. Its goal is to facilitate analysis of observational data by teams creating requirements documents. Pieces of *evidence* – quotes, images, video clips, etc. – and *keywords* are grouped using *patterns* and *hypotheses* that are related to the *themes* of a project. All of these items are represented as pieces of paper that are arranged on whiteboards and walls (Figure 17). The paper notes have barcodes that, when scanned, cause the associated media to be displayed on large screens. A team discusses and organises the notes, reviewing associated media to gain context, and storing any structure decided upon by scanning barcodes to group items.

The i-LAND project [54] is based on a vision of workspaces of the future that support cooperative work of dynamic teams. Integration of information and architectural spaces is envisioned.

Several *roomware* components have been developed – computer-augmented objects integrating room elements with information technology. The spread of the desktop computer, its role as the primary point of access to up-to-date information, and the fact that it has become a bottleneck due to, for instance, limited



Figure 18 The *CommChair* with integrated pen-based computer [54]

screen space and complex handling of windows, are cited as antecedents for the new methods of interaction with information spaces. The set of *roomware* components presently consists of an electronic wall (*DynaWall*), an electronic table (*InteracTable*), and mobile and networked chairs with integrated displays (*CommChairs*). The *DynaWall* is a 4.5x1.1m touch sensitive display that allows teams to display and interact with large information structures. The *InteracTable* has a horizontal 65x85cm display that supports drawing and gestures with a finger or pen, and text input via a wireless keyboard. Because up to six people can simultaneously use the *InteracTable* by standing around it, it requires new human-interaction techniques for a display with no predefined orientation. The *CommChairs* (Figure 18) have wireless network access, independent power supplies, and integrated pen-based computers which provide a private workspace and support interaction with public workspaces by, for example, allowing annotations to be added to the *DynaWall* remotely.



Figure 19 A visualisation of the iRoom with its many projected displays [62]

The Interactive Room (iRoom) at Stanford is being used to research the infrastructure and HCI issues pertinent to a dedicated computer-augmented space. The room, which is similar to that used for i-LAND, has three touch-sensitive SmartBoard displays, a bottom-projected table, a front projected (output-only) screen, a variety of wireless mice, keyboards and PDAs for interacting with the screens, and approximately eight PCs providing computing, rendering and display server capabilities [61]. The intent is to augment existing applications to use the many devices. PowerPoint has been modified so that the speaker can make use of the many screens, and the audience, using laptops and PDAs, can browse through the information presented, annotate it and post questions to be integrated into the presentation. A central feature of the iRoom is the event heap, which is like a tuple space that enables entities in the room to communicate. Multiple entities can subscribe to a given event, and events auto-expire so senders do not have to verify that a receiver is present. A display called an Interactive Mural has been developed for the iRoom [63] that combines eight overlapping 1024x768 pixel projectors with intensity blending along the seams to form a large display that uses a laser pointer as input. The purpose of the display is to investigate new interaction paradigms for large displays and environments such as the iRoom.

3 Work Done

3.1 Pen Input

The Polhemus FastTrak [8] is a 6 degree-of-freedom tracking device that uses magnetic fields to track the position and orientation of up to four receivers. The receivers have been mounted in cylindrical casings like pens, which have two buttons each. The devices thus created are referred to here as crayons (Figure 20). The tracker hardware has been attached to a computer – a PC running Windows 2000 – via the serial ports, and software has been written to configure the tracker and capture the data it produces. The buttons have been attached to the computer's parallel port via a small piece of specially made hardware (Figure 21), and software has been written to poll their state via an extra kernel-mode driver for Windows. The software was written in C++.

Initially a program was written to allow the crayons to be used to 'write on the wall'. A projector was pointed at the wall to display a drawing surface, and a system to calibrate the crayon with the projected image was created using the four point warping method used by the DigitalDesk [3 p66-69]. A point on the drawing surface can be calculated by creating a perpendicular line to the surface from the crayon tip, or by intersecting a ray from the crayon with the surface (Figure 22). The latter method was deemed easier to use as it allows both pointing from a distance and touching the surface directly, although it is sensitive to errors in the angle as well as the position data from the tracker.

The four point warping method transforms a point (x,y) on the wall to a point (x',y') on the projected image using the formulae below:

$$\begin{aligned}x' &= c_1x + c_2y + c_3xy + c_4 \\y' &= c_5x + c_6y + c_7xy + c_8\end{aligned}$$

This warping method compensates for translation, scaling, shearing and keystoneing of the projected image – the major transformations that will occur due to the relative positioning of the projector and the surface. To calibrate the system, four points are projected in turn onto the drawing surface, and the crayon is used to select each one. The four pairs of points from the projector and the tracker are substituted into the warping equations, and Gaussian elimination is used to obtain the constants c_1 to c_8 .

3.2 Projected Display

Code has been developed using the DirectDraw API on Windows 2000 to access multiple display devices and display double-buffered images on them. A Java interface to the crayon software was created and a Java program was written that allows multiple crayons to be used to write on the wall and move graphical objects around using direct manipulation. Java objects are used to store the data for the graphical objects, so that the safety and



Figure 20 The magnetic transmitter (left) produces a field that is detected by the receivers in the crayons (right)



Figure 21 The buttons on the crayons are connected to the computer's serial port via a special piece of hardware




Figure 22 To get a point on the surface we can use a perpendicular (A), or a ray (B)

functionality of Java is available on top of the hardware access and speed provided by the C++ routines. The buttons on the crayons trigger various drawing functions.

The Matrox G400 dual-head graphics card and DirectDraw code currently in use will be used to drive two projectors from a single card when the *foveal* display described below is constructed. DirectInput has also been used because it will be useful for obtaining input from keyboards, mice and other devices

3.3 VNC

Virtual Network Computing (VNC)[4] is a remote display system developed by AT&T Research Laboratories Cambridge. It allows a computer 'desktop' environment to be viewed from anywhere on the Internet and from a wide variety of machine architectures (Figure 23). The VNC server program resides on the computer that would normally display its output directly on a screen. In the case of a UNIX machine the VNC server acts as an X server. The VNC client runs on the computer with the screen, displaying the desktop of the remote computer for the user to see. A *remote framebuffer protocol* allows the server to update the client's framebuffer. The input side of the protocol is based on a standard workstation model of a keyboard and a multi-button pointing device. Input events are sent to the server by the client whenever the user presses a key or moves the pointing device.

VNC, whose source code is made available by AT&T, was adapted to make the projected display into a VNC client that is controlled using a crayon. Figure 25 depicts the main parts of the system. The projected display and crayon input subsystems were created as described above, and a standard VNC client was modified to reroute display updates to the projector and synthesise mouse events from the crayon data. An unmodified VNC server allows the display to be used as the X Windows display of a UNIX machine. Figure 24 shows a standard X Windows drawing program being displayed on the wall and controlled with a crayon.



Figure 24 An unmodified X Windows drawing program controlled by pointing a *crayon* at the wall




Figure 23 An X desktop being viewed from Microsoft Internet Explorer on a PC [4]




Figure 25 Elements of the VNC system with projected display. The shaded boxes are new or modified.

3.4 Hardware

A survey of input devices was conducted. The Polhemus tracker was used initially because it was available and because it offers simultaneous use of four crayons, with position and orientation data – the crayons can act like both the pens of a digitiser tablet and like pointing devices such as laser pointers. It is desirable for the pens to be cordless and it is anticipated that the full six degrees of freedom provided by the Polhemus will not be necessary and that pen tip location on a flat surface will suffice. Large format, 3x4' digitisers with cordless pens and an accuracy of 0.01" are available for around \$2000 but these only allow the use of one pen at a time. Tablets in Wacom's Intuos range support multiple pens but their maximum size is 12x18". Development can proceed using the Polhemus, but before the interaction techniques are tested more accuracy will be needed. The large movement of the non-dominant hand and the smaller, detailed movements of the dominant hand could be accommodated using a combination of a large format digitiser and a tracker such as Mimio by Virtual Ink. A cordless pen would be available for each hand and the devices pen would not interfere with each other because one is based on electromagnetism while the other uses ultrasound and infrared.

The projector currently being used has been mounted above a desk with a mirror positioned to reflect the projected image downward (Figure 26). This creates a horizontal display on which normal application programs can be operated (Figure 27). Another projector will be mounted in a similar way to the first to create a display that fills the whole of the desk.



Figure 26 Projector on a shelf with image its reflected downward




Figure 27 The projected display turns a desk into a computer desktop.

4 Research Proposal

4.1 Objectives

The aim of this research is to remove the monitor from the desktop: to make a system where a person interacts with the documents they are using, as they do with paper documents, rather than manipulating them indirectly through devices that are the focus of attention.

The areas involved are:

- Multi resolution *foveal* display
Projector arrangement, calibration, image warping, intensity blending
- Bimanual input
Communicating with devices, tool palette, bimanual techniques such as rotating objects
- Virtual paper
Generation from postscript, positioning pages on desk, piles, annotation
- Collaboration
Establishing a connection, voice channel, interaction of participants' actions, gesturing

The target application will be the University of Cambridge supervision system. Groups of undergraduates, typically containing two students, are regularly given supervisions on lecture courses by postgraduate students and members of staff. A supervision is arranged by e-mail or telephone, the supervisor sets some work in advance, the students hand in their answers, and the supervisor and students will meet for a one-hour face-to-face meeting to discuss the work and any other questions the students have about the course material. The participants in a supervision combine verbal interchange with work and learning in a small-group situation indicative of many commercial and academic congresses. The discussion, in computer science at least, is centred around the written answers, notes and diagrams made during the supervision, and the course notes, text books and other written material. The supervision system is well understood and provides a benchmark for a computer-based medium designed to accommodate it.

4.2 Foveal Display

A structure will be assembled that allows two digital projectors to cast their images down onto a desk, the two images being overlapped with the larger covering the whole desk, and the smaller covering just the area in front of the user. Using a front-projected rather than a back-projected display should not cause problems: it was noted during testing of the DigitalDesk that users were not bothered by the projection system and the shadows cast by hands and other objects [3 p26]. If the smaller display has 1024x768 pixels and the dimensions shown in Figure 28 making it roughly the size of an A3 piece of paper, and the larger display has 800x600 pixels and the dimensions in the diagram, pixels on the larger display will have 15 times the area of those on the smaller display. The double projector system will create a *foveal display* with a high detail focus accommodating the objects to which the user is attending, and a lower detail periphery holding objects that may be brought quickly to the 'fovea'. An object may be transferred to the fovea after being deemed by the user to be useful to the current task and hence located rapidly using instinctive spatial memory, or after providing a visual cue that it requires attention which urges the user to move the object to the fovea for closer inspection.

Two 'crayons' will be used, one in each of the user's hands. A mechanism for calibrating the crayons and the two displays will be developed. The four point warping method already used alters the co-ordinates provided by the tracker, to calculate the correct position on the rectangular image that is projected. It will




Figure 28 High and low resolution overlapping projected images

be more desirable, but more computationally expensive, to specify the projected images in a vector format and warp them before they are displayed to compensate for imperfect projector positioning. Necessarily perfect positioning would be undesirable in a system intended to be constructed from standard parts and operate robustly in a normal office or home environment. Given this, in the absence of advanced warping mechanisms in the projectors, at least one of the projected images will have to be distorted to ensure registration between the high and low detail sections of the display. Intensity blending may also be necessary to smooth the transition between the two projected images. This type of warping and blending has been investigated by groups at MIT, Stanford and UNC [48, 47, 46]. A precise method for calibrating the crayons and projectors will be chosen following experience with the system.

4.3 Virtual Paper

Sheets of virtual paper will be placed on the desk, and moved and marked with the pens. A new sheet will either be blank or have contents generated from a specified postscript file: rendering from postscript will enable any document in a standard format to be converted to sheets and arranged on the desk.




Figure 29 (a) Grouping pages in *piles* (b) Visualisation techniques

Various tools will be available to annotate a

sheet, and highlight information already on it. Methods for grouping sheets related to the *pile metaphor*, described above, will be explored (Figure 29a). The appearance of the groups and interaction methods with the two hands are important. The contents of a pile will be browsed using one of the techniques described in section 2.6. Visualisation techniques for fitting more sheets onto the desk will be tested, such as scaling down pages outside the fovea (Figure 29b), or scaling sheets by their distance from the fovea. The pen in the dominant hand will be used to annotate the sheets by writing or highlighting.

4.4 Bimanual Input

The non-dominant hand will make long sweeping movements and provide a frame of reference for the work of the dominant hand, which will require higher precision to perform intricate movements. A tool palette will be designed (Figure 30). Attention will be paid to the details of this interface: how it is invoked, the displacement from the ND hand, the tools available and their representation in the palette. The palette will initially contain different types of drawing tool, and a tool to create new sheets. Experience with the new interface will suggest desirable functionality that will either be added as new tools in the palette, or as new techniques using one or two hands.




Figure 30 Tool palette positioned by the left hand and used by the right



Figure 31 The user reads from a sheet held by the left hand while making notes with the right hand

The virtual paper will be annotated with the D hand, translated with the ND and perhaps the D hand, possibly rotated, annotated with the D hand, placed in piles, located in and removed from piles, and manipulated in other ways (Figure 31). Fast and intuitive interaction techniques for these actions will be formulated. A balance will have to be struck between having actions initiated from the tool palette where their availability is advertised by icons, and thus the cognitive load on the user is reduced, and having them invoked via specific interaction techniques that are fast and direct, but force the user to remember their details.

4.5 Collaboration

The desk system will be duplicated or linked to a second computer that simulates the interface in an impoverished way using a conventional monitor, mouse and keyboard. Communications capability will be added, probably using Microsoft's NetMeeting 3 API that allows links carrying audio, video and data streams to be set up between computers over an IP network.

Various issues will be raised by the introduction of another user into the workspace. Simultaneous access to resources, such as when two users try to move the same sheet of paper, will have to be handled. MMM was designed with such interactions in mind [70], and the lessons learned from that project will be considered when undertaking this one. The sheets of virtual paper and any other graphical objects on the projected desktop will be able to accept inputs simultaneously from multiple devices and multiple users. They will be similar to *cooperatively controlled objects* [71, 72], which will be studied during the implementation of the single-user system so that adding another user should not involve major changes to the interaction techniques for the sheets. Each of the physically distributed collaborating users will be able to talk to the other, and see the work the other is doing, but extra, fine-grained gestures that allow subtle communication and synchronisation of activities will be investigated as it is anticipated that these will be crucial to providing a fluid collaborative experience.

The Internet currently treats all data streams equally, providing a 'best-effort' service, but much work is underway to implement quality of service (QoS) guarantees through systems such as the RSVP protocol [73]. ATM is becoming the technology of choice for high speed networks. Switches can be made larger and latency guarantees are easier to give than with Ethernet, for example. Access to the Internet is generally charged at a flat rate – a subscription fee – but research is also addressing methods of charging for resource usage or provision based on expected usage [74], especially as link capacity for subscribers and demand for multimedia applications increases. The collaborative application described in this proposal will provide fine-grained simultaneous interaction that relies on low latency communication. It is expected that this will be more cheaply supported by future networks, than the more course-grained interaction mediated by social protocol that is provided by video-conferencing applications, which requires reservation of greater link capacity.

4.6 User Study

Holding Computer Science supervisions with the completed system will provide user feedback from educated participants who know how a supervision can go right or wrong. The undergraduates will be asked to submit work in an electronic format that will be converted to postscript and represented as sheets of virtual paper on the desk. Following the supervision they will answer a set of questions about the new system and invited to make comments. The testing will result in suggestions for improvement, which will be used to hone the interface.

4.7 Timetable

Year	Quarter	Tasks
1	1	Assemble hardware for foveal display. Write DirectDraw software to drive projectors. Create Java interface that synthesises events from the tracker data. Create drawing program with multiple crayons. Write calibration routines for two-projector display. Implement intensity blending.
	2	Implement sheets of virtual paper that can be manipulated and annotated. Add routines to render to the sheets from postscript. Devise a set of useful unimanual and bimanual manipulation techniques. Implement a tool palette in the non-dominant hand, with some useful tools.
	3	Replace the crayons with a more accurate input system – possibly the large format electromagnetic digitiser and sonic tracker described above. Implement a version of the pile metaphor for grouping and browsing sheets. Implement a number of visualisation techniques to make better use of the space available on the desk.
	4	Write routines to set up an audio and data connection between two computers using the NetMeeting API. Create another large display system or simulate one with a program on a normal computer with a keyboard and mouse. Implement distributed access to the objects held by the system Evaluate and enhance the collaborative experience.
2	1	Perform user testing. Obtain feedback and make enhancements Produce a thesis plan.
	2	Survey new published results. Investigate additions such as gesture recognition and marking menus. Start writing work chapters of thesis.
	3	Perform quantitative analysis of processor and network use to characterise resource requirements of the new system. Review changes in display and input technology. Based on trends in capability and price of technology, define areas for future work. Start writing remaining chapters of thesis.
	4	Complete and submit thesis.

References

1. Sellen A., Harper R., Paper as an analytic resource for the design of new technologies, *Proceedings of CHI '97*, pp 319-326, ACM, 1997.
2. O'Hara K., Sellen A., A comparison of reading paper and on-line documents, *Proceedings of CHI '97*, pp 335-342, ACM, 1997.
3. Wellner P. D., Interacting with Paper on the DigitalDesk, PhD Thesis, *University of Cambridge Computer Laboratory Technical Report 330*, 1994
4. VNC: Virtual Network Computing, AT&T Laboratories Cambridge, <http://www.uk.research.att.com/vnc/>
5. Hinckley K., Pausch R., Proffitt D., Kassell N. F., Two-Handed Virtual Manipulation, *ACM Transactions on Human-Computer Interaction* 5(3), pp260-302, September 1998
6. Henderson D.A. Jr., Card K., Rooms: The Use of Multiple Virtual Workspaces to Reduce Space Contention in a Window-Based Graphical User Interface, *ACM Transactions on Graphics Vol. 5 No. 3*, pp 211-243, July 1986
7. Bolt R. A., *The Human Interface*, Lifetime Learning Publications, Belmont, California, 1984.
8. Polhemus, <http://www.polhemus.com/>
9. Miller, G. A. Psychology and Information, *American Documentation* 19(3), pp 286-289, July 1968
10. Weiser M., The Computer for the 21st Century, *Scientific American* Sept 1991, pp 94-104
11. Smith D.C, Irby C., Kimball R., Verplank B., Designing the Star User Interface, *Byte Magazine April 1982*. Reprinted in Preece J., Keller L., *Human-Computer Interaction*, pp237-259, Prentice Hall, 1990
12. Johnson J. et al, The Xerox Star: A Retrospective, *IEEE Computer* 22(9), Sept 1989, pp11-29. Reprinted in Baecker R. M. et al, *Human-Computer Interaction: Toward the Year 2000*, pp53-70, Morgan Kaufmann, San Francisco CA, 1995
13. Microsoft Task Gallery, <http://www.research.microsoft.com/ui/TaskGallery>
14. Robertson G. et al., The Task Gallery: A 3D Window Manager, *Proceedings of CHI 2000*, pp494-501, ACM, 2000
15. Robertson G. et al., Data Mountain: Using Spatial Memory for Document Management, *Proceedings of CHI 2000*, pp153-162, ACM, 2000
16. Mynatt E.D. et al, Flatland: New Dimensions in Office Whiteboards, *Proceedings of CHI '99*, pp346-353, ACM, 1999
17. Ishii H., Ullmer B., Tangible Bits: Towards seamless interfaces between People Bits and Atoms *Proceedings of CHI '97*, pp234-241, ACM, 1997
18. Ullmer B., Ishii H., The metaDESK: Models and Prototypes for Tangible User Interfaces, *Proceedings of UIST '97*, pp223-232, ACM, 1997
19. Bederson B.B., Hollan J. D., Pad++: A Zooming Graphical Interface for Exploring Alternative Interface Physics, *Proceedings of UIST '94*, pp17-26, ACM, 1994
20. Bolt R. A., "Put-That-There": Voice and Gesture as the Graphics Interface, *Computer Graphics* 14(3) (Proceedings of SIGGRAPH '80), pp262-270, ACM, 1980
21. Watson B., Walker N., Hodges L.F., Worden A., Managing Level of Detail Through Peripheral Degredation: Effects on Search Performance with a Head-Mounted Display, *ACM Transactions on Computer-Human Interaction* 4(4), pp 323-346, 1997
22. Bolt R.A., Gaze-Orchestrated Dynamic Windows, *Computer Graphics* 15(3) pp109-119, August 1981
23. Keahey A.T., The Generalized Detail-In-Context Problem, *Proceedings of IEEE Visualization '98*, 1998
24. Robertson G.G., Mackinlay J.D., The Document Lens, *Proceedings of UIST '93*, pp101-108, ACM, 1993
25. GreenBerg S., A Fisheye Text Editor for Relaxed-WYSIWYG Groupware, *Proceedings of CHI '96 conference companion*, pp212-213, ACM, 1996
26. Yoder L.A., The State of the Art in Projection Display: An Introduction to the Digital Light Processing Technology, Texas Instruments, Available at <http://www.ti.com/dlp/resources/whitepapers/>
27. Streitz N.A. et al., DOLPHIN: Integrated Meeting Support across Local and Remote Desktop Environments and LiveBoards, *Proceedings of CSCW '94*, pp345-358, ACM, 1994
28. Buxton W., Two-Handed Document Navigation, *Xerox Disclosure Journal* 19(2), pp103-108, March/April 1994, Available at <http://www.dgp.toronto.edu/people/BillBuxton/index.html>
29. Buxton W., Myers B., A Study in Two-Handed Input, *Proceedings of CHI '86*, pp321-326, ACM, 1986
30. Rekimoto J., A Multiple Device Approach for Supporting Whiteboard-Based Interactions, *Proceedings of CHI '98*, pp344-351, ACM, 1998
31. Rekimoto J., Saitoh M., Augmented Surfaces: A Spatially Continuous Work Space for Hybrid Computing Environments, *Proceedings of CHI '99*, pp378-385, ACM, 1999
32. Rekimoto J., Pick-and-Drop: A Direct Manipulation Technique for Multiple Computer Environments, *Proceedings of UIST '97*, pp31-39, ACM, 1997
33. Ayatsuka Y., Matsushita N., Rekimoto J., HyperPalette: A Hybrid Computing Environment for Small Computing Devices, *Proceedings of CHI 2000 Extended Abstracts*, pp133-134, ACM, 2000

34. Stafford-Fraser Q., Robinson P., BrightBoard: A Video-Augmented Environment, *Proceedings of CHI'96*, pp134-141, ACM, 1996
35. Stafford-Fraser Q., Video-Augmented Environments, PhD Thesis, *University of Cambridge Computer Laboratory Technical Report 419*, 1996
36. Moran T.P., Chio P., van Melle W., Pen-Based Interaction techniques for Organizing Material on an Electronic Whiteboard, *Proceedings of UIST '97*, pp45-54, ACM, 1997
37. Nakagawa M. et al., Interactive Dynamic Whiteboard for Educational Applications, *Proceedings of International Conference on Virtual Systems and Multimedia (VSMM '96)*, pp479-484, IEEE, 1996
38. Bier E.A., Buxton W., et al., Toolglass and Magic Lenses: The See-Through Interface, *Proceedings of SIGGRAPH '93*, pp73-80, ACM, 1993
39. Kurtenbach G., Fitzmaurice G., Baudel T., Buxton W., The Design and Evaluation of a GUI Paradigm based on Tablets, Two-hands and Transparency, *Proceedings of CHI '97*, pp35-42, ACM, 1997
40. Leganchuk A., Zhai S., Buxton W., Manual and Cognitive Benefits of Two-Handed Input: An Experimental Study, *Transactions on Human-Computer Interaction 5(4)*, pp326-359, ACM, 1998
41. PowerWall, <http://www.lcse.umn.edu/research/powerwall/powerwall.html>
42. Cruz-Neira C., Sandin D.J., DeFanti T.A., Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE, *Proceedings of SIGGRAPH '93*, pp135-142, ACM, 1993
43. Humphreys G., Hanrahan P., A Distributed Graphics System for Large Tiled Displays, *Proceedings of Visualization '99*, IEEE, 1999
44. Fuchs H. et al., The Office of the Future: A Unified Approach to Image-Based Modeling and Spatially Immersive Displays, *Proceedings of SIGGRAPH '98*, pp179-188, ACM, 1998
45. Raskar R., Welch G., Fuchs H., Seamless Projection Overlaps using Image Warping and Intensity Blending, *Proceedings of Fourth International Conference on Virtual Systems and Multimedia (VSMM 98)*, IEEE, 1998
46. Raskar R. et al., Multi-Projector Displays using Camera-Based Registration, *Proceedings of Visualization '99*, IEEE, 1999
47. Humphreys G., Hanrahan P., A Distributed Graphics System for Large Tiled Displays, *Proceedings of Visualization '99*, IEEE, 1999
48. Surati R.J., Scalable Self-Calibrating Display Technology for Seamless Large-Scale Displays, *PhD Thesis, Massachusetts Institute of Technology*, 1999
49. Buxton W.A.S., Telepresence: Integrating Shared Task and Person Spaces, *Proceedings of Graphics Interface '92*, pp123-129, 1992
50. Väänänen K., Hübner W., ShareME – Shared Multimedia Environments: Some Issues on Interaction in Distributed Multimedia Information Environments, *Proceedings of the HCI'92 Conference on People and Computers VII*, pp369-388, BCS, 1992
51. Krueger M.W., Videoplacement, in *Artificial Reality*, Addison Wesley, 1983
52. Moran T.P., Saund E., van Melle W. et al., Design and technology for Collaborage: Collaborative Collages of Information on Physical Walls, *Proceedings of UIST '99*, pp197-206, ACM, 1999
53. Ellis C.A., Gibbs S.J., Rein G.L., Groupware: Some Issues and Experiences, *Communications of the ACM*, 34(1), 1991
54. Streit N. et al., i-LAND: An Interactive Landscape for Creativity and Innovation, *Proceedings of CHI '99*, pp120-127, ACM, 1999
55. Elrod et al., LiveBoard: A Large Interactive Display Supporting Group Meetings, Presentations and Remote Collaboration, *Proceedings of CHI '92*, pp599-607, ACM, 1992
56. Pedersen E., McCall K., Moran T.P., Halasz F., Tivoli: An electronic Whiteboard for Informal Workgroup Meetings, *Proceedings of InterCHI '93*, pp391-398, ACM, 1993
57. Agrawala M. et al., The Two-User Responsive Workbench: Support for Collaboration Through Individual Views of a Shared Space, *Proceedings of SIGGRAPH '97*, ACM, 1997
58. Tang J.C., Minneman S.L., VideoDraw: A Video Interface for Collaborative Drawing, *ACM Transactions on Information Systems 9(2)*, pp170-184, April 1991
59. Ishii H., Kobayashi M., Grudin J., Integration of Inter-Personal Space and Shared Workspace: ClearBoard Design and Experiments, *Proceedings of CSCW '92*, pp33-42, ACM, 1992
60. Greenberg S., Hayne S., Rada R., *Groupware for Real-Time Drawing: A Designer's Guide*, McGraw-Hill, London, 1995
61. Fox A., Johanson B., Hanrahan P., Winograd T., Integrating Information Appliances into an Interactive Workspace, *Computer Graphics & Applications*, IEEE, May/June 2000
62. Stanford Interactive Workspaces Project, <http://graphics.stanford.edu/projects/iwork/>
63. Humphreys G., Hanrahan P., A Distributed Graphics System for Large Tiled Displays, *Proceedings of Visualization '99*, IEEE, 1999
64. Elliott A., Hearst M.A., How Large Should a Digital Desk Be? Qualitative Results of a Comparative Study, *Proceedings of CHI 2000*, pp165-166, ACM, 2000
65. Lange B.M., Jones M.A., Meyers J.L., Insight Lab: An Immersive Team Environment Linking Paper, Displays and Data, *Proceedings of CHI '98*, pp550-557, ACM, 1998
66. 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), pp99-112, January 1983
67. Pemberton S., Views and Feelings: Metaphorically Speaking, SIGCHI Bulletin 27(4), ACM, October 1995
 68. Mander R., Salomon G., Wong Y., A 'Pile' Metaphor for Supporting Casual Organization of Information, *Proceedings of CHI '92*, pp627-634, ACM, 1992
 69. Harrison B.L., E-Books and the Future of Reading, *Computer Graphics & Applications May/June 2000*, IEEE, pp32-39
 70. Freeman S., An Architecture for Distributed User Interfaces, PhD Thesis, *University of Cambridge Computer Laboratory Technical Report 342*, 1993
 71. Bricker L.J., Baker M.J., Tanimoto S.L., Support for Cooperatively Controlled Objects in Multimedia Applications, *Proceedings of CHI '97*, pp313-314, ACM, 1997
 72. Bricker L.J., Baker M.J., Fujioka E., Tanimoto S.L., Colt: A System for Developing Software that Supports Synchronous Collaborative Activities, *University of Washington Technical Report UW-CSE-98-09-03*, September 1998
 73. RSVP, <http://www.isi.edu/div7/rsvp/rsvp.html> and <http://www.ietf.org/html.charters/rsvp-charter.html>
 74. *Journal of Electronic Publishing Volume 2 Issue 1*, Issue on Internet Economics, University of Michigan Press, May 1996, Available at <http://www.press.umich.edu/jep/>