Rethinking the Book

David L Small

B.S., Cognitive Science, MIT (1987)
S.M., Visual Studies, MIT (1990)

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning,
in partial fulfillment of the requirements for the Degree of
Doctor of Philosophy, Massachusetts Institute of Technology

January 1999
Massachusetts Institute of Technology

author	David L Small	Program in Media Arts and Sciences January 8, 1999
accepted by	William J. Mitchell	Dean, School of Architecture and Planning
certified by	Stephen A. Benton	Chair, Departmental Committee on Graduate Students, Program in Media Arts and Sciences

Rethinking the Book

David L Small

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning, on January 8, 1999
in partial fulfillment of the requirements for the Degree of
Doctor of Philosophy.

abstract

Electronic media have lagged behind their paper progenitors in the clear, usable display of large bodies of information. New visual languages have been created for information display which exploit the computer's unique ability to render dynamic and three-dimensional typography. These languages demonstrate that the use of three dimensional form, expressive movement, visual focus and layering, in harmony with human perceptual abilities, improve navigation and contextual understanding of complex written documents. This thesis shows that graphic displays can be combined with physical interfaces to create interactions with purely typographic information that are rich, tactile and humane.

Thesis Advisor: William J. Mitchell
Dean, School of Architecture and Planning, MIT

Doctoral Dissertation Committee

committee chair	William J. Mitchell	Dean, School of Architecture and Planning MIT
member	John Maeda	Assistant Professor of Design and Computation Sony Corporation Career Development Professor of Media Arts and Sciences MIT, Media Arts and Sciences Program
member	Suguru Ishizaki	Assistant Professor of Communication Design School of Design, Carnegie Mellon University

acknowledgments

Michael McKenna, my friend and partner, for his love and support.

Howard and Phyllis Small, for their encouragement and patience.

Muriel Cooper, for setting me on this path and keeping me from straying off it.

Bill Mitchell, for his guidance and friendship.

John Maeda, who always looked out for me.

Suguru Ishizaki, for leading the way.

Chris Pullman and Paola Antonelli, for their criticism.

Yin Yin Wong, who can accomplish more in three weeks than most do in a year.

Tom White, for being a scholar and a gentleman.

Hiroshi Ishii, for his interest and advice.

Whitman Richards, for explaining how to interpret data.

All my fellow students from the Visible Language Workshop. I'm the last one, finally. I'll turn off the lights and lock the door on my way out.

Elise Co, Peter Cho, Bill Keays, Rich DeVaul, Ben Fry, Golan Levin, Reed Kram, Matt Grenby and Chloe Chao for making ACG the place to be.

Jared Shiffman, without whom aménagér would have been spelled amenager.

Josh Smith, for his generosity and skill in hand sensing.

table of contents

CHAPTER 1 *introduction*

Electronic media have lagged behind their paper progenitors in the clear, usable display of large bodies of information. New visual languages have been created for information display which exploit the computer's unique ability to render dynamic and three-dimensional typography. These languages demonstrate that the use of three dimensional form, expressive movement, visual focus and layering, in harmony with human perceptual abilities, improve navigation and contextual understanding of complex written documents. This thesis shows that graphic displays can be combined with physical interfaces to create interactions with purely typographic information that are rich, tactile and humane.

There have been two primary forms which books have taken through history - the scroll and the codex book. In recent years, with the advancement of computer technology, new electronic forms have begun to emerge as tools for managing large corpora of typographic information. To date, most of these attempts have neither achieved the formal refinement of printed books nor have they fully exploited the capabilities of modern computers. Through a series of design studies this thesis tackles such basic requirements as form, navigation, movement, layering, scale and interface. Detailed working prototypes were used to provoke a critical examination of these issues. The goal is to create a compelling vision of how computers can redefine the paradigm of the book. It is not simply about producing a series of experiments or design sketches, but to use them to ask fundamental questions about how reading, writing and expression are changing and about how written language will evolve in response to computer technology. The result is a personal vision of the future of computer mediated typography. The approach taken is perhaps not typical of the traditional thesis, but necessary in order instantiate these ideas so that they can be presented and experienced.

Although the structure of this document divides the discourse into subjects such as

three dimensional form

layering, juxtaposition and scale

expressive movement

tangible issues

there are no such divisions in the work itself. It is crucial that the designs produced address these issues concurrently. After all, design is about the synthesis of interrelated and competing constraints. The title of the thesis is rethinking the book. Instead of applying the scientific method, the approach is to provoke myself and my audience into a fundamental restructuring of ideas about how we read and our relationship to text. The vision is one of an organic, reactive construct which supports both amorphous, liquid typographic forms and more rational, architectural mechanisms.

In the field of architecture and design, progress us usually measured in the form of a critique. The presentation of work and intense discussion with experts provide both specific suggestions for improving the design as well as lay the intellectual groundwork for the theory of the design. I had the good fortune in the preparation of this work to have the council of my committee, William Mitchell, John Maeda and Suguru Ishizaki. They were joined in several critique sessions by Christopher Pullman, head of design for WGBH (Boston's public television station), and Paola Antonelli, curator in the design department of the Museum of Modern Art in New York. The critique sessions were documented and concerns and suggestions from these sessions were incorporated into the design. Excerpts from the critique will appear periodically throughout the thesis in the form of sidebars.

CHAPTER 2

some history

Since the invention of writing systems, there have been a variety of methods used to gather together and present written material. While the complete history of scrolls and the codex book are beyond the scope of this thesis, this chapter will cover some of the highlights of their evolution, in particular focusing on those inventions which help orient the reader. These will be contrasted with various experiments in reformulating the book in light of the computer.

FIGURE 1. Why books?

They are better than tying papers together with string.

Before the advent of the codex book, scrolls were one of the basic forms used to gather together and protect individual leaves of paper. Rolling the scroll and placing it into a tube or other container protected the contents and allowed their owners to handle entire books as single objects. Once unrolled, the navigation scheme was straightforward. The reader could unroll with one hand, advance one column at a time, and roll the excess material up with the other hand. A simple experiment to contrast this form with an electronic analog proved to be quite illuminating and illustrates the basic methods followed throughout this thesis.

The Scroll of Frolicking Animals

Consider a long continuous painting mounted on a scroll. For my example, I chose "The Scroll of Frolicking Animals", a Japanese painting which uses animals to satirize court society [TOBA12]. You can lay the scroll on the floor, rolling it open on one side and picking up slack on the other. Your own body and the physical qualities of the paper and cloth medium determine how the information is designed and how you find your way through it. In this case you only see portions of the painting at any one time and the artist has depicted scenes which occur at different times as one connected landscape.

In designing digital analogs to this experience, we must examine the boundary between the real world and the virtual world created by the computer. Any display creates an internal visible logic in the way the user controls the movement of the information in the space within its surface. By connecting aspects of the virtual world to real world objects, we allow the user to literally feel his way through the computer generated world. The objects which inhabit our working spaces should be legible in their function, provide clear feedback to the user, and be flexible in their application.

In traditional computer interfaces the graphics display is separate from the input devices. The user manipulates objects out on the desk and views the resulting changes on a display surface which has no clear relation to the work space itself. This means that there is always a feeling of disconnect between what the hand does and what the eye sees. Although head mounted displays or immersive systems provide a well integrated experience, the user cannot perform tasks in the virtual space while engaging in activities in the

real world. My approach is similar to systems such as Fitzmaurice [FITZMAURICE93], in which a palmtop display reveals virtual information associated with real objects in space. His system used 3D location to bring relevant information to the display. The scroll experiments use location in only one dimension to let people browse a large information space.

The Scroll of Frolicking Animals, which dates back to 12th century Japan, is a narrative painting in the form of a scroll. Scrolls have the property of being both continuous and linear and so impose interesting constraints on the manner in which information is revealed. Two different designs were considered.

FIGURE 2. The hand-held scroll below is controlled by two cylinders.

In the first, the display was placed on a platform between two cylinders which rotated in tandem (Figure 2). To use this virtual scroll, the user holds the display in both hands, rotating the cylinders with her thumbs. The image of the scroll on the display moves with the turning of the cylinders. The virtual scroll seems to unroll from one side and roll up on the other. The user has fine control over the movement of the painting, and like the actual scroll, must work through the entire painting in sequence.

FIGURE 3. A display can only show a portion on the painting at any instant. The virtual frame can slide back and forth along the painting to reveal different sections

In the second design, the cylinders are removed and wheels are added to the display platform. The assembly is placed on the floor and rolled back and forth. To move the painting forward one foot, the display has to roll one foot. The image appears to be held by friction to the floor as the display slides back and forth above it. This has the benefit of impressing on the user, in a physical way, the unusual length of the painting. To view the entire scroll, one has to roll the display eight feet along the floor. This design was very engaging for the user and generated much useful feedback.

People are used to holding books, newspapers and even scrolls in their hands and moving them about. When the user's hands and body are engaged in an appropriately physical manner, we can bring the user closer to the virtual world contained in the computer. The virtual information on the display can appear to be held directly in the hands of the user. Physical objects can become powerful interface objects if they are legible in their purpose and use. Our tendency to invest meaning into simple objects allows them to be used as symbolic tokens for the manipulation of abstract information. And, by combining the interface with the display we can give meaningful form to the space which resides on the other side of the screen and its relationship to the body.

While this has been a rather long digression for an introductory chapter, I hope that I have shown, in a small way, the basic methodology used throughout this thesis. In trying to duplicate in an electronic form the qualities of traditional media, we can come to a better understanding of both the traditional media and how they function in the human context. We can also generate new methods of designing interactive, electronic analogs which, while they don't always attempt to exactly copy the old media, can still provide the same underlying affordances.

We can also see from this example some of the limitations that the scroll must have suffered. Random access to any particular section of the scroll could be difficult and unless one had the space to unroll the entire scroll it was difficult to get an overview of the entire contents. The codex book was a major advancement in random access over the scroll form. The reader could nearly instantly access any page or could browse pages by flipping them with the thumb. The hard cover protected the delicate pages and made the whole assembly easy to carry around.

Within this basic format, many kinds of books of various complexity have been produced. Perhaps one of the richest examples is the Talmud. Multiple threads of text and commentary share space on each page and can be read as separate streams from page to page or as complex pages of multiple voices. Each part of the text has a different typographic treatment, so that it is possible to easily read the construction of the page. The book form, along with magazines and newspapers, has proved so practical that there has been a tremendous amount of inertia in moving to electronic media. While this can in part be explained by the poor reproduction quality of the CRT in comparison to paper, it must also be recognized that there is a complexity barrier that must be surmounted. The interface to a book is taught early and is easily learned.

FIGURE 4. A page from the Talmud [STEINSALTZ89]

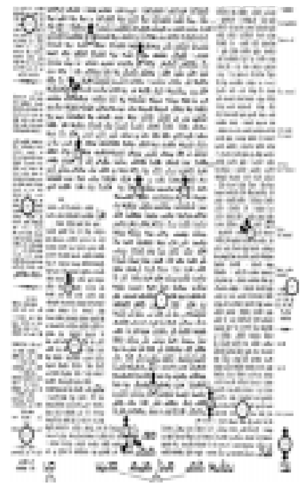

With the invention of the Gutenberg press in mid-Fifteenth century, the worldwide production of books began to increase dramatically. The manuscript book, which its graceful hand lettering, rich illumination and elaborate bindings had been raised to a fine art. The earliest printed books attempted to be facsimiles of these manuscripts, and it took nearly a century for them to become a distinct art.

Early Books of the incanabula

FIGURE 5. Gutenberg. Luke 10-11.

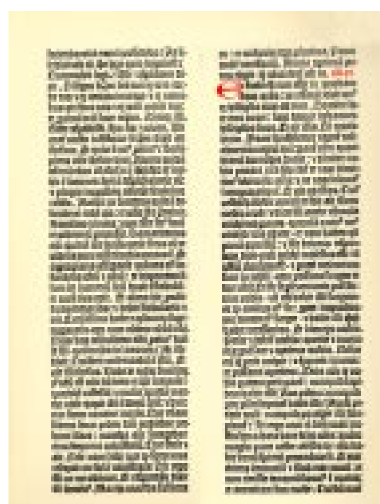

Aldus Manutius was one of the first great book designers and was responsible for several early innovations [ORCUTT28]. A well educated scholar, he settled in Venice in order to pursue his dream of publishing fine books of the Greek and Roman classics. He was the first to develop greek typefaces and he developed a set of SMALL CAPITALS to work with the Roman face he had adapted from Jenson's. He took the inclined, cursive handwriting of Petrarch, translated it into metal type and called the face *Italic*. He used his new small capitals to make running heads in a few of his volumes, so the reader didn't have to flip the book over to check what the title was. This innovation was later adopted by French publishers, who also added numbers to each page of the book.

While Venice was an important early center of bookmaking activity, in the sixteenth century social upheavals there enabled the French bookmakers, encouraged by Francois I, to surpass them in quality. Bindings were simplified so that books could stand side by side on a shelf and the printer Jean Golier was the first to us a stamped title on the shelf-back of a book. Tables of contents, title pages and other innovations appeared. In each case, they moved away from the imitation of manuscript texts, and improved the handling and usage of increasingly large collections of books by ordinary people.

Gill, Tschichold and the modern book

Moving ahead to the industrial age, the mechanization of printing began to have a strong influence. Eric Gill (1882-1940) set out his philosophy of typography and the industrial age in a beautiful, slim volume entitled "Essay on Typography" [GILL36]. He writes at the end of the industrial revolution at a time when machines have firmly taken over most of the tasks which had previously been hand done by skilled craftsman. His concern is not so much for the loss of the hand craft, as he appreciates the near-perfection of machine made papers and machine printed pages, but for the fact that the new worker has no chance to excel at his craft and is merely a tender of the machine.

In the first essay, he discusses this dilemma and make the following observation:

> It is no longer permissible to design things with no reference but to our own pleasure, leaving it to engineers to design things capable of making them; our business is now to design things which are suitable for machines to make. And this is not to say that we accept the limitations of machines as they are to-day, but that we accept the limitations of machinery as such. Moreover, and this is even more important, we are not saying that the machine is the arbiter in design: the mind is always that. The shape of A cannot be changed at the bidding of any machine that is or could be made. But, taking the shape of A to be that which the judgement of the mind lays down, we have to conform it to the nature of the machine, and not attempt to impose upon mechanical production either those ornamental exuberances which are natural and proper enough to human beings working with their hands or those peculiarities of detail which are proper to the pen, the chisel, and the graver.

We can see how this works in the discussion of book sizes and the determination of margins. He talks about four basic categories of book sizes: books to be held in the hand (a novel), or at a table (books of reference with large illustrations), or at a desk or lectern (a missal), or kept in the pocket (a travellers' dictionary). Book sizes used to be related to the nearly standard sizes of printing papers. Large oblong sheets when folded in half made a narrow folio, in half again for a wide quarto, and again to make a narrow octavo, etc. Machine made papers on continuous rolls freed the designer to consider any book shape that would sell, but the use of the book (hand, pocket, table) still guides the designer to a size appropriate for the task.

FIGURE 6. A double page spread from An Essay on Typography.

In the discussion of margins, Gill lays out a rationale for the layout of an octavo page:

> With a normal octavo page of 5 inches wide and $7\frac{1}{2}$ inches high, & supposing that we allow margins as follows : inner, $\frac{1}{2}$ inch; top, $\frac{3}{4}$; outer 1; & bottom, $1\frac{1}{6}$; we shall get a type page 3 inches wide by $5\frac{2}{3}$ inches high (i.e. 34 lines of pica type, 12 pt., set solid). This allows for a line of an average length of 10-12 words in pica, & pica is a good ordinary size for a book held in the hand.

The wide margins on the bottom and outer edges gives the reader a comfortable place for the hand to hold the book open without obscuring and text. The two inner margins seen together will be in good proportion to each outer margin. These proportions also result in a comfortable number of words per line - much longer involves more and longer eye movements and is less comfortable to read; any shorter and the lines will be difficult to justify. So we see that despite the freedom that technology had given designers in determining book sizes and layout, those constraints which were ultimately humanistic and not technological were retained. This is an important lesson for designers of new media. Screen sizes vary greatly, and yet we are still well served to have line lengths of 10-12 words and proportions which minimize the distance of eye travel we expect from the reader.

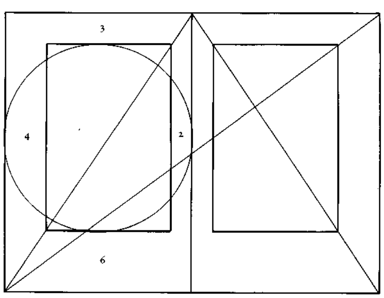

FIGURE 7. The secret cannon used by many late medieval manuscripts and incanabula derived by Jan Tschichold in 1953.

Jan Tschichold (1902-1974) is another good source when looking at the design of books. His text, The Form of the Book [TSCHICHOLD91], presents a series of essays on topics from title pages, to the use of quotation marks to the criteria for selecting paper. In one chapter he has a particularly interesting discussion of the ratio of the page and the proper placement of the text. The law, or canon, for determining the form for a 2:3 ratio page results in margins of the proportion 2:3:4:6. The height of the text is equal to the page width. This canon can also be derived by dividing the diagonal of the page into 9ths. He was intent on deriving mathematical proportions in design and felt that these ratios (e.g. 2:3, 1:$\sqrt{3}$, etc.) would have a visual harmony missing from accidental ratios (such as $8\frac{1}{2}$ x 11 paper). Unfortunately, the programming standards used in computer graphics have made these types of proportional canons difficult to follow. The designer usually has no control over even the proportions of the web browsing window, let alone control over placement of the type within that frame.

Critique
Here are some comments by Chris Pullman about the structure of the book which was missing in the virtual environment:

Chris Pullman
Is there the equivalent of this (holding a closed book) in this [virtual] world - a book with its cover closed? How do you depict the not in use state? Do you come to it out of the box? Open? There is a long tradition of what the sequence is from the cover to the first page...there is a lot of other stuff that has to be accounted for. I see this as an environment for reading into which you bring other texts.

Tschichold also discusses at length his view on the proper use of typography. Good typography is the result of long experience and careful study of the forms which have been in use for generations. “A good example of print must be of noble design and be pleasing to the eye. Beyond that, it should not attract particular attention.” Unlike other arts, typography achieves perfection when it recedes into invisibility. It must not call attention to itself, because in doing so it must fight against the words it must convey. Any unusual variation in the shape of a letter or in the spacing of letters and words will hurt legibility and make reading difficult. This is not the only definition of good typography, but it is the one used in this thesis.

The Memex

Vannevar Bush was perhaps one of the first to propose a totally new way of storing pages of information. In his essay, “As We May Think”, he describes a machine which he termed the Memex [BUSH46]. His proposal can at last be realized with current technology. Bush points out that we will not develop new ways of storing and retrieving information simply because it will be an improvement over current methods, but because the vast increase in stored human knowledge will make it a necessity.

> It consists of a desk, and while it can presumably be operated from a distance, it is primarily the piece of furniture at which he works. On the top are slanting translucent screens, on which material can be projected for convenient reading. There is a keyboard, and sets of buttons and levers. Otherwise it looks like an ordinary desk. In one end is the stored material. The matter of bulk is well taken care of by improved microfilm. Only a small part of the interior of the memex is devoted to storage, the rest to mechanism. Yet if the user inserted 5000 pages of material a day it would take him hundreds of years to fill the repository, so he can be profligate and enter material freely... There is, of course, provision for consultation of the record by the usual scheme of

> indexing. If the user wishes to consult a certain book, he taps its code on the keyboard, and the title page of the book promptly appears before him, projected onto one of his viewing positions. Frequently-used codes are mnemonic, so that he seldom consults his code book; but when he does, a single tap of a key projects it for his use. Moreover, he has supplemental levers. On deflecting one of these levers to the right he runs through the book before him, each page in turn being projected at a speed which just allows a recognizing glance at each. If he deflects it further to the right, he steps through the book 10 pages at a time; still further at 100 pages at a time. Deflection to the left gives him the same control backwards. A special button transfers him immediately to the first page of the index. Any given book of his library can thus be called up and consulted with far greater facility than if it were taken from a shelf. As he has several projection positions, he can leave one item in position while he calls up another. He can add marginal notes and comments... just as though he had the physical page before him.

FIGURE 8. Vannevar Bush

Now, discounting some of the technical specifics of Bush's description, it is easy to imagine such a system existing today with current technology. We can now place papers and photographs into "storage" and we can retrieve them and create links from one "page" to another. I think he somewhat underestimated how often one would have to look at one's "codebook" in order to find the appropriate address or command to find a piece of stored information, but we can essentially take this description as accurate and buildable. I think that there is one standard that he makes, however, which has not yet been met. He writes, "Any given book of his library can thus be called up and consulted with far greater facility than if it were taken from a shelf." Although there are many on-line books and the virtual storage capacity greatly exceeds the physical storage one could have on a bookshelf, it is still often faster to grab a book near your desk rather than to start searching on the World Wide Web. This is in part due to the sheer speed with which one can access a physical book. But, it also has something to do with the way book design enables you to quickly find and retrieve information from specific sections within the book. The goal of this thesis is to propose new designs for electronic texts which could in fact be “consulted with greater facility”

Book without pages

In 1978 Nicholas Negroponte, Richard Bolt and Muriel Cooper published a proposal to the NSF (National Science Foundation) entitled Books without Pages [NEGROPONTE78]. This proposal was concerned with *soft copy* and an attempt to bring knowledge from the fields of psychology, engineering and design together in order to retain some of the facilities of *hard copy.*

In coining the term soft copy, they were distinguishing between text as marks of ink on real paper, from the more malleable and yet insubstantial images of type on TV or digital video displays. They discussed “media fiducials”, those subtle qualities of media which provide the user information beyond the literal text.

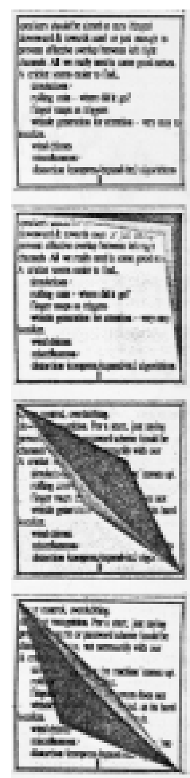

FIGURE 9. An illustration of page flipping from the SDMS.

If you read with both hands, the left hand holds what has been read and your right hold the information yet to be covered. By weight and feel, you know where you are, even if it never percolates up to the conscious level. Dog ears, bookmarks and other artifacts of the paper book were felt to be lost in the new digital media.

An example is given from the infamous Spatial Data Management System [BOLT79]. In the SDMS, rather than scroll up and down through the text, the text is segmented into pages and an animated flip from the corner shows each page lift and reveals the following page (Figure 9). They examined the whether this method would improve the subjects ability to correctly report where they were in the text (1/8th through, 1/2 through, etc.).

Two short excerpts from the proposal are particularly relevant to this thesis. In the first, they discuss the loss of the paper page and its relevance for research into new forms.

> Books without Pages is not a slogan advocating the abolition of the page, nor a celebration of the falling away of the pages in the context of text displayed on the face of the monitor, and scrolled up and down. It is both an observation that indeed the physical basis of the page as we know it is gone (replaced in only its "window"-like aspect by the rectangle of the typical monitor display face), and a call for re-examination of the opportunities and deficits that arise as a result. What do we gain in terms of release from a structural frame that (seemingly) bears no necessary relationship to the understanding and comprehensibility of the subject matter? What do we lose in terms of the basis of orientation, an extra-semantic gauge of size and/or progress? How do we reinstate the functional equivalent of any lost positive benefits inherent in the traditional page?

This insight into the basic challenge the computer was bringing to information management was apparent, even given the monochrome displays and limited computer power available in 1978. They also recognized that technology would bring necessary changes to the static typographic forms of printed typography.

> Not only can dynamic text depart from the statically usual, but it can modulate in real-time before the reader, even as a function of where he is looking (Cf. Gould, 1974). Consider a text line where any word may dynamically change in size, shape, color, luminosity. This additional dimension, call it "Z," can be an instrinsically "non-spatial" dimension (color, luminosity), or can be along spatial dimensions into or across the page, exhibiting actual movement. Why should text move or change? We see at least five reasons: to convey information that itself is changing, to pace the observer, to save "real-estate", to amplify, and to be attention getting. Each reason can be addressed in terms of kinds of change, which include (but are not limited to): 2-D/3-D translation and rotation; color changes; shape transformations; transparency; and transfigurations between icons and symbols.

These basic concerns remained part of the research of the Visible Language Workshop, founded by Cooper and Ron MacNeil, for many years. In the fifteen years that followed, computer graphics technology improved radically and those ideas that were only hinted at in the earlier research became possible.

Information Landscapes

In 1994, Muriel Cooper and her students presented their work on Information Landscapes at the TED5 conference in Monterey, California. Using a new Silicon Graphics computer, the half-hour presentation was a continuous flight through a changing three-dimensional space which explored 3D and dynamic typography in the display of hierarchical, financial and geographic information. In this demonstration professor Cooper finally realized some of the design goals she had set years earlier in the Visible Language Workshop.

FIGURE 10. Geographic information and 3D typography.

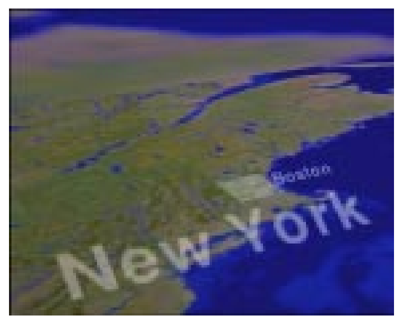

The purpose of the information landscapes research was to apply typographic techniques developed in two-dimensional graphic design to the design of three-dimensional information graphics. Previous studies had shown that three-dimensional presentation is effective in visualizing large and complex information, which continues to be increasingly available to ordinary workspaces (for example, [CARD91] [FAIRCHILD88][MACKINLAY91][ROBERTSON91][STAPES93]). However, although typography have been effectively used as a means of visually clarifying information in traditional graphic design, issues in typography had not yet been explored in the field of three-dimensional visualization.

Suguru Ishizaki and I implemented an experimental software tool in order to investigate the use of interactive three-dimensional graphics as a medium for typographic communication. Working with Muriel Cooper and other students of the Visible Language Workshop we sought to clarify issues in three-dimensional typography.

The software, named Typographic Space, enabled us to position text at any size, position and orientation in an extremely large three-dimensional space. Typographic Space allows us to change a basic set of typographic attributes in order for us to experiment with various designs. A simple interface (a mouse and keyboard) was provided for the user to travel through a three-dimensional space. The mouse is used to change view distance and eight keys are assigned to set rotation, translation of the viewpoint.

The use of typography in three-dimensional information graphics requires special care which may not be inherited from traditional typography. This study clarified issues which are important in the design of three-dimensional information graphics. Here, we will discuss four major issues.

FIGURE 11. An example of perspective distortion and its effect on legibility

First, an obvious problem in three-dimensional typography is the distortion of typographic form caused by the perspective and the arbitrary movement of a viewpoint. Three-dimensional display does not guarantee the visual quality of the original typeface design since the view angle and perspective change the form of the text on a display. Furthermore, it is possible to move a viewpoint behind a text so that the text is displayed in mirror image. A text can also become a line, or disappear when it is viewed from its side. One way to solve this problem is to display text always facing to the viewer. However, we have found that it is only effective when it is used with a small number of texts in a small area. This is because this technique can destroy the overall visual structure of the display, which is often effective to present the underlying structure of complex information.

Second, since the apparent size of text changes as the user travels through three-dimensional space, the type size may not be used as an effective visual cue. In a two-dimensional design, the type size is used as an effective cue for indicating particular kinds of information as well as navigating the reader's eyes over a page. In a three-dimensional display, small text that is located close to the viewpoint may have the same perceptual size as a text further away. However, when there is another clear visual structure in the space, the type size can be used as an effective visual cue. In our example, since orientation and color are used to indicate city names (tilted 45 degrees and transparent white), the type size is used effectively to differentiate large cities from smaller ones.

FIGURE 12. Small, or just further away? Other clues are needed to resolve the distinction.

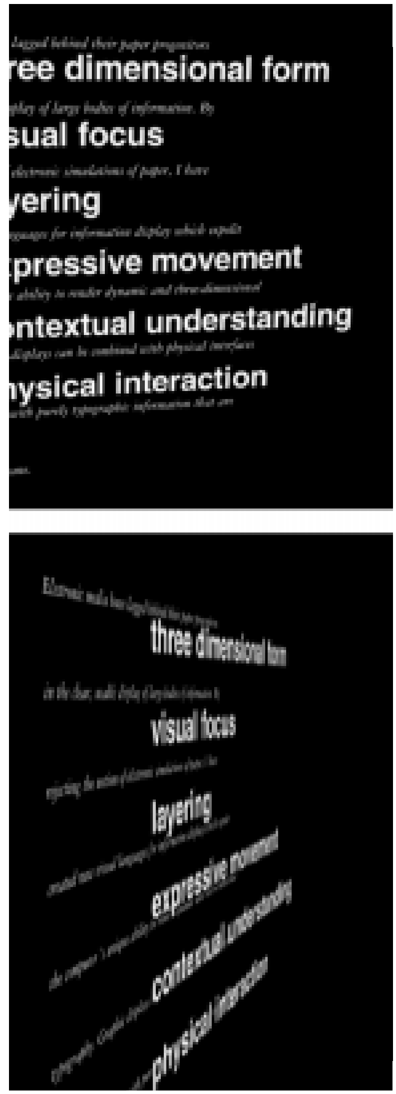

The type size also confuses the reader's sense of space, since the size of objects are one of the important cues for the perception of depth of field. In order to solve this problem we have applied the stereoscopic display technique. The stereoscopic technique clearly eliminated spatial confusion that is created by the differences of text size. Also, it enhanced the differentiation of overlapping translucent texts, which was one of the problems of multi-layering technique in two-dimensional graphics [COLBY92].

Finally, three-dimensional motion can be a new expression in typography. In two-dimensional graphics, various temporal expressions, such as movement, blinking and flashing are used as a means of conveying messages. Three-dimensional graphics increases the complexity of expressions the designer can create. Though three-dimensional motion has not been explored in terms of visual communication, as we can find the use of dynamic motion in a bee's gestural language, it opens an interesting area to explore.

We have seen how the study of traditional forms, such as the scroll, can inform the design of electronic information displays. In order to derive new insights into the design of book-like machines, we examined the development of the book form since the first printed volumes. Thought experiments, such as Bush's Memex prefigured the way in which computers would change the design of and interaction with information. The work of the Architecture Machine Group and the Visible Language Workshop culminated in a radical but unproven technique for breaking the two-dimensional conventions that the computer had adopted from print.

The main body of this thesis which follows will examine in greater detail the many issues which were uncovered in the design of Information Landscapes. The work will be organized into chapters which each cover a specific design issue, such as form, layering, navigation, movement and interface. Examples and related works will be discussed, although you will see that most of the design experiments that are presented could naturally fall into any of several chapters. We begin by discussing in more detail the issues of three-dimensional typographic form.

CHAPTER 3 *form*

It is night and you are dreaming. The sky is dark and you are floating lightly above the earth. As you cast your thoughts about, you flit and fly above the landscape from place to place and then zoom out into space. As soon as a constellation appears ahead, you skim instantly to that place, exploring, moving as fast as thought.

Navigation of information might be just as I described—smooth, simple, and as fast as your thoughts. Professor Muriel Cooper, founder of the Visible Language Workshop within the MIT Media Laboratory, coined the phrase information landscape [SMALL94] to describe this sort of space, where information "hangs" like constellations and the reader "flies" from place to place, exploring yet maintaining context while moving so that the journey itself can be as meaningful as the final destination.

FIGURE 13. 3D views of hierarchical information.

The idea that typography could break the two-dimensional plane (Figure 13) of the CRT and form three-dimensional volumes was first described by Muriel Cooper. At the TED5 conference in 1994 she talked about the issues involved. "How do you retain the integrity of the information, and at the same time, retain the context and clues that allow you to traverse complex information? You are, in a sense, in an architectural construct, but you don't have the constraints of having to believe a physical building. So you can both use the abstract conceptual issues, as well as the physical cues that people are accustomed to."

In architecture, the form of a building is usually a natural extension of the building's purpose, site and "contents". When pursuing clear information design, we must also build forms which are a natural extension of the information elements, the purpose or use to which the information will be put and the site or visual context within which the information will be seen.

Although lacking in the tactile and solid forms found in architecture, we can make use of the ephemeral qualities of screen typography to create forms which are fluid and mutable. As sections of text are read their forms can be simplified and abstracted into recognizable chunks.

Traditional graphic design has always been bound by the underlying form of the page. Within the initial constraints of the shape of paper used, the various design elements had to find their place. In designing for the screen, we still must work within a rectangular frame- visual elements must respect and coexist with the edge of the display. However, this frame is also a moveable window which can reveal and contain the larger forms of the whole corpus.

We are still feeling our way through a new discipline with few precedents and conventions. Through the careful development of a variety of design studies, some sense of the basic vocabulary of forms can be constructed as part of a visual language of three-dimensional typography.

Virtual Shakespeare

A landscape, whether real or virtual, provides an experience in which context is continuous and meaningful. It is through context that we can understand new information and can relate it to what is already known. Drs. Stephen and Rachel Kaplan wrote about the experience of mystery in landscapes in their book Cognition and Environment: Functioning in an Uncertain World [KAPLAN82].

> In the case of mystery, the new information is not present; it is only suggested or implied. Rather than being sudden, there is a strong element of continuity: the bend in the road, the brightly lighted field seen through a screen of foliage--these settings imply that new information will be continuous with, and related to, that which has gone before. Given this continuity one can usually think of several alternative hypotheses as to what one might discover.

By escaping the confines of the flat sheet of paper, we can arrange information into meaningful landscapes that exhibit qualities of mystery, continuity, and visual delight.

The plays of William Shakespeare are used to explore the design of an electronic information space that maintains the qualities of a meaningful landscape. A large image from the system is shown on the following two pages. Each character in the play A Midsummer Night's Dream is marked in a different color, and different typefaces are used to distinguish stage direction, names, dialog, and commentary. Each scene is laid out in a single column of text. These scenes are aligned at their top edge and separated horizontally by a small gutter. The gutter, or space between columns of text, is increased between acts so that each act forms a distinct

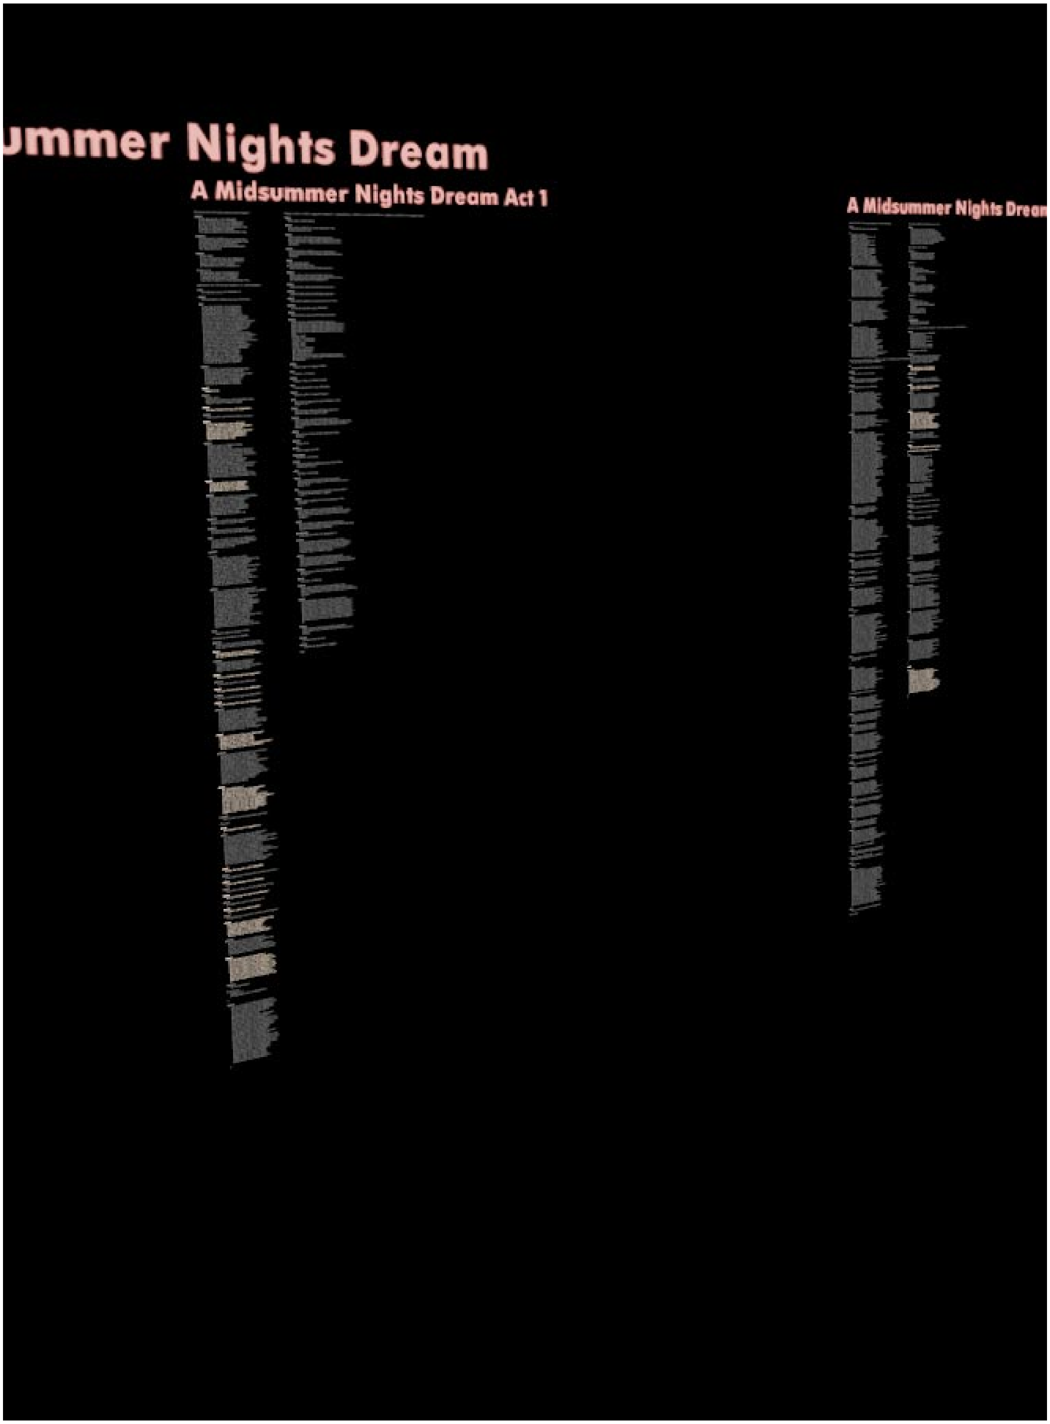

A large scale view from the Shakespeare project. This view shows the entire contents of A Midsummer Nights Dream.

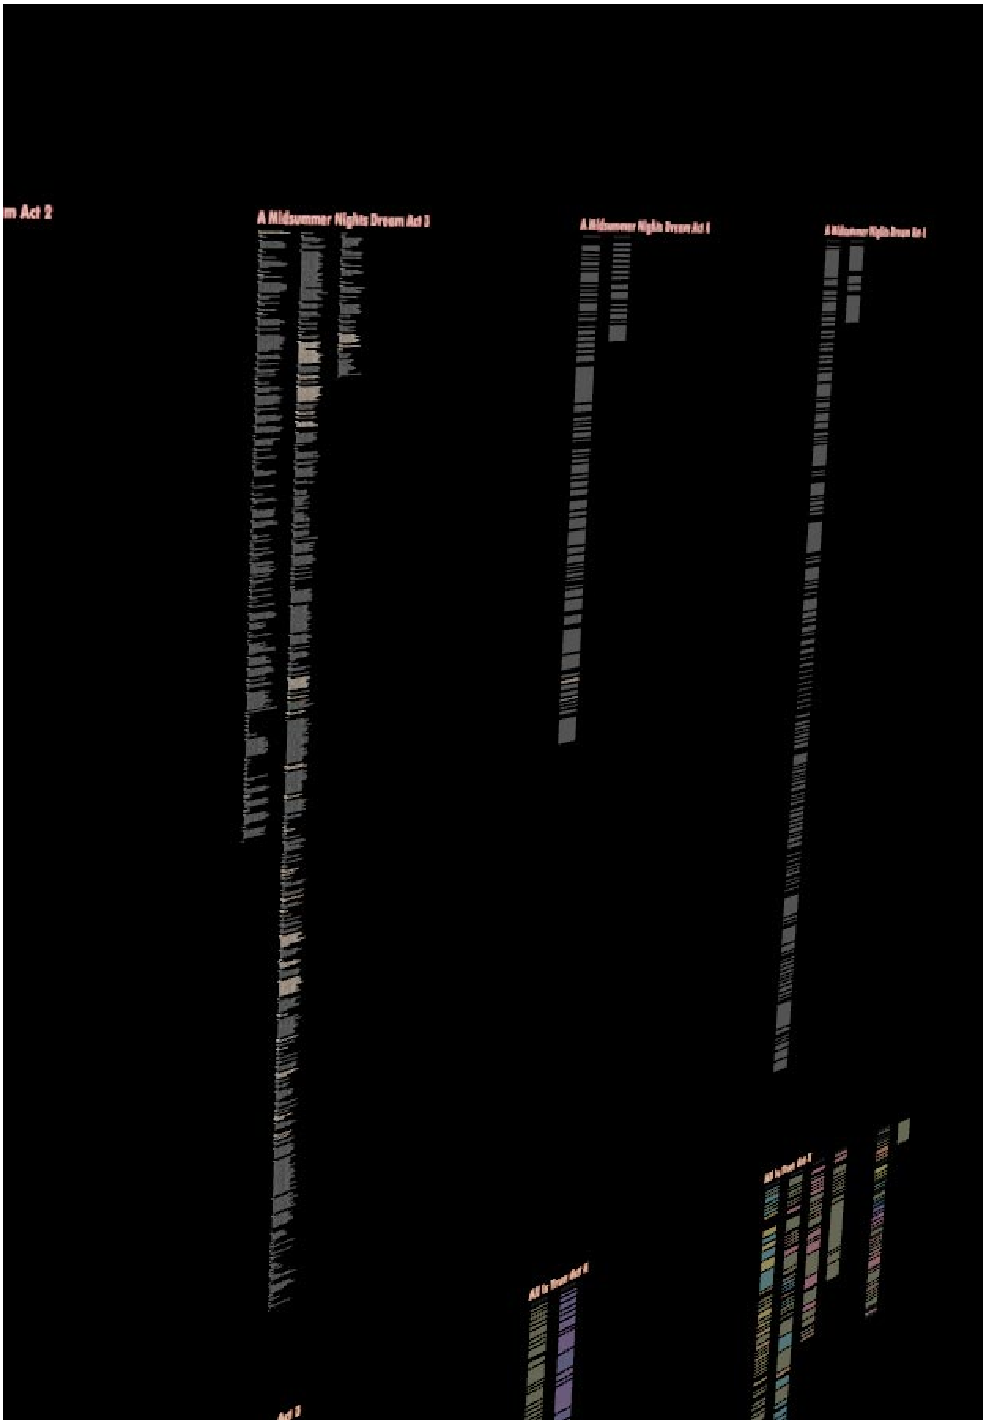
m Act 2
A Midsummer Nights Dream Act 3

visual chunk. Each of the five acts of the play are arrayed from left to right, and finally, each play is arranged one above the other.

In the information landscapes work we developed a method of displaying typographic forms at any size, position, and orientation in three-dimensional (3D) space. A virtual camera is then moved through the space, exploring the information, both text and images, which inhabits the space. By adapting and expanding on techniques developed for two-dimensional graphic design to the mostly unexplored realm of three-dimensional design, a number of visual experiments were produced. These interactive sketches addressed a number of design issues, including the use of perspective, scale, space, and interaction to create simple and flexible visualizations of information.

First, we must remember that letters were designed to be viewed directly on a flat two-dimensional (2D) surface and, by allowing arbitrary viewpoints, perspective distortion is created. Although this is correct for 3D perception, it is less than ideal for reading. Since it is not always possible to guarantee the angle of the view relative to the angle of the text, one cannot be certain of maintaining the integrity of the letterform. Each new angle will result in a differently shaped letter and at extreme angles the text image can be reduced to a line. Furthermore, when the camera moves behind the text, it looks reversed as though seen in a mirror. While certain word shapes can still be recognized in less than ideal circumstances, in general there are few views from which text holds its legibility. One can solve this problem by constantly rotating all of the text objects so that they face the viewer, but this has the problem of destroying the overall structure of a complex three-dimensional space. Another solution is to constrain the movement of the camera to maintain a minimum legibility of text in the scene, but such constraints are not always acceptable.

In Figure 14 a map of North America is labeled with the country names Canada and the United States (partially seen), which are easily read when the map is viewed from the common orientation with north at the top. Still, nothing prevents the viewer from moving to the North Pole, from where the text will appear reversed.

FIGURE 14. Geographic labels in 3D.

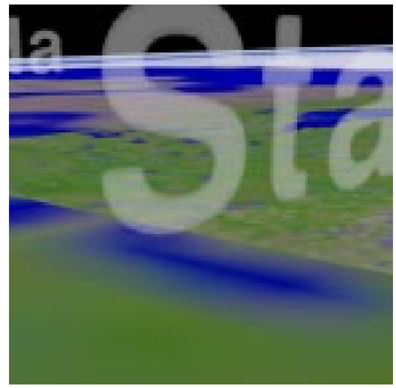

A graphic designer can use size differences to visually distinguish certain elements in a text, such as the headline of a newspaper story or the fine print on a contract. In a three-dimensional space, you cannot always resolve the relative size of two objects. If one object appears smaller in the picture plane, it could actually be smaller, or it could be the same size and farther away, or it could even be much larger and very far away. So, in the design of an information space, one must be careful about using size as a differentiating variable. One interesting advantage of designing an information space is that the designer can use an almost unlimited range of scale to represent information. In print, the difference in size between the largest and smallest element is limited by the resolution of the printer and the physical size of the paper, but in a virtual space, typography can have almost unlimited variations in scale.

We can also examine the use of space and how that differs in the design of digital media. Traditional graphic design has always been concerned with the disposition of pictorial space. There is a constant tension between the desire to include as much content as possible and using white space to create a harmonious and uncluttered image. In two dimensions, the designer is constrained by the limits of physical space and the static nature of the medium. In three dimensions, despite the easing of those constraints, the problem remains to create clear, legible relationships. Because of the ease with which one can create content at different scales and orientations, it is possible in an electronic landscape to present massive amounts of information while giving an impression of low visual density. This capability can work against the design as easily as it can help. Vast amounts of empty space can lead to an environment with very low legibility. A legible landscape is one that is meaningful, rich, and clear. Kevin Lynch wrote in his book The Image of the City [LYNCH60]:

> By this [legibility] we mean the ease with which its parts can be recognized and can be organized into a coherent pattern. Just as this printed page, if it is legible, can be visually grasped as a related pattern of recognizable symbols, so a legible city would be one whose districts or landmarks or pathways are easily identifiable and are easily grouped into an over-all pattern.

Because typographic elements can appear at any scale, an information landscape can create a good sense of overview and context while losing a clear understanding of the density of the content. In a paper book, we can understand at a glance the amount of text by the size of the book, the width of its spine, and so forth. As we dynamically shift the scale of an electronic text, we may not be able to have a constant yardstick or scale against which to understand the size of a text. This can be seen in the large spread on page 32, where an entire Shakespearian play is visible. It is difficult to get a sense of how many words are in the play or how long it will take to read.

Finally, we can examine how we can use either increased spatial resolution or time varying images to increase the density of information displays. The constant dimensions of the computer screen and its low resolution (100 dots per inch, or dpi) when compared to print (over 600 dpi) greatly limit the amount of text that can be simultaneously presented to the user. One solution to this problem is simply to increase the resolution of the display (Figure 15). The Visible Language Workshop has built an extremely high-resolution display that enables the simultaneous display of large amounts of information [MASUISHI92].

FIGURE 15. 6,000x2,000 pixel display prototype

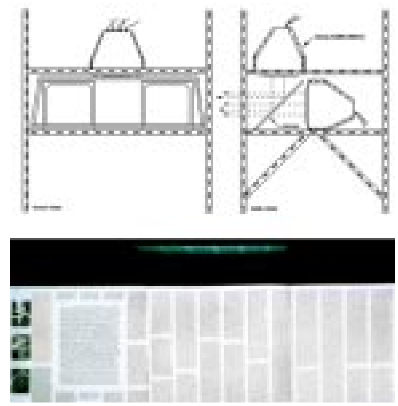

This approach, however, makes extreme demands on display technology, compute power, and data bandwidth. Instead we can use a standard size display and dynamically shift our viewpoint around a larger virtual information space. Although the resolution at any one moment in time is still limited, we can smoothly move from an overview to a detailed view in a manner that helps to maintain the all-important context of the larger body of information.

In addition to the issues of perspective, scale, use of space and spatio-temporal resolution, there are some basic perceptual issues which need to be considered when typography is used in three-dimensional space. In order to maintain legibility, the speed of movement must be precisely chosen. Also, the designer must be careful about the deformation of a type caused by the change of the view angle, as described earlier.

The purpose of the Virtual Shakespeare Project was to explore the design of a large body of textual information. The amount of text is on the order of one million words and the work itself has many structures that can be made visible: speeches, scenes, acts, and so forth (Figure 16).

FIGURE 16. The complete works of Shakespeare in paper form. 896 pages, 876,406 words, 7 lbs.

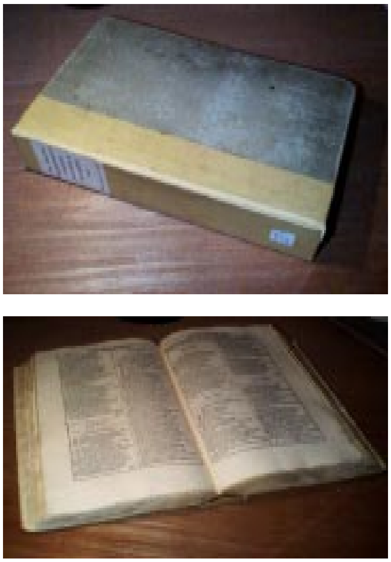

FIGURE 17. Different scales

A rendering model was developed that is optimized for rapid navigation and changes in scale. If your viewpoint is close to the text it will be fully rendered. If it is farther away, and therefore smaller on the display, a simplified texture is used in place of each line of text. This technique, called greeking, maintains the overall shape of each line, although individual words are lost. As distance increases to the point where each line of text blurs into the next, each block of text is drawn as a simple rectangle of the same size and overall density. Breaks between the dialog of different characters are used as the delineator for the larger text blocks. This means that even at a great distance, the reader can still follow who was speaking and how much was said. The final stage comes when the dialogs become so small as to merge together. At this point each scene is rendered as a simple rectangle. These different views are shown in Figure 17. As we move back to include ever larger amounts of information in our view, the display of the information becomes more abstract while maintaining visual continuity.

It is important that all transitions from one level of detail to another be as smooth and inconspicuous as possible. The reader should believe that all the information is there on the screen. A simple cross-fade is used to blur the transition from one state to another. This works quite well; however, there are still some problems associated with color and the typography itself. As typographic elements change size, it is not always possible to maintain a consistent perceived color for the text. As an object becomes smaller in the visual field, its surroundings have a greater effect on its perceived color. In the case of the rendering engine used in the Virtual Shakespeare Project, the text becomes darker as it gets smaller. This becomes a problem when color, or even brightness, is used to distinguish one object from another. Figure 18 shows how highlighting can work effectively at two different distances. Because it is easy to see Titania's dialog while viewing the entire play, we can readily explore her thread through the narrative.

FIGURE 18. Highlighting

In addition to perceived shifts in color, the typographic forms themselves appeared somewhat unstable. That is, the thickness of the stems appears to change, the serifs wriggle and fade, and the counters tend to clog up when the letterforms shrink. Typefaces are generally designed to be used in only a small range of sizes and always so that they are flat to the page. One problem that needed to be overcome was the fact that letters were being used that were much larger or smaller than had ever been intended and that some view angles created such perspective distortion as to render the typeface illegible. Through experimentation, we have found that some typefaces are more sturdy in this respect than others; however, no single typeface is adequate for all situations. What will be required is to design a new kind of typeface that can dynamically adjust its form to its environment. For example, in the days of lead typefaces, each size was designed independently. Designers knew that a letterform that looked clean and elegant at 12 points would be tall and spindly looking at 6 points, so letterforms became squatter and thicker as they grew smaller. New work in multimaster typefaces by Adobe Systems Inc. [ADOBE95] allows the generation of a range of faces from a single master; however, they have not been used in a dynamic display. Future work will explore the generation of variable typefaces that can adapt to suit their environment.

Despite these problems, it is possible to use cues such as color or change in typeface to visually highlight portions of the text. For example, one may be interested in seeing all of the dialog for a specific character. Whatever visual technique is used, it should clearly distinguish the selected text at a wide range of scales. If a change in typeface, such as boldface or italic, is used it can be difficult to see when viewing an entire scene or act. Dynamic highlighting, such as blinking can be effective when the selected text is small and could be swamped by other information; however, it can also render illegible just that information that one wishes to make visible. To avoid these problems, I used brightness to cue the dialog of a character. The contrast between the selected and unselected text was continuously adjusted to account for changes in scale. As the distance from the text increases, bright objects become surrounded by more and more black space and must be made brighter to seem to maintain a consistent visual distinction from the unselected text. The ability to visually filter out some portions of the text enables the reader to see patterns and structure that were impossible to find in the traditional book format. For example, when Titania's dialog is highlighted, you can immediately see her role in the narrative structure. She is introduced in the second act, has a rather long soliloquy, and then comes and goes a few times during the rest of the play. By allowing us to read within a meaningful context, the computer can fundamentally change the kinds of understandings we can glean from a text.

The use of space in an information landscape is fundamentally different from that of traditional design. One example of this can be seen in the presentation of footnotes or supplementary material. In traditional book design there are few options for visually treating such related materials. Footnotes can be placed at the bottom of the page or in the margin and referenced by number or asterisk. The length of the footnote is quite limited, unless it appears in an extremely small and barely legible typeface. Tschichold carefully enumerates the many typesetting problems associated with footnotes in The Form of the Book [TSCHICHOLD91].

Hypertext systems, such as those used to access the World Wide Web, allow the designer to tag a text with footnotes of arbitrary size. However, when the reader selects a link the footnote appears and completely obliterates the original text. The use of 3D space gives the designer new possible solutions to this problem, two of which are shown in Figure 19.

FIGURE 19. Footnotes

One obvious solution is to take advantage of the ability to rapidly change scale and place the footnote next to the referring text, but much smaller, as shown in the first image. Since size is arbitrary, you can even put a footnote in the dot of an i or in the period at the end of a sentence. The problem with this solution, which this extreme example makes clear, is that it is difficult to see both the footnote and the referring text at the same time. In the second image, the footnote is shown at the same size as the main text, but at ninety degrees to it. When looking directly at the text, the footnotes, being infinitely thin, disappear, but with a quick twist they can be read. This solution has the advantage of providing quick, yet unobtrusive access and allowing the simultaneous display of both texts.

Although these new rendering techniques allow many different views of a large-scale text, the visualization is only useful if the user can easily navigate about the text. A number of new methods of navigation were developed to address this problem. Most current interface paradigms (windows, buttons, mice) were based on a two-dimensional screen. A three-dimensional model requires new kinds of controls that allow for easy manipulation in space.New navigation approaches, using physical controls and positioning sensing were developed and are described in more detail in “Symmetric interfaces” on page 85.

All of these methods increase the ease with which the reader can navigate the text and greatly increase the utility of the system. For example, giving the user a physical handle on the text makes it easy to quickly reorient it to see footnotes, which may be placed at right angles to the main text. The most important lesson learned from these experiments was that it is impossible to separate the visual design from the design of the interface. Subtle interactions between the visual design and the physical controls may facilitate many actions but make others more difficult.

Other forms

We've looked at ideas of form in the design of information, in particular the kind found in books. There are, of course, other kinds of information content which could benefit from this treatment. Lisa Strausfeld's work in the display of financial information [STRAUSFELD95] is a good example of how different kinds of content will effect the use of three dimensional space. She developed a tool for examining spreadsheet data of mutual funds. It took the form of an open grid of data points in X, Y and Z dimensions. Each orthogonal axis was used to display a different type of information (fund, time and performance measures). The volume could be viewed from each of three directions and the numeric values would rotate into the correct plane for reading as the viewpoint shifted.

FIGURE 20. Financial Viewpoints

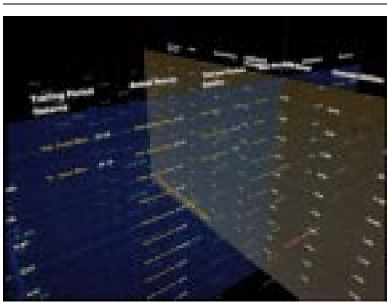

This particular form allows for the easy reading of rows and columns of numerical data, no matter which pair of axes are viewed. The form follows from the data and the function of the tool. The adjustment of the form to the function required of the display is the task of he designer, whether human or computer. In their Geospace project [LOKUGE95], Lokuge and Ishizaki, defined a meta design which was interpreted by the computer to dynamically adjust the form of the design in response to the evolving functions required of the user.

Through a discourse with the user, the computer would build a model of the requirements of the user and redesign the information display. Since the information was primarily geographic, the data objects were positioned on a two dimensional grided surface. Adjustments to the typographic forms were primarily those of size and opacity. The plane could be rotated about the view angle and the text and other data would adjust their orientation to conform to the new viewpoint.

We can see that the form is necessarily tied to the type of information displayed, but there are certain constants, such as the need to orient text to face the reader, which will be common to any design. Another constant is the need to provide, in the form itself, a logical path through the information.

Navigation and wayfinding

The use of three-dimensional typography has fundamentally changed the way we think about the use of space in graphic design and how the surface of the computer screen is understood by the designer. When the design of three dimensional typography is cast as a landscape design problem, we understand that it has more to do with creating compelling views than with the strict arrangement of elements. As a garden design will lead one through a series of vantage points which hide, reveal and accentuate a series of features [MOORE88], a journey through an information landscape should provide a meaningful context for the information elements.

Just as one reads the physical landscape in order to navigate the world, so too must people be able to find their way about information spaces. The designer should be careful, however, not to confuse the abstract spaces of typographic information with the roads, subways and buildings of our built environment. In the visual design of information spaces it is much more important to understand relative and ultimately fluid relationships between shifting and mutating information chunks than the fixed elements of the real world.

For example, in navigating the urban environment, we make use of fixed signs and landmarks, transportation systems and placed based addressing schemes, such as street addresses in the US or chome (postal neighborhoods) in Japan. In information systems, the data itself can have an inherent address (e.g. Exodus 4:12 or Romeo and Juliet, act 3, scene 1). Since we can move instantaneously from one location to another, getting the "lay of the land" may be less important than having a clear view of your current location and meaningful jump points from there.

The tools for getting from one location to another are unclear, but we can understand what is required for a usable interface. It is always easier to find your way to something that you can already see. Through use of scalable text which can be layered, we can keep much more information simultaneously visible than was previously possible. Nonetheless, we still require clear indications of where we are within a space and what lies just out of our view.

We must also consider that any journey through space is also one through time. No movement is ever truly instantaneous and the way in which we move and how the journey unfolds through time can be of great help in revealing the underlying structure of a landscape.

The correspondence between narrative space and architectural space or the landscape is a natural one. There are many examples, from the friezes in which the story is organized along the lines of the architectural structure [BRILLIANT84] to the songlines of the Australian aboriginal people [CHATWIN87]. In the book Learning from Las Vegas [VENTURI72], Robert Venturi describes the strip as a textual event. In her design of the book, Muriel Cooper, visually demonstrated the sign-map of the strip.

In this chapter we have examined the creation of three-dimensional forms for information display. Early work in Information Landscapes was expanded in scale in the Shakespeare project. Other kinds of information will have different requirements that guide their form. Finally, we can look towards landscape design and architecture for hints on creating navigable spaces of information. In the following chapter, we will look in more detail at typographic forms allowed by the computer. In particular, the use of focus for layering information, the placement of elements relative to each other, and the effect of changing scale on the design.

CHAPTER 4

layering, juxtaposition and scale

The primary problem designers encounter when moving to the digital medium is that the resolution they are accustomed to in print is completely lacking on the computer screen. Often, it is desirable to show more information at one time than can reasonably fit onto the display. We can take advantage of the computer's ability to create multiple dynamic layers of information and to rapidly change the scale of information elements to overcome that constraint and to go beyond anything that was possible in the realm of ink on paper.

The context within which we find information often tells us as much as the information itself. Although new markup languages, such as HTML, allow designers to link many pieces of information together, the information elements are still viewed as isolated, fragmentary bits. It is now possible to control the focus and transparency of information objects, as well as color, typeface and other variables dynamically. This gives us the opportunity to concurrently display multiple threads of information and dynamically shift visual focus from one to another.

Through the use of transparency and focus we can effectively layer multiple threads of text. One layer can recede into a blurry cloud while another will suddenly “pop” into focus and float above the other layers. If the computer can deduce which layer is of current interest, this focus shifting can be, at least partially, automatic.

Scale has always been implicit in design because one designed real, physical objects which had a certain size and relationship to the human form. Objects are designed to fit the body, such as a book which one can hold in one's lap. Architectural design is likewise intimately connected with the human form. This extends even to the more abstract realm of graphic design. Typefaces are made to be read from a certain distance and occupy a certain size on the retina.

In the virtual space of the computer screen we are free to explore a vast range of scales. In particular, I am interested in how to design for both the reading scale, where a display can hold five hundred or a thousand words, and a contextual scale, where a million or more words can be in some way visible. While we understand a great deal about how people read characters which occupy a hundred or so pixels, it is unclear how to abstract text for display at single pixel or subpixel resolutions.

The Talmud project

Unlike the Shakespeare project described in the previous chapter, the Talmud project directly addresses the issue of working with multiple texts simultaneously. Because it was necessary to show several texts and the relationships between them in the same space, the Talmud, in its complexity, helped clarify the visual and interaction problems involved.

The Talmud is a collection of sacred writings on the Torah or old testament. This project was built around an essay by the philosopher Emmanuel Levinas [LEVINAS94], [LEVINAS82] who's commentary on a tract of the Talmud, which itself is a complex, nested series of references to the Torah, forms an intricate web of text and references. This style of writing is called hermeneutics, the reference of scripture to support an argument.

The primary goal of the Talmud project was to create a workspace in which the relevant texts could coexist and interact. The fact that these texts are themselves complex and carry a long history of study brings the issues into a sharp relief. The system should be fast and responsive - after all it only takes a moment to flip through a book and find a particular passage. It should give the sense that all of the material is close by and accessible. It should also reward further study, meaning that even though a novice should be able to quickly orient himself in the texts, an expert should be able to "perform" with a degree of precision that is evidence of his or her knowledge and experience.

The chosen texts deal with the subject of the Cities of Refuge. When one has caused the accidental death of another, the law recognizes that this is not the same as murder. For example, if a man is chopping wood in his yard and the axe-head flies loose and strikes dead a person walking down the street, the law recognizes that there was no intent and the manslaughterer is 'subjectively innocent'. Nonetheless, the family of the slain man has the right of blood vengeance. This paradox, of existing on a state of both guilt and innocence, forms the basis for the reading. Levinas, in his

From "Learn Talmud" by Judith Abrams [ABRAMS95]: ...Get used to having many volumes of books out at one time. By the end of a study session you could have several books spread out on the table: the volume you are studying, a Bible to look up the verses that are cited in the Talmud, the Reference Guide for this set of Talmud, volumes of the Encyclopaedia Judaica to provide additional historical background information, various Hebrew or Aramaic dictionaries, and other volumes of rabbinic literature. This isn't messiness. This is the traditional mode of study and it really feels great... The talmud is studied in a particular manner, which although not unlike the method in which most would approach a scholarly work, has been made explicit over many centuries of study. The talmud should be studied with another person (hevruta) [Stone98] and one of the two should be more experienced than the other. The act of reading should be punctuated with argument and discussion of the issues raised by the text. In a typically Talmudic expression, this is described with the prescription - there should be crumbs that fall into the binding of the Talmud because the scholars will be so engrossed in discussion that their lunch will find its way into the book.

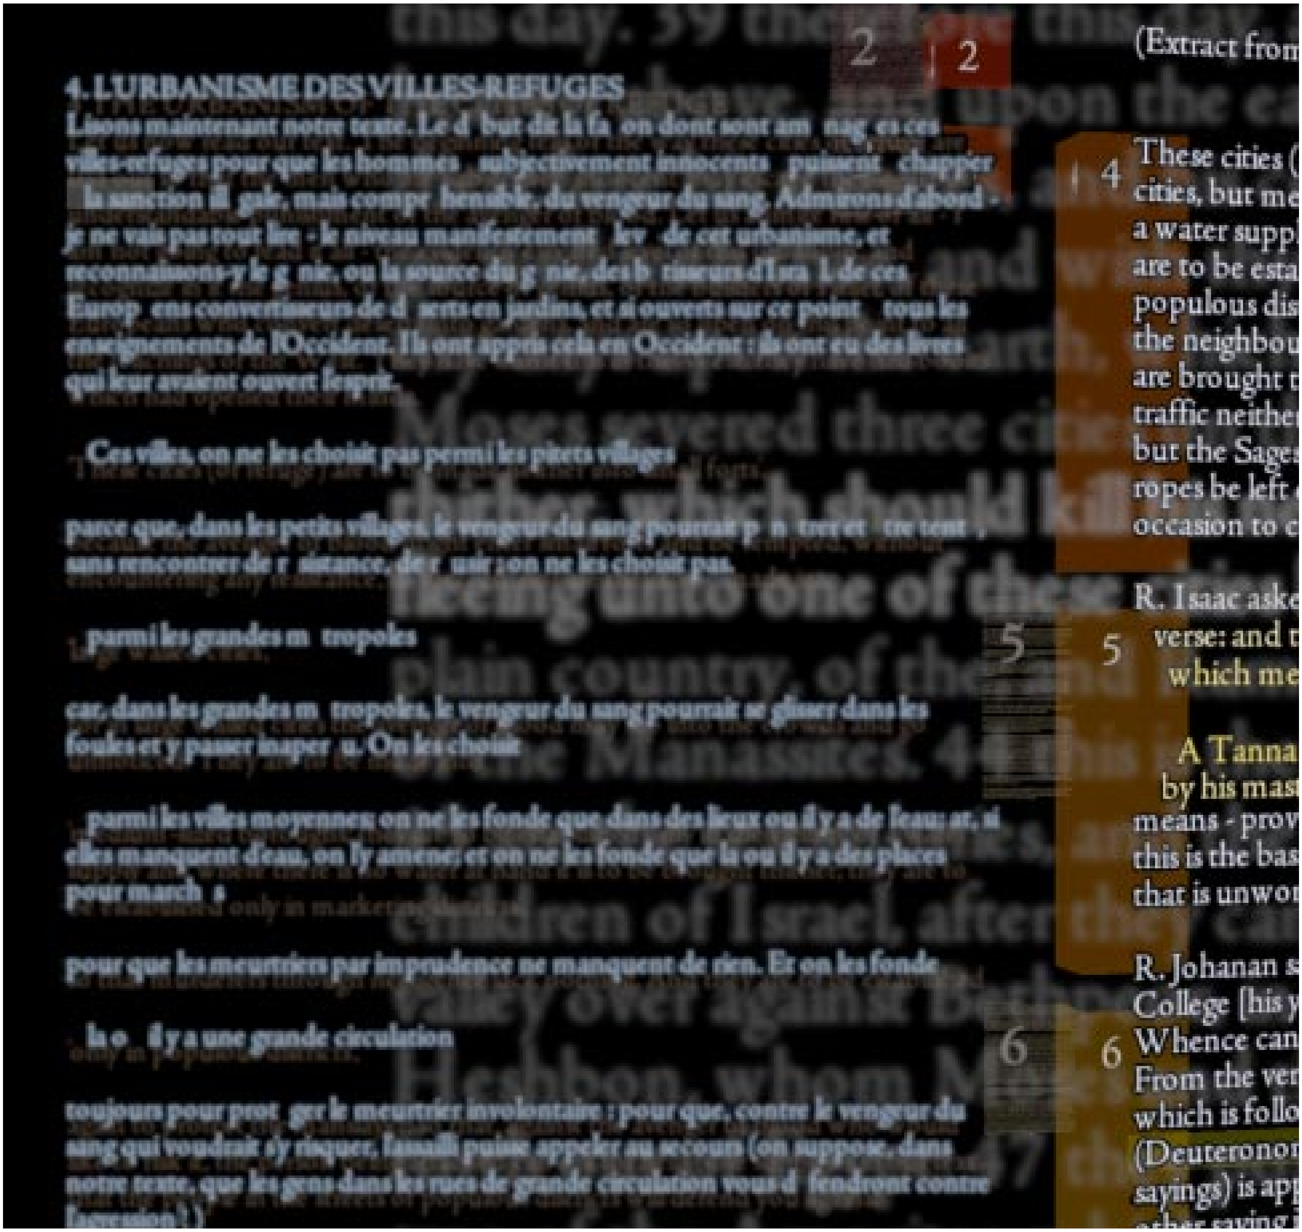
4. L'URBANISME DES VILLES-REFUGES
(Extract from
These cities (
cities, but me
a water supp
are to be esta
populous dis
the neighbou
are brought t
traffic neithe
but the Sages
ropes be left
occasion to c
R. Isaac aske
verse: and t
which me
A Tanna
by his mas
means - prov
this is the bas
that is unwo
R. Johanan s
College [his y
Whence can
From the ver
which is follo
(Deuterono
sayings) is ap

A screen from the Talmud project. The three texts are all on the display, but only the Talmud (Makkoth 10a) is in focus.

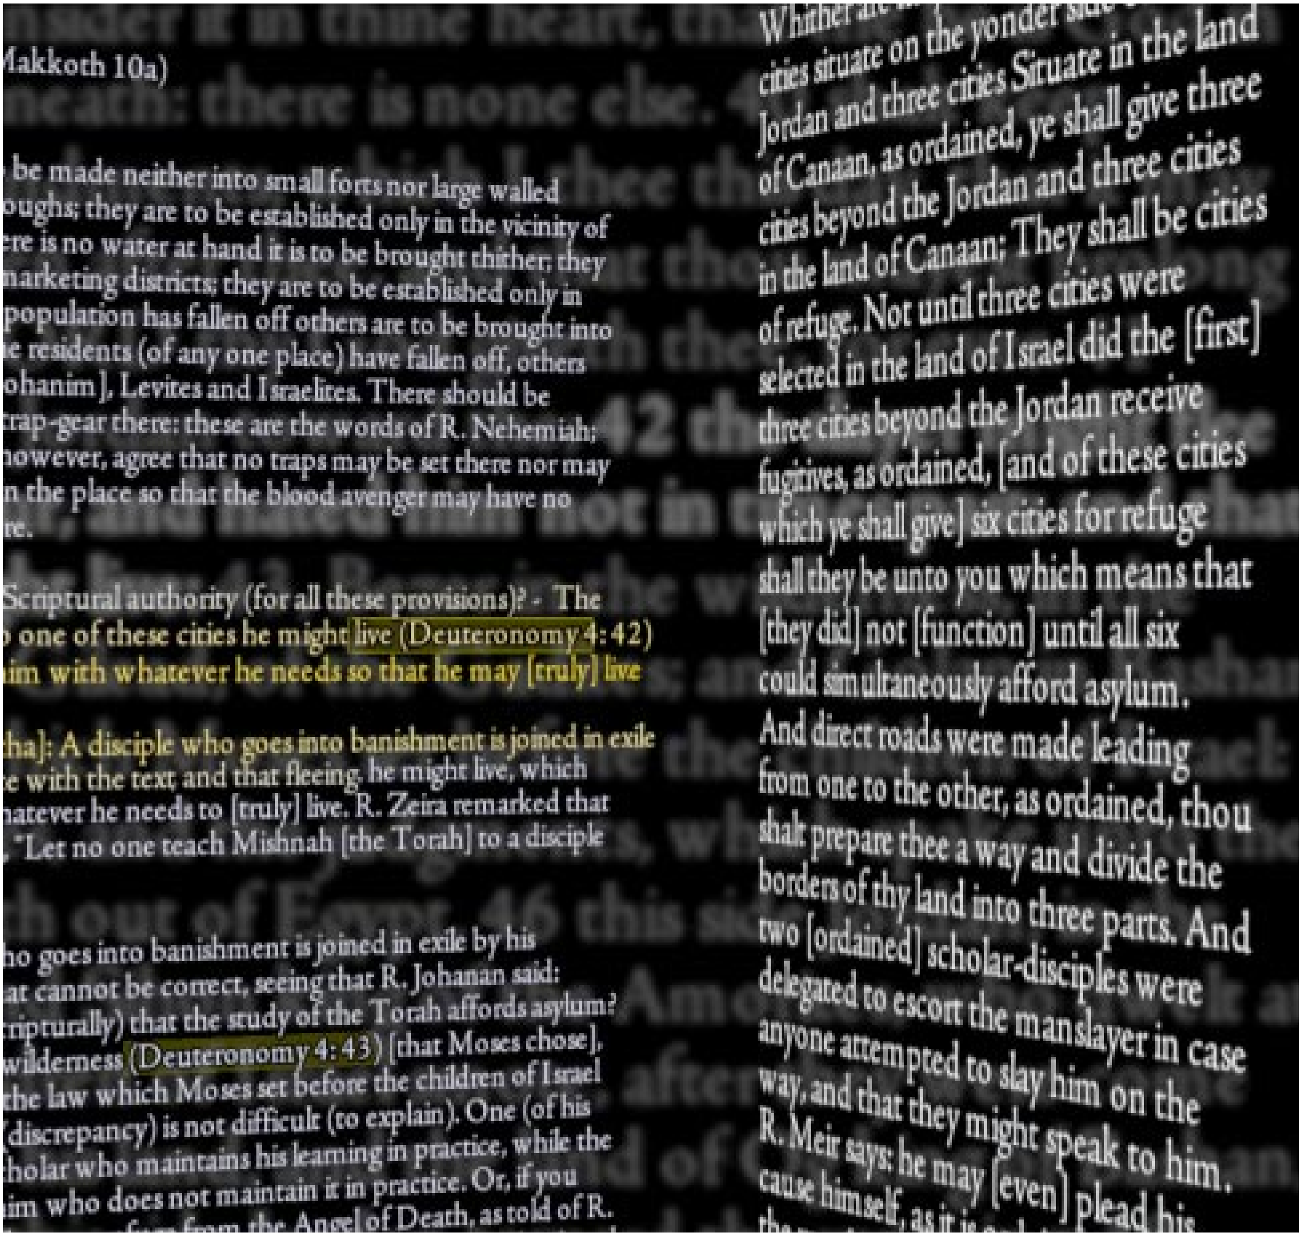

approach to the Talmud, brings this paradox into the context of 20th century life and tries to give the text the widest possible reading.

A visual representation of these interconnected texts should construct a space for discussion and argument in which scholars can pull and push the words as they dissect the intellectual issues posed by the text. Some of the initial designs for this project used graphic controls for navigating and controlling the three layers of text (the Levinas text, the Talmud and the Torah). This proved to be unsatisfactory for a variety of reasons. First, the controls existed in the same visual space as the data being manipulated and it was often difficult to keep one from visually conflicting with the other. Also, because graphic widgets rely on the computer mouse for control, it was difficult for more than one person to have control at any one time. Finally, there simply wasn't a good feel to the controls - they lacked the tactile quality of leafing through an actual book.

The solution to these problems was to create physical controls which exist in the space immediately surrounding the display. These controls will be discussed in greater detail in the chapter entitled *tangible issues*. This chapter will concentrate on the visual problems posed by the requirement that the various texts, under the immediate physical control of one or more readers, can visually coexist in a smooth and natural manner.

The primary problem faced by the designer in electronic media is the lack of resolution and space afforded by paper. The average computer display has about one million pixels and can display perhaps one thousand words. The resolution of paper allows for a larger number of words in the same space and, because the resolution and contrast of ink on paper is much higher, the type itself is of higher quality. And, because paper is thin and inexpensive, many sheets can be bound into a book which can easily contain over one million words.

Despite these limitations, electronic media have some distinct advantages over paper which we can exploit. The electronic display is a dynamic surface which can change and adapt over time. More important however is the fact that the computer processor can manipulate and understand the underlying model of the information. Unlike paper, which knows not what is printed on it, the computer can be programmed to intelligently react to changing inputs and models of both information and the user.

Layering with focus control

Layering is defined as the simultaneous display of two or more information objects within the same two-dimensional space of the projected display surface. This can occur when two objects are in fact occupying the same space or when a particular view into a three-dimensional landscape of information causes one layer to occlude, or pass in front of, another.

Even though you may want to display several information objects at the same time and in the same space, the reader's attention will only be focused on one at a time. If we know which layer is of interest at the moment, we can adjust the display such that the various layers appear to either "pop" out to the front or recede into the background. This is accomplished through a combination of focus and transparency controls. Different colors can also be used in this way, although it may not be desirable to constantly shift the color of an object. The goal is to make the minimum change which allows the selected layer to be easily read without

Critique
As mentioned in the introduction, there were three critique sessions during the development of the Talmud project. In discussing the early designs for the Talmud, there was some general comments about the positioning of the various elements and the way that they moved. Here are some key points by John Maeda and Paola Antonelli:

John Maeda
I was thinking that right now its not making use of the 3D - of the depth. Your elements are really 3D; those elements are the key. The depth isn't really coming out. It feels very planar to me.

Paola Antonelli
It needs clarity of purpose. Innovation per se doesn't really do anything unless it works. Always keeping in mind the practical goals is important. [talking about movement in 3D]...might be a little confusing. In a way that is the conflict in the design...you might have some innovation, but you also have to take into account people's habits. If you disorient people too much it might be counter-productive. I'm not saying you shouldn't have movement...but whatever movement you have should be smoother. It is really about clarity - at every single moment can I see what I am doing?

giving the impression that any of the layers have really changed.

We can see this in the treatment of the three basic information layers of Levinas, Talmud and Torah. Although all three reside in the same space, only one is ever fully in focus at any moment. So, for example, you might be reading a passage from the Torah while the referring page from the Talmud appears to hang behind the Torah and just enough out of focus so as to maximize the legibility of the Torah. This goal of "just enough" is quite difficult to achieve. The optical size of the two texts, their orientation to each other, their color and transparency all have an effect on legibility.

Although a fully rigorous study was not done, some general observations can be made. First, the smaller the text, the less of an effect that blurring will have. Once text gets below a few pixels in width it is too small to see the blur. In general, the smaller the text, the more you want to increase the transparency to give the same overall density to blurred text. Second, the more that the texts align with each other, the greater the blur needed to keep them separated. So, if two texts are at the same point size and the lines appear to be nearly on top of each other, the background text will tend to impair the legibility of the foreground text. If they are at different angles or significantly different point sizes, there will be less interference and consequently less blur may be needed to keep the foreground text legible.

By using changes in both blur and transparency it was possible to create dynamic shifts between the three information layers which provided clear. legible text in the foreground layer. Because this could be done without otherwise changing the position or scale of any of the layers, there was far less chance of confusion about where specific areas of text were on the screen, even those that were temporarily illegible. Before going on the discuss juxtaposition, lets talk in some more detail about the method used to create smooth transitions of focus.

A method for real time focus control

To blur an image, such as a letter A, you multiply it against a filter, called a kernel. This process is known as convolution and requires a multiply and add for each element of the kernel at each pixel. The greater the blur, the larger the kernel, and the number of computations increase by the square of the filter width. Early experiments by Laura Scholl and Grace Colby with the Media Lab's Connection Machine supercomputer took advantage of parallel processing for this computations, but each image took minutes to blur and display [SCHOLL91]. Even with today's fast processors, it is not possible to filter images in real time. So, to create text which could smoothly change from perfectly sharp to bleary-eyed blurry, a new technique had to be developed.

Of course, we are not dealing to with an infinite number of images. We only need to consider the hundred or so characters commonly used in each typeface. So, the idea of precomputing the required blurry characters and caching the results is compelling. Still, there is a limit to the amount of memory that can be allocated for all of the copies of each typeface, especially when you consider that the most efficient rendering can only occur when the images are cached in the graphics pipeline.

Even if you cache ten or more images per character, the reader will be able to see discrete changes in focus. In order to provide truly smooth gradations, and to limit the number of images to be cached, a weighted average of just two images is used. This is done by drawing two versions of the character directly atop each other and varying the transparency of each. You can think of this as a process of compression, where the image of the letter is broken into discrete spatial frequencies and then reassembled. The sharp image corresponds to the highest spatial frequencies and the blurred image has only lower spatial frequencies. Surprisingly, this method produced results indistinguishable from individually filtered text. It was decided to use just three master images - the original, a 3x3 Gaussian filtered image and a 9x9 Gaussian. If an even blurrier image was needed, a 27x27 filter would be used. To display an image corresponding to any filter in-between, the two neighboring filters are weighted and combined.

Dynamic juxtaposition

In traditional graphic design the space of the paper is used as a kind of map to the underlying information. Elements which are related are located on the page in proximity to each other. These spatial relationships are fixed once the page is printed. In the dynamic context of the computer, the elements are in a continuous state of change. In order to maintain a specific relationship between two typographic elements, they have to constantly adjust to changing conditions.

Lets look at two examples from the Talmud project. Each devises a solution to the problem of keeping related texts next to each other. In the first, two versions of the same text must be compared line by line. In the second, blocks of text maintain proximity despite changes in size, position and orientation.

Levinas' writing only reveals itself after long and careful study. Even after several readings a passage may not be entirely clear. As with any translated work, the next step is to re-read in the original language. This can be difficult when using printed books. The two versions are often in separate books and it is difficult to go back and forth without losing the essence of the passage. The pages do not correspond to each other and as the books usually will have different publishers, the type design is different. There is the approach in which different languages run in different columns through the pages of the book. Otl Aicher's Typographie [AICHER88] is a good example of this technique.

A recently published Dutch monograph by designer Harry Ruhé took a different approach. The Dutch and English texts were directly superimposed in transparent inks - red for Dutch and green for English. A pair of colored gels provided with the book selectively reveal each layer. This has the advantage of putting the control directly into the hands of the reader as well as conserving space. Still, it is difficult to go back and forth between the languages and colored text is not nearly as easy on the eye as black text.

FIGURE 21. "The red and green theme is totally unnecessary in terms of translation, but it works. It's not user friendly, but its a nice game" [ID98].

The goal in providing both English and French versions of the Levinas text was to allow the reader direct control over which language was primary and to show both when needed. A simple dial moves the display from one language to the other and various levels in-between. At first the texts were superimposed and transparency varied to fade one language into another. One could read the English, turn the dial and then read the corresponding French. This proved unsatisfactory when you wanted to see both texts at the same time - it wasn't enough to hold the image of one in the mind and switch quickly between them. The image of the one text tended to obliterate the memory of the other (this effect is called *masking*).

To address this, I tried to put the French text in the spaces between the English lines (leading). It wasn't legible, unless the leading was increased dramatically, which was undesirable. Finally, I tried increasing the leading dynamically as one text faded into the other. In this way, each text seen solo was set nice and tight, but when the dial was halfway between, the French text sat just above each line of English with some leading in between each pair of lines. The two languages were different colors, so it was easy to track either one from line to line or to glance at both simultaneously. Since the French text contained accents above characters, it was placed above the English. This increased the legibility of the French, without compromising the English. This dynamic juxtaposition of the two texts gives particular affordances to the task of reading Levinas in the original and in translation.

FIGURE 22. French (blue) and English (tan) text together.

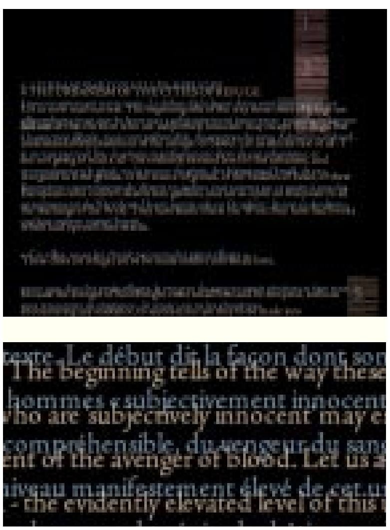

The second example has to do with the relationship between the Levinas essay and the tractate of the Talmud to which his essay refers. The goal is to be able to jump back and forth between the two texts easily as well as scale each larger or smaller. Each chapter of the Levinas text corresponds to a group of lines from the Talmud. To keep them separated the Levinas text hugs the left edge of the display and the Talmud tend to the right. A control shifts the boundary between the two texts. In the original design, this was seen as a dynamic margin - it could move so that one text had the majority of the screen space and the other shrank to fit into the remaining margin. By adjusting the scale of each text to fit the available space, it ensured that each text was always visible, even if it shrank to minuscule proportions.

As the prototype developed, this simple scheme became difficult to maintain. The Talmudic text could rotate to reveal either the *mishnah* (oral text c.200BC) or the *gemara* (commentaries c. 150 AD). Also, the overall scale of the Talmud (several volumes) dwarfed the scale of the Levinas essay. And it became clear that each chapter of the Levinas essay should be in proximity to those few lines of the Talmud to which it referred. This juxtaposition is accomplished by connecting the texts with springs when maintain the spatial relationships despite any rotation or scaling of either text and keeps the Levinas text always to the immediate left of the Talmud. Although the texts exist in three dimensional space, the visual relationship refers to the two dimensional projection of the texts. This was accomplished by creating special spring anchors which were outside of the normal object hierarchy, at the intersection of a plane parallel to the view plane and a ray from the viewpoint to the anchor point of the text. So, as the Talmud was scaled larger or smaller, or rotated or scrolled up and down, each Levinas chapter would dynamically adjust to maintain a specific juxtaposition-aligned to the top of the referring section, its right edge to the right of the Talmud and its left edge to the left edge of the display. Using a spring model for the constraint added a slight dynamic to the adjustment - following any user input which changed an element, the other elements would "catch up".

In each of the two preceding examples, we see how a dynamic constraint can maintain a useful juxtaposition between multiple texts. The visual relationships that are clearly apparent to the reader must match the structural relationships of the information. Moreover, we must be aware of the ways in which reader's eyes wove about the display and design accordingly.

Scale

The method used to render typography at various scales and the visual issues involved are discussed at some length in the previous chapter in the context of the Virtual Shakespeare project. Here, I would like to concentrate on problems associated with presenting several texts of different scales simultaneously. One of my early design goals was to keep all of the elements visible at all times. One of the great advantages books have had over electronic displays is their persistence. If I have three books laid out on my desk and I start reading one of them, the other two do not “disappear” off the desktop or behind an opaque window. They remain, in the background, but available.

Although computer window systems allow you to keep several views open at a time, only the information in the topmost window is visible. This places the burden on the user to remember what was on the various other windows. Furthermore, there is no smooth transition from one view to the next. So, it is difficult for the user to maintain a consistent mental model of the relationships between the different information objects. By allowing the various texts to scale between reading size down to postage stamp size, it is possible to keep everything in the current visual space and still have enough room to work. As well, the smooth transitions allow the reader to track where they are going within a text and what the relationships are between texts.

One difficulty encountered in the smooth scaling of columns of text is that one would often move close to a specific passage only to find that the neighboring columns would be equally strong typographically, despite the fact that they may be inconsequential. It is analogous to reading a story in a newspaper and accidentally jumping over a column into an unrelated story. When seen at a small scale, columns of text form an meaningful image which can “read” as the underlying information structure. You can tell at a glance how the stories on the front page of a newspaper are constructed and plan a path through them [SILVERSTEIN90]. You may pick up a paper, read its structure and then bring it closer to a comfortable reading distance and work you way

Critique
Some comments by Suguru Ishizaki about jumping from book to book:

Suguru Ishizaki
To pick one scenario: some researcher is looking in the library at different papers, one after another from the citations, keeping one and looking at the new, then reading some parts and making connections. Keeping one; discarding some...

Bill Mitchell, “And that is incredibly hard on paper”

...Yes. Its really hard to find the answer, like when you say, “I really like the argument that someone said, but I don’t remember where that was...”

Something like this could be useful.

through and around the page. You adjust the scale simply by moving the paper relative to your body.

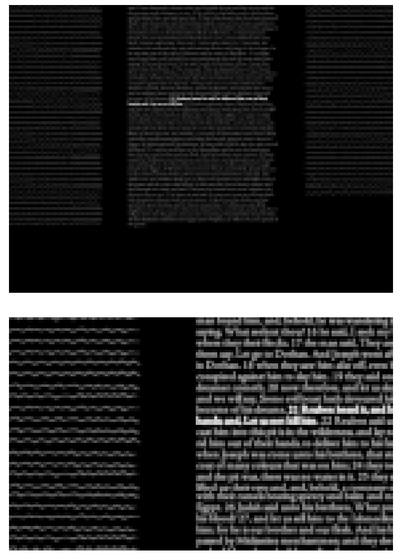

FIGURE 23. Sketchy and solid text side by side.

In the case of the Talmud project, it was possible to scale the text many more orders of magnitude than our newspaper example. Different elements had to respond differently at different sizes in order to maintain a legible focus on the section in question. In order to clearly indicate which was the column of interest, the neighboring columns would be rendered an a kind of sketchy way to show that they were there, but not the primary focus. At a far distance, all the columns would still look like quite the same. There are many other ways in which different elements of a text could respond to scaling - in the case of the Talmud, it might be sensible to scale the *gemara*, which comprises a smaller fraction of the total text, at a different rate then the *mishnah* with which it is interwoven.

In addition to the problem of showing texts at the same scale in different ways in order to guide the readers attention, the Talmud project, with its multiple texts, had to allow the reader to work between texts that may be at different scales. For example, a notation sketched by the user over the Levinas text may also direct the reader to a particular phrase in the Talmud or Torah. As the three texts vary in scale, the notation may lose any visual sense. Future work in this area could examine what types of structure add distinctiveness to scaled text so that small images of large works could present a more complete picture from a distance.

In this chapter, we have used the design of a Talmud study tool to provoke issues in the layering of information, appropriate juxtaposition of those layers and the use of scale. A method for real-time focus control of typography was described, including its use to simultaneously keep more visual information available and reduce the complexity of the display. Several methods of juxtaposing related textual material were discussed, with their advantages and limitations. Finally, the relation of scale to the multiple layers of information was examined. While this chapter was primarily about the disposition of space in the information display, the following takes a closer look at the role of time-varying displays. The sequential presentation of typography over time is, in a sense, another method of layering.

CHAPTER 5

expressive movement

The use of motion on the computer screen changes how we design, how we read and how we interact with typography. Previous work in this area has examined the expressive qualities of moving typography and a variety of computational approaches. In the Brain Opera, Yin Yin Wong and I explored the interaction between reading and listening. This led to an experiment in the pacing of expressive, dynamic typography. The use of physically-based modeling to create naturalistic animation was explored in the Stream of Consciousness project with Tom White. The use of dynamics in an interactive spreadsheet and the Talmud project demonstrate how these findings can be applied to information design problems.

While we may think of dynamics as being a property confined to media, such as film, television and the computer, there has always been a dynamic quality to graphic design. Each page of a book can create a dynamic effect in the way that the eye is moved across the visual space. And as each page is turned to the next, the changes create a kind of animation. Muriel Cooper, in her design of the Bauhaus book, spoke of the pages not as individual double-page spreads, but as a filmic sequence. In fact, she made a movie of the book by shooting a few frames of each page with a 16mm Bolex camera. The resulting film showed the entire contents of the book in a few minutes. This temporal overview reveals the stable, underlying grid which is implied, but not visible in each page. The changes from page to page direct the eye to what is important and provide a visual overview of the entire work.

FIGURE 24. Bauhaus design by Muriel Cooper for the MIT Press.

Recent work in dynamic typography by Suguru Ishizaki, Yin Yin Wong and John Maeda emphasizes the possibility of bringing expression to the written word through simple motion and careful attention to timing. I have examined how physically based motion can provide clues to the underlying information structure. Graphic designers often indicate the structural relationships between information chunks by placement, typeface, size, etc. These relationships can also be shown by carefully rendering the dynamic qualities of the various visual elements.

The Brain Opera, a production of the MIT Media Lab composed and directed by Tod Machover, provided a unique opportunity to examine the relationship between typographic voice and the musical voice. In the creating of an animation which would act as a visual counterpoint to the music, we conceived of the typographic design as a musical score, with various typographic voices interacting with each other and with the musical voices.

This project and others play off of the idea that there are parallels between the typographic image and the spoken voice. In order to probe these parallels more deeply, an experiment was run which examined the effect of speech-like timing on the perception of serially presented text. The results showed that despite the reported experience that they "heard" the words as they were read, subjects appear to process written text at rates unrelated to the way that they are spoken.

Motion base on physical simulation can give abstract objects the feeling of physical properties such as mass, volume and energy. These techniques were used in the Stream of Consciousness to create typography which seemed to float on moving water and with which the audience could interact in a natural way.

There are many kinds of forces we can create to act upon our typographic objects. Beyond the simple force of gravity, we can create wind, springs, jointed constructions and viscous drag. Through the careful design of the physical properties of the information, and our interaction with these dynamic systems, it should be possible to create spaces where the objects and their structural relationship to each other can be quickly "read" and understood.

Some related work

In the book Dynamic Form [MAEDA93], John Maeda differentiates between a dynamic skin and a dynamic form. In the past, larger objects were more complex than small ones. This meant that the industrial artist had space proportional to the complexity that needed to be expressed. With miniaturization, the surface available to the designer has shrunk, while the complexity of products has increased. The designer has used the unlimited dimension of time to give objects changing, *dynamic skins* which reveal the interior function. Maeda seeks to expand on the dynamic of the skin by modifying the volumetric form over time as well. He proposes a programming tools in which the software is created by the physical arrangement of modular components in which changes to the form create changes in the software. Thus a dynamic form visually indicates the interior function.

FIGURE 25. A calendar design by John Maeda for Shisedo.

In his dynamic designs, Maeda expresses the power of computation with great clarity. He describes his art as, "complex weaves of temporal graphics that appear simple because all the details have been hidden along the axis of time" [MAEDA98]. In his Reactive Books series, he uses dynamics to express *reaction* at a sensorial level. He has used the JAVA programming language as a platform to transmit those reactive systems. In a series of calendars for Shisedo, the computation becomes visible through interaction and the motion of the display.

FIGURE 26. Yin Yin Wong

Animated typography can also be used to enhance the emotional content of a message. Yin Yin Wong has explored the use of dynamic forms to create expressive visual narratives (see Figure 26) [WONG95]. By using moving typography, she was able to overcome the resolution limits of the display to create type that appears to be speaking to the reader with an incredible density of meaning. Her work hints that the computer cannot only faithfully reproduce text, as in a book, but can express the meaning of the text, as an actor performing a role.

Suguru Ishizaki took that notion of typographic performance on step further in his Ph.D. thesis [ISHIZAKI96]. He looks at dynamics as an emergent behavior of interacting software agents. Using concepts from the performing arts, he develops a design language in which performers, each responsible for the display of a segment of information, interact and negotiate the changing design.

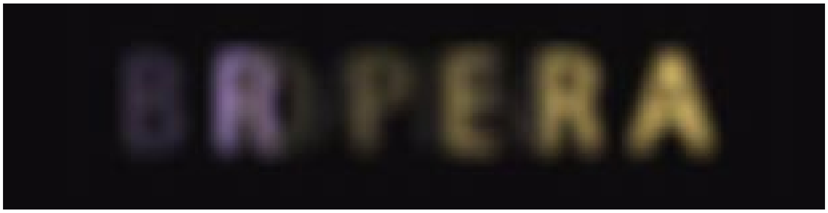

Minsky Melodies

Minsky Melodies is a typographic and musical feast for the mind. This animation was produced as part of the Brain Opera, developed at the MIT Media Laboratory under the direction of Tod Machover. During the second movement of the opera, the music revolves around the words of Marvin Minsky as he muses about music and the mind. Because of the importance of the lyrics to the composer, he wanted the words to be understood, even when sung in allegro form. A characteristic in this type of singing is that it is sometimes difficult to both listen to the musical composition and fully understand the words. Animated typography was used to visually prime the audience to hear the lyrics. As the words are vocalized, the same words dance and perform on a large screen behind three musicians. Visually rich treatment of typography

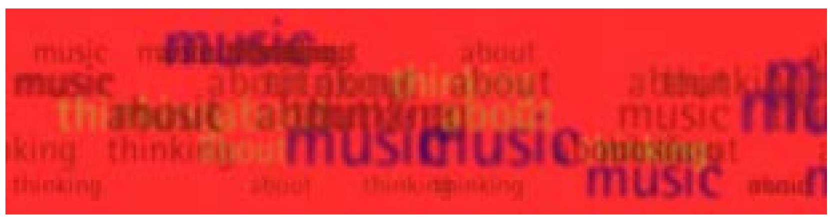

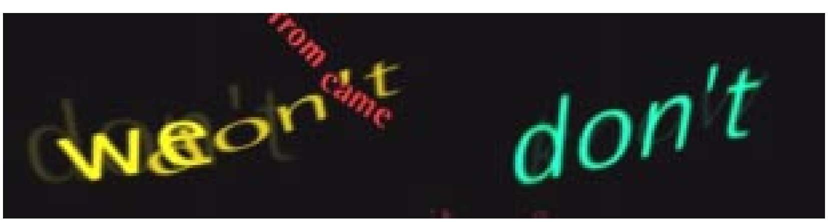

played as a kind of musical counterpoint to the words being sung.

In making the Minsky Melodies, we confronted a set of design issues and problems unique to combining motion typography with music. One problem was perceptual synchronization between seeing and hearing. For example, when a word is presented at the precise moment it was sung, it was perceived as being too early or too late. Careful treatment on timing led to enhanced perception of simultaneity between the visual and aural.

Another concern was readability of temporal typography. Unlike spoken language which can be heard from any direction, a reader's eyes have to focus on the words to read them. The composition was designed to visually flow from one word to the next, priming the reader to attend to the space where each word would appear. Readability is also an issue when typographic forms change and move over time. For example, in three-dimensional space, type forms may become distorted depending on the perspective and view angle. The text can even disappear when viewed from its side. Depending on the speed of the word as it moves across the screen, the word may not be perceptually distinguishable. To solve this problem, techniques such as blur were explored to enhance the readability of moving type forms.

In addition to addressing the issues, our design goal was to create an overall visual narrative that enriches the inflection, tone of voice, mood and meaning of the lyrics sung. The addition of the temporal dimension extends the expressiveness of traditional typographic language. Attributes such as size, color and form were enriched with temporal attribute such as motion. For example, motion was used to elucidate multiple overlapping voices. In the score shown in Figure 27 the soprano is singing four different and overlapping lines. One of the properties of music is that, as Marvin Minsky says, it "lets you think of three or four things at a time."

FIGURE 27. A section of the score which shows four interwoven voices.

We were unsure if we could give the audience the idea that they could read more than one thing at a time. By using motion paths to distinguish the four typographic streams, we were able to create the sense of four simultaneous voices, even if the audience only thought that they read the actual text.

Perception of temporal typography

New media are enabling designers to think of typography not as a static form, but as a dynamic event. Words can dance on the computer or television screen, enriching the communicative power of text. As smaller and more powerful computer devices proliferate in our lives, we need to develop new techniques for presenting the maximum amount of information in the most efficient manner on displays of limited size. Early work in the Rapid Serial Visual Presentation (RSVP) of text has shown that reading speeds can increase when words are presented sequentially to the reader at a fixed location.

As type is animated, the written text begins to take on some of the qualities of spoken language. An experiment was designed to examine the use of spoken rhythm in the visual presentation of text in comparison to a rhythm based on predictable word durations. If reading temporal typography has a cognitive relation to those processes which handle spoken speech, then we might expect that the rhythms of spoken speech would increase the readability or perceived quality of the presentation. On the other hand, if reading text is fundamentally limited by our visual processing, we should find that reading speed is optimized by giving the visual system the best possible chance to decode each word.

In addition to examining the tempo of the presentation, a variety of design parameters that increase the sense of voice were explored to see if they would effect the presentation.

These designs ranged from one which gave no indication of speaker or stress, and merely presented each word in the same location at the same size, color, typeface, etc., to one which indicated the speaker by position and color and stressed words through changes in size.

This work has a strong connection to earlier research in Rapid Serial Visual Presentation (RSVP). RSVP is a technique developed to study cognition in reading. It works by presenting one or several words at a time in a sequence where the experimenter has control over how long each word or grouping is displayed. Overall rate, relative duration of particular words, duration of pauses, alignment, masking, and number of characters shown are some of the parameters that can be varied in RSVP [POTTER84]. The task can test memory, comprehension, comparison tasks, or interactions with another task.

Speech recoding is one theory which is of particular importance to this experiment. This refers to the experience readers have of hearing what they are reading. Words are read and then recoded and interpreted as if they were auditory input. An experiment by Petrick and Potter [PETRICK82] showed that even at RSVP rates that were as fast as 12 words per second, they could show interference for a probe word which was acoustically similar to the expected word. RSVP rates do not however correspond to listening rates. Speech compressed to more than 8 words per second becomes very difficult to understand [WALLACE82], while RSVP shows comprehension at rates up to 12 words per second.

Yin Yin Wong [WONG95] proposed a model of dynamic visual treatment of written language in which the temporal presentation of typographic forms enhanced the communicative ability of text. Several of her experiments examined the time each word would remain on the display. In some cases, important content words were given greater expression through animation and remained for longer periods of time on the display. In other design studies, important words would remain in the visual field simultaneous with the main text. In this way, she found she could maintain a high word rate (100 msec. per word) while enhancing comprehension of the main content ideas.

Suguru Ishizaki's recent dissertation [ISHIZAKI95] proposed a model of design in which words are actors that engage in a performance for the reader. His focus is on the interactions between words on the display. Both his and Wong's work suggest that the one • word • at • a • time • reading used in this experiment may not be the most effective method of presentation.

The experiment consisted of two variables - tempo (word duration) and design (graphic representation of voice). Each of eleven test subjects were shown fifteen different presentations of the same text, one for each combination of the three tempos and five designs. Subjects were asked to interactively vary the speed of each presentation until they found the fastest comfortable speed. In addition, each subject was asked to rate the quality of each presentation.

There were three variations on tempo that were presented: one (T1) based on a spoken "performance" of the text, one (T2) which gave each syllable the same duration and included pauses for each punctuation mark and change in speaker, and one (T3) which gave each word equal weight without any pauses. In musical notation, the three tempos looked like those in Figure 28. In the notation shown, a section of the test data shows that the same text required 48 six-

FIGURE 28. A portion of the text shown in music notation.

teenth notes for the spoken tempo, 37 sixteenth notes for the even syllable tempo with between speaker pauses and a mere 25 sixteenth notes for the even word tempo without pauses.

There were five variations on the design of the text which were matched to each of the three tempos described. The generic design (D1), did not give any indication of speaker or emphasis. It merely presented the text in the center of the visual field. The text was set in Century Schoolbook in white on a black field. The height of a capital letter was 13/8 inches and viewed from a 36 inch distance (2.19° angle subtended). The words were centered rather than left or right justified, which allowed the reader to maintain a one fixation point, reading all but possibly the longest words without any eye movement [O'REGAN84].

The second design (D2) is identical to the first, with the exception that the words of each speaker (mother, Pierre, and narrator) are drawn in a slightly different color. Three colors were chosen that were as close as possible to the default white while still clearly different from each other. One of the important qualities of dynamic typography is that the position of a word in the visual field can have more meaning than simply indicating the order in which the text should be read. Position can add nuance to the literal meaning of a text and in particular can be used to set up a relationship between different speakers. In the third design (D3), the mother occupies a position above and on the opposite side of the screen from Pierre. The narrator occupies a position in between and below the two characters. The narrator is also rendered in an italic face, since its voice is not as literal as the two people. Figure 29 shows a transparent overlay of several moments in time to indicate the relative spatial locations used in the experiment.

One of the problems encountered when placing text in various locations over time is that one cannot always be sure that the reader's gaze is in the best position to read a word when it appears. Although a change in location more clearly showed a change in speaker as well as the relative size and power differences in the mother/child relationship, the refocusing of ones attention might slow down the reader, or at least cause the first word of an utterance to be lost. The fourth design (D4) attempts to solve this problem by preceding each change in location change with a with a small fixation point at the center of where the text would appear. So, if there was any reduction in speed between D2 and D3 caused by constant relocation of gaze, D4 would counter that effect.

FIGURE 29. Mother asks "Would you like to stay right here?".Then Pierre, shorter and to her left replies, "I don't care!".The narrator occupies the space between and below them.

The final design (D5) looked at the visual display of another important part of the rhythmic structure of the text. When read, there is generally a strong stress mark on the rhymed syllables. In this design, the stress is indicated by a change in the size of the displayed text. Stressed words change size as they are displayed, expanding outward until they fade away. It acts as a visual analog to the sonic experience of the stressed word, and helps to emphasize the rhyming pattern. Each of the five designs increased the information flow to the subject.

We saw a strong effect from the variations in tempo on both speed, which increased almost 50% ($p < 0.0001$) between T1 and T3. Surprisingly, they different designs tested had little effect on the speed chosen ($p = 0.57$), although there was an effect on the quality reported ($p = 0.002$). Subjects reported a 12 percent increase in quality from the generic design as more graphic information was provided.

The experiment shows that the use of spoken rhythms does not appear to improve the reading speed or perceived quality of the text. In fact, reading speed decreased when the spoken rhythms were used.

The results of this experiment suggest that designers can increase the visual complexity of temporal typography without fear of reducing their audiences ability to read the text quickly. Increasing the amount of visual information that was displayed did not confer any advantages in the speed of reading, but did increase the subjects rating of the quality of the presentation.

However, the high variability of preferred speed between subjects, all of whom were college students used to reading moving type on computer displays, suggest that it may not be possible to find a common preferred speed for the design of temporal typography. Those who design such systems should account for this variance by giving more control to the reader in the pacing of the information.

As technology improves and designers create new, dynamic forms for communication, they will have to better understand the perception of dynamic visual events. Hopefully, through that understanding will come a new visual language for temporal typography. The next section looks at the aesthetic qualities of dynamic typography.

Stream of Consciousness

The Stream of Consciousness project, developed in collaboration with Tom White, attempts to bring the computer into the garden in harmony with stone, water, and plant materials. The computer is used to drive a video projector, creating the illusion of text floating on the surface of the water as it flows through the garden. This relaxing computational environment lends itself well to several open ended active and passive modes of interaction.

FIGURE 30. The garden with rocks, vegetation, flowing water carrying words, and an interface mapped onto the large pool.

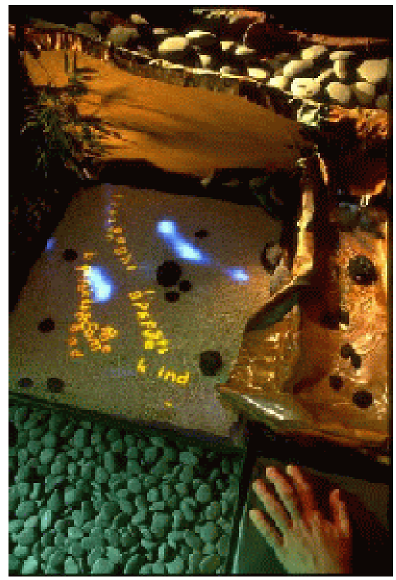

The interactive poetic garden is literally a fountain of words. Water flows briskly down a series of cascades into a glowing pool. A tangle of words projected on the surface of the pool float like leaves in a stream. Sitting on the edge of the pool-but without getting your hands wet-you can control the flow of words, blocking or stirring them up, causing them to grow and divide into new words that are eventually pulled into the drain, then pumped back to the head of the stream, only to tumble down again.

Although it measures only six feet square, our garden contains all of the elements of a classical garden: flowing water, river stones, bamboo, and a bench for people to sit. The garden design is based on a square recursively divided into a series of smaller squares, a design which can be traced back to the earliest formal gardens of Persia [MOORE88]. Water enters at the back of the garden and cascades down a series of pools until it reaches a large square pool. This larger pool is lined with crushed white coral and here the water moves slowly until it spills out the back edge. Words appear to tumble down the rocks along with the water, calmly pull themselves through the shallow pool, and then magically reappear at the top of the stream along with the water. The words mimic the physical behavior of objects floating in a real fountain. The person sitting at the bench can interact with the words through a special hand interface letting her stop the word flow, push and pull words, and over time change the content of the words themselves.

These physically modeled words are projected from above onto the rocks and coral. The computer computes the word image as well as managing the camera based input device, which lets the person control the word flow.

There are obvious and subtle differences in fluid dynamics and typographic layout. It was our goal to have the words appear to flow naturally along with the water as if they were leaves floating downstream. An important design concern was to have the words maintain an orientation and interletter spacing such that the words are legible most of the time. We decided on a mass/spring system in which each letter of the word is a point mass connected to its neighbor letters with springs. Additionally, the first and last letter of the word is connected with a separate spring that pulls the letters of the word into a line. A force is applied to the first letter of each word that propels it through the stream. This force is defined as a stable vector field that is designed to match the real water flow through the garden as well as appropriate forces from the input device. Of course, this complexity was completely transparent to the person arriving at the garden. To him, the words were naturally flowing and bouncing through the pools along with the water.

FIGURE 31. Three fingers hold back a stream of words

Initially, we wanted our garden to be a place for meditation. We also wanted people to be able to affect the overall content of the words as they circle through the garden, as well as influence the dynamics of the digital content. We called our installation Stream of Consciousness because we hoped to evoke the fluid contents of conscious memory and shifting focus of attention with word association [STILLINGS95].

Through the hand interface, a person can reach into the pool of words and create a blue aura behind them. The words are repelled by the hand, making it easier to create blockages in the flow (Figure 31). But if the person presses directly onto a word so that the glow is directly behind it, it begins to swell larger than the other words. Eventually it bursts into two words, the original along with a related word. If the hand is not moved, more related words will continue to be generated from the seed word. As the pool circulates, old words are removed, so that over time the words in the water are the words that have been chosen as interesting.

It is not difficult to affect words given the rich amount of information provided by the specially designed interface. The forces on the word slide away from the pressure gradient, and the pressure information is presented directly as the blue glow for instant feedback. Giving any word the ability to divide into a rich set of related words was trickier, but we were able to implement this functionality thanks to resources of Princeton's WordNet project [FELLBAUM98]. For example, the word stone might give birth to such words as: pumice, grit, substance, and pit.

The overall effect of the garden was foremost as a quiet, contemplative space. Over time, several hundred people experienced this interaction, and the response was warm and enthusiastic. We were also pleased to see the many unique experiences that people left with because the interaction was so open ended. Some people were content to passively watch the words, others world repeatedly damn up the words into clumps and then release them, and others would attack the words so that they divided out of control and filled the water with hundreds of words. Even very young children were able to explore the water and stones and the "lights" which shone on the water. The computer is an integral part of the design, but it is also only one element among many. The interface allows a wide range of open ended interactions and creates a harmonious environment that is often lacking in computer based environments.

Applications of dynamics to information design.

The previous examples have shown how dynamic typography can greatly increase the expressive range available to the designer. This final section looks at how we might apply these expressive techniques in a meaningful way to traditional information design problems.

In the MetricViews project, another collaboration with Yin Yin Wong, a Java applet creates multiple spreadsheet views and presents them in a singular dynamic context. Instead of publishing a large number of discrete spreadsheets and graphs, each of which would show a particular subset of information, we designed an interactive display with which users can construct a wide variety of views into the information space. A structured information tool enables users to construct meaningful views and shift between them without losing context.

FIGURE 32. The result of switching from the tabular view to the graph view.

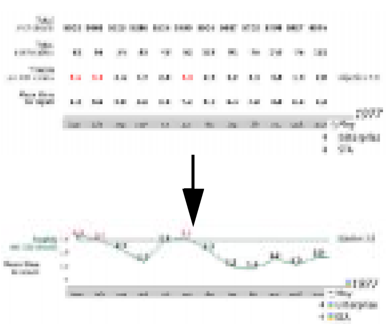

The interaction model for MetricView was simplified by allowing only progressive changes to the visual display. Users construct metric views by changing one dimension at a time resulting in views that flow from one into another. This method avoids jarring changes to all the mappings at once which would result in the loss of visual continuity. Animation of typographic elements from one view to the next ensured that the user could literally see where information was coming from. If the view was changed to a new spreadsheet, those elements which were common to the two views would animate into their new position, unused elements would slide off the display to the right and new elements would slide in. The user could easily keep track of what had changed.

Unlike most systems, which present tabular and graph views as separate objects, MetricViews treats them as the same object with two representations. At any time, users can transition back and forth between these two representations. A crucial aspect of this design is the animation technique we developed to clearly depict the transition from tabular to graph view while maintaining visual continuity. Numbers in the tabular view animate from their position to their graph position, visually appearing to “morph” between the tabular and graph view (Figure 32). Since the numbers moved smoothly, each number appeared in the graph was seen as the same object as it had been as an entry in the spreadsheet. Animation, in this case, revealed the source of the data points and therefore helped the user to quickly read the structure of the graph.

Almost all the animation in the Talmud project was under the direct control of the user. This sort of motion is very easy for the user to read, as long as the interaction loop between the controls and the graphics are fast and responsive. The user would turn a dial, for example, and an element would rotate in response. Because the feedback is immediate, and the user can always return the dial, and therefore the information, back to its original position, there is little training required. Moreover, the user feels in control of the interaction, because he can effect change in the display when and how he wants. Unfortunately, if you make a control for every possible movement of each and every object, you soon will have a control system which overwhelms the information itself. By building dynamic constraints between information objects and limiting controls to those dimensions which are meaningful, we can keep the interaction manageable without overly restricting the user.

FIGURE 33. As the user slides the control to the right, the Levinas text enlarges to the right, while the Talmud scales smaller and slides right to make room.

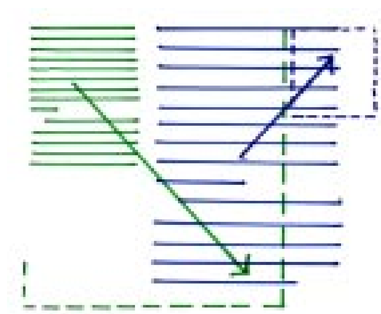

Simple constraints were easy to implement because of the message passing system used by the acWindows system (for a more detailed description of the acWindow system, please see Appendix A). Each acWindows object can both send and receive messages. So, for example, if an object receives a message from a control to change scale, it can, on the receipt of that message, send a message to another object to change its position. Any of the messages can also include a duration, which will direct the recipient to smoothly animate from its current state to a new state in a specific amount of time. Because these animations were built into the system at such a low level and were very easy to use, it meant that the designer could choose to animate any change (e.g. color, focus, position, etc.) without any extra work. While these simple animations meant that transitions were smooth and continuous, there was need for more complex motion.

As discussed in Chapter 4, the Talmud project used a physical spring model for maintaining spatial relationships during interaction. The goal was to keep certain texts next to each other even as they changed in scale and position. As the user interacted, they would introduce energy into the system and the springs, acting in a damping field, would move to reduce the energy. The users could, through the motion, literally see the underlying constraints and how they were solved. Initial designs damped the springs on slightly, so it would take a considerable amount of time (10 or more seconds) for the system to reach equilibrium. This motion was felt to be far too distracting, so the damping was increased to nearly cancel out the stiffness of the springs. The system would now reach equilibrium in about a second. This brief motion was still enough to reveal to the reader the underlying constraints.

Animation was also used when the reader followed a link in the text to a particular verse of the Torah. If you think about how links work in most web applications, clicking a link wipes the current web page from the screen and then (slowly) paints the new web page. The starting location is gone and there is no indication of how it might be related to the end result. To find a verse in the Torah, the text would smoothly slide from the current verse to the new book, chapter and verse. This worked best when the view allowed both the start and end points to be visible so that the path from one to the other could be easily traced. If the camera was close in and the verses far apart, the animation blurred into a meaningless smear, like the view out a window of a speeding car. In his thesis, Tinsley Galyean [GALYEAN95] proposed a method for planning path through virtual spaces, such as a museum. While some of these techniques are applicable to a more abstract and purely typographic space, future work might examine how to plan paths through information spaces that are most meaningful and keyed to the legibility of typography.

This chapter has described a number of techniques for creating expressive movement on typographic forms. Previous experiments in dynamic typography inspired the use of motion in expressing a musical voice in the Brain Opera. Perceptual experimentation showed how different rhythms and visual forms effected the legibility of serially presented text. Physically-based motion gave compelling movement to liquid typography in the Stream of Consciousness. Finally, we saw the application of movement to a variety of information design problems and the relationship between interaction and the movement of typography on the screen.

The following chapter will examine in the interaction models and the design of control systems for information systems. It begins by looking for ways to replace traditional graphic widgets with physical controls.

CHAPTER 6

tangible issues

Interaction design should be intimately connected to any information design process. The physical interaction with the display and its controls has the potential to deepen our relationship to information and is perhaps the main improvement that digital media can give to the traditional book. This chapter emphasizes the physical over the graphical interface and examines the relationship between the display surface and the interaction device. By simultaneously designing the graphical display as well as the physical controls it is possible to completely rethink the interaction model and make tangible abstract typographic constructs.

If computers are tools for manipulating information, they have been notoriously poor at using the hands (and bodies) of the people who use them. By engaging the hands of the user in real space, it is hoped that a more practical, productive and fluid interface will be created.

The approach taken was examine specific information design problems and then to create sketches which examine how a new interface technique would enable richer interaction with an information space. Each new technique looks at the boundary between the real world and a virtual world created by the computer. By connecting aspects of the virtual world to real world objects, we allow the user to literally feel his way through the computer generated world. The objects which inhabit our working spaces should be legible in their function, provide clear feedback to the user and be flexible in their application.

People are used to holding books, newspapers and even scrolls in their hands and moving them about. When the user's hands and body are engaged in an appropriately physical manner, we can bring the user closer to the virtual world contained in the computer. The virtual information on the display can appear to be held directly in the hands of the user.

Physical objects can become powerful interface objects if they are legible in their purpose and use. Our tendency to invest meaning into simple objects allows them to be used as symbolic tokens for the manipulation of abstract information. And, by combining the interface with the display we can give meaningful form to the space which resides on the other side of the screen and its relationship to the body.

If you hold a hammer in your hand, everything in the world begins to resemble a nail. Computer interfaces do not have to be monolithic. By engaging the user's hands and body in appropriately meaningful ways, we can bring the user closer to the virtual world contained in the computer.

Early experiments

Most computer interfaces function through the manipulation of objects which control the graphic display, even if they are as simple as the handling of a mouse which controls the position of a graphic cursor. This research is concerned with understanding the variety of relationships that are possible between the held interface and the virtual world within the graphic display. We can classify such interfaces according to the symmetry of the physical action and the graphic reaction, the number of dimensions that are encoded in an interface and the spatial relationships between the two spaces.

A flexible toolkit of interface components, including 6D position/orientation sensors, portable displays, and LEGO components (bricks plus touch, resistance and rotation sensors) was used to quickly design, build and evaluate a number of interfaces between real and virtual worlds. A wide range of studies were undertaken which demonstrate a variety of techniques. The relationship between the real world and the virtual world can be completely symmetric, with the virtual world reflecting objects and actions in the real world and augmenting them in some way. It can be more abstract, where real world objects are used in a symbolic way to manipulate abstract information. And finally, we can completely integrate the display object with the interaction model. By connecting aspects of the virtual world to real world objects, we allow the user to literally feel his way through the computer generated world. The objects which inhabit our working spaces should be legible in their function, provide clear feedback to the user and be flexible in their application.

Symmetric interfaces

This work evolved out of a general frustration in working with three dimensional information spaces [SMALL94, SMALL95]. Although the computer was capable of displaying dynamic information landscapes, the spaces proved to be enormously difficult to navigate. By constructing a simple virtual world of LEGO objects and designing LEGO objects for both navigating and manipulating that world, we were able to study interface techniques that would then be applicable to more complex and abstract environments. LEGO bricks are rigid, manufactured to precise specifications and can be configured in innumerable ways. Simple LEGO elements can be combined to form complex and satisfyingly solid models which can be tracked with surprising precision. These properties encouraged us to redesign often and quickly while engaging in ever more complex tasks.

FIGURE 34. The virtual LEGO project

The virtual LEGO project consists of a real-world LEGO scene and a three dimensional computer model of the same scene (Figure 34). Objects in the virtual space can be controlled though interactions with the physical LEGO. The LEGO scene consists of a police station, parkland, a coast guard dock, water and a variety of people and vehicles. It sits on a small table in front of the computer display so that one or more people can easily see both the real and virtual LEGO scene. The interaction model is quite simple- there is a one to one correspondence between the worlds, with the computer acting as a sort of mirror which both reflects actions in the real world and augments them in meaningful ways.

Navigation is accomplished through a special brick which can report back to the computer its exact position and orientation. This was implemented with a Flock of Birds six degree of freedom tracker, housed in a custom LEGO brick. Data from the sensor is sent to the computer which transforms it from millimeters to LEGO units (one LEGO unit is equal to the distance between two studs). The virtual model is constructed in the LEGO coordinate system by specifying the LEGO elements in a simple text file.

FIGURE 35. The view from the tracker with the synthetic view on the display.

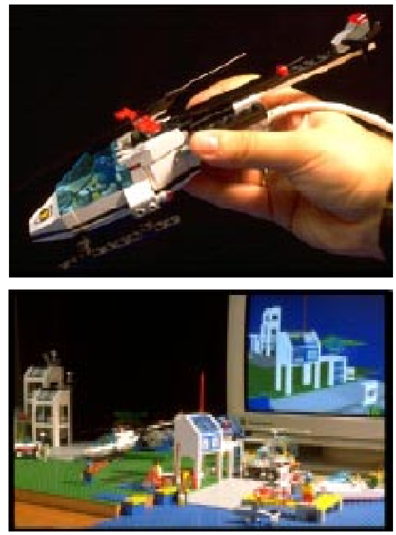

The position and orientation of the virtual camera is set to correspond to the position and orientation of the tracker (Figure 35). Since the dimensions of the real world exactly match the dimensions of the virtual world, the display looks like the view from an eye positioned on the sensor. Moving the tracker to the right by four bricks shifts the graphic display by the same amount. While this system is quite simple, it proved difficult for beginning users to use the tracker object to specify a particular view. Often, the brick would be held in the hand in such a way as to turn the virtual world side-ways, backwards or upside-down. An attempt was made to solve this problem by drawing an arrow on the tracker indicating the direction of view. The arrow was generally ignored. Finally, the tracker was built into a LEGO model of a helicopter. This had several advantages which slowly became apparent. First, the helicopter has an implicit orientation, so it is easy to tell front from back, up from down. There is a strong inhibition to turning the helicopter upside down as this would cause it to “crash”. Second, the helicopter fits into the user's hand more easily than the brick alone. This is not surprising because the helicopter is designed to be used in this manner (hold helicopter in hand, then run around the room making helicopter noises). Finally, the helicopter offers a simple explanation of how the virtual view is being produced: the display shows the view pilot's view. We found that with this simple mental model, users could more easily put themselves in the LEGO scale. By building a tracker into a LEGO element, we had been able to reconfigure the interface on the fly without recalibrating or losing precision.

One inherent benefit to our system over head mounted display or immersive systems [FIENER93], [WELLNER93] is that in our design, the user can perform tasks in the virtual space while engaging in activities in the real world. This is most similar to systems such as Hinckley [HINKLEY94], in which props are used for neurosurgical visualization. In both cases, meaningful graspable objects are used to specify parameters in a virtual space.

FIGURE 36. The real and virtual LEGO gear trains.

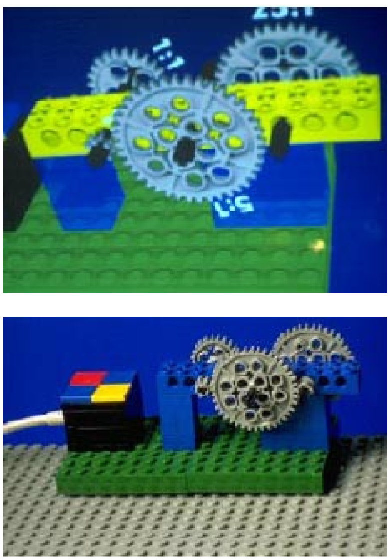

Simply moving through a space, while important, was not the only goal of the LEGO system. We wanted the user to be able to perform tasks in the virtual space. The task was a simple construction job - follow the instructions and build a gear train. Two LEGO switches moved through the instructions forward and backward and the user could build the model step by step. This is similar to the normal LEGO building procedure, where drawings of each stage of assembly are matched by the builder. In the first step, a sensor is attached to the base plate of the model so that the computer can track the construction and match the graphic of the instructions to the growing assembly (Figure 36). In this way the construction is always in the same orientation as the instructions. The user can pull in the helicopter to inspect de-tails of the assembly while he or she is working. Because the graphics and the physical model are so tightly linked, you lose the sensation of referring back and forth between the instructions and the work. The instructions follow the work around and are one and the same with the model.

The biggest advantage to using our system instead of printed instructions is that our virtual LEGO can display the dynamic characteristics of the model. As the gear train is constructed, the gears themselves can move and the function of the mechanism begins to be revealed. Although the printed image fully specifies the position of each element, it does not show how the gears work together. For the builder, this means that it is difficult to be sure how the machine should function and how the gearing mechanism functions. In order to have the virtual construction mimic the functionality of the real construction, a LEGO rotation sensor is included in the model. When the physical gears are turned, the computer can sense this and rotate the virtual gears. If the model has been assembled correctly, then the turning of the virtual gears will match the physical ones. In this way, the builder not only sees that the gears are in the correct location, but that they are functioning correctly. The virtual gear train displays the gear ratio alongside each gear, giving a numerical quantity to the gear's function. Thus, using the virtual LEGO system, the user can see not only that the gear at one end turns the gear at the other end more slowly, but exactly twenty-five times more slowly.

In traditional user interfaces, the user interacts by moving a mouse corresponding to an image of the pointer on the screen. By selecting images of buttons and objects with the image of the pointer, the user can make changes in the virtual world. In this design study, we wanted to directly manipulate real-world objects in order to affect our virtual system. One of the advantages that virtual LEGO bricks have over real ones is that it is possible to change their color as needed. In a simple painting task, real world "buttons" were used to select colors and virtual LEGO bricks were "painted" by touching their real world counterparts. Two by two flat LEGO bricks in different colors were used as a palette. The user could select a color by touching the plate with a "pen" which was snapped onto the 6D LEGO tracker. Since the location of the sensor in the tool is known exactly and the model is rigid, it is simple and reliable to compute the exact position of the tip of the "pen". Once a color is selected, the virtual bricks are colored by touching the real bricks with the pen. This is similar to the system used by Fitzmaurice and Ishii [FITZMAURICE95], except that we augmented the generic input device so as to enhance its use and the legibility of its function. Simple additions to the basic tracker enable it to assume new identities and uses.

A more elaborate painting system was developed in conjunction with Nike, Inc. In this case, our goal was to enable shoe designers to paint directly onto the three dimensional surface of the shoe mold, called the "last". The last is a real, three-dimensional model of the interior of the shoe. The last is used as a model in the design process and the shoe is built up around it. One problem faced by designers creating the graphic elements for the shoe is in mapping them onto the three dimensional surface. Traditionally, graphic elements are drawn on a flat tablet and mapped onto the shoe surface in a 3D modeling package.

FIGURE 37. Shoe painting system

By tracking a stylus and a last, in much the same way as we tracked LEGO bricks, we were able to develop a painting system wherein both the brush and the last could be moved in space (Figure 37). Whenever the brush touched the surface of the last (according to the position sensors), virtual "paint" would flow. In this manner, graphics could be built up directly on the last. Because of the tracker's precision and the one-to-one mapping between real and virtual space, no artificial collision detection or force feedback were required.

Another experiment in the symmetry of physical controls and the graphic display was inspired by a model of a hyperbolic paraboloid in the Collection of Historical Scientific Instruments at Harvard University. Built of brass and wood, with strings held taught by lead weights, it allowed students to create an examine a variety of hyperbolic paraboloids.

We built a smaller scale model with potentiometers mounted at key mechanical points. The change in resistance at those points is used to determine the state of the model as users modify it in real time. On the computer display, a synthetic model is shown along with dynamic equations which show how variables such as volume and surface area change as the model changes (Figure 38).

FIGURE 38. Hyperbola model

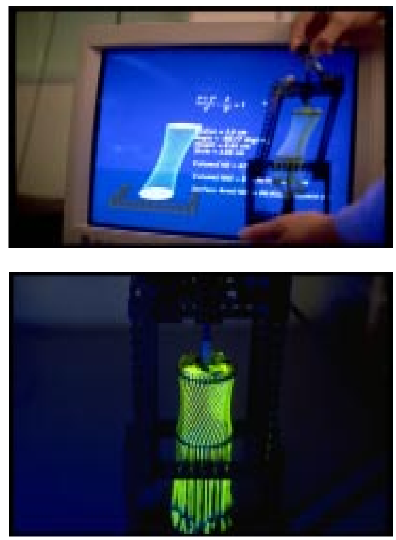

The symmetry in this interface is clear - both the real and virtual model match so smoothly that the user feels as if they are the same object. What was important in designing the symmetry was deciding which were the most salient variables which would be controlled by the interface. Mechanically, there were three variables which could be measured, but they did not correspond easily to the variables in the classic form of the mathematical formula. However, the mechanical model set up the users understanding of the shape and users found that they could make the shape change as they expected with that interface. It was much more difficult to make desired changes to the shape by modifying variables in the mathematical formula. Future work on this project will include using better graphic representations of the underlying math in the virtual space which should clarify the connection between the form and the equation.

Abstract interfaces

In each of our previous examples, the virtual world nearly matched the real world, so it was easy to understand how they were related to each other. This model breaks down when one wishes to manipulate that which has no counterpart in the real world. Type has no mass, no thickness, no absolute dimension and is only meaningful when the letters can be perceived in particular combinations. Nonetheless, words must be written, manipulated, analyzed, stored, retrieved and read. Abstractions such as type are fundamental to our information culture.

In Chapter 3 we discussed the use of three-dimensional typography in the display of the complete works of Shakespeare. Although these new rendering techniques allow many different views of a large scale text, the visualization is only useful if the user can easily navigate about the text. A number of new methods of navigation were developed to address this problem. Most current interface paradigms (windows, buttons, mice) are based on a two dimensional screen. A three dimensional model requires new kinds of controls that allow for easy manipulation in space.

FIGURE 39. LEGO controls.Viewing Shakespeare using a helicopter and brick for navigation.

The first approach was to use one position/orientation sensor to act as a handle for the text and another as a virtual camera. By holding the first sensor in one hand it is possible to easily control the position and orientation of the text. A graphical representation of the “handle” is displayed along with the text to help orient the user. The other sensor is built into a small LEGO helicopter that can easily fit into the other hand. A helicopter was used because it has an implicit sense of pointing and an implicit orientation (rotor above, tail behind, windshield in front). The helicopter pilot “looks” at the handle and determines the position of the virtual camera in the information space. In addition to these controls, some simple LEGO machines were constructed that allow the user to zoom in and out, and to position the text horizontally and vertically relative to the virtual handle (Figure 39).

The second approach consisted of a small LEGO stage. The positioning gears were built into the structure of the stage in two orthogonal orientations (Figure 40). To move the text left and right, the wheel above the stage is turned left and right. To move up and down, the wheel on the side of the stage is used. The angle of the text is controlled by rotating the entire stage assembly. The text moves in space as if it were locked to the stage. To rotate the text around the vertical axis and reveal the footnotes, for example, the user rotates the stage about the vertical axis. Although it is possible to pick up the stage and rotate it about any angle, it is more confortable to keep the stage, and therefore the text, upright.

FIGURE 40. LEGO stage

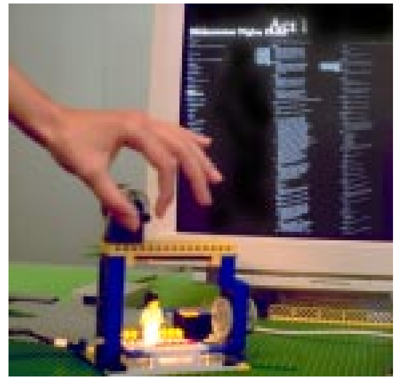

The last approach was to merge the display with the navigation controls. A small display is held in the hands and moved about in real space to navigate the plays. This led to other experiments which are described in the next section.

In addition to this interface for navigation, we wanted to give the users tools for filtering the text. To this end, we soldered resistors to the feet of LEGO figurines, and used them to represent Shakespeare's characters. To select different characters in the play, the corresponding LEGO actor is placed on the stage. The system then recognizes each actor when his or her figurine is snapped into place by its resistance. The footlights come on and the text for the character is "lights up". In the virtual LEGO example, the virtual objects were tokens (or symbols) for the real ones, here the LEGO people become symbolic abstractions in relation to the text. Their lack of unique features encourage a fluidity in the meanings associated with them. Just as an actor can take on a role in the context of a performance, a LEGO figurine can assume a certain identity relative to one task and then shed that identity for others as the need arises.

Merging the interface and the display

In the virtual LEGO and Shakespeare projects, the graphics display is separate from the interface devices. The user manipulates objects out on the desk and views the resulting changes on a display surface which has no clear relation to the work space itself. This means that there is always a feeling of disconnect between what the hand does and what the eye sees. To resolve this the display itself can become the input device. The display is held in the user's hands and manipulated to control the virtual space. This was initially achieved through a slight modification to the virtual LEGO system. The helicopter was replaced by a small, flat panel display and the user saw in the display what was hidden behind it. The conceptual model was straight forward and the feedback immediate. This design was quickly adapted to more abstract virtual spaces. In the introduction to this thesis, we looked at the design for a scroll in which the user could move the display back and forth to reveal different sections of the painting. This was idea was expanded into multiple dimensions in the design for a daily newspaper.

Imagine holding a newspaper. You move the paper up and down, closer and further from your eyes, and manipulate the pages as you read your way through it. Your own body and the physical qualities of the paper medium determine how the information is designed and how you find your way through it. In this design we erase as much as possible the division between input (control) and output (feedback). The display itself is aware of its location in space and relative to the users body. The graphics can appear to slide beneath the display as it rolls back and forth, or be pulled across its surface by gravity as the display is tilted.

In traditional computer interfaces the graphics display is separate from the input devices. Although head mounted displays or immersive systems provide a well integrated experience, the user cannot perform tasks in the virtual space while engaging in activities in the real world. Our approach is similar to systems such as Fitzmaurice [FITZMAURICE93], in which a palmtop display reveals virtual information associated with real objects in space. His system used gross location to bring relevant information to the display. Our system provides more refined gestures, such as tilt, to let people browse a large information space.

While respecting the physicality of the scroll was important to that application, many tasks neither require, nor benefit from such a literal interpretation. In the newspaper study, a mock up of a news reading device was modeled in balsa wood and acrylic (Figure 41). It contains a portrait display, a single thumb button and a storage area for news stories that have been "downloaded" into LEGO bricks. Users navigate the paper by engaging the thumb button, which acts like a clutch, and moving the display relative to their own body. Several different motions are recognized. Tilting the paper up and down scrolls the text vertically, tilting left and right moves the text horizontally, and pushing the whole display away from or close to the body zooms the text in and out. We studied several functions which related the tilt of the display to the movement of the text. When the display is tilted around the X axis, "gravity" pulls the text up or down across the display. This was seen as the primary motion in normal use as it corresponds to reading down a vertical column of text. Several equations, such as: velocity $= \theta$, acceleration $= \theta$, velocity $= \theta^2$, and velocity $= \theta^3$ were tried. The last equation proved to be the most effective, because it allowed the user to hold the text still at any given point, scroll very slowly one line at a time with a small tilt and still be able to scan past many stories in a second with a greater angle of tilt. Even this rather complex mapping is simple for users and requires no specific explanation. In a

FIGURE 41. Newspaper design prototype. The user engages the clutch with his right thumb and moves the display about in space to navigate around news stories.

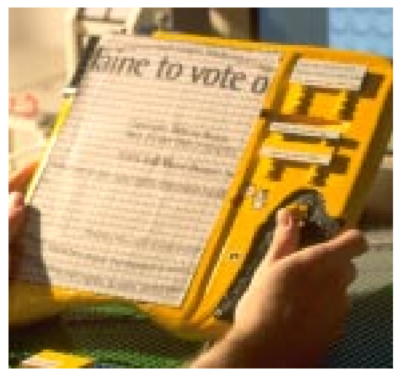

short period of use, the user can browse and read easily without much conscious thought about navigation.

In addition to scrolling up and down one newspaper column, the user can slide the text left and right to move to parallel columns of related news stories. Pushing the paper away from the body results in a smooth zoom out, giving the reader and overview of many news stories. The head lines becomes more visually prominent and the body of the news story becomes lighter to show general massings of text. By pulling the display close in to the body, the user can zoom back into a particular story.

Although just a balsa-wood prototype, the hand held newspaper demonstrated the feasibility of a display which was aware of its position in space. Later, when a lightweight, flat display became available, this interaction technique could be combined with others shown here to create a unified interaction experience.

Physical design for the Talmud

Traditionally, the study of Talmud is specific about the environment and the tools used. Several volumes must be available and right at hand. The workspace should allow two people to study and work together. A tool called a *yod* is used to point out passages in the text. Everything is arranged to facilitate study and discussion.

In designing an electronic analog for Talmud study, I wanted to bring together as much as possible what had been learned about form, layering, navigation, movement and interaction. The graphic tools devised were discussed in previous chapters, but it is in the interaction where everything comes together.

I actually began by using on-screen graphics to interface with the evolving tool. It was immediately apparent that there was a problem with this approach. There was so much information to display that I was using every square inch of the display surface. As I began to add more and more controls - and there was no lack of ideas for new widgets - they began to consume more and more of the screen real estate. I hoped to use the same scaling technique which allowed more and more text to dynamically fit into a smaller and smaller area to rein in the interface. When the mouse approached the interface widgets, they would expand until they were at a usable size. The user could interact and then, upon leaving the space, the controls would shrink out of the way. This solution was abysmal. When the controls were large enough to use, they obscured the information they were meant to control, and even shrunk down to a postage stamp, they didn't seem to fit into the same space as the information.

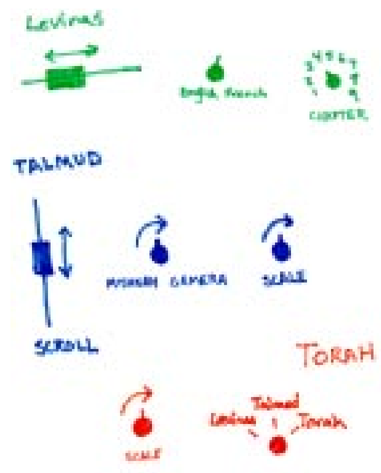

FIGURE 42. A sketch of the early graphic controls for the Talmud.

The solution was to remove the controls from the display entirely and put them elsewhere. A small set of physical controls, such as dials, buttons and sliders would be able to replace the on-screen menus. Rather than create a very specific interface, a toolkit of generic devices was developed. These few controls - a smooth dial, a discrete dial, a slider and a 4-key keypad - could be readily manufactured and used interchangeably to generate input messages to the acWindows system. This meant that the interface could be changed nearly as quickly as the previous graphic interface. The ability to rapidly modify the prototype was crucial to keep up with the tight critique and design loop used to develop the project.

FIGURE 43. The set of input controls.

FIGURE 44. The final interface design with the laser cut panel.

Each control would produce a variable resistance which was converted to an digital reading by the LEGO Control Lab interface. The interface communicated the data back to the host computer via the serial port up to 100 times a second. The actual data was only sampled at the display rate, which varied between 5 and 30 frames per second. A panel which fit over the display was laser cut to accept the various components. As the design changed, either in the specific components used (slider, dial, etc.) or simply in their layout or labeling, a new panel could be cut. The laser both cut through the material and etched text and graphics. After testing several kinds of plastics, the final material chosen was a two-ply mat board. It was very inexpensive, cut cleanly, and gave the finished panel a softer touch. The panel design was done with a standard drawing package and new designs only took about four minutes to cut.

The controls were divided into three main groupings - one each for the Levinas, Talmud and Torah texts. The following chart shows each control and its use:

Grouping	**control**	**type**	**function**
Levinas	dial	continuous	English to French
Levinas	dial	12 position	section selector
Levinas	slider	continuous	scale
Talmud	slider	continuous	up/down selection
Talmud	dial	continuous	rotation, *gemera* to *mishnah*
Talmud	dial	continuous	scale
Torah	dial	continuous	scale
	keypad	4-key	layer selection

The Talmud is traditionally read by two people so that it can be argued and debated. The controls were designed to facilitate this style of polemic. While discussing some fine point of logic, either of the readers can grab a control and modify the visual relationships between the texts in order to support their argument.

In addition to the eight controls, two sensing devices were built into the interface/display. First, a positioning sensor was used to report back the absolute position and orientation of the display. Second, and electric field sensor was used to locate the users hand above the display.

The positioning sensor was used in much the same way as the hand-held newspaper described earlier. In this case, the user would select the particular layer of interest and then move the display. The selected layer would be locked into its projected position in real space and the display surface could then me moved about the information. So, for example, you could lock down the Talmud and move the display to the left. The Levinas and Torah layers would move with the display, but the Talmud layer would slide right, appearing to remain in the same location in real-world space. This technique worked well for both position control and changing the orientation of a particular layer. Because of the weight of the display, two handles were bolted to the interface to make motion easier.

The idea for this came out of looking at people using books in a library. There would be a rough spatial mapping of information on the desk. Notes would usually occupy the central position and various books would be arranged in concentric circles with the most used volumes closest to the user. Similarly, the user of the digital Talmud could locate various texts around the area of the display and access them

Critique

Bill Mitchell
You want to read back and forth among multiple texts; you want to keep them all visible like having a bunch of books open on a desk. But you want to be able to shift attention and focus. Start with a situation; literally, a person with three books...

Think of reading as a social act and as an interactive thing rather than just an individual. You want to lay out a typology of reading situations...one reading situation is the rabbi or priest standing there and reading out, the word from on high...another one is the individual reader, the individual scholar...another one (the really interesting one) is the discussion where you are using the texts to structure a discussion. You need a different interface for each one. a different spatial arrangement. You are engaging the text in a different way. To lay that out as a starting point and say that you are engaging interaction of a couple of people in discussion around multiple texts is the most interesting.

It gets you away from the business of the *personal* computer world. *Personal* tells it. It is non-social, just you and the machine.

FIGURE 45. A typical library work session (MIT).

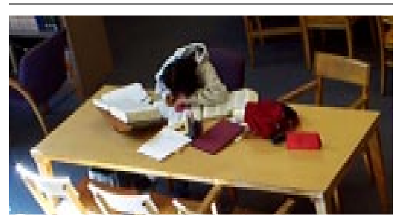

through the movement of the display surface. Other possible uses of this technique include placing divergent arguments at different locations, so that two scholars could literally jump back and forth between their individual points of view during the course of an argument.

Electric field sensing technology, developed by Josh Smith [SMITH99], was used to sense the position of the users hands over the control surface. The technology was not quite precise enough to exactly track, for example, a fingertip over the display, but it was perfect for detecting when the user’s hand came into the general area of a set of controls. This was used to anticipate which layer the user wanted to interact with. As the hand approaches a particular control, the layer affected by that control will come into focus and the other layers will move into the background. This simple technique was very effective, because it freed the user from explicitly controlling which layer was in focus. By the time you thought to manipulate a layer and your hand had touched the appropriate control, that layer was already in focus. By sensing, even in a rudimentary manner, the intentions of the user, we can greatly facilitate the interaction.

In this chapter we have looked at the physical interface in relation to information design. Physical objects can be used to control their virtual counterparts. Mechanical controls can allow the user to directly touch abstract information in the display. Also, by tracking the display itself, we can use it as a navigational device. All of these techniques were brought together to create a unified interface for manipulating the Talmud.

CHAPTER 7

conclusion

The goal of this thesis was to create a compelling vision of how computers can redefine the paradigm of the book. It is not simply about producing a series of experiments or design sketches, but to use them to ask fundamental questions about how reading, writing and expression are changing and about how written language will evolve in response to computer technology.

The current paradigm of a computer with a fixed, heavy display, a keyboard, and a mouse is not nearly as comfortable to use as a simple book. Nonetheless the ability of the computer to gather, analyze, and filter vast amounts of data makes it indispensable in today's world. Any attempt to improve the design of information in electronic media should address not only the visual display of the information, but the design of the computer itself and how one interacts with it in the context of the real world.

In making information accessible to people, it is necessary for designers to rethink current design paradigms. The computer screen is not a piece of paper and should not be treated as such. By taking advantage of the ability of the computer to display dynamic, flexible, and adaptive typography, we can invent new ways for people to read, interact with, and assimilate the written word. Like a garden, well-designed information should be legible, inviting, and comfortable, and its exploration should and can be a true delight.

CHAPTER 8

references

ABRAMS95 — Judith Abrams. *Learning Talmud; How to Use the Talmud The Steinsaltz Edition*. Jason Aronson Inc.; Northvale, New Jersey, 1995.

ADOBE95 — *Font & Function*, No. 16. Adobe Systems Inc., 1585 Charleston Road, P.O. Box 7900, Mountain View, CA 94039, Fall 1995.

AICHER88 — Otl Aicher. *Typographie*. Ernst & Sohn; Berlin, 1988.

BARBACHETTO87 — Barbachetto, Gianni. *Design Interface: How man and machine communicate. Olivetti Design Research by King and Miranda*. Arcadia, Milano. 1987.

BENEDIKT92 — *Cyberspace First Steps*. M. Benedikt, Editor, MIT Press; Cambridge, MA, 1992.

BOLT79 — Richar A. Bolt. *Spatial Data Management*. MIT, Cambridge, 1979.

BRILLIANT84 — Brilliant, Richard. *Visual Narratives: Storytelling in Etruscan and Roman Art*. Cornell University Press; Ithica, 1984.

BRINGHURST92 — Robert Bringhurst. *The Elements of Typographic Style*. Hartley & Marks; Vancouver, B. C., 1992.

BUSH46 — Vannevar Bush. “As We May Think”, *Atlantic Monthly*. July 1946. pp. 104-107.

CARD91 — Card, S.K., Robertson, G.G., and Mackinlay, J.D. “The Information Visualizer, An Information Workspace”. *Proceedings of CHI’91*. 1991. pp.181-188.

CARTER90 — Mathew Carter. “Typography and Current Technologies”, *Design Quarterly 148*, 1990. pp.55-64.

CHATWIN87 — BruceChatwin. *The Songlines*. Penguin Books; NewYork, 1987.

COLBY92 — Colby, G., and Scholl, L. “Transparency and Blur as Selective Cues for Complex Information”. *Proceedings of SPIE* 1992.

CORBU46 — Le Corbusier. *Towards a New Architecture*. Architectural Press; London, 1946.

EISNER85 — Will Eisner. *Comics & Sequential Art*. Poorhouse Press; Tamarac, FL, 1985.

FAIRCHILD88 — Fairchild, K.M., Poltrock, S.E., and Furnas, G.W. “SemNet: Three-Dimensional Graphic Representations of Large Knowledge Bases.” In *Cognitive Science and its Application for Human-Computer Interface*. R.Guindon (ed.). Lawrence Erlbaum. 1988.

Christiane Fellbaum, Ed. *Wordnet : An Electronic Lexical Database*, MIT Press; Cambridge, MA. 1998. FELLBAUM98

Feiner, S., MacIntyre, B., and Seligmann, D. “Knowledge-based augmented reality”. *Communications of the ACM*, 36(7), July 1993. pp.52-62. FIENER93

George Fitzmaurice. “Situated Information Spaces and Spatially Aware Palmtop Computers”, *Communications of the ACM*, July 1993. pp.38-49. FITZMAURICE93

George Fitzmourice, Hiroshi Ishii & William Buxton. “Bricks: Laying the Foundations for Graspable User Interfaces”, *CHI'95 Proceedings*, Denver, CO, 1995. pp.442-449. FITZMAURICE95

Tinsley Galyean. *Narrative Guidance of Interactivity.* PhD Thesis, Massachusetts Institute of Technology, 1995. GALYEAN95

Karl Gerstner. *Designing Programmes*. Teufen AR, Switzerland: Arthur Niggli, 1968. GERSTNER68

Gerstner, Karl. *The Forms of Color: The Interaction of Visual Elements*. MIT Press; Cambridge, MA, 1986. GERSTNER86

Eleanor J. Gibson and Harry Levin. *The Psychology of Reading*. The MIT Press; Cambridge, MA, 1975. GIBSON75

Eric Gill. *An Essay on Typography.* David R. Godine; Boston, 1936 (1988). GILL36

Gould, J.D.; “*Looking at pictures*”, paper presented at *National Symposium on Eyemovements and Psychological Processes,* August 19, 1974. GOULD74

HINKLEY94 Hinckley, K., Pausch, R., Goble, J. C. and Kasell, N. F. . "Passive Real-World Interface Props for Neurosurgical Visualization". *Proceedings of CHI '94*, 1994. pp. 452-458.

ID98 *ID Magazine*, Annual Design Review, July/August 1998. page 131.

ISHII94 Ishii, H., Kobayashi, M. and Arita, K., "Iterative Design of Seamless Collaboration Media", *Communications of the ACM* (CACM), ACM, Vol. 37, No. 8, August 1994. pp. 83-97.

ISHIZAKI96 Suguru Ishizaki. *Typographic Performance: Continuous Design Solutions as Emergent Behaviors of Active Agents*. PhD Thesis, Massachusetts Institute of Techology, February, 1996.

ITTEN75 Johannes Itten. *Design and Form*.Van Nostrand Reinhold; New York, 1975.

KAPLAN82 S. Kaplan and R. Kaplan, *Cognition and Environment: Functioning in an Uncertain World*. Praeger, 1982.

KEPES61 Gyorgy Kepes. *Language of Vision*. Paul Theobald; Chicago, 1961.

KLEE53 Paul Klee. *Pedological Sketchbook*. Faber and Faber Ltd.; England, 1953.

KOONS91 Koons, D., Sparrell, C., Thorisson, K. *Integrating Simultaneous Input from Speech, Gaze, and Hand Gestures*, in Maybury, ed. 1993.

LAUREL90 Laurel, Brenda, Ed. *The Art of Human-Computer Interface Design*, Addison-Wesley; Reading, MA, 1990.

Emmanuel Levinas. *L'Au-Dela Du Verset: Lectures et Discours Talmudiques*. Les Éditions de Minuit; Paris, 1982. LEVINAS82

Emmanuel Levinas. *Beyond the Verse: Talmudic Readings and Lectures*. Indiana University Press; Bloomington, Indiana, 1994. LEVINAS94

Ishantha Lokuge and Suguru Ishizaki, "GeoSpace: An Interactive Visualization System for Exploring Complex Information Spaces." *Proceedings of CHI'95*, Denver Colorado, 1995 LOKUGE95

Kevin Lynch. *The Image of the City*. MIT Press; Cambridge, MA, 1960. LYNCH60

Mackinlay, J.D., Robertson, G.G., and Card, S.K. "The Perspective Wall: Detail and Context Smoothly Integrated." *Proceedings of CHI'91*, 1991. pp.173-179. MACKINLAY91

John Maeda. "The South Face of the Mountain." *Technology Review*, July/August, 1998. MAEDA98

John Maeda and Kevin McGee. *Dynamic Form*. International Media Research Foundation; Tokyo, 1993. MAEDA93

T. Masuishi, D. Small, and R. MacNeil, "6,000 x 2,000 Display Prototype," *SPIE Proceedings*, Vol. 1664 (1992), pp. 202-209. MASUISHI92

Lazlo Moholy-Nagy. *Vision in Motion*. Paul Theobald; Chicago,1947. MOHOLY-NAGY47

Lazlo Moholy-Nagy. *Painting Photography Film*. MIT Press; Cambridge, MA, 1969. MOHOLY-NAGY69

MOORE88 Moore, Charles W., Mitchell, William J. & Turnbull, Jr., William. *The Poetics of Gardens*. MIT Press: Cambridge, 1988.

MÜLLER-BROCKMANN88 Müller-Brockmann, J. *Grid Systems in Graphic Design*. 3rd revised edition ed., Verlag Gerd Hatje; Stuttgart, 1988.

NEGROPONTE78 Nicholas Negroponte, Richard Bolt and Muriel Cooper. *Books without Pages*. Proposal from the Architecture Machine Group, MIT to the Office of Information, Science and Technology, NSF. Collection of the author. 1978.

NISHIMURA85 Nishimura, Y., Sato, K. "Dynamic Information Display", *Visible Language* 19, 1985. pp.251-271.

NORMAN90 Donald Norman. *The Design of Everyday Things*. Addison-Wesley; Reading, MA, 1990.

O'REGAN92 O'Reagan. "Optimal viewing position in words and the strategy-tactics theory of eye movements in reading," in *Eye Movements and Visual Cognition: Scene Perception and Reading*. K. Rayner (ed.), Springer-Verlag, New York, NY, 1992. pp.333-354.

PETRICK82 M. S. Petrick and M. C. Potter. *Acoustic and semantic encoding during rapid reading.* 1982.

PINKER85 Steven Pinker, ed. *Visual Cognition*. MIT Press; Cambridge, MA, 1985. Chapters 1, 2 & 4.

POTTER84 Mary C. Potter, "Rapid Serial Visual Presentation (RSVP): A Method for Studying Language Processing." in D. E. Kieras and M.A. Just (eds.) *New Methods in Reading Comprehension Research*. Erlbaum, 1984.

Robertson, G.G., Mackinlay, J.D. and Card, S.K. "Cone Trees: Animated 3D Visualizations of Hierarchical Information". *Proceedings of CHI'91*. 1991. pp.189-194. ROBERTSON91

Richard Rubinstein. *Digital Typography,* Addison-Wesley: Reading, MA, 1988. RUBINSTEIN88

Laura Scholl. *The Transitional Image*, MSVS Thesis, Massachusetts Institute of Technology, September, 1991. SCHOLL91

Ben Shneiderman. "Direct Manipulation: A Step Beyond Programming Languages." *IEEE Computer*, New York; IEEE Computer Society, August 1983. SHNEIDERMAN83

Louis Silverstein. *Newspaper Design for the Times*. Van Nostrand Reinhold; New York, NY, 1990. SILVERSTEIN90

D. Small, S. Ishizaki, and M. Cooper. "Typographic Space," Human Factors in Computing Systems, *CHI '94 Conference Companion*, ACM, 1994. pp.437–438. SMALL94

Small, D., Cooper, M. in *Information Architects*, R.S. Wurman (ed.), Graphis Press, 1995. pp.202-211. SMALL95

Joshua Smith. *Electric Field Imaging*. PhD Thesis, Massachusetts Institute of Technology, February 1999. SMITH99

Erik Spiekermann & E. M. Ginger. *Stop Stealing Sheep & Find Out How Type Works*, Adobe Press; Mountain View, CA, 1993. SPIEKERMAN93

Staples, L. "Representation in Virtual Space: Visual Convention in the Graphical User Interface." *Proceedings of InterCHI'93*. 1993. pp.348-354. STAPLES93

STARNER95 Starner, T., Mann, S., et al. *Wearable Computing and Augmented Reality*, MIT Media Lab, 1995.

STEINSALTZ89 Adin Steinsaltz. *The Talmud; Reference Guide*, Random House; New York, NY,1989.

STILLINGS95 Stillings, Neil A., et al. *Cognitive Science*. The MIT Press, Cambridge, MA, 1995. pp.37-42.

STRAUSFELD95 Lisa Strausfeld, “Financial Viewpoints: Using Point-of-view to Enable Uderstanding of Information.” *CHI ’95 Conference Companion*. Denver, CO, May 1995. pp.208-209.

THOMAS81 Frank Thomas and Ollie Johnston. *The Illusion of Life*, Hyperion; NY, 1981.

TOBA12 Toba. *Scroll of Frolicking Animals* (Chojugiga). Benrido, Kyoto, Japan (12th century).

TSCHICHOLD67 Jan Tschichold. *Asymmetric Typography*. Reinhold Publishing; New York, 1967.

TSCHICHOLD91 Jan Tschichold. *The Form of the Book*. Hartley & Marks; Washington, 1991.

TSCHICHOLD95 Jan Tschichold. *The New Typography*. University of California Press, 1995.

TUFTE90 Tufte, E.R. *Envisioning Information*. Graphics Press; Cheshire, Connecticut, 1990.

VENTURI72 Robert Venturi. *Learning from Las Vegas*. MIT Press; Cambridge, MA, 1972.

W. P. Wallace and G. Koury. "Comprehension following speed listening at 500 words per minute." Paper presented at the Psychnomic Society Meeting, Minneapolis, 1982. WALLACE82

Pierre Wellner, "Interacting With Paper on the Digital Desk," *Communications of the ACM*, July 1993. pp.87-96. WELLNER93

Whitney, Patrick, Ed. *Design in the Information Environment*. Southern Illinois University Press, Carbondale, 1985. WHITNEY85

Winograd, T. and Flores, F., *Understanding Computers and Cognition*, Addison-Wesley, 1986 WINOGRAD86

Yin Yin Wong. *Temporal Typography: Characterization of time-varying typographic forms*. M.S. Thesis, Massachusetts Institute of Technology, Cambridge, MA. 1995. WONG95

Ygge, J. and Lennerstrand, G. Ed. *Eye Movements in Reading*. Pergamon, Oxford, 1994. YGGE94

APPENDIX A AcWindows

The acWindow system is a generic toolkit for building interactive 3D applications. acWindows is built on Silicon Graphics' Performer API and many acWindow classes are derived from Performer classes. By using an existing system, it was possible to leverage a large amount of previous effort - in picking, the maintenance of a display list and in tuning the application for optimal performance.

The most important base class is the acObject. It is derived from Performer's pfDCS class, which is an abstraction for a coordinate frame. acObjects can be attached to parent objects (coordinate frame) and can have any number of child objects attached. The acObject has methods for translation, rotation, scale, switching (on and off), and has an associated color and transparency. In addition to these direct methods, acObjects can receive messages (acMessage objects) which can trigger the built in methods. An acObject also contains a list of messages which can be triggered upon the receipt of a message. So, for example, using the attach-

Message method, you can set up and acObject such that it will send a pre-defined message to another acObject when it receives and successfully handles a particular message.

Messages

An important aspect of the acMessage object is that it contains not only a message tag (i.e. TRANSLATE, SETCOLOR, JUSTDOWN) and a pointer to data associated with that tag, it also contains a float value for the **strength** of the message and for its **duration**. the following is a table of system messages:

Event	**value**	**data**	**function**
NOEVENT	0	NULL	NULL
JUSTDOWN	0x00000100	acEvent	Mouse down
JUSTUP	0x00000200	acEvent	Mouse up
BEENDOWN	0x00000400	acEvent	Mouse down/over
BEENUP	0x00000800	acEvent	Mouse up/over
KEYPRESS	0x00001000	acEvent	Keyboard event
TICK	0x00002000	acEvent	Clock tick
ENTER	0x00004000	acEvent	Enter object
LEAVE	0x00008000	acEvent	Leave object
SETCOLOR	0x00010000	pfVec4	Set object color
SETALPHA	0x00020000	float	Set transparency
TRANSLATE	0x00040000	pfVec3	Move object
ROTATE	0x00080000	pfVec3	Rotate object
SCALE	0x00100000	float	Scale object
SCALE3	0x00200000	pfVec3	Scale object
HIDE	0x00400000	NULL	Hide object
SHOW	0x00800000	NULL	Show object
BLUR	0x01000000	float	Set blur (for type)
HIGHLIGHT	0x02000000	pfVec4	draw bounding box
UNHIGHLIGHT	0x04000000	NULL	erase box

Durations are handled automatically for most messages - the object will arrange to subscribe to clock events and will compute for each frame the interpolated value between the starting value and the goal value. At the end of the time allotted, the object will set itself to the goal value and unsubscribe to clock events. These animations are actually stored on a stack and can overlap with each other.

Object hierarchy

There are a variety of basic graphic types which are derived from the base acObject class. Here is the hierarchy:

Here is the hierarchy of acWindow classes

- acObject
    - acText
        - acTextString
            - acFloatRegister
            - acStringRegister
            - acIntRegister
            - acPopupLine
            - acPopupTrigger
        - acParagraph
        - acWrappedParagraph
    - acRectangle
    - acScene
    - acLoadable
    - acParticle
    - acObjectEditor
    - acSketch

In addition to sending messages, it is also easy to attach function callbacks to acObjects.Each object maintains a linked list of callbacks. When an acObject receives a message (e.g. JUSTDOWN) it checks to see if the message matches any callback functions and if so they are executed.

All messages that are handled and all functions that are called return a float value. The system accumulates the returns and halts message propagation when the value exceeds 1.0.

Input

Input from the keyboard and mouse is handled automatically. In fact, any device which generates messages can be considered an input device. The clock, for example is an input device which produces TICK events. Any input device which produces x-y information, such as a mouse or a tablet, has messages sent to the objects it covers. The coordinate of the device determines a ray through the viewpoint and the x-y location of the device on the view frustum. All objects are checked for intersection with this ray and a list of objects, sorted from front to back, is returned. These objects are sent the appropriate mouse event (JUSTDOWN, BEENDOWN, JUSTUP or BEENUP) and the first one to completely handle the message will block further propagation.