

Beyond Paper: Supporting Active Reading with Free Form Digital Ink Annotations

Bill N. Schilit, Gene Golovchinsky, Morgan N. Price

FX Palo Alto Laboratory, Inc.

3400 Hillview Ave., Bldg. 4

Palo Alto, CA 94304

+1 650 813-7322

{schilit, gene, price}@pal.xerox.com

ABSTRACT

Reading frequently involves not just looking at words on a page, but also underlining, highlighting and commenting, either on the text or in a separate notebook. This combination of reading with critical thinking and learning is called *active reading* [2]. To explore the premise that computation can enhance active reading we have built the XLibris “active reading machine.” XLibris uses a commercial high-resolution pen tablet display along with a paper-like user interface to support the key affordances of paper for active reading: the reader can hold a scanned image of a page in his lap and mark on it with digital ink. To go beyond paper, XLibris monitors the free-form ink annotations made while reading, and uses these to organize and to search for information. Readers can review, sort and filter clippings of their annotated text in a “Reader’s Notebook.” XLibris also searches for material related to the annotated text, and displays links to similar documents unobtrusively in the margin. XLibris demonstrates that computers can help active readers organize and find information while retaining many of the advantages of reading on paper.

Keywords

Paper-like user interface, reading online, affordances of paper, pen computing, dynamic hypertext, document metaphor, information retrieval

INTRODUCTION

Computers, once expected to create a paper-less office, have instead produced ever-increasing quantities of paper documents. Dataquest predicts that 1,344 billion pages will be generated by printers and copiers in the US in 1997 [15]. This statistic suggests that people are not using computers to read. Whereas paper is lightweight, inexpensive, and easy to annotate, interfaces for reading online typically involve clumsy interactions with bulky desktop monitors.

Although reading online presents a number of problems, we will show that integrating computation with reading also presents novel opportunities for improving the reading process. Thus there is a tension between the advantages provided by computation and the advantages provided by paper: the choice depends on the reader’s goals. For reading a romance novel at the beach, low weight and portability are essential, and it is unlikely that computation could provide any real benefit. For more dynamic types of reading, however, computation may be desirable.

Active reading is the combination of reading with critical thinking and learning, and is a fundamental part of education and knowledge work. Active reading involves not just reading *per se*, but also underlining, highlighting and scribbling comments, either on the text itself or in a separate notebook [2]. Readers use these marks to organize their reading for later review and retrieval. In addition, active reading often requires readers to move from one text to another to satisfy their information needs.




Figure 1: XLibris prototype in a reader's lap

We have built an “active reading machine,” XLibris, to explore the premise that computation can enhance active

To appear in Proceedings of CHI 98

reading. To make reading online acceptable, XLibris implements a paper document metaphor (see Figure 1) that provides key affordances of paper. To go beyond paper and make reading online desirable, XLibris helps readers organize and retrieve information. By exploiting the existing practice of free-form ink annotation, XLibris provides these benefits with minimal additional effort from the reader. To apply these advantages to existing documents, XLibris allows input of both paper and digital documents. Indeed, an active reading machine would likely be used together with desktop computers (e.g., for writing) and with paper (e.g., for some forms of collaboration).

The rest of this paper is organized as follows. We describe active reading practices, and the problems with reading online. We then describe the XLibris system in detail, and conclude with a discussion of related and future work.

ACTIVE READING

The combination of reading with critical thinking is a fundamental part of education and knowledge work. This activity is practiced by students learning a textbook discipline, by lawyers processing discovery documents, researchers keeping “up to date,” by analysts finding answers to specific questions, by executives cleaning out their inboxes, and by citizens deciding how to vote. Active reading is a skill that most literate people share.

Today, active reading is a “craft skill” rather than a formal academic discipline. Even though critical reading is a key part of learning from a text, reading is rarely included in school curricula beyond the elementary level. Nevertheless, improving the skills of learning from and finding information in a text has occupied educators for centuries. More recently, Mortimer Adler’s *How to Read a Book* [2]—part of the Great Books campaign—presented a methodology for active reading in great detail.

Adler and other educators recommend pigeon-holing a text, skimming the contents, reading at different speeds, marking (underlining, highlighting, starring) important information, and in general, “carrying on a conversation” with the author. For example, scribbling “Norman would disagree!” in the margin is part of a dialog between reader and author. These practices have been observed in real use in a number of studies [21, 16, 1, 22] that are summarized in the remainder of this section.

Note Taking Practices

Readers write. Marks made while reading aid understanding, memorization and later retrieval. Readers generate marks in different places: on the document, in a separate notebook, or sometimes on loose-leaf paper or note cards. We use *annotation* to refer to marks on (or attached to) reading matter, and *notes* to mean marks that are not co-located with the text to which they refer. The tradeoffs between annotations and notes are discussed below.

Annotations

Many people annotate directly on paper documents while they are reading. When people do not want to mark in their

books (or in borrowed reading matter), they may write on removable Post-It™ notes or on slips of paper. Readers tend to scribble short comments in margins; they draw asterisks to make an important statement stand out; they highlight, underline or circle key words or phrases [16].

The advantages of annotating directly on the page are convenience, immersion in the document context, and visual search. Annotating documents does not disrupt the flow of reading. Also, annotations contain more information because they occur within the “context” of the document. Marks such as “!” or “?” or “not true!” may represent various levels of agreement with adjacent text. Finally, annotations stand out visually, allowing readers to scan for them.

The major disadvantage of annotations is that they become lost in the piles of papers and books. Annotations only work well when you *see* them on the document page.

Notebooks

Some readers use a bound notebook to store their notes in one place. People use notebooks to record facts, figures, and key points, to copy interesting quotes, and to jot down ideas. Readers can review their notebooks to study and to understand what they have read. Over time, a notebook becomes a reader’s *catalogue raisonné*, containing annotations and thoughts about documents that have been read. Additionally, flipping through the pages in a notebook is a rapid way to review information across documents.

Although notebooks provide the reader a broad overview of their reading, they often lack details and depth. Readers must recreate the note’s context—the document name, page number, idea, quotation, etc.—in the notebook. Writing “yes!” in the margin is a useful annotation, but in a notebook it must be accompanied by substantially more information. Copying this context into a notebook can be tedious, time-consuming and error-prone. Incomplete notes may not be comprehensible later, and may also create problems of attribution.

Note Cards and Loose-leaf Paper

People also use index cards and loose-leaf paper to write notes while reading. The advantage of unbound notes is that they can be arranged and reorganized flexibly. For example, authors often organize their notes by topic or to fit an outline.

The problem with unbound notes is the effort required to create and to manage them. As with a notebook, there is no context available from the document, but there is also no context available from surrounding notes. Furthermore managing dozens of note cards can be a chore in itself.

In summary, all three places for marking—on the page, in notebooks, on unbound paper—are used not just to record information but also to organize it. Annotations highlight key information but tend to be lost in piles of paper. Notebooks are reviewed easily, but taking notes is tedious and error-prone. Unbound notes can be reorganized flexibly, but require even more effort by the note taker.

Finding related material

Finding related material and moving from one text to another is also part of active reading. For example, an ecology student studying the effects of acid rain in the Appalachian Mountains decides to search the Internet for information on the economy of West Virginia. A doctor reading up on a drug she's never prescribed decides to follow a reference to a study of its side effects.

This activity of finding related materials while reading is often disconnected from the main activity. Typical information retrieval interfaces force users to suspend their reading activity to identify and then to retrieve related documents. References found in paper documents are even more time-consuming to track down. Although hypertext was designed to address some of these problems, authors cannot anticipate the information needs of all readers. Ironically, although readers' interests change fluidly, interfaces for moving from reading to finding and back to reading again tend to be time consuming and disruptive.

THE CHALLENGES OF ACTIVE READING ONLINE

People typically print long documents to read them, rather than reading them on the computer. Why does this happen? What affordances does paper provide that are lacking online? A number of researchers [5, 12, 18, 21] have studied these issues; we summarize applicable results below.

Tangibility

Tangibility refers to the physical properties of sheets or stacks of paper that allow them to be held, moved, folded, rotated, etc. Readers move their documents around in space to improve legibility [5], to adjust perspective or gain objective distance [12], and to ease and speed up handwriting [11]. In contrast, reading online typically means reading from a monitor on a desk.

The physical nature of paper also affords effortless navigation [18]. For example, readers can look at the previous page while marking on the current one, sense the length of a document by its thickness and weight, and sense location in a book by the thickness of the sides.

Free-form ink annotation

Free-form ink annotation refers to the ease of marking on paper. Annotating with a pen requires little cognitive overhead compared to typing or to selecting text with a mouse and issuing a command. Annotation is an important part of active reading, yet it tends to be very difficult online [1, 18].

Furthermore, annotations on paper are conceptually and visually separate from the underlying document. In contrast, text annotations online often blend in with surrounding text or change document layout.

Page orientation

Page orientation refers to the fixed division of paper documents into pages and the fixed spatial layout of those pages. Page layout conveys the organization of the document and supports navigation, skimming, and spatial memory [6].

This "sense of the document" is often lost online [12, 18]. One problem is that the landscape orientation of typical monitors does not allow a full page of text to be displayed. This leads to awkward view scrolling and zooming. In addition, scrolling views (e.g., Web browsers, text processors) that do not respect fixed page layout can make it difficult to read a document non-linearly.

Multiple displays

Reading at work often involves using several documents—paper and/or digital—together, each with their own display [1]. Paper provides an unlimited number of physical displays because it is cheap. Most computers provide multiple virtual displays, but existing windowing systems are too cumbersome to support reading from several documents [18].

Sharing

Paper is directly sharable. More than one person can view and interact with a single piece of paper at a time; people can hand over a paper document in person; and people can see what others are doing with paper documents [21].

Legibility

Although legibility is a common explanation for why people prefer reading on paper, legibility does not appear to be a significant factor. Many studies in the 1980's found significant differences in reading speed or proofreading accuracy between CRT displays and paper, but these differences shrink or disappear with modern high-quality CRTs [5, 10].

MEETING THE CHALLENGE OF READING ONLINE

Supporting all of the advantages of paper described above in an active reading machine seems to be impossible. It is also unnecessary: our goal is to improve active reading, not to replace paper. Readers will use active reading machines together with paper and traditional computers, and documents will move back and forth as best supports people's work. This suggests that active reading machines need not necessarily provide multiple displays or support sharing directly.

Instead we have focused our attention on tangibility, free-form ink annotation, and page orientation. To provide these affordances, XLibris combines a pen tablet display with a paper document metaphor. Together, these allow XLibris to emulate the physical experience of reading and marking on paper: the reader can hold a scanned image of a page in their lap and mark on it with digital ink.

Pen Tablet Display

XLibris runs on a paper-like tablet computer. A tablet is a lightweight (1–2 kg) device that combines a high-resolution flat-panel display and a pen. The devices come in two varieties: tablet displays connected to conventional computers, and standalone, portable pen computers.

Tangibility

Unlike CRT monitors, tablets are graspable physical objects. Although 1–2 kg is heavier than most books, tablets are light enough to put in one's lap, comfortable to

write on, and movable. By movable we mean that they can be tilted easily, shifted around, etc., within a small space such as an office. This distinguishes tablets from portable computers (e.g., laptops) that can be packed up and transported from place to place, but may not be conveniently movable once unpacked for use. Conversely, a tethered device—a tablet connected to a stationary computer or source of power, or a book on a chain in a library—may be movable but not portable.

Physical widgets incorporated into the tablet can further enhance tangibility. For example, Harrison *et al* [14] have mounted pressure sensors on the top corners of the XLibris prototype’s case (Figure 1). Readers turn pages by thumbing across these sensors (from left to right or from right to left), mimicking page turning in a paper document.

Paper document metaphor

XLibris uses the *paper document metaphor* to capture several important characteristics of paper documents. The metaphor includes a display that emulates the appearance of a sheet of paper, an interface for viewing pages in the linear order of the document, the ability to mark in any place on any page, and cues about the size of a document and about the reader’s location within it.

XLibris reinforces this user interface metaphor through the form factor of the device. The tablet and the pen complement the software interface to create a richer, more tangible document reading experience.

XLibris differs radically from standard WIMP interfaces. This is reflected not only in the document metaphor, but also in fundamental differences in interaction between a desktop computer and a tablet. Although the mouse and the pen serve similar purposes, they have quite different affordances. Furthermore, screen space in XLibris is at a premium because tangibility requires a small form factor, and also because document pages require all the pixels that current display technology provides.

User interface design principles

As we designed the XLibris system, we identified several design principles that should pervade and structure the user interface. These principles included a transparent (“No UI” UI) user interface and modeless interaction.

XLibris provides a minimal interface with a few commands for turning pages, changing views and selecting ink and highlighter colors. The richness of the interaction is derived from the form factor of the device and from the ease and familiarity of making marks with a pen. The intent was to create a transparent user interface that users would accept implicitly without the typical and frustrating process of “learning the UI.”

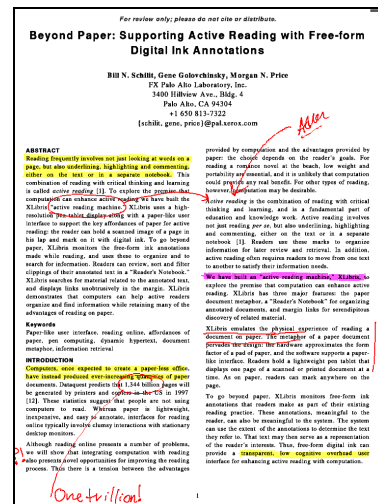


Figure 2. An annotated document page in XLibris

An essential aspect of ink on paper is its lack of modality: you can write anything you want, anywhere on a page, in any order. In XLibris, therefore, users can scribble notes, draw figures, and annotate text, all without switching modes or applications. In contrast, conventional graphical user interfaces use widgets to produce a highly structured interaction, and typically force users to adapt to the structure of applications.

Similarly desirable is the ability to perform the same task in different ways. Users prefer interfaces that provide multiple navigation options (i.e., queries, hypertext links, buttons, etc.) to interfaces with fewer navigation mechanisms [8]. Although the destination in each case may be the same, different interface actions may be appropriate for different goals, or the choice may depend on a user’s prior experience. Thus we designed XLibris to support a variety of navigation mechanisms interfaces; they will be discussed in detail in the “Support for navigation” section, below.

Page orientation

XLibris maintains page layout by using tablets in portrait mode and devotes the entire display to a single page (Figure 2). XLibris uses page turning rather than scrolling to move between pages of a document. Page turning is preferred to scrolling because it preserves the visual characteristics of pages (enhancing visual memory) and because it is more consistent with the paper document metaphor.

XLibris provides visual feedback about the length of the document and about the reader’s location within the document using “location guides” (Figure 2). These are small widgets at the top corners of the page that indicate visually the approximate size of the document, the reader’s position in the document, and whether there is a following or a preceding page. Location guides are similar to Book Emulator’s black page edges [3] but take up far less space; their design will be refined based on experimental evaluation.

Free-form ink annotation

XLibris supports free-form ink annotation directly on the page, anywhere on the page. We believe that the ability to make unstructured, idiosyncratic marks is a key aspect of a paper-like user interface. Although such marks may not have explicit meaning to the computer, they have rich meaning to the user and thus support episodic memory. XLibris provides different colors of highlighters and pens to increase users' flexibility of expression.

Operating XLibris

In the real world, we envision our active reading machine coexisting with personal computers and with paper, and we see documents moving freely between these three media. In our current system, readers can "print" from their desktop computer into XLibris and can scan paper documents into XLibris. Similarly, we envision XLibris printing to paper or exporting text and images for use in desktop applications.

Navigation

One of the design guidelines we identified was that XLibris should support multiple navigation mechanisms. We have implemented the following: readers can turn pages using the pressure sensors until the desired page is reached; a "document view" showing thumbnails of all pages simultaneously can be used to select the desired page (Figure 3); or the "Reader's Notebook" view (see next section) can be used to find a page with a familiar annotation. A view history allows user to backtrack through the reading session in time. We also envision using the location guides to riffle through pages.



Figure 3: The document view shows page thumbnails, including annotations, for the entire document.

XLibris uses small page images, or thumbnails, instead of generic icons to represent links to pages and documents. Although almost entirely illegible, the thumbnails retain the layout information and reveal the idiosyncratic annotations when they are present. Thus readers can identify marked-up pages or pages with distinctive figures or unusual layouts from the thumbnail, without following the link.

XLibris also provides a "workspace view" to help readers navigate between documents. The workspace view (Figure 4) displays a thumbnail (with a shadow that indicates document size) for each of the reader's documents. Each thumbnail is linked to the first page of the corresponding document.




Figure 4: The workspace view shows all of the reader's documents. Each document has a drop-shadow that indicates its approximate length.

Command Invocation

Readers using XLibris can ink on the page with a pen, turn pages by thumbing on the case, follow links by tapping with the pen, and invoke commands from a menu. Currently the commands include selecting a pen, highlighter or eraser, toggling margin links, or moving to one of the views described above.



Figure 5: "Silk screen" menu at bottom of page

We have explored three styles of command menus. Initially we implemented a circular pie menu that popped up when users tapped while holding down the pen barrel button. Since users found this awkward, we replaced the button with a "dwell" gesture: pressing the pen down and holding it stationary brings up the pie menu. Dwell menus should not interfere with writing because dwell is not part of paper practice: holding a pen in one spot on paper causes bleeding. We also designed the menu as a set of silk screen buttons, emulated by a button bar at the bottom of the display, as shown in Figure 5.

ENHANCING READING WITH COMPUTATION

XLibris provides a paper-like user interface based on the document metaphor, but some functions of an active reading machine should not be entirely paper-like. For example, active links do not exist in paper documents, and paper notebooks cannot be reorganized dynamically. Although these features are not present "verbatim" in the document metaphor, they are not inconsistent with it. For example, people have physically cut out sections of documents to generate clipping files, and there is a scholarly tradition of putting references to related works in the margins of a document. These existing practices are reinterpreted in XLibris in a computational manner, as described below.

The Reader's Notebook

The Reader's Notebook is a notebook of clippings derived from readers' annotations. The Reader's Notebook combines the advantages of annotating directly on the page, of taking notes in a separate notebook, and of organizing index cards. Furthermore, the Reader's Notebook overcomes the limitations of its paper cousins: as with paper documents, readers mark on the page, in the context of the document, and without tedious and error-prone copying of information. As with a bound notebook, readers can review concise annotations by time. Finally, as with note cards, flexible filtering and sorting of the view allow readers to reorganize their information as their needs change.

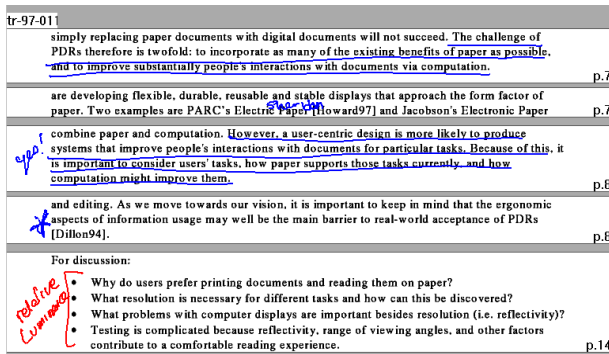


Figure 6: The Reader's Notebook shows annotated clippings of documents laid end-to-end. Clippings are labeled with document title and page number, and are linked to the corresponding pages.

The Reader's Notebook (Figure 6) extracts annotated clippings and lays them end-to-end in a separate, multi-page view. Each clipping is linked to the annotated page, so the reader can move fluidly between notes and documents. Each clipping includes some surrounding text and is labeled with document title and page number to help the reader understand the meaning of the marks.

How much text should be shown with each annotation? Clippings are derived from the bounding boxes of annotations by expanding them horizontally to the width of the page, and by expanding them vertically to include complete words. Clippings that overlap are merged together.

The Reader's Notebook can display clippings from one document or from all documents, and can be sorted and filtered. By default, clippings are sorted by time, so that new information appears at the end, as in a paper notebook. Clippings can also be sorted by page number. Searching through clippings sorted by page number is analogous to rummaging through a pile of paper documents, but should be faster because readers need not deal with less important (un-annotated) information.

Finally, readers can filter the clippings by ink color to search for different kinds of marks or to group related items together. Many readers already use different pens to mark different types of information. For example, some lawyers highlight "pro" information in green and "con" information in red.

Margin Links

XLibris supports a novel type of *serendipitous* retrieval. Consider a library containing large collections of documents arranged by topic. As people walk to the shelf for a particular book, walls of related books surround them, making it possible to find interesting material accidentally. This serendipity in finding books in a library is a common and often rewarding experience. Magazines and journals sometimes provide the same function. Unfortunately, in both paper and online environments, finding related

materials while reading is often disconnected from the main activity.

Margin links provide serendipitous access to related documents *during* the reading process. As readers mark up documents, the system finds documents related to annotated passages and presents links unobtrusively, in the margin. Each anchor is a thumbnail image of the target page (see Figure 7).

Link creation does not disrupt the reading process. The reader is free to disregard the suggestion and to continue reading. The intent is to provide a modeless link suggestion mechanism, and to allow readers to follow the links at their leisure, rather than at the time that the links are identified.

Query-mediated links

Query-mediated links derive a query from a user's interaction with a document and use that query to identify related documents. Golovchinsky [9] has shown that query-mediated links based on explicitly selected words and passages are effective in supporting information exploration tasks.

Similarly, XLibris computes margin links from the words, phrases and passages that are implicitly selected by the reader's marks. Each mark is converted in turn into a text selection, a full-text query, and a best-matching passage. The system adds a margin link to the best match if its similarity value (or belief score) is above a threshold.

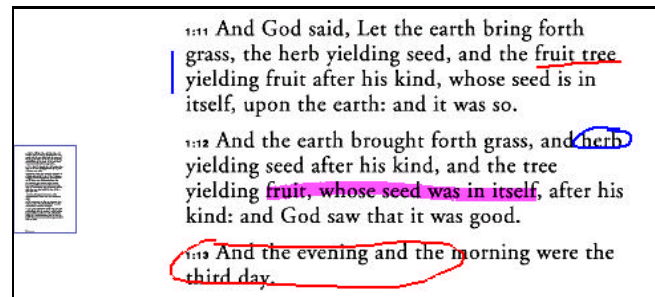


Figure 7. Examples of highlighting, underlining, circling, and margin annotations. Each annotation generates a query. If a good match is found, XLibris adds a margin link (rectangle on the left) that shows the thumbnail of the destination page.

Query-mediated links require sources of appropriate documents for retrieval. Although our prototype implementation relies on the reader's private collection, commercial digital document repositories are becoming available. The ACM and the IEEE already publish their journals and conference proceedings electronically, and the Columbia University Press has announced online access to their academic publications. Institutional memory repositories and other private collections can also serve as sources of documents.

Ink annotations as queries

XLibris recognizes several distinct ink patterns from which queries are computed. These include underlined words,

circled words, circled passages, and margin annotations (Figure 7). Each type of annotation results in a different query for the search engine. Marks that select specific words translate into queries that emphasize those words over others in the same sentence. Marks that select longer passages generate queries that search for similar phrases. Currently, a separate query is computed for each stroke; we envision grouping ink by time and type of annotation.

In summary, margin links allow the reader to stumble into information serendipitously—without any explicit interaction—in the context of their reading. This contrasts with interfaces for browsing or information exploration, which separate information gathering from in-depth reading and require readers to describe their interests more explicitly.

IMPLEMENTATION

XLibris is implemented in C++ and the Win32 API, and runs on Windows 95/NT. The program takes over the entire display, and does not use any Windows widgets. Instead, XLibris defines a small set of widgets, renders them to an in-memory bitmap, rotates the bitmap, and copies it to the display device. In this way, XLibris uses a landscape device in portrait mode.

Our current hardware platform is a Mutoh 12.1" (30.7 cm) diagonal 1024x768-pixel tablet display tethered to a PC [17]. XLibris displays documents as 100-dpi 64-color bitmap images shrunk slightly to fit into the aspect ratio and size of this display. Although the margins are clipped slightly on each side, enough space remains for annotations. Documents containing 10 point text or larger are quite readable.

In order to treat scanned-in paper documents and printed digital documents identically, XLibris uses an “image + text” file format [19, 4]. The image supports rendering independent of the source of the document. The text, inferred from print commands or from OCR, is tied to the image by a bounding box for each word. This “text” supports information retrieval, and is never shown to the reader directly.

RELATED WORK

Most previous attempts to improve reading with computation have focused on creating new media—hypertext or hypermedia—that redefine the documents being read as well as the act of reading, placing them in a separate world from existing books and paper documents. This contrasts with our approach of supporting paper practices and paper-based documents.

A number of systems have explored digital ink annotation. Wang Freestyle [7] provides free-form ink annotation on a fixed page of information, but does not use computation to enhance the ink. Wolf et al. [23] built a “Paper-like Interface” consisting of a pen tablet display that supported traditional computer applications rather than active reading, and used gestures rather than annotations. The collaborative text editor MATE [13] interprets free-form ink revision marks as editing commands at the user’s request, or allows

users to specify their own interpretation. XLibris’ Reader’s Notebook is similar to the ink property index in Dynamite [22], but XLibris organizes the underlying document, not just the ink.

The Lectrice project at DEC SRC [4] designed and built a document reader intended to be an alternative to reading on paper. However, Lectrice did not provide free-form digital ink annotation or support active reading.

The Remembrance Agent [20] finds documents related to the last few lines typed into a text editor. Whereas the Remembrance Agent supports writing, XLibris supports reading by generating margin links from ink annotations and the underlying document.

CONCLUSIONS AND FUTURE WORK

We began with the premise that computation can enhance active reading. Toward this end we explored the task of active reading and the affordances of paper for reading. We formulated a set of challenges for reading online that must be overcome, and also a set of opportunities that computation provides.

We have designed and implemented the XLibris “active reading machine.” XLibris implements a paper document metaphor that includes a page-oriented display, document-like navigation, ink annotation, and cues indicating document size and location within the document. These software features are complemented by the affordances of the pen tablet device on which XLibris runs. The tangibility of the device reinforces the document reading experience.

XLibris allows users to mark up documents in a manner analogous to writing on paper, but then goes beyond paper by letting readers review, sort and filter clippings of their annotated text in a “Reader’s Notebook.” In addition, XLibris monitors the free-form ink annotations made while reading, and uses them to discover the reader’s interests. XLibris searches for material related to the annotated text, and displays links unobtrusively as “margin links.”

We have begun a paper-based study comparing reviewing annotations in the Reader’s Notebook with reviewing annotations on the document. In the future, we plan to deploy XLibris and evaluate its use in a real work setting. The main obstacle that must be overcome first is seamless scanning and printing into XLibris.

Although we have not completed our user evaluation, we are encouraged by our observations thus far. XLibris provides a natural, paper-like user interface and people find the computational enhancements powerful and compelling.

ACKNOWLEDGEMENTS

We thank Beverly Harrison, Anuj Gujar, Roy Want and Kim McCall for integrating the pressure sensors into XLibris. We thank DEC SRC for making their Lectern document format publicly available. We thank members of the Xerox PARC PDR team for comments and suggestions. We thank Joe Sullivan and Jim Baker for supporting this research.

REFERENCES

1. Adler, A., Gujar, A., Harrison, B.L., O'Hara, K., and Sellen, A. (1998) A Diary Study of Work-Related Reading: Design Implications for Digital Reading Devices. In *Proceedings of CHI98* (Los Angeles, CA, April 1998), ACM Press.
2. Adler, M.J. and van Doren, C. (1972) *How to Read a Book*. Simon and Schuster, New York, NY.
3. Benest, I.D. (1990) A hypertext system with controlled hype. In *Hypertext: State of the Art*, R. MacAleese and C. Green (Eds.), Intellect, Oxford, UK, pp. 52-63.
4. Birrell, A. and McJones, P. Virtual Paper. Available at www.research.digital.com/SRC/virtualpaper/home.html
5. Dillon, A. (1992) Reading from paper versus screens: a critical review of the empirical literature. *Ergonomics*, 35(10), pp.1297-1326.
6. Dillon, A. (1994) *Designing Usable Electronic Text*. Taylor and Francis, London, UK.
7. Francik, E. (1996) Rapid, Integrated Design of a Multimedia Communication System. In *Human-Computer Interface Design*, M. Rudisill, et al., (Eds.), Morgan Kaufman, San Francisco, CA.
8. Golovchinsky, G. (1997a) Queries? Links? Is there a Difference? In *Proceedings of CHI97*, (Atlanta, Georgia, March 1997), ACM Press, pp. 407-414.
9. Golovchinsky, G. (1997b) What the Query Told the Link: The integration of hypertext and information retrieval. In *Proceedings of Hypertext '97* (Southampton, UK, April 8-11), ACM Press, pp. 67-74.
10. Gould, A., Alfaro, L., Barnes, V., Finn, R., Grischkowsky, N., and Minuto, A. (1987) Reading is slower from CRT displays than from paper: attempts to isolate a single-variable explanation. *Human Factors*, 29(3), pp. 269-299.
11. Guiard, Y. (1987) Asymmetric Division of Labor in Human Skilled Bimanual Action: The Kinematic Chain as a Model. *Journal of Motor Behavior*, 19(4), pp. 486-517.
12. Haas, C. (1996) *Writing Technology: Studies on the materiality of literacy*. Lawrence Erlbaum Associates, Mahwah, NJ.
13. Hardock, G., Kurtenbach, G., and Buxton, W. A. (1993) A Marking Based Interface For Collaborative Writing. In *Proceedings of UIST '93*. ACM Press, pp. 259-266.
14. Harrison, B.L., Fishkin, K., Gujar, A., Mochon, C., and Want, R. (1998) Squeeze me, Hold me, Tilt me! An Exploration of Manipulative User Interfaces. In *Proceedings of CHI98* (Los Angeles, CA, April 1998), ACM Press.
15. Johnson, A. (1997) Fax Trends. *Office Equipment and Products*, July 1997, p. 25.
16. Levy, D.M. and Marshall, C.C. (1994) Washington's White Horse? A Look at Assumptions Underlying Digital Libraries. In *Proceedings of Digital Libraries '94*. Texas A&M University Press, College Station, TX.
17. Mutoh (1997) Mutoh America Inc. Available at www.mutoh.com
18. O'Hara, K. and Sellen, A. (1997) A Comparison of Reading Paper and On-Line Documents. In *Proceedings of CHI97* (Atlanta, GA, March 1997), ACM Press, pp. 335-342.
19. Phelps, T.A. and Wilensky, R. (1996) Towards Active, Extensible, Networked Documents: Multivalent Architecture and Applications. In *Proceedings of Digital Libraries '96*, ACM Press. Pp. 100-108.
20. Rhodes, B.J. and Starmer, T. (1996) A continuously running automated information retrieval system. In *Proceedings of The First International Conference on The Practical Application of Intelligent Agents and Multi Agent Technology (PAAM '96)*, London, UK, April 1996, pp. 487-495.
21. Sellen, A. and Harper, R. (1997) Paper as an Analytic Resource in the Design of New Technologies. In *Proceedings of CHI97* (Atlanta, GA, March 1997), ACM Press, pp. 319- 326.
22. Wilcox, L.D., Schilit, B.N., and Sawhney, N. (1997) Dynamite: A Dynamically Organized Ink and Audio Notebook. In *Proceedings of CHI97* (Atlanta, Georgia, March 1997), ACM Press, pp. 186-193.
23. Wolf, G.C., Rhyne, J.R., and Ellozy, H.A. (1989) The Paper-like Interface. In *Designing and Using Human-Computer Interfaces and Knowledge-Based Systems*, G. Salvendy and M.J. Smith (eds.), Elsevier Science Publishers, Amsterdam, The Netherlands, pp. 494-501.