Linking By Inking: Trailblazing in a Paper-like Hypertext

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

FX Palo Alto Laboratory
3400 Hillview Avenue, Bldg. 4
Palo Alto, CA 94304 USA
+1 650 813-7233
{price, gene, schilit}@pal.xerox.com

ABSTRACT

"Linking by inking" is a new interface for reader-directed link construction that bridges reading and browsing activities. We are developing linking by inking in XLibris, a hypertext system based on the paper document metaphor. Readers use a pen computer to annotate page images with free-form ink, much as they would on paper, and the computer constructs hypertext links based on the ink marks. This paper proposes two kinds of reader-directed links: automatic and manual. Automatic links are created in response to readers' annotations. The system extracts the text near free-form ink marks, uses these terms to construct queries, executes queries against a collection of documents, and unobtrusively displays links to related documents in the margin or as "further reading lists." We also present a design for manual (ad hoc) linking: circling an ink symbol generates a multi-way link to other instances of the same symbol.

KEYWORDS: Dynamic hypertext, information retrieval, paper-like user interface, pen computing, document metaphor, digital ink

INTRODUCTION

Implicit in the concept of hypertext is the assumption that people will read from computers rather than from paper. Yet, Dataquest predicts that in 1997 US office workers will print or copy over 10^{12} pages of paper [25]. This statistic suggests that, despite predictions of a paper-less office, US office workers continue to use paper, rather than computers, for much of their reading. Thus, we believe that consumers of today's electronic documents (including hypertexts) rely on two display technologies: computers for browsing and printed paper documents for reading.

Unfortunately, using paper and electronic hypertexts together can lead to disruptive transitions between reading and browsing. A user who finds an interesting article in a hypertext, and wishes to study it, will probably have to wait for it to print. Conversely, someone who reads a paper copy of an article has no convenient way to find additional information.

In *Proceedings of Hypertext'98* (June 20-24, Pittsburgh, PA), ACM Press, pp. 30-39

We propose a paper-like hypertext system that supports reading *and* browsing, and is compatible with paper documents. We do not suggest that a paper-like hypertext can provide all of the benefits of paper, or that computers should replace paper. To quote Andrew Dillon [13]:

> "… one should avoid seeing electronic text as a competitor to paper in some form of 'either-or' challenge for supremacy. It is not inevitable that electronic text will replace paper … but it might displace it as more and more human activities become mediated by information technology." (p. 3)

Thus our goal is not to replace paper, but to support information seeking better than the combination of paper and electronic texts does today.

When we say paper-like, we refer to paper documents, not to books. Information gatherers work more with piles of paper documents than with books [29, 35, 37]. Our introductory statistic and numerous studies of broader work practices (e.g., [43]) suggest that paper documents are an important means of working with information.

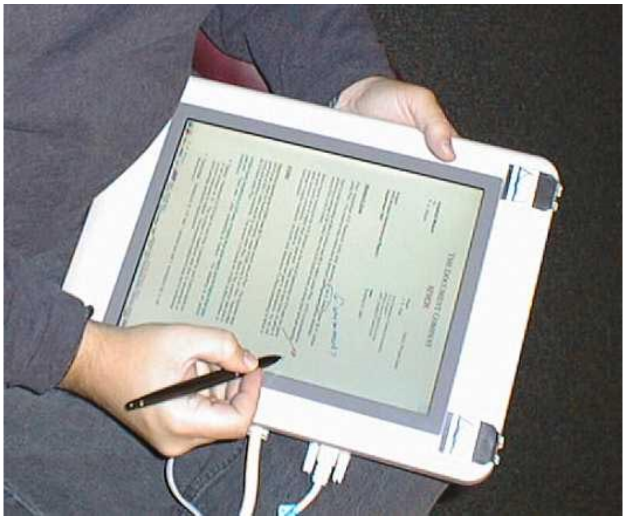

Figure 1: XLibris prototype in a reader's lap

Paper-like hypertexts can help information gatherers in several ways. First, support for reading and browsing in the same medium can eliminate disruptive transitions between paper and computer. Second, integrating reading with browsing makes it possible to enhance both activities. Finally, the benefits of hypertext become available to paper documents and to digital documents that can be printed onto paper (e.g., PostScript files).

We are exploring the concept of a paper-like hypertext in the context of the XLibris "active reading machine" [42]. XLibris supports reading with a rather literal paper document metaphor: XLibris displays a full image of a page on a device that readers hold in their hands and mark on with a pen (Figure 1). In addition, XLibris is compatible with both paper and digital documents: people can scan their paper documents, and "print" their digital documents (through a print driver that generates XLibris documents), into the system.

Where will the links in our hypertext come from? Static authored links will not suffice: many of the documents come from paper, or are designed for paper and lack explicit links. Therefore our goal is to support *trailblazing* [7], or reader-directed linking, without requiring much effort from the reader. Our solution is to create links based on the free-form ink annotations that readers make.

This paper describes several techniques for "linking by inking"—constructing links from readers' free-form ink marks. Linking by inking requires minimal effort from the reader, and augments existing work practice of annotating paper documents while reading. Readers' interests can be inferred from their annotations, translated into queries, and used to find, and to link to, related passages of text. Readers can also construct links explicitly with "ink anchors," circled marks that are matched against each other to compute many-to-many links.

The rest of this paper is organized as follows. We first discuss information seeking and the gap between reading and browsing in more detail. We then discuss how XLibris addresses problems with reading online, and how free-form ink can be used to support query-mediated and explicit linking. We then describe the status of our implementation and contrast paper-like hypertexts to open hypermedia systems. We finish with future work, related work, and conclusions.

COLLECTING INFORMATION VS. DEEP READING

People gather information to help them solve problems. For example, methodical consumers study *Consumer Reports* and compare manufacturers' specifications. ("Which fridge should I buy?") Lawyers search for legal precedents and conduct discovery. ("Why should the judge rule for my client?")

Information gathering combines searching, browsing, skimming, and deep reading. New information can be collected through browsing or through directed search. Readers tend to skim documents to select the most useful information or to get a rough sense of content. Finally, readers use deep reading to understand how they can solve their problem.

Information gathering is an iterative process [28, 35]. New information changes users' expectations of what they might find in the future. Conversely, deep reading can change users' understanding, and often leads to searching or browsing with a new topic in mind.

Interfaces for information gathering should support the constituent activities and the transitions between them. In practice, however, people move between browsing and searching online, and deep reading on paper:

> "Many of them [information analysts] relied strictly on annotative notes; that is, they'd make hardcopies of source materials, and mark up the pages…" ([29], p. 11)

> "Initially, library clients read through the pile of material returned from the search. In most cases, they read paper copies, even when the material had been delivered electronically." ([35], p. 443)

Although these studies are several years old, our informal observations suggest that information gatherers still do most of their reading on paper. As David Levy points out, digital libraries are being designed to support "shallower, more fragmented, and less concentrated reading" ([26], p. 202); deeper reading has been left out.

Transitions between online and paper media can create problems for information gatherers. Moving to paper is inherently slow and disruptive. People often need to go to another room to pick up printouts, and often spend time waiting for the printer. They can lose their focus while waiting, or may switch to an unrelated task to fill the time.

Moving from paper to the computer, on the other hand, loses valuable information. Annotations made on paper may reflect readers' interests, but the computer cannot use this information to help people when they shift from deep reading back to collecting information.

Shifting between media also changes how people gather information. Because computers support searching and browsing so well, people are encouraged to gather large amounts of information. Because computers do not support deep reading, however, information gatherers typically do not evaluate that information immediately. The result is a work practice called *information triage*: people collect many documents online, print them out, and evaluate them on paper [31]. Unfortunately, when people actually read what they have collected, they often realize that it is not what they were looking for.

READING ONLINE WITH XLIBRIS

Why is there a gap between reading and browsing? Why do people print out documents to read them on paper? Although the most commonly cited (and studied) reason for preferring paper is its higher display quality, research shows that reading from modern displays can be as fast as reading on paper [12]. Small differences in legibility or reading speed may exist, but this does not explain the strong preference for paper [12], and seems to us a poor explanation for why people wait for documents to print when they can view them online instantly.

Indeed, studies of work practice show that people do much more with paper documents than look at the words. Paper has many useful properties, or affordances, which are almost effortless to take advantage of but are awkward to replicate online [43]. In particular, an experimental comparison of deep reading online and on paper found numerous differences in how subjects worked with the text: in general, working on paper seemed far easier [36]. These affordances explain why information gatherers print documents to read them.

To address these problems, XLibris implements a *paper document metaphor* that supports three key affordances for deep reading. The metaphor includes both the hardware—a high-resolution pen display tablet that imitates the form factor and appearance of a paper document—and the software. The subsections below describe these affordances and how XLibris supports them.

Tangibility

Paper documents are tangible artifacts that can be manipulated directly.[1] Readers often move paper around to avoid glare, to speed up handwriting [21], or to adjust their perspective of a text [22]. In contrast, most computer monitors and laptop displays are stationary while in use, so that readers must move themselves—their heads, their bodies, and their arms—instead of the display. Even though an XLibris tablet is much heavier than most documents, readers can put it in their lap, tilt it to avoid glare, and hold it in their hands.

The tangibility of paper documents also enhances navigation. Turning paper pages seems easy and natural, and the weight and thickness of a paper document convey length and location: people often do not use page numbers when navigating around a document [36]. In XLibris, pressure strips on the case of the tablet (described in [23]) provide a tangible interface for page turning: pressing on one side of a sensor moves to the next page, and holding down starts riffling through pages. The harder you press, the faster you move. A quick animated transition indicates the direction of turning.

In addition, visual "location guides" drawn at the top of the page (see top right of Figure 2) give some feedback as to length and location. The location guides are similar to the page edges in the Book Emulator [5] but take up far less space. They were intended to be "obvious" (require no effort to perceive, as recommended in [34]), but we did not achieve this goal: the graphic design needs work, and we may need to add parameterized sounds [17] for page turns.

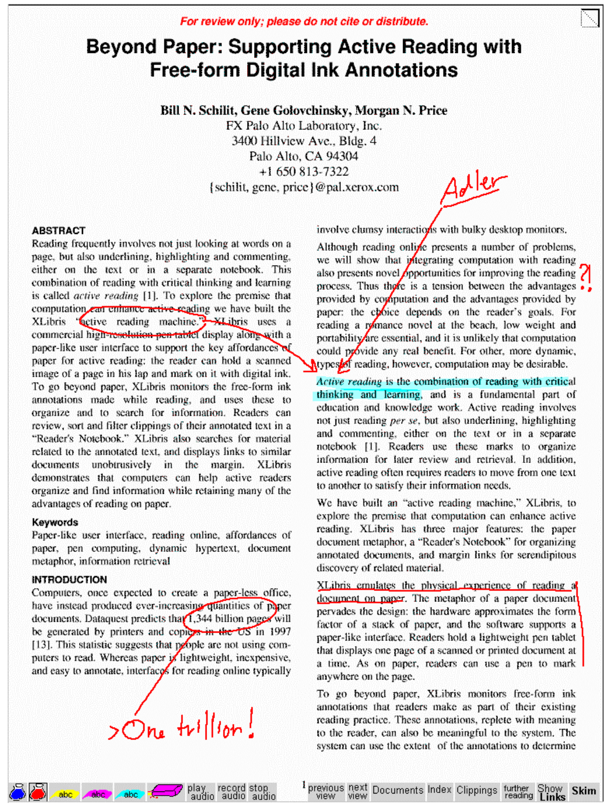
For review only; please do not cite or distribute.

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

Adler

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* [1]. 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 [13]. 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.

>One trillion!

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 other, 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 commenting, either on the text or in a separate notebook [1]. Readers use these marks to organize information for later review and retrieval. In addition, active reading often requires readers to move from one text to another to satisfy their information needs.

We have built an "active reading machine," XLibris, to explore the premise that computation can enhance active reading. XLibris has three major features: the paper document metaphor, a "Reader's Notebook" for organizing annotated documents, and margin links for serendipitous discovery of related material.

XLibris emulates the physical experience of reading a document on paper. The metaphor of a paper document pervades the design: the hardware approximates the form factor of a stack of paper, and the software supports a paper-like interface. Readers hold a lightweight pen tablet that displays one page of a scanned or printed document at a time. As on paper, readers can use a pen to mark anywhere on the page.

To go beyond paper, XLibris monitors free-form ink annotations that readers make as part of their existing reading practice. These annotations, replete with meaning to the reader, can also be meaningful to the system. The system can use the extent of the annotations to determine

abc abc abc play audio record audio stop audio 1 previous view next view Documents Index Clippings further reading Show Links Skim

Figure 2. An annotated page in XLibris

Free-form Ink Annotation

Readers of paper documents can mark anywhere on the page, at any time. Annotation is an essential part of deep reading: the act of annotation with pencil, pen or highlighter helps readers to think, to remember, or even to stay awake [2, 27, 30, 36, 37]. In short, readers write [1,2].

Annotation on paper requires very little effort. Although interfaces for annotating with a keyboard and mouse do exist, they typically require more effort than scribbling with a pen. Furthermore, annotations and documents are often less distinct visually online than on paper: text annotations are not as salient as ink marks, and their presence may cause the document text to be reformatted [36, 43].

Although annotation is the most common way for information gatherers to record information during reading [29, 35], many people also take notes separately, on paper[2] [29, 37]. Readers of library books usually take notes, although they also use removable annotations—colored Post-It notes [29]—and photocopy pages out of the books

---

1 Tangibility is distinct from portability, which is the ability to carry information around, and which many computers do support. Although portability is essential for student researchers collecting information [37], we believe that tangibility is a more general requirement.

2 Two of 25 subjects in [37] took notes on laptops; the rest used paper.

so that they can annotate them or so that they can carry the information around [37]. Notes are more powerful than annotations in some ways [37]: taking notes forces readers to summarize their reading, which typically requires more effort than annotation, but may lead to deeper thinking about the material. Also, notes turn material from disparate sources into a single source for later review, and notes give people more space to work with their ideas.

XLibris, like paper, allows readers to make free-form ink marks, anywhere on the page, at any time. XLibris provides two colors of ink, three highlighter pens, and a stroke-based eraser. XLibris confers some of the benefit of notes onto annotations with a "Reader's Notebook" that displays a list of "clippings" computed from annotations [42]. Finally, XLibris includes a blank notebook in which readers can write when they need more space.

Fixed Page Layout

Paper documents are laid out on fixed-size pages. Page layout communicates the document "genre" (e.g., business letter versus technical article) and the location of important information such as the title of the document. The fixed layout also supports spatial memory that helps readers find old information [22, 12]. In summary, paper pages give readers an excellent overall "sense of the document."

When reading online, however, page layout information is often lost. Web browsers reflow the text on a "page", resulting in global changes to the layout. Hypertext versions of paper documents lead to disorientation and a sense of being "lost" [14]. Finally, most monitors cannot display a full page legibly, so documents in page-oriented formats cannot be read online without awkward scrolling and zooming [36].

XLibris respects the benefits of paper and dedicates the entire display to the image of a single page.[3] (The location guides at the top and a button bar at the bottom of the page used to invoke commands are minor exceptions.) The tablet is used in portrait mode to match the aspect ratio of typical paper documents.

LINKING BY INKING

XLibris supports reading and browsing over a wide variety of documents. Many of these documents were designed for presentation on paper, and lack hypertext links. One way to create a browsing environment is to automatically detect topics, references and other static document structure, and then to generate the links automatically [3]. Unfortunately, only a small fraction of all possible links are useful to a particular reader at a particular time, and deciding on the right links ahead of time—without any knowledge of a reader's activities—is difficult.

[3] Although displaying one page at a time limits readers' ability to compare multiple pieces of information, the alternative is to divide the display into windows, breaking the paper document metaphor. Furthermore, the effort of managing windows can interfere with reading [36]. We will return to the issue in the future work section.

Our approach is to give readers tools for constructing links. Links computed from readers' input should be more useful than statically computed links because they can reflect readers' idiosyncratic interests. At the same time, link construction must not interfere with readers' primary activity—reading.

XLibris minimizes the effort required to create links by computing them from free-form ink marks. As observed above, annotation is already a common reading practice, so link construction need not disrupt reading. Free-form ink is also a modeless interface: readers do not need to switch between reading and linking. The rest of this section describes three types of reader-directed links—margin links, further reading lists, and ink anchors—in more detail.

Margin Links

Serendipity is a common and rewarding experience when browsing in libraries. People looking for a particular book on the shelf often find unintended yet valuable material nearby. Though this may seem accidental, it actually reflects the organization of books by topic, and the large amount of information available. Paper magazines and journals sometimes offer this benefit as well.

XLibris' margin links provide this kind of serendipity to readers as they read. As a reader annotates a document, the system performs queries, and displays links to related pages (see Figure 3). The interface is designed to avoid disrupting reading: only links to highly related material are presented; the links appear in the margin; and at most one margin link results from each annotation. Readers who are not interested in finding related material can simply ignore the links or they can turn off automatic link creation.

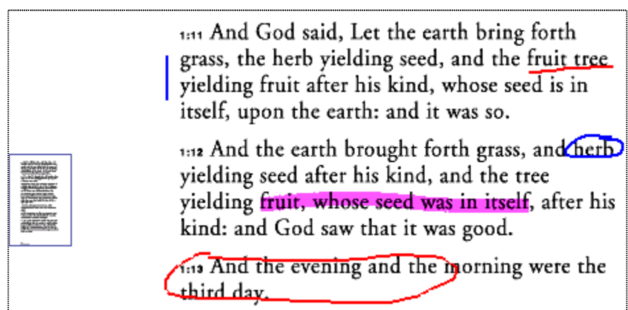

Figure 3. A margin link (rectangle on the left) computed from a highlighted phrase

Presentation. Anchors are drawn as thumbnail images of target pages, including shrunk annotations (if any). This allows readers to recognize familiar landmarks, and to find links that they previously followed.[4] The topic of the link

[4] Thumbnail anchors are less effective for links to new documents: although thumbnails communicate the document genre, they take too much space. Footnotes with titles of new documents, or smaller, purely symbolic anchors, might be more appropriate for this case.

is implicitly described by the annotated text that produced it. On the target page, a reverse link is presented near the passage that best matches the query, and terms from the query are underlined in color wherever they appear.

Query Construction. XLibris classifies annotations into several categories. Circled phrases, underlines, and highlights select specific phrases; circled passages select corresponding sentences; and margin bars select nearby passages.

These selections are then converted into queries. Annotations that select a phrase result in weighted queries based on the entire surrounding sentence, with the emphasis on the selected words. (Padding the query with context words helps to avoid irrelevant passages that include the selected words but use them differently.) Annotations that select one or more sentences result in queries based on the entire selected passage.

Similar techniques for query-mediated linking (computing links through full-text search) have been shown to be effective. VOIR, a dynamic hypertext system, selected terms from users' queries to create link anchors; clicking on an anchor caused the system to expand the previous query with the terms from the anchor's context. The system provided a browsing interface with two– to three–second response times for a 250MB collection (approximately 75,000 nodes). Experimental evaluation of query-mediated linking showed that such automatically-constructed queries performed as well as queries typed by users (recall and precision means for automatic queries were higher, but not reliably so), and better than un-weighted queries based on passages selected by users [19].

Further Reading Lists

When readers reach the end of a document, they often need to know more. The document may not emphasize the most interesting aspect of the topic, or the document may spark an interest in a new topic, without providing enough depth or detail to satisfy that interest. For example, articles in *Scientific American* typically describe an area of research at a high level, without providing much technical detail. These articles include a further reading list to help readers go into more depth (Figure 4).

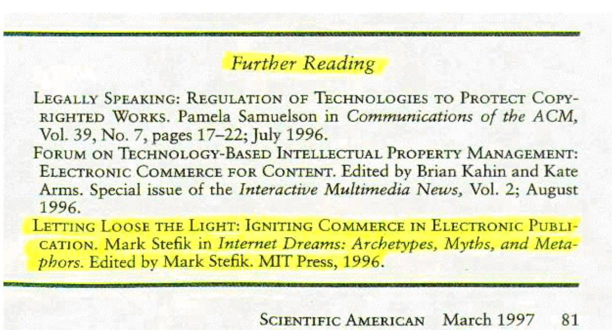

Further Reading

LEGALLY SPEAKING: REGULATION OF TECHNOLOGIES TO PROTECT COPYRIGHTED WORKS. Pamela Samuelson in *Communications of the ACM*, Vol. 39, No. 7, pages 17–22; July 1996.

FORUM ON TECHNOLOGY-BASED INTELLECTUAL PROPERTY MANAGEMENT: ELECTRONIC COMMERCE FOR CONTENT. Edited by Brian Kahin and Kate Arms. Special issue of the *Interactive Multimedia News*, Vol. 2; August 1996.

LETTING LOOSE THE LIGHT: IGNITING COMMERCE IN ELECTRONIC PUBLICATION. Mark Stefik in *Internet Dreams: Archetypes, Myths, and Metaphors*. Edited by Mark Stefik. MIT Press, 1996.

SCIENTIFIC AMERICAN March 1997 81

Figure 4: A further reading list from a *Scientific American* article

This example illustrates the limitations of authored further reading lists. The article describes a technological vision for protecting digital copyrights. Consider a reader with a narrower interest, such as "how digital copyright affects librarians." The reference to a six-page article on digital copyright law may help track down this information, but the reference is unlikely to answer their question directly, and may well be out of date.

XLibris automatically generates further reading lists for each document (Figure 5). Unlike static references, these lists reflect the interests of a specific reader. As with margin links, readers' interests are inferred from annotations, and no additional effort is required to construct these links.

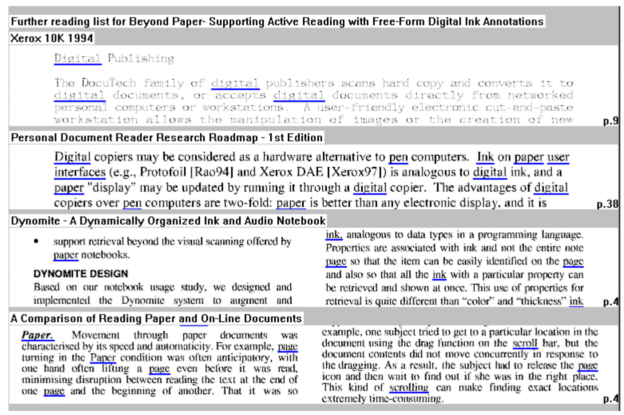

Further reading list for Beyond Paper- Supporting Active Reading with Free-Form Digital Ink Annotations

Xerox 10K 1994

Digital Publishing

The DocuTech family of digital publishers scans hard copy and converts it to digital documents, or accepts digital documents directly from networked personal computers or workstations. A user-friendly electronic cut-and-paste workstation allows the manipulation of images or the creation of new p.9

Personal Document Reader Research Roadmap - 1st Edition

Digital copiers may be considered as a hardware alternative to pen computers. Ink on paper user interfaces (e.g., Protofoil [Rao94] and Xerox DAE [Xerox97]) is analogous to digital ink, and a paper "display" may be updated by running it through a digital copier. The advantages of digital copiers over pen computers are two-fold: paper is better than any electronic display, and it is p.38

Dynomite - A Dynamically Organized Ink and Audio Notebook

- support retrieval beyond the visual scanning offered by paper notebooks.

DYNOMITE DESIGN

Based on our notebook usage study, we designed and implemented the Dynomite system to augment and

ink, analogous to data types in a programming language. Properties are associated with ink and not the entire note page so that the item can be easily identified on the page and also so that all the ink with a particular property can be retrieved and shown at once. This use of properties for retrieval is quite different than "color" and "thickness" ink p.4

A Comparison of Reading Paper and On-Line Documents

Paper. Movement through paper documents was characterised by its speed and automaticity. For example, page turning in the Paper condition was often anticipatory, with one hand often lifting a page even before it was read, minimising disruption between reading the text at the end of one page and the beginning of another. That it was so

example, one subject tried to get to a particular location in the document using the drag function on the scroll bar, but the document contents did not move concurrently in response to the dragging. As a result, the subject had to release the page icon and then wait to find out if she was in the right place. This kind of scrolling can make finding exact locations extremely time-consuming. p.4

Figure 5: A further reading list in XLibris

Presentation. Further reading lists are presented in a separate multi-page view associated with the source document. Readers can access this view at any time, whether or not they have annotated the current document. Of course, with more annotations, the links reflect readers' interests better.

Each view shows a list of links to related documents. Target anchors are displayed as clippings that retain visual cues from target documents, and also include some surrounding context. Clippings help readers understand the destination of the link, and also make the passage more recognizable if the reader follows the link. As with margin links, terms matching the query are underlined. Readers need not follow the link to one of the target documents if the "clippings" view provides sufficient information: a tap of the pen will return the system to the previous view.

Query Construction. Each annotation is interpreted as a text selection and transformed into a list of word weights. The weights for each word are then summed across all annotations of the document. If several annotations select the same occurrence of a word, then the maximum weight for that occurrence is used. Queries longer than about thirty words are truncated because the quality of search results may degrade if too many terms are added [24], and also to minimize the query's running time. If the query is short, then terms that are most characteristic of the document are added to the query with a low weight. The padding prevents a query based on a few annotations from returning documents that are entirely unrelated (due to polysemy, for example).

We have begun to evaluate the effectiveness of further reading lists. We found that queries derived from ink annotations had significantly higher precision and recall than relevance feedback operations performed by the same users on the same documents. Although these results are preliminary, and cannot be generalized to all documents and all annotations, they suggest, nonetheless, that links derived from users' ink marks can be an effective interface for some information exploration tasks.

Ink Anchors

Readers often link documents together for their own particular reasons. One typical practice is to mark places on different pages with the same circled symbol, thereby creating a logical link between them. We propose a design for "ink anchors" to help readers create idiosyncratic many-to-many links with minimal effort. Circling any mark converts that mark into an anchor. The system maintains multi-way links between similar ink anchors.

Ink anchors may have a number of advantages over traditional link construction interfaces. Ink anchors are personal and idiosyncratic. Readers can use the flexibility of free form ink to make anchors meaningful. Furthermore, ink anchors are an extension of current practice: readers already use ink marks, post-it notes, colored stick-ons, and other types of annotations to categorize information efficiently [27, 31].

Ink anchors should require relatively little effort to create. Readers should be able to create links without selecting from a list, navigating, or typing text. They would merely need to recreate their own, personal marks. Although readers may forget anchors that they use infrequently, we argue that this is preferable to alternatives such as presenting readers with a large number of choices every time they build a link.

An evaluation of ink matching for the task of finding names in a personal phone book supports these predictions [16]. Users rated ink matching higher than browsing or text search, and used it more frequently when given the choice, because it required less effort. Users preferred ink matching despite problems remembering which "name" they gave an entry.

Presentation. Tapping on a source ink anchor should produce a list of clippings that contain matching target ink anchors (Figure 6). This design is based on our interface for reviewing annotations in a Reader's Notebook view [42], which shows clippings of documents that correspond to annotations made by the reader. The clippings include enough context to make the ink marks meaningful, but still present a concise view of annotated documents.

Clippings include nearby annotations, and therefore may show multiple ink anchors in the same clipping (e.g., third clipping in Figure 6). This allows readers to observe connections between different sets of anchors without following the links.

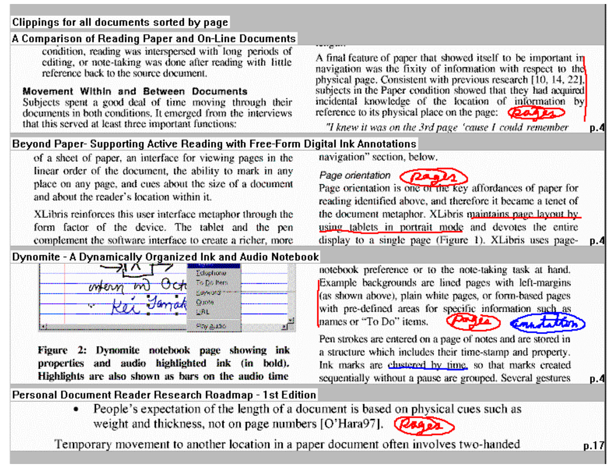

Figure 6. A mockup of the result of following an ink anchor

Ink Anchor Matching. Ink matching has been shown to work fairly reliably for modest quantities of one person's handwriting. Error rates as low as 3% on a database of 200 words have been reported [16].

Our design provides feedback on the results of the ink matching by replacing the newly-drawn mark with the match. Erroneous matches can be corrected by erasing and rewriting the mark; unintended anchors can be removed by erasing the circle that created the mark. We plan to indicate the first instance of an ink anchor by increasing the thickness of the circle that identified it.

Finally, we plan to help readers remember the "names" they've been using for their links by providing a list of clippings over all ink anchors, sorted by time. For succinctness, this view might contain only the most recent occurrence of each anchor.

IMPLEMENTATION

All of the features described in this paper, except ink anchors, have been implemented in XLibris. XLibris runs on Windows 95/NT, uses the Win32 API, and is implemented in C++. 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 can use a landscape device in portrait mode.

XLibris uses an "image + text" file format [39, 6]. With this format, XLibris can handle most digital documents, and also documents that came to the user on paper. The text, inferred from print commands or from optical character recognition (OCR) applied to the scanned image, is tied to the image by a bounding rectangle (in pixels) for each word. This text is indexed to support information retrieval, but is not shown to the reader directly.

In principle, the “image + text” file format lets users load documents into the system easily, by “printing” or by scanning. Although we have a working Windows 95 print driver that can print into XLibris, the image quality is not yet acceptable. (The figures in this paper are based on the output of a PostScript processor [6] that requires too much effort for real use.) Scanning paper documents into XLibris is far too slow (perhaps a minute per page) because we do not display the document until the OCR completes. Once we can load many documents into XLibris easily, the system will need scalability improvements, both to the document management system and to the simple interface for organizing documents described in [42].

Hardware

XLibris runs on two kinds of display tablets: standalone pen computers, and tablets connected to conventional computers by a tether. Our current platform is a Mutoh pen tablet display tethered to a PC (shown in Figure 1) [33]. The display has a 12.1" (30.7 cm) diagonal and a resolution of 768x1024 when rotated. XLibris displays a 100 dpi full-color image of an 8.5 x 11" page at roughly full size with clipped margins. 100 dpi is high enough resolution for 10-point Times fonts (such as you are reading now) to be legible (see Figure 2, which shows a 768x1024 screenshot, albeit reduced).

Although pen computers will not match the display quality or portability of paper documents in the foreseeable future, we believe that far more comfortable devices will soon exist. Much lighter, standalone pen computers and 140-dpi color LCD displays are already available, and both technologies are improving rapidly.

Link Construction

The Verity search engine is used to index and retrieve documents. Although the current implementation stores all indexed documents on local disk, queries could be run on remote databases as well. Running queries on external databases and fetching documents as needed would allow the use of XLibris as an interface to a digital library. In addition, some obvious improvements to the margin links implementation are required. Queries should run in the background, and queries should be derived from groups of nearby ink strokes, rather than from single strokes.

PAPER-LIKE HYPERTEXT VS. OPEN HYPERMEDIA

The hypertext community has placed some emphasis on open hypermedia systems in the last few years. The definition of open hypermedia systems calls for native data formats, interoperability with all tools on a given platform, support for distributed processes, lack of artificial distinction between authors and readers, and the ability to incorporate new functionality [11].

In contrast to open hypermedia systems, we require data to be imported into our system. We do not expect to handle every kind of document because not all documents are meant to be *read*. Documents that are appropriate, however, can be imported easily from any application that uses standard print driver interfaces. Reliance on this interface assures us of full compatibility with all well-behaved applications.

By taking this stance, we consciously deny the user the ability to interact with dynamic aspects of documents, as would be possible if native data formats were respected. Although this assumption makes our approach less than universal, we believe that a large proportion of what people read in their daily work is, in fact, static. Published research papers, on-line documentation, legal briefs, newspaper articles, patents, and documents filed with government agencies are just a few examples of documents that, once created, remain fixed. When presented in electronic form, most of the interaction with such documents is navigational in nature, rather than some manipulation or transformation of logical elements of a document.

The focus on static page-oriented documents has advantages for users. The paper document metaphor supports a uniform style of interaction. Thus, any document may be marked up, any portion of any document may be linked manually to any portion of any other document through ink anchors, and full-text queries can be used to search for any document.

Open hypermedia systems, on the other hand, while imposing requirements on the infrastructure, are mute about interfaces for following (or even seeing!) the links that users are likely to encounter. Davis *et al.* [10] describe several levels at which linking has been integrated into applications in support of open hypermedia systems. Each of these levels allows the application complete freedom, but introduces some mechanism for selecting and traversing links. Thus, instead of a consistent, predictable interface that integrates documents, the user is confronted with an endless variety of interfaces, many of which are poorly suited to their task due to the constraints imposed by existing applications.

FUTURE WORK

We plan to implement ink anchors and to correct the problems with our prototype described above. We will then be ready to evaluate the use of a paper-like hypertext in a realistic setting. We would also like to extend XLibris to support information exploration and multiple document tasks directly.

Although XLibris supports the transitions between deep reading and information exploration, the prototype does not support information exploration directly. For example, there is no explicit search interface. And, as we have said, the current prototype cannot access remote document repositories.

We would also like to extend XLibris to support multiple document tasks. Document work often involves the use of several documents together [1]. In the context of gathering information, people need to compare sources [35], to take notes [37], and to write reports [29, 35]. Standard windowing systems are too cumbersome to support these activities effectively [36], so we must consider other alternatives. One possibility is the use of additional displays in conjunction with an XLibris tablet. To support authoring, we will need to integrate XLibris with desktop word processors.

RELATED WORK

The main focus of this paper is the creation of query-mediated and ad hoc links from readers' free-form ink annotations. Previous interfaces for query-mediated linking have not used ink annotation, and were not designed to support reading. VOIR [20, 19] computed links from keywords or passages selected with a mouse to support information exploration. The Remembrance Agent [41] computed links from the last few lines typed into a text editor to support writing.

There has been little work in the area of ad hoc links and pen computing. GO PenPoint [9] included a pen-based interface for ad hoc linking: a link gesture created a "button" that pointed to the current page; it could then be dragged to another page to complete the link. PenPoint also included another interface for linking, where a gesture on the source page created a link to the most recent selection [9,32]. Although PenPoint and XLibris share similar goals of reducing the effort required to create links, PenPoint did not create ad hoc links from free-form ink.

XLibris' use of circling to identify ink anchors is similar to ink keywords in Marquee, a real-time tool for video logging [45]. A handwritten word could be circled to copy it to a palette of ink keywords, from which it could then be dragged onto other pages. These ink keywords were later converted to text manually in order to support search. Although a later version used ink matching to highlight matching notes [40], ink matching was neither combined with the keywords nor used to create links.

Dolphin creates hierarchical links from an ink gesture: boxing an area on the page creates a "button" or anchor to a new page that is beneath the current page in a hierarchy [44]. This hierarchical link differs from the query-mediated and ad hoc links described here.

Dynomite [47] lets users apply properties to free-form ink in a notebook, and displays "ink index" pages showing all notes with a particular property, linked to the full pages. The ink index is similar to the Reader's Notebook feature of XLibris [42] described in passing, except that Dynomite supports note taking rather than reading documents. The ink index cannot support ad hoc linking because only a few properties are available directly.

The XLibris paper document metaphor is related to the book metaphor used in many hypertext systems, such as the Book Emulator [5] and the Web Forager [8] (see [8] for a review). There are some subtle yet important differences between these metaphors: documents are annotated, disposable parts of workflow; whereas books are clean, organized, permanent, and are more likely to be used for reference or for pleasure. Although some hypertexts based on a book metaphor do support annotation (e.g., [5]), the purpose of a metaphor is to shape user's expectations, and we believe that a paper document metaphor may be more appropriate than a book metaphor for a free-form ink-based system.

SuperBook [15] and DynaBook [46] use hierarchical indexes and table of content views to support fact-finding, or answering questions from reference material. They share XLibris' high-level goal of improving a particular kind of work by emulating the existing paper artifacts and then improving on them. These systems, however, were designed for very different tasks, did not use free-form ink, and did not implement the paper document metaphor.

Finally, paper-based hypertexts are a complementary approach to bringing links to paper documents. For example, in the PaperLink system a highlighter pen containing a camera can be used to select a word or phrase on a paper document [4]. The user can then create and follow links from the paper highlight to electronic documents.

CONCLUSIONS

We began with the premise that a paper-like hypertext can support reading and browsing. We have described how a paper document metaphor overcomes several problems with reading online: readers hold the image of a page in their lap and mark on it with digital ink, much as they would on paper. With better support for reading online, we hope to minimize the disruptive transitions between collecting information online and deep reading on paper.

We have shown how free-form digital ink annotations can be used as an interface for both query-mediated and *ad hoc* linking. In this way, annotations made while reading enhance browsing. In addition, linking by inking allows hypertexts to include paper-like documents that do not contain static links.

We have begun to evaluate our designs experimentally. This process will include some controlled experiments, but we expect to obtain much useful feedback about our designs through more naturalistic observations. Only through open-ended evaluation can we understand the impact of paper-like hypertext on real work.

ACKNOWLEDGMENTS

We thank Mark Weiser for discussions of the "further reading list" concept, and Anuj Gujar, Beverly Harrison, and Roy Want for integrating the pressure sensors, and Bernard Mont-Reynaud for early discussions of ink anchors. Comments from Cathy Marshall, Eleanor Rieffel, and from the anonymous reviewers improved this paper. We thank the PDR group at PARC and at XRCE Cambridge for many interesting discussions. 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. Allan, J. (1995) Automatic Hypertext Construction. Ph.D. thesis, Cornell University, January 1995.

4. Arai, T., Aust, D. Hudson, S.E. (1997) PaperLink: A Technique for Hyperlinking from Real Paper to Electronic Content. In *Proceedings of CHI97* (Atlanta, GA, March 1997), ACM Press. pp. 327-334.

5. 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.

6. Birrell, A. and McJones, P. *Virtual Paper*. Available at http://www.research.digital.com/SRC/virtualpaper/home.html

7. Bush, V. (1945) As We May Think. *Atlantic Monthly*, July, pp. 101-108.

8. Card, S.K., Robertson, G.G., York, W. (1996) The WebBook and the Web Forager: An Information Workspace for the World-Wide Web. In *Proceedings of CHI '96*, ACM Press. pp. 111-117.

9. Carr, R. and Shafer, D. (1991) *The Power of PenPoint™*. Addison-Wesley, Reading, Massachusetts.

10. Davis, H.C., Knight, S. and Hall, W. (1994) Light Hypermedia Link Services: A Study of Third Party Application Integration. In *Proceedings of Hypertext '94* , (Edinburgh, Scotland, September 1994), ACM Press. pp. 41-50.

11. Davis, H.C., Hall, W., Heath, I., Hill, G., Wilkins, R. (1992) Towards an Integrated Information Environment with Open Hypermedia Systems. In *Proceedings of ECHT '92*, pp. 181-190.

12. Dillon, A. (1992) Reading from paper versus screens: a critical review of the empirical literature. *Ergonomics*, 35(10), pp. 1297-1326.

13. Dillon, A. (1994) Designing Usable Electronic Text: Ergonomic Aspects of Human Information Usage. Taylor & Francis: London, 1994.

14. Dillon, A., Richardson, J., and McKnight, C. (1990) Navigation in hypertext: a critical review of the concept. In Diaper, D., Gilmore, D., Cockton, G., and Shackel, B. (Eds) *INTERACT '90*. North Holland, Amsterdam.

15. Egan, D.E., Remde, J.R., Gomez, L.M., Landauer, T.K., Eberhardt, J., and Lochbaum, C.C. (1989) Formative Design-Evaluation of SuperBook. ACM TOIS 7,1, January 1989, pp. 30-57.

16. Frohlich, D. and Hull, R. (1996) The usability of scribble matching. In *Conference Companion of CHI96*, ACM Press. pp. 189-190.

17. Gaver, W.W. (1993) Synthesizing Auditory Icons. In *Proceedings of INTERCHI 93* (24-29 April '93), pp. 228-235

18. Golovchinsky, G. (1997a) Queries? Links? Is there a Difference? In *Proceedings of CHI97*, (Atlanta, Georgia, March 1997), ACM Press, pp. 407-414.

19. 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.

20. Golovchinsky, G. and Chignell, M.H. (1997) The newspaper as an information retrieval metaphor. *Information Processing and Management*, 33 (5), pp. 663-683.

21. 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.

22. Haas, C. (1996) *Writing Technology: Studies on the materiality of literacy*. Lawrence Erlbaum Associates, Mahwah, NJ.

23. 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.

24. Harman, D. (1992) Relevance feedback revisited. In *Proceedings of SIGIR '92*, ACM Press, pp. 1-10.

25. Johnson, A. (1997) Fax Trends. *Office Equipment and Products*, July 1997, p. 25.

26. Levy, D. (1997) I Read the News Today, Oh Boy: Reading and Attention in Digital Libraries. In *Proceedings of DL97*, (Philadelphia, PA), ACM Press. pp. 202-211.

27. 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, pp. 163-169.

28. Marchionini, G. (1995) *Information Seeking in Electronic Environments*. Cambridge University Press, 1995.

29. Marshall, C.C. (1990) Work Practices Study: Analysts and Notetaking. Unpublished report, 31 May 1990.

30. Marshall, C.C. (1997) Annotation: from paper books to the digital library. *In Proceedings of Digital Libraries '97* (Philadelphia, PA, July 1997), ACM Press, pp. 131-140.

31. Marshall, C.C. and Shipman, F.M. (1997) Spatial Hypertext and the Practice of Information Triage. In *Proceedings of Hypertext '97,* Southampton UK, ACM Press, p. 124-133.

32. Meyrowitz, M. (1991) Hypertext and Pen Computing. In *Proceedings of Hypertext '91*, p. 379.

33. Mutoh (1997) Mutoh America Inc. Available at www.mutoh.com

34. Nygren, E., Lind, M., Johnson, M., and Sandblad, B. (1992) The Art of the Obvious: Automatically processed components of the task of reading frequently used documents. Implications for task analysis and interface design. In *Proceedings of CHI 92* (May 3-7, ACM), pp. 235-239.

35. O'Day, V.L. and Jeffries, R. (1993) Orienteering in an Information Landscape: How Information Seekers Get From Here to There. In *Proceedings of INTERCHI '93*, ACM Press, pp.438-445.

36. 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.

37. O'Hara, K., Smith, F., Newman, W., Sellen, A. (1998) Student Reader's Use of Library Documents: Implications for Digital Library Technologies. To appear in *Proceedings of CHI98*, ACM Press.

38. Pearl, A. (1989) Sun's Link Service: A Protocol for Open Linking. In *Proceedings of Hypertext '89*, (Pittsburgh, PA, November 1989), ACM Press, pp. 137-146.

39. 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.

40. Poon, A., Weber, K., Cass, T. (1995) Scribbler: A Tool for Searching Digital Ink. In *CHI95 Conference Companion*, ACM Press. pp. 252-253.

41. Rhodes, B.J. and Starner, 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.

42. Schilit, B.N., Golovchinsky, G. and Price, M. Beyond Paper: Supporting Active Reading with Free-form Digital Ink Annotations. (1998) In *Proceedings of CHI98*, ACM Press.

43. 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.

44. Streitz, N.A., Geißler, J., Haake, J.M., and Hol, J. (1994) DOLPHIN: integrated meeting support across local and remote desktop environments and LiveBoards. In CSCW '94, *Proceedings of the Conference on Computer supported cooperative work*, pp. 345-358.

45. Weber, K. and Poon, A. (1994) Marquee: A Tool for Real-Time Video Logging. In *Proceedings of CHI94*, (April 1994, Boston, MA), ACM Press, pp. 58-64

46. Weyer, S.A. (1982) The design of a dynamic book for information search. Int. J. Man-Machine Studies, 17, pp. 87-107.

47. Wilcox, L.D., Schilit, B.N., and Sawhney, N. (1997) Dynomite: A Dynamically Organized Ink and Audio Notebook, In CHI 97 Conference Proceedings, ACM Press, pp. 186-193.