Digital Library Information Appliances

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

FX Palo Alto Laboratory
3400 Hillview Avenue
Palo Alto, CA 94304, USA

E-mail: {schilit,price,gene}@pal.xerox.com
+1 650 813-7220

ABSTRACT
Although digital libraries are intended to support education and knowledge work, current digital library interfaces are narrowly focused on retrieval. Furthermore, they are designed for desktop computers with keyboards, mice, and high-speed network connections. Desktop computers fail to support many key aspects of knowledge work, including active reading, free form ink annotation, fluid movement among document activities, and physical mobility. This paper proposes portable computers specialized for knowledge work, or *digital library information appliances*, as a new platform for accessing digital libraries. We present a number of ways that knowledge work can be augmented and transformed by the use of such appliances. These insights are based on our implementation of two research prototype systems: XLibris, an “active reading machine,” and TeleWeb, a mobile World Wide Web browser.

KEYWORDS: Digital library, information appliance, paper document metaphor, active reading, browsing, information exploration, digital ink, pen computing, mobile Web browser, mobile computing, network computer.

INTRODUCTION
The standard platform for accessing today’s digital libraries is a desktop computer with a keyboard, a mouse, information retrieval software, and a network connection. Unfortunately, desktop computers are not suitable for many types of knowledge work. Studies of people using electronic and paper documents show that desktop systems fail to support four important activities:

- Active reading
- Free form ink annotation
- Fluid movement among document activities
- Physical mobility

As a result of focusing on the desktop, today’s interfaces for accessing digital libraries support retrieval in isolation, without addressing these broader work practices.

In Proceedings of Digital Libraries ’98 -
Third ACM Conference on Digital Libraries
Pittsburgh, PA, USA, June 23-26, 1998

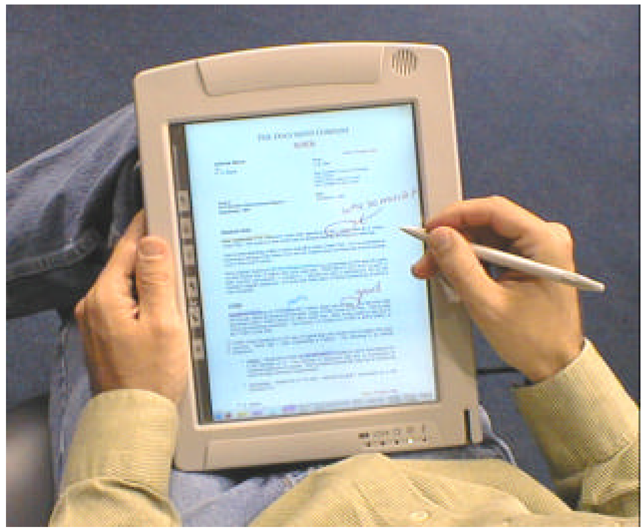

Figure 1: XLibris prototype in a reader's lap.

We propose a new platform for accessing digital libraries: specialized portable computers designed for knowledge work, or *digital library information appliances.*[1] In this paper, we draw on our experience designing and building two research prototype information appliances—the XLibris “active reading machine” and the TeleWeb mobile Web browser—to show how these devices can support a broad range of activities in the digital library.

READING AND DIGITAL LIBRARIES
The first and most fundamental document activity is reading. “Reading” encompasses a broad range of complex and poorly understood practices involving documents, including skimming, searching, browsing, speed-reading, surfing, reviewing, and rereading. Although we often take it for granted, reading can be hard work. We use the term *active reading* [2] to distinguish this rich collection of activities from simply looking at words on a page.

Desktop computers do not support active reading. Once expected to create a paperless office, computers have instead produced ever-increasing quantities of paper documents. Dataquest predicts printers and copiers will generate over

[1] Jeff Raskin coined the term “information appliance” as a small system mainly intended to perform one task [22].

10[12] pages in the U.S. in 1997 [13]. This statistic suggests that people do not use computers to read. The reason for this is obvious: paper supports extended, focused, deep reading practices better [23].

Current research on digital libraries does not address a broad range of reading activities. Instead, as David Levy points out, digital libraries support shallower, more fragmented, and less concentrated reading [15]. The research focus in this field is on tools for search, selection, and distillation, tools that promote identification and extraction of information fragments.

Because digital library interfaces do not support active reading, people search online and then print documents to read them. This “search and print” reality contrasts with the “search and read” ideal held by researchers (e.g., [8]).

Although a dual online/paper system is not *a priori* undesirable, we will show the many advantages of an online system that enables a variety of reading practices. The XLibris active reading machine supports a broad class of reading activities by employing a “paper document metaphor” (as opposed to the desktop metaphor) that imitates the physical experience of reading and marking on paper (Figure 1). The paper document metaphor follows lessons learned from a number of studies comparing reading online to reading on paper, summarized below.

Tangibility

Readers often move paper documents to avoid glare, to speed up handwriting [10], or to adjust their perspective of a text [11]. In contrast, most computer displays are stationary while in use, so that readers must move themselves—their heads, their bodies, and their arms—rather than their display. XLibris supports paper document-like tangibility by running on tablet displays.[2] Although XLibris has not been deployed for real use, people say that holding a page-sized display in their lap and being able to easily reposition it changes their online reading experience.

The tangibility of paper documents also supports navigation. Turning paper pages seems easy and natural, and the weight and thickness of a paper document convey length and location [23]. In XLibris, pressure strips on the case of the device (see [12]) provide a tangible interface for page turning: pressing on one side of a sensor moves to the next or previous page, and holding down initiates riffling through pages. The harder you press, the faster you move. A quick animated transition indicates the direction of turning. Finally, “location guides” of varying thickness drawn in the top cor-

[2] Over the course of the project, we have used a number of hardware platforms. XLibris prototype 2 runs on a commercial Mutoh “pen tablet display” that is a 1024x768 color LCD panel packaged with a digitizing pen tablet [21]. This device has a cable that connects to a traditional computer, and replaces the monitor and mouse. XLibris prototype 3 is based on a self-contained Fujitsu Point 510 pen computer with an 800x600 color display, shown in Figure 1.

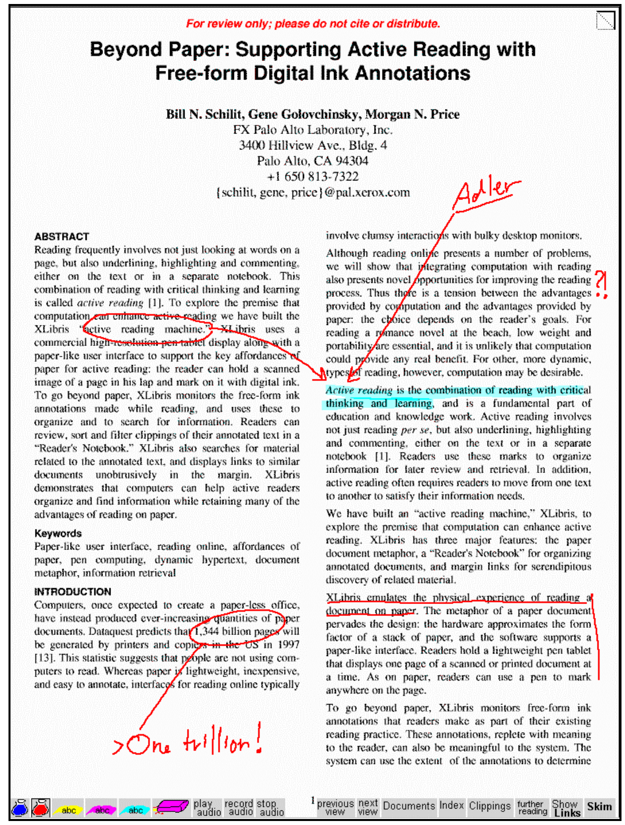
For review only; please do not cite or distribute.

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

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.

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

1

Figure 2: The entire XLibris display shows an annotated document page

ners of each page (see top right of Figure 2) provide feedback about the length of the document and the reader’s current location.

Page Orientation & Fixed Layout

Paper documents are laid out on fixed-size pages. The page layout often communicates the type of the document (e.g., business letter versus technical article) and where to find important information (e.g., a return address). The fixed layout also supports spatial memory and helps readers find old information [11, 6]. In these and other ways, paper pages give readers an excellent “sense of the document.”

This sense of the document is often lost online. Most CRT monitors cannot display a full page of text legibly, leading to awkward scrolling and zooming [23]. Word processors and Web browsers commonly re-flow the text on a “page.” And hypertext versions of paper documents can lead to disorientation and a sense of being lost [7].

XLibris respects pages by displaying the image of a single page at a time. Because the LCD panels we use are designed for landscape viewing, XLibris rotates the image 90º to approximate the portrait orientation of a typical printed page. XLibris stores the image of each page and the underlying text; this “image + text” file format is compatible with printed and with scanned documents [5, 24].

Multiple displays

People often work with multiple pieces of paper simultaneously, either with several pages from one document or with several documents [1]. Most of today's computers provide multiple virtual displays through use of a windowing system. Unfortunately, the effort of managing window layout and switching among windows can interfere with reading [23]. Therefore XLibris displays a single page image and does not support windows.

Since limiting the digital library information appliance to a single page image is problematic, we have considered a number of ways to address the need for collateral displays. Presenting two full-page images side by side, as in Book Emulator [3], makes text unreadable on current tablet displays, but may be feasible as LCD technology evolves. A form factor with multiple physical displays is another approach. (Bush's vision of a Memex desk included multiple displays [4]). In our case, to maintain mobility, a book (instead of a tablet or desk) with two opposable LCD panels may be appropriate. Weiser [32] presents an enticing vision of ubiquitous computing with devices so cheap they can be scattered and used as paper. In the near term, however, a small number of pen computers in conjunction with nearby desktop displays or even paper may provide the most practical solution to the problem of collateral document work.

In summary, XLibris employs tangibility, page orientation, and fixed layout to provide an online reading experience comparable, in many ways, to reading from paper.

FREE FORM INK ANNOTATION

> "Until we [provide support for smooth integration of annotation with reading], it is likely that people will continue to annotate paper materials, even as they read materials in a digital library." – Cathy Marshall [17]

The second key document activity is annotation. Knowledge work involves reading combined with categorizing, speculating, remembering, judging, or, more broadly, critical thinking. This active reading process is facilitated by underlining, highlighting and making notes on the text or in a separate notebook. Marking on paper helps readers "make the text their own" and is a common practice associated with deep understanding of written information [2]. In a nutshell, readers write [1, 16, 23].

Unfortunately, annotation online is quite different from its paper-based cousin. Interfaces for annotating often involve selecting a command, pointing with a mouse, and typing on a keyboard. They generally require much more effort than scribbling with a pen. Text annotations are not as visually distinct as ink marks, and online annotations often cause changes in the layout of the document [23, 30].

We believe that the ability to make unstructured, free-form, idiosyncratic ink marks is a crucial feature of any interface to digital libraries. Although such marks may not have explicit meaning to the computer, they have rich semantic meaning, supporting visual and episodic memory. An essential aspect of ink on paper is its lack of modality: users can write anything they want, anywhere on the page.

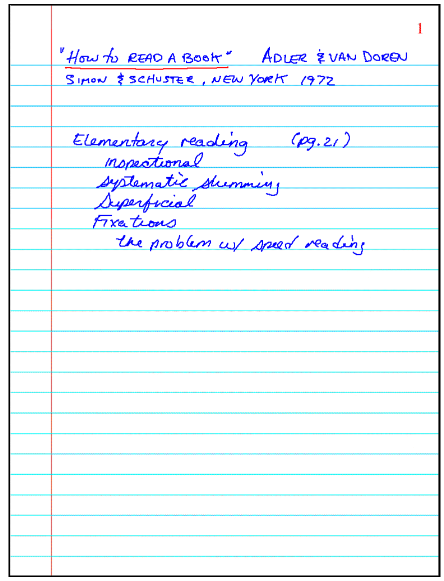

Figure 3: A notebook page from XLibris.

In XLibris, we followed this principle by letting users scribble notes, draw figures, highlight, or annotate text, all without switching modes or applications. Readers highlight and mark-up document page images as they read (Figure 2) or flip to a blank page in a lined notebook for more space (Figure 3). XLibris provides several colors of ink, several highlighter pens, and a stroke-based eraser.

FLUID MOVEMENT AMONG DOCUMENT ACTIVITIES

The third key aspect of knowledge work is fluid movement among different styles of reading and different document activities. As an example, consider the process of "information triage" [19] where an analyst skims through a pile of search results and reduces them to a manageable set. The current practice is to fire off a query to a search engine, retrieve the 100 top-ranked documents, print them out, mark them up with highlighters, and put them into piles. Analysts then read the most relevant passages to answer the original question.

Consider an alternative scenario where the analyst owns a digital library information appliance. After conducting a search, the analyst skims the documents online. The appliance enhances skimming by emphasizing important material (Figure 4). As on paper, the analyst marks up important information as it is found (Figure 2). Instead of reviewing the collected information by thumbing through piles of paper, the analyst views a concise collection of his or her text annotations (Figure 5). If any of the annotated material is worth reading in greater depth, the analyst can do so comfortably online. If the analyst needs more information, the text that has already been highlighted can be used to generate the next query (Figure 6 and Figure 8).

Figure 4. A skimming view of a document: the darker the term, the more representative it is of the document.

At a high level, this scenario demonstrates fluid movement among document activities. The analyst quickly moves among searching, skimming, annotating, reviewing, and deep reading. There is no waiting for a printer. Because skimming and deep reading are tightly integrated with retrieval, analysts can determine if the information they are collecting is useful before they invest much effort.

This scenario also demonstrates how computation can enhance skimming, reviewing, and searching. Because the interface is based on a paper document metaphor, the reading and annotation activities remain much the same as on paper. In our design, the "computer" remains in the background and, when requested, rises to the surface to assist the reader with the tasks of skimming, reviewing, and searching. Below we describe how XLibris supports and augments these three activities.

Skimming Mode

Skimming is one way readers gain a quick impression of a text. When skimming, the reader's eyes typically search for and alight on key words or sentences and take in short passages before moving on. Sometimes the reader becomes engaged and shifts into deep reading, or is distracted and moves into quicker skimming or into riffling[3].

[3] As described earlier, XLibris supports riffling with pressure sensitive strips mounted on the case that allow page turning at several speeds.

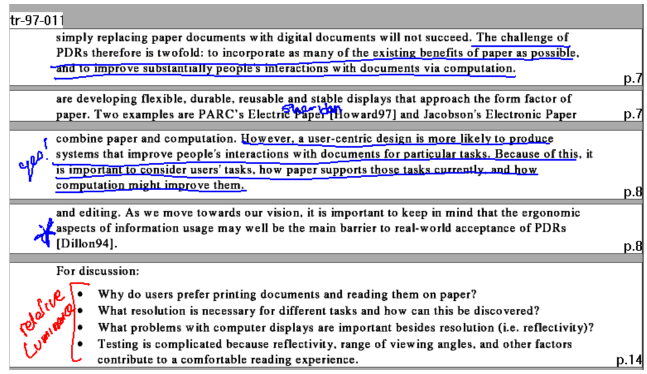

tr-97-011

simply replacing paper documents with digital documents will not succeed. The challenge of PDRs therefore is twofold: to incorporate as many of the existing benefits of paper as possible, and to improve substantially people's interactions with documents via computation. p.7

are developing flexible, durable, reusable and stable displays that approach the form factor of paper. Two examples are PARC's Electric Paper [Howard97] and Jacobson's Electronic Paper p.7

combine paper and computation. However, a user-centric design is more likely to produce systems that improve people's interactions with documents for particular tasks. Because of this, it is important to consider users' tasks, how paper supports those tasks currently, and how computation might improve them. p.8

and editing. As we move towards our vision, it is important to keep in mind that the ergonomic aspects of information usage may well be the main barrier to real-world acceptance of PDRs [Dillon94]. p.8

For discussion:

- Why do users prefer printing documents and reading them on paper?
- What resolution is necessary for different tasks and how can this be discovered?
- What problems with computer displays are important besides resolution (i.e. reflectivity)?
- Testing is complicated because reflectivity, range of viewing angles, and other factors contribute to a comfortable reading experience. p.14

Figure 5: 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.

XLibris' "skimming mode" highlights phrases and sentences that are characteristic of the document being skimmed. We call out key phrases because they can be read at a glance and tend to reflect the topic of the nearby text. This assists with the activity of scanning the text for relevant portions to read. The *Wall Street Journal* and *People Magazine* use a similar technique of boldfacing company names and Hollywood celebrities' name to help readers find information.

We also highlight summary sentences to support a speed-reading–like activity. Skimming mode uses a commercial text summarizer [30] to identify summary sentences that are emphasized. Skimming through a document of highlighted summary sentences provides a narrative overview that is not available by just reading key phrases. Both of these highlighting techniques permit a smooth transition from skimming to deep reading: key terms and sentences can help readers decide whether a passage is worth reading and are likely to be appropriate places to transition to deep reading.

Skimming mode uses shades from gray to black to indicate term importance (Figure 4). Meaningful terms are identified by heuristics that select noun phrases. The shade for each term is then based on a statistical information retrieval measure: terms that occur frequently in the document but occur rarely in other documents are colored black, while terms common to many documents are colored light gray.[4] For example, the term "digital library information appliance" would be colored black when it appears in this document.

Skimming is a common practice: a study of college students' reading habits reported that skimming occupied 32% of their reading time, and that "86% of the descriptions of skimming strategies mentioned a selection process based on key words and/or particular sentences and paragraphs in certain text locations" [20]. Although our design requires

[4] *tf·idf* is a standard information retrieval measure of terms importance [27]. We found that the repetition of highlighted terms in the document already increases their visual impact, so we use *log(tf)·idf* instead.

empirical evaluation to confirm its utility, from our experiences thus far, we are optimistic that computers can improve the skimming activity and that augmented skimming is a useful function for a digital library information appliance.

Reviewing the Reader's Notebook
Reviewing paper documents is facilitated by the marks made during reading: marks not only record information but are also used to organize it for later review [17, 24]. There are three common venues for marking: annotating on the page, taking notes in a notebook, and writing on loose-leaf paper. Annotations on the page highlight key information but tend to be lost in piles of paper. Notebooks are compact and can be reviewed quickly, but taking notes can be tedious and error-prone. Unbound notes can be reorganized flexibly, but require even more effort by the note taker.

XLibris' Reader's Notebook combines the best features of annotating directly on the page, of taking notes in a separate notebook, and of organizing index cards. As with paper documents, readers mark on the page in the context of the document, yet without the laborious and imprecise step of copying. 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 needs change.

The Reader's Notebook (Figure 5) extracts "clippings" of annotated text and lays them end to end in a separate, multi-page view [29]. Each clipping is linked to the corresponding annotated page, so readers can move fluidly between notes and documents. Each clipping includes some surrounding text, and is labeled with document title and page number to help readers understand the meaning of the marks.

In designing clippings, we had to decide how much of the document should be shown for each annotation. Starting with the bounding rectangle of each ink stroke, we expand it horizontally to the width of the page and vertically to include complete words. Snippets that overlap are merged together resulting in reasonably sized clippings of annotated text (Figure 5).

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

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.

Finding Information Related to Readers' Annotations
Readers often search for related material and move from one text to another. 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 has never prescribed decides to follow a reference to a study of its side effects.

Unfortunately, in both online and offline situations, finding related materials and reading are not well integrated. Typical information retrieval interfaces force users to stop reading, to identify related documents, and then to print them. 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, despite the fact that readers' interests may change rapidly, interfaces for moving from reading to finding and back to reading again are time consuming and disruptive. Clearly this is another example where fluid movement across document activities is desirable.

XLibris facilitates fluid transitions among reading, searching and browsing by creating a dynamic hypertext from readers' marks [26]. The system observes readers' annotations and the text of the underlying document and generates hypertext links (margin links) to related documents: no additional effort from readers or authors is required. Because XLibris suggests links of interest to particular readers at particular times—based on the passages they are marking—people can discover useful information *serendipitously*.

This type of serendipitous retrieval is a benefit of paper libraries: as people walk to the shelf for a particular book, walls of related books surround them, making it possible to find interesting material accidentally. Although serendipity in the library is often a rewarding experience, it is generally lacking in online systems. XLibris provides two user interfaces that integrate reading, searching, browsing, and serendipity. These are described below.

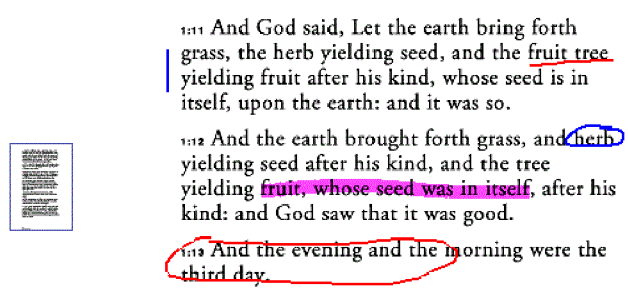

Figure 6: Examples of highlighting, underlining, circling, and margin annotations. Each annotation generates a query. If a good match is found, XLibris adds a margin link (rectangle on the left) that shows the thumbnail of the destination page.

Margin links
Margin links provide serendipitous access to related documents *during* the active reading process. As readers mark up passages, the system finds related documents and presents

links unobtrusively in the margin; our intent is to provide a modeless link suggestion mechanism. Because margin links persist, readers can follow links at their leisure. We use thumbnail images of the target page as anchors (see Figure 6). The reader taps on an anchor to move to that page.

For the most part, link creation does not disrupt the reading process since the reader is free to disregard the suggestion and to continue reading. In the current design, however, margin links intrude on the reading process more than we would like. One problem is that XLibris computes a separate query for each stroke (although strokes often do not result in margin links because of a similarity threshold). We have begun to address this by adding a button to turn links off and by designing ways to manage the frequency with which margin links appear.

Further reading lists

When readers reach the end of a document, they often want to know more. The document may not emphasize the topic they are most interested in, or it may spark an interest in a new topic without providing enough depth or detail. Sometimes this need is addressed. For example, articles in *Scientific American* typically describe an area of research at a high level, without providing much technical detail. To help readers go into more depth, the editors create a further reading list at the end of each article (see Figure 7).

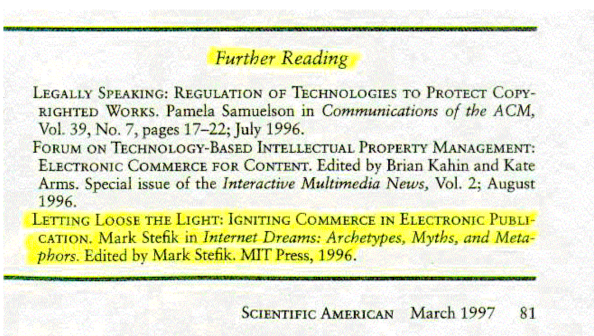

Further Reading

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 7: A further reading list from *Scientific American*

The *Scientific American* example illustrates the limitations of authored 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 the question directly. Furthermore, by the time the reader scans the further reading list, more appropriate articles may have been written.

XLibris augments this traditional editorial practice by automatically generating further reading lists for each document (see Figure 8). Unlike static references, these lists reflect the interests of a specific reader at a specific point in time. As with margin links, the reader's interests are inferred from annotations, and no additional intervention from the reader is required.

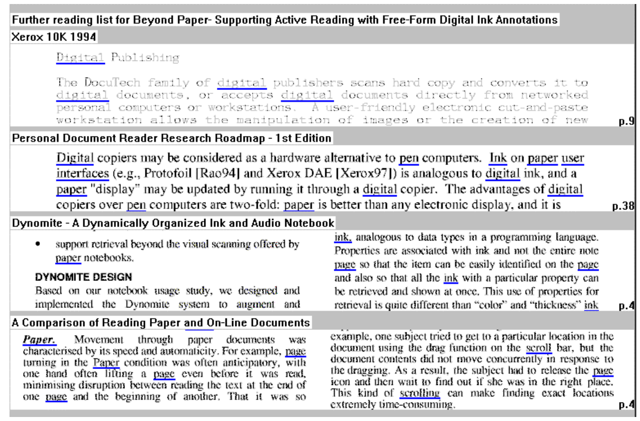

Further reading list for Beyond Paper- Supporting Active Reading with Free-Form Digital Ink Annotations

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 8: A further reading list in XLibris. Each related document is presented as a clipping of the most relevant sentence with key phrases underlined.

Further reading lists are presented as a separate set of pages, or *view*, associated with each document. Readers can access this view at any time, whether or not they have annotated the current document. Of course, the more annotations the user has made, the more focused the resulting list is on the reader's interest as opposed to the document as a whole. Visually, the further reading list looks like a Reader's Notebook of annotations except that each clipping is a segment of a related document obtained from a search engine. These clippings, with matching terms underlined, help readers understand the destination of the link and also make the target passage more recognizable if they choose to follow the link.

Query-mediated links

Both margin links and further reading lists use a technique called *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.

XLibris computes margin links from the words, phrases and passages that are implicitly selected by the reader's marks. (We do not expect readers to make marks for querying explicitly, but rather as part of their existing annotation practice.) Marks are converted into text selections, which are then expressed as full-text queries that yield 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.

For further reading lists, each annotation is interpreted as a text selection and is transformed into a list of word scores. The scores for each word are then summed across all annotations. If several annotations select the same instance of a word, then the maximum score for that instance is used. The words along with their scores are used to generate a query[5]. If not many words have been included, then terms that are most characteristic of the document can be added. This padding reduces the chance that a query based on a small num-

[5] Weights are used in a query to specify the degree of importance for each search term.

ber of annotations will return documents that are entirely unrelated.

Ink annotations as queries
XLibris recognizes several distinct ink patterns from which queries are computed. These include underlined words, highlighted words, circled words, circled passages, and margin annotations (see Figure 6). We emphasize that there is no "vocabulary of marks" that readers must learn; rather, our heuristics are based on general annotation practices. Each type of annotation results in a slightly different query for the search engine. Marks that select specific words translate into queries that retrieve other instances of the same term. Marks that select longer passages generate queries that search for similar passages.

PHYSICAL MOBILITY
The fourth and final document activity supported by digital library information appliances is mobility. Until now we have discussed reading in XLibris; in this section we describe mobile information access in TeleWeb [28].

Desktop access to digital libraries is not adequate for mobile workers. A doctor making hospital rounds has no desktop with her to access a digital library. This could have grim repercussions: a recent study concluded that patients for whom prompt MEDLINE searches were conducted have significantly lower costs and shorter hospital stays [14]. Mobile work is not limited to physicians and traveling salespeople; even office workers often work with documents away from their desks [1].

For these people, portable information appliances can provide access to digital libraries anywhere, anytime. Unfortunately, it is not possible to give mobile people the same network access that they would have in their office. Compared to wired desktop networks, wireless networks are slow, expensive (with different billing schemes), and only intermittently available. As workers move from place to place, they are likely to experience three levels of connectivity: high quality connectivity from an office network, lower quality connectivity from a wireless network, and no connectivity at all when disconnected. How can people use digital library appliances to access distributed information sources across this range of operating conditions?

Aside from the variable networking conditions, there are two other issues that affect mobile access to digital libraries. First, downloading a document over a wireless network can be very expensive, and current interfaces for search engines and Web browsers do not help people manage these costs. Second, user interaction is punctuated by surprises: an information request may return right away (when data is locally stored), take a while (when data is remote), or take forever (when the network is down).

The TeleWeb mobile Web browser was designed and built to explore these three issues [28]. TeleWeb runs on a laptop computer with multiple communication channels: it may be docked to an office Ethernet, used from home over a telephone line, used while traveling over a wireless network such as a cell phone, or operated while not connected to a network. The lessons learned from TeleWeb apply not only to the Web but also to digital libraries.

Figure 9: A document annotated with road sign icons to show hypertext anchors absent from the cache.

Networking not ubiquitous
In the design of TeleWeb, we assumed that at least some of the time users would be disconnected but would still want to browse the Web. Some mobile networks are so limited in bandwidth compared to office networks that users may feel like they are disconnected. How can people browse the Web when disconnected?

TeleWeb uses "caching for availability," the standard solution for this problem. This means that the system loads documents into a local cache when fully connected. Later, if the user needs the document while the network is not available, the locally cached copy is returned.

Caching for availability only works if the system knows the user's information needs ahead of time. The system can use past history to predict future needs. In addition, the user can tell the system what documents belong in the cache; unfortunately, in practice people find it difficult to articulate which documents they need. Furthermore, both techniques share the drawback that information needs often change in unpredictable ways. Overall, caching for availability often works, but is frustrating when it does not.

An alternative solution is to replace synchronous interaction (a user request followed immediately by a response from the network) with asynchronous interaction. For example, when the user requests a document not in the cache, TeleWeb asks the user if the system should download it later. If so, TeleWeb downloads the information the next time the user connects and adds the document to a "to read" list that the user can browse. This interaction style is analogous to sending a request for a book from an out of state library, and eventually receiving it in the mail.

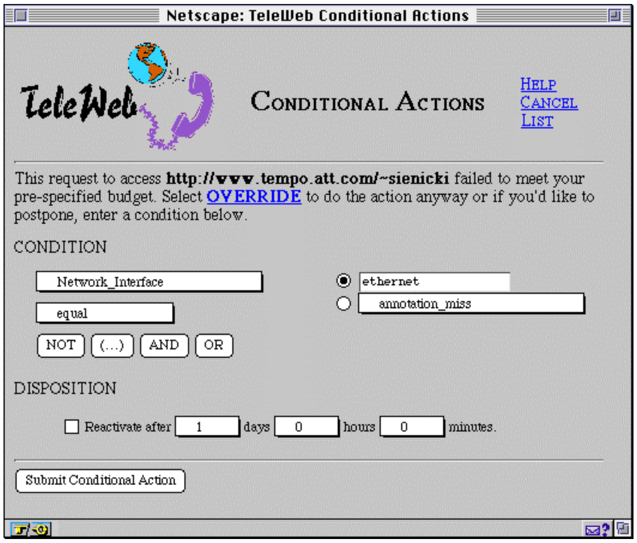

Figure 10: Creating a conditional download to fetch documents when reconnected.

No feedback on cost

There are other reasons to postpone downloading a document. For example, if a user is connected over a cellular telephone, downloading a large document is going to be expensive, and can tie up the network connection for hours. Given the choice, the user might be happy to download this document overnight or upon returning to the office. There is a tradeoff here between how urgently people need information and how much they are willing to pay for it in time and money. With a mobile information appliance, users should be able to decide how much they are willing to pay for each document and how quickly they need that document.

With this principle in mind, TeleWeb includes a budget monitor that sits between the user and the network. If a request is made for a document that is not in the cache and downloading the document will break the user's budget, then TeleWeb presents a "conditional download" form instead of the document (see Figure 10). Users can override their budget to fetch the document immediately, or they can request that the document be fetched at a later time.

Unpredictable systems are frustrating

The third problem of mobile appliances is unpredictable behavior. Those using computer networks already know this problem: sometimes they hit a key and nothing happens. In wireless computing, this experience is magnified many-fold.

TeleWeb's solution is to show the user what is going on inside the cache and network manager. Specifically, TeleWeb labels links that will be slow to download with a "road work" icon (see Figure 9). If there is no road sign, then either the link destination is cached, or a high-speed connection to the information superhighway is available. This simple change—essentially exposing the state of the system—makes mobile Web browsing much more predictable. This contrasts with the traditional approach of trying to simplify the user's view of the system by concealing the state of the network.

THE FUTURE

The development of tomorrow's digital library information appliances is an iterative process. Our current cycle, however, is not complete until we evaluate the ideas presented in this paper. Towards this end, we have begun a series of laboratory user interface experiments and are in the planning stage for a real-world deployment. One experiment suggests that free-form ink marks made while reading are better for query formulation than traditional relevance feedback. People who have tried XLibris agree that our approach of taking a work practice and augmenting it—rather than redefining it—makes the system immediately accessible. Information *appliances* must be accessible.

The digital library community must decide whether digital libraries should, as we propose, support users throughout the process of turning information into knowledge.[6] If so, there are many issues that must be addressed. One reviewer of this paper asked: "Where is the digital library?" We purposely avoided describing protocols, services, and architectures because we believe infrastructure should *support* the vision of working with a digital library information appliance, not *define* it.

The next step is to understand how users of digital library information appliances should interact with external repositories. For example, should users pay authors to cache documents on their devices, even when they do not read them? How should users search repositories that are not always accessible? How should searches be distributed between local and remote databases? These and other issues require further investigation.

CONCLUSION

We have shown how information appliances based on the paper document metaphor can support active reading. Because these systems allow people to work on digital documents much as they would on paper, this provides an alternative to the standard "search and print" model of digital libraries. Furthermore, by integrating a wide variety of document activities, and by allowing fluid movement among them, disruptive transitions between paper and digital media can be eliminated. The combination of the paper-like and the digital allows us to augment these activities without redefining them. Finally, the mobility of a digital library information appliance supports work away from the desk. With all these benefits, digital library information appliances can create a rich, universally accessible, digital library experience that improves the way we work. Clearly there is an exciting future for digital library information appliances in the digital library.

---

6 We thank one of the anonymous reviewers for this succinct summary of our goal.

ACKNOWLEDGMENTS

We would like to thank our colleagues at Xerox PARC and Xerox Research Centre Cambridge for valuable discussions. Mark Weiser suggested the Scientific American analogy for further reading lists. Beverly Harrison, Roy Want, Ken Fiskin, and Anuj Gujar provided pressure sensor hardware for XLibris. The TeleWeb research project at AT&T Bell Labs included Fred Douglas, Dave Kristol, Paul Krzyzanowski, Jim Sienicki, and John Trotter. We thank the anonymous reviewers for many excellent suggestions. We would also like to thank Joe Sullivan and Jim Baker at FXPAL 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. To appear in *Proceedings of CHI98* (Los Angeles, CA, April 18-23), 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 R. McAleese and C. Green (eds) *Hypertext: State of the Art.* Oxford: Intellect. pp. 52-63.
4. Bush, V. (1945) As We May Think: A Top U.S. Scientist Foresees a Possible Future World in Which Man-Made Machines Will Start To Think. *LIFE*, 19(11), pp. 112-124.
5. Birrell, A. and McJones, P. Virtual Paper. Available at www.research.digital.com/SRC/virtualpaper/home.html
6. Dillon, A. (1994) *Designing Usable Electronic Text*. Taylor and Francis, London, UK.
7. 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.
8. Entlich, R., Olsen, J., Garson, L., Lesk, M., Normore, L., and Weibel, S. (1997) Making a digital library: the contents of the CORE project, ACM TOIS, 15(2), pp. 103-123.
9. Golovchinsky, G. (1997) 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. 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.
11. Haas, C. (1996) *Writing Technology: Studies on the materiality of literacy*. Lawrence Erlbaum Associates, Mahwah, NJ.
12. Harrison, B.L., Fishkin, K., Gujar, A., Mochon, C., and Want, R. (1998) The Design and Use of Squeezable Computers: An Exploration of Manipulative User Interfaces. To appear in *Proceedings of CHI98* (Los Angeles, CA, April 18-23), ACM Press.
13. Johnson, A. (1997) Fax Trends. *Office Equipment and Products*, July 1997, p. 25.
14. Klein, M.S., Ross, F.V., Adams, D.L, and Gilbert, C.M. (1994) Effect of online literature searching on length of stay and patient care costs. *Academic Medicine* 69, June 1994, pp. 489-95.
15. Levy, D.M. (1997) I Read the News Today, Oh Boy: Reading and Attention in Digital Libraries. In *Proceedings of 2nd ACM International Conference on Digital Libraries*, (Philadelphia, PA, July 23-26), ACM Press, pp. 202-211.
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. Marshall, C.C. (1990) Work Practices Study: Analysts and Notetaking. Unpublished report, 31 May 1990.
18. Marshall, C.C. (1997) Annotation: from paper books to the digital library. In *Proceedings of 2nd ACM International Conference on Digital Libraries,* (Philadelphia, PA, July 23-26), ACM Press, pp. 131-140.
19. Marshall, C.C. and Shipman, F.M. (1997) Spatial hypertext and the practice of information triage. In *Proceedings of Hypertext '97* (Southampton, UK, April 6-11), ACM Press, pp. 124-133.
20. Masson, M. E.J. (1982) Cognitive processes in skimming stories. Journal of Experimental Psychology: Learning, Memory, and Cognition, 8(5) pp. 400-417.
21. Mutoh (1997) Mutoh America Inc. Available at www.mutoh.com
22. Nielsen, J (1988). Hypertext '87 Trip Report. *ACM SIGCHI Bulletin* 19**,** 4 (April 1988), pp. 27-35.
23. O'Hara, K. and Sellen, A. (1997) A Comparison of Reading Paper and On-Line Documents. In *Proceedings of CHI97* (Atlanta, GA, March 22-27), ACM Press, pp. 335-342.
24. 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.
25. Phelps, T.A. and Wilensky, R. (1996) Towards Active, Extensible, Networked Documents: Multivalent Architecture and Applications. In *Proceedings of Digital Libraries '96* (Bethesda, MD, March 20-23) ACM Press, pp. 100-108.
26. Price, M.N., Golovchinsky, G., and Schilit, B.N. (1998) Linking By Inking: Trailblazing A Paper-like Hypertext. To appear in *Proceedings of Hypertext '98*, (Pittsburgh, PA, June 20-24), ACM Press.

27. Salton, G. and McGill, M. (1983) *Introduction to Modern Information Retrieval*, McGraw-Hill.

28. Schilit, B.N., Douglis, F., Kristol, D.M., Krzyzanowski, P., Sienicki, J., and Trotter, J.A. (1996) TeleWeb: Loosely Connected Access to the World Wide Web, *Computer Networks and ISDN Systems*, 28(1996) pp. 1431-1444.

29. Schilit, B.N., Golovchinsky, G. and Price, M. (1998) Beyond Paper: Supporting Active Reading with Free-form Digital Ink Annotations. To appear in *Proceedings of CHI98* (Los Angeles, CA, April 18-23), ACM Press.

30. *Search '97 Developer' Kit: Advanced Features.* Verity, Inc.

31. Sellen, A. and Harper, R. (1997) Paper as an Analytic Resource in the Design of New Technologies. In *Proceedings of CHI97* (Atlanta, GA, March 22-27), ACM Press, pp. 319- 326.

32. Weiser, M. (1991) The Computer for the 21st Century. *Scientific American*, 265(3), September 1991, pp. 94-104.