

Collaborating over portable reading appliances

Catherine C. Marshall, Morgan N. Price, Gene Golovchinsky, Bill N. Schilit

FX Palo Alto Laboratory
3400 Hillview Avenue, Building 4
Palo Alto, CA 94304 USA

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

Abstract

Reading appliances or e-books are a new kind of personal technology that hold substantial promise as interpersonal technology, devices that help us collaborate. In this paper, we use a study of a group activity – a reading group that meets to discuss articles of mutual interest – to explore four scenarios for collaborating with e-books: (1) meeting room and face-to-face discussions; (2) serendipitous sharing of annotations, as when we borrow a document from a colleague or buy a used book; (3) community-wide use of anonymous annotations to guide future readers; and (4) e-books as a basis for initiating interaction between people. In so doing, we describe some methods for implementing these facilities, and introduce design guidelines.

Keywords: collaboration, e-book, reading appliance, annotation, qualitative study, handheld CSCW

1. Introduction

People use a great variety of handheld personal technologies – from Palm Pilots to Cross Pads – in an increasing number of work and leisure settings. It is natural for people to interact and work with each other in these settings, in spite of the fact that they may be using devices designed as personal, not *interpersonal*, technologies. At the conclusion of a meeting, for example, many of the attendees may pull out their Palm Pilots to come up with a scheduled date and time for the next meeting.

Reading appliances or *e-books* are a new kind of personal technology that shows great promise for supporting interpersonal collaborations. Reading appliances are a realization of the long-held vision of interactive electronic books. Most generally, they give readers the familiar experience of reading and writing on paper documents. Then they go beyond paper to offer readers capabilities we associate with computers, such as keyword search, hypertext links, and the ability to quickly obtain new materials from online sources like digital libraries.




Figure 1. XLibris prototype (implemented on the Fujitsu Point 510) in a reader's lap. Higher resolution devices are also available.

How can a portable reading appliance support collaboration? At first glance, this seems an ill-conceived question. Reading, many assume, is an individual act, not a collaborative activity, and it is best conducted in private, quiet places. Portability enters into the equation only peripherally. Books or documents are portable, but once we sit down to read, we are essentially anchored to the spot, and have no cause to move further. Furthermore, we expect that any material reflection of reading – our annotations, dog-ears, and bookmarks – are private, artifacts that no-one else will ever use.

When we look more closely at reading, however, these assumptions break down. The act of reading is situated within a larger sphere of human activity – of reaching a shared understanding, of writing, of functioning in the world – and often transcends the individual. Furthermore, people read everywhere [1], and they read opportunistically, carrying documents with them until they find the right time and place to read [10]. The marks we make when we are reading may cross into more public realms. Annotations and bookmarks are shared as documents pass from hand to hand, either as we borrow document from our colleagues, or buy used books at the bookstore.

Thus portable reading appliances introduce new possibilities for collaboration: they can support existing practices such as sharing annotated documents, and they can introduce new kinds of capabilities such as recording the reading patterns of many readers, and using these patterns to help other readers. Further investigation is required to turn e-books from personal technologies into interpersonal technologies.

In this paper, we use a recent study of a reading group – their individual reading practices, and how they come together in both planned and ad hoc ways – to further our understanding of how reading appliances may be used in collaborative settings. The study relies on an *intervention* in their usual reading practice; in this case, we intervened by introducing a portable document appliance, XLibris, into an existing situation, researchers preparing for a reading group meeting [10]. We will briefly describe the study, but first we will describe our portable reading appliance, the device that is used as a basis for the intervention.

1.1 XLibris

For the past two and a half years, we have been investigating how portable reading appliances can combine the mobility and affordances of paper¹ with computational augmentation [14]. By contrast to the recent spate of electronic books [17], we intend reading appliances to support a range of intellectual activities associated with active reading [2].

Our explorations to date have resulted in a research prototype, XLibris, that uses a *paper document metaphor* to support analytic reading activities [13]. The device reproduces the physical experience of working with paper: readers can hold electronic documents on their laps, moving the e-book as appropriate to avoid glare; they can mark on the electronic documents with a variety of pens and highlighters; and they can turn from one page to the next by thumbing on the device. In essence, XLibris attempts to capture the materiality associated with reading physical documents (see Figure 1) and the unselfconscious engagement people have with paper documents. Figure 2 shows a page of an electronic document that a reader has annotated with freeform ink using XLibris.

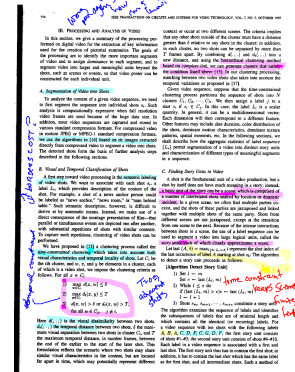


Figure 2. A page annotated by a reading group member using XLibris.

Why would people use an e-book if what it does is imitate paper? A primary emphasis of XLibris research has been to go beyond the capabilities of paper documents. For example, the XLibris Reader's Notebook helps readers review their reading by gathering marked-up passages into a list, which they can use to navigate to source documents [13].

The study we describe below was initially aimed at understanding how usable the reading appliance was, given an analytic reading task, and what would make it useful to readers. A key question in this study – since it was of a collaborative activity – was how reading appliances might fit into a setting in which reading leads to further collaboration.

1.2 A study of reading and collaboration

Our investigation of reading centered around a group of six researchers selecting papers for, preparing for, and participating in an hour-long discussion. We observed and interviewed the group over the course of two consecutive readings – a conference paper about video indexing and a longer journal article about video summarization – and the associated meetings. The first meeting allowed us to get a sense of how the group normally selects, reads, annotates, and discusses a paper. We asked group members to prepare for the second meeting by reading the article using XLibris. Thus, to prepare for the first meeting, the group members read paper copies that they had produced themselves. We did not constrain where or when people read the first paper.

The reading group members prepared for the second meeting by reading the article on XLibris using a tablet (the Mutoh MVT-12) tethered to a PC; a tethered version was used because of its slightly higher display resolution (1024x768) than the portable XLibris shown in Figure 1 (800x600). They read in

¹ The affordances of paper are described in [15]

one of two offices in which we set up the appropriate hardware. For the second meeting, we printed color copies of each group member's annotated paper and Reader's Notebook [13] after they finished reading on XLibris, and group members used their paper copies during the discussion. Both meetings took place in the usual location for the reading group's weekly get-togethers, a small conference room.

We collected data before, during, and after the two meetings. We interviewed participants individually about their experiences of reading each paper – from the mechanics of how and where they read it, to what they thought of the contents of the paper and why it was chosen by the reading group, to what they would do with it after the reading group discussion was over. The interviews were semi-structured and open-ended; reading group members gave long accounts of their experiences and readily introduced new topics. We videotaped each of the reading group's meetings from two vantage points; this enabled us to see how the group members were using the papers in the meetings. Figure 3 shows the reading group's second meeting.




Figure 3. The reading group's meeting. Note that they all have their own copies of the paper at hand.

We also collected copies of the annotated papers before and after the meetings, so that we could use the papers during interviews. Figure 4 shows the comparable page marked up by each reading group member. We use these marked up copies, coupled with our analysis, to provide examples for this paper. Finally, logs of user interaction events during the XLibris sessions were maintained so that we could analyze patterns of annotation and navigation.

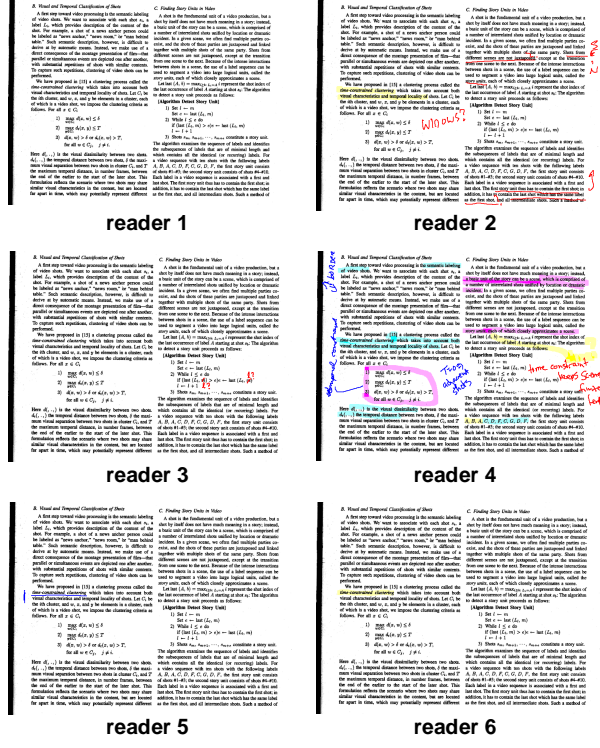


Figure 4. A portion of a page of the second article annotated by 6 reading group members. The pages were annotated using XLibris.

What we saw in our study was a rhythm of collaboration and individual reading practices, coupled with mobility as reading group members (using paper) read when and where the opportunity arose. People also used the papers in a variety of situations – in their scheduled meeting, as well as in impromptu conversations about specific topics of interest in the papers. Our analysis showed that individual reading and annotation practices did not vary significantly with the introduction of XLibris [10]. Thus we could turn our attention to the reading group's collaboration.

Three central observations motivate our scenarios for collaborating over portable reading appliances:

- The marks correspond (roughly) to their expected use in the reading group's meeting and in other settings. They also depend on the individual reader's expected role in this discussion. This should be no surprise: people read for a purpose, and when they read for a similar purpose in the same work setting, the annotations they make reflect this.
- Readers discovered other uses for what they read in addition to the more formal collaborative functions they originally anticipated. Ad hoc

meetings between subsets of group members took place before and after the scheduled meeting. For example, two reading group members got together to follow up on selected references. Another pair discussed a brief, but important, portion of the paper that pertained directly to their own work.

- Paper's mobility and versatility are important for supporting seamless connections between reading and associated activities. People carried, saved, and used their annotated copies to remind them of their questions, reactions, topics to introduce, and references to retrieve; the annotations bridged the gap between reading and collaboration. Although paper is mobile and tangible, it does have notable limitations. One reading group member said he did not save papers; if he did, he soon would be overwhelmed by them.

2. Four scenarios for collaboration with portable reading appliances

We use the results of our studies to explore four different kinds of collaboration that may be supported using portable reading appliances. The first is based on what we observed: meeting and discussing a document, the synchronous sharing of documents that is typical of face-to-face collaboration; this kind of collaborative activity has been investigated at some length over the past decade or more [18,4,19,20]. The second applies the study results to explore serendipitous sharing, the unintentional sharing of marks that is typical within workgroups. Both kinds of support for collaboration rely on the reader's ability to engage with a document unselfconsciously at the time that they are reading – to read and mark as if they were using paper.

The final two scenarios are more speculative; they rely on computational methods of harvesting annotations over a larger group of readers. The first uses patterns of readers' marks to guide other readers through a document; it is predicated on the anonymity of the harvested marks. The second allows people to locate others with similar interests through their reading habits; it relies on readers' desires to meet like-minded people.

To support these scenarios, reading appliances must first be usable: a reader should be able to engage naturally with documents, to carry them around, and to mark on them unselfconsciously. They must also support sharing in some way: the reading appliance should have the ability to collect and compute over annotations. Finally, these scenarios are predicated on the ability to manage the social aspects of using

annotations by limiting the circulation of marks or by guaranteeing anonymity. Taken together, our strategies for making e-books into interpersonal devices rely on the observation that an individual's marks are often reliable indicators of what piqued the reader's interest while he or she was reading.


In each of the four explorations, we will first present our observations of the phenomenon we are using to define the scope of the exploration. Next we will show how these characteristics lead to a design. Finally we will discuss the challenges that arise from the design. Each exploration may take advantage of a variety of methods for presenting marks. We discuss the methods first, since each method appears in several scenarios.

2.1 Methods for using multiple readers' annotations


Sample methods for using multiple readers' annotations will help us illustrate subsequent discussions of how e-books may support collaborative activities. These methods fall into several general categories:

- Overlaying several sets of annotations in a single display so a reader can see multiple annotations at once (see Figure 5.1);
- Allowing readers to access the annotations and underlying text apart from the whole document to create a compact representation of passages of interest (see Figure 5.2); and
- Processing the readers' annotations to emphasize annotated passages without showing the annotations themselves (see Figure 5.3).

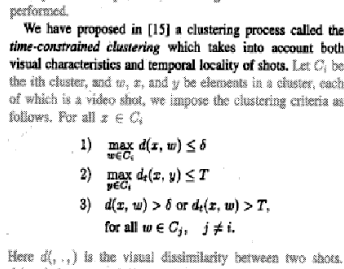
Figure 5 uses examples from the article the reading group members annotated using XLibris to illustrate these methods. We will refer to them briefly in each of the scenarios, primarily to show examples of how to implement the capabilities we describe. There are many ways to implement each of the general methods we list above besides the ones shown in the figure. Note that the final of the three methods does a more effective job of preserving anonymity; this characteristic will be important to the final two scenarios.



(a) An example of overlapping marks. The reading group's marks have been merged into a single document.



(b) An example where readers access annotations and underlying text apart from the whole document. In this case, only the text of particular interest is displayed.



(c) An example of processing the readers annotations to change the presentation of the document. In this case, the reading group's combined marks have been converted into gray level shading on marked words.

Figure 5. Methods for using multiple readers' annotations.

2.2 Scenario 1: Meeting and discussion (synchronous sharing)

People who read a document for roughly the same purpose in roughly the same circumstances identify some of the same regions of interest – an equation that's the "meat" of a technical paper, or a key point in a legal case. For example, in Figure 3, the paragraph immediately preceding the inset equations in the first column of text has interested the majority of the reading group members. In fact, they discussed the contents of this paragraph in their meeting; their annotations helped guide the introduction of topics of interest and focus the discussion. Observing their meeting showed that reading group members used the paper copies of the article to manage and coordinate their conversation.

What, then, would happen if everyone brought their reading appliance to a meeting rather than their paper copy? The effect might be minimal: people using a paper-like interface, coupled with personal reading devices, may recreate the situation the devices imitate. On the other hand, the known viewing angle problems introduced by LCD displays may simply make it more difficult for people to coordinate as easily. With paper, they can readily – and peripherally – see what page(s) the others are on, or where their marks overlap, all without distracting from the meeting. Thus, a workable interpersonal technology must provide this awareness of where others are in a document.

There are at least two general ways to do this: first, we might hypothesize a meeting room setup that includes a large shared display as well as individual e-books. Rooms that use large shared displays in concert with individually controlled computers are now commonplace; this variation simply replaces the individual computers with handheld devices as in [11]. As in other kinds of computer-supported meetings, the large central display can be used as a locus for shared reference [3], or to help the participants anticipate what is coming next.² Either the leader's marks or a sensible combination of the group members' marks (using one of the methods described earlier) may be projected for use in this way.

The second kind of hypothetical meeting room setup depends on the group members bringing their own e-books, and takes advantage of their existing practice of using peripheral awareness of each other's marks to guide discussion. As above, the leader's marks or combined marks of group members can be propagated to each participant's copy of the article on their reading appliance; this level of sharing would enable the readers to anticipate the direction of the discussion.

If we choose this strategy of using the marks to provide a shared focus, we must also be aware of what happens in a meeting. When we observed the reading group in action, it became apparent that although the marks – especially the discussion leader's marks – coincided to a great extent with what was discussed, the reading group members also seamlessly drifted in and out of sync with the discussion without interrupting the shared focus. Figure 6 shows a

² Our observations showed that the pace of the reading group's meeting more or less reflected the density of their markings on the article. For example, sparse markings near the end of the article corresponded to very brief coverage during the one of the meetings.

situation in which five of the six reading group members were discussing a point on the third page of the paper, while the sixth person was temporarily distracted, checking a reference. The group member's paper that is circled in white is open to the last page of the article; other reading group members are looking at the first. Thus, each reader needs to retain individual control of what is on his or her own display at any given time, even as they are given access to a shared representation.




Figure 6. One member of the reading group examines a different page than the others.

2.3 Scenario 2: Serendipitous sharing

When people read paper documents – books, reprints, and other documents – they may mark on them. These marks are transferred to others in a more or less seamless way whether the original reader means them to be or not. Thus we may ask a respected colleague if we can borrow a paper from him, and get not only the paper, but the main points that he's highlighted. These, in turn, may help us interpret the paper, because we know he's an expert in the area. We refer to this seamless transfer of annotations from hand to hand as “serendipitous sharing.”

In the reading group study, one situation stands out as a good example of the possibilities for serendipitous sharing. In this case, a figure was labeled in a way that made it difficult to interpret. Two reading group members annotated the document in much the same way to render it comprehensible (see Figure 7); the figure was later explained according to these labels during the reading group's discussion. It would have been helpful for readers to have seen these marks while they were preparing for the reading group.




Figure 7. A figure in the reading group's article is relabeled in the same way by two different readers.

This situation is not uncommon. Readers frequently inherit other people's markings. For example, students buy used textbooks that other students have marked up. On occasion, they may find these marks to be valuable, and may even seek them out [8]. However, one of the advantages cited for *digital* documents is the ability to remove this kind of annotation, or – more to the point – to decide upon reflection to not share personal markings.

Thus, for serendipitous sharing to work on a reading appliance, appropriate conventions must be negotiated that would enable sharing to take place. Technologically, of course, the solution is relatively simple: e-books readily support the overlay of another's marks on one's own document, or the ability to choose among one's colleagues' annotated copies.

The challenge for serendipitous sharing is that not all marks are valuable. There are many kinds of annotations that serve non-interpretive functions. Typical annotations include reactions to the text; for example, one reading group member marked a confusing passage with a question mark. They also may be reminders of something to do; for example, reading group members marked references or URLs that they would like to follow up on. Or they may simply reflect the material circumstances of reading; for example, one reading group member wrote a note about her frustration with the figure quality in the reproduction of the article: “Huh? Can't read it!”

Furthermore, the interpretations may not be universally intelligible. We have seen that readers' marks are not necessarily easy for the reader him or herself to understand at a later date. For example, our interviews of reading group members revealed that the participants often forgot their specific intent in making the marks in the interval between the meeting and the final interviews. One reading group member admitted that a diagram he had drawn on the last page of the

paper was now a mystery, but was “sure it had some significance at the time.”

This problem is exacerbated when the marks pass from one context of use to another. Yet we know that serendipitous sharing takes place. To support this kind of serendipitous sharing in the face of these challenges, it may help to preserve marks with context: who made the marks (preserving anonymity, we can record what kind of person – another student? the faculty member teaching the course? an expert in the field? – made the marks, and some aspects of the circumstances of reading).

2.4 Scenario 3: Reading roads

Annotations on paper documents and books reflect a reader’s unselfconscious engagement with the text. Freeform annotations on e-books enable this same sort of engagement with digital media. The question then becomes: Can we parlay these marks into something more? How can freeform annotations become a resource for a larger community of annotators reading the same document?

If we examine the kinds of marks that readers make, there is great variability among individuals. However, one thing is clear: the annotations, highlights, and other marks often demonstrate a reader’s interest. By looking at *where* people mark, rather than *how* they mark, we can identify areas of consensus: places in the text that all or many of the readers found sufficiently important to annotate [9]. In essence, annotations may create paths or roads through a document that others may benefit from.

In the reading group’s case, we observed that overlapping annotations indicated group interest, an interest that was demonstrated by discussion in their meeting. Not only were the locations of the marks meaningful, but also passages marked by more annotators received correspondingly more attention in the meeting.

To establish that the overlap of different readers’ marks was significant, we compared the number of overlapping marks that readers made with those calculated on the basis of probability. Table 1 shows the number of paragraphs annotated by exactly 1, 2, 3, and 4 participants in each of the two articles they read during our study. Since two of the readers made few marks, even modest amounts of four-way overlap are significant. We computed the expected numbers by assuming that each person’s marks are independent, and that participants are annotating randomly at different rates (estimated from the actual number of

paragraphs each participant marked).³ There was more 3- and 4-way overlap – and fewer paragraphs annotated by only a single participant – than expected, indicating that the readers’ marks overlapped in a meaningful way. Because we were observing a small number of readers, Table 1 is not intended to show statistical significance, but rather to give a more quantitative picture of the degree of overlap.

# paragraphs annotated by <i>n</i> participants	0	1	2	3	4
First paper (observed) out of 95 paragraphs total	58	22	12	2	1
First paper (expected)	50	36	9	1	<<1
Second paper (observed) out of 157 paragraphs total	86	37	22	9	3
Second paper (expected)	67	64	22	4	<1

Table 1. Number of paragraphs annotated by exactly *n* participants, observed and expected.

As we mentioned earlier, during the meetings, points of 4-way overlap signaled topics that received considerable discussion time and group engagement. At the other of the spectrum, passages that none of the group members annotated were rarely mentioned.

Earlier we looked at how overlapping marks might be used to focus a synchronous meeting. How might a reading appliance use these overlaps to support asynchronous large-scale sharing and interpretation of documents? First of all, the marks can serve to *summarize* what others reacted to as they read using a technique such as the one shown in Figure 5b; these passages of greater interest can be presented to a future reader as alternative entry points to the document. Second, using an interface technique like the one presented in Figure 5c, they may show regions of high interest within the context of the original document. Finally, they may be used for progressive disclosure of a long document, or show an alternative structure to the explicit document structure (see [9] for examples of these other methods). Much like ratings systems, reading ruts can draw people’s attention to what is popular [5,16]; unlike ratings systems, they rely on implicit indicators of interest, not explicit assignment of values.

³ For example, the expected number paragraphs not annotated by anyone is the total number of paragraphs times the product over all reading group members of one minus the number of paragraphs annotated by that member divided by the total number of paragraphs:

$$N_{none} = N_{total} \prod_{i=1}^6 \left(1 - \frac{N_i}{N_{total}}\right)$$

Reading roads also present some very real social and implementational challenges. For example, for this strategy to be effective, widespread harvesting of annotations will have to take place; distributed annotations will need to be accessible, and appropriately protected to preserve the privacy of the individual annotator. When paper materials move from hand to hand (say, as used books), the divide between private and public markings is seamless. Readers, when they relinquish their annotated copies of books do not say, “I hereby publish these annotations.” Instead, they simply pass the annotated books (possibly through an intermediary, like the used book store) to the next reader. It is rare to encounter such seamlessness in electronic media; the act of passing along annotations is far more intentional, especially since it is so easy to remove them.

However, we must remember that reading roads may, in fact, represent ruts. In the long run, they may reinforce patterns of access that are detrimental to the comprehension and interpretation of the entire document. Perhaps the most read paragraph is not the most important paragraph, or the best summary.

Furthermore, a harvesting strategy will need to account for the limits of scalability: to what extent will the collected annotations remain meaningful to others? A collection of student annotations on *Moby Dick* – the reading roads of a college literature course – are of uncertain value to a community of Melville scholars, for example. Although reading roads is a powerful idea, this whole set of social and implementational entailments must be addressed.

2.5 Scenario 4: Connecting people

Handheld devices are beginning to be seen as a way of bringing people together. For example, the somewhat fanciful Lovegety [7] was designed to help people locate potential partners for conversation, karaoke, or something more. The Hummingbird is an experimental device that allows people to be aware of co-location with members of their own circle in crowded places or in distributed office settings [6]. These handheld devices work on simple notions of group members in proximity, or very rough profiles.

Can we connect proximate (or distant) colleagues by what they are reading?⁴ After all, people have long recognized that what we read is a reflection of who we are and what we are doing. Figure 8 illustrates a very

⁴ This strategy has been implemented as an Intranet library application. See [12].

literal interpretation of this insight: books help us find other single booklovers.




Figure 8. We are what we read: a personals ad aimed at bringing together readers.

The books and documents people read along with the ink annotations they make on reading appliances creates a new way to identify potential collaborators. For example, a reader has expressed interest in finding collaborators in his work. He is reading and marking-up a company report on low-emission vehicles and his reading appliance notifies him that “someone in manufacturing has been reading the same report. They are also interested in nitroglycerine fuel cells. Would you like to send e-mail?” A system for connecting people may work best with a *quid pro quo* arrangement: a reader must publish his or her reading and annotation data before he or she can search other people’s reading data for collaborators.

Bringing together people with similar interests and goals as they perform knowledge work can be extremely powerful. The key advantage of reader appliances is taking people’s existing practice of making marks and then inferring interest and matching one persons reading interests to another.

The challenges to this scheme are readily apparent. We have strong cultural expectations of privacy about what we read. Readers’ privacy is protected, for example, in libraries, bookstores, and video rental establishments. Furthermore, people aren’t just what they read. Encounters between people in the same section of a bookstore or library do not often result in long term relationships of any sort (except, of course, in the movies).

3. Caveats and conclusions

E-books exist within a broader setting of work and technology, something that designers of e-book supported cooperative work must account for. In particular, the e-books will likely be used in tandem with paper and traditional desktop computers. Any collaborative strategy for e-books must consider the case in which some people are reading on e-books, and others are reading on paper; for example, some of the reading group members might choose to read this week’s selection on an e-book, and some on paper. It is also possible that an individual might use paper in concert with her e-book; for example, as is evident in

Figure 3, one of the reading group members routinely uses a paper lab notebook.

Furthermore, in a social setting, documents are used in ways beyond just conveying information; placing a paper on your manager's keyboard is different than putting it in her mailbox [15]. In the reading group meeting, members can glance quickly at one another's copy of the article to see if they are "on the same page."

Finally, as with most CSCW technologies, collaborative e-reading will require up-front social negotiation about privacy and reciprocity. Because we assume reading to be an essentially solitary activity, and the marks we make while we are reading to be private, any sharing scheme, whether it is anonymous or not will require agreement about the appropriate situations and scope.

We use paper books and documents in our daily collaborations. Reading appliances have the potential to enhance these paper-based collaborations and also to create new opportunities and ways of working together. Our initial focus has been on people's freeform ink marks as the principal means for initiating collaborative activities with little additional overhead. Collaborative practices using paper books and documents are rich, complex, and varied, and invariably reading from electronic devices will increase this diversity. This article has explored but a few aspects of how portable reading appliances may help us collaborate.

References

1. Adler, A., Gujar, A., Harrison, B.L., O'Hara, K., Sellen, A. A Diary Study of Work-Related Reading: Design Implications for Digital Reading Devices. In: *Proceedings of CHI98*. ACM Press, New York, 1998. pp 241-248.
2. Adler, M.J. and van Doren, C. *How to Read a Book*. Simon and Schuster, New York, NY, 1972.
3. Bly, S. A use of drawing surfaces in different collaborative settings. In: *Proceedings of CSCW'88*. ACM Press, New York, 1988. pp 250-256.
4. Elrod, S., Bruce, R., Goldberg, D., Halasz, F., Janssen, W., Lee, D., McCall, K., Pedersen, E., Pier, K., Tang, J., and Welch, B. Liveboard: A large interactive display supporting group meetings, presentations, and remote collaboration. In: *Proceedings of CHI '92*. ACM Press, New York, 1992. pp x-x.
5. Hill, W., Stead, L., Rosenstein, M., and Furnas, G. Recommending and Evaluating Choices in a Virtual Community of Use. In: *Proceedings of CHI'95*. ACM Press, New York, 1995. pp 194-201.
6. Holmquist, L.E., Falk, J., and Wigstrom, J. Supporting Group Awareness with Inter-Personal Awareness Devices. *Personal Technologies*, March 1999; Vol. 3, No. 1.
7. Iwatani, Y. Love: Japanese Style. *Wired News*, June 11, 1998. <http://www.wired.com/news/news/culture/story/12899.html>
8. Marshall, C.C. Annotation: from paper books to the digital library. In: *Proceedings of Digital Libraries '97*. ACM Press, New York, 1997. pp 131-140.
9. Marshall, C.C. Toward an ecology of hypertext annotation. In: *Proceedings of Hypertext '98*. ACM Press, New York, 1998. pp 40-49.
10. Marshall, C.C., Price, M.N., Golovchinsky, G., Schilit, B. N. Pen, Paper, E-Reader: a study of a reading group. Technical Report FXPAL-TR-99-008, 1999.
11. Myers, B., Stiel, H., and Gargiulo, R. Collaboration Using Multiple PDAs Connected to a PC. In: *Proceedings of CSCW'98*. ACM Press, New York, 1998. pp 285-294.
12. Robertson, S., Jitan, S., and Reese, K. Web-Based Collaborative Library Research. In: *Proceedings of Digital Libraries '97*. ACM Press, New York, 1997, pp 152-160.
13. Schilit, B.N., Golovchinsky, G., and Price, M.N. Beyond Paper: Supporting Active Reading with Free Form Digital Ink Annotations. In: *Proceedings of CHI98*. ACM Press, New York, 1998. pp 249-256.
14. Schilit, B.N., Price, M.N., Golovchinsky, G. Digital Library Information Appliances. In: *Proceedings of Digital Libraries '98*. ACM Press, New York, 1998. pp 217-226.
15. Sellen, A. and Harper, R. Paper as an Analytic Resource in the Design of New Technologies. In: *Proceedings of CHI97*. ACM Press, New York, 1997. pp 319-326.
16. Shardanand, U. and Maes, P. Social Information Filtering: Algorithms for Automating 'Word of Mouth.' In: *Proceedings of CHI95*. ACM Press, New York, 1995. pp 210-217.

17. Silberman, S. Ex Libris: The joys of curling up with a good digital reading device. *Wired*, 7/98: 98-104.
18. Stefik, M., Foster, G., Bobrow, D., Kahn, K., Lanning, S., and Suchman, L. Beyond the Chalkboard: Computer support for collaboration and problem solving in meetings. *CACM* 30(1): 32-47.
19. Streitz, N., Geissler, J., Haake, J., and Hol, J. DOLPHIN: Integrated meeting support across local and remote desktop environments and liveboards. In: *Proceedings of CSCW'94*. ACM Press, New York, 1994. pp 345-358.
20. Tatar, D., Foster, G., and Bobrow, D. Design for conversation: lessons from Cognoter. In *International Journal of Man-Machine Studies* 34(2): 185-209.