Active Reading and Its Discontents: The Situations, Problems and Ideas of Readers

Craig Tashman
GVU Center, Georgia Institute of Technology
85 5th St, Atlanta, GA 30308
craig@cc.gatech.edu

W. Keith Edwards
GVU Center, Georgia Institute of Technology
85 5th St, Atlanta, GA 30308
keith@cc.gatech.edu

ABSTRACT
The increasing popularity of personal reading devices raises the question of how best to support so-called active reading, which involves acts like annotation, note taking, etc. Prior research addressed this question by observing the active reading process, and identifying disparities between computers and paper as a reading medium. We extend this research by 1) investigating the problems that readers experience in their real world tasks, 2) inquiring about their requirements for an ideal reading technology, and 3) updating earlier studies of naturalistic reading behavior, which are several years old now. We present here the results of our investigation, which included a diary study, interviews, and participatory design workshops.

Author Keywords
Active reading, diary studies, participatory design.

ACM Classification Keywords
H.5.2 User Interfaces.

General Terms
Design, Human Factors.

INTRODUCTION
Reading is undergoing a revolution. Specialty devices, tailored to helping people consume text documents, are gaining popularity at a substantial rate, as are electronic versions of texts themselves [25]. Schools are even beginning to deploy these devices in the classroom, with some expecting e-textbooks to occupy 15% of the market by 2012 [6]. And as this revolution develops, with users exploring new form factors and reading experiences, we have an opportunity to carefully reconsider what types of technology readers truly need for their task. One category of reading that has traditionally commanded particular attention in the interactive systems community, and which has been notoriously difficult to adequately support, is the deep, focused reading process known as active reading [13].

CHI 2011, May 7–12, 2011, Vancouver, BC, Canada.

In contrast to how one might read a novel, active reading can be characterized by the greater demand it places on the reader and her media and tools. Active reading (AR) frequently involves searching, highlighting, comparison, non-sequential navigation, and the like.

Over the past two decades, there has been considerable research into AR. In addition to developing systems to support computer-based active reading (e.g., [7, 11]), many have sought to understand the requirements and phenomena of AR itself. This includes broad studies of AR [1], analyses of why paper tends to support AR better than computers [17], and focused research into processes related to AR (e.g., [9, 18, 21]). Together, this research tells us much about how people perform AR, and ways that computers can become more "paper-like" in supporting AR. But it does not address the problems and frustrations readers face that extend across both computers and paper. And while it suggests design recommendations for closing the paper-computer gap, it does not tell us how to move beyond that; it does not tell us what kinds of functionality readers want in order to improve their AR experience beyond what either paper or computers can provide now.

Therefore, to provide added grounding to the design of future systems, we sought to build on this work in three ways:

1. Problems: Using a diary study and interviews, and generally in contrast to prior work, we asked readers directly what difficulties they face in AR, and how they would like their current AR process to change.
2. Ideas: Using participatory design workshops, as well as the diary study and interviews, we asked readers to propose speculative ideas for the digital functionality they need: both to support current AR practices, and to go beyond them to provide a more ideal AR environment.
3. Update prior work: Much of the AR literature is over a decade old; since that time, computing has taken on a broader role in content consumption, and new computing form factors have become common. We thus re-assess how people perform AR tasks, the media they use, the role of computing, etc.

The components of this study help provide different perspectives on readers' behaviors, needs and ideas. The diary study let us gather in-situ, naturalistic data about readers' AR situation over an extended period of time, telling us about their context, their choices of tools, and the like—which would be

difficult to obtain using laboratory methods. Our interviews offered a way to get retrospective feedback from readers, as they generalized over the data they collected though their diaries. And the workshops gave readers a chance to think collaboratively, integrating their personal experiences and ideas with those of others.

While AR possesses different connotations in different contexts, we took a broad lens to identify requirements for AR reading media across a wide range of scenarios, in order to be as inclusive to design opportunities as possible. The broad definition we took in our study was to focus on reading activities that involve more interaction with the reading media than simple sequential advancement through the text (like page turning). This includes, for example, searching a text, comparing texts, annotating, bookmarking, etc., and is consistent with the usage of the term *active reading* in related literature [7, 13, 22]. We note, however, that the boundary between AR and the many other tasks that accompany reading is not always clear. For example, AR often occurs intertwined in a larger workflow such as composing a report, studying for a test, or attending a meeting. To maintain a clear distinction, we define AR *per se* as the part of the workflow involving the actual reading medium itself. And while this is our focus, we address tangential tasks as well where it provides insight into the design of future AR technologies.

RELATED WORK

Several research communities have explored AR, including human-computer interaction, digital libraries, and psychology. We begin by discussing literature related to AR, followed by other areas of research that intersect with AR.

Active Reading

Active reading research. Much prior work has sought to understand the processes that go on in active reading, and to identify ways in which computers fall short of supporting them. To this end, Sellen and Harper conducted an ethnographic study of how knowledge workers interact with documents [23], describing the complex role of paper in subjects' socio-technical environments. Through this understanding, they identified various ways in which computers were inadequate for supporting subjects' paper-oriented behaviors, and suggested design guidelines for changing computers to better support paper-based practice. Focusing on a broader sample population, Adler, et al. used a diary study to explore how, where, what, and why people read at work, and likewise identified many subtleties of paper that people rely upon, and which computers fail to provide [1]. Thus they suggested design ideas for helping computers to better emulate those affordances.

Others have explored the disparity between paper and computers in a lab setting. In 1997 O'Hara and Sellen described how people use paper when doing an AR task, and the ways their behavior changes when trying to do the same task on a computer [17]. They again detail the many places where efficiency declines as people are forced to awkwardly change their behavior to suit the computer—and thus suggest ways for computers to better match the provisions of paper. Morris et al. ran a similar study in 2007, finding that tablet computers were at least as good as paper in supporting AR [13]. And in contrast to the earlier study, they identified a small number of areas where computers provide support paper lacks (e.g., copy/paste and spell-check).

Together, the above studies provide 1) a detailed description of processes people engage in while performing AR tasks, as well as 2) where computers fall short of paper or, rarely, vice versa. But the problem, from a user centered design perspective, is that neither of these gives us a clear picture of the problems readers actually face in AR, independent of medium. In practice, the descriptions of current AR behavior do not tell us where readers feel frustrated or how they would change their tools. Comparisons between computers and paper are valuable, but ultimately only tell us where one medium is deficient relative to the other; if users need something offered by neither, these studies do not show it.

Many systems that have been designed to support AR are in a similar situation. Projects like XLibris [22] insightfully work to replicate, and even surpass, the key properties of paper using computers. PapierCraft goes further, building a hybrid paper-computer AR environment to try to offer the best of both worlds [7]. These projects present compelling design ideas, but typically, these ideas come from 1) replicating the properties of paper on computers, and 2) the designers' insights into what new functionality readers need. This may result in novel AR technologies, but it remains unclear whether the novel functionality offered really addresses the functions readers want and the difficulties they face in their real world AR tasks.

AR-related research. Rather than exploring AR per se, other research has investigated subsets of, or areas related to, AR. Both Marshall and Wolfe, for example, looked at how students use annotation, in the hope of supporting it digitally [8, 27]. Others have explored how people use annotation for personal versus collaborative purposes [10, 20]. Past work has likewise revealed much about other aspects of AR, including: non-linear navigation (e.g., [9, 16]), the need for bimanual control [17], the importance of skimming [15], cross-document navigation, and the ability to create short-term bookmarks [2, 4]. Other studies considered how readers use notes or outlines created outside the document being read [5, 11, 18].

But while this research provides insights into many aspects of the AR process, this again does not tell us where this current process frustrates readers, or is inadequate. And while this research often contains design requirements for supporting the current process, it does not reveal to us the kinds of additional functionality users want in order to improve upon their current processes.

There are exceptions, though. For example, in creating an intervention specifically to support note taking, Wilcox et al.

observed note takers' current behavior, but also enquired of subjects both what difficulties they face with their current tools, and what an ideal tool would be like [26]. But these exceptions are rare—thus leaving a gap in the existing AR research: we do not actually know how readers would like their current AR process to change, or what tools they would like for support.

This paper is intended to be a step toward addressing this gap: We thus sought to engage readers in the design process, using diaries, interviews and workshops to let them tell us about their real-world AR frustrations, and letting them tell us the kinds of new functionality they want in an ideal AR system. Additionally, much of our current understanding of real-world, naturalistic AR behavior is well over ten years old (e.g., [1, 23]); this research also helps to update aspects of these earlier findings.

STUDY OF AR PRACTICES AND DESIRES

Participants

With notable exceptions, many studies of AR focus on fairly narrow subject pools (e.g., students). So to investigate the subject generally, we recruited from a large, vertically integrated, design and manufacturing firm—which offered us a broad pool, ranging from woodworking engineers to researchers. We spoke with participants before enrollment to screen out people who were not frequent active readers, and offered participants a $5 lunch voucher for their time.

Of the 24 people who initially volunteered, 20 completed the study, 13 of whom were women. The jobs of these 20 were as follows (note that the students were also interns):

Manager	P1
Design student	P2, P3, P5
Design researcher	P4, P6
IP Lawyer	P7, P9
Math/Computing student	P11
Mgmt info systems student	P12
MBA student	P13, P16
Engineering student	P15, P24
Ass't project coordinator	P17
Corporate strategist	P19, P22, P23
Marketing, mgmt student	P20
Senior executive	P21

Table 1. Participant Job/Role

Methodology

Our study consisted of four phases: 1) A diary task where readers logged all of their AR activities, 2) a semi-structured interview about readers' AR behaviors and problems, 3) exposure to an early prototype AR support system we developed, and 4) a series of participatory design workshops where users collaboratively designed future AR environments. In this section, we explain the methodology and reasons for each of these phases.

Diary Task

To provide a picture of real-world AR tasks and behavior, we asked participants to record a diary entry each time they completed an AR task over the course of generally one week (but up to three weeks). Each entry included answers to ten short questions addressing how, why, and where the task was performed, as well as the most difficult part of the task and how it could have been improved. We gave participants physical journals to act as tangible reminders to record entries, and we sent them daily reminder emails as well. Eighteen of the twenty participants turned in their diaries.

The first objective of the diary task was to understand the AR process as it actually occurs, from readers in or near the situations where it occurs. This diary approach is consistent with earlier AR studies [1], and was intended to reduce our dependence on retrospective reconstructions of AR events.

The second objective was to make readers aware of their own AR behaviors, difficulties, etc. This was to keep ideas and feedback in the subsequent phases of the study as closely grounded in the reader's actual AR situation as possible.

Interview

In contrast to the very structured diary task, our interviews gave participants a way to direct the conversation and tell us about the areas of AR most important to them—as well as a way for us to collect additional information from them, such as their professional experience, what kinds of reading they did, the problems they faced in AR, and their ideas for solutions. All twenty volunteers participated in the interviews, which we recorded and took notes on as well.

Having been made aware of the details of their behavior through the diaries, the interviews also acted as a way to let participants generalize over those details and give us their broader perspective on AR. It also served as an opportunity for participants to reinforce their memories of their AR situation to better prepare them for the design workshops.

Prototype Exposure

In our prior experience with participatory design, we observed participants tending to create incremental ideas, largely drawing upon the kinds of interactions they had already experienced. In order to broaden our participants' thinking in advance of this study's workshops, we had them perform an AR activity using a prototype system we had developed. We chose this prototype system as it supports interaction different from either paper or traditional computers, and would thus potentially broaden participants' interaction experiences. To mitigate the risk of unduly influencing participants' ideas, we waited until after both the diary task and the interviews before exposing participants to the prototype (with the exception of P1, P7, and P22, who briefly saw the prototype prior to the study).

While the details of the prototype are outside our scope (see [24]), the system provided users with multitouch gestures for manipulating documents, and was running on a Dell XT2

multitouch tablet PC. The prototype included functions for collapsing document areas to bring disparate text into proximity, and supported various gestures for creating, interconnecting and aggregating document excerpts and reader comments. Parallelism was central to the system's design, and allowed users to engage many navigation widgets at once. It likewise supported various forms of linking within and across one's documents, particularly to connect comments and excerpts to their original contexts. Together this collection of interactions provided an experience different that of paper or traditional computer-based AR systems.

Naturally, intentionally exposing participants to a technology incurred a risk of artificially influencing them. But as we could readily distinguish ideas that may have been influenced by the prototype from those that came beforehand in the interviews and diaries, we concluded the risk was small, especially if this gave participants a broader base of ideas from which to draw.

Design Workshop

The final part of our study was a series of participatory design workshops, which allowed our readers to build on their awareness of their AR process, and the ideas they had thought about and been exposed to so far. To keep the workshops of manageable size, we divided participants into two groups: Group One had seven participants and Group Two had five. Since the structure of the workshops required even numbers of people, and some of the original 20 participants could not attend for scheduling reasons, we recruited two extra participants just for this phase.

The three-hour workshops were structured to gradually bring participants from scenarios to design: We began by having participants collaboratively describe the current state of AR, first listing scenarios where AR takes place, then the types of processes that go on within AR tasks, and finally the problems faced during the AR process. Next, participants brainstormed broadly about the ideal AR environment of the future—coming up with either specific ideas or general concepts as they wished.

For the remainder of each workshop, we broke participants into teams of two, asking each to think of a list of concrete ideas for creating an improved AR environment. After sharing their ideas, each team selected a concrete idea to incorporate into a physical mockup AR system, which they assembled out of various materials we supplied. We also provided each team a digital camera setup enabling them to create simple flip-book movies showing how their mocked-up system would be used.

Data Gathered

In total, we received 106 diary entries, with an average of 5.9 per participant (s.d.=3.4). We took a liberal view of what constitutes AR, and accepted diary entries involving non-sequential navigation, annotation, the production of output documents or notes, the use of multiple documents at once, or content searching within or among documents. With these criteria, we excluded nine of the above entries. We also obtained 200 minutes of audio recordings from the interviews, and 6 hours of video form the workshops.

Analysis Methodology

Drawing from the diaries, interviews, and workshops we coded the data from each phase of the study independently, in order to allow us to distinguish in-situ data (the diaries) from retrospective data (the interviews), and both of those from collaboratively generated data (the workshops). We also coded the two workshops separately to allow us to identify differences between their results.

Within each separately coded part of the study, we principally relied on open coding as it readily identifies larger categories and themes in the data. In cases where many different dimensions of categorization were possible, we drew from prior AR literature (e.g., [16]) in establishing first-level categorizations.

FINDINGS

In exploring people's AR experience, we observed several themes crossing the different phases of the study. Certain well known problems, such as the navigation difficulties of computers, were strongly present. Other issues though, like document visualization and difficulties maintaining awareness, came as more of a surprise. So to set the context for discussing these and other issues, we begin by considering how people perform AR today, as revealed by our diary study and interviews.

The State of Active Reading

Reading goals. To understand the problems readers face in AR, we first wanted to understand their motivations for engaging in AR. Our findings here came principally from the diaries and were in mostly line with existing literature (e.g., learning, condensing, etc. [14, 16]). However, we did identify two additional goals. The first, *organizing content*, was observed 11 times (out of 63 unambiguous entries), such as when, P2 was clustering her interview transcripts by topic, or P3 was clustering car repairs by priority. The second, *comparison*, was observed twice, such as when P15 was evaluating different configurations while shopping for a personal computer.

Personal reading. This shopping example is also significant as an example of personal AR. Although AR is generally taken to be a work-related phenomenon (e.g., [1, 23]), we found four examples to the contrary, including P17's highlighting key quotes from a novel, or P1's claim to annotate virtually everything he reads except fiction. And expectedly, personal AR tended to be performed in personal, informal spaces, including lawn chairs and couches.

Active reading locations. While the importance of mobile AR is known [1], we were surprised by the breadth of mobile settings used. For example, we found participants doing

work-related AR in bed, at picnic tables, at the kitchen table, and in a car. Generally, of the 73 diary entries with unambiguous locations, 19 (or 26%) involved AR in mobile locations. So, in line with earlier work [1], this suggests readers broadly require a relatively portable means to support AR.

Computers and paper. Complicating the need for portability was readers' continued reliance on multiple media for performing AR tasks. Similarly to earlier work [23], we found that approximately 25% of unambiguous diary entries described using paper-plus-computers for a given AR task. But while [23] found that a majority of the time, document tasks were paper-only, we found 63% of AR diary entries were exclusively performed on the computer—suggesting that computers are indeed overtaking paper as the AR medium of necessity. This is not to say that it is the medium of choice—participants generally continued to express a preference for paper in AR tasks. Even the extremely paper-like tablet-PC was described by P21 as requiring too much time to access the inking functionality.

One reason people mixed computers and paper was corporate policy—as in filling out an online expense form. But in other cases, computers were simply better at handling some parts of a task. For example, P23 described printing documents to read them, before going back to the computer to highlight them. Generally, when comparing media, tasks demanding more "pure" reading were seen as better suited to paper. Part of the reason for this was that paper could show more text at once than a computer.

Significantly, this shortcoming of computers was noted in 1997 [17], but was suspected to have been resolved by 2007 [13]. And while the benefits of larger displays have been seen for specific AR tasks under laboratory settings [3], we find that it is still a problem in actual AR activity as well. And the most obvious solution—increased screen space—conflicts with the importance of portability.

Besides display space, others described a preference for the tangibility and navigability that paper afforded; one further noted how that physicality can lead to feelings of guilt over printing to paper.

But paper had drawbacks as well. P12 and P6 both describe frustration with physical books because of the difficulty reading near the crease of the page, and the weight of many hardcovers. In line with prior lab studies [19], P1 and P16 discussed paper's lack of margin space, while others felt entering/editing information is harder using paper than computers.

In a similar manner to [23], our participants saw computers as perhaps better suited to linear, structured AR tasks. P2 described them positively as, "forcing me to be neater." But paper appeared superior for tasks with many parallel components, without clear structure, and where added flexibility was needed. P2 continued,

"When I work in digital form, I want to be more codified in how I approach things...with pencil and paper, I like the ability to feel loose. It doesn't feel like as much of a process as it does when it's ...on the computer."

But as computers and paper each have domains where they are superior, some research has found people comfortably using them together (e.g., [11]). But by contrast, our participants revealed considerable frustration using paper and computers together. Part of the problem was that readers felt frustrated that paper does not support computer-like functions like drag-and-drop. The difference in spatial orientation caused regular frustration as well when participants had to switch from looking up at their monitors to down at their paper. Part of this disparity may be resolved by putting computer monitors and paper at a common viewing angle, but the differences in contrast and light reflection properties noted by participants may make a solution more complex.

Output materials. AR is known to involve collaboration in certain domains [23], and we found that over 73% of the 57 diary entries with clear outputs (i.e., materials resulting from the AR task) were collaborative. And while the outputs varied between broadcast (e.g., blogs) or narrowcast (e.g., presentations), and in creating content new or modifying existing material, the single most common type of output was email, which was found as an input and an output in about 20% of entries. By contrast, earlier studies pointed to email as a way to send AR documents [23], but did not suggest emails themselves as the focus of AR tasks. Besides email, we found a variety of other collaborative materials were produced, including summaries, analysis reports, PowerPoint slides, etc. Generally, we found that AR is not confined to documents designed for careful, deep reading. P21, for example, describes trying to make sense of all the different authors involved in several long chains of email.

By contrast, 27% of outputs were not intended for sharing, but these were usually intermediate results, such as lists of notes, outlines, or sketches.

Active Reading Problems

Readers listed numerous AR breakdowns related to many of the process's major phenomena, such as navigation, annotation, search, content awareness, etc. Most of these problems came from the diaries, where we asked readers about the hardest parts of their AR experience. Analyzing the results yielded 92 clear issues, which we organized into 21 categories. We also noted 30 issues identified in the design workshops, and several more in the interviews.

Some of the most prominent categories of challenges are discussed above. Visibility, for example, comprised 11% of issues identified in the diaries, and included seeing many lines of text, documents or notes at once. Media switching, between paper and computers, comprised about 8%. Here we discuss several additional categories of issues.

Navigation. Inter-document navigation issues constituted over 16% of those identified in the diaries. One well-known example of this is switching among many application windows

(and as we found, spreadsheet tabs) [17] involved in an AR task. This was partly because readers forgot their places in their documents, as P19 described this in his diary,

"[It is hard]…jumping back and forth between documents and trying to remember where you were. I would do this much differently physically. I'd have my five documents spread on the table and make a mark where I was at."

Likewise, readers struggled with just keeping track of all the different materials and their significance. P13 explained this, "[It's hard] managing opening multiple documents, understanding the context, relevance, and usage of the documents."

Users' struggles with multiple-document scenarios have several implications. First, the fact that multi-document usage is both common and especially problematic strongly suggests that researchers exploring AR must consider multi-document studies to be the "default" scenario; many studies (e.g., [13, 17]) have only considered the case where a user is reading from one document at a time. Second, multi-document reading further complicates the design of future AR technologies by putting even greater demands on already limited visual real-estate, perhaps demanding more creative visualization solutions.

Besides inter-document navigation, about 5% of issues identified related to *intra*-document navigation. For example, P16 notes,

"…especially when the print-outs are not on the same page; printed on different pages and you've got to flip back and forth which drives me nuts."

While prior research described how readers freely arranged pages of paper documents to view disparate document areas in parallel [17], this clearly breaks down for bound documents, as the quote illustrates. In some ways then, intra-document paper navigation is similar to inter-document navigation on computers: one can often see only one page at a time and so must rely more on working memory than perception to integrate disparate content. That is, paper's spatial flexibility is important to readers, allowing them to arrange things freely, but this seems to be where that flexibility reaches its limit.

Generally, one implication of readers' navigational difficulties is to emphasize the point that AR navigation is non-linear. Readers do not consume a document or a page and move on to the next one. Rather, they cycle between documents and portions of documents often in complex integration tasks.

Annotation. In the diaries, annotation issues were less common, comprising about 5%, and largely relating to software limitations, such as applications which did not allow comments or highlights, forcing the reader to put thoughts in separate documents. The workshops and interviews though, revealed several additional issues. In the interviews 5 readers cited difficulties making margin notes on paper, especially due to the lack of physical space for doing so. Other comments concerned managing annotations—such as re-finding one's annotations after the fact, easier editing of sticky notes, and a desire for annotations to intelligently make themselves known at the times and places where they are needed.

In an interesting counterpoint to readers who found themselves unable to annotate at all, one of the workshop discussions considered the opposite—the temptation to annotate too much. This effectively makes things like highlights meaningless leaving the reader uncertain of why she annotated that material. This may partly be due to readers' use of media that can not support the full breadth of different annotations (with concomitant semantics) that readers produce.

While annotations have been studied before [21], our analysis suggests that non-text annotations can be categorized along two dimensions (Figure 1): *inter-page* (i.e., marks visible across pages, such as a dog-ear or bookmark) versus *intra-page* (like an underline), and *generative semiotics* (like sketches, where the reader can put marks together to construct more complex meanings) versus *simple semiotics* (again, like underlining). Different media then, support different points along these two dimensions. For example, a traditional word processor supports simple, intra-page semiotics like highlights and underlines, but makes it awkward to create marks complex enough to be generative. Likewise, a book supports intra-page as well as inter-page, simple semiotics, such as dog-earing and sticky notes. And notably, while books also support generative semiotics within the page, no commonly used medium seems to offer support for semantically rich markings that can cross page boundaries (inter-page, generative semiotics). So the use of media that do not adequately support this space may be a way to look at why readers could not express their annotations in a way that would be richer and thus more memorable.

	Simple semiotics	Generative semiotics
Inter-page	*Dog-ear*	*[unknown]*
Intra-page	*Underline*	*Sketching*

Figure 1: Two dimensions of non-text annotation, with examples for each quadrant.

Awareness. One challenging aspect of AR was maintaining an awareness of all of the background information relevant to what one was reading at the moment, which comprised about 8% of issues in the diaries. This information included context established elsewhere in the documents at hand (such as who said what in an email), as well as one's earlier thoughts and reflections, and even location within the material. For example, P21 described doing frequent flipping through a document because of the difficulty of trying to keep so much data in one's head at once. By contrast, P24 described taking notes from multiple sources but losing track of where each note came from. This latter issue is interesting as it highlights a challenge of using paper in multi-document work, as it appears to offer little support for capturing or viewing cross-document relationships.

Note taking. While the diaries did not include challenges relating to note taking, several were noted in the interviews. Many of the issues cited came from creating new notes, ranging from the mechanics of the copying text (P23), building and managing outlines (P11), and anticipating what will be important in the context of a text in order to decide what to copy (P2).

Participants also described issues such as determining how best to organize their notes, or how to integrate together notes from multiple sources. And, because notes are separate form the source text, a document switching task was often involved as well. This was explained by P2 during the design workshops:

"[Annotation] gets noisy because I'm highlighting or tagging different things in the same document; and [with notes] I'm creating new documents but then I've got a slew of new documents and I'm constantly going back between the originals and the new documents…"

Notes then, in contrast to annotations, appear to be at a greater risk of reducing the efficiency of the AR process. Nonetheless, their function as a means of centralization and aggregation appears to be essential.

Retrieval. Even after their creation, readers continued to struggle with notes—and annotations—in the retrieval process. Like annotations, retrieval of notes was hard in part because readers needed a thought or excerpt about a document at a particular moment in life when it was relevant, and it was hard to make these connections naturally. As we discuss below, readers mentioned some ideas to address this, including integrating an AR system with a to-do list.

The design workshops also included considerable discussion on the problem of recalling needed context when returning to an entire AR task. This observation led to several participant ideas we discuss in detail in the next section.

In summary, we briefly review some of the problems we have discussed so far. First, paper and computers each have shortcomings—paper is flexible but often laborious, while computers can be efficient but overly structured. Among the major problem areas we considered (navigation, annotation, awareness, note taking and retrieval) one frequent theme was re-acquiring needed mental state—as in recalling the meanings of annotations or notes, or even of entire documents. This difficulty has led multi-document AR tasks to be especially challenging because of the increased memory burden, but even single-document tasks can have this problem when the text is a large document. Mechanical issues are less common, but still significant when creating notes. In the next section, we discuss our participants' ideas and solutions for addressing these and other AR problems.

PARTICIPANT IDEAS

Through the study, participants generated over 170 "ideas" for the future of AR, ranging from vague suggestions and incremental changes to radical, novel ideas. These ideas also ranged from addressing active reading itself (such as ways to annotate) to tangential tasks (such as to-do list functionality). Notably, several of these ideas have appeared in other research or commercial systems that our participants were not familiar with. We report on all of these however, because they offer insight into the types of tools that people want integrated into their AR environments, and thus have bearing on the design of future e-readers and other AR technology.

We first present ideas generated in the diary phase, before participants were exposed to the prototype, followed by those coming afterward, during the design workshops.

Diaries

The most situated of participants' ideas came from the diaries—where participants focused on improving their immediate AR experience, as opposed to supporting entirely new experiences. We classified the 65 relevant diary ideas into 14 themes; the most prominent of which we present here.

Visual layout of content. Likely in response to the significant problems in navigation and visualization, this category focused often on getting disparate content into visual proximity. Suggestions sometimes took the form of changes to existing applications, such as P19, who wanted Excel to freeze more than one column at a time and PowerPoint to show more then one slide at a time. Other suggestions included ways to view all the highlighted portions of a document at once, or various visualization ideas such as P13's using visual links to show document relationships, or P5's suggestion to use focus-plus-context views.

Increased directness. Several participants described using more direct ways to interact with, and especially copy, text. While one offered a pen-based approach, most were interested in touch-based methods to drag content, as with P21,

"Would be great for touch functionality…just touch the screen to highlight something, flip between pages, etc.; That'd be much easier than dealing with the tablet pen."

It is important to note that our recruitment instruments mentioned our touch-based AR prototype, so even before seeing the prototype itself, our sample might have been especially inclined to the idea. Still, their ideas are worth noting as ways touch could be used in support of AR.

Linking. Also likely as a response to navigation problems, several participants described improved ways of linking pieces of content together. P3 described "threads" that could connect paper documents together, and then extend to link to digital material as well. Others described linking audio notes to a location in a text, or a centralized log book linking notes to many different documents.

Annotation expressiveness. Potentially as a response to the above problems of inadequate annotation semantics, several readers also described richer ways to embellish content, such as more colors of markers or color-changing paper. Likewise, P9 described dimming irrelevant parts of a document as a

way to reduce visual saliency—effectively the opposite of a highlighter, which increases saliency. Participants also suggested multimodal annotation, in which voice could be used to annotate content.

Collaboration. In a reflection of the often collaborative nature of the inputs and outputs of the AR process, three participants described synchronous and asynchronous collaboration ideas. For example, P21 wanted long emails to be color-coded to show what text was written by whom.

Workshop Discussions

Unlike the diaries, the workshop discussions afforded an opportunity to build on one another's ideas, allowing participants to focus on developing the thoughts that were more interesting to the group. We discuss several of these areas of focus here.

Collaboration. Participants in both workshops independently showed a strong interest in collaborative AR, both synchronous and asynchronous. The scale of sharing, though, varied between quick excerpts and one's entire AR workspace—suggesting such variability would be important to an actual collaborative AR system.

State capture. Both workshops also discussed a desire to capture the state of the AR environment. Workshop 1 (WS1) was mainly interested in recording one's AR history in order to build a corpus of cues to help readers regain their train of thought when returning to an AR task after an interruption. WS2, by contrast, explored state capture as a way to switch between and transport AR "workspaces". They then imagined a special pen that, besides annotation, could also store and transfer all or part of a workspace to support collocated collaboration.

Hardware platform. Developing a large-screen workspace, participants in WS2 considered several details about their hardware platform. First, participants expected a multitouch-plus-stylus work surface. Touch was seen as better for manipulation, like scaling and moving materials, whereas the stylus was seen as better for selection and highlighting.

Addressing the cross-media problems described above, participants considered embedding scanners in their workspace to easily digitize papers (as in [12]), which they refined into a paper-sized tablet computer sitting flush in a large table-top surface. The tablet could then be removed to focus on one document in a different setting (e.g., at a café), and returned to the table-top when the user needed more space. We found this an interesting resolution to the paper-computer dichotomy: replacing both media with a unified hybrid, especially since modern tablets do appear comparable to paper [13].

By contrast, WS1 was primarily interested in exploring additional input and output modes. Besides haptic and audio feedback, they sought parallelism, such as letting the user read text visually while the computer read text aloud. For input, they envisioned the use of gaze, chorded keyboards (which they borrowed from pianos), and a stylus for creating and manipulating content. Generally, they wanted to maximize efficiency while providing what they called "natural interactions" for controlling the AR environment.

Sense-making. Of particular interest to WS1 were ideas to help readers better understand a text, especially using visualization, such as semantic zooming, to show in-place keywords for a given region of text. Recognizing that many of their documents may have been read and reviewed by others, participants also suggested automatically searching the web for summaries and presenting tag clouds of their most frequent terms. To help readers identify important areas of a text, WS1 also suggested readers could tag some portion of the text as important, then the system would identify other similar text—either at the word or conceptual level.

Annotation. In line with the interest in more flexible annotations noted in the diaries, WS1 considered ways to annotate annotations. They discussed organizing annotations into layers, and assigning annotations levels of importance. They also wanted to use these rich annotations as a search mechanism, allowing the user to find all parts of a text that have been tagged a given way.

Rather than add more structure, WS2 was interested in free-form annotation. They envisioned, on their large, pen-touch workspace, the reader could freely write not only on documents, but across them, depicting more complex inter-document relationships. This fits well with the importance of inter-document scenarios.

Organization. Participants differed in the level of organization they desired, as evidenced by an outright disagreement in one workshop about whether excerpts of content should be organized in well-structured folders or in freeform piles. The result is instructive—though many readers may need organizational affordances, the amount of structure these impose needs to vary depending on the person and the situation.

Workshop Mockups

Having brainstormed an array of ideas, the prototype mockup activity gave participants a chance to integrate and refine a selection of those ideas into simple system designs. We review the general direction taken by several groups.

WS1, team 1. Participants sought to present a history of one's AR processes by using a multitouch, interactive timeline to show when notes were taken relative to known events, like the items on one's outlook calendar.

WS1, team 2. Team 2 stressed tangibility and portability, designing a tabletop plus e-paper hybrid environment to allow the user a seamless, synchronized transition between the two media.

WS2, team 1. Also building on a tabletop-plus tablet hybrid, this team proposed tighter integration with other PIM tools, including using a to-do list as a workspace manager. This allows one to click a to-do item to move to the associated workspace.

WS2, team 2. This team designed a system to support the collaborative design workflow, with AR as just one piece of that process. The system thus supported sketching, video-conferencing, and could link parts of designs to documents.

DISCUSSION

Cumulatively, many of our findings are consistent with earlier results, such as the importance of mobile and collaborative AR scenarios, and the difficulties of navigating documents using computers and of switching between computers and paper. But extending and updating this prior work, we have observed that computers are now a more common AR medium than paper. Consequently then, people are now performing AR on a wide variety of computer-based content. Email, for example, is now one of the most common AR outputs. We also found AR to extend past documents designed for deep reading, including PowerPoint presentations, spreadsheets, web forums, etc. Notably though, only one participant listed an example of new media in AR—P11, who was creating a blog entry. And only P1 performed AR using a small mobile device (a Blackberry). Generally, it seems AR tasks and habits may evolve more conservatively than computing as a whole.

Besides observations of how AR is done, we have extended earlier work to also show that users have numerous frustrations beyond the deficiencies of computers, such as the weight and slope of paper books, the difficulty comparing parts of a bound document, paper's lack of margin space, the challenge of finding annotations, the tendency of annotations to become meaningless, the difficulty expressing and maintaining awareness of cross-document relationships, and the like. So, even a perfect mix of modern computers and paper would still leave readers with significant challenges.

Readers though, had numerous ideas for possibly addressing some of these challenges and, more generally, revealing requirements for a more ideal AR system. Some of these have been noted in the AR literature, such as pen-plus-touch input. But others, such as history recording, better multitasking support, annotation across documents, and various visualizations for sense-making are not often examined in the context of AR.

From a design standpoint, we also observed that participants did not so much borrow ideas from the prototype system, but rather built on its general mode of interaction. We saw this with the workshops making use of multitouch interfaces, but adding stylus input as well. Likewise, participants' designs often included linking between pieces of content, but extended this idea to include drawings and to-do lists. Generally, participants took their designs in directions quite distinct from the prototype, such as collaboration, workspace management, text visualizations, and the like. Ultimately then, we believe the prototype exposure served its purpose in giving participants a broader base of interaction techniques to draw from, but without stifling their own ideas.

Tensions

Our observations however, lead to several tensions for designers of AR systems, one of which relates to scale. Readers made it clear that current computers do not offer enough space to support real world multi-document tasks—like comparing and integrating information between texts. Yet at the same time, readers need to be able to carry their AR workspaces with them, quickly set up at a café or kitchen table, and equally quickly remove them. And the natural solution to this—switching among multiple document windows on a small screen—was found to be an important source of frustration in computer-mediated AR.

Of course, one solution may be to provide better visualizations involving focus-plus-context views and semantic zooming, which participants suggested, to help promote awareness of multiple documents. But even this is complicated by participants' desire to flexibly annotate on, between, and outside documents. Interestingly, readers did not seem to expect a single solution to these problems. They envisioned hybrid systems, with different, but tightly integrated, platforms for different parts of the AR process.

A second tension results from the role of AR in readers' larger workflows. Rather than being isolated, participants' workshop designs show AR to be closely connected to one's calendar and to-do list; the diaries show it to happen in email, web browsers and in Microsoft PowerPoint and Excel. But yet, participants hated reading from computer monitors, and were very interested in exotic, often touch-based, input modalities not common on traditional PCs. So presently, it is unclear how to maintain the level of integration participants expect, while providing the rather unusual input and output technology they want for doing AR.

But while the above issues raise questions for designers, other AR issues can be addressed more readily:

- Support collaboration. To our surprise, participants indicated that collocated and distant, as well as synchronous and asynchronous, collaboration are important parts of AR. Current systems (and paper) have very limited support for this aspect of AR.
- Support flexible annotation. Besides being freeform, annotations should be able to span page and document boundaries, and should have properties like priority or category. Annotations and annotated content should be able to be aggregated together and organized, but with a level of structure chosen by the user.
- Support memory. Like many tasks, AR is complex and imposes a substantial memory burden. Readers are looking for possible cues that might refresh their sense of what a document was about, or what they were doing when returning to an AR task. This is a general problem, but because the AR process is well understood, it may be possible to design these cues more effectively for AR tasks than in the general case.

- Visualization. Corollary to supporting memory, participants wanted ways to take in data as efficiently as possible. Visualizations of document content (e.g., tag clouds) and relationships (e.g., links among documents) are part of this. Another part is helping users maintain awareness of original content and their own thoughts by letting them show several parts of several documents or notes at once.
- Multiple workspaces. As is true in the general case studied in the window manipulation literature, readers engage in multiple AR tasks at once, and need ways to efficiently switch among them.
- Directness. For manipulation tasks like copying text, efficiency and a sense of directness seem to be important. As seen in the workshops, readers were interested in both pen and touch input for this purpose, but generally seem to want a hybrid pen-plus-touch input model.

In conclusion, it has been well understood that AR is difficult to support, and our study has sought to more deeply explore the reasons for this, identifying the problems readers face across media. But by bringing readers themselves into the design process, we have a sketch of the types of solutions they are looking for. And while some of our readers' requirements present serious technological and design challenges, many of the things they are looking for with input modalities, visualization, linking and annotation, are practical with technology today.

ACKNOWLEDGEMENTS

We would like to thank Steelcase, Dell and the NSF (award #IIS-0705569) for their support of this research.

REEFERENCES

1. Adler, A., Gujar, A., Harrison, B.L., O'Hara, K. and Sellen, A. A diary study of work-related reading: design implications for digital reading devices *CHI*, ACM Press, 1998.
2. Askwall, S. Computer supported reading vs. reading text on paper: a comparison of two reading situations. *International Journal of Man Machine Studies*, *22*. 425-439.
3. Costabile, M.F., Paternò, F., Bae, S., Badi, R., Meintanis, K., Moore, J.M., Zacchi, A., Hsieh, H., Marshall, C.C. and Shipman, F.M. Effects of Display Configurations on Document Triage. in *Human-Computer Interaction - INTERACT 2005*, Springer Berlin / Heidelberg, 2005, 130-143.
4. Dillon, A., Richardson, J. and McKnight, C. Human factors of journal usage and design of electronic texts. *Interacting with computers*, *1*. 183-189.
5. Golovchinsky, G. Reading in the office *Proceeding of the 2008 ACM workshop on Research advances in large digital book repositories*, ACM, Napa Valley, CA, USA, 2008.
6. Laxer, M. With e-books Still Evolving, Professors Remain Unsure of the Advantages, Northwestern University, Chicago, IL, 2010.
7. Liao, C., Guimbretiere, F., Hinckley, K. and Hollan, J. Papiercraft: A gesture-based command system for interactive paper. *ACM Trans. Comput.-Hum. Interact.*, *14* (4). 1-27.
8. Marshall, C.C. Annotation: from paper books to the digital library *Proceedings of the second ACM international conference on Digital libraries*, ACM, Philadelphia, Pennsylvania, United States, 1997.
9. Marshall, C.C. and Bly, S. Turning the page on navigation *Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries*, ACM, Denver, CO, USA, 2005.
10. Marshall, C.C. and Brush, A.J.B. Exploring the relationship between personal and public annotations *Proceedings of the 4th ACM/IEEE-CS joint conference on Digital libraries*, ACM, Tuscon, AZ, USA, 2004.
11. Marshall, C.C., Price, M.N., Golovchinsky, G. and Schilit, B.N. Designing e-books for legal research *Proc. of ACM/IEEE conf. on Digital libraries*, ACM, Roanoke, VA, USA, 2001.
12. Microsystems, S. Starfire, 1992.
13. Morris, M.R., Brush, A.J.B. and Meyers, B.R., Reading Revisited: Evaluating the Usability of Digital Display Surfaces for Active Reading Tasks. in *Tabletop*, (2007), 79-86.
14. Murray, T., Applying Text Comprehension and Active Reading Principles to Adaptive Hyperbooks. in *Cognitive Science*, (Boston, MA, 2003).
15. Muter, P., Latremouille, S.A., Treunit, W.C. and Beam, P. Extended reading of continuous text on television screens. *Human Factors*, *24*. 501-508.
16. O'Hara, K. Towards a Typology of Reading Goals *RXRC Affordances of Paper Project*, Rank Xerox Research Center, Cambridge, UK, 1996.
17. O'Hara, K. and Sellen, A. A comparison of reading paper and on-line documents *CHI 1997*, ACM, 1997.
18. O'Hara, K., Smith, F., Newman, W. and Sellen, A. Student readers' use of library documents: implications for library technologies *CHI 1998*, ACM Press, 1998.
19. Pearson, J., Buchanan, G. and Thimbleby, H. Improving annotations in digital documents *European conference on Research and advanced technology for digital libraries*, Springer-Verlag, Corfu, Greece, 2009.
20. Qayyum, M.A. Capturing the online academic reading process. *Inf. Process. Manage.*, *44* (2). 581-595.
21. Renear, A., DeRose, S., Mylonas, E. and Dam, A.v. An Outline for a Functional Taxonomy of Annotation, Presented at Microsoft Research, Redmond, WA, 1999.
22. Schilit, B.N., Golovchinsky, G. and Price, M.N. Beyond paper: supporting active reading with free form digital ink annotations *CHI 1998*, ACM, 1998.
23. Sellen, A. and Harper, R. Paper as an analytic resource for the design of new technologies *CHI 1997*, ACM, 1997.
24. Tashman, C. and Edwards, W.K., LiquidText: A Flexible, Multitouch Environment to Support Active Reading. in *CHI 2011*, (Vancouver, Canada, 2011), ACM.
25. Wahba, P. E-readers gain traction, spur sales: poll, Reuters, 2010.
26. Wilcox, L.D., Schilit, B.N. and Sawhney, N. Dynomite: a dynamically organized ink and audio notebook *Proceedings of the SIGCHI conference on Human factors in computing systems*, ACM, Atlanta, GA, United States, 1997.
27. Wolfe, J.L. Effects of annotations on student readers and writers *Proceedings of the fifth ACM conference on Digital libraries*, ACM, San Antonio, Texas, United States, 2000.