

Navigation Techniques for Dual-Display E-Book Readers

Nicholas Chen¹, Francois Guimbretiere¹, Morgan Dixon¹, Cassandra Lewis¹, Maneesh Agrawala²

¹University of Maryland
3452 A.V. Williams Building
College Park, MD 20742

²University of California
635 Soda Hall
Berkeley, CA 94720

{nchen, francois, mdixon3}@cs.umd.edu, clewis@umd.edu

maneesh@cs.berkeley.edu

ABSTRACT

Existing e-book readers do not do a good job supporting many reading tasks that people perform, as ethnographers report that when reading, people frequently read from multiple display surfaces. In this paper we present our design of a dual-display e-book reader and explore how it can be used to interact with electronic documents. Our design supports embodied interactions like folding, flipping, and fanning for local/lightweight navigation. We also show how mechanisms like Space Filling Thumbnails can use the increased display space to aid global navigation. Lastly, the detachable faces in our design can facilitate inter-document operations and flexible layout of documents in the workspace. Semi-directed interviews with seven users found that dual-displays have the potential to improve the reading experience by supporting several local navigation tasks better than a single display device. Users also identified many reading tasks for which the device would be valuable. Users did not find the embodied interface particularly useful when reading in our controlled lab setting, however.

Author Keywords

E-book, reading, multiple display devices, embodied interfaces, motion sensing

ACM Classification Keywords

H5.2. Information interfaces and presentation (e.g., HCI): User Interfaces.

INTRODUCTION

Although electronic versions of books, papers, and periodicals are becoming ubiquitous, people commonly print these documents in order to read, annotate, and interact with them. While laptop and tablet screens are growing in size and resolution, they, unlike paper, consume significant amounts of power, can be difficult to read in sunlight, and their size and weight make them difficult to hold in comfortable orientations for reading and annotation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2008, April 5–10, 2008, Florence, Italy.

Copyright 2008 ACM 978-1-60558-011-1/08/04...\$5.00



Figure 1 - Our prototype e-book reader with faces in the attached, side-by-side configuration (top), and the detached configuration (bottom).

In response, various e-books, devices designed specifically for reading electronic documents, have been offered. The most recent of these devices, which include products like the Amazon Kindle [3], use E-Ink [11], one of several *bi-stable* display technologies that do not require power to maintain an image on the screen and can be read in bright sunlight. Since these devices require far less power than those with LCD displays, they can be made thinner and lighter, resulting in a more mobile platform for reading.

However, current e-books provide only a single screen and therefore lack some important affordances of physical reading materials. In particular, embodied navigation in the form of turning, flipping and folding pages is limited. Moreover, the reduced screen real estate of a single-display e-book reader makes it difficult to configure the device to simultaneously show multiple regions of a document.

The low power consumption of bi-stable displays make it practical to build portable multi-screen e-books readers. The availability of dual-display readers will be a significant development since ethnographic studies have consis-

tently reported that a large proportion of reading activities involve the use of multiple display surfaces [1, 30]. Although companies like Jinke and iRiver have announced dual-display products [6, 18], details about these devices are not available and little is known about how to leverage dual-display devices to enhance the reading experience.

In this paper, we explore the design of interfaces that support reading activities with a dual-display e-book reader. To do so, we built a prototype dual-display reading device with integrated motion sensors (Figure 1) as a platform to explore two unique aspects of dual-display systems:

Navigation using an embodied interface. Reading paper documents requires physically turning, flipping, and sometimes folding back pages. A bound, dual-display system offers a similar set of physical interactions and we extend strategies for creating embodied interfaces in prior systems, like tilting [12] and synchronous gestures [14] to a system with two displays. Our device allows readers to turn pages by fanning one face toward the other and, in the back-to-back configuration, by flipping the device over.

Flexible display configurations. A significant drawback of current e-book readers is that they offer very limited screen space. While it is possible to construct devices with bigger displays, the increased size would generally make them more cumbersome to use and carry around. Dual-display e-book readers offer a compromise between increased screen real estate and a flexible and convenient form factor.

On our device, when the two displays are attached, the increased screen space can be used to present a broader overview of the document. Users can also fold one face behind the other, which provides an alternative, more compact form factor. When the displays are detached, multiple documents can be viewed simultaneously. In such a setting, users can arrange the displays in convenient and ergonomic positions in their workspace, a crucial feature when working with paper documents [26].

Using our prototype, we examine how these different features can be combined to offer a rich reading experience. Our evaluation comparing a single-display device to a dual-display device indicates that several advantages arise from the addition of an extra display. Users reported having an easier time re-reading and finding new content with two faces. Users also found the ability to view multiple documents simultaneously directly applicable to their daily reading activities. Areas that require improvement include reducing the clumsiness when operating the two-display system, and developing more compelling embodied interactions. We believe that our findings will help inform future e-book designers on how to provide a more enjoyable and efficient reading experience.

RELATED WORK

How People Read

The most direct motivation for the use of a two screen reading device comes from Adler, et al. [1] who report on the characteristics and requirements of work-related reading activities. They emphasize that close to half of all such activities involve multiple display surfaces and conclude that: multiple displays “need to be supported in digital reading devices if they are to satisfactorily replace current practice.” O’Hara and Sellen [26] compare reading on paper with reading online and similarly report that people commonly use multiple documents at the same time. Morris, et al. [24] extend O’Hara and Sellen’s study using more modern hardware and software, like dual-monitors and TabletPCs. One of their key conclusions was that reading systems should offer several highly configurable displays. These studies’ analysis of how people work with multiple documents in the workspace motivate our design allowing for both bound and detached modes of operation.

Marshall and Bly [21] offer an extensive analysis on the way people navigate magazines. They show that people use a wide variety of navigation techniques like skimming, flipping and glancing. They also highlight the importance of lightweight navigation techniques—those that require only a small amount of conscious thought to perform—for serendipitous discovery of the document content. These findings motivate our use of embodied interfaces.

Software and Hardware for Electronic Book Reading

Many portable devices such as laptops, PDAs and TabletPCs include electronic document reading software, but the most comprehensive document reading system is probably XLibris[28]. XLibris was developed with the goal of preserving the most important affordances of paper while introducing the benefits of electronic reading. For example, XLibris provides tools for navigation, annotation and summarization. These tools support *active reading* [2] in which the user mixes reading and critical thinking. XLibris was then used by law students [23] and a reading group [22] to develop several conclusions about digital reading. Among these include the need to support mobility, and the reading strategies that people perform on paper.

Early attempts to build specialized reading devices were unsuccessful, in large part due to poor display quality and short battery life. Schilit, et al. [29] presents an overview of many of these devices. However, these technological limitations are disappearing as new display technologies, such as reflective bi-stable displays [11, 25], become available. In a bi-stable display, the pixels can be in one of two stable states. Consequently, the screen can retain a given image without power, only requiring power when the display refreshes. The availability of bi-stable displays has spurred the introduction of products like the Sony Reader and Amazon Kindle [3, 31], which offer impressive display quality and long battery life. These systems all use one




Figure 2 - Estimated power usage as a function of framerate for Nemoptic displays (adapted from [25]).

display, however. The idea of using dual displays can be traced back to the Everybook [29], an early e-book prototype containing two LCD screens set side by side in portrait mode. However, our work is the first to contribute interaction techniques for dual-displays systems and to evaluate how these systems may enhance the reading experience.

Embodied Interfaces

Our work draws heavily on work in embodied interfaces for document navigation. Harrison, et al. [12] augment a computer with a set of sensors to provide a rich set of navigation tools. Users can turn pages with a flicking gesture, or navigate a list by tilting the device up and down. Tilt sensors and accelerometers have also been used for scrolling [5, 12], text entry [27, 34], collision detection to establish a link between devices [14] and activity detection [15]. Our work explores how these interface ideas can be applied in the context of a dual-display e-book reader.

DESIGNING A DUAL-DISPLAY ELECTRONIC READER

The main goal of our work is to explore how dual-display systems can improve the reading experience for digital content. Since reading occurs in diverse environments, our approach was to focus on a set of features that would span a wide variety of reading scenarios. Drawing directly from ethnographic work on reading, we establish three design goals for a dual-display e-book reader:

- **Improving local navigation.** When reading a document, users often need to consult material on or around the current page, often repeatedly. Local navigation has been observed in technical reading, as well as in casual magazine reading [21, 22]. Dual-display e-book readers can facilitate these activities by increasing the reading area and providing lightweight, gesture based controls.
- **Improving global navigation.** Global navigation involves users obtaining an overview of a document and then moving to a location i that document. Global navigation frequently involves the visual search of a document, like when rapidly flipping through a reference book to get to a specific point. Dual-display e-book readers can use the increased display area to implement

improved global navigation techniques to make performing these tasks easier and more efficient.

- **Improving multi-document navigation.** Many reading activities require simultaneous access to more than one piece of paper. For example, note-taking involves reading from one sheet while writing on another and perhaps referring to a third. In such situations, it is important that the documents can be positioned freely with relation to one another [26]. Dual-display systems with detachable faces can improve multi-document navigation by allowing users to arrange and interact with each display separately.

Since bi-stable display technology seems to be a good fit for reading, our designs are sensitive to one of the key tradeoffs associated with the technology: while they do not consume energy to maintain an image on the display, bi-stable displays require significantly more energy than a standard LCD display to refresh an image. Thus, highly interactive techniques are very costly, power-wise (Figure 2). Current bi-stable displays also have relatively slow refresh rates, requiring 0.5s to 1s to update. We believe this is a temporary artifact of the current technology and that interactive refresh rates will be possible in the near future [4]. Thus, our designs assume interactive refresh rates.

We used an iterative design approach to design our dual-display reader. We began with paper-based prototypes to evaluate control placement. We then built an instrumented two-face prop to evaluate the feasibility of different embodied interactions. Finally, we built two high fidelity prototypes. As low power consumption was not a goal of these prototypes, we used LCD displays because they are inexpensive, readily available, and have fast refresh rates. A practical final system will require using bi-stable displays.

Improving local navigation

Devices like the Sony Reader, Amazon Kindle, and XLibris all support turning a single page, with buttons and




Figure 3 - In the back-to-back configuration, flipping the device over automatically turns the page.

a pressure sensor, respectively. However, Marshall and Bly [21] describe several other types of local navigation that occur with paper documents.

Marshall and Bly report that magazine readers dynamically expand or contract the viewing area of their document, by folding and unfolding pages behind each other. Our system allows a similar interaction by automatically transitioning between displaying one or two faces worth of information based on the relative configuration of the two faces.

When reading a two-sided document, like a magazine, with the pages folded back, one can quickly change pages by flipping the document over. Our dual-display reader provides a similar interaction and changes the page as soon as the device is flipped (Figure 3). However, unlike paper documents, where two consecutive flips bring back the original page, our flipping gesture continuously advances or rewinds the pages depending on the orientation of the flip (flipping clockwise advances while flipping counter-clockwise rewinds). To facilitate the common task of flipping back to reread a previous page, the device contains a hysteresis mechanism that preserves last viewed page on the other face. Keeping the last viewed page ready allows users to rapidly switch back and forth between the most recently viewed pages in a lightweight fashion.

For documents in the side-by-side configuration, a user can quickly shift a page from one hand to the other, like when thumbing through a magazine. Our device provides an analog in the form of a fanning gesture (Figure 4) that removes the need to search for small hardware controls to change pages. Not having to acquire a control is useful as people




Figure 4 - The right face is brought towards the left face (“fanning”) to move forward in the document.




Figure 5 - Each slate has two trackballs, and below them, a menu button to trigger the command system.

shift the position of their hands while reading, either for comfort, or to perform other activities [21].

To complement the embodied navigation interface, our design also includes more traditional navigation widgets in the form of one small clickable trackball on each side of the display (Figure 5). These widgets are essential in the detached mode but are also used in the attached mode when flipping and fanning are not convenient, such as when the device is used on a desk.

Since our device can be operated in a variety of configurations, laying out the controls in a way that is comfortable to use and logical in all settings, had to be considered. We explored many possible locations for the trackball, including below the screen and on the side of the device, and we found that positioning the trackball on the side of the screen provided the most comfortable grip. We also adopted a symmetrical layout to complement flipping in the back-to-back configuration. Since flipping the device causes the faces to alternate being on top, the controls are laid out identically on the two faces for the flipping to be transparent. Similarly, giving each face an identical layout keeps the interface uniform when the faces are detached.

Even though the trackball detects two-dimensions of motion, we have observed that when the device is held with a single hand, the range of motion of the thumb is limited primarily to left and right movement [19]. Consequently, our interface is designed so that all controls and menus can be operated using the horizontal motion of the trackball. The trackball allows users to smoothly move forward and backward through multiple pages, but we also allow users to push down on the trackball to move one page at a time. Clicking on the trackball on the right advances the document clicking on the trackball on the left moves back. Other navigation controls are possible, such as touch-strips [7, 12], but trackballs offer a good compromise between ease of use and ease of implementation.




Figure 6 - Using SFT for rapid document navigation. Inset shows the red area in the photo in more detail.

Early in the design process, we also considered using the tilt of each slate to control scrolling [5, 12]. However we found users had difficulty holding the device still at a particular orientation. Since the problem would likely be exacerbated in the more unsteady environments outside of the lab, we omitted the feature.

Improving global navigation

The ability to quickly browse for and access a specific section of a document is a key feature of paper documents. Although table of contents and hyperlinks in electronic devices address some of the issues related to global navigation [8], visual search is essential. In electronic devices, the predominant method for performing visual navigation is scrolling. Unfortunately, scrolling interfaces such as the standard scrollbar, Speed-Dependant Automatic Zooming [17], Rapid Serial Visual Presentation [10] and Flipper [32] all require high display refresh rates, an inefficient mode of operation for bi-stable displays (Figure 2).

To prevent global navigation from negating the power consumption advantages of bi-stable displays, we rely on the fact that dual-display configurations offer increased screen space. The extra space makes it possible to effectively implement Cockburn et al.'s Space Filling Thumbnails (SFT) [9] as an alternative to scrolling, or the multi-page thumbnail view in XLibris[28]. In SFT, page thumbnails of the entire document are tiled across both displays (Figure 6). This setting allows users to see an overview of the entire document at once, and takes advantage of users' spatial memory. SFT has the added benefit of requiring very few display refreshes, making it ideal for use with bi-stable displays. In our current prototype, about 70 pages of graphic-heavy content can be comfortably viewed when using both screens. As in Cockburn's system, we provide an enlarged view of the highlighted thumbnail when the sizes of the thumbnails are below 154 x 205 pixels.

Bookmarking is another important feature for helping users get to areas of a document they have accessed in the past




Figure 7 - By having two independent displays, users can view two pages with different aspect ratios without having to constantly rotate the device.

[8, 22, 23, 28]. Here too, the use of a dual-display configuration can significantly improve the navigation experience. Upon activation of the bookmark menu the bookmark/comparison tool uses one face to display a bookmarked page alongside the normal view of a page in the document. The use of two faces allows the bookmarked page to be viewed simultaneously with an arbitrary page in the document. Viewing two pages side-by-side aids comparisons, which are crucial to sensemaking [8].

Improving inter-document navigation

Certain types of reading are characterized by a significant amount of inter-document operations, in which multiple documents are referenced and compared. While such a task can be performed on a single display using multiple windows, a single screen means that inter-document operations must be performed serially, which is cognitively demanding, slow, and frustrates users [26]. Two displays bound together can mitigate the situation slightly by increasing the screen real estate available. Yet, a rigidly bound dual-display device is restrictive because it precludes custom arrangements that are best suited for the task at hand.

To support inter-document interactions, the slates of our reader can operate independently (Figure 7). Upon separating the two slates, each slate becomes independent and can show two different documents, different sections of the same document, or a document and associated notes¹. Since each face provides independent control, they also enable bimanual interaction and the parallelism commonly observed in paper practice [26].

When the slates are reattached, the reattachment process can either use a default side (i.e. the left slate) as the source of the common document, or prompt the user to choose a

¹ Currently, the user is able to add ink onto the page. There is not yet a dedicated utility for aggregating and managing notes.

side. The remaining slate is then automatically synchronized to return to a coherent side-by-side display.

Menu and Command Invocation

The different functions of the device, such as managing the bookmarks, or calling up the SFT view, are accessed using a menu system controlled by the trackballs. Pressing either of the command buttons situated below the trackballs (Figure 5) launches a linear menu of options. The aforementioned restrictions on thumb movement precluded the use of directional menus, like marking menu [20]. Once in the menu, the trackballs are used to move through the selections. Clicking the trackball enters submenus and invokes commands. Pressing the command button a second time deactivates the menu.

High Fidelity implementations

Although it would have been possible to simulate a subset of the behavior of our device using a system like PaperWindows [16], we decided that a physical device would allow evaluations in a wider variety of settings, in and out of the lab. Another option was to build our prototype around an off-the-shelf computer or PDA. We bypassed these options because PDAs generally have small screens and even very light TabletPCs are still too heavy to bind together. Furthermore, we were concerned that it might be difficult to modify the physical configuration of an off-the-shelf product to support a custom control layout and to incorporate additional sensors.

Our solution was to build high fidelity prototypes based around a Gumstix single board computer. Each slate includes a Gumstix Basix 400xm-bt (400MHz XScale PXA255 processor with an integrated Bluetooth module) that drives a color LCD, and is connected to a PIC microcontroller that manages data from the trackball, command button, face attachment switch, and motion sensors.

Sensor data and device state are shared between the faces using a Bluetooth Personal Area Network, via a context server [14]. The context server currently runs on a separate

laptop computer but can also run on one of the faces of the device. The context server tracks the state of the reading device, synchronizes sensor readings and button states, and performs gesture detection. The context server then determines what page should appear on each face and sends that information to the slates.

First prototype

To simplify our design, our first high fidelity prototype² did not include advanced power management and was powered by 4 AAA NiMH batteries. As a result, each slate was heavy (520g) and thick (18mm), and the system could only run for about 30 minutes per battery charge. Nevertheless, this prototype was able to convey our interface concept to users. We implemented the design described above using accelerometers to sense slate movement, but did not include a touchscreen to capture ink. Each of the slates ran the Windows CE 5.0 operating system, with the application software written in C#.

Using this device, we conducted a preliminary feedback session. The evaluation consisted of a semi-directed interview format lasting 25 minutes. We alternated between demonstrating the navigation techniques and allowing the participants to try them. After all the techniques were demonstrated, the users were asked to explore a sample document using any combination of techniques.

Users reported the fanning gesture to be natural, with one saying that it “feels just like a book.” For global navigation all users agreed that the SFT were useful for global navigation. The primary complaints were about the size and weight of the device, which made it difficult to perform the embodied flipping gesture. Another issue was that gestures did not work reliably, because the accelerometers did not produce stable data in vertical orientations of the faces. The short battery life also precluded the users from fully exploring the device. Based on the feedback of the first three users in our evaluation, it was clear that various issues needed to be addressed before we could capture more meaningful comments from additional users.

Second prototype

We developed a second high fidelity prototype (Figure 8) that addressed the issues raised by our first set of users. This prototype uses Li-Polymer batteries, and an LED backlit display that consumes less power than the one in our first prototype. Consequently, the second prototype is thinner (12mm) and lighter (400g), and offers more than 2.5 hours of run time. The new prototype also uses gyroscopic motion sensors, which improves the reliability of the embodied navigation controls. Since users thought that the bookmarking feature was too heavyweight in the first device, we did not include it in the second prototype.




Figure 8 - An internal view of the reader. The internal electronics are on the right. The bottom half of is the main board and the top half is the sensor board. The system is powered by two Li-Polymer battery packs.

² This device was described in the UIST 2007 extended abstracts.

For the second prototype, we switched the operating system on the slates to Linux and rewrote the software in C++ to use the Qtopia embedded development library. We also rewrote the context server in C++ to allow it to run on any POSIX compliant operating system.

EVALUATION

The main focus of our evaluation was to compare our new dual-face reader with a single-faced reader in a typical magazine reading task.

The single face reader was implemented as one face of the dual-display reader in detached mode, without any embodied interface. A direct comparison of our dual-display device to a commercial implementation such as the Sony Reader [31] would have been difficult, as those devices have extremely slow screen refresh rates and provide a different navigation interface. We believe our approach significantly reduces confounding factors in the results.

We used a within-subject design for this experiment. The format of the experiment consisted of our demonstrating the controls on either the single or dual-face device, then having the users try out the device, and finally having the users read uninterrupted for 10 minutes using the device. For the uninterrupted reading, participants were instructed to find and read the articles that they found most interesting. The only other instruction given to participants was to try out all the navigation features at the beginning, but to use whatever navigation techniques they felt most comfortable for the remainder of the reading time. Following the uninterrupted reading session, we conducted an in-depth semi-directed interview structured around a standardized set of 41 questions covering diverse areas as comparing the device they used to reading on paper and reading on a computer, control performance, and size and weight of the device. The procedure was repeated for the other device. We fully balanced the content users viewed on each device as well as the order the devices were presented across the participants. After the participants had used both devices, we asked them to compare the single-display and dual-display device (specifically: "Please compare the pros and cons of the single-display and dual-display reader.").

Lastly, at the end of the experiment, we demonstrated the face separation capability for displaying two different documents (or two different sections of the same document), allowed users to try the feature, and asked them whether such a feature was applicable and useful to the reading activities they personally perform.

The content for the experiment was the aggregate of articles in the "U.S." and "Technology" sections of the September 9, 2007 and September 14, 2007 editions of the New York Times. The articles were formatted for the resolution of the screens on the devices using the New York Times' digital reader, which dynamically lays out content for arbitrary screen sizes. Each set of articles contained 92

pages, and individual articles ranged in length from one to five pages. Users read one set on the single screen device and the other on the dual-screen device.

We recruited eight participants from our University. One participant encountered a hardware malfunction during her experiment so we discarded her data. The results presented below are for the seven remaining participants (P1-P7). Participants received \$20 for 1.5 hour sessions.

RESULTS

One Display vs. Two Displays

Users were all able to report some differences between reading with a single face and reading with two faces. P1 and P4 noted that it was much easier to re-read content with two displays. Examples include people's names and geographical locations that were first presented in a previous page. P2 said that two displays were helpful in helping her gain a better idea about how long an article was. Also, for P2 and P7, the second face helped them find new articles to read. P2 remarked that, "sometimes I would see something interesting on the second face and then I would go and read that." Users also cited increased fluidity (P5) and efficiency (P1, P6, P7) when using the dual-display device. One unexpected advantage that P4 identified was that presenting the content on two faces forced her to periodically change the angle she looked at the device, alleviating neck-strain she experienced when reading from a single face, or when reading on the computer.

Many of the advantages attributed to two displays mirror the examples of lightweight navigation that Marshall and Bly [21] describe. Operations like flipping back to re-read content, gauging the extent of an article, and serendipitous discovery of new content all appear to be supported by the inclusion of a second display.

The disadvantages of two faces mainly centered on the lack of flexibility in positioning the device. Five users (P3, P4, P5, P6, P7) mentioned that two faces limited the freedom they had in holding the device, compared to a single face. P4 and P7 specifically said that the arrangement felt flimsy, so we believe a different hinge design that provides more stability could improve the situation.

Four users (P2, P3, P4, P5) believed that the single face configuration was simpler and was less confusing to use. Users' main source of confusion with two displays came from being unclear about where the next page would appear, when pages changed at the same time, as there was no constant point of reference. However, all thought displaying page numbers would greatly lessen the confusion.

When asking about the face separation feature, all of the participants were able to give specific personal examples where having faces that operate independently would be useful. Of the variety of examples provided, three users (P1, P4, P5) mentioned the applicability to having to work

with multiple research papers at once, such as to compare two articles in detail. Four users (P2, P3, P4, P6) mentioned it would be useful to show different sections of textbooks. P4 also liked the separable faces because, as she explains, “two faces is better for serious reading,” but for casual reading on the go, she preferred the single-face device. With separable faces, she said she would keep the second face at home or at the office for the serious reading, and only carry around the single face. The ease with which our users could find examples in their daily reading that could take advantage of having two independent displays strongly parallels Adler et al.’s finding [1] that more than half of the tasks they observed people performing required operating on two or more separate display surfaces.

Embodied Interactions

All of the users tried the fanning and flipping gestures at the beginning of the reading sessions but none continued using them for the rest of the free-reading task.

Four users (P1, P2, P4, P7) reported that the fanning gesture worked reliably. Two users (P2, P7) said the fanning was a “cool feature” and three users (P1, P4, P5) mentioned that it made the device feel like reading a printed book or magazine. However, in the end, all the users agreed that they thought that the effort required to fan a face was greater than that of the standard controls, especially when traversing multiple pages, but did not provide any marked advantages over the standard controls.

For the flipping gesture, even with our slimmed down design, all but one (P7) of the users complained that the thickness of the faces and having to hold the faces together made flipping in the back-to-back configuration clumsy. “It was awkward, like juggling,” said P5. Again, an improved hinge that better retains the state of the faces may help. Our observations also indicated that at least three users accidentally brushed the standard controls while performing the gesture, causing undesired shifts in location. P4 commented that while she often read paper magazines with the faces held back-to-back, the primary reason was because the paper was flimsy. On our device, with rigid faces, having the faces back-to-back felt superfluous. P7 commented that she hesitated when flipping because she had to consciously think which direction to rotate the device.

Unlike fanning, some users believed flipping to have some untapped potential. Two participants said that they would be more receptive to the flipping if the device was lighter. P5 said that flipping was advantageous because it gave him a better sense of moving through the document than a button. P3 thought the embodied interface would serve to relieve the monotony of pressing a button repeatedly when reading long documents.

Thumbnail-based Navigation

There was a wide range of responses to the SFT navigation system. While nearly all the users agreed that the thumb-

nails should be bigger, they disagreed about how useable the system was in its current form. P1 and P7 thought that both the thumbnails and enlargement were unintelligible, but the remaining users were at least able to make out which pages contained headlines. Most users used the headlines in the thumbnails to find where articles began, at which point they would jump to that page. Another use for the thumbnails that we observed was to determine where an article ended, in order to skip over boring articles.

One interesting observation was that 5 out of the 7 users did not realize that the thumbnails were larger in the two-display configuration, the result of having a larger display area. When asked about the thumbnails on the two-display device, they said they thought there were just more pages. One user even thought the smaller thumbnails on the single-face device were easier to see and use.

Size and Weight

Three users (P2, P3, P6) believed that the two-display assembly was too heavy for comfort. On the other hand, no users felt that the single-display configuration was too heavy. In both the single and dual-display devices, all users applauded the flexibility in positioning the device compared to reading on a laptop. Five users (P1, P2, P3, P5, P7) immediately brought the device close to their body and leaned back in the chair to read. All of these users had read from laptops and added that they could not comfortably do the same thing with the laptop. Several users made comments about how they felt they could take the device to places where they normally would not bring their computer. For example, P5 remarked, “I don’t have to sit at my desk, if I want to lay down, I can lay down.” P6 incorrectly assumed he had to keep the device on the table, and when informed at the end of the session that he could pick it up, mentioned that he would prefer to pick it up to optimize his reading angle. P4 preferred to have the device on the table to read, but, even then, she said that both the one and two-face devices were superior to reading from her computer because the device laid flat on the table, allowing her to assume her preferred reading position of looking straight down at the table. All thought that a slight increase in display size would not hurt, but were wary of making each display the size of a letter-sized (8.5” x 11”) paper, citing concerns about portability.

DISCUSSION AND FUTURE WORK

The data we gathered during our evaluation, allows us to make several recommendations to designers of next-generation electronic document readers.

Number of Displays

Our results suggest that providing two displays in an electronic reader would probably improve the electronic reading experience, given the benefits users cited in the areas of local navigation and applicability to multi-document interactions when using two displays.

Many of these improvements could also be realized through a single, large display capable of displaying several pages simultaneously. We believe, however, that having separate displays is preferable. While the total number of pixels available to display information is important, additional advantages are gained from how these pixels are arranged. Multiple separate displays provide more flexibility and adaptability to different reading situations. For example, users raised concerns about the portability of devices with very large displays. Also, to read multiple documents of different aspect ratios simultaneously, a single-display device would need to be considerably larger than a dual-display device. Furthermore, a single-display device does not allow users lay out the electronic document alongside paper documents or books in a custom manner, such as grouping relevant documents together. However, further studies comparing a single large display with multiple small displays are needed to test our hypothesis.

With bi-stable displays, the power consumed is a function of how many different pages one views, *not* the number of displays. As a result, it is possible to consider devices with more than two displays. Several directions for future work promise to further improve the interaction with multiple displays. One unexplored question for multiple-display devices is the point of diminishing returns: where adding extra displays imparts no additional interaction advantage. The interface users will use to manage and operate these multiple displays will play a large part in determining that point. Thus, techniques for specifying how several displays work together to present content will need to be developed.

Also, thinner and lighter devices will be increasingly important as the number of devices increases. Reducing the size and weight of the devices can be achieved by further reducing power consumption. One way to achieve an additional reduction is through energy-efficient interfaces [13, 33]. For example, an interface might reduce power consumption by using more optimal refresh strategies that work to maximize the efficiency of the display technology.

Embodied Interfaces

We initially believed that embodied interfaces would provide a more natural, and hence less distracting, means of interacting with the device. We also believed that these embodied interactions would better facilitate the lightweight navigation that occurs on paper documents. However, our results suggest that the physical differences between an electronic device and paper documents are significant enough that familiar interactions with paper require much more effort to perform on the electronic device.

Embodied interactions should not be discounted entirely, however. One possible issue is that our interactions borrowed too literally from paper interactions, which ultimately did not translate well to our devices. Most of our users agreed that if the device were lighter, the interactions might be more palatable. But since electronic reading de-

vices will probably never provide the same set of affordances as paper, it will be wise to examine embodied interactions that go beyond existing reading conventions. Also, our experiment did not include a condition where the conventional controls were not readily accessible, for example, when operating the device with a single hand. Thus, under different reading scenarios and with other types of interactions, an embodied interface may prove to be useful.

Need for a specialized reading device

One valid question is whether dedicated electronic readers are necessary, given the availability of more general devices like laptop computers, which are also portable. Our users' feedback about the greater comfort and flexibility of our device argues in favor of a specialized device. Specialized reading devices may not provide all the capabilities found on a computer, but in return, minimize size and weight, which positively affects the reading experience.

To maximize the utility of a specialized reading device, methods to seamlessly integrate with the existing electronic and paper tools people use will be essential. For example, transferring a document from a computer to be viewed on the device will need to be as simple as clicking on the "print" button is to output to paper. Also, past experiences with digital documents suggests that paper will continue to occupy a role in our reading activities. Therefore, it will be important to investigate ways to perform mixed-media interactions between the reading device and paper.

CONCLUSION

Thanks to advances in display technology, electronic reading devices with two displays are possible. In this paper, we explored various affordances of a two-display electronic reader as they relate to navigation tasks. To accomplish our exploration, we constructed a series of prototype dual-display electronic readers, demonstrating key features such as an embodied navigation interface, and flexible reading configurations stemming from the ability to detach the two faces. We conducted an evaluation of our system to understand what the advantages and disadvantages of a dual-display system were compared to a single-display system. User feedback from our study indicates benefits from providing an additional display. The advantages include better support of lightweight navigation, which made finding content to read, as well as re-reading easier. The embodied interactions we implemented for the dual-display device generally were not perceived to aid reading, but the feedback we received suggests several areas where we can make improvements. We believe that our results show that further developing the design of dual-display devices and devices with more than two displays will be essential for improving the electronic reading experience.

ACKNOWLEDGEMENTS

This project was funded through NSF Grants #0447730 and #0414699. Agrawala is partially supported by the AI-

fred P. Sloan Foundation and the Okawa Foundation. We would like to thank Ben Bederson, Bobby Bhattacharjee and Bill Pugh for hardware support, David Vescovi for his assistance in getting the software for our prototypes up and running and Farzam Yazdanseta for his contributions to the initial prototypes. We would also like to thank the New York Times for the content used in our evaluation.

REFERENCES

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. In *Proc. of CHI 1998*, ACM Press (1998), 241-248.
2. Adler, M. J. and Doren, C. V. *How to Read a Book*. Touchstone, 1972.
3. Kindle: Amazon's New Wireless Reading Device. <http://www.kindle.com>
4. Angelé, J. and Emeraud, T. BiNem electronic paper. *Gekkan Display - Techno-Times Japan October*, (2006),
5. Bartlett, J. F. Rock 'n' Scroll Is Here to Stay. *IEEE Comput. Graph. Appl.* 20, 3 (2000), 40-45.
6. iriver's E-BOOK reader. <http://www.engadget.com/2007/01/08/irivers-e-book-reader/>
7. Buxton, W. Two-handed document navigation. *XEROX Disclosure Journal* March/April (1994), 103-108.
8. Card, S. K., Hong, L., Mackinlay, J. D. and Chi, E. H. 3Book: A 3D Electronic Smart Book. In *Proc. of AVI 2004*, (2004), 303-307.
9. Cockburn, A., Gutwin, C. and Alexander, J. Faster document navigation with space-filling thumbnails. In *Proc. of CHI 2006*, (2006), 1-10.
10. De Bruijn, O. and Spence, R. Rapid serial visual presentation: a space-time trade-off in information presentation. In *Proc. of AVI 2000*, ACM Press (2000), 189-192.
11. E-Ink. <http://www.eink.com>
12. Harrison, B. L., Fishkin, K. P., Gujar, A., Mochon, C. and Want, R. Squeeze me, hold me, tilt me! An exploration of manipulative user interfaces. In *Proc. of CHI 1998*, ACM Press (1998), 17-24.
13. Harter, T., Vroegindewij, S., Geelhoed, E., Manahan, M. and Ranganathan, P. Energy-aware user interfaces: an evaluation of user acceptance. In *Proc. of CHI 2004*, ACM Press (2004), 199-206.
14. Hinckley, K. Synchronous gestures for multiple persons and computers. In *Proc. of UIST 2003*, ACM Press (2003),
15. Hinckley, K., Pierce, J., Horvitz, E. and Sinclair, M. Foreground and background interaction with sensor-enhanced mobile devices. *ACM Trans. Comput.-Hum. Interact.* 12, 1 (2005), 31-52.
16. Holman, D., Vertegaal, R., Altosaar, M., Troje, N. and Johns, D. Paper windows: interaction techniques for digital paper. In *Proc. of CHI 2005*, (2005), 591-599.
17. Igarashi, T. and Hinckley, K. Speed-dependent automatic zooming for browsing large documents. In *Proc. of UIST 2000*, (2000), 139-148.
18. Hanlin eReader V2d. http://www.jinke.com.cn/compagesql/English/embedpro/Hanlin_eReader_V2d_pamphlet.pdf
19. Karlson, A. K., Bederson, B. B. and Contreras-Vidal, J. L. *Studies in One-Handed Mobile Design: Habit, Desire and Agility*. HCIL-2006-02, Computer Science Dept, University of Maryland, 2006.
20. Kurtenbach, G. *The Design and Evaluation of Marking Menus*. Ph.D Thesis, University of Toronto, 1993.
21. Marshall, C. C. and Bly, S. Turning the page on navigation. In *Proc. of JCDL 2005*, ACM Press (2005), 225-234.
22. Marshall, C. C., Golovchinsky, G. and Schilit, B. N. Introducing a Digital Library Reading Appliance into a Reading Group. In *Proc. of ACM Dig. Lib. 1999*, (1999),
23. Marshall, C. C., Price, M. N., Golovchinsky, G. and Schilit, B. N. Designing e-books for legal research. In *Proc. of JCDL 2001*, ACM Press (2001), 41-48.
24. 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 *Proc. of TABLETOP 2007*, (2007), 79-86.
25. Nemoptic. <http://www.nemoptic.com/>
26. O'Hara, K. and Sellen, A. A comparison of reading paper and on-line documents. In *Proc. of CHI 1997*, ACM Press (1997), 335-342.
27. Partridge, K., Chatterjee, S., Sazawal, V., Borriello, G. and Want, R. TiltType: accelerometer-supported text entry for very small devices. In *Proc. of UIST 2002*, ACM Press (2002), 201-204.
28. Schilit, B. N., Golovchinsky, G. and Price, M. N. Beyond paper: supporting active reading with free form digital ink annotations. In *Proc. of CHI 1998*, (1998), 249-256.
29. Schilit, B. N., Price, M. N., Golovchinsky, G., Tanaka, K. and Marshall, C. C. As We May Read: The Reading Appliance Revolution. *Computer* 32, 1 (1999), 65-73.
30. Sellen, A. J. and Harper, R. H. R. *The Myth of the Paperless Office*. MIT press, 2001.
31. *Sony Operation Guide, PRS-500, portable Reader System*. Sony Inc, 2006.
32. Sun, L. and Guimbretiere, F. Flipper: a new method of digital document navigation. In *Proc. of CHI 2005*, (2005), 2001-2004.
33. Vallerio, K. S., Zhong, L. and Jha, N. K. Energy-Efficient Graphical User Interface Design. *IEEE Transaction on Mobile Computing* 5, 7 (2006), 846-859.
34. Wigdor, D. and Balakrishnan, R. TiltText: using tilt for text input to mobile phones. In *Proc. of UIST 2003*, ACM Press (2003), 81-90.