

ABSTRACT

Title of Dissertation: **THE USE OF MULTIPLE SLATE DEVICES TO
SUPPORT ACTIVE READING ACTIVITIES**

Nicholas Yen-Cherng Chen, Doctor of Philosophy, 2012

Dissertation Directed By: Professor François Guimbretière
Department of Computer Science

Reading activities in the classroom and workplace occur predominantly on paper. Since existing electronic devices do not support these reading activities as well as paper, users have difficulty taking full advantage of the affordances of electronic documents.

This dissertation makes three main contributions toward supporting active reading electronically. The first contribution is a comprehensive set of active reading requirements, drawn from three decades of research into reading processes. These requirements explain why existing devices are inadequate for supporting active reading activities.

The second contribution is a multi-slate reading system that more completely supports the active reading requirements above. Researchers believe the suitability of paper for active reading is largely due to the fact it distributes content across different sheets of paper, which are capable of displaying information as well as capturing input. The multi-slate approach draws inspiration from the independent reading and writing surfaces that paper provides, to blend the beneficial features of e-book readers, tablets, PCs, and tabletop computers.

The development of the multi-slate system began with the Dual-Display E-book, which used two screens to provide richer navigation capabilities than a single-screen device.

Following the success of the Dual-Display E-book, the United Slates, a general-purpose

reading system consisting of an extensible number of slates, was created. The United Slates consisted of custom slate hardware, specialized interactions that enabled the slates to be used cooperatively, and a cloud-based infrastructure that robustly integrated the slates with users' existing computing devices and workflow.

The third contribution is a series of evaluations that characterized reading with multiple slates. A laboratory study with 12 participants compared the relative merits of paper and electronic reading surfaces. One month long in-situ deployments of the United Slates with graduate students in the humanities found the multi-slate configuration to be highly effective for reading. The United Slates system delivered desirable paper-like qualities that included enhanced reading engagement, ease of navigation, and peace-of-mind while also providing superior electronic functionality. The positive feedback suggests that the multi-slate configuration is a desirable method for supporting active reading activities.

THE USE OF MULTIPLE SLATE DEVICES TO SUPPORT ACTIVE READING
ACTIVITIES

By

Nicholas Yen-Cherng Chen

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2012

Advisory Committee:

Professor François Guimbretière, Chair
Professor Matthew Kirschenbaum
Dr. Abigail Sellen
Professor Ben Shneiderman
Professor Neil Spring

©Copyright by

Nicholas Chen

2012

Acknowledgements

The successful completion of this dissertation was only possible with the support of the many people who have been a part of my life. Through their encouragement, beneficial peer-pressure, and advice, my family, friends and colleagues helped me make it through this most challenging, but also most rewarding, period of my life.

First, I would like to thank my advisor, François Guimbretière, with whom I have worked since my first day as a graduate student. I am grateful for his continuous support as we embarked on an ambitious and unconventional program of research. His high standards and unwavering confidence in me ensured I did not shy away from the hard problems, and helped me produce research that I am intensely proud of.

I would also like to thank Abigail Sellen, whose research is the foundation on which mine rests. The expertise and prior experience she brought were indispensable in my work. The results of my research are stronger and more relevant from her contributions.

The rest of my committee deserves recognition as well. Ben Shneiderman's insights into what makes for good research (e.g., "two parents and three children") has been tremendously helpful in helping me frame my work. Neil Spring taught me important lessons about being more detail-oriented that I will not soon forget. Matthew Kirschenbaum was instrumental in connecting this research with its applications in the humanities. Vibha Sazawal was part of my proposal committee and helped me formulate the central thesis of my research and also encouraged me to consider the broader implications of my work.

The financial and equipment support of people and organizations that saw value in my work were critical. Bill Buxton and Microsoft Research provided the generous gift that made it possible to build the slate devices used in my research. A significant portion of the

work in this dissertation was completed while I was funded through a Google Ph.D. Fellowship in Human-Computer Interaction. Hod Lipson loaned us the Optitrack system used in Chapter 7. NSF Grants IIS-0447703, IIS-0749094, and IIS-0916217 provided the remaining resources for this research.

My research has given me the opportunity to collaborate with a number of incredible people. The origins of this dissertation can be traced back to when Maneesh Agrawala hosted me at UC Berkeley. Morgan Dixon and Cassandra Lewis helped me realize my initial hardware prototypes at UMD. John Barrus and Ricoh Innovations, Inc. supported me with a summer internship but also provided me with the tremendous talents of Ken Gudan, Sergey Chemishkian, Ed Schwartz, and Rich Kosoglow, from whom I learned a great deal.

I must also acknowledge the participants in my user studies and fieldwork. In particular, Kari Kraus and the students in her Spring 2010 ENGL428Y course at UMD, and my Ph.D. student participants at Cornell. The need for anonymity prevents me from recognizing these people by name. Thank you all for being so understanding about the fickleness and shortcomings of prototype hardware.

My studies have taken me through both the University of Maryland and Cornell University. Hyunyoung Song has been my officemate, friend, and confidant at both places. From her, I received a constant stream of encouragement, knowledge, entertainment, and good ideas throughout grad school. I am also fortunate to have met Elizabeth Murnane during my final year in grad school. Her amazing work ethic inspired me to stay focused during the last mile of my dissertation. My friends and labmates Chang Hu, Yit Phang Khoo, Chunyuan Liao, Alexander Quinn, and Sureyya Tarkan in the UMD CS Department, and Jaeyeon Kihm, Shenwei Liu, Victoria Schwanda, Madeline Smith, and Shaomei Wu in the Cornell IS Department provided all-important social outlets inside and

outside of the lab. Lastly, I had no idea my weekly dinners with David Blocher, Rocio Esmeralda Ruelas, Matthew Holden, and Joel Nishimura would evolve into the helpful grad school support group that they ended up becoming.

To my family: Mom, Dad, and Stephen, thank you for always being so supportive. I value your constant reminders that as central a part school and work are in my life, there are things that are even more important, like health and happiness. Most of all, I am grateful to my love, and my best friend in the world, Megan, who has stuck by me continuously from our first week of college until now. Her constant support and companionship have helped me get through the bad times and made the good times that much sweeter. I am so lucky I has a you.

Table of Contents

1	Introduction	1
1.1	Thesis Contributions	4
1.2	Solution Overview	5
1.2.1	Analysis of the requirements for work-related and academic reading	5
1.2.2	Design and Implementation of a Multi-Slate Reading System	6
1.2.3	Evaluation of multi-slate reading systems	8
1.3	Dissertation Roadmap	9
2	Related Work	12
2.1	Supporting Reading Activities	12
2.2	How People Read With Paper	14
2.2.1	Annotation and Note-Making	15
2.2.2	Micro-Mobility and Physicality of Materials	16
2.2.3	Rich Navigation Within Documents	17
2.2.4	Working With Multiple Documents	18
2.2.5	Document Organization and Management	19
2.2.6	Macro-Mobility and Portability	19
2.3	Providing “Paper-Like” Functionality	20
2.3.1	Augmenting Paper	21
2.4	Simulating Paper	23
2.4.1	Electronic Reading Devices	23
2.4.2	PCs	24
2.4.3	PDAs and Smartphones	24
2.4.4	Tablet Computers	25
2.4.5	E-book Readers	28
2.4.6	Large Displays and Surface Computing	30
2.5	Reading Requirements and Technologies At A Glance	31
2.6	Distributing Content Across Screens and Surfaces	33
2.7	Interactions Between Displays/Devices	34
2.8	Broader Implications of this Work	37
3	The Dual-Display E-book	40
3.1	Introduction	40
3.2	Design Of The System	43
3.2.1	Improving local navigation	44
3.2.2	Improving global navigation	47
3.2.3	Improving inter-document navigation	49
3.3	Iterative Device Design	50
3.3.1	First Prototype	50
3.3.2	Second Prototype	51
3.3.3	Prototype Implementation Details	52
3.4	Evaluation	53
3.5	Results	54
3.5.1	One Display vs. Two Displays	54
3.5.2	Embodied Interactions	56
3.5.3	Thumbnail-based Navigation	57
3.5.4	Size, Weight, and Mobility	58
3.6	Lessons from the Dual-Display E-book	58
3.6.1	Desirability of Multiple Screens	58
3.6.2	Embodied Interfaces	59

3.6.3	Dedicated Reading Device versus Computer.....	60
3.7	Summary.....	60
3.8	Epilogue	61
4	Multi-Slate Reading Platform.....	63
4.1	Introduction	63
4.2	Design Rationale and Specifications.....	65
4.3	Slate Hardware	67
4.3.1	Key Components.....	67
4.3.2	Main Electronics Architecture.....	69
4.3.3	I/O Control Strips.....	72
4.3.4	Enclosure Design	73
4.4	System Software.....	75
4.4.1	Operating System	75
4.4.2	Device Drivers	76
4.4.3	System Scripts.....	77
4.4.4	Application Development.....	78
4.5	Design Discussion and Assessment	79
4.6	Summary.....	80
5	The United Slates	82
5.1	Introduction	82
5.2	Technology Probe.....	84
5.2.1	Equipment and Interface for the Probe.....	84
5.2.2	Findings From the Probe	85
5.2.3	Implications from the Probe	87
5.3	The United Slates System	88
5.3.1	Mapping Information to Slates.....	88
5.3.2	Viewing Documents at Different Semantic Levels.....	90
5.3.3	Moving Information Around.....	91
5.3.4	Controlling Slates From The PC.....	98
5.3.5	Document Management With Stacks.....	100
5.3.6	Putting Everything Together	103
5.4	Implementation.....	105
5.4.1	System Architecture.....	105
5.4.2	Software	106
5.5	Preliminary Evaluation	107
5.5.1	Observation Procedure.....	108
5.5.2	Results.....	111
5.5.3	Evaluation Summary	115
5.6	Discussion.....	117
5.6.1	Multiple Devices Help Resolve Reading Tensions.....	117
5.6.2	Opportunities for Improvement	119
5.7	Summary.....	122
6	Cloud-based System Infrastructure	124
6.1	Introduction	124
6.2	Challenges of Real-World Deployment.....	126
6.2.1	Challenge 1: Availability of Resources	126
6.2.2	Challenge 2: Robustness.....	127
6.2.3	Challenge 3: Scalability	128
6.3	Added Benefits of a Cloud-based Infrastructure.....	129

6.3.1	Simplified System Architecture	130
6.3.2	Inherently Cross-platform	130
6.3.3	Receiving More for the Same Amount of Work	130
6.4	Implementation Details.....	131
6.4.1	Dropbox File Transfer and Synchronization	131
6.4.2	Routing Data Through the Cloud.....	132
6.4.3	Cloud-Based Document Processing.....	132
6.4.4	Optimizations	133
6.5	Summary.....	134
7	The Role of Physicality in Electronic Reading.....	136
7.1	Introduction	136
7.2	Comparing Tablets, Large Horizontal Screens, and Paper.....	139
7.3	Experiment Devices	141
7.3.1	Tablet Condition	142
7.3.2	Large Horizontal Screen (LHS) Condition	143
7.3.3	Paper Condition	146
7.4	Experiment Tasks.....	146
7.4.1	Linear Reading.....	147
7.4.2	Cross-document Comparison.....	147
7.4.3	Cross-device Reading	148
7.5	Experiment Methodology	149
7.5.1	Training.....	150
7.5.2	Starting Setup	151
7.5.3	Data Collection.....	152
7.5.4	Participants	152
7.5.5	Analysis	153
7.6	Results	153
7.6.1	Task 1: Linear Reading	154
7.6.2	Task 2: Cross-document Comparison	156
7.6.3	Task 3: Cross-device Reading	160
7.6.4	Summary of Layout Behaviors	164
7.6.5	Overall User Preferences	165
7.6.6	User Sub-populations	166
7.7	Discussion.....	170
7.7.1	General Efficacy of Multiple Reading/Writing Surfaces	170
7.7.2	Making Sense of the Differences Between Conditions	171
7.7.3	Functionality Cuts Both Ways.....	172
7.7.4	Suggested Improvements for the Tablet.....	172
7.7.5	Suggested Improvements for the LHS	173
7.7.6	Other Important Improvements.....	175
7.7.7	Tablets Occupy the “Middle Ground”	175
7.8	Summary.....	176
8	Real World Deployment of the United Slates System	177
8.1	Introduction	177
8.2	Study Goals and Setup.....	179
8.2.1	Hardware Configuration	179
8.2.2	Reading Interface	180
8.2.3	Managing Documents	182
8.2.4	Inter-slate Interactions.....	183
8.2.5	PC Connectivity.....	184

8.2.6	User Feedback and Journal	185
8.2.7	Web Interface	185
8.3	Methodology	186
8.3.1	Photo-Diary	186
8.3.2	Training Sessions	187
8.3.3	Extended Deployment	189
8.3.4	Analysis	190
8.4	User Profiles and Activities	191
8.5	Findings and Key Themes	198
8.5.1	Writing.....	198
8.5.2	Ubiquitous Access to Data.....	201
8.5.3	Extra Displays.....	205
8.5.4	Increased Micro-Mobility.....	209
8.5.5	Partitioning Information Spaces.....	212
8.5.6	Navigation Support	214
8.5.7	Inter-Slate Interactions	219
8.5.8	Connecting with the PC	222
8.6	Discussion.....	225
8.6.1	Comparing United Slates with Other Technologies	225
8.6.2	New Interactions Resulting from the United Slates	229
8.6.3	Improving the United Slates.....	230
8.7	Summary.....	231
9	The Big Picture.....	233
9.1	Supporting Reading Using Multiple Slates.....	233
9.1.1	Definite Advantages of Multiple Slates.....	233
9.1.2	Other Beneficial Features	238
9.1.3	Limitations of the System	240
9.1.4	United Slates In Context	243
9.2	Comparing to Augmented Paper Approaches.....	246
9.3	Recommended Approach for Building and Testing HCI Systems	248
9.4	Summary.....	249
10	Conclusions and Future Work.....	251
10.1	Contributions	252
10.2	Future Work.....	255
	Appendix A: Dual-Display E-Book Interview Questions	258
	Appendix B: United Slates Preliminary Evaluation Materials	262
	Appendix C: Physicality of Reading Surfaces Questionnaire	276
	Appendix D: Deployment Final Interview Questions	282
	Bibliography	285

List of Tables

Table 1 Comprehensive requirements of reading drawn from the literature.	32
Table 2 Degree to which existing devices support reading requirements.	33
Table 3 Technical specifications for the slate hardware platform.....	69
Table 4 Procedures for the preliminary evaluation of the United Slates system.....	110
Table 5 United Slates functionality and feedback placed in context with existing technologies.	226
Table 6 Updated comparison of how well different electronic reading devices support active reading activities.....	244

List of Figures

Figure 1 Contemporary e-book reading devices. From left to right: The Sony Reader PRS-650, Amazon Kindle, Barnes and Noble Nook Simple Touch.	28
Figure 2 The Dual-Display E-book in the connected configuration.	41
Figure 3 Flipping gesture on the Dual-Display E-book advances the page.	44
Figure 4 The fanning gesture being performed on the Dual-Display E-book.	45
Figure 5 Close up view of one display of the Dual-Display E-book. Trackballs and menu buttons are to the right of the thumb.	46
Figure 6 Space filling thumbnail navigation on the dual-screen configuration.	48
Figure 7 Two sides of the Dual-display operating independently and showing two different documents.	50
Figure 8 First (left) and second (right) prototypes of the Dual-Display E-book.	51
Figure 9 Block diagram of electronic components.	70
Figure 10 DC/DC conversion stages in the slate hardware.	71
Figure 11 Flexible PCB for touch sensing.	73
Figure 12 Internal view of the slate device. Note the two flex input boards along the sides and Gumstix carrier board at the top. The orange squares along the bottom edge are the radio antennas.	74
Figure 13 Front and back views of the device enclosure.	75
Figure 14 System software stack on the slate device platform.	76
Figure 15 The four levels at which users can interact with documents. Our multi-slate interactions enable combinations of these views to support reading activities. A row of view selection buttons (top) is placed at the top of the screen and is used to switch between levels.	90
Figure 16 Distributed thumbnail navigation using the Conduit interaction technique. Left: Changing the page on the center slate via a thumbnail on the right slate by depressing the command button (1), and then tapping on the thumbnail (2). Right: Result of the operation; note that the user can be thought of as a conduit through which the data flows (arrow).	93
Figure 17 Performing a Stitching operation using stitchable hyperlink elements. Users have a choice of selecting from full menu of commands (left, top) or using the	

- shortcut of a long cross (left, bottom). While the system waits for a destination endpoint for the endpoint for the command, the indicator square next to the selected command is filled. (Figure is taken from a screenshot from our device simulator and uses desktop cursors for clarity).....94
- Figure 18 The slate remote control interaction. (1) Press and hold command button on local device. (2) Click command button on remote device (enhancement), (3) Remote screen is frozen and transferred (note the message informing the user that the session is active), (4, 5) Using local device as a proxy, (6) Double click command button to end remote control session.....96
- Figure 19 Left: Single screen margin callout window. Right: Proximity links used to expand the view of a page to include margin notes.....98
- Figure 20 Sample layout for PC Remote Control. Left: PC showing overview of slate devices. Right: Detail view of screen on PC, after selecting a slate to control. Insets: Lassoing text, then selecting the copy command, at which point the text is copied to the distributed clipboard.....99
- Figure 21 Screen shot of the Stack Manager component of the Document Library. Users select a combination of system-added and user-defined tags along the sidebar to filter for documents. Bottom inset: Using buttons at the bottom of the interface, the selected documents can either be added to the stack, replace existing stack contents, or be saved with a new attribute tag. 101
- Figure 22 Using tag operations to manipulate, save and restore document stacks in order to save and resume reading activities across different environments and times. Although Slate A is always present in all of the venues in the above example, there is no requirement that a common slate be present. 103
- Figure 23 Multi-slate interactions support the reading requirements presented in Chapter 2. 104
- Figure 24 Example lifecycle of a document when using the cloud-based infrastructure. 127
- Figure 25 The flow of data through the cloud-based United Slates infrastructure. 131
- Figure 26 The slate condition. Three slates are arranged on the desk. Each slate had four buttons that corresponded to previous and next page controls. Note that the entire desk is covered with a wooden board to minimize any impact from the horizontal display embedded within the desk. 142
- Figure 27 Side by side comparison showing the appearance of pages in each of the device conditions. Note the tabs to access pages on the right edges of the electronic reading surfaces..... 143
- Figure 28 Close-up of the LHS condition with hardware controls labeled. Identical controls (not labeled) are on the other side of the screen. The pen, and the finger-

tracking marker are on the left of the screen. Note that the laptop is raised using a piece of foam board to reduce electromagnetic interference with the pen digitizer in the horizontal screen.....	144
Figure 29 Left: Diagram depicting the camera setup for the LHS condition. The OptiTrack tracking cameras (red) are at the far left and right corners of the desk as well as above the working area of the desk. Dotted lines show the field of view of each camera. Right: The reflective marker is affixed to the first knuckle of the user's index finger on the non-writing hand.....	145
Figure 30 Starting configuration of the paper, tablet and LHS conditions.....	151
Figure 31 Range of layouts users produced in the linear reading task (Task 1). Conditions from top to bottom: Tablet, LHS, Paper.....	154
Figure 32 Range of layouts users produced in the comparison reading task (Task 2). Conditions from top to bottom: Tablet, LHS, Paper.....	157
Figure 33 Range of layouts users produced in the cross-device reading task (Task 3). Conditions from top to bottom: Tablet, LHS, Paper.....	161
Figure 35 Frequency of each type of laptop arrangements.	162
Figure 34 Diagram of layouts users created to accommodate the laptop. Rectangles labeled "S" denote a surface (window, tablet or paper).....	162
Figure 36 Most preferred devices for each reading activity, by user.	166
Figure 37. Distinct subpopulations based on user preference. (Note: Y axis labels are tablet, tablet, LHS; X axis labels are LHS, paper, paper).	167
Figure 38 A diagram of the beneficial properties identified in each device condition. Benefits that conflict with each other are bolded and highlighted in red.	171
Figure 39 Left: Full page view of a document. Right: Space-filling thumbnail view of a document.	180
Figure 40 Scriboli-style radial menu for selecting commands in the deployed system....	181
Figure 41 Left: Writing in a note panel (middle) is scaled by 75% and placed at the panel anchor point (bottom). Right: User-created hyperlinks are indicated by a hollow square. Annotations can be placed beside the hyperlink as a reminder about the destination.....	182
Figure 42 Left: Active stack view; the widget on the right edge of each document entry removes the item from the active stack; "Save" button on bottom tags the stack with a new attribute. Right: Library view with user-created attribute tags listed in the left sidebar.....	183

- Figure 43 Macintosh version of the Slates application. The application shows the active document on the slates in the environment. Users turn pages using the arrow keys on the keyboard and launch a document on the PC by clicking. 184
- Figure 44 Top: Diary entry dialog on the slate. Bottom: Corresponding entry in the web diary interface showing that a user's notes are transferred to the web interface. 185
- Figure 45 Hannah's two laptop reading environment. One laptop was for reading and the other for taking notes. Hannah preferred this arrangement to switching between windows. 192
- Figure 46 Jane's notebook for taking written notes when reading with the iPad. 195
- Figure 47 Left: Anne would often read and annotate on a paper book. Right: Notes from paper books and digital documents would be transferred into the Scrivener software package, which provided tools for writing, organizing ideas, and research management. 196
- Figure 48 Annotations on one of Gina's document. Gina reported having to write larger than she was used when using the slates. 200
- Figure 49 Hannah would rest a slate on her lap to more easily read while writing on the PC. 211
- Figure 50 Left: Jane would keep a slate on the side of the PC for reference. Right: Jane would periodically hold the slate to read in detail. 212
- Figure 51 Left: Anne preferred to copy and paste from her slates directly into her notes on the computer. Right: Jane would use the Slates application (note application running on PC) to open documents she was viewing on the slates in order to access the digital content. 224

1 Introduction

While paper use in the workplace has been declining for the past decade [The Economist 2008] paper continues to pervade our reading activities. The reliance on paper should not be surprising, though. Paper possesses a number of qualities that makes it suitable for reading activities. Paper is high contrast, and thus easy to read. Individual sheets of paper are thin and light, affording the physicality and tangible manipulations that accompany reading [Marshall 2005; Marshall and Bly 2005]. Reading on paper requires no power, meaning that few restrictions are placed on when, where, or for how long one can read. The ability to write on paper provides a fast and fluid way to capture information. And since paper is so inexpensive, it is possible to place information across many sheets of paper, which function as independent information surfaces. These surfaces can then be spread out and viewed simultaneously in one's workspace [O'Hara and Sellen 1997; Adler et al. 1998; O'Hara et al. 2002].

In spite of all of paper's positive qualities, there are two big reasons why a shift toward electronic reading technology is desirable for reading tasks in the classroom and the workplace. First, reading tasks in these domains are well placed to benefit from electronic functionality. Many of these tasks involve documents that are stored and accessed digitally, and occur in support of work occurring on electronic devices. Second, paper is saddled with some fundamental limitations that are intimately connected to its physical characteristics. For instance, paper materials quickly become unwieldy as they grow in size, limiting the amount of information we can have available at any given time. Managing paper documents is also clumsy, requiring large amounts of effort to store, find, and access these documents [Sellen and Harper 2003]. Finally, generating all of these documents is costly from both financial and natural resource standpoints. At the scales with which

students, academics, and working professionals use information, these limitations of paper become significant burdens.

For the above reasons, a variety of research and commercial efforts have aimed to transition work on textual documents from paper to the computer. Beginning with the Wang Freestyle Tablet, which introduced the ability to annotate digital documents, devices like XLibris [Schilit, Price, and Golovchinsky 1998], and TabletPCs, have provided increasing support for various tasks associated with reading. Software solutions like the 3Book [Card et al. 2004] have attempted to do the same on the PC.

One common factor that has held these previous systems back has been the physical limitations of computing hardware technology. Until only a few years ago, even the most portable computing devices have lacked the mobility and physicality of paper-based documents; most have been comparable in weight to a laptop computer, or have had small, low-resolution displays. Recent advances in display technology and miniaturization now enable the creation of e-book and tablet devices that feature high fidelity displays and long battery lives, in thin and light packages. Since then, these devices have made significant inroads toward making electronic reading the norm; in February 2011, the Association of American Publishers reported that, for the first time ever, the number of electronic books sold surpassed those in paper formats [Sporkin 2011].

Yet, these new reading devices and tablet appliances still do not support reading as fully as paper, as evidenced by the fact these devices have found most of their success in leisure reading activities. Active reading activities encountered in the workplace [Adler et al. 1998; Tashman and Edwards 2011a] and in the classroom [Pugh 1978; Thayer et al. 2011] are considerably more complex and make higher demands on the reading technology. These activities include reading across multiple documents, a variety of navigation strategies, and the tight interleaving of reading and writing subtasks. Thus, attempts to use these devices

beyond leisure reading have been met with mixed responses at best [Behler 2009; Young 2009; Thayer et al. 2011]

The problem this dissertation addresses is the creation of a reading system that provides the comprehensive support for reading activities that paper presently offers. The main insight upon which our work is based is that existing reading devices up to now have generally provided only a single display surface for reading. Ethnographers have observed that tasks associated with active reading, such as information extraction [Adler et al. 1998], note taking [O'Hara and Sellen 1997; O'Hara et al. 2002; Thayer et al. 2011], document comparison [Adler et al. 1998], and cross-referencing [Marshall et al. 1999; Marshall et al. 2001; Thayer et al. 2011], all benefit from multiple reading surfaces. Consequently, the central hypothesis this dissertation seeks to test is whether an electronic reading device with multiple reading and writing surfaces supports active reading activities better than existing electronic devices, while providing beneficial digital functionalities that surpass those of paper.

Our solution takes the form of a novel reading environment with multiple displays. These displays consist of several slate devices plus a PC. Although this hardware configuration is itself unique in the literature, our goal was to use it to provide expanded reading functionality. Thus, our key contributions center around developing suitable infrastructure and interactions that enable the displays to work synergistically in support of reading activities, while reducing the overhead of working across these devices.

We found that the multi-slate configuration enables users to complete complex active reading tasks that are currently performed with paper. Furthermore, the system provides distinctly electronic advantages such as superior interoperability with computing devices, increased accessibility to document resources, and improved portability when compared to paper.

1.1 Thesis Contributions

This thesis offers contributions in three areas.

1. An analysis of the requirements for work-related and academic reading and the degree to which existing computing devices support these requirements.

The analysis consists of:

- A) A survey of work-related reading and academic requirements that had previously been scattered through the literature. The synthesis of these requirements provides a holistic view of the requirements surrounding reading.
- B) An analysis of the degree to which contemporary devices support the aforementioned reading requirements.

2. Design and implementation of a multi-slate reading system. Our work produced artifacts at the infrastructure, hardware, and user interface levels.

- A) The Dual-Display E-book [Chen et al. 2008] combined two physically independent screens into a unified reading device. The Dual-Display E-book addressed difficulties with navigation in digital documents. The Dual-Display E-book also validated initial hypotheses about the feasibility and utility of a multi-screen reading device.
- B) The development of custom slate devices to satisfy reading requirements rooted at the hardware level. The custom slate devices form the fundamental component of the multi-slate reading system.
- C) The United Slates system paired the custom slate devices with interactions designed to support active reading. The United Slates enabled the use of multiple slates with each other, and with other computing devices to satisfy reading requirements.

D) The creation of a cloud-based infrastructure for the United Slates system. Shifting the United Slates system to the cloud enabled the practical use of the United Slates in real-world deployments outside of the lab.

3. In-lab and *in situ* evaluations to characterize the operation of multi-slate reading systems.

A) A comparison between slate, paper and on-screen reading surfaces for reading activities. We found that the physicality of the slates had a positive effect on reading comfort and the ability to work alongside a PC. However, there were situations where having a tangible surface was not necessarily desirable.

B) An *in situ* deployment of the United Slates system with Ph.D. students in the humanities. Our users successfully used the system in support of the class work, research, and teaching responsibilities.

1.2 Solution Overview

In this section I outline how the dissertation addresses the thesis contributions.

1.2.1 Analysis of the requirements for work-related and academic reading

Our analysis involved enumerating reading activities described in the literature and assessing how well existing computing devices like PCs, tablets, e-book readers, and large tabletops support these reading activities. The analysis suggested that the various options users currently have for reading electronically individually cater to only a subset of these requirements. As a result, paper functionality surpasses any single one of these devices.

While performing our analysis, we noticed that reading activities occurred at many different levels. Previous work has described activities ranging from low-level operations on each page of a document, to interactions across different documents, to macro-scale phenomena like reading in different locations. The combined support for all of these different components determines the overall reading experience. Seeing reading activities

in this holistic light helps to explain why existing technologies can be suitable for certain reading tasks while failing to support others.

Our solution is the multi-slate reading system on which this dissertation focuses. We hypothesized that this novel configuration would enjoy a unique combination of characteristics that include: having more information available at a glance, comparatively high amounts of micro and macro mobility from each slate, the ability to partition information and activities across different physical spaces, and tight integration with existing computing devices. What makes these characteristics exciting is that they form the union of beneficial properties of existing electronic devices. That is, a multi-slate reading system combines the separate advantages of existing devices to finally offer a complete enough set of reading functionality to present a compelling alternative to paper.

1.2.2 Design and Implementation of a Multi-Slate Reading System

The origins of the multi-slate system can be traced to our efforts to enhance navigation on electronic reading devices. In this research, we sought to address three kinds of navigation challenges encountered while reading: local navigation between adjacent pages of content, random-access navigation within a document, and between-document navigation. We believed that a device with two displays could better address these challenges. Since such a device did not exist, we made our initial foray into building custom hardware. The Dual-Display E-book prototype we constructed confirmed that a dual screen configuration was able to support navigation interactions possible on paper better than a single screen. Our experience with the prototype also made clear that the physical characteristics of the reading device (and not just the software UI) affect the reading experience.

Following the work on the Dual-Display E-book, the aim was to test the applicability of the multi-screen reading configuration toward a broader range of reading activities, especially those occurring in the real world. For these activities, it was necessary to improve upon the

Dual-Display E-book hardware by providing a larger number of displays that offered tighter integration. At a practical level, these displays also needed to be big enough to view commonly used letter-sized documents, to have improved inking fidelity, and to be robust enough to use outside of the lab. When we started this phase of the project, approximately three years ago, no devices on the market delivered the hardware functionality we believed to be necessary. In response, we partnered with industry to develop a new hardware platform that did satisfy our requirements. The resulting device featured a smartphone class processor, reflective electronic paper display, pen input, and wireless communications. From a prototyping standpoint, the hardware created is of interest because it demonstrated that addressing design goals through hardware was no longer the exclusive domain of large engineering teams in industry. As the hardware we created demonstrates, rapid prototyping techniques have made developing custom hardware a more practical proposition for academic researchers.

Around this hardware platform, we created a new reading environment, the United Slates, which capitalizes on the potential of having multiple slates in the environment and reduces the overhead of managing these additional devices. Our interactions were in response to Morris, Brush and Meyers' [2007] finding that a collection of TabletPCs conveyed some benefits, but the lack of connectivity and interoperability between devices hindered the system. These interactions focus on supporting navigation activities, resolving physically awkward slate interactions, and finding more integrated ways of dealing with documents within the context of the overall active reading workflow.

Given the complexities of real-world reading, it was essential for the United Slates system to be tested under ecologically valid conditions. In order to make the United Slates feasible for use in real-world reading environments, a large portion of the system was re-implemented to use a cloud-based infrastructure. The cloud-based infrastructure simplified

the process of transferring data on and off the slates, increased system robustness, and allowed the system to scale to the reading demands of an extended deployment outside of the lab. The infrastructure employed services like Dropbox, Google App Engine, and Amazon Elastic Compute Cloud to provide an always-on system that enabled users to easily move documents onto their slates, and for annotations to propagate throughout their devices. The cloud-based approach helped ensure connectivity in a wide range of environments and ensured that users had ubiquitous access to their documents and annotations.

1.2.3 Evaluation of multi-slate reading systems

To better understand the implications of employing a multi-slate reading system, we tested the configuration using multiple evaluation methods. The role that physicality of the reading media plays in active reading processes was tested in a controlled laboratory study where reading from multiple slates, a single large horizontal display, and multiple sheets of paper was compared. The study showed that the slates afforded more micro-mobile interactions versus the large display, and improved organization versus paper. At a higher level, these results suggest that differences across reading activities make it impossible to support all reading activities perfectly with any single device. Rather, a major challenge revolves around creating an environment that is versatile and adaptable enough to transition between these activities.

As mentioned above, the reason reading is difficult to support is because of the wide variety of activities that it encompasses. Controlled studies, as a matter of necessity, present users with tasks that are only a rough facsimile of what occurs in the real world. Moreover, lab studies cannot effectively simulate the broader contexts for which reading activities are performed. Therefore, we concluded our examination of multi-slate reading with an extended *in situ* deployment of the reading system to doctoral students in the humanities.

The deployment provided an opportunity to observe whether the United Slates system would be useful in the real world, whether it would integrate with existing tools and practices, and what new practices would emerge from its introduction. On average, users had the devices for approximately 3 weeks, during which time the reading system was used for activities that included: reading for research, dissertation editing, and teaching. The results of the deployment illustrate many of the benefits that additional display real estate and physical separation of reading and writing surfaces contribute to active reading. Moreover, the electronic functionality of the system increased the availability of documents and annotations as well as streamlining reading and writing tasks that occur on the PC.

For the most part, the findings from the deployment confirm the hypothesis that a multi-slate reading system would better support reading activities than existing electronic devices, while surpassing what is possible when reading with paper. However, certain features of the United Slates that preliminary studies indicated would be essential, such as interactions between slates, did not receive extensive use during the deployment. This unexpected result suggests that these features may cater to the needs of a different population of users or that the deployed implementation may suffer from some design shortcomings. In any case, further studies will be necessary to gauge the true efficacy of these features.

1.3 Dissertation Roadmap

The first part of the dissertation covers the preparatory work we engaged in to inform the design and construction of the multi-slate reading system. Chapter 2 covers the literature that exists about reading on paper and reading electronically. We begin by detailing the various processes and needs surrounding active reading in the classroom and workplace. Next, we present various electronic tools that are currently available for reading. We then place existing devices in context with each other, as well as the reading requirements researchers have identified. Chapter 2 ends with a presentation of prior work on computing

across multiple devices and screens and the techniques that have been developed to support interactions in these hardware configurations. The interactions present in the United Slates system draw from these techniques heavily. Chapter 3 covers the Dual-Display E-book, which was an important precursor to the more sophisticated and full-featured multi-slate reading system we eventually created.

The second portion of the dissertation deals with the technical aspects of creating the multi-slate reading system. Chapter 4 presents the slate hardware platform on top of which the remainder of the multi-slate reading system is implemented. The chapter details the physical properties, electronics, and intermediate software layers on each slate device. Chapter 5 describes the variety of reading interactions the slates support, with an emphasis on functionality that capitalizes on the unique multi-display aspects of the system. We refer to the combination of slate hardware and interactions as the United Slates system. Chapter 6 presents the cloud-based component of the United Slates system, which increased the robustness of the United Slates and provided improved integration with users' existing computing tools—crucial characteristics when we ask users to use the United Slates for real work.

The final portion of the dissertation focuses on the evaluations we performed with the multi-slate reading system. Chapter 7 covers the laboratory study where we compared slates to paper and a large horizontal display. It starts with a comparison of these three reading systems, then describes the experimental setup, and ends with findings and implications from the study. Chapter 8 presents the final *in situ* deployment we performed with the United Slates system. In it, we describe the system we provided to users and the methodology we employed. We then give extensive profiles of our users and how they used the reading system. We end by providing key themes from our users' experiences with the

system and discuss the success and shortcomings of the United States during the deployment.

Chapter 9 ties together the findings made across all of the previous chapters and reconnects these to the initial hypotheses of the dissertation. The chapter places an emphasis on providing broad lessons about the use of multiple slates for reading. Chapter 10 concludes the dissertation with a summary of the contributions plus future directions for multi-device computing configurations to better support reading and other activities.

2 Related Work

This dissertation includes elements of user-centered research, hardware design and interaction design. As such, it draws from and builds upon prior research in those three areas. In this chapter, we first survey and unify results from research into what people do when they read, and extract out a set of common reading requirements. Then, we examine previous systems developed for electronic reading and place them in context with these reading requirements. Finally, we present the interactions for multi-screen and multi-device computing environments that inspire or form the basis of the interactions that make reading across multiple slates practical and powerful.

2.1 Supporting Reading Activities

The availability of computing devices capable of displaying text electronically has motivated investigations into whether computing devices can be effectively used to support these reading activities. Dillon [1992] provides a comprehensive survey of the various factors that can affect the usability of electronic texts versus paper texts. Dillon's conclusion was that no single factor dominates in explaining the differences between reading on paper and reading on electronic media. For instance, even though low-level properties of electronic reading media like the display resolution have been shown to affect reading performance [Mills and Weldon 1987; Muter and Maurutto 1991], the overall suitability of the electronic medium is dependent on many other properties as well. So while creating new display technologies is one approach for supporting reading, there are other avenues open to researchers. The approach this dissertation takes is to design hardware and interactions that address the process differences between reading electronically and reading on paper.

This dissertation focuses on how to increase electronic devices' support of the reading practices associated with reading in the classroom and in the workplace. The importance of improving these higher-level interactions is supported in studies that report no inherent performance differences between reading on electronic screens and on paper, but still find that there are significant differences in user preference. That user preferences can vary so dramatically despite outcome measures, like reading speed and accuracy, being identical suggests that the interactions possible with electronic media are perhaps more important in determining how well the medium supports reading. For instance, Holzinger et al.'s [2011] comparison of paper to electronic documents in the medical domain notes that:

“The professionals clearly stated after the experiment that the advantage of paper for their work in the hospital is much more than the mere visual quality of paper. It is its flexibility, ubiquitous availability and mobility and the possibility of carrying reports from one place to another and sharing them with colleagues.”

With regards to the processes encountered in the workplace and the classroom, the literature documents a wide range of reading practices. For instance, Pugh's [1978] study of how university students read identified five distinct styles of reading in which students engage. These include: receptive reading, the process of linearly progressing through the text without interruption; reading to search for a specific piece of information; acquisition of information without a set goal; reading to get an overview about the general structure of the material; and responsive reading, where the user engages with the material in the form of note taking, annotation, and cross-referencing. Adler et al.'s [1998] diary study of how professionals read in the workplace also identifies several categories of reading that include things like identification, learning, skimming, reminding, reviewing, and cross-referencing. Reading activities where readers are actively engaged and interacting with the text, which include many of the types of reading mentioned above (e.g., responsive reading, cross-referencing, critically reviewing text), have been collectively referred to in the literature as “active reading” [Adler and van Doren 1972; Schilit, Price, and Golovchinsky 1998;

Morris, Brush, and Meyers 2007]. It is important to investigate how electronic devices can better support active reading activities because they are generally the ones for which people heavily employ paper documents. It bears noting that although active reading tasks are defined by the presence of things like responsive reading and cross-referencing, these tasks do not occur in isolation [Thayer et al. 2011]. Rather, active reading is intimately intertwined with the receptive reading and skimming activities that also make up the reading task. Thus, it is impossible to support active reading while ignoring the requirements for other types of reading.

To get a sense of what an electronic device must do to support reading, we draw from a number of different sources. Ethnographers have produced detailed observations of how people make use of paper materials. Researchers have also compared the differences in reading process between using paper and electronic materials. Lastly, studies of electronic reading device use in real-world situations provide additional insights about requirements that are specific to electronic reading. The aggregate of this literature forms the intellectual basis upon which the dissertation rests.

2.2 How People Read With Paper

The central thesis in Sellen and Harper's *The Myth of the Paperless Office* [Sellen and Harper 2003], is that the heavy reliance on paper in the workplace stems from the fact that paper provides many beneficial features that help support reading activities. Thus, observing how people read with paper can be used to identify essential functionality that needs to be incorporated into an electronic device seeking to provide an alternative to paper. The following sections detail the key aspects that characterize how people perform active reading on paper.

2.2.1 Annotation and Note-Making

In professional and academic settings, one of the most commonly observed features of reading on paper is that of writing and in conjunction with reading [Marshall 1997; O'Hara and Sellen 1997; O'Hara et al. 1998]. The process of annotating and note taking are important components of actively engaging and interacting with textual matter [Adler and van Doren 1972].

Marshall's [1997] examination of annotations made in academic books is perhaps the most comprehensive treatment on this subject. Marshall specifies that freeform in situ annotation, placed within, but distinguishable, from the main text as an important design goal. Marshall also notes that annotations go beyond recording information for later use. For instance, Marshall describes the use of non-interpretive markings to help focus attention in difficult passages of text. Overall, Marshall regards the ability to smoothly integrate annotation with reading to be one of the most essential challenges for any reading.

More extensive notes made alongside reading are also a feature of work-related [Adler et al. 1998] and academic reading [O'Hara et al. 1998]. O'Hara et al. [1998] detail the reasons for and information recorded during note-making in academic reading. O'Hara et al. note that notes offer a way for readers to engage in a dialogue with the text and offer increased flexibility over the nature of information recorded. Note making also offers an opportunity to unify disparate sources of information and to condense information. These needs are echoed in Adler et al.'s reports of reading in the workplace. O'Hara et al. additionally note that readers prefer to make handwritten notes for their convenience over photocopying, portability, and ease of use later on in the writing process.

The literature as a whole suggests that it makes little sense to think of reading in isolation from writing [Adler et al. 1998]. Moreover, since markings and information contained

within are highly variable and devices must support free-form input from the user. Reading technologies that are designed to support active reading almost certainly need to provide annotation and writing functionality.

2.2.2 Micro-Mobility and Physicality of Materials

Micro-mobility involves the physical adjustments that occur with the reading material even though the reading activity occurs in the same venue. For instance, while reading on paper, sheets of paper can be picked up, folded, or moved around. Although primarily concerned with the reading of fictional texts, Mangen [2008] provides a broad survey about the role of tangibility and phenomenological differences between reading on tangible media versus reading electronic media. Mangen, in particular, notes the importance of materiality for maintaining engagement with the text.

The physical aspects of active reading were documented in O'Hara et al.'s [2002] description of readers' spatial rearrangements of reading surfaces. For example, they noted that readers picked up and dropped documents multiple times, lifted up documents for closer inspection, and held documents side-by-side. Marshall and Bly [2005] observed magazine readers repeatedly folding and unfolding the magazine to selectively cull away distracting content or to broaden their focus.

The notion that physicality introduces differences in task performance was tested in Terrenghi et al.'s [2007] photo sorting task. Although not a reading task, *per se*, the study illustrated that users uniquely took advantage of the physicality of objects. For instance, users in the study picked up physical photos to focus on them more closely. The overall conclusion that Terrenghi et al. arrived at was that graphical representation and multi-touch cannot recapture everything that is possible with physical objects. Instead, they believe that it would be prudent to support the physical aspects of interaction by retaining some tangible components.

The importance of physicality has been investigated from a cognitive science perspective, Guiard [1987] noted that the ability to use the non-dominant hand to manipulate a page prior while writing increased spontaneous writing speed. Guiard suggested that physical adjustments help establish reference frames that aid in the insertion of content.

Overall, the literature provides strong evidence about the importance and micro-mobility for reading. These results make a case for a tangible reading solution that is lightweight and easy to configure in a variety of orientations.

2.2.3 Rich Navigation Within Documents

The navigation processes associated with active reading involve more than a linear progression through the document being viewed. Many reading activities require users to jump between different parts of the same document, or glance back to re-read content. For example, readers in the research group that Marshall et al. [1999] studied exhibited skimming navigation and also flipped back to re-read materials. These readers also bent the page they were viewing back to see it alongside the references section at the end of the document. In their study of magazine reading, Marshall and Bly observed many examples of unselfconscious navigation around the current position in the text [Marshall and Bly 2005]. The style of navigation observed enabled readers to get an overview of the content and supported the serendipitous discovery of new content. O'Hara and Sellen [1997] noted the speed and automaticity of navigating on paper. For instance, the properties of the physical page made it possible to perform anticipatory page turning, save locations of interest with one hand, and to gauge one's position in the document based on feel. Also, the fixity of information on the printed page aids spatial memory, accelerating the random access of content. Based on these reports documenting the diversity of navigation within documents, a reading system that only support simple page turning is insufficient.

2.2.4 Working With Multiple Documents

Active reading activities additionally rely on the use of multiple documents simultaneously. Adler et al. [1998] studied a broad range of professionals' work-related reading activities. In their study, Adler et al. found that these professionals read for diverse purposes and adopt a number of different reading styles. However, commonalities across all professions include the use of several documents at once and a high frequency of reading occurring in conjunction with writing. Key reasons for reading from multiple documents include extracting or integrating information from one document to another, comparing documents, and following cross-references. Adler et al. conclude that "[activities involving multiple documents], need to be supported in digital reading devices if they are to satisfactorily replace current practice." [Adler et al. 1998, 246].

O'Hara et al. [2002] observed people of different professions performing real-world writing tasks while reading from multiple source materials. O'Hara et al. specifically mention the gains realized from being able to quickly glance across documents. Writers frequently shifted their attention across the different information sources. Moreover, the shifts in attention were extremely brief and marked by a great deal of back and forth movement, making the simultaneous presentation of information crucial. O'Hara et al. also described how multiple reading surfaces afforded laying out information spatially in the workspace. Spatial arrangements were used to support the readers' mental models of the content, guide attention, and to facilitate closer inspection of specific documents.

O'Hara and Sellen's [1997] controlled laboratory study comparing a reading and summarizing task on paper and on a computer revealed that participants frequently made use of different reading and writing surfaces while writing. The electronic setting's single display limited participants' abilities to quickly navigate across documents, or to arrange documents spatially on the desktop. Another finding of interest was that participants

preferred their reading and writing surfaces to be placed at different angles, to maximize comfort.

Paper's suitability for reading tasks involving multiple documents can largely be attributed to how it distributes content across multiple pages. Having distinct pages enables documents to be laid out in space. Moreover, spreading content across different pages also affords the separate reading and writing surfaces that allow reading and writing activities to be smoothly interleaved. For those reasons, O'Hara and Sellen concluded that additional reading surfaces, even just two, one each for reading and writing, would be beneficial.

2.2.5 Document Organization and Management

Related to multi-document use in reading are the strategies people employ for organizing and managing large collections of documents. Sellen and Harper [2003] observed that office workers often have a set of 'hot' documents that they switch between and also keep a set of 'warm' documents ready for use. Malone's [1983] study of the organization of documents on the physical desktop details how layout structures like files and piles speed up the process of finding and accessing documents. Documents that are laid out in the workspace also remind users about tasks they need to attend to. Bondarenko and Jannsen [2005] add that the ability to quickly reorganize and regroup documents with paper is critical, as tasks and needs can shift over time.

2.2.6 Macro-Mobility and Portability

Lastly, the macro-mobile elements of reading reflect the fact that reading does not necessarily occur in a fixed location. Adler et al. [1998] observed various work-related reading occurring outside of the office. Marshall et al.'s [2001] study of legal scholars reported that the students read in a variety of unpredictable locations such as the library, dorm rooms, and offices, to meet their resource needs. Even the computer science researchers that Marshall et al. studied, who do not fit the traditional profile of mobile

workers reported carrying around reading materials to different locations away from the office [Marshall et al. 1999]. Tashman and Edwards' [2011a] diary study of active reading further corroborates these results based on data gathered from an even broader set of users. Tashman and Edwards noted that over a quarter of active reading activities occurred in mobile locations and identified a strong need for reading workspaces to be portable. The requirement for portability suggest that reading are probably best supported using portable devices, rather than on a stationary system.

2.3 Providing “Paper-Like” Functionality

Transitioning activities that are traditionally done on paper to electronic alternatives has long been a goal of technologists and researchers. Electronic systems can address the shortcomings associated with paper documents such as the difficulty of accessing, storing and sharing paper documents [Sellen and Harper 2003]. Reading electronically can also supplement reading activities with digital functionality like interactivity, text search, and content extraction.

The central challenge in the design of these systems has been in figuring out how to provide additional functionality without sacrificing beneficial paper affordances. System designers have employed two general strategies in their attempts to tackle this challenge [Mackay et al. 2002]. The first strategy, augmenting paper, preserves paper documents as the primary working medium but provides auxiliary electronics to provide computer-like functionality. Paper augmentation can also be used to minimize the cost of switching to an electronic version of the document if necessary. The second strategy, taking inspiration from paper, starts with a computing device and adds hardware and software elements to enable the device to more closely match the functionality of paper. Below, we discuss the relative merits of these approaches.

2.3.1 Augmenting Paper

A number of systems take advantage of the existing qualities of paper and then augment it with additional capabilities. The DigitalDesk [Wellner 1993] was one of the earliest systems to employ this approach. DigitalDesk augmented paper documents by projecting a dynamic display alongside or on top of the paper. Another early example of augmenting paper was Ariel [Mackay et al. 1995], which connected engineering drawings on paper to a corresponding digital representation on the computer. More recently, the A-Book [Mackay et al. 2002] bridged paper and electronic content by using a PDA as a magic lens for viewing layers of digital information within a paper notebook. Paper++ [Luff et al. 2004] also explores the idea of augmenting paper by connecting printed documents with digital resources. Other projects that contain examples of how digital functions have been added to paper documents include: The Listen Reader [Back et al. 2001], Audio Notebook [Stifelman, Arons, and Schmandt 2001], and Books With Voices [Klemmer et al. 2003].

The recent introduction of Anoto Digital Paper¹, which employs specially printed paper along with a camera-equipped digital pen to identify where a user is writing has enabled the electronic augmentation to be moved to the pen itself. Commercial products employing Anoto technology include the Fly pentop computer² and LiveScribe³ family of products, both of which capture user writing and allow similar interactive tasks to be performed on paper.

Additionally, the availability of Anoto technology has also made possible cohabitation strategies such as PADD [Guimbretièrre 2003] and PapierCraft [Liao et al. 2008], which allow the user to easily move between working with the paper and electronic versions of a document. In the case of PapierCraft, digital functionality generally associated with

¹ <http://www.anoto.com>

² <http://www.flyworld.com>

³ <http://www.livescribe.com>

computers like text search can be performed directly on the paper document. Other examples of this approach include Prism [Tabard, Mackay, and Eastmond 2008], a hybrid paper-digital laboratory notebook, as well as ProofRite [Conroy, Levin, and Guimbretière 2004] and PaperProof [Weibel et al. 2008], digital-paper based systems where proofreading marks made on paper are incorporated into the digital version of the document. These cohabitation approaches collectively recognize that paper may not always be the best medium for a task, and that certain activities benefit from a fully electronic representation of a document.

One feature in common to augmented paper solutions is that they retain almost all of the properties of paper, regardless of whether they are good or bad. So while users enjoy the positive aspects of reading on paper, they must also contend with the fact that paper documents are static and have physical mass. For example, when using these systems users must go through the process of printing out documents and regenerating them if the digital versions change. Paper can also be cumbersome to transport and use when there are many documents, or when documents large in size.

For scenarios that rely on properties unique to paper such as the fact it is extremely cheap, disposable, foldable, etc., (e.g., for biological fieldwork [Yeh et al. 2006]) augmented paper systems are generally preferable. However, for other situations the downsides of working with paper must be weighed when deciding whether augmentation is the right approach. Reading tasks in the classroom and the workplace often involve the use of many large documents. Also, documents are often stored and distributed in digital form, meaning that generating paper printouts involves extra steps. These kinds of documents are the ones where the weaknesses of paper can be most easily detected. Therefore, we focus our efforts on the alternative strategy of simulating paper.

2.4 Simulating Paper

The alternative to augmenting paper is simulating paper [Mackay et al. 2002]. Simulating paper is predicated on the notion that with the right features, an electronic device can recapture the functionality paper enables without the use of paper. Although the terminology employs the term “simulating” it is important to note that these approaches do not necessarily involve replicating what people do on paper. It is possible, and often desirable, to draw inspiration from paper and to *transpose* paper interactions into a form that is more appropriate for the device.

The idea of an electronic device that simulates paper media was first presented in the Dynabook, Alan Kay’s vision of a portable device used to access digitized information. Kay’s thesis was that “[computers] can be all other media if the embedding and viewing methods are sufficiently well provided.” [Kay and Goldberg 1977, 31]. These “suitable embedding and viewing methods” are what efforts at simulating paper attempt to attain. The Dynabook also pioneered the idea that reading could be enhanced through computation.

Since the Dynabook, systems supporting various reading requirements have emerged, the multi-slate reading we present in this dissertation being one of them. Given the number and variety of devices that have existed, we provide the following section to present these devices in more detail.

2.4.1 Electronic Reading Devices

A variety of research and commercial devices have been developed for reading activities. The availability and use of these systems in recent years has inspired new designs for devices. Through the use of these devices, people discover important aspects of reading that they had previously taken for granted on paper.

2.4.2 PCs

The traditional personal computer, which includes both desktop and laptops, is a versatile tool well suited for composition and editing tasks that often go hand in hand with reading. PCs have robust multi-tasking capabilities along with comparatively large screens making them better suited for cross-referencing, re-reading, and sorting tasks. But, the PC's focus on being a very general, all-in-one tool results in software interfaces that are more complex and hardware that tends to be comparatively bulky and difficult to handle.

Researchers have nevertheless created techniques to enhance the PC reading experience. 3-Book [Card et al. 2004] provides a 3-D visualization that emulates the look of a paper book. 3-Book also offers navigation support by allowing rapid zooming and also the ability to “slide” out pages of interest for revisiting content quickly. Space-Filling thumbnails (SFT) [Cockburn, Gutwin, and Alexander 2006] are a navigation alternative to scroll bars developed to facilitate random access navigation while leveraging spatial memory. Systems that graphically simulate the results of physical interactions on paper [Hill et al. 1992; Alexander et al. 2009] can help users navigate to locations they had previously visited. As a whole, these techniques offer compelling solutions to the navigational requirements of reading and can be used in reading interfaces for devices other than PCs. The United Slates system we present, for instance, uses SFT for navigation. However, these software techniques do not yet address reading requirements like physicality and writing, which are associated with hardware functionality.

2.4.3 PDAs and Smartphones

Modern day smart phones and their predecessors, Personal Digital Assistants (PDAs), have been used for reading, especially in mobile environments, because of how portable they are. However, the limited amounts of information that can be displayed at any given time on their small screens have made them less useful for active reading activities. Marshall and

Ruotolo [2002] reported on the deployment of PDAs for use with course materials for university students. In the deployment, Marshall and Ruotolo found that the portability and availability of the devices to be advantageous. The ability to perform text searches was also found to be beneficial. However, the small screens on the devices meant that layout-sensitive texts were difficult to use and random-access navigation on the devices was difficult, particularly in long documents. Two ways for overcoming the limited display space of PDAs are to use multiple devices to implement space-sharing and content-splitting [Fails, Druin, and Guha 2011]. Although Fails, Druin and Guha studied the use of these techniques in multi-user environments where users each have a PDA, they are equally applicable to a single-user environment with multiple slates. As such, these techniques are employed heavily in the multi-slate reading systems we created.

2.4.4 Tablet Computers

The Gridpad and Wang Freestyle were the first commercial systems in a tablet form factor. The Freestyle system, of which the Freestyle tablet was a component, was particularly notable as it was one of the first systems to attempt to implement paper practice, such as annotation on documents, in the digital world. Since the introduction of those initial tablet devices, a variety of tablet devices have appeared in research and commercially. The XLibris system, built on various commodity hardware platforms, is the primary example in the research domain. Commercially, devices range from the TabletPC, which preserve many of the conventions of the PC, like multi-tasking to tablet appliances like the iPad, which adopt an app-style interaction model.

XLibris

XLibris [Schilit, Price, and Golovchinsky 1998] was a research system for supporting active reading on pen-enabled hardware devices. Versions of XLibris were created using a variety of hardware devices. Earlier versions ran on desktop computers with an attached pen

enabled screens and later iterations of XLibris were built using pen-enabled tablet computers. XLibris primarily consisted of software designed to convert a general-purpose tablet device into a machine dedicated to supporting active reading tasks. XLibris specifically aided active reading tasks with features like free-form pen annotation, full-page views of documents, and on-screen controls to perform intra and inter-document navigation. XLibris additionally offered special digital functions that improve annotation retrieval and hyperlink backtracking [Marshall et al. 2001], and enable the automatic creation of hyperlinks to related content [Price, Golovchinsky, and Schilit 1998].

XLibris was tested with various groups of people including researchers in a reading group and legal scholars [Marshall et al. 1999; Marshall et al. 2001]. Users in these studies found that XLibris' document presentation, annotation, information extraction, and non-linear and cross-document navigation features facilitated reading activities. On the other hand, the researchers found themselves less likely to follow references because facilities to view the main material and references simultaneously were not available on XLibris. The legal scholars who evaluated the prototype expressed a need for the device to better integrate with the work on PCs, leading researchers to wonder whether a laptop with annotation capabilities would be more suitable.

The United Slates system we created has similar goals to XLibris and incorporates some of XLibris' advanced functionality in support of reading requirements. However, by virtue of the multi-display design, the United Slates system we created can address issues like the limited writing area on the margin of documents, and the inability to view different locations at the same time that were brought up during XLibris usage. Also, the United Slates system offers enhanced capabilities to interoperate with existing electronic devices like PCs whereas XLibris did not include these provisions.

TabletPC

TabletPC devices were introduced to address the lack of writing capabilities on the regular PC. Therefore, TabletPCs provide a high-resolution pen digitizer that can capture ink markings from a stylus. Beyond that, TabletPCs offer similar functionality as a regular PC, and are comparable in size and weight to PC laptops. Marshall and Bly's [2005] comparison of magazine reading on paper and on a TabletPC device found that glancing back to re-read, glancing ahead, and the presence of helpful physical interactions with paper magazines were absent in the TabletPC editions. This research lend furthers credence that reading consists of more than visiting pages in a document in a linear sequence. Morris, Brush and Meyers [2007] employed multiple TabletPCs in studying the role of additional displays in active reading. Morris, Brush and Meyers found that the mobility of the TabletPC was beneficial when compared to stationary displays. The independence and isolation of each TabletPC in the experiment proved to be problematic, however. Also, the lack of a keyboard on the TabletPCs hindered users' ability to perform composition tasks.

Tablet Appliances

Contemporary tablet appliances such as the iPad, Kindle Fire, and Nook Color count reading as one of their features. These devices depart from the standard laptop/Tablet PC interface by subscribing to a full-screen application model, much like XLibris. Most tablet appliances these days forego a pen digitizer, however, and rely instead on multi-touch technology for interactions. By virtue of their low weight, high-resolution screens, and fast processors, they are proving to be good reading platforms for single documents. Moreover, the multi-touch capability found in almost all of these devices offers many possibilities for enhancing navigation, such as those demonstrated in LiquidText [Tashman and Edwards 2011b].

2.4.5 E-book Readers

Portable dedicated electronic reading devices have existed since the 1990s. One well-known example from that era was the Rocket eBook from 1998. A broader overview of these first generation e-book devices can be found in Schilit et al.'s [1999] survey of these devices. These early e-book readers generally featured smaller screens, were not particularly lightweight and had limited memory. These devices were also the subject of Wilson and Landoni's [2003] survey of several different electronic reading devices. Wilson and Landoni's study identified portability and organization features as the main positive aspects of these devices. Other benefits they identified were text search and the ability to follow hyperlinks. However, device size, weight, battery life, display resolution, and navigation speed were areas where users at the time felt these devices could be improved. In the decade that has elapsed since the first reading devices were introduced, hardware technologies have improved dramatically. The current generation of e-book readers, which include products like the Sony Reader, Amazon Kindle, Barnes and Noble Nook (Figure 1) offer highly readable screens, low weight, and long battery lives, making them far more suitable for linear (receptive) reading tasks. However, the slow refreshing black-and-white displays (from which many of their beneficial physical properties derive) limit the

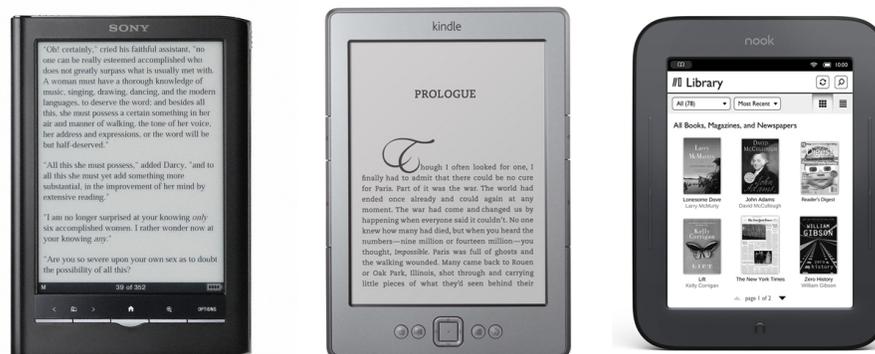


Figure 1 Contemporary e-book reading devices. From left to right: The Sony Reader PRS-650, Amazon Kindle, Barnes and Noble Nook Simple Touch.

navigation capabilities as well as the content that these devices can reproduce. With the exception of select devices from iRex (now defunct) and Sony, these devices only support basic highlighting and typed, rather than freeform, annotations.

Recent deployments of commercial e-books into the classroom [Behler 2009; Young 2009] (the Sony Reader in both cases) have noted serious deficiencies in the navigation systems in these devices for classroom reading. These navigation issues often stemmed from the slow refresh rates on the displays. Another problem these studies identified was that illustrations, color ones in particular, were not reproduced with sufficient fidelity.

The Kindle DX, which is another commercial reading device, has recently been tested in several universities. The Kindle DX is somewhat unique in providing a larger than average sized electronic paper display (9.7" diagonal), that can comfortably display full paged documents. Thayer et al. [2011] tracked Kindle DX use among Computer Science students. As with previous experiences with e-readers in the classroom, Thayer et al. found that the device they tested lacked support for the diverse navigation needs of student readers. In particular, Thayer et al. noted that the lack of tangibility in the pages of an e-reader removed important physical, or kinesthetic cues that helped in giving users a sense of where they were in the text, and how much material remained. Thayer et al. also reported that the absence of freeform ink annotation capabilities severely limited students' abilities to read responsively. A report detailing the deployment of Kindle DX devices at Princeton University in International Affairs and Classics courses mirrored many of the sentiments regarding the difficulty of navigating and the lack of annotation facilities [The Trustees of Princeton University 2010]. These experiences with e-book readers in the classroom highlight the importance of writing and supporting rich navigation.

2.4.6 Large Displays and Surface Computing

Computers that consist of large tabletop-sized displays have also started to enter the mainstream. One of the first implementations of such a system was built using the DiamondTouch [Dietz and Leigh 2001], a large capacitive sensing surface on which a software interface can be projected. A number of new techniques for creating large horizontal, multi-touch screens have since been developed. Two prominent systems that have received a great deal of attention are the Microsoft Surface and various systems based on Han's research into Frustrated Total Internal Reflection (FTIR) multi-touch sensing [Han 2005].

Although we are not aware of a reading system specifically developed for a tabletop configuration (or any subsequent evaluations of the technology for reading), multi-touch computing surfaces like the Microsoft Surface have positive characteristics that support reading for knowledge work. Tabletops, by virtue of allowing direct manipulation of objects spread across a large surface are suitable for spatial layout and working with several documents at once. Certain tabletop configurations, like the FLUX system [Leitner et al. 2009], are also able to receive high-resolution pen input.

While the vast majority of these systems are large, stationary devices, which may limit their use as portable reading systems, the PlayAnywhere system [Wilson 2005], packages up cameras and a projector into a compact unit, and demonstrates that this is not necessarily the case. Other technologies on the horizon, such as large rollable or flexible displays [Geller 2011], could result in large displays that are more portable. But, if these potential developments are taken into account, tabletop displays still restrict where and how reading is performed. For one, tabletops make it difficult to support the physicality associated with reading because documents are confined to the surface. Additionally, with a single large

screen, these devices would necessarily require environments that provide a correspondingly large and flat working area.

2.5 Reading Requirements and Technologies At A Glance

To set the stage for the reading system described in this dissertation, we recap the reading requirements we obtained from the literature about how people read in Table 1. In this table, we group the requirements into five major semantic levels. At the lowest level are requirements about the physical nature of the reading medium. Addressing the physicality of reading and writing must be done at this level. On the next level are the activities users do when reading one page of content. Above that are requirements concerning the document as a whole. Then, there exist requirements that concern the use of multiple documents in the workspace. Finally, macro requirements reflect the needs of users to read in different venues, and to read in across multiple reading sessions.

	Where Described in the Literature
Hardware / Physical Requirements	
Mobility; thin, lightweight, graspable	Marshall and Ruotolo 2002; Marshall 2005; Wilson and Landoni 2003
Readability	Wilson and Landoni 2003; Young 2009
Support for writing	O'Hara and Sellen 1997; Adler et al. 1998; O'Hara et al. 2002; Sellen and Harper 2003; Thayer et al. 2011
Page Level Requirements	
Linear reading	Adler and van Doren 1972; Pugh 1978
Graphics: Page Layout / Illustration	O'Hara and Sellen 1997; Marshall and Ruotolo 2002; Young 2009; Behler 2009;
Superimposed ink annotations	Adler and van Doren 1972; Marshall 1997; Thayer 2011
Text search	Wilson and Landoni 2003; Young 2009
Glancing back to re-read	Marshall et al. 1999; Marshall et al. 2001
Physical page cues	O'Hara and Sellen 1997; Thayer et al. 2011
Document Level Requirements	
Non-sequential navigation	O'Hara and Sellen 1997; Behler 2009; Tashman and Edwards 2011a
Building cognitive map of content	O'Hara and Sellen 1997; Wilson and Landoni 2003; Thayer et al. 2011
Skimming to get overview	Pugh 1978; Adler et al. 1997
Discovery of topical knowledge	Pugh 1978; Adler et al. 1997
Switch between navigation styles	Thayer et al. 2011
Workspace Level Requirements	
Sorting and triage of documents	Adler et al. 1997; Marshall and Shipman 1997
Spatial layout	O'Hara and Sellen 1997; O'Hara et al. 2002
Extracting information	Adler et al. 1997; Morris Brush and Meyers 2007
Reading from multiple documents	O'Hara and Sellen 1997; O'Hara et al. 2002; Wilson and Landoni 2003; Tashman and Edwards 2011a
Integrating with PC workflows	Marshall, 2001; O'Hara et al. 2002; Morris, Brush and Meyers 2007; Tashman and Edwards 2011a; Thayer et al. 2011

Macro-Mobility Requirements	
Reading in different venues	Sellen and Harper 2003; Thayer et al. 2011; Tashman and Edwards 2011a
Filing and archiving	Malone 1983; Sellen and Harper 2003
Restoring reading workspace, resuming reading activities	Malone 1983; Wilson and Landoni 2003; Tashman and Edwards 2011a;

Table 1 Comprehensive requirements of reading drawn from the literature.

Table 2 shows how well existing reading solutions map to the requirements identified above. For the analysis, we try to consider the newest software developments for each platform. For many of these requirements, we consider paper to be the target experience, given the paper's support for those requirements. However, it is important to note that paper is not perfect and that there are several requirements for which paper is distinctly ill suited.

	KEY					
	PC	Tablet	Tabletop	Tablet Appliance	E-book Appliance	Paper
Hardware / Physical Requirements						
Mobility; thin, lightweight, graspable	○	●	○	●	●	(●)
Readability	●	●	○	●	●	(●)
Support for writing <i>(Note 1)</i>	○	●	●	○	○	(●)
Page Level Requirements						
Linear reading	●	●	●	●	●	(●)
Illustrations	●	●	●	●	○	(●)
Superimposed annotations	○	●	○	○	○	(●)
Text search	●	●	●	●	●	○
Glancing back to re-read <i>(Note 2)</i>	○	○	●	○	○	(●)
Kinesthetic cues	○	○	○	○	○	(●)
Document Level Requirements						
Non-sequential navigation	●	●	●	●	○	(●)
Building cognitive map of content	○	○	○	○	○	(●)
Skimming to get overview	○	○	○	○	○	(●)
Discovery of topical knowledge	○	○	○	○	○	(●)
Switching between reading styles <i>(Note 3)</i>	○	○	○	○	○	(●)
Workspace Level Requirements						
Sorting documents	○	○	●	○	○	(●)
Spatial layout	○	○	●	○	○	(●)
Extracting information	(●)	●	●	○	○	○
Reading from multiple documents	○	○	●	○	○	(●)
Integrating with PC workflows	(●)	○	○	○	○	○
Macro-Mobility Requirements						
Reading in different venues	(●)	●	○	●	●	○
Filing and archiving	(●)	●	●	○	○	○
Restoring reading workspace, resuming reading activities	○	○	○	○	○	○

Notes on the table: 1. Glancing back to re-read is a function of how much content can be viewed at once, hence better scores for devices with more screen area. 2. Ability to reconfigure the interface for different reading styles and requirements.

Table 2 Degree to which existing devices support reading requirements.

One product of our analysis is the hypothesis that a system combining the physical properties of a dedicated e-book reader, with the pen input of a Tablet PC, and the multi-document functionality of a tabletop system would begin to match the capabilities of paper. One problem is that creating a single device to support all of these characteristics is impossible given inherent conflicts in the nature of the devices. For instance, the size of the tabletop is at odds with the lightweight graspable e-reader. Instead, we take a cue from how paper documents distribute contents across many sheets of paper that are individually thin and lightweight but powerful when used in aggregate. We advance the idea that a multi-slate solution where each slate is individually useful for reading, but can also be used as part of a larger ensemble [Schilit and Sengupta 2004], would better support the range of reading requirements than existing devices.

2.6 Distributing Content Across Screens and Surfaces

The United Slates system adopts the strategy of distributing content onto several independent slate devices in light of the importance of supporting freeform annotation alongside multiple simultaneous document use. The idea of a workspace containing several independent electronic surfaces was put forth in Weiser's [1995] vision of the future of computing. Weiser envisioned electronic reading devices that were as plentiful and disposable as scratch paper.

One of the more familiar configurations where content is distributed across multiple electronic displays is that of multi-monitor environments [Grudin 2001]. Grudin made the surprising finding that physical divisions between screens can actually be beneficial. Grudin

observed that the physical partitions help users multi-task and work with multiple sources of information.

Recently, researchers have started to build systems that consist of more than one tablet to take advantage of the additional screen real estate and new interaction possibilities they bring. We built and compared a Dual-Display E-book (Chapter 3) with a single-display device and found that the dual-display device exhibited benefits that were previously found to be characteristic of paper documents [Chen et al. 2008]. Hinckley et al. [2009] explored the design space of interaction techniques for a dual-screen tablet device. Neither the Dual-Display E-book nor the Codex examines the logistical and interactional issues that arise when the number of devices that are available is increased or reduced.

Paper windows [Holman et al. 2005] explores interaction techniques that would be possible if tablets have the properties of paper (using paper-based mockups). The United Slates system takes a more restrained view of the future and as such is designed to be sensitive to the limitations that we expect slate-type devices to impose in the immediate future.

2.7 Interactions Between Displays/Devices

One of the problems with using several devices together comes from the fact that these devices are computationally and interactionally isolated. The isolation is problematic because it limits the devices' capability to work together in a collaborative fashion. Morris, Brush and Meyers' [2007] studied a system composed of three independent Tablet PCs that were not functionally linked. This study was the first instance where more than two independent tablets were used together to complete a reading task. Users indicated the spatial layout and annotation capabilities of the tablets were valuable. However, users reported issues from their inability to move information between tablets. Other studies of multi-device use has described problems associated with working with multiple devices like setup overhead, ensuring access to data, and making sure data was synchronized

[Oulasvirta and Sumari 2007]. Thus, central to the goal of providing an effective multi-slate reading system is the dual challenge of creating interactions that capitalize on the configuration while also addressing the complexities associated with managing multiple independent devices.

While the literature does not provide specific examples of how to construct an interface to unify slate devices for a reading application, there are several techniques that can be adapted to the purpose of connecting reading devices together. One frequently encountered example of using multiple screens is that of multi-monitor computing. These screens can be attached to the same computer or to different ones using software packages like Xinerama⁴. However, multi-slate systems are slightly different from traditional multi-monitor systems in that the relative positions of displays in a multi-monitor setup tend to be static and can share a unified 2-D coordinate space. In contrast, a single 2-D coordinate space is impossible with multiple slates that can be freely positioned in 3-D. Therefore, in the United States, slates are treated as separate devices and separate spaces; the goal is instead on loosely bridging these spaces.

One important requirement for bridging these spaces is to simplify the movement of information across the spaces. A great deal of work has been done in streamlining data transfer between devices. The examples most relevant to our work are SyncTap [Rekimoto 2004], Pick-and-Drop [Rekimoto 1997], Synchronized clipboard [Miller and Myers 1999] and Stitching [Hinckley et al. 2004]. The one aspect all of these techniques have in common is that they offer users a fast and direct way of specifying the endpoints of a transfer operation.

Another strategy for unifying disparate devices is by allowing one device to control another remotely. Remote control spans the spectrum of functionality provided by Virtual Network

⁴ <http://sourceforge.net/projects/xinerama/>

Computing [Richardson et al. 1998], which provides a full proxy to a different computer system, to systems like Pebbles [Myers 2001] and PointRight [Johanson et al. 2002], where portions of the input and output stream are redirected to control other devices. These techniques are valuable in situations where it is physically awkward or undesirable to directly operate a device. For example, our United Slates system uses remote control in situations where one wishes to view content up close but also wants to preserve a particular spatial layout.

Finally, with improvements in sensing technology, physical connection of devices as a means of specifying relationships between devices has also been explored. Examples include Siftables [Merrill, Kalanithi, and Maes 2007] and ConnecTables [Tandler et al. 2001]. Although we believe that these interactions are valuable, slate devices have appreciable mass and activities that result in having to repeatedly connect or disconnect devices can be cumbersome. The lesson here is that although physical connection can be useful, it may not always be appropriate.

Inter-device communications are possible only with the appropriate infrastructure for linking devices. Although network connectivity is ubiquitous these days, CloudBooks [Pearson and Buchanan 2011] is perhaps the first to describe a cloud-based infrastructure for linking together multiple iPads in a reading application. The infrastructure lets devices exchange messages through the cloud as well as provides synchronization functionality. We use a similar cloud-based infrastructure to support the multi-slate operation of the United Slates system. However, the United Slates provides considerably more functionality in terms of document synchronization, device management, and robustness.

The overarching goal in assembling this combination of inter-device interaction techniques in our slate reading system is to establish a unified set of multi-slate interactions. To provide an analogy, these interactions serve a similar purpose as interactions on desktop GUIs that

enable information to be exchanged between windows. The vision is that establishing a consistent set of interaction conventions between devices we provide an extensible platform on top of which useful tools that leverage the multi-device configuration can be created.

2.8 Broader Implications of this Work

A recurring theme in the history of reading technologies is that emerging needs prompt innovations in the ways for creating, distributing and presenting textual material. These technological innovations then permanently alter the face of reading. Two notable examples in history (out of many) are the advent of silent reading and the invention of the movable type printing press.

The first example of how supporting emerging reading practices can have a significant knock-on effect is the development of silent reading. Silent reading of texts, did not become commonplace until well into the middle ages [Manguel 1996]. Prior to that, books were intended for reading aloud. As a result, there was no need to add separation between words and sentences. However, a need to more quickly peruse texts prompted scribes to gradually add features to the text like punctuation marks, upper case characters, and eventually spaces between words and sentences. The addition of spacing made silent reading much easier to perform and completely revolutionized readers' relationships with their texts in a way that continues to today.

The creation of the printing press is another example of how a technology supporting new reading needs has shaped human history. Motivation for the printing press grew out of increased demand for books caused by growing literacy in the fifteenth century. Clanchy [2007] writes, "A vigorous book-using culture was the precursor to the invention of printing rather than its consequence." The invention of the printing press has been revolutionary; the printing press has been credited as making possible the scientific

revolution, and sparking the Protestant Reformation [Eisenstein 1979]. It is, of course, at the root of contemporary print culture as well.

Electronic reading might be viewed as another point in the constant evolution of reading technology. As such, electronic reading can similarly have a profound effect on how people read. One way electronic reading can shape reading is that it can offer new ways to experience existing texts. The example Alan Kay gave to motivate the Dynabook was the ability to seamlessly switch between the four different viewpoints presented in each volume of Durrell's *Alexandria Quartet* [Kay and Goldberg 1977]. Electronic reading also spurs the creation of new kinds of textual materials. For example, the ability to proceed non-linearly through documents via hypertext has changed the way modern readers access information. Another potential application of electronic reading is in enabling new forms of literature that challenges the traditional notion that texts need to have an implicit ordering. Johnson's *The Unfortunates* [Johnson 1969] was an experimental work distributed in an unbound state so that the reader could read the "book" in any order. Non-linear texts are clumsy to consume on paper but are ideally suited for the electronic format.

Beyond advancing the content that is consumed, electronic reading technology can catalyze major societal changes. Electronic publishing, distribution and consumption of texts have resulted in the democratization of publishing and new authorship models. For instance, open-access journals and open-source textbooks have the potential to dramatically extend the dispersion of knowledge. However, these developments are only possible now that electronic reading technologies have removed the barriers of entry imposed by traditional print materials.

In tackling contemporary challenges associated with moving reading from paper to electronic screens, this dissertation has the practical effect of allowing users to enjoy the significant interactional advantages associated with paper while enabling new reading

experiences. Although this dissertation focuses on supporting a relatively well-defined set of requirements associated with active reading, the past history of reading suggests there will probably be broader implications that arise from these incremental improvements.

3 The Dual-Display E-book

This chapter describes the Dual-Display E-book, which marked our initial efforts into developing custom electronic devices in support of reading activities. The Dual-Display E-book was used to test whether additional displays might better support navigational tasks encountered during reading activities. The Dual-Display E-book confirmed hypotheses that multiple reading surfaces were beneficial for navigation activities and demonstrated the feasibility of creating custom reading device hardware.

3.1 Introduction

Although electronic editions of books, magazines, journal papers and newspapers are becoming ubiquitous, people still read from paper versions of these materials. Although laptop and tablet computers offer electronic displays that can display content with high fidelity, physical attributes like size and weight limit these devices' mobility and preclude them from being held in comfortable orientations.

In response to these limitations of standard PCs, specialized devices like e-books have been offered. The most recent of these devices, which include products like the Amazon Kindle, use electronic paper displays, one of several new display technologies that require minimal power to maintain an image on the screen and can be read in bright sunlight. Since these devices require much less power than those using LCD displays, they can be made thinner and lighter, resulting in a more mobile platform for reading. However, almost every contemporary e-book reader provides a single screen and therefore lacks important affordances of paper-based reading materials. In particular, embodied navigation in the form of turning, flipping and folding pages is limited. Moreover, the single display e-book

Portions of this chapter were originally published by the author, François Guimbretière, Morgan Dixon, Cassandra Lewis, and Maneesh Agrawala in [Chen et al. 2008]

reader makes it difficult to configure the device to simultaneously show different parts of a document, or different documents simultaneously. These features are important because ethnographic studies have shown that the extra display area can be beneficial for navigation activities [Marshall and Bly 2005] and that a large proportion of reading activities involve the use of multiple display surfaces [Adler et al. 1998; Sellen and Harper 2003].

A reading device with two displays could potentially enable the creation of navigation interfaces that support reading activities. Although prototypes of dual-display devices had been created (e.g., EveryBook [Martin 1998], iRiver E-BOOK [Block]), these prototypes had not been evaluated. Moreover, these prototypes were not available to us. Therefore, we built a prototype dual-display reading device with integrated motion sensors (Figure 2) as a platform to test how a dual-display e-book could benefit navigation tasks. The prototype was designed to study two specific qualities of a dual-display e-book: the use of an embodied interface and the ability to support flexible display configurations.

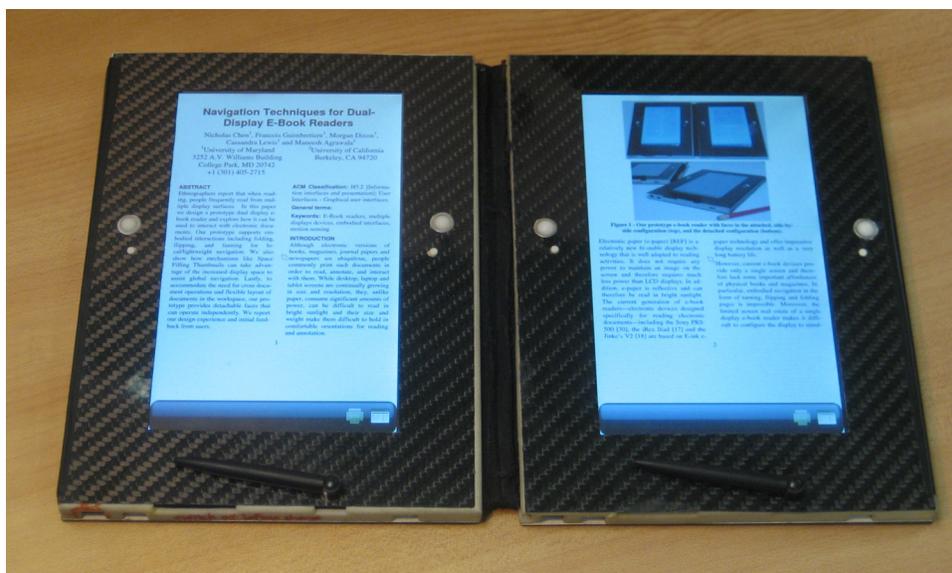


Figure 2 The Dual-Display E-book in the connected configuration.

1. Use of an embodied interface. Reading paper documents requires physically turning, flipping, and sometimes folding back pages. A bound, dual-display configuration offers a similar set of physical interactions and we extend prior strategies for creating embodied interfaces in prior systems, like tilting [Harrison et al. 1998] and background interactions [Hinckley et al. 2005b] to a system with two displays. The Dual-Display E-book allows readers to turn pages by fanning one face toward the other and, in the back-to-back configuration, by flipping the device over.

2. Flexible display configurations. A significant drawback of many e-book readers is that they offer 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. The Dual-Display E-book offers a compromise between increased screen real estate and a convenient form factor. 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 [O'Hara and Sellen 1997].

Using our prototype, we examined how these different features can be combined to offer a richer reading experience. An evaluation comparing a single-display device to a dual-display device found 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. These results have added significance because the advantages we found with the dual-display system mirror many of the benefits gained from

reading with paper materials. These results lend credence to the notion that additional displays are a way to convey capabilities available with paper to the electronic domain. Our evaluation also revealed that physical characteristics play a major role in the reading experience, a lesson with significant implications for our subsequent work.

3.2 Design Of The System

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 [Marshall et al. 1999; Marshall and Bly 2005]. Dual display e-book readers can facilitate these activities by increasing the visible 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 within 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, mainly work-related reading, 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 [O'Hara and Sellen 1997]. Dual-display systems with detachable faces can improve multi-document navigation by allowing users to arrange and interact with each display separately.

3.2.1 Improving local navigation

In addition to the act of turning a page, Marshall and Bly [2005] describe several other processes at work when navigating locally. For example, with magazines, people switch between side-by-side and back-to back configurations of the pages. Marshall and Bly also explain that dynamically expanding and contracting the viewing area is also a common form of unselfconscious navigation. Broadening the focus allows for a better overview of the material and narrowing the focus prunes away distracting material. The Dual-Display E-book allows a similar interaction. It provides an automatic transition between displaying one or two faces worth of information based on the relative configuration of the two slates.

Flipping The Device Over

When reading any two-sided document, such as a magazine, with the pages folded back, one can change pages by flipping the document over. The Dual-Display E-book 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, the flipping gesture continuously advances or rewinds the pages depending on the orientation of the flip (flipping clockwise advances while flipping counterclockwise 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

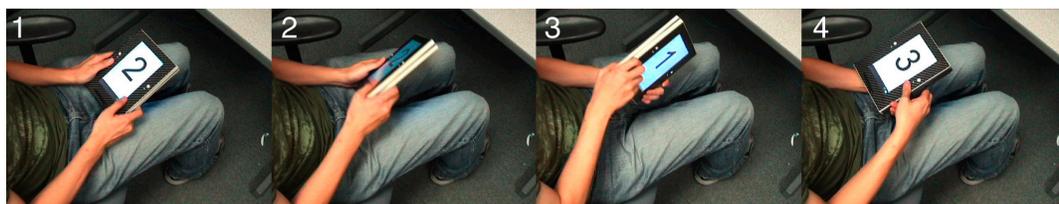


Figure 3 Flipping gesture on the Dual-Display E-book advances the page.

the page last viewed ready allows users to rapidly switch back and forth between the most recently viewed pages in a lightweight fashion.

Fanning The Faces

For a document 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. The Dual-display E-book 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 shift the position of their hands while reading, either for comfort, or to perform other activities [Marshall and Bly 2005].

Early in the design process, we also considered using the tilt of each slate to continuously turn pages [Harrison et al. 1998; Bartlett 2000]. However we found users had difficulty maintaining tilts, making it difficult to control. We omitted the feature after we decided the problem would be even worse if this sort of device was used outside of the relatively vibration-free confines of the laboratory.

Hardware Controls

To complement the embodied navigation interface, our design also includes hardware

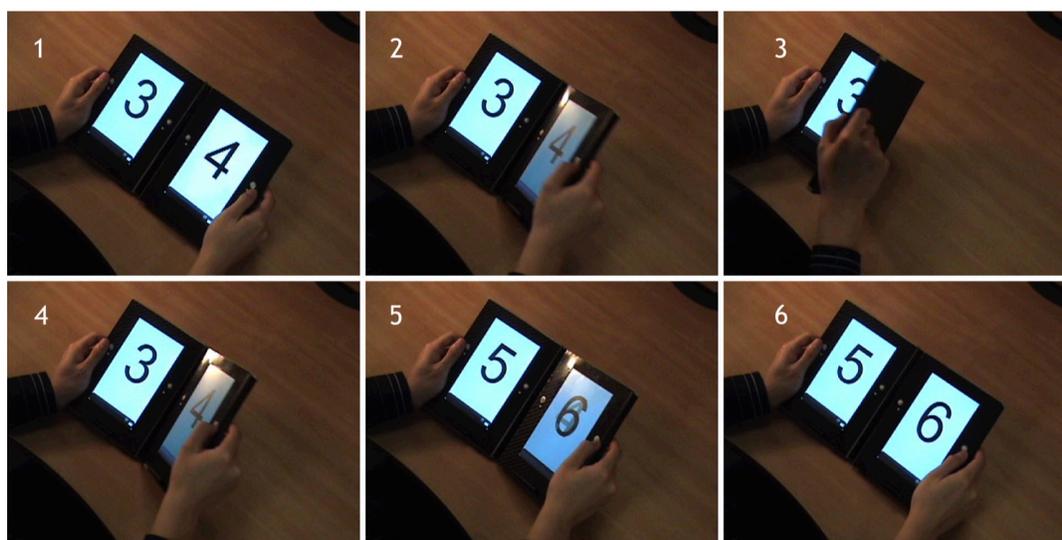


Figure 4 The fanning gesture being performed on the Dual-Display E-book.

controls in the form of one small clickable trackball on each side of the display (Figure 5). These controls must be used when the screens are detached but can also be used when screens are attached in situations where flipping and fanning are not convenient, such as when the device is placed on a desk.

Since the reading device could be used in a variety of configurations, control layouts needed to be ergonomic and logical for all of these settings. We explored several potential locations for the trackball, including below the screen or on the side of the device, and 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 switches the side of the device facing the user, the controls are laid out identically on the two faces to make the two sides indistinguishable. Giving each face an identical layout also keeps the interface uniform when the displays are detached.

Even though the trackball detects two-dimensions of motion, when the device is held with a single hand, the range of motion of the thumb was limited primarily to left and right

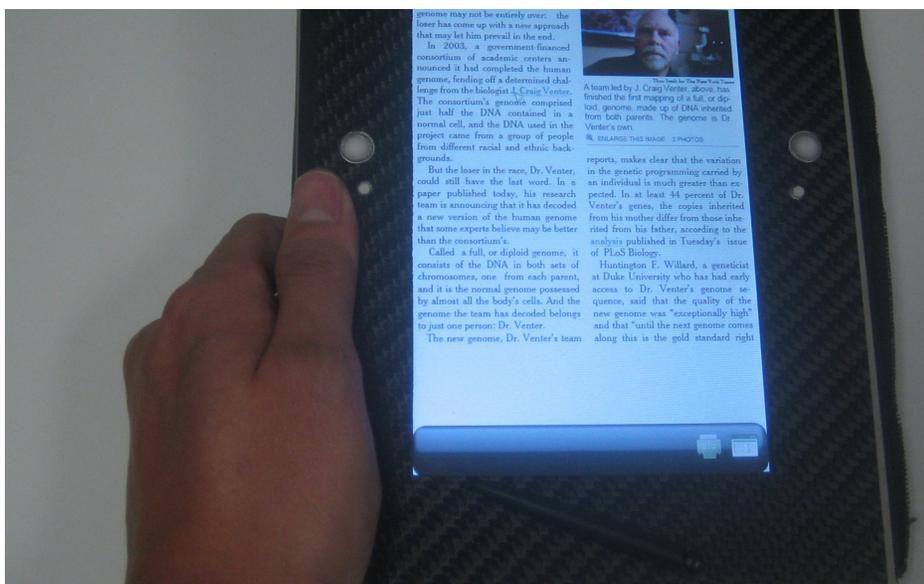


Figure 5 Close up view of one display of the Dual-Display E-book. Trackballs and menu buttons are to the right of the thumb.

movement [Karlson, Bederson, and Contreras-Vidal 2008]. Consequently, the 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 click 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 were possible, such as touch-sensitive strips [Buxton and Myers 1986; Harrison et al. 1998] but trackballs offered a good compromise between ease of use and ease of implementation.

3.2.2 Improving global navigation

The ability to quickly jump to a specific section of a document is a key feature of paper documents. Although a table of contents and hyperlinks in electronic devices address some of the issues related to global navigation [Card et al. 2004], visual navigation remains important. In electronic devices, the predominant method for performing visual navigation is scrolling. Relying solely on scrolling or related techniques [de Bruijn and Spence 2000; Igarashi and Hinckley 2000; Sun and Guimbretière 2005] can hinder users' abilities to leverage spatial memory about paginated documents [O'Hara, Sellen, and Bentley 1999]. Therefore, the Dual-Display E-book employs Space Filling Thumbnails (SFT) [Cockburn, Gutwin, and Alexander 2006] to provide a global overview of the document being viewed.

In SFT, page thumbnails of the entire document are tiled across both displays (Figure 6). SFT takes advantage of the dual-display configuration by leveraging the added screen space of the second display to show thumbnails at a larger size. With the current device prototype the about 70 pages of graphic-heavy content can be comfortably resolved 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 154×205 pixels. SFT also has an additional advantage of requiring comparatively fewer display refresh operations than alternative navigation techniques like scrolling. The practical importance of this behavior is that SFT is better suited for use with contemporary devices with electronic paper displays, which generally have a slower refresh rate. However, in future applications where bi-stable displays (screens where power is only consumed when the image displayed changes) are employed, the fewer refreshes SFT requires would provide power consumption savings.

Bookmarking is another important feature for supporting revisitation [Schilit, Price, and Golovchinsky 1998; Marshall et al. 2001; Card et al. 2004]. Here too, the use of a dual-display configuration can significantly improve the navigation experience. Upon activation



Figure 6 Space filling thumbnail navigation on the dual-screen configuration.

of the bookmark function, one screen displays the bookmarked page (and a list of other bookmarked pages) while the other screen continues to show the last position the user was at in the document. The use of two screens allows the bookmarked page to be viewed simultaneously with an arbitrary page in the document, which enhances sense making [Card et al. 2004].

3.2.3 Improving inter-document navigation

Certain styles of reading are characterized by a significant amount of cross-document navigation and involve cross-referencing and comparisons across multiple documents. Although a single screen can show multiple windows of information, most screens—especially those on mobile devices—can show only one document at a time comfortably. In this situation, cross document operations must be performed serially, which is cognitively demanding, slow, and frustrates users [O’Hara and Sellen 1997]. 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 prevents users from arranging the screens in space in way that better supports the task at hand.

To support inter-document interactions, the slates that form the faces of the Dual-Display E-book can operate independently (Figure 7). Upon separating the two faces of the device, each face can show different documents or different sections of the same document. Since each face provides independent controls, the system can support the bimanual interaction and the parallelism commonly observed with paper [O’Hara and Sellen 1997]. When the faces are reconnected, the e-book returns to showing a single document on the two screens. The device can be configured to always use a particular face as the “master” or prompt the user to pick the face to determine what to show across the two screens.



Figure 7 Two sides of the Dual-display operating independently and showing two different documents.

3.3 Iterative Device Design

In developing the Dual-Display E-book we used an iterative design approach to gradually refine the device. We began with paper-based prototypes to evaluate control placement. We then built electronically instrumented mockup devices to evaluate the feasibility of different embodied interactions. Finally, we created two working prototypes of the device.

3.3.1 First Prototype

To simplify the electronics design, our first prototype (Figure 8, left) did not include power management and was powered by 4 AAA batteries. With the batteries installed, each slate was heavier (520g) and thicker (18mm) than our initial targets. Also, the system could only run for about 30 minutes per battery charge. Nevertheless, this prototype was useful to convey our interface concept to users. We implemented embodied gestures using accelerometers to sense motion.

Using this device, we conducted a preliminary feedback session. The evaluation consisted of a semi-directed interview format lasting 25 minutes. We alternated between



Figure 8 First (left) and second (right) prototypes of the Dual-Display E-book.

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.

3.3.2 Second Prototype

We developed a subsequent prototype (Figure 8, right) that addressed the issues raised by our first set of users. This prototype uses lithium-ion 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 about 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 found the bookmarking feature to be too heavyweight for random access navigation in the first prototype, we did not include it in the second prototype.

3.3.3 Prototype Implementation Details

Both prototypes were 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. A PIC micro-controller manages data from the trackball, command button, face attachment switch, and motion sensors and forwards the information to the Gumstix.

Sensor data, and device state are shared between the faces using a Bluetooth Personal Area Network, to a context server [Hinckley 2003]. The context server can either run on a separate computer or 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 uses the information to determine what page should appear and communicates that information to the display application running on each side of the device. In the first prototype each face of the device ran the Windows CE 5.0 operating system, and all software was written in C#. In the second iteration, we switched the operating system to Linux to take advantage Linux's improved support of the Gumstix platform. In doing so we rewrote the software in C++ to use the Qtopia embedded development library. We also rewrote the context server in C++ so that it can run on any POSIX compliant operating system.

3.4 Evaluation

The evaluation of the Dual-Display E-book compared the Dual-Display E-book with a single-display device in a typical newspaper or magazine reading task. The goal of the evaluation was to test whether the embodied navigation functionality and the addition of the second display enhanced local navigation. Also the evaluation examined whether the use of Space Filling Thumbnails [Cockburn, Gutwin, and Alexander 2006] across two screens improved global navigation tasks.

For the evaluation, the single-display reading device was realized using the detached face of the dual-display reader, without any embodied interface. A direct comparison of our dual-display device to a commercial e-book would have been difficult, as these devices have 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. In the experiment, we first demonstrated the controls on either the single or dual-face device, and then had the users try using the device on their own. After users had familiarized themselves with the controls, we had them read a series of newspaper articles uninterrupted for 10 minutes using each 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. The interview covered topics like 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). Then, we allowed users to try the feature, and asked them whether such a feature was applicable and useful to the reading activities they personally performed.

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 the University of Maryland. 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.

3.5 Results

3.5.1 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 screens forced her to periodically turn her head as she read. The motion of her head helped to alleviate the neck-strain she reported when reading from a single screen, or when reading on the computer. Many of the advantages attributed to two displays mirror the examples of lightweight navigation that Marshall and Bly [2005] 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’ chief source of confusion stemmed from being unclear about on which face the next page would appear. The problem was especially pronounced when both faces changed images at the same time, as there was no common point of reference. However, all thought having page numbers visible 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 [1998] finding that more than half of the tasks they observed people performing required operating on two or more separate display surfaces.

3.5.2 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 the effort required to fan a face was greater than that of the standard controls, especially when traversing multiple pages. Moreover, the fanning gesture did not provide any marked advantages over the standard controls.

For the flipping gesture, even with the 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 the Dual-Display E-book, 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.

3.5.3 Thumbnail-based Navigation

There was a wide range of responses to the SFT navigation system. While nearly all the users agreed that the thumbnails 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 notice that the thumbnails were larger in the dual-display configuration (the result of distributing thumbnails across more screen space). 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.

3.5.4 Size, Weight, and Mobility

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, saying that both the single and dual-face configurations were superior to reading from the computer because the device could be placed flat on the table. As a result, P4 was able to assume her preferred reading position of looking down at the table. All users 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.

3.6 Lessons from the Dual-Display E-book

3.6.1 Desirability of Multiple Screens

Our results indicate that including two displays in an electronic reading device has the potential for supporting reading activities in similar ways to paper. Users cited various benefits in the areas of local navigation and applicability to multi-document interactions when using two displays.

Many of these improvements might also be realized through a single, large display capable of displaying several pages simultaneously. We believe, however, that having separate displays is still 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, in order to read multiple documents, especially those of different aspect ratios, a device with a single large screen would need to be considerably larger than a dual-display device. Furthermore, a single-display device prevents users from laying out electronic document alongside paper documents or books in a customizable manner, such as when grouping relevant documents together. The studies we conduct with a multi-slate reading system comparing a single large display with multiple small displays go on to test this hypothesis.

3.6.2 Embodied Interfaces

We initially believed that the use of embodied user interfaces—interfaces that treat the body of a portable device as part of its user interface [Fishkin et al. 2000]—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. This result is an important reminder that blindly mimicking how we read on paper cannot be relied upon to produce desirable solutions.

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 the Dual-Display E-book. Most of our users agreed that if the device were lighter, the interactions might be more palatable. But since electronic reading devices 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.

3.6.3 Dedicated Reading Device versus Computer

One valid question is whether dedicated electronic readers are necessary, given the availability of more general devices like laptop computers, which are also portable and can do more. Some researchers have advocated the use of a “document laptop” that includes some functionality in support of reading [Marshall et al. 2001]. However, our users’ feedback about the greater comfort and flexibility of reading from the Dual-Display E-book argues in favor of a specialized device. Specialized reading devices might not provide all the capabilities found on a computer, but the reduced size and weight would positively affect the reading experience.

To maximize the utility of a specialized reading device, methods to better integrate with the existing electronic and paper tools people use are required. Technology trials in the real world that we later conduct confirm this need for superior integration. The United States system that we subsequently created includes features that ensure more integrated operation with existing electronic devices like PCs.

3.7 Summary

We explored various affordances of a two-display electronic reader as they relate to navigation tasks. To accomplish the 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 fold or detach the two faces. We conducted an evaluation of the Dual-Display E-book to understand what the advantages and disadvantages of a dual-display system were compared to a single-display system. The study found that advantages of the second display include better support of lightweight navigation, which made locating content to read, as well as re-reading easier.

The use of multiple screens in the multi-slate configuration obviously has its roots in the Dual-Display E-book. However, other lessons from the Dual-Display E-book pervade the remainder of this dissertation. For example, the embodied interactions we implemented for the Dual-Display E-book generally were not perceived to aid reading. As such, they reinforced the idea that the process of moving interactions possible with paper to the digital domain involves more than merely emulating these interactions. In later work, we draw from these lessons and avoid copying the interactions on paper. Instead, we take inspiration from these paper interactions and transpose them into a form that is friendly to the new domain. Also, the general approach of using devices that are individually functional and give the option of being linked is an important and beneficial characteristic of the multi-slate reading configuration.

3.8 Epilogue

Since our initial work on this device, a number of dual-screen reading devices have appeared as research prototypes, product concepts, and commercially available products. Since the initial publication of our work, similar devices have been introduced. The Codex [Hinckley et al. 2009], which explored interactions on and different poses for dual-display devices. The Microsoft Courier concept showed how a dual-screen device would support note taking and reading. The Entourage Edge, and KNO reader are commercial products

that were released and which have been geared towards enhancing reading in the classroom environment. Lastly, the Toshiba libretto W100 and Acer Iconia 6120 are commercial laptops that have also adopted a dual-screen approach. We believe the interest and activity in this domain confirm the notion that a dual-screen configuration offers an attractive balance between increased display area and a portable form factor.

4 Multi-Slate Reading Platform

After our work in supporting navigation using the Dual-Display E-book, we set our sights on supporting the other facets of reading activities using a similar strategy of employing several electronically linked displays. A key lesson we learned from our experiences with the Dual-Display E-book was that the physical properties of the device (e.g., size, weight, screen size) influenced the reading experience. Therefore, supporting these other reading activities would need to start at the hardware level. We begin the presentation of the multi-slate reading system with this chapter, which concentrates on the hardware platform over which the remainder of the reading system is created.

4.1 Introduction

Many computing tasks are agnostic to the physical form of the computer because most modern GUIs rely on indirect input through the keyboard and mouse. Reading is not one of these activities. For instance, reading involves a great deal of physicality: users find it helpful to hold up their reading material and to write on it, among other physical interactions [Marshall and Ruotolo 2002; Marshall 2005]. As a result, the hardware plays a large role in determining the reading experience. If these low-level physical aspects of reading are not supported, then the more involved types of reading—the activities we specifically wish to support—also suffer.

Although we successfully used The Dual-Display E-book for laboratory studies, there were limits to applying it to a more general set of reading activities. The Dual-Display E-book did not have provisions for precise writing and its screens were too small to display many types of documents that are distributed on paper. Moreover, the Dual-Display E-book prototypes were quite fragile and had limited battery life.

To address these issues with the Dual-Display E-book, we sought out a new hardware platform. We first looked towards modifying commercial devices but found our options limited. When we started the project in 2009 the devices we could choose from were TabletPCs, PDAs and E-book Readers. PDAs and TabletPCs were attractive because they provided stylus-based support for writing out of the box. Unfortunately, PDAs were too small, which affects how much content one could see at a time, and TabletPCs, too bulky, which reduces mobility. Since these were hardware properties, they would have been difficult to work around. E-book appliances had the right form factor, especially the Amazon Kindle DX, which was announced at the same time we started the project. The problem with E-book appliances was their limited options for extensibility in software and hardware. For instance, there was no end-user development support and also no way to augment the device with new electronics. We decided that our best option was to develop a new reading device platform around which we would build the multi-slate reading system.

We partnered with Ricoh Innovations, Inc., who was interested in developing an electronic paper-based slate device at the same time as we were, to jointly develop a new device platform that enabled reading and writing. The platform development consisted of the mechanical design, electronics, and system software for the device platform. We provide details about each of those aspects in this chapter. These implementation details are of interest because it has recently become feasible for small research groups (like those in academia) to prototype new device platforms, as we have done. We believe our efforts can serve as a guide for how to build and test more realistic and robust devices in HCI research.

4.2 Design Rationale and Specifications

When determining what we would use to implement the component slates for the multi-slate reading system, we looked towards the positive qualities of paper for guidance. Each sheet of paper has the combination of high micro and macro mobility, high readability, ability to show a full page of information, and support for freeform annotation and writing.

We looked for a hardware platform that offered the combination of the above four characteristics. What we discovered was that devices at the time possessed some of these characteristics, but not all at once. PDAs offered high mobility and support for annotation, but could not show very much content; E-book readers had good mobility characteristics and readable screens, but was missing support for pen input⁵; TabletPCs allowed for pen input, but lacked mobility. When we decided that no suitable device could be purchased, we set out to create our own slate device that had the desired combination of mobility, readability, screen size and writing support. Since we would be specifying the hardware components, it was possible to draw up a “wish-list” of features for this device.

Mobility - The mobility aspect of the hardware was perhaps the most important design goal. Based on our previous experiences with the Dual-Display E-book, we knew size and weight were crucial factors in the overall usability of the device. Previous feedback suggested a target thickness of 6 mm and weight of 200g per surface to ensure it could be easily handled. Since one eventual goal was to deploy the slates outside of the laboratory, the construction needed to be able to withstand the rigors of real world use. Lastly, we wanted to make sure we did not compromise the mobility of the device by requiring it to be tethered to a wall outlet. Therefore, extended operation using batteries was also critical.

⁵ Certain devices did support pen input, such as those from iRex. However, these devices had an additional problem of having nearly no end-user development support.

Screen-size and Readability - The display on the device needed to accommodate standard paper documents. Repeated panning and scrolling due to the display being too small is inefficient and annoying. Since we did not want to arbitrarily limit the types of materials that could be used on the system, a screen that was large enough to show a Letter/A4 page was preferable. The choice of display technology was a toss-up. Electronic paper displays (EPDs) offer superior readability but LCDs provide better interactivity. In the end, several factors contributed to this choice. First, LCDs of the desired size that could be acquired at the time were considerably thicker and heavier than comparable EPDs. Second, promising developments in the area of accelerating EPD refresh rates [Rhodes 2009] offered the possibility of narrowing the speed differences between EPD and LCD screens. In particular, the technology provided faster page flipping and responsive inking on EPD screens.

Writing - In order to support writing while reading, we needed to provide the capability to ink electronically on the device. The best-in-class electronic inking experience is currently provided through inductive pen digitizers like those found in graphics drawing tablets and TabletPCs. We alternatively considered the use of capacitive overlays that would allow touch sensing as well. Touch overlays were deemed impractical because the reduced optical transparency would lower the contrast of the non-backlit EPD and adversely affect readability.

In addition to the above concerns influenced by the properties of paper, we also knew that wireless connectivity between slates would be necessary to ensure data would not be isolated on a device. Although we were somewhat agnostic about the specific details of the connectivity, we considered our options nonetheless. In the Dual-Display E-book the connectivity was provided via Bluetooth, and we had good experiences with the

technology. However, generic connectivity to more devices can be achieved with Wi-Fi and the Internet. An ideal solution would be to have both types of connectivity.

4.3 Slate Hardware

The hardware for the multi-slate system was developed jointly with Ricoh Innovations, Inc. (RII). RII had a central role in the mechanical design for the devices and provided significant expertise in the design of the core electronics. My roles in the creation of the hardware were specifying mechanical requirements, performing the electronics design of various portions of the core system and the performing the complete electronics design and firmware development for the input-sensing module. The specifications of the slate devices is summarized in Table 3 and detailed below. The slate devices largely satisfied the requirements described above.

4.3.1 Key Components

Computation - The slate devices are powered using a Gumstix Overo Air computer-on-module. The Overo Air modules feature a Texas Instruments OMAP3503 embedded application processor, which integrates an ARM Cortex-A8 class CPU core with on-board peripherals like an LCD controller. The Overo Air also integrates on-board 802.11b/g Wi-Fi and Bluetooth wireless radios. The Gumstix was deemed the most suitable module because of its small size relative to other available modules and built-in wireless connectivity. The Overo Air module also provides signal lines for serial communications, Serial Peripheral Interface Bus (SPI), and USB to allow the module to interface with other electronic components and hardware devices.

Display – The slate hardware uses 9.7” E-Ink electronic paper displays (EPD). With a resolution of 150 dpi, the EPD offered readability nearly matching paper and considerably lower power consumption compared to standard LCD displays. Furthermore, the use of an EPD helped reduce the thickness and weight of the device by virtue of a thinner screen and

smaller battery requirements. We operate the display using an E-Ink Metronome EPD controller chip. The Metronome display controller chip receives 8-bit per pixel display data over the standard LCD output from the Overo, along with special control signals at the start of each frame of video data. Based on the image data provided to the Metronome chip, the Metronome computes the appropriate voltages to apply across each pixel of the EPD to produce the desired image on the display.

One tradeoff the use of an EPD entailed was that it limited the types of interactive visualizations of the document that could be used. Interface elements that are normally possible with LCD screens would not be possible due to the 250ms to 750ms (depending on the size and bit-depth of the image that needed to be drawn) required to refresh the EPD. These limitations were alleviated somewhat through the use of specialized algorithms for updating the EPD [Feng et al. 2008; Rhodes 2009]. The responsive inking that these algorithms made possible were a critical aspect of the slate devices. We also took advantage of the fast page flipping technology in initial prototypes of the slate devices, which demonstrated the significant improvement over what was possible with the stock Metronome controller. However, complexities associated with pre-processing document files to make them compatible with fast page flipping prevented us from using it for later deployments of the slate devices.

Pen Input - The EPD is overlaid over a Wacom inductive pen digitizer backplane. The pen digitizer is identical to those used in TabletPCs and provides performance in terms of tracking accuracy and resolution (2000 dpi). The pen digitizer requires a special inductively powered stylus to detect writing. One of the advantages of the inductive digitizer over other touch/writing technologies is that contact with the display from the user's hand is not reported as an input event. Since people write with their palm resting on the writing surface, the inductive digitizer allows for a more natural and ergonomic writing experience.

Input and Sensing –When designing the device, we knew that the lack of a touch screen would limit the expressiveness of the input, particularly for interactions where continuous adjustment would be beneficial (e.g., navigation). In response, we wanted to explore whether a touch-sensitive bezel could be used for these types of tasks. Since our collaborators did not share this particular design goal, we designed the main board so that it could be extended with a variety of input schemes. To support our goal of bezel-based touch sensing, we developed a custom input circuit board to detect touch while also providing standard tactile buttons. We provide details about this auxiliary board in Section 4.3.3. The devices also provided multiple sensors to match the functionality of the Dual-Display E-book reader. These sensors include accelerometer and gyroscopic motion sensors to detect embodied interactions and Hall-effect magnetic sensors to detect when slates are against each other. Lastly, to help characterize the power consumption of the device, we placed several current-sensing Coulomb counters in several circuits.

Device Specifications

Size and Weight	Height: 270 mm Width: 205 mm Depth: 9.90 mm
Storage	8GB (expandable via Micro-SD)
Wireless	Wi-Fi (802.11 b/g), Bluetooth
Display	9.7-inch (diagonal) E-Ink Vizplex Electronic Paper Display (EPD), providing 825×1200 resolution at 150 ppi
Processor	600 MHz TI OMAP3503 on a Gumstix Overo Air Computer-on-Module
Power and Battery	11-watt-hour rechargeable lithium-polymer battery
Input	Wacom pen digitizer, capacitive input overlay over tactile buttons
Sensors	Three-axis gyro, Accelerometer, Ambient Light Sensor,

Table 3 Technical specifications for the slate hardware platform.

4.3.2 Main Electronics Architecture

The electronics architecture is can be divided into four major subsystems. The graphical summary of how these subsystems connect together is shown in (Figure 9). Below, we provide insight into the design of these subsystems.

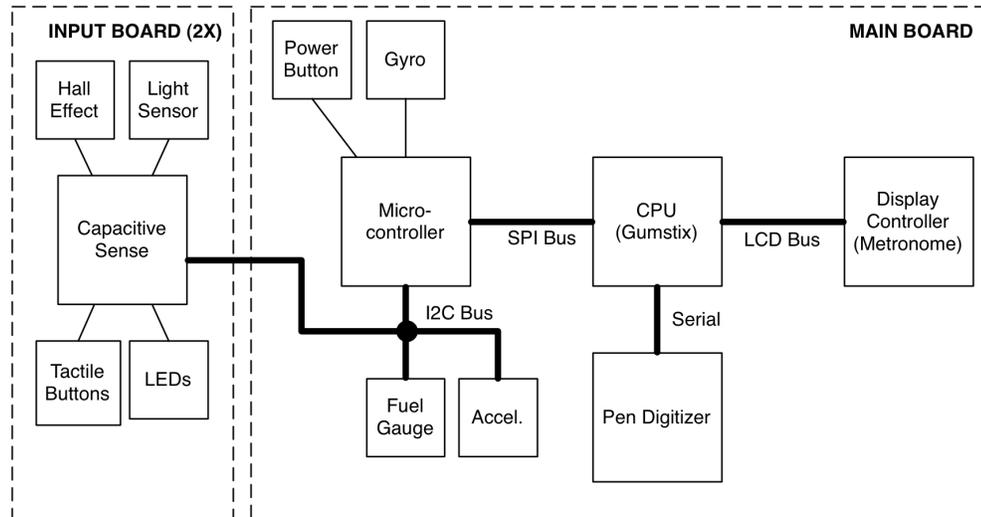


Figure 9 Block diagram of electronic components.

Display Subsystem - The display subsystem uses the LCD controller and bus lines on the Gumstix to communicate with the Metronome display controller. In addition to the data and clock signals, the Metronome display controller requires out-of-band control lines to specify standby and reset conditions, and to indicate when it is ready to accept new data.

Power - The slate electronics required a variety of operating voltages. The Gumstix module is powered at 1.8V, the display controller and micro-controller are 3.3V parts, the USB ports output 5V and the EPD requires +15V, +22V, -20V and -15V lines. Complicating matters was the fact that the microcontroller always had to be powered and that lithium batteries can swing between 2.7V and 4.2V depending on remaining capacity. We show the DC/DC conversion architecture we employed in Figure 10. This solution allows for efficient generation of the always-on 3.3V line that is able to use the entire capacity of the battery. The 1.8V supply is generated from the 3.3V supply to minimize conversion losses. For other parts of the circuit that require more exotic voltages we provide a separate 5V supply from which to generate these other voltages.

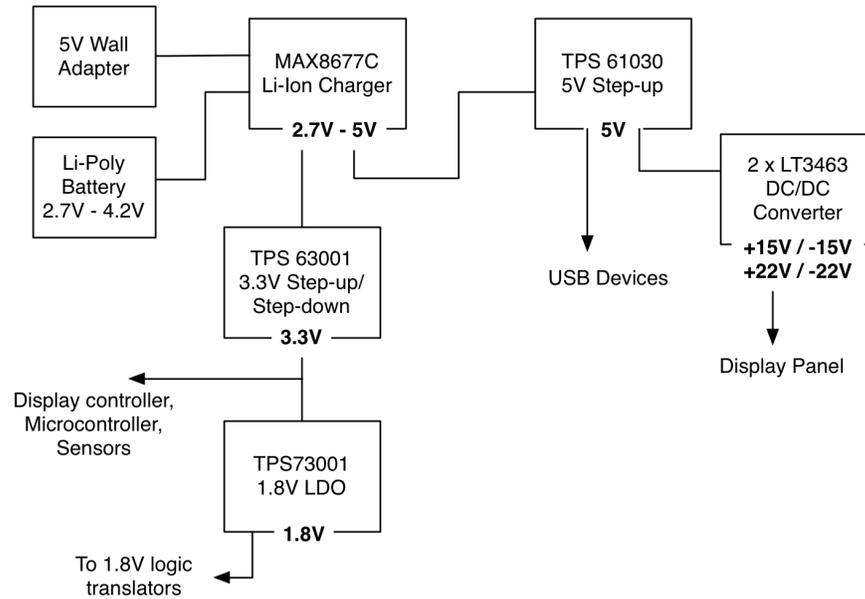


Figure 10 DC/DC conversion stages in the slate hardware.

Battery and Power Management – The slates use off-the-shelf 1500mAH Lithium Polymer rechargeable batteries. The specific packs we use are aftermarket replacement packs for the second-generation iPod. The battery charging is managed by a MAX8677C Battery Charge controller to ensure the correct charging current and voltages and also a DS1337 Battery Fuel Gauge to track battery charge and estimate remaining capacity. The Fuel Gauge can be read via i2C while the charge controller status is communicated over discrete signal lines.

Microcontroller and Hardware Input - All input and sensor data is first routed to an always-on TI MSP430F5436 micro-controller. The micro-controller periodically polls devices over the I2C bus as well as monitors signal lines connected to on-board sensors. The benefit of this architecture is to ensure that system monitoring can still be performed even if the Gumstix module is powered down or suffers from a software fault. For instance, this configuration allows the power switch to turn the device on and off and for it to act as an out-of-band method to reset the Gumstix (via a long press). The micro-controller is also responsible for monitoring the battery charge status and can initiate shutdown procedures

should the battery drop to critical levels to prevent data loss. When the Gumstix is powered up, it actively queries the MSP for sensor and input data from the micro-controller over the SPI bus (the Gumstix acts as the SPI master). When the Gumstix is powered down, the micro-controller enters a power-savings mode and only periodically wakes to check the control and sensor status. The microcontroller is also responsible for powering on and off nearly all of the components on the main board, allowing fine-grained control over what subsystems are active.

Pen Input - The digitizer electronics are self-contained and transmit pen input with the Gumstix module over a standard asynchronous serial communications link to one of the Gumstix's 3 on-board UARTs.

4.3.3 I/O Control Strips

As described above, the main system could be extended using custom input circuit boards to offer alternative control schemes. For the slates, we developed one such control scheme that combined the use of capacitive touch sensing with tactile switches. The combination of capacitive elements with tactile switches allows for continuous value adjustments along with tactile confirmation of input events. The system also enables richer interactions, like having buttons that respond differently depending on where they are pressed, as well as gesture-based interactions on the surface of the button.

Since we wanted to sense touch at the surface of the bezel, the sensing electronics needed to be placed close to the upper surface of the device. However, spatial constraints meant that the components on the main board needed to be mounted facing the bottom of the device. To harmonize between these two requirements, we used a flexible PCB ("flex") to route between the two planes (Figure 11). For reasons of cost, we designed the flex so that the single design could be mounted on both sides of the device (by flipping it upside down).

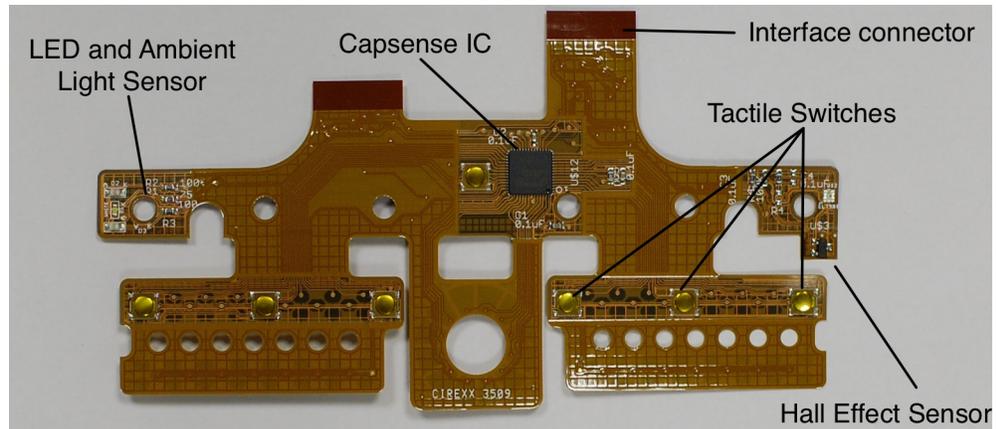


Figure 11 Flexible PCB for touch sensing.

The flex is fabricated with capacitive electrodes on the surfaces where it rests on top of the bezel. Touch sensing is accomplished using a Cypress CY8C24894 capacitive sensing controller. Also on the flex were LEDs to provide users with feedback about power and battery charging states, Hall-effect sensors to detect slate proximity, and an ambient light sensor that could be used to automatically put the device to sleep or to wake it up.

4.3.4 Enclosure Design

The mechanical design of the enclosure had to satisfy multiple requirements simultaneously. In addition to holding all of the electronics components inside the casing, the design needed to match the thickness and weight goals, be ergonomic to hold, and to support the input scheme we selected. We prototyped various bezel sizes early on in the design process by affixing laser-cut acrylic sheets (that represented the bezel) over foam board pieces that approximated the expected thickness of the devices. We eventually settled on a bezel width of 30 mm, which offered a good balance of providing a graspable region that did not obscure the display while not making the device footprint too large.

We added much-needed structural integrity to the fragile and brittle glass EPD by placing it in a tight stack against the rigid PCB, with the digitizer backplane sandwiched in the middle (Figure 12). To minimize device thickness, the Gumstix module was moved to a



Figure 12 Internal view of the slate device. Note the two flex input boards along the sides and Gumstix carrier board at the top. The orange squares along the bottom edge are the radio antennas.

second PCB connected with flexible cabling so that the module could be placed co-planar with the display. Figure 12 also shows how the interior of the slate accommodates the batteries and radio antennas.

Figure 13 depicts finished enclosure of the device. The case construction was done using cast polyurethane, with molds generated from a 3-D printed master. The input sensing flex PCBs roll up along the sides of the device. The Hall-effect sensors are also placed on the flexible input PCB. The casing accommodates magnets in complementary positions on the other side of the slate so that when two slates are brought together, the magnet lines up with and triggers the Hall-effect sensor.

The front view (Figure 13, left) of the casing shows the control layout. The controls on the device include two large round buttons in the center plus four additional tactile buttons above and below the round buttons. The input modules consist of a polycarbonate self-adhesive label glued over the capacitive sensing PCBs. The capacitive sensing electrodes are able to detect finger contact through the polycarbonate label. A clear window on the



Figure 13 Front and back views of the device enclosure.

label allows the user to see the feedback LEDs on the input board. We placed controls on both sides of the device to support a variety of grips and reading positions.

In the rear view (Figure 13, right) of the system, one can see the docking points for the magnetic latching system we designed to bind two slates together. The docking points each hold a strong neodymium magnet. The magnetic latching system was designed to allow the slates to be connected in a side-by-side orientation, or to enable one slate to fold behind the other. The latching straps, which also have a magnet at the end, fold back into the cutout on the back of the device for storage, or mate with the cutout on an adjacent device. When two slates are attached, four latching straps in total hold them together.

4.4 System Software

4.4.1 Operating System

The Gumstix computer module we use has a standard ARM-based processor, for which several standard distributions of Linux are available. The slates run a minimal version of Debian Linux and the system operates much like any other Linux computer. We try to use

standard Linux utilities as much as possible for controlling system services like networking. The major areas of customization we made were the addition of device drivers to use the specialized electronics we included in the system, system startup and shutdown scripts, and background monitoring daemons.

4.4.2 Device Drivers

To enable applications in user space to interface with the custom hardware on the devices we created a set of device drivers for the hardware. The three components that we needed to interface with were the Wacom digitizer, the graphics controller and the micro-controller (Figure 14).

Pen Input

Pen input is routed to a proprietary Ricoh Innovations, Inc. (RII) inking and graphics-processing software layer [Feng et al. 2008], the specific implementation details of which we were not privy. The software decodes digitizer data packets and makes it available at the application level over local UNIX sockets. The proprietary layer also computes the necessary updates to the display from the digitizer data so that ink strokes appear on the EPD with lower latency (about 250ms).

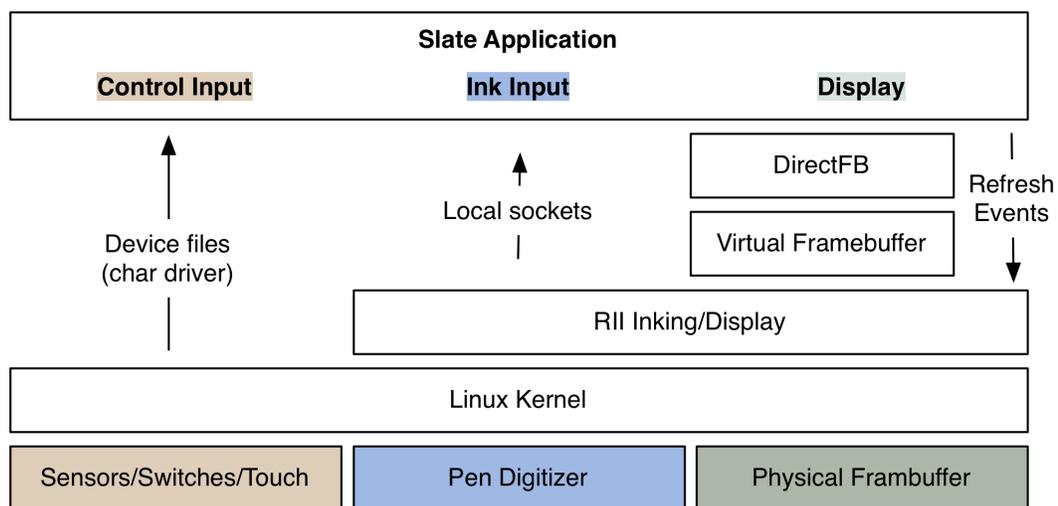


Figure 14 System software stack on the slate device platform.

Graphics Controller

From the Gumstix's standpoint, the Metronome display controller receives input from the LCD signal lines and can be treated like a standard LCD panel. However, the controller differs from a standard LCD panel by requiring special commands to be sent in-line with the image data on screen. To accomplish this, a software middle layer needs to intercept the graphics data that would normally be drawn to the framebuffer, append the control data at the start of the frame, and then write the resulting data to the physical framebuffer. To intercept the graphics data, we developed a Linux framebuffer driver that creates a virtual framebuffer device that Linux graphics libraries know how to draw to. The virtual framebuffer then notifies the RII inking/graphics layer, which processes the graphics data, adds the control data and writes the result to the physical framebuffer.

Microcontroller

The microcontroller/input driver pulls sensor and input device data from the microcontroller. The microcontroller driver queries the microcontroller over SPI for the requested data and makes the data available through standard Linux character devices. Since we use standard character devices, software in userspace can read data and send commands to the microcontroller using the same mechanisms for reading and writing to files.

4.4.3 System Scripts

At startup, the OS runs a startup script that powers up the different subsystems on the device and starts a system-monitoring daemon that periodically queries the microcontroller for information like battery state and power button pushes. One of the advantages of exposing the microcontroller (and by extension the sensor data, battery charge status and input state) via character devices was that standard UNIX utilities could be used to monitor the status of the system. Thus, the monitoring daemon is implemented as a looping UNIX shell script.

If the power switch is pressed or if the battery drops to a critical level, the daemon sends a message over the system message bus (DBus) to warn running applications of impending shutdown. When this occurs, applications have a window to save out important data. Then, the system begins the system shutdown process.

When the shutdown process begins, a special shutdown script informs the microcontroller that the slate is entering the shutdown state. At the end of the shutdown sequence the script asserts a signal pin to inform the microcontroller that it is safe to cut power to the computer module. In the event the shutdown process hangs, the system will not send the shutdown complete signal. Therefore, the microcontroller is programmed to wait 10s for the assertion, if it does not see the assertion in that time window, it will assume the system has hung and cut power anyway.

4.4.4 Application Development

The slate platform includes an application and graphics stack to speed up the process of developing applications. At the bottom level, the slate platform uses a lightweight DirectFB graphics layer to simplify drawing to the virtual framebuffer. DirectFB has fewer dependencies and lower overhead than a more generic graphics layer like X. A number of Linux GUI toolkits like GTK+ and Qt can be compiled to output to DirectFB rather than X so the use of DirectFB is mostly transparent. The system leaves it up to the developer to choose what specific GUI toolkits to employ for application development. The one unconventional thing the developer needs to do when creating a program is to add a hook to the screen invalidation and refresh functions so that the inking/display layer knows when to redraw the screen.

The slate platform provides a device simulator for programs developed using the cross-platform GTK+ toolkit. One of the advantages of using a commonly available cross-platform toolkit like GTK+ is that we were able to quickly create a device simulator that

runs on a PC running Linux. Our simulator works by running the program that needs to be simulated against the version of GTK+ built for the PC. On the PC, the application runs in a window and input comes from the mouse and keyboard instead of the stylus and hardware buttons. Other aspects of the program behave the same. Creating the PC based device simulator increased development speed by allowing us to program, test and debug primarily on the PC, reducing the time for an implementation and debug cycle.

4.5 Design Discussion and Assessment

The final hardware platform largely satisfied our design goals. The design of the enclosure and structural rigidity requirements resulted in a slate that was 10 mm thick. The resulting weight of the slate was approximately 500g. Although these were higher than our ideal targets, they still compare favorably to contemporary devices like the Apple iPad (the slates are lighter) and Amazon Kindle DX (similar size and weight), both of which have comparably sized displays. The subsequent introduction of consumer tablet appliances like the iPad has addressed some of the factors that led us to create custom hardware. These devices now provide development tools and hardware expansion capabilities. For reasons of weight and missing support of high-precision writing, these tablet appliances are still not perfect for reading activities. However, for certain activities, particular those where writing is not as essential, commercial tablet appliances offer a good alternative to building custom hardware.

A major weakness in our slate hardware is power management. Linux kernel support for power management on the slates is not yet robust and therefore the slates cannot enter a low power sleep mode. The main consequence is that battery life was not as long as our designs had predicted. We are able to strike a balance and achieve our full day runtime if the slates shut themselves off after a long period of inactivity (currently set to 45 minutes). The downside is that if the slate does turn off, there is a lengthy start up process where the

slate boots up and the operating system starts. The issues with power management also prompted us to add the optional second lithium polymer battery, which increased the weight of the slate by about 30g.

One other design element that did not turn out as successfully as planned was the magnetic latching system. While the magnets were strong enough to hold the straps in place, we did not foresee that the super glue used to attach the magnets to the enclosure would be the point of failure. As a result, we largely avoided the use of the latching system to reduce the possibility of devices unexpectedly disconnecting.

There were some concerns regarding the proprietary nature of the RII graphics and inking layer. The RII layer was provided as a black box with implementation details omitted, which could negatively impact the reproducibility of our work. However, we believe that the rate at which EPD display technology is advancing makes these concerns less serious. The major roles the RII layer played was to address limitations of the Metronome chip; specifically de-ghosting (removing previously displayed images on the EPD), fast drawing (for ink) and partial screen refresh. Newer EPD display controller chips implement the functionality in hardware and the system can be reproduced using these chips without the need for the proprietary ink/graphics layer. Therefore, the proprietary RII layer, while important to the current slate platform, is not in the critical path of implementing similar hardware.

4.6 Summary

In this chapter we described the design goals going into our decision to create a custom hardware platform to support reading. Our decision to create a new hardware platform was driven by the dearth of hardware options when we initially started the project. Devices available at the time had undesirable physical properties that would have been difficult to work around. The key elements our believed our reading devices needed to support, which

existing devices did not, were high physical mobility, readability, ability to capture writing, approximately letter-size screen, and facilities to communicate with each other. We were successful in satisfying these design goals, and had unique access to a hardware platform customized to support reading requirements.

5 The United Slates

A set of lightweight and readable slates that can also capture electronic writing to users provides functionality that is unmatched in existing devices. But this hardware in itself is not sufficient. In this chapter, we present a unified set of interaction techniques to maximize the potential of the unique hardware platform we developed, while streamlining the management of an ensemble of slate devices.

5.1 Introduction

It would appear that providing users with more display area, in the form of additional slates, would be enough to address many of the pressing requirements of reading. Although having additional displays is an important first step, it alone does not completely address all of the requirements associated with active reading.

An additional challenge comes from the overhead associated with operating a multi-device configuration. With fully independent devices, there is the problem of each device acting as an isolated island of information. Users who use more than one computing device readily attest to the challenges of working across separate electronic devices [Oulasvirta and Sumari 2007]. Participants in Morris, Brush and Meyers' [2007] study complained that data was isolated on a single device while using multiple TabletPCs together. In settings with large amounts of display space, users also mentioned that they also did not have an easy way to efficiently make use of the available space. These problems arise because isolated devices do not take full advantage of digital affordances like electronic connectivity, which would allow the devices to better operate in unison.

In response, we created a comprehensive multi-slate reading system to address the dual concerns of lowering the barriers of working in a multi-device environment and providing capabilities in support of reading. To guide the design of the system and to test out key interaction ideas, we first ran a trial where a technology probe was deployed to students in an undergraduate class at the University of Maryland. We found that the physical attributes of the slates, such as the support for freeform inking, and the additional display to be beneficial. However, users in the probe reported difficulties in the areas of navigating between different documents, integrating with existing computing devices, and supporting varying numbers of devices.

On the basis of this trial, we went on to design the interactions that underpin the multi-slate reading system in more detail. The resulting system, United Slates, pairs the unique slate hardware platform with a set of UI conventions and interactions to more completely support the reading needs of knowledge workers and students when compared to existing electronic solutions. The interactions we designed for the United Slates system focuses on supporting navigation activities, resolving physically awkward slate interactions, and finding more integrated ways of dealing with documents within the context of the overall active reading workflow. Of particular significance is how the system demonstrates that new functionality can be derived from the electronic connectivity between devices.

We followed up by conducting an evaluation of the United Slates system by having participants compare using a set of slates that incorporated the multi-slate reading tools to isolated devices that did not have that functionality. Based on our results, we believe that the United Slates' ability to accelerate non-sequential navigation across many documents, to assist in working with large collections of documents, and to fluidly integrate with PCs, are critical improvements that make it a compelling system for knowledge workers and

students. The suitability of the United States system solidified its selection as the basis of the multi-slate environment we eventually deployed to real-world users.

5.2 Technology Probe

With the exception of Morris, Brush and Meyers' [2007] work, knowledge of the interactional needs for multi-slate reading comes solely from studies about what people do when working with paper. Since slates and paper differ significantly in terms of physical characteristics like weight and thickness, not to mention their interactional capabilities, users most likely would not use slates in the same way they use paper. To better understand how slates would be used for active reading in a real classroom environment, we conducted a 2.5-month long trial with a technology probe with 13 undergraduate English students in the Spring 2010 semester. The use of a technology probe reflects the fact that reading activities in the classroom are complex and can involve styles and preferences that vary across users. It was judged that the probe was superior to a more controlled approach because it would better capture the richness and diversity of how the slates would be used in the classroom.

In this trial, we first assessed students' reading activities with a single slate. Starting with a single slate gave us the opportunity to test for usability issues and bugs prior to deploying the more complex dual-slate configuration. Then, after 1.5-months, we provided a second slate and demonstrated additional features. During the course of the study the students were asked to use one or both devices to carry out the majority of their course readings. We conducted bi-weekly interviews with students to gather feedback. When possible, we modified the software on the slates to reflect the feedback.

5.2.1 Equipment and Interface for the Probe

The probe used the slate hardware presented in Chapter 4. Each slate displayed one full page of a document at a time, without margins, so that pages of letter-sized documents

could be viewed without scrolling. We pre-loaded the devices with course readings after consulting with the instructor. Students did not have the capability to place their own documents on the devices. For each document, we generated a page of Space Filling Thumbnails [Cockburn, Gutwin, and Alexander 2006] that users could tap on to jump to a specific location on the document. On each page, users could apply freeform ink on the pages with a stylus without entering any special inking modes. When more than one device was in use, ink marks made on one slate were transmitted over Bluetooth to the other slate so that the annotations on both slates remained consistent. The system also allowed students to create “notebook” documents, which were special documents with 50 blank pages. Users could ink on these pages and they could also copy and paste regions of pages into the notebook. The clipped regions linked back to the location from which they were extracted so that the user could view the context of the clipping. The clipping behavior was inspired by similar functionality found in research systems like InkSeine [Hinckley et al. 2007] and PapierCraft [Liao et al. 2008].

5.2.2 Findings From the Probe

Students were particularly enthusiastic about the portability aspects of the device, opting to carry the device in lieu of the course readings on paper. Students also mentioned that the form factor and weight of the device were often preferable to a book because it was easier to operate with a single hand and did not flop around. The free-form ink annotation also received a large number of positive comments. Students reported writing more on the electronic devices because they felt they were not “defacing” the original document due to the reversibility of electronic ink marks. The difference reported was particularly pronounced for books, which many students said they wanted to keep pristine. Furthermore, students mentioned that the rigidity of the device actually made annotating easier than on paper. In fact, the simple presence of a pen was a plus in itself; like readers in

O'Hara and Sellen's [1997] study, many students reported that reading with a pen was useful for helpful for maintaining one's place in the text. Based on this feedback, we were reasonably confident that the hardware design of the devices—the support for annotation, in particular—were in line with our goals.

Perhaps the biggest shortcoming of the system we deployed was the fact that the devices did not integrate well with the existing electronic devices students employed. Students desired a way to put their own documents on the device and also be able to move data off in support of their writing activities. Additionally, many of their regular daily activities revolved around the PC. Users often reiterated their desire to use the slates in conjunction with their computers. This result corroborates a large number of studies that emphasize the integral role PCs play in many reading tasks [Marshall et al. 2001; Morris, Brush, and Meyers 2007; Tashman and Edwards 2011a].

In addition, while users reported that receptive reading was satisfactory on the devices, many users ended up working on the paper versions of documents when they needed to complete writing assignments and review for exams. This was particularly pronounced for large books. Students' experiences working with one large text that was several hundred pages long proved particularly illustrative of the challenges they faced. For this particular book, we split up each chapter of the book into a different document in order to simplify access to each chapter. So while SFT could be used to navigate within the chapter, navigating across chapters required a lengthier process of changing documents. Student comments indicated that navigating with SFT was generally bearable. However, for writing and reviewing activities, students frequently needed to consult disparate information that was spread across different chapters. As a result, it was easier to perform that type of navigation for these tasks with the paper book. These comments highlight the fact that a multi-slate system needed to not only support rapid non-sequential navigation within a

single document, but also needed a strategy for providing equally quick access across different documents. The wide range in navigation practices we observed is similar to those that Thayer et al. [2011] found in their deployment of electronic reading devices.

Finally, even though we only provided users with two slates in the probe, we were able to observe interesting differences between how the students made use of the additional slate. The number of slates a user employed was heavily dependent on the context, task and individual preferences. For most tasks students used a single slate. The reasons users preferred a single slate include higher mobility while on the go and also the fact that many tasks did not benefit from a second slate. However a handful of students (3) found specific tasks where multiple slates were beneficial. These tasks were things like clipping out content into the electronic notebook, note taking using the second slate, and comparing two documents side by side. These comments confirm the fact that practical issues like size and weight will limit how many slates a user will be able to use at any given time.

5.2.3 Implications from the Probe

The technology probe directed our attention to critical parts of a multi-slate system that we needed to refine in order to produce a viable reading platform. We identified three big lessons as a result of conducting the probe.

- A) Slates are not just “hard paper”. The physical size and weight of a slate device, plus the fact that each device multiplexes several pages on a single physical object means that slates are used differently from paper. As a result our design should seek to *achieve the same goals* of paper-based interactions, not replicate them on the new platform.
- B) Slates should be tightly interconnected with each other and other tools like laptops. We believed we had addressed Morris et al.’s recommendations by implementing a distributed clipboard function (which users used, and liked), but this was not enough. A successful system must take extensive advantage of the available screen real estate

to simplify interactions with each document. It is also important that the multi-slate systems be easy to use during the text creation phase when users are mostly typing text on the computer.

- C) The “ideal” number of slates changes over time. Whether a user employed one or two slates was heavily dependent on task, environmental constraints, and individual preferences. As such, a reading system should not force users into using a fixed number of screens. Rather, it is preferable for the system to offer a way to reconfigure itself and remap content as the number of devices change. This ability is particularly important in light of the related research describing the importance of allowing users to read in a variety of venues [Adler et al. 1998; Tashman and Edwards 2011a].

5.3 The United Slates System

Our treatment of the United Slates system can roughly be divided into four parts. The first two parts are concerned with establishing the basic concepts behind the multi-slate environment and consist of how to distribute information to slate devices and how to move this information around efficiently. These functions are mostly targeted at reading tasks at the page, document, and partially the workspace level (as outlined in Chapter 2). Then, we turn our attention to additional functionality that enables slates to work with PCs in the workspace and finally, we discuss how our document and workspace model supports reading activities that span multiple reading sessions, venues and devices.

5.3.1 Mapping Information to Slates

The mapping of document information onto the slates is a matter of striking a balance between the many layout options afforded by independent sheets of paper against the simplicity of managing more structured collections of documents. Similar tradeoffs occur in the paper world when readers choose between loose-leaf paper and bound volumes. Fortunately, the dynamic display of electronic devices means that on any given slate, users

can operate at several different structural levels when working with their documents. Here, we describe how these different levels relate to one another.

As with most commercial reading appliances, all of a user's documents are aggregated into a single library. However, having a single collection for all of a user's documents can present a problem when the number of documents is large: as with stacks of paper, increasing the number of items in the stack can make managing the collection or switching between documents tedious. To address this problem, the mapping of information in the United Slates treats each slate as representing an active stack documents drawn from the library, along with several inactive stacks that can be swapped in at will. This model resembles how knowledge workers employ a small set "hot documents" that are actively in use, and an archive containing the remainder of the documents [Sellen and Harper 2003]. The inactive documents continue to be available through the library common to all slates. The active stack and global library metaphor is also conducive to how the slate environment interfaces with the outside world. The shared library can be presented to the user as a globally shared folder across all devices. Thus, the user can populate the contents of the library by simply moving document files into the shared folder.

The document at the top of the active stack is the active document. Navigation within the active document changes the active page, which is what the user sees on the screen in the default, single-page view of the content. Navigation between documents changes the active document in the stack (and changes the active page, by extension).

One major difference between the document mapping in the United Slates system and paper documents is that the United Slates does not allow pages of a document to be split apart from one another as can be done with paper. The practical effects of this policy are seen in some of the operations we present later in this section that allow pages to be moved between slates. When a page is moved from one slate to another, the document to which

the page belongs is placed at the top of the active stack of the destination slate. The reason for this restriction is to retain rich layout capabilities, and avoid the annoyance of having pages of different documents becoming intermingled.

5.3.2 Viewing Documents at Different Semantic Levels

The user interface (UI) of the slates contains a row of view selection buttons at the top of the screen. These buttons allow users to switch between one of four standard views (Figure 15), corresponding to the different information mapping levels outlined above. To take full advantage of the available screen real estate, each view is always shown full screen. This approach simplifies the interface by avoiding inefficient tiling, and complicated panning or scrolling maneuvers that would be necessary if multiple documents or views shared the same screen. The first view is a single-page view of the document. The second is a space filling thumbnail (SFT) [Cockburn, Gutwin, and Alexander 2006] overview of the active document. The third view shows the active stack of the slate. The fourth view shows the contents of the shared library available across the devices.

The key principle behind the United Slates system is that users will often use of more than one slate. Using multiple devices allows information to be multiplexed in space, which

4 Levels of Views in Our Reading System

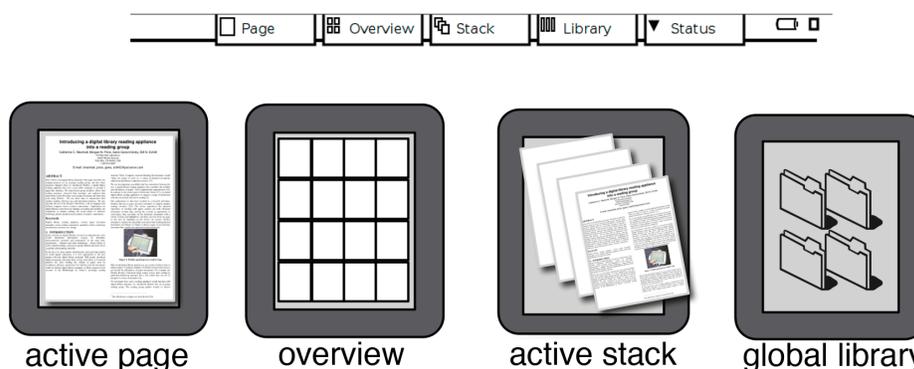


Figure 15 The four levels at which users can interact with documents. Our multi-slate interactions enable combinations of these views to support reading activities. A row of view selection buttons (top) is placed at the top of the screen and is used to switch between levels.

allows more material to be viewed at once. Spatial layout also helps the slates to conform to users' mental models of the content, which is important for organizing and keeping track of information during reading activities [O'Hara et al. 2002]. The cognitive benefits are potentially reinforced by the fact that the information is on physically separate devices, which allows users to take advantage of physical partitioning between devices [Grudin 2001] to better isolate distinct groupings of information.

Having multiple slates in the environment naturally implies multiple views will be visible to the user at any given time. Consequently, most of our interactions are designed to leverage the opportunities that are possible when multiple views are used in support of each other.

5.3.3 Moving Information Around

While the initial distribution of information across the workspace is important, reading activities additionally require this information to be manipulated and moved around. For instance, on paper, the process of moving information around allows users to access, organize, and compare information. The key realization is that there are many ways information is moved around in the environment. For paper interactions, similar tasks may involve completely different physical manipulations. For instance, while annotating a document placed far away, a user might reach and annotate in-place, temporarily bring the document closer and then put it back, or more permanently reconfigure the workspace. Depending on their needs, users decide on the fly which movements to employ. Every method of moving information imposes some type of cost/benefit tradeoff making certain movement strategies better suited for a particular task than others. Factors that can influence the appropriateness of a particular transfer technique include the speed to perform the transfer, the physical effort required to specify the relationship, and the duration of the interaction (across which the setup costs are amortized).

One important characteristic of paper is that moving information around is synonymous with moving the physical sheets of paper. In the case of slates, moving the physical devices is not only unnecessary, since one can simply move the information over electronic channels, but undesirable, because slates are heavier and thicker than sheets of paper. Instead, the United Slates system strives to enable users to achieve a similar degree of information movement, but decouples the movement from the need to physically manipulate the slates.

With this in mind, our reading environment provides a range of interactions for moving information electronically between the views on our slates. Through these electronic links, it is possible to create a wide range of electronic tools that enhance reading. We describe these interactions and the features they enable in more detail below.

Point-to-Point Connections Between Slates

To provide an unambiguous and fast way to move information between views, we aimed to give the user a direct method of specifying the slates participating in the operation while also taking advantage of bimanual processes present in reading [O'Hara and Sellen 1997]. Having the user directly specify the source and destination is well suited for our application, where slates are in the vicinity of one-another. Some advantages of direct interaction over a symbolic approach (like picking devices from a list) include reducing confusion about where data is going, side-stepping the complexity of mapping a device to a name, and supporting the use of anonymous displays [Rekimoto 1997]. The last point about anonymous displays is important because it abstracts away the idea that each new screen is a separate computer.

The Conduit interaction we developed achieves the bimanual point-to-point target specification described above. In the Conduit interaction the non-dominant hand is first used to designate a target and then the dominant hand, which offers more precision, selects

the item to be transferred. One way to visualize the process is that the user's body forms a conduit through which the information is transferred (hence the name). Conduit operations are particularly well-suited for activities where content needs to be transferred to a device with which the user is actively interacting, since the user's non-dominant hand is already close to, or holding that device.

One way we use the Conduit technique to move information between devices is to allow users to navigate to the page corresponding to a thumbnail on a device other than the one where the thumbnails are displayed. Retargeting the thumbnail navigation makes it possible for users to have devices that are dedicated for thumbnail overviews. To navigate, the user first depresses one of the large circular buttons on the device where he wishes to view the thumbnail in detail (which we will subsequently refer to as the command button), which specifies the destination slate of the command. Then, the user selects one of the page thumbnails shown in the overviews (Figure 16). The slate on which the command button is depressed shows the selected page. In this way, the user has the ability to quickly jump to any page to see it in detail, without having to repeatedly switch between page and overview views. With multiple panels of thumbnails, each corresponding to a different document, the user can rapidly access pages across different documents. Another situation where the Conduit technique is employed in the United States is for transferring documents from the

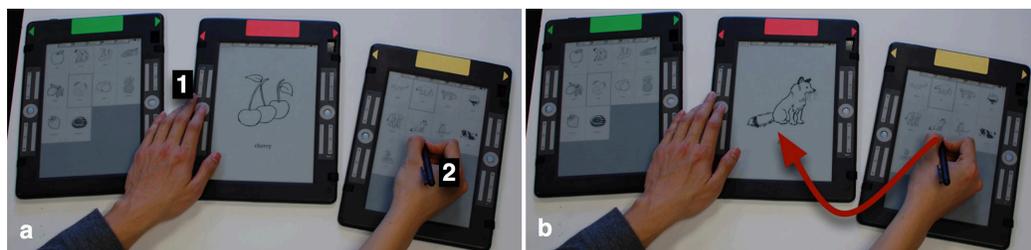


Figure 16 Distributed thumbnail navigation using the Conduit interaction technique. **Left:** Changing the page on the center slate via a thumbnail on the right slate by depressing the command button (1), and then tapping on the thumbnail (2). **Right:** Result of the operation; note that the user can be thought of as a conduit through which the data flows (arrow).

active stack of one slate to the active stack of another.

Although Conduit makes quick transfers possible, one limitation is that the transfers are asymmetric toward the slate the non-dominant hand controls. Moreover, given its simplicity, it is also not very well suited for commands where parameters are needed at both the origin and destination, such as copy and paste. For these cases the United Slates offers a variant of Stitching [Hinckley et al. 2004] (Figure 17). The basic mechanics of Stitching are the following: First, a stitchable element on the source slate is crossed through with the command button depressed, which causes a context menu to pop up. Then, the stitching icon corresponding to a command is picked from the menu. Finally, the user taps on the destination slate, optionally specifying additional parameters for the command, at which point the command is executed. For interactional consistency, the source and destination slates can be the same in a Stitching operation.

To further speed up stitching operations, the United Slates provides an additional shortcut by allowing the user to cross through UI elements with a long cross to automatically select a default command, rather than having to pick from the menu [Dixon, Guimbretière, and Chen 2008].

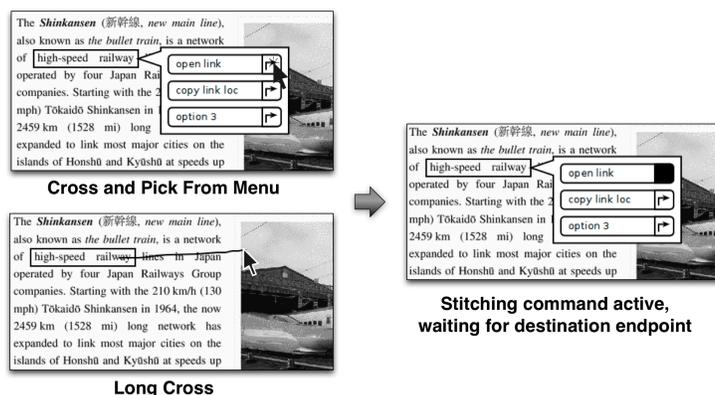


Figure 17 Performing a Stitching operation using stitchable hyperlink elements. Users have a choice of selecting from full menu of commands (left, top) or using the shortcut of a long cross (left, bottom). While the system waits for a destination endpoint for the endpoint for the command, the indicator square next to the selected command is filled. (Figure is taken from a screenshot from our device simulator and uses desktop cursors for clarity).

At the page view level, Stitching hyperlinks to another device offers a way to explore hyperlinks on a page without having to navigate away from that page. By keeping the origin of the hyperlink visible, this feature helps maintain context for the newly opened link, supports comparisons between linked documents, and speeds up backtracking by not requiring extra navigation to return to the source of a link.

Lastly, we want to clarify that Stitching and Conduit can be used interchangeably for all of the point-to-point connections so far. For instance, even though our example above demonstrates opening hyperlinks on a different slate using Stitching, Conduit could have been employed instead. Likewise, it is possible to stitch a thumbnail from one slate to another.

Proxying Interactions Through Slates

One interesting feature that can be exploited in a multi-slate system is that since individual slates are functionally identical, one slate can easily double as another. The chief way we apply this functionality is in a remote control feature to provide interaction at a distance.

Since users have trouble stacking up and overlapping slates, they have a tendency to spread them out across a large area on the desktop. Unfortunately, certain interactions such as navigating, writing, or examining a page in detail are difficult to perform when a device is far away. Users can stretch to operate a slate, but that would be uncomfortable for lengthy operations. Users could also pick up the distant slate and move it closer. Unfortunately, if a user has a slate in hand already, this requires extra work because it is difficult to hold two slates at once given their size and weight. The remote control feature (Figure 18) addresses these problems by allowing the user to use a nearby device as a proxy for a device that is far away. The functionality leverages the fact that users frequently have a device in hand while reading. By enabling users to temporarily change a slate's identity, the remote control function gives the illusion that the user has picked up a distant device and brought it closer.

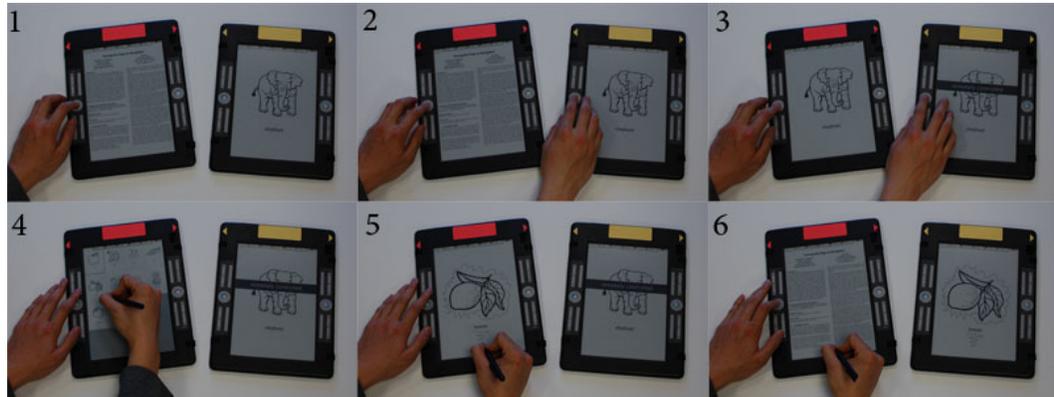


Figure 18 The slate remote control interaction. (1) Press and hold command button on local device. (2) Click command button on remote device (enhancement), (3) Remote screen is frozen and transferred (note the message informing the user that the session is active), (4, 5) Using local device as a proxy, (6) Double click command button to end remote control session.

In the system, a Conduit interaction is used to activate the remote control feature. As with all Conduit interactions, the command button on the closer device is depressed. Next, one of the view selection buttons at the top of the screen is picked on the distant device, which sends that particular view to the nearby device. When that occurs, the screen on the distant device is frozen and displays a message reminding the user that it is being remotely controlled. One enhancement we added to make the remote control session easier to start was to let users push the command button of the distant device instead of having to tap the view selection button on the distant slate. This enhancement provides users a larger target to acquire, as well as tactile confirmation, which is helpful when the user must reach to access a distant device.

When a remote control session is active, all of the functionality of the system is retained, with the exception of being able to invoke an additional remote control session. The reason for disallowing the invocation additional remote control sessions is because we thought that it would be difficult for users to keep track of multiple simultaneous remote control sessions, or chains of remote control sessions. In order to enable the normal controls to work transparently during the remote control session, a special input sequence is needed to terminate the remote control session. Therefore, the user ends the remote control session

by double clicking the command button on the nearby device. When this occurs, the distant device updates to match the changes made while it was being remotely controlled, and the local device returns to its original state (Figure 18, panel 6).

Links Based on Slate Proximity

More permanent connections between slates can be useful for reducing repetitive actions to set up relationships between slates. However, invisible electronic links between slates run the risk of confusing the user if the user is not aware of the connection. For instance, Chen et al. [2008] reported that users were confused when displays on a device changed unexpectedly while the user was working with a different device. These issues are an example of mode confusion. In order to reduce the chance of this type of mode confusion, the system only allows more permanent connections to be formed based on slate proximity. This provides strong visual reinforcement that there is a relationship between devices, reducing the chance that slate state changes result in confusion. Although the cost of establishing a proximity-based connection between devices is relatively high, given the physical manipulation required, the resulting implicit connection between devices saves user effort later. Therefore, for connections that are expected to last a long time, proximity links are preferable.

Like the Dual-Display E-book [Chen et al. 2008], we use proximity links to create a two-page, side-by-side view of a document with synchronized page turns. The two-page view effectively increases the amount of content one sees from a document at a given time. Proximity links are also used in the United Slates system to streamline the creation and viewing of margin callouts on the page view of a slate. These callouts are a means to provide additional writing space, which help to complement superimposed ink marks [Pearson, Buchanan, and Thimbleby 2009]. Normally, when using one slate, users can get extra writing space by creating a margin callout with an angle bracket gesture (>) against the edge of the screen. With a single slate, margin callouts need to be called up and viewed

one at a time (Figure 19, left). But when the slates are linked, the system expands the margin to the second slate; the right slate displays all of the margin callouts associated with the page shown on the left slate (Figure 19, right). Creating new callouts, or editing existing ones also occurs on the right side, eliminating issues of the callout occluding the main text in the single slate case. When slates are initially brought together the system defaults to the two-page view, but the user can easily switch the behavior to expand the margin using a system setting.

5.3.4 Controlling Slates From The PC

We saw during our technology probe that it was important for users to have their reading system well integrated with their main content creation tool, which was generally laptop or desktop PC. PCs provide a keyboard and vertical display that faces the user and are generally preferred for composition tasks [Morris, Brush, and Meyers 2007]. Other tasks for which a PC is superior include things like accessing data from digital libraries. As such, we expect that many documents will be marshaled together on the PC first, before making their way onto the slates. While it is an option to attempt to duplicate PC functionalities on the slates themselves, like a Tablet PC, that would necessarily entail adding a great deal of complexity to the interface. Wilson and Landoni [2003] point out that one desirable quality for a reading device is maintaining its simplicity so that operating the device does not begin



Figure 19 Left: Single screen margin callout window. **Right:** Proximity links used to expand the view of a page to include margin notes.

to interfere with reading activities. Thus, we thought a better approach is for the slate system is to complement, rather than replace, the PC in reading environments. One way we handle the requirements of federating with a PC is through a synchronized distributed clipboard [Miller and Myers 1999], of which the PC is a member. The synchronized clipboard supports information extraction tasks by ensuring that content on a slate can be shared between both slates and PCs.

The concurrent use of PCs and slates presents an interactional challenge arising from differences in the devices' input modalities. PCs are primarily operated via keyboard and a pointing device, while our slates are controlled via hardware buttons and a stylus. Tasks that involve repeated switches between PCs and slates are slowed considerably from the user having to re-home on different input devices after every switch. To address this problem, we decided to extend the remote control mechanism so that slates can be controlled from a PC. Upon activation, the PC interface presents a grid of all of the slates in the environment (Figure 20). From this visualization, the user can select a particular slate, and then control the slate using the keyboard and mouse. This feature uses the same mechanism that powers the remote control feature described in the previous section. The PC remote control function allows a user to have full control of a slate environment (including text entry and text selection), but treat each slate as if it was another window on

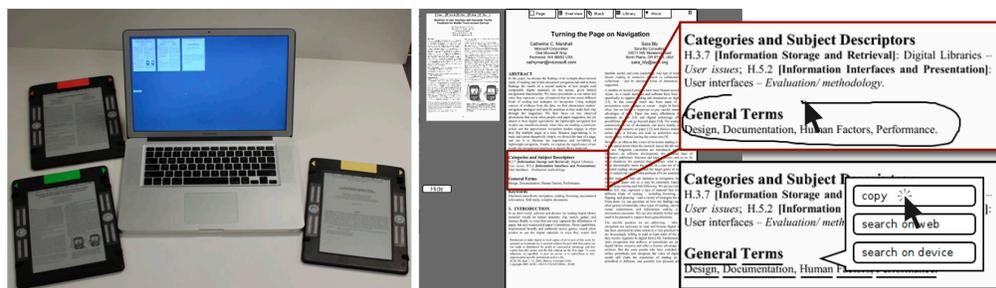


Figure 20 Sample layout for PC Remote Control. **Left:** PC showing overview of slate devices. **Right:** Detail view of screen on PC, after selecting a slate to control. **Insets:** Lassoing text, then selecting the copy command, at which point the text is copied to the distributed clipboard.

the computer desktop. For example it is possible to excerpt a piece of text very quickly using this feature (Figure 20, right). A visualization that matches the positioning of the slates in the environment would be possible if we had more sophisticated position localization techniques at our disposal.

5.3.5 Document Management With Stacks

The functionality we have presented up to this point has been in the context of reading activities that use the small set of documents that are in the active stack of each slate.

However, reading activities will at times require documents that are not in the active stack, forcing users to modify, save, or replace the contents of the active stack on a slate. The availability of robust stack management capabilities has far-reaching effects with regard to aiding document organization, providing the ability to suspend and resume reading activities, as well as enhancing the mobility of the system.

Aiding Document Organization

One of the design choices we made in the system is to map several documents to the active stack of each slate, which we detailed in the first part of this section. Recall that the reason for this decision is our belief that users will probably need to work with more documents than slates available at any given time. Presumably, like paper stacks, the electronic stacks comprise logical organizations of documents so that users can quickly find and access a desired document. However, with a limited number of slates, a slate may need to be repurposed to view a document that may not necessarily be related to the other documents in the stack of that slate. To illustrate, suppose there are two slates, one with a stack of documents about different types of cats and one about different types of dogs. At some point, a user would like to compare two different dogs, side-by-side, meaning he would have to move a dog document onto the stack of cat documents. After several rounds of this, the stacks on each slate lose their organizational significance.

The Stack Manager (Figure 21) component of the document library addresses the above problem with an attribute tagging mechanism, like the one found in Presto [Dourish et al. 1999], to group together documents in a more permanent fashion. Tagging of documents is similar to placing them into folders, except that a document can belong to several tag groups. The Stack Manager allows users the ability to quickly specify complex groupings of documents using unions of tag groups.

In one operation, documents in a tag group (or union of tags) can be added to, or replace, the active stack of a slate. The tagging system allows users to quickly and concisely select documents to put on the active stack of a slate, ensuring that users can rapidly get at the

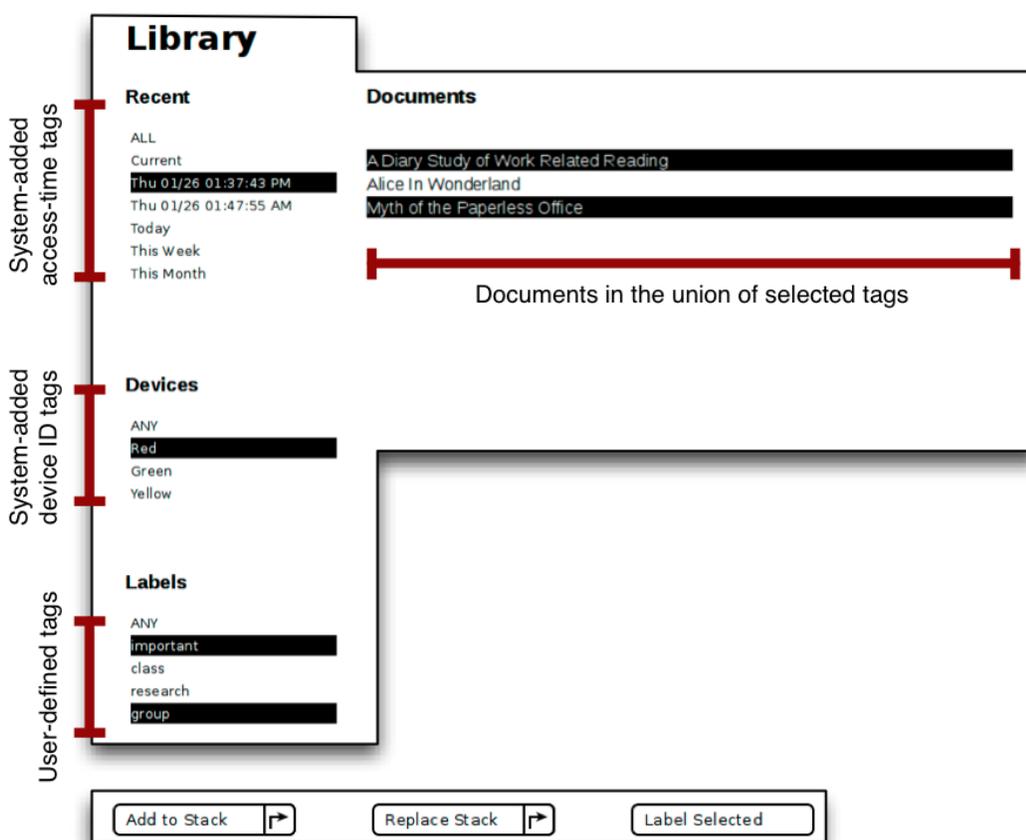


Figure 21 Screen shot of the Stack Manager component of the Document Library. Users select a combination of system-added and user-defined tags along the sidebar to filter for documents. Bottom inset: Using buttons at the bottom of the interface, the selected documents can either be added to the stack, replace existing stack contents, or be saved with a new attribute tag.

documents they need. As for the above example, it is simple for the user to "clean" the stack of each slate by reloading the contents with documents that have been appropriately tagged. The reverse process, where a user assigns a tag to the documents in the active stack, is useful as well, especially for sorting tasks. Suppose a user is going through a large set of uncategorized documents and wishes to group related documents together. The user can move documents from the slate with a stack of uncategorized documents, to other slates that represent the document categories, which initially have empty stacks. After moving the documents to the appropriate slates, the user can tag the documents in each stack in a single operation. This method of sorting allows the user to leverage spatial layout to sort documents, much like what is possible with paper, and avoids repetition by allowing several documents to be tagged together in batch.

Suspending and Resuming Reading Activities

We have shown how the Stack Manager is good for organizing documents, but its functionality also makes it possible for reading activities to be suspended and resumed at a later time, at a different venue, and with different devices. If a user already has slates available in the different places where reading occurs, the ability to restore stacks frees the user from having to transport slates to those venues.

The mechanism that makes this additional aspect of the Stack Manager practical and transparent is the application of automatic, system-added attributes to the documents. As documents are placed onto the active stack of a slate, the Stack Manager automatically tags them with a timestamp as well as the ID of the device on which it is being placed. Using these automatic tags, it is possible for users to recall the specific documents used on a device at a particular time, without needing to explicitly tag the document ahead of time. If a user wishes to restore a stack of documents he was previously working with, he can filter for the documents by time and by slate, and quickly resume working with those documents. More powerful, however, is how the user can leverage the functionality to restore a reading

session using completely different devices. For instance, if a user is using three slates at work, he can recreate his reading workspace at home using slates he already has at home (Figure 22, rightmost panel). The ability to filter by slate makes it easy to map the contents on each work slate to a home slate.

Dealing With Resource Mismatches

The Stack Manager further increases mobility by allowing an environment with many slates to gracefully scale down to one with a smaller number of slates. As a result, users are able to continue working in environments where fewer slates are available (assuming the user does not want to carry a full set of slates), or environments where working with many slates is not practical (e.g., public transit). The method to achieve this is simple: the user can take the union of tags corresponding to documents on different slates and redistribute them onto a smaller number of slates (Figure 22, middle panel). While functionality may suffer from the reduced number of displays, this technique helps reduce the disruption caused by a resource mismatch between reading environments.

5.3.6 Putting Everything Together

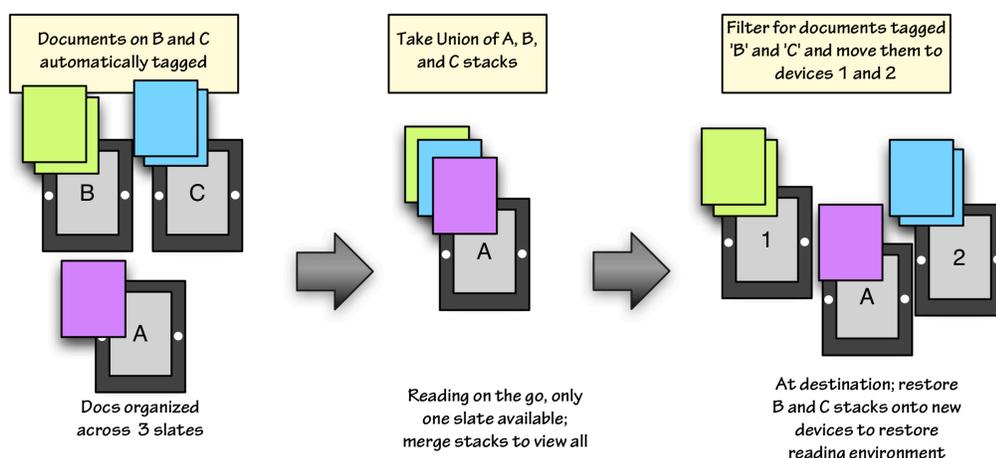


Figure 22 Using tag operations to manipulate, save and restore document stacks in order to save and resume reading activities across different environments and times. Although Slate A is always present in all of the venues in the above example, there is no requirement that a common slate be present.

In the sections above, we have described many examples of interactions between slates and explain how these interactions can benefit reading activities. To give a better idea about how the multi-slate interactions presented above fit in with the broad requirements of reading, we place these multi-slate in context with the different reading requirements identified in Chapter 2. Figure 23 organizes the features we presented in this section of the paper in a structure matching that of Table 1. As one can see, the beneficial enhancements that the United States system spans the full gamut of reading requirements, from the page level to the multi-session reading level. The features made possible through multi-slate use, coupled with the unique functionality of the individual slates (e.g., light weight, readability,

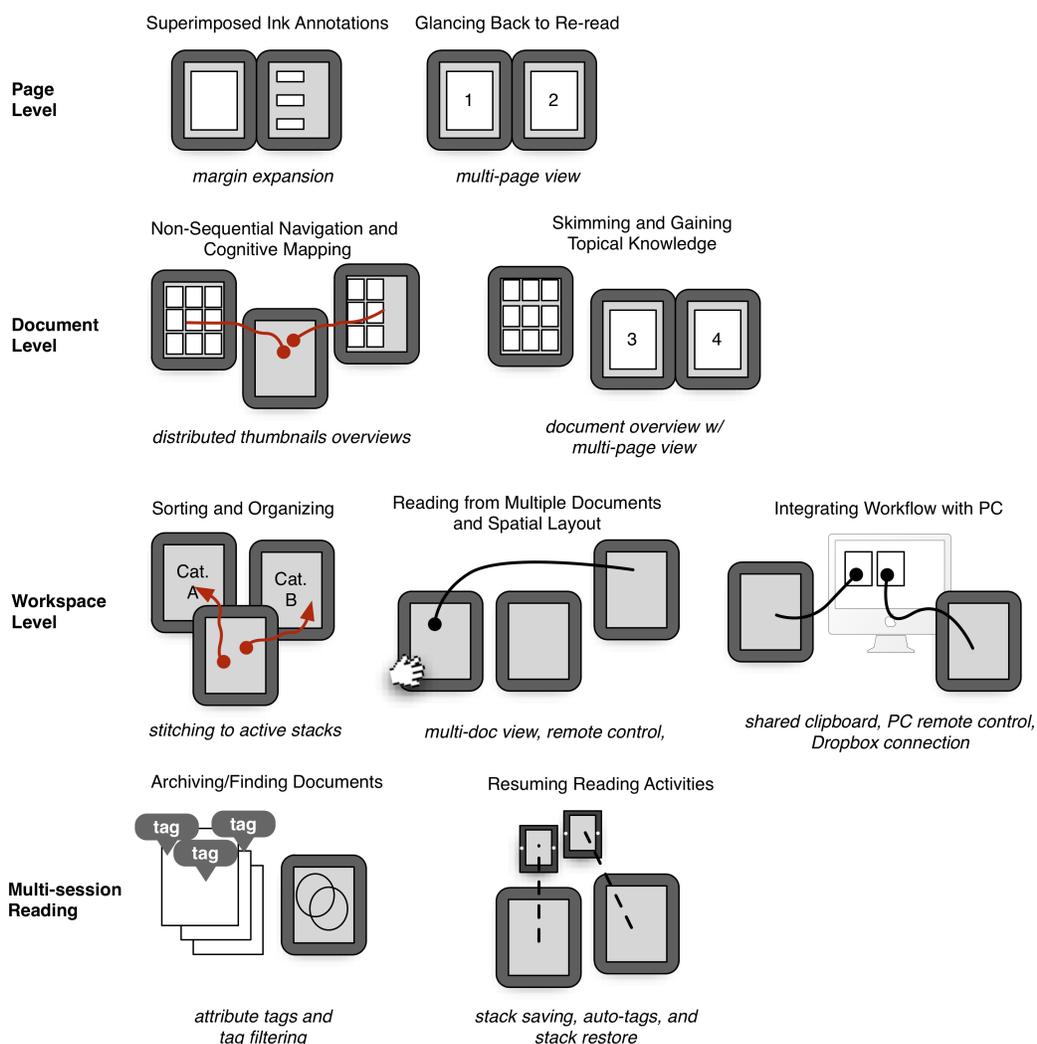


Figure 23 Multi-slate interactions support the reading requirements presented in Chapter 2.

and support for writing), translate to more complete support of reading requirements than existing electronic reading technologies.

5.4 Implementation

5.4.1 System Architecture

The United Slates system uses the slate hardware described in Chapter 4. In the initial implementation for laboratory testing, the slates use the Bluetooth radios to establish a Bluetooth Personal Area Network (PAN). The Bluetooth PAN allows each device to connect the others through an IP network abstraction. The advantage of using the PAN is that no external network infrastructure is required to establish the network. The major disadvantages of the PAN are: (1) there is a maximum of 8 devices; (2) one special device, the Gateway Network Controller (GN), always needs to be present to route data between the other devices; and (3) the capability and UI to connect to a PAN is not available or poorly implemented on many common operating systems. In the laboratory, these were not a problem, since we simply specify that a particular device will be the GN and connect the PC to the PAN ahead of time.

In order to implement the stack tagging, saving and restoring features described, we use a central server that stores the contents of every device's stack. In the current implementation, we simply placed the server on the Bluetooth GN, because we knew for sure that that particular device would be present in any configuration.

In the laboratory implementation, we also make the assumption that there exists a method for users get documents onto the slates and to ensure that the documents on each slate are identical. However, for testing, we did not implement the functionality and simply placed matching documents into the local storage on each slate. Annotations and other changes to the documents are propagated, however. As annotations are applied, the device on which

the annotation is made is responsible for sending out the annotation data to all over the other devices. We use a non-optimistic serialization scheme [Greenberg and Marwood 1994] (last input in wins) across devices. This scheme works for our application because the scenario is of a single user operating several devices. Also, the modifications users make to the document are simple and can be applied quickly. Strokes generally appear on remote devices within a second.

The implementation described above makes a number of assumptions about the availability of certain devices and omits details of how document data initially gets onto the slates. These assumptions mean the system functions well only in a controlled lab environment. In the next chapter, we describe additional system components that provide the data synchronization, device state sharing, and connectivity required for practical real world use.

5.4.2 Software

The application software used in the United Slates is written using Python and uses the GTK+ toolkit, along with PyGTK Python bindings. Inter-device communication was performed using PyRO⁶, a Python framework that provides facilities for object serialization and remote procedure calls. The tablet federation application on the PC is also written in Python and shares much of the same codebase as the application running on the slates.

The application employs the Model-View-Controller [Reenskaug 1979] pattern heavily. The application is split into several view-controller and model classes, each responsible for one level of the system interface (e.g., page view, overview, active stack, library). The view-controller class handles drawing and input while the model interfaces with underlying data and handles inter-slate operations (e.g., document and annotation data, thumbnail data, stack data). A view manager class routes input events to the view that is active. The model

⁶ <http://irmen.home.xs4all.nl/pyro/>

classes maintain all variables related to the state of their associated view, and can export the state on demand or temporarily load a new set of state variables. It is through this capability that we implement the remote control features of the system. The United Slates can use a number of different document formats. For the evaluation presented in this chapter, a custom document format consisting of Portable Grayscale (PGM) page images organized into directories on the filesystem is used. Strokes are stored in separate files and overlaid on top of the page images. We use the $\$1$ Recognizer [Wobbrock, Wilson, and Li 2007] and ShortStraw [Wolin, Eoff, and Hammond 2008] algorithms for recognizing the carats, lines, and angle bracket selection gestures used in the system.

5.5 Preliminary Evaluation

We carried out an evaluation study in which we asked participant users to describe some of their typical work practices involving active reading, and then to try out some of the features of the United Slates system. The goal was to have these users assess whether the features the United Slates system would be applicable to the real world reading tasks they usually carry out. This was not a controlled study aimed at comparing our system against others. Rather, this was an exploratory study in which we aimed to provoke discussion and feedback from potential real users against a grounded understanding of their real work practices.

For this study, we recruited 12 students from Cornell University. We did not recruit for specific majors or class level. Our participants consisted of 11 undergraduate students and 1 graduate student. Participants received \$20 compensation for their participation.

The study was set up as a semi-directed interview in which the participants had the opportunity to perform and compare different kinds of interactions on the slates. We started each session with an interview about the reading (e.g., class-related, work-related, or research-related), which they had undertaken in the previous semester. This was not only to

give us an understanding of their practices, but also to get the participants to reflect on the ways in which they really do their work prior to using the United Slates. We asked participants to describe things like: the reading materials they used; what and why they printed to paper; additional resources they employed when reading and writing; situations where they used multiple documents together; and instances where the documents they used were on different types of media.

5.5.1 Observation Procedure

Following the interviews, we explained that the purpose of the study was to assess the efficacy of features that would take advantage of situations where multiple slates were available. First, we outlined the basic capabilities of a single slate and let participants familiarize themselves with the controls and inking capabilities of the device. Then, we tested the multi-slate reading support tools we developed in the following order: margin expansion, overview navigation, hyperlink navigation, screen teleport, shared clipboard, PC remote control, and stacks and reading sessions.

The general process we used for testing the tools was to have users perform short tasks that simulated portions of real-world reading activities using each tool. We first had users perform the tasks with the United Slates condition, which employed the interactions we have presented. Then we had users use a single slate baseline condition that represented the experience of using a contemporary single-screen reading device or the user experience of a multi-slate system that did not provide features integrating the slates. The purpose of the naïve condition was to give users a point of comparison against which to discuss the features in the United Slates system. There were two features for which we did not provide a naïve condition: copy and paste to the PC using the distributed clipboard, and stack management across slates. These did not have a naïve condition because there was no plausible method of performing those tasks efficiently without the functionality the United

Slates provides. Table 4 provides specific details about the task that we used to test each feature, the procedure employing the features of the United Slates system, and the procedure for the baseline condition.

Feature	Task	United Slates Procedure	Baseline Procedure
Margin expansion	Make margin callout annotations on several pages, return and review the annotations	Use two slates that are attached using proximity links so that annotations are created and displayed alongside the main text	Use a single slate to create margin annotations and cycle through annotations one at a time.
Distributed Thumbnail Navigation	Navigate to a different page in Document A (A), navigate to a page in Document B (B), and then navigate back to A	Use a total of three slates, with two slates showing thumbnail overviews of A and B, and one slate for viewing a page in detail. Use the distributed thumbnails to jump to target pages.	Use a single slate to go between A and B and back to A. Users needed to go to the stack view to switch documents, and then use the thumbnail overview in each document to navigate to the target page.
Hyperlink Navigation	Visit a sequence of hyperlinks. The sequence includes backtracking navigation where the next hyperlink to follow is on a previously viewed page.	Perform the navigation using three slates, using a combination of Conduit and Stitching interactions to open links on other devices in the environment.	Use a single slate to visit the sequence of links. Backtracking was performed using a dedicated “Back” button.
Remote Control	With a slate in hand, change the page on a different device to a specific target page. Once at target, make ink markings on that page.	Use Remote Control function to move the display of the distant device to the slate in hand. Perform navigation and inking tasks using slate in hand and then send the display back to distant device.	Reach and pick up the distant slate, perform navigation and inking task and then return the distant slate back to its original location.
Distributed Clipboard	Copy and paste text passages from the slate to a word processing application on the PC	Copy selected text to the distributed clipboard. Paste text on PC using either mouse or keyboard (up to the user).	N/A
Remote Control from PC	Navigate on the two slates alongside the PC so that they are displaying specific target pages.	Use the PC interface to remotely control the slates. Perform navigation without touching the slates.	Navigate by directly operating the two slates.
Stack Manager, Tagging	Place a selection of documents into a newly created stack. Shut off a device, use stack manager to restore that stack.	Use stack management interface and perform tasks using tags.	N/A

Table 4 Procedures for the preliminary evaluation of the United Slates system.

After testing the multi-slate and single slate conditions for each feature, we asked users to compare the single slate method of accomplishing that task and the multi-slate approach. Then we asked participants to describe whether the feature that was just tested would be useful for their personal reading tasks, whether the feature was confusing, the advantages and disadvantages of the feature, and how the feature might be improved to make it better suited for the style of reading they typically performed.

5.5.2 Results

Users appeared to understand the operation of all of the techniques with minimal intervention and explanation on our part. For nearly all of the features, users had no trouble jumping in and performing the interactions we asked them to perform during the evaluation. Therefore, it would appear that these techniques could readily be used in a system outside of our lab.

Overall, users were positive about all of the multi-slate functionality we tested with them. The positive feedback is perhaps not surprising given the fact that participants in such studies are generally inclined to be positive about prototype systems in which they see the investigator as having a stake. Therefore, the more interesting takeaways are in the specific comments that participants made with respect to each feature, including their suggestions for improvement.

Margin expansion – Participants reported that expanding the space allocated to the margin callouts by physically attaching two devices made creating and viewing annotations easier than the process when using a single screen. Participants mentioned that using this technique meant that callouts no longer occluded the main text, allowing one to read while writing (P4, 7, 9, 10, 11, 12) and that showing all of the notes at once on the second screen made it easy to find pertinent notes (P4, 5, 7, 8, 10, 12). The most requested improvement was to provide facilities to either expand the writing area or to allow writing on the entirety

of the second device (P6, 7, 8, 9, 11). Participants added that the proximity link system could also be used for other features like automatically searching selected text on a second screen (P7) or automatically displaying documents that the user has associated with a particular page (P2, 8, 12).

Distributed thumbnail navigation – Every participant except for P8 mentioned that overview navigation would be very useful for working with multiple documents or locations in tasks such as writing papers, comparing articles, and browsing lecture notes. Most participants (P1, 2, 3, 5, 7, 8, 9, 11, 12) pointed out that they liked the simplicity and speed of using the thumbnails compared to the alternative we had participants try, which involved going to the current stack, switching documents, and then navigating to the desired page. Improvements that participants suggested included showing annotations and margin notes more prominently in the thumbnails (P2, 4), and making the process of calling up a thumbnail on a different slate more like drag-and-drop by providing visual feedback showing the process of moving the thumbnail from one device to another (P6). P11 mentioned that he probably would not want to use a separate device for each document overview as that could become difficult to manage. Instead, he said he would prefer to “swipe” or otherwise cycle between SFT overviews on one device, while employing an additional device to view the page in detail.

Hyperlink navigation – Participants found several reasons why they would want to open a link on a different device. First, opening a link on a second device made sure both the source and target documents are visible, ensuring the user did not get lost. It also reduced the amount of flipping that needed to happen to work with both documents (P1, 2, 6, 9, 10, 11, 12). Cross-device hyper linking also helped participants focus on reading material, by allowing them to move content from a peripheral slate onto the “main” device that they were holding (P6, 8, 9, 10, 11). One problem that some participants encountered

when we introduced Stitching after demonstrating Conduit (for thumbnails in the previous task) was confusion about which technique to use when attempting to open a link on a different device. When asked to open a link, some users proceeded to pick up the target device to use Conduit. In these situations, we reminded the participant to use Stitching. Most users, when asked about this behavior, attributed it to having the (incorrect) impression that information could only move from the dominant hand to non-dominant hand. They added that once they became more familiar with the interaction, Stitching made sense as well. Based on this, we believe that a more consistent way to view the Conduit interaction is as a *shortcut* for special cases of Stitching. That is, every operation is a Stitching operation, but by virtue of the command button, Stitching to the slate the user is holding does not require an endpoint to be specified. A few participants (P11, 12) expressed concerns that spreading documents among too many devices could be confusing. Two participants (P10, 12) suggested that a simple list based approach for selecting a target, rather than interacting with the target device itself, might be preferable.

Slate Remote Control –In the non-remote control condition, where users had to manually operate the remote slate, all participants placed the slate they were holding down on the table before picking up the other slate. This behavior confirms the difficulty of holding on to more than one slate at any given time. However, the actual utility of using one slate to control another split user opinion. Most participants found that using a nearby device to control a remote device to be preferable than the alternative of reaching over to pick up the remote device (P5, 6, 7, 8, 9, 10, 11, 12). P8 described it as being “more efficient—3 presses of a button did the same as reaching over.” However, a contingent of participants (P1, 2, 3, 4) found the feature to be confusing, mentioning that it would be simpler to operate the devices directly, even if it meant physically shifting devices around. These participants said that the prototypes were light enough to move around without being awkward, and indicated that the experience of using the screen teleport feature to be

slightly worse than directly moving the slates. Like with the hyperlinks, participants mentioned it would be helpful to pick from a list (P10, 12) rather than having to stretch to interact with the device to proxy to. P12 specifically mentioned that picking from a list would be useful for situations where she moves from reading at her desk to reading on the bed and wants to access a resource that was collocated in the environment, but just out of reach.

Shared PC clipboard – All participants thought the shared clipboard feature would be helpful, although P8 believed that working with electronic documents directly on his PC would still be preferable. Not surprisingly several participants (P1, 2, 7, 9, 10) pointed out that it was annoying to switch between the slate and the PC. There were some differences in how our participants employed the clipboard, however. Instead of using pull-down menus on the PC, P5, and P11 used keyboard shortcuts to perform the paste operation on the PC side. Since it was possible to press the shortcut keys while holding the pen, these two users did not report issues with having to switch devices.

Remote Control from PC – All participants, with the exception of P12, found controlling slates remotely through the PC to be the more preferable option for operating the slates when we asked them to control the slates while performing a simulated writing task on the PC. P12 found it easier to operate the devices directly. Many participants cited the fact that they did not have to switch devices as an advantage (P3, 4, 5, 6, 8, 11). P4 and P8 remarked that the remote control allowed the slate alongside the computer to act as an extra, portable monitor. Another advantage participants cited as a plus for the remote control functionality was the fact that they were familiar with using a traditional PC (P1, 2, 3, 9, 10). P7 and P10 mentioned that the PC interface made controlling many devices in rapid succession easier. One unforeseen benefit of the PC remote control feature was that turning pages and switching documents on the PC was *much* faster, owing to the faster

processor and display on the computer. Participants had mixed feelings about the lack of an accurate rendering of slate location in our visualization. A majority believed that the visualization actually felt that it provided a more organized view of the slates in the environment, allowing one to more easily see and access the slates (P1, 2, 4, 5, 6, 9). But, several participants mentioned that it was somewhat confusing to identify a target slate because the ordered grid of slates in the PC UI did not match the spatial layout of the slates on the desk.

Stacks and reading sessions – The vast majority of participants thought that tagging documents and organizing them into tag groups to be a good translation of how they currently organize documents. P9 made the comment that the organizational system made going through the documents easier by making it “feel like checking email.” Some users additionally mentioned that they sort documents by time as well, which the time-based tagging automatically supported. Users also mentioned that the stacks would allow them to go beyond what they currently do. A number of participants mentioned that they frequently lose track of documents, so the automatic session tracking feature would be helpful (P3, 6, 7, 9, 12). Finally, participants confirmed that the ability to save reading sessions to be resumed later on other devices would increase portability (making it possible to not have to transport devices) and would allow them to just carry a single “main/master device” (P5, 6, 10, 11, 12). The participants who were not as enthusiastic about the stacks and sessions (P4, 8) mentioned that the use of the tags resulted in too many grouping options. Instead, these users said that they would just prefer to organize documents using traditional folders.

5.5.3 Evaluation Summary

Participants found the interoperability of the slates with the PC to be the most compelling feature, since the isolation of information on paper documents from tasks on the PC was a

problem familiar to all of our users. The other feature of our reading system that participants judged to be directly applicable to their current reading needs is the ability to manage large collections of reading material we provide in our stack manager. Users though other aspects of the United Slates would be helpful as well, at least in the context of the experiment tasks. But, for these other features users were more restrained. The restraint can be attributed to the fact that users did not have as clear of an idea about whether these other features would actually be useful in their actual reading activities.

One particularly interesting observation was that while users were split about the utility of screen teleport between slates, they were considerably more positive about the application of the same core functionality to control the devices from the PC. This is an interesting finding on several levels. For one, the fact that many users did not find picking up a slate and moving it around onerous provides a data point about the threshold of how heavy a reading device can be before it becomes burdensome to handle. The slates, which weigh in at 500g each, appear to be under this threshold. The discrepancy between slate and PC remote control is enlightening as well. We believe there are two main reasons for the discrepancy. First, there is more work involved to switch from using a PC to using slates than to switch from slate to slate. Second, many users found that performing navigation tasks on the slates using the PC interface provided a better experience than working on the slates directly: on the PC all the slates can be controlled from a single location, and using a keyboard and mouse was fast and familiar.

Finally, we noticed that there was a practical limit in the number of slates that could reasonably fit in a user's immediate working area, especially when a computer was present. The main reason for the limit is because users do not overlap slates: the thickness of the slates causes slates to wobble when partially stacked on top of each other. Therefore, P11's

suggestion to allow cycling through multiple panes of thumbnails on one device is a particularly useful suggestion.

5.6 Discussion

5.6.1 Multiple Devices Help Resolve Reading Tensions

Tashman and Edwards [2011a] identified two major tensions in active reading technologies. The first of these tensions was that users wanted a large amount of space to work with many documents, but also wanted the reading environment to be portable so that they could set up in arbitrary locations. The second tension was that users liked the experience of reading from a dedicated device, given their unique affordances, but required the use of their PC for many of the support activities connected to that reading. We believe that the multi-device strategy employed by the United Slates system offers an effective solution for resolving both of these active reading tensions.

One of the fundamental motivations underpinning the United Slates is to provide more space to work with documents, in ways that are both flexible and optimized for cross-document use. At the same time, we believe that the United Slates is also a more portable solution than other approaches. At the most basic level, transporting a set of slates is a far more practical proposition than transporting a large multi-touch surface. However, with the ability to save and restore reading sessions, users may not need to transport any slates at all, as they can capitalize on slates that are already present at the venue where they wish to read. One way to further increase the mobility of a multi-slate system is for public venues to provide slate devices as infrastructural tools, much like a projector. The reason this vision is compatible with a multi-slate system is because the stack manager blurs the lines between public and private devices. Work performed on a shared public device is not lost once the user returns the public device. Furthermore, physical proximity links might be used as a way to temporarily associate public devices with a user in a fast and unambiguous manner.

The physical nature of this type of authentication may also be useful for dealing with rights management issues.

How United Slates integrates reading tasks with activities that occur on other tools in the reading workspace was another requirement that we expressly designed our interface to support. We recognized that PCs and reading devices have unique niches owing to the significantly different input and output modalities that each support. As such, we believed it was unwise to force users to constrain their reading tasks to the PC, or relegate standard computing tasks to a device that offers a smaller screen, and lacks keyboard and indirect pointing control. This belief led us to implement features that streamlined the use of both slates and PC simultaneously. In our study, users expressed great enthusiasm about the ability to copy and paste between the slates and the PC rather than a desire to write papers on the tablet. Therefore, we believe that designers should think of slate reading devices not as a separate computing device to be used in isolation, but as a peripheral where certain activities can be offloaded. In some sense, reading systems should reduce the costs of choosing the “best” device for a particular task. To support this pattern of use, we believe it will be important (and probably not too difficult given advances in cloud computing) to offer even tighter integration with the PC. For example, files and annotations on the PC should be readily available on slates and changes made on either type of device should immediately propagate.

During our evaluation an even more interesting phenomenon involving the PC emerged: we discovered that the PC served as a good platform for performing meta-tasks in the multi-slate environment. For example, the PC provided a central place to quickly set the state of multiple slates. Thus, we believe that connecting the multi-slate reading environment with external devices goes beyond making sure that it is possible to move data

between slates and PC. Instead, researchers should consider new interactional options that further take advantage of the union of slates and PCs.

Overall, the United Slates system presents a multi-device solution that enables users to work with many documents at once, provides the annotation qualities and form factor found in dedicated reading devices, and allows users to fluidly work across PC and reading slates. This is a particularly interesting result because it indicates that a composition of devices, rather than a single hybrid device, is a good way to combine the desirable functionalities found across different devices. Don Norman once observed that, “When one machine does everything, it in some sense does nothing especially well, although its complexity increases.” [Rheinfrank 1995] Our results indicate that the strategy of using device ensembles [Schilit and Sengupta 2004] is an effective way to maintain the unique qualities of a device while increasing overall functionality.

5.6.2 Opportunities for Improvement

Although the United Slates, as presented, attempts to address many of the outstanding problems in supporting reading activities on digital systems, there are still many enhancements that could be made to the system to improve its performance.

Faster, Touch Capable Displays

We were forced to make a number of tradeoffs about the technologies employed for the slate hardware used in the United Slates system. In the end, we elected to use slower, but lighter, and more readable e-paper displays. Practical considerations for things like battery life, arising from our plans for future deployments with these devices, shaped our decisions as well. A consequence of this choice was that not only was our screen refresh rate limited, we could not employ direct touch sensing because touch sensing overlays considerably degraded the screen image. Fortunately, these tradeoffs are not fundamental, as a number

of viable solutions are around the corner. For instance, the Mirasol display⁷ promises a fast, full-color display with energy consumption characteristics of current electronic paper. It is also important to point out that elements of the United Slates could readily be implemented on commercial tablet appliances like an iPad. For situations where writing and weight are not critical, these commercial devices could provide similar functionality as the United Slates.

Faster screens and touch control would both be valuable for improving the interactions on the United Slates. Faster displays would enable the system to more closely match PCs and tablet appliances in reading activities like skimming, where rapid presentation of content is beneficial. Furthermore, more interactive displays translate to more visualization possibilities for improving users' cognitive awareness of the state of the document, stack or reading environment.

Touch input would open the door to richer bimanual operation of the devices, allowing them to more closely approximate what is possible in a tabletop-computing environment. Projects like LiquidText [Tashman and Edwards 2011b] have further demonstrated how the combination of highly interactive displays and multi-touch can be used to support active reading. These techniques would be great additions to the system we present in the paper. Moreover, touch input can further extend the expressiveness of the techniques we presented. For instance, with touch, one can directly specify screen positions at both the source and destination slates of a Conduit operation. By adding enhanced bimanual operation, the fast mode of transfer Conduit provides can be employed in a wider variety of cross-slate operations.

⁷ <http://www.mirasoldisplays.com>

Slate Localization and Identification

For certain interactions like opening thumbnails and remote control, participants in our evaluation mentioned that indirect interaction that did not require physical reaching, like picking a device from a list, would be faster. This preference was further reflected in users' enthusiasm about controlling many devices indirectly from the PC, which could be done quickly, without any reaching required. However, in an environment where slates are similar in appearance (and/or showing similar-looking content), relying on indirect interaction could be confusing since users would not be able to reliably correspond a choice in a list to a slate in the environment. Although providing additional feedback or identification cues through multi-monitor window notification techniques [Hoffmann, Baudisch, and Weld 2008] could help alleviate this problem, we earlier mentioned it would be desirable to recover the position of the devices in the environment.

Having the positions of the devices in the environment would, for one, allow the remote control interface to render the virtual slates in a way that more closely matches the physical layout of the actual slates. But, slate localization has additional benefits. Recovering the relative positions of devices would allow richer indirect interactions, like those seen in Augmented Surfaces [Rekimoto and Saitoh 1999] which might further address problems of reaching. For instance, users might be able to select the slate on which to open a link by merely gesturing in the general direction of the destination device.

We omitted this functionality in our current implementation because existing positioning systems require infrastructural support at the room level, which would limit where the slates could be used. However, one compromise is to have the functionality only be active in select environments. A better solution, of course, is to have a localization system that does not require any environmental infrastructure. One possibility would be to use acoustic localization technology like that demonstrated in BeepBeep [Peng et al. 2007] to obtain

centimeter-level positioning accuracy or to adapt localization techniques developed for sensor networks [Broxton, Lifton, and Paradiso 2006].

Integration With Other Computing Platforms

Although the United Slates system has shown promise integrating beneficial features from a variety of existing computing configurations, it does not reproduce all of their beneficial functionality. Therefore, there should be an emphasis on seeing how multiple slates might co-exist with these other types of computing systems. For example, an actual tabletop computer would be a great complement to slates when they are available. Tabletops make working with very large quantities of information tractable. Moving things around, creating and deleting items is fast and easy on a tabletop computer. Slates, as we have shown, combine a dynamic display with high mobility. The two together would support a wide range of reading situations that include working with very large numbers of documents and multi-user reading with a shared display. Another technology should not be discounted is paper. Paper is both inexpensive and dispensable. Furthermore, there do not appear to be any technologies on the horizon that match the writing experience on paper. Therefore, digital paper technologies like PapierCraft [Liao et al. 2008] would be a good candidate to add to this integrated, multi-device reading environment.

5.7 Summary

In this chapter, we presented the design of interactions for a multi-slate environment that supports a full range of the reading activities of knowledge workers and students. These interactions concern mapping of documents to slates, a range of interactions that support of flow of information between slates, and techniques that support reading activities that occur across different venues and using different slates. From these principles, we provided new capabilities to support the reading requirements we identified.

Pairing these interactions with the slate hardware we developed allows the reading environment to acquire the positive traits of several different reading technologies. These include the portability and physicality of e-book appliances, inking capabilities of Tablet PCs, and the spatial layout possible with a tabletop computer. The electronic aspects of the system also simplify the process of working with large quantities of documents, reading in different venues, and interfacing with PCs, which are activities that tend to be difficult when reading with paper.

The evaluation of our prototype demonstrates that the new functionality the system provides would be applicable to the reading activities in the classroom and workplace. Given these positive results, we refined the United Slates system so that its use can be observed in real-world deployments.

6 Cloud-based System Infrastructure

Moving a research system from the laboratory to the real world traditionally involves a serious investment in effort. In this chapter, we describe how using cloud-based services made deploying the multi-slate system in the real world feasible. From our experiences, we believe that leveraging the widespread availability of cloud services can result in simpler system architectures that also allow for a transparent transition from the lab to the wild.

6.1 Introduction

The use of a system in real-world settings introduces issues concerning resource availability, system robustness and scalability. Addressing these issues tends to be difficult, time-consuming and labor-intensive. As a result, many systems languish in the lab when their use in the real world can provide valuable information. Cloud-based services offer scalability, redundancy, and other beneficial features that can potentially address the issues arising from the real-world deployment of systems. However, until recently, it has been impractical to assemble the necessary hardware and software infrastructure to deploy robust cloud-based resources for research prototypes. Moreover, ensuring that prototype devices had access to these cloud-based resources was also a tricky proposition.

Over the course of developing the United Slates system, a number of new cloud computing platforms have appeared. Platforms like the Google App Engine, and Amazon's Elastic Compute Cloud (EC2) have streamlined the process of deploying a cloud-based service. Crucially, these platforms free researchers from having to invest in and set up computers, network connectivity, and software packages. Furthermore, the widespread availability of Wi-Fi and cellular connectivity in most locations means that devices can easily leverage these cloud-based services.

The effects of these developments have mostly resulted in an explosion of new web services, but we believe these developments are also significant for system builders in the HCI community. The comparative ease of prototyping systems with a cloud component nowadays argues for implementing a significant portion of system architecture using cloud-based services from the start. Doing so incurs relatively low development overhead, can potentially simplify the architecture, provides useful functionality and, most importantly, automatically readies the system for real world use. Basically, with minimal added effort, it is possible to build a system that can be used in both laboratory and real world settings.

For the United Slates system, we first identified several real-world deployment challenges that are representative of real world use of prototype systems. These challenges necessitated changes to the system architecture we employed in laboratory testing. In response, we developed several cloud-based support components that enabled the United Slates to be deployed outside of the laboratory setting. These components employ cloud services like the Google App Engine⁸, Dropbox⁹ and Amazon EC2¹⁰ to add accessibility, resiliency and scalability to the United Slates system. We discovered that due to existing technological constraints the cloud could not be used exclusively for highly interactive systems.

Therefore, the description of the architecture also provides examples of optimizations that can work around these issues.

Our efforts in this area demonstrate that a multi-slate reading system like the United Slates can be realistically and readily implemented using off-the-shelf technologies. The broader implication of our work is that the cross-platform interconnectivity afforded through the Internet coupled with the ease of deploying always-on services in the cloud are helpful developments for solving issues arising from the real-world deployment of research systems.

⁸ <http://appengine.google.com>

⁹ <http://www.dropbox.com>

¹⁰ <http://aws.amazon.com/ec2/>

6.2 Challenges of Real-World Deployment

The United Slates system was initially used for controlled lab studies. In these studies, we could do things like stipulate which devices are available, ensure there was network connectivity and specify what documents are going to be used. Also, in the lab, a well-defined experimental protocol along with piloting minimized the chance of unforeseen problems. When problems did occur, we could immediately offer a technical workaround. Finally, experimental tasks were short, which means that problems of scale generally do not have a chance to manifest themselves. When we needed to deploy the system with users in the wild, we encountered major challenges that required us to re-architect the system architecture.

6.2.1 Challenge 1: Availability of Resources

With the multi-slate system, we wanted to make sure that updated document data was available on all of the slates. With versions of the United Slates system used for in-lab testing, we got around this problem by preloading documents we needed onto the slates and ensuring that slates were always connected and receiving changes.

The situation when using the devices in the real world is radically different. First, we wanted to ensure that users could place documents of their choice on the devices, making pre-loading impossible. We also wanted users to be able to use slates interchangeably. Forcing the user to check whether a device contained desired materials would have added considerable overhead to the operation of the devices. One solution was to have users manually copy documents to each slate, but that would be repetitive and tedious. Manually populating the contents of the slates would also be problematic if slates are stored at a remote site.

To deal with these problems, we use the cloud-based Dropbox service as the mechanism for transferring and synchronizing data between the computer and slates (Figure 24).

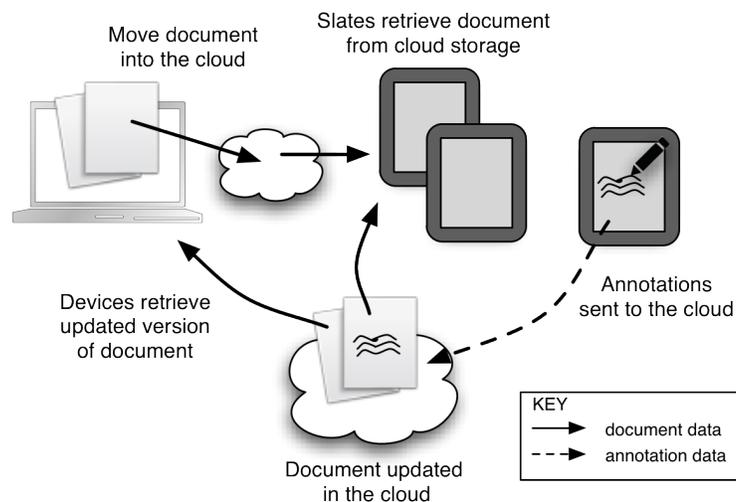


Figure 24 Example lifecycle of a document when using the cloud-based infrastructure.

Dropbox works by synchronizing designated files and folders against a central copy stored on the Dropbox servers. A client on PCs and other devices synchronize against the copy in the cloud. The Dropbox approach maximizes availability of data in the system. By automatically synchronizing slates against a single master copy of a file in the cloud, we ensure that the user does not need to perform extra work to copy files onto the devices regardless of how many devices there are. Moreover, devices have access to document data without other devices having to be present or turned on.

6.2.2 Challenge 2: Robustness

In the wild, how users operate the reading system cannot be anticipated ahead of time. Making matters worse is the fact that users are geographically distributed and can use the system at unconventional times, meaning that researchers cannot resolve problems in person. When problems occur, they disrupt the user's work. If disruptions happen frequently and interrupt the user's work significantly, it is likely that the user will decide to stop using the system. Therefore, ensuring that the device is robust against unexpected configurations and problems is of paramount importance.

One of the areas where robustness was a big factor was in providing reliable inter-device communication. In the real world users would employ a variable number of devices for different tasks. Our previous solution where a special device would route traffic between the others in a Bluetooth network was impractical. We also needed a mechanism to allow us to diagnose and fix problems quickly and from a remote site in order to minimize downtime when problems did occur.

We dealt with connectivity issues in the system by moving the communications from Bluetooth to Wi-Fi. Wi-Fi access points are largely ubiquitous and users have familiarity connecting to these access points. For simplicity, and to get around firewalls and routing policies, we route data through the cloud. Slates exchange messages by making HTTP requests with a server running on Google's App Engine (GAE) platform. Slates receive streaming updates from the cloud via a persistent XMPP channel. To make the cloud-based system even more robust to the limitations in the latency and bandwidth of current networks, we introduced a few optimizations that bypass the cloud to guarantee good system performance.

For situations where no Wi-Fi access points are available, we have two fallback mechanisms. First, we provide users with a mobile hotspot that provides connectivity for all of a user's devices to the cloud. Second, even if there is no connectivity, user interactions are cached and are sent to the cloud when connectivity is restored.

Additionally, we provided a mechanism to store and distribute system updates via the system running on GAE. We were able to correct problems remotely by uploading a new version of the system software onto the server and then asking users to update their devices.

6.2.3 Challenge 3: Scalability

Deployments operate on a longer time scale than lab studies, which introduces issues of scalability. Having the devices for longer would likely mean that users will put more

documents onto the device, add more user-generated to each document, and encounter large documents that the slates cannot easily handle. Moreover, the fact that deployments take a long time means that running one participant at a time can be impractical from a scheduling perspective. The alternative, deploying to several users in parallel, requires the system to support simultaneous users.

To cope with issues of scalability, we implemented a cloud-based document processing service that used computers on Amazon's Elastic Compute Cloud (EC2). The document processing service was responsible for resource intensive activities like applying annotations to documents, sending the modified documents to Dropbox, generating thumbnails for documents, and pre-rendering documents that are slow to render on the slates. Since the server running on Amazon's EC2 was more powerful than the slates many documents could be processed in parallel, reducing the amount the user needed to wait for documents to be updated with the latest marks. Overall, moving these activities into the cloud helped to ensure slate responsiveness even if the user was working with a large number of highly complex documents.

To allow us to more easily manage multiple users that each had multiple devices, we created a web-based dashboard for performing administrative tasks. Tasks performed through the dashboard including adding new users, assigning devices to users, checking the status of each device, and monitoring the other components in the cloud infrastructure. The dashboard stores all of the configuration data in a centralized location, which reduced the need to manually edit configuration files on each device.

6.3 Added Benefits of a Cloud-based Infrastructure

In addition to addressing the issues with real world deployments covered above, shifting portions of the United Slates system to the cloud-based infrastructure provided other desirable qualities that were conducive to both in-lab and real world use.

6.3.1 Simplified System Architecture

Decentralized peer-to-peer systems require extra components to allow devices to discover each other and coordinate their roles. Using a centralized server in a known location to route data means that devices can communicate without the need for these additional components. The always-on nature of cloud services also means that there is a reliable data store for maintaining system state, precluding the use of a more complicated a distributed storage scheme.

6.3.2 Inherently Cross-platform

There is significant diversity in the mix of computing platforms that potential users employ. Developing custom software solutions for each of these platforms can be error prone and time-consuming. Moving more functionality onto the cloud means that a single component can be shared across all of these platforms. Most computing platforms now provide robust library support for commonly used protocols for communicating with cloud services. The cross-platform aspects of the cloud also mean that we can more easily reuse portions of the system even if the hardware platforms change. Finally, since the code is in the cloud, fixing bugs is simplified because deploying a new version of the service can be done transparently and fixes problems for all users of the system at once.

6.3.3 Receiving More for the Same Amount of Work

Although similar functionality could be realized with a server we personally administer, the amount of effort required write the code would have been nearly identical. However, setting up a server of our own would have incurred additional time and work. Moreover, relying on the robust infrastructure that underlies Google's App Engine conveys more stability, ability to dynamically expand capacity if needed, and superior connectivity to the rest of the Internet. For these reasons, we believe that designing systems from the outset to

use cloud-based services, even for systems that may not need to be deployed in the real world, is a prudent approach.

6.4 Implementation Details

The revised United Slates system architecture for real-world deployment is shown in Figure 25. PCs and slates in the system access document data by syncing their local filesystems with Dropbox. The slates and PC communicate with the Google App Engine (GAE) service (annotation and system state data) and amongst each other (UI operations requiring low latency). Annotation, thumbnail creation, and preprocessing commands are first sent to GAE, which then posts tasks to the EC2-based document processing service. The document processing service then retrieves, updates, and uploads documents back to the shared Dropbox storage.

6.4.1 Dropbox File Transfer and Synchronization

On the user's PC, we use the native Dropbox client, which creates a special folder inside which all files are synchronized onto Dropbox's cloud storage system. The slates use the Dropbox Mobile API to access the files in the user's Dropbox folder. Dropbox was particularly convenient to use because they provide clients for all modern operating systems. Also, many users already have the client installed on their computers.

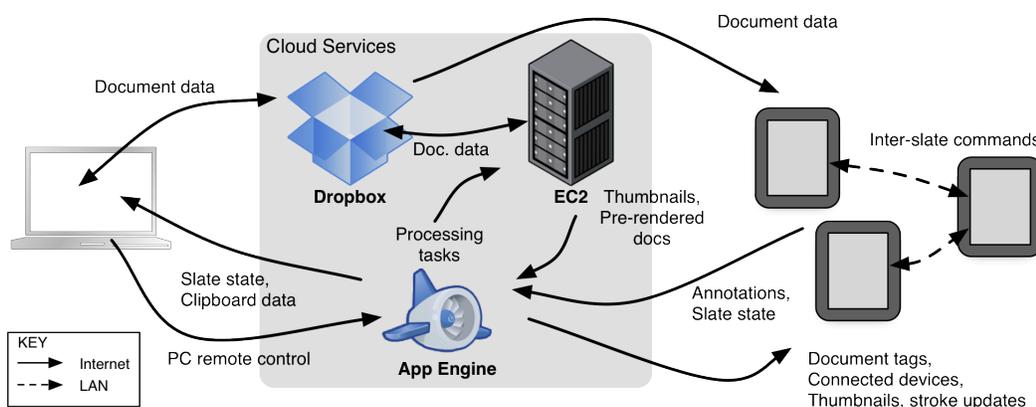


Figure 25 The flow of data through the cloud-based United Slates infrastructure.

6.4.2 Routing Data Through the Cloud

We created a RESTful [Fielding 2000] web service running on the Google App Engine.

The specific pieces of data that go through the cloud are device state changes like document changes and page turns, annotations and document modifications, shared input events like command button presses and magnetic switch activations, and command parameters like hyperlink destinations or the text for a copy and paste operation.

In web-based services it is easy for the client to send data to the server but more complex for the server to push data to the client. The usual solutions involve opening a persistent connection between client and server over HTTP server push or some other protocol. Our solution is to open a second communication channel using the XMPP [Saint-Andre 2004a; Saint-Andre 2004b] instant messaging protocol because of weak support for HTTP server push when running a program outside of a web browser (many implementations rely on cookies). Each user has a unique XMPP user name and each device they use is given a random resource ID as a slate first powers on and connects to the App Engine service. The use of XMPP to communicate between devices is based on the implementation of Personal Information Environments [Pierce and Nichols 2008].

6.4.3 Cloud-Based Document Processing

In the specific implementation for the cloud-based United Slates system, annotations are relayed to the App Engine service and serialized into a work queue. The cloud-based processing service running on the Amazon Elastic Compute Cloud (EC2) polls the App Engine for tasks, retrieves the document from Dropbox, applies the annotations and then uploads the modified file back to Dropbox. The annotation system assigns annotations timestamps and UUIDs, which are embedded with the annotations inside the PDF. This allows slates to unambiguously communicate to the annotation service which items to add,

modify or delete from the document. If the annotation service crashes, annotations are journaled in the App Engine and can be applied after the annotation service is restarted.

Our implementation has several nice properties. First, the process of adding annotations to PDFs involves making changes to binary data, which can be difficult to reconcile if the same document is being modified at the same time. Having a single process modifying documents even though there are multiple slates bypassed numerous synchronization problems. Second, most networks tend to be asymmetric: it is most often the case that downstream bandwidth is significantly higher than upstream bandwidth. Therefore, having slates upload modified PDF documents, particularly very large documents, after every set of changes would be prohibitive. The implementation means slates only need to send annotation data upstream, rather than the full document. An unforeseen benefit of our decision to use EC2 is the fact that Dropbox uses Amazon's S3 service for storage. Since our annotation service and Dropbox's storage servers are both on Amazon's network, the update process is nearly instantaneous.

For thumbnail processing, we generate the thumbnails once and then store the resulting image in the GAE, which is shared to all slates, reducing repeated work. The same strategy is used for creating pre-rendered page images for problematic documents that were slow to render on the slates.

6.4.4 Optimizations

To avoid the latency that comes from having to communicate with the GAE server and then waiting for the server to forward the message via XMPP to the destination, we allow devices to directly open connections to each other to invoke remote procedure calls. We moved our RPC system to JSON-RPC to ensure cross-platform compatibility with future devices. Since broadcast based discovery services like Bonjour can be blocked by network policies, we use the GAE component as a discovery service to find connected

devices. Also, the original cloud-based routing mechanism is still used if it is not possible to create a connection between two devices. The heuristic employed is that if two devices are on different subnets or if the RPC connection fails, then the communications will fall back to the cloud-based routing.

The second obstacle we encountered was the fact that Dropbox's mobile API forces devices to download a complete copy of a modified file. For smaller documents, this is not an issue, but we discovered that many users were using PDF documents that were around 15-30 MB. Although the document sizes were insignificant for the cloud-based services, repeatedly downloading several copies of a 30MB file over the mobile hotspot stressed the network. Our workaround for this was to store all annotation data in a separate file made available through the App Engine. The annotation file would be updated every time the document processing service modified the PDF. Slates that supported the optimization would download the data file instead of the full PDF. This workaround is not ideal however, because changes to the PDF made on the user's PC do not go through the annotation service. However, since most users do not annotate PDFs on their PC, this was not a significant problem.

These optimizations were designed to work around real-world constraints. We believe that as the communications infrastructure improves, making things like bandwidth less of an issue, the need for these optimizations will decrease.

6.5 Summary

Given the low barriers to entry and ease of implementing cloud-based services these days, we believe it is a good idea to consider incorporating cloud-based components into prototype systems from the outset. Cloud-based components have the potential to simplify system architecture, increase flexibility and ready the system for real-world deployment. By using the cloud to share and distribute files, support system interactions, process

documents, and manage users and devices, we enabled the United States system to be reliably used outside of the lab. The fact that users successfully used the devices at home, the office, coffee shops, conferences, road trips and during cross-country air travel attest to the reliability and robustness of the system.

7 The Role of Physicality in Electronic Reading

In the United Slates system, we selected hardware slate devices as the principal surfaces for reading and writing. We made our decision based on the premise that providing physicality would be beneficial for reading activities. Although previous research [Terrenghi et al. 2007] hinted at differences arising from physicality (or the lack thereof), these effects have so far been untested for reading in a laboratory setting. In this chapter we detail an experiment we conducted to determine the role physicality of reading surfaces plays for a broad spectrum of active reading activities.

7.1 Introduction

One of the central reasons why the United Slates reading system uses the multi-slate configuration is to allow users to distribute content across multiple independent reading and writing surfaces in a manner similar to what is possible with paper. Having these different reading and writing surfaces helps users to lay out information spatially, interleave reading and writing, and view multiple pieces of information at once, all of which are essential for active reading.

Independent reading and writing surfaces are not exclusive to paper and the multi-slate configuration. Large horizontal screens (LHS), like the Microsoft Surface¹¹ and other tabletop computing devices also enable one to work with multiple pieces of content on independent surfaces and to spatially arrange the surfaces bimanually on a horizontal plane.

The properties of these surfaces (sheets of paper, windows, and slates) have important differences. The primary difference between these three implementations has to do with the

Portions of this chapter have been submitted for publication by the author, François Guimbretière and Abigail Sellen to the *ACM Transactions on Computer-Human Interaction (TOCHI)* and are under review.

¹¹ <http://www.microsoft.com/surface/>

physicality of the reading surfaces. On paper, surfaces are physical sheets of paper that are tactile and can be picked up. With slates, the surfaces consist of a collection of physical tablets that are tangible like paper, albeit thicker and heavier. Windows on an LHS have no appreciable mass or thickness, and are fixed to the plane. The differences in physicality extend to the mapping of content to objects. On paper, there is a one-to-one mapping between a page of content and a physical object. With slates, each slate can represent multiple pages of content. Finally, the LHS is one physical object with content mapped to virtual divisions on that object.

Ethnographic research has suggested that physicality and mobility play strong supporting roles in the reading experience, which is why we opted to develop a reading system based on slate devices. However, the effects of physicality have not been tested in a systematic fashion. The research that is most relevant, Morris, Brush and Meyers' [2007] comparison of reading technologies, did not attempt to isolate and analyze the effects of physicality. In response, we set out to more closely analyze how these distinct technical approaches to providing multiple reading and writing surfaces may or may not impact active reading.

The specific goals we hope to achieve are the following: (1) to compare the user experience of reading from these electronic implementations with paper, (2) to characterize the differences between these electronic implementations, (3) to determine what aspects of each system still require improvement, and (4) to highlight the tradeoffs of each system so that designers can make informed choices about the most appropriate technology.

To accomplish our goals, we conducted a controlled laboratory experiment in which we invited participants to perform three reading tasks using all three active reading configurations (Paper, LHS, and Tablets¹²). We collected user ratings of the systems and

¹² At the time this work was conducted, we referred to slates as tablets. Therefore, for the remainder of this chapter, we will be using the term tablet to keep the terminology consistent with the original experiment materials and user feedback.

conducted a detailed interview about the strengths and weaknesses of each system. We also analyzed the video recordings of our participants performing these reading tasks to find similarities in the way participants completed the experimental tasks. For instance, we noted commonalities between the layouts of materials that participants created. Although, overall, users preferred the Paper and LHS configurations, our evaluation identified distinct advantages for every type of surface. For instance, Tablets and LHS both were judged neater to operate, whereas Tablet and Paper surfaces were superior for layout. For example, Tablet and LHS allowed users to lift the surfaces off the plane of the desk, and also supported eyes-free rearrangement. Finally, the lightweight surfaces in the Paper and LHS conditions were more easily and quickly moved around. Our analysis also showed that although LHS and Paper were the preferred configurations, users were polarized into pro-Paper and pro-LHS camps. For the pro-Paper camp, Tablets were preferred over the LHS. This result suggests that the physicality was beneficial and that Tablets are better than LHS as an alternative to paper when users prefer paper. The pro-LHS camp tended to prefer Tablets to paper because they were more similar to the LHS. Although Tablets did not take the top spot, we believe that Tablets do combine positive aspects of Paper and LHS. In that light improvements to the Tablets could produce a compromise solution that would appeal to the broadest audience.

One of the more surprising results was that there exist properties that were judged to be beneficial by some users and detrimental by others. One example is our finding that certain users enjoyed the increased layout space afforded by physical surfaces. Other participants preferred the restricted space of the LHS because it was more compact and did not overload the participants with layout options. The fact that subsets of users regard what others consider to be limitations to be beneficial (and vice versa) indicates that there is no single “best” system. Rather, designers should create systems that are flexible enough to selectively occupy either regime.

Our findings suggest that there is no single “best” platform for performing reading activities. Although Tablets offer a good middle ground, providing functionality that mimics virtual surfaces can be beneficial. For instance, we propose interactions for tablets that preclude the need to physically move tablets around given that tablet movement was judged to be more difficult than paper or windows. Additionally, our analysis suggests that hybrid devices that mix Tablets and LHS and Paper are a promising area for further exploration.

7.2 Comparing Tablets, Large Horizontal Screens, and Paper

With their support for multiple reading and writing surfaces that can be arranged independently, both Tablets and large horizontal screens (LHS) are better candidates for a paper substitute than more traditional technologies like desktop and laptop PCs. However, it is important to understand the tradeoffs inherent to each system. To begin, we examine some of the key differences between Paper, Tablets and LHS.

Physicality - Paper distributes content on discrete tangible surfaces that can individually be picked up, moved, and otherwise manipulated. Since tablets are larger and more expensive than paper, they cannot provide a one-to-one mapping between pages of content to devices. Rather, multi-tablet setups multiplex several pages onto a small number of tangible devices. Lastly, LHS devices provide virtual surfaces in the form of windows that can be laid out spatially, but lack any tangible characteristics.

Portability - Paper, Tablets and LHS have major differences in terms of portability and how that portability scales as the number of documents increases. Small sets of paper are extremely portable but collections of large documents quickly become difficult to carry around. The portability curve for tablets is different in that each tablet incurs more size and weight than a sheet of paper, but the number of sheets that one can carry is virtually limitless. LHS hardware can span the gamut of portability, ranging from stationary systems

to being mildly portable as in the case of PlayAnywhere [Wilson 2005]. Large, foldable displays may potentially offer another method for implementing a LHS in a portable package. Like Tablets, LHS also allow the user to access a near limitless number of documents.

Versatility (Adaptability to Space) - The devices we are interested in also differ in how well they support reading in different environments. One of the key strengths of paper is the fact that it can be used in a variety of situations. Sheets of paper can be made compact but can be spread out if necessary. Tablets, by virtue of being individually quite portable, allow the user choices about how many devices to use for a given situation. LHS give maximum flexibility in terms of presentation because virtual surfaces are completely dynamic. However, LHS are more restrictive about where they can be set up and operated. Stationary LHS devices require a large empty space, while portable LHS implementations require a flat plane on which to project or lay out the display. In contrast, Paper and Tablets can be operated even when a large flat area is not readily available, such as when reading in an airplane.

Scalability of Reading Surfaces - The number of reading surfaces that a user can realistically have and comfortably operate is yet another difference between these devices. LHS arguably are the most scalable, allowing users the ability to create and delete surfaces at will. Paper comes next because paper surfaces each have a direct and fixed mapping to content. Tablets are less scalable than paper as their size and weight will make it difficult to have a large number in use at once. Surface scalability is connected to the physicality of the reading surface, but is not the same thing.

Digital functionality - LHS and Tablets offer digital functionality for which plain paper has no equivalent. Both Tablet and LHS offer dynamic electronic displays that can not only display different content but also provide visual feedback. Furthermore, LHS and

Tablets are able to easily support digital operations like copy and paste, text search, and connectivity with other computing devices. Paper can, however, be augmented with electronics (see Section 2.3.1) to provide some digital functionality.

Looking at these differences, we believe that the effects on the reading experience stemming from the physicality of the surface are the most important to determine because physicality will always be present as a factor to some degree. Whereas portability or adaptability may not apply to scenarios in which people are relatively stationary, it is impossible to think of a scenario where the differences in tangibility are invisible to the user—it is the difference in tangibility that defines the devices we are testing. Thus, we wanted to design an experiment to isolate the role physicality plays as much as possible. We note that some of the other differences we identified are intertwined with physicality and cannot be separated entirely. For instance, digital functionality is clearly what the electronic conditions rely on to provide equivalent functionality to paper. The best resolution for these situations is to ensure that the experimental task does not artificially favor one condition due to one of the ancillary factors we identified.

7.3 Experiment Devices

The goal of our study was to investigate how differing degrees of physicality in reading surfaces can affect reading tasks, which is the defining difference between implementing multiple reading and writing surfaces on paper, on several tablets, or on multiple windows on a large horizontal surface (LHS).

To control for differences between these three device setups, we wanted to ensure that the different device configurations had some important commonalities. In particular, we wanted to make sure that users were able to easily spread out documents so that they could view multiple pages simultaneously. To ensure that participants could easily adjust the layout, we made sure that all conditions allowed the users to use their non-dominant hand

to adjust the position of pages/tablets/windows. We also ensured that surfaces displayed pages of content in full to remove the need to pan and scroll. Finally, we provided controls to allow users of the two electronic implementations to randomly access pages of content like they are able to do in the paper condition. The following sections provide specific details about the device configurations we tested.

7.3.1 Tablet Condition

The Tablet condition (Figure 26) represents a multi-slate reading system similar to the United Slates. For the experiment, we used the slate hardware we described in Chapter 4, but *not* the full United Slates system in order to maintain functional parity with the other conditions. This was done so that the effects from physicality can be studied in isolation. The tablets we used were wirelessly linked and presented a view into a single, synchronized data store. That is, annotations and other writing made on a page on one tablet would be propagated to all the other devices.

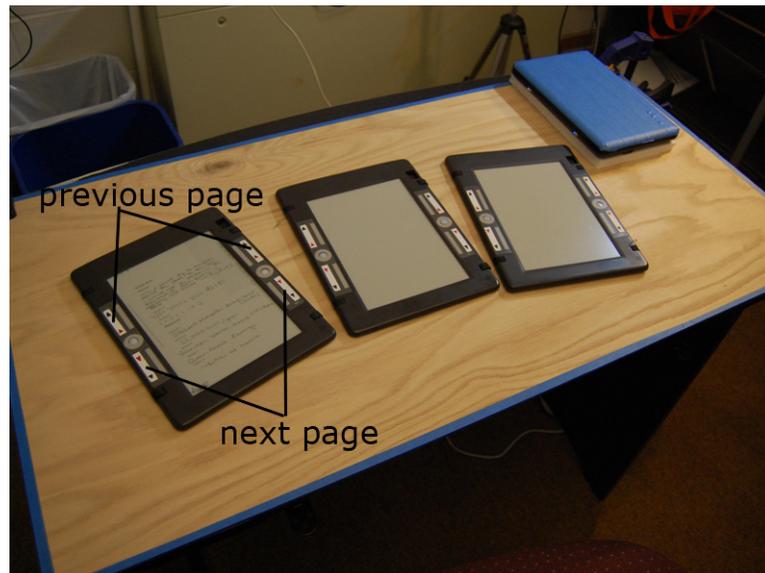


Figure 26 The slate condition. Three slates are arranged on the desk. Each slate had four buttons that corresponded to previous and next page controls. Note that the entire desk is covered with a wooden board to minimize any impact from the horizontal display embedded within the desk.

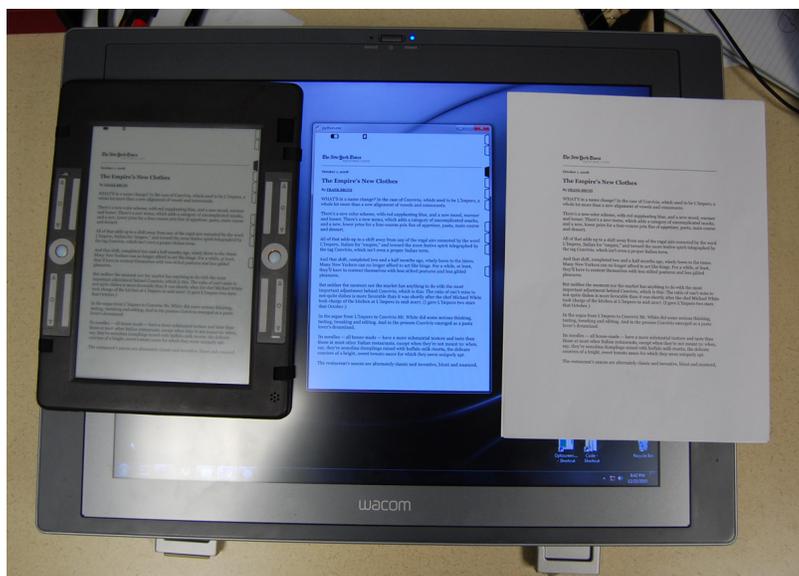


Figure 27 Side by side comparison showing the appearance of pages in each of the device conditions. Note the tabs to access pages on the right edges of the electronic reading surfaces.

With regard to controls, every tablet contains four rectangular shaped buttons, placed on the edges of the tablet. The upper buttons on each edge turn a page backward, and the lower buttons turn a page forward. Alongside each page is a row of tabs that allow users to jump to any other page when the tab is tapped with the pen (Figure 27, left).

Documents were presented by providing participants using three tablet devices; the page content on each device was scaled to 825 pixels \times 1200 pixels. At the slates' native pixel density of 150 dpi, the physical size closely matched that of documents in the Paper and LHS condition (approximately 139.7mm \times 203.7mm).

7.3.2 Large Horizontal Screen (LHS) Condition

The LHS condition is designed to mimic a horizontal display environment, like a tabletop computer, in which the desk surface on which the user works is also the display surface. We considered several LHS technologies to use for our experimental condition. Many systems in the literature rely on rear-projecting the visual content. The main downside of a projector-based approach is that the projected images tend to have low pixel density (i.e.,

low DPI), reducing both the legibility of textual documents and the precision with which users can annotate documents. Another technique we rejected was a large LCD unit with multi-touch capabilities, which had higher resolution, but still lacked pen input. In the end we chose to sacrifice display area for pixel density and writing fidelity.

Our LHS setup used a 21" Wacom Cintiq 21 UX pen-enabled display as the display surface. We embedded the display into a desk so that the surface of the display was flush with the rest of the desktop (Figure 28). Since the Cintiq did not support touch input, we used three Optitrack V100 cameras in conjunction with a retro-reflective sphere placed on the index finger of the user's non-dominant hand (Figure 29, right) to track the position of the user's finger in 3D. Two cameras placed at table level detected whether the finger was in contact with the display (Figure 29, left). Two cameras were used to provide redundancy if the view was somehow occluded. One camera suspended above the work area tracked the X-Y position of the finger (Figure 29, left). Following a short calibration procedure, which maps the finger from camera coordinates to screen coordinates, our optical finger

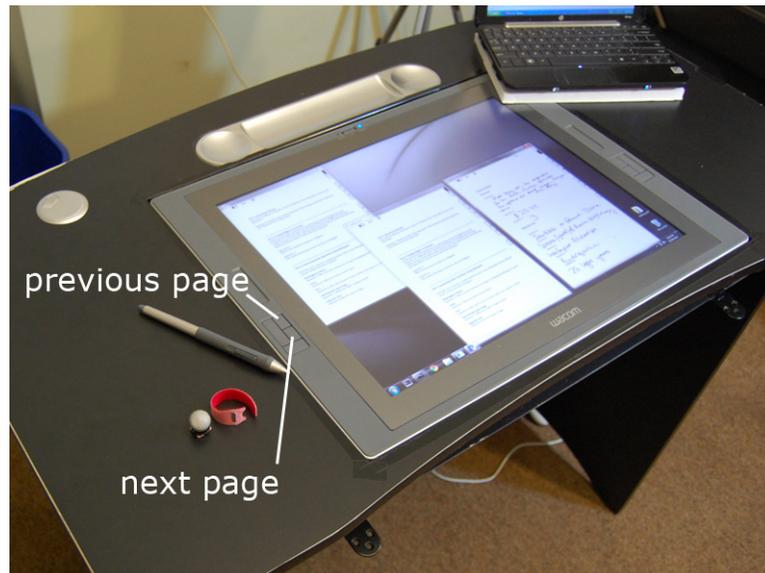


Figure 28 Close-up of the LHS condition with hardware controls labeled. Identical controls (not labeled) are on the other side of the screen. The pen, and the finger-tracking marker are on the left of the screen. Note that the laptop is raised using a piece of foam board to reduce electromagnetic interference with the pen digitizer in the horizontal screen.

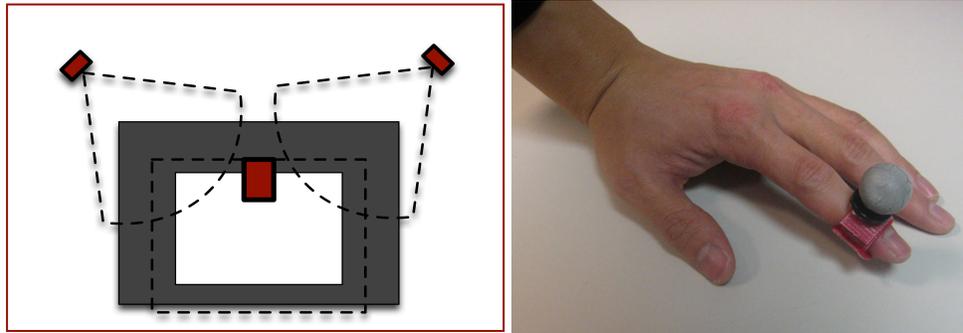


Figure 29 Left: Diagram depicting the camera setup for the LHS condition. The OptiTrack tracking cameras (red) are at the far left and right corners of the desk as well as above the working area of the desk. Dotted lines show the field of view of each camera. **Right:** The reflective marker is affixed to the first knuckle of the user's index finger on the non-writing hand.

tracking allowed the user to perform translation of on-screen windows with the non-dominant hand.

We use the hardware controls on the Cintiq display to turn pages (Figure 28). The controls turn the pages of the window that currently has focus. Users can bring windows into focus by starting to write on the window, tapping the title bar with the pen, or tapping the window with their non-dominant index finger. Alongside every page is a set of tabs that allow random jump to other pages.

The LHS condition displayed pages in three windows on the GUI desktop. Each window was fixed in size, but the positions of the windows could be adjusted. The computer ran Windows 7 and windows had corresponding window decorations. The screen had a resolution of 1600×1200 at 96 dpi; page images had margins removed and then scaled so that they were 530×770 pixels (approximately $139.7\text{mm} \times 203.7\text{mm}$) so that 3 pages could be viewed side by side. Like the tablet condition, the content presented in the three windows was synchronized. We included a 500 ms page jump delay to match the refresh delay the user experiences when jumping between pages in the tablet condition.

7.3.3 Paper Condition

The Paper condition consisted of the experiment documents printed on paper. In addition to the experimental reading materials, we provided participants with a basic pen.

The documents were printed single-sided on standard white paper. Pages that related to the same experimental subtask were stapled together to match the tab grouping in the other conditions. Pages were scaled to 75% and centered using the copy machine; the size of the text in the resulting copies matched the size of the content in the other two conditions, and resulted in final top and bottom margins of 57 mm and left and right margins of 45 mm (Figure 27).

7.4 Experiment Tasks

Selecting tasks for the evaluation was of paramount importance to maintain external validity. As we have discussed, previous research has emphasized that reading comes in many forms. Depending on the requirements of any given reading task, different features of each device configuration would be taxed. Thus, selecting a broad range of reading tasks ensures that we obtain a more complete picture of the strengths and weaknesses of a system for specific reading activities. Based on the literature describing reading activities in the workplace, we thought it important to include tasks that required extended linear reading, comparison between documents, finding information, working with an external computer, and writing. It was important that the variety of tasks be diverse enough to take advantage of the positive aspects of each condition. Also, tasks needed to be realistic and grounded in real world reading activities for maximum external validity.

We eventually settled on an evaluation process that required users to perform three reading tasks: (1) linear reading, which simulates reading a single, lengthy document plus some writing; (2) comparison reading, which requires comparison across different pages and information seeking, and (3) hybrid reading, which requires reading and writing on the

device in conjunction with a typical laptop computer. We describe the tasks in more detail below.

7.4.1 Linear Reading

Our first task, linear reading, is probably closest to what most people generally think of as “reading”. We selected a linear reading task to produce reading behaviors associated with focused reading of a specific document. The tangibility of devices may play a role in readers’ success as prior studies have shown that picking up or otherwise manipulating the media can help readers concentrate. To encourage participants to focus on reading the content as they would a longer piece of text, we kept the navigation requirements to a minimum: participants could simply proceed through the material in a linear fashion. The material we selected for the subtask consisted of restaurant reviews from the New York Times, each containing approximately 1200 words. Participants read one review per condition. Restaurant reviews were well suited for an experimental task because they were similar in length, subject matter, and were written by the same author. To ensure participants actually read the review in depth, we asked participants to write several sentences describing what aspects of the restaurant appealed to them and what aspects did not.

7.4.2 Cross-document Comparison

The second task, involving cross-document comparison, reflects the widespread activity of comparing information across documents. We observe comparison behavior when readers cross-reference, examine different versions of a document, or extract information across disparate documents. In this task, participants were presented with three lists of flights and asked to find a set of three flights (one flight from each list) that satisfied a specific criterion: the flights must all depart within a one-hour window. The flight-planning task forced cross-document comparison by scattering the flight data across three different pages.

Furthermore, the nature of the task meant that participants had to backtrack to find the correct set of flights, requiring several comparisons to find the solution. Since viewing all three sets of flights simultaneously was crucial for completing the task, we expected the task to encourage users to create layouts where several reading surfaces were visible. The flight information we used was artificially generated using a simple randomization script to ensure that participants saw unique flight schedules for each experimental condition, but that each dataset had an equal number of possible solutions.

7.4.3 Cross-device Reading

Our third reading task focuses on reading activities that are performed in conjunction with a more traditional computing device like a laptop computer. For instance, people often perform reading activities that span paper and computer. The use of heterogeneous reading technologies seems likely in light of Morris, Brush and Meyer's [2007] finding that people prefer vertical displays and keyboards for tasks that require text input. Our third subtask required each participant to use a laptop computer connected to the Internet to answer various questions using articles on Wikipedia. For this task, the questions and space for answers were on the device being tested, but the query had to be performed on the laptop. Cross-device operation introduces some unique challenges not present in the other subtasks such as requiring the user to arrange their work area to accommodate the laptop. Since the question sheet and answer sheet were distinct, users also needed to set up in a way that makes viewing the question sheet, typing and reading from the laptop, and writing answers comfortable.

One feature found in almost all electronic systems is the ability to copy and paste information. In our case users could have moved information from the laptop into the Tablet or LHS using a Distributed Clipboard [Miller and Myers 1999] or a similar mechanism. Since copy and paste would have provided the electronic conditions a

disproportional advantage that could obscure any differences arising from the tangibility of surfaces, we omitted the functionality. An alternative approach could have been for us to employ an interactive paper system like Papiercraft [Liao et al. 2008], which would have provided the paper condition with similar capabilities. However, we decided against that approach given the learning curve required for interactive paper, and the potential that it could alter the experience of reading from paper.

7.5 Experiment Methodology

Our experiment was a within-subjects design. Participants used all device configurations but the order of devices was counter-balanced using a Latin square. We created three distinct sets of documents so that, over the course of experiment, participants would not encounter the same reading materials. The order of the documents that the participants encountered was identical and was not dependent on the device ordering. This ensured that a particular document set was not linked to a specific experiment condition.

To ensure that participants could successfully complete the reading task, we performed a short training session ahead of each device configuration. If a participant was using either of the two digital systems, we first explained the basics of operating the device being tested. Specifically, we explained the two ways participants could navigate through the pages in the reading materials (sequentially using hardware buttons, or random access by tapping on tabs). We also explained how the windows/tablets provided a synchronized view into the reading material. That is, all the content used during the experiment were available on every tablet/window. Also, ink marks made on one tablet/window would show up on the others.

Using a special training document set, which was identical in composition to the materials used in the experiment, we first showed participants the first page, which corresponded to the task description page. Then, we showed participants the answer sheet on the second

page. We told participants to supply answers by writing directly on the answer sheet. We then moved back to the task description page to explain the subtasks in order.

7.5.1 Training

Linear Reading

Keeping the task description sheet visible, we brought in a sample restaurant review from the training set on a different tablet/window/page. We explained that the main objective of the task was to read and comprehend the restaurant review. We pointed out that the task required answering a single question about what they liked and disliked about the restaurant.

Cross-document Comparison

For the digital conditions, we began the explanation of the comparison task by switching each tablet/window to a different page of the three flight lists. Placing the flight information side by side served the purpose of showing the participant that one strategy for completing the flight-planning task was to have all the pages visible, rather than having to flip back and forth between pages. We did not explicitly tell the participant to use this strategy to complete the task, however. For the Paper condition, we brought out the stapled flight information pages, but did not remove the staple. With the flight data visible, we explained what it meant for a set of flights to depart within a 1-hour window. Then, we walked the user through the process of identifying a set of flights in the training set that satisfies that requirement.

Cross-device Reading

For the cross-device reading task, we switched back to the task description page to view a list of sample questions the participant must answer using Wikipedia. In the first training session, we familiarized the user with the laptop by placing the laptop in front of the participant and going to the Wikipedia front page. Using one of the questions in the training set, we demonstrated how to find an article in Wikipedia. We also explained that

the participant could use the web browser’s text search capability to speed up the process of finding answers on a page. We then asked the participant to operate the laptop so that they can assess the sensitivity of the touchpad used for pointing. In the training for the second and third conditions, we omitted the laptop familiarization process except to reassure users that it was fine to place the laptop on top of the display in the LHS condition.

After explaining the tasks, we reminded users about the ordering of the documents (task description, followed by answer sheet, followed by the documents required for each subtask, followed by a blank page of scratch paper). We also informed users of the tablet and paper condition that they did not need to keep the tablets or paper sheets on the desktop. Additionally, we told participants using paper that they were free to do anything they wanted to help complete the task including removing the staple.

7.5.2 Starting Setup

For the Paper condition, the initial setup provided all the materials for the experiment assembled in a stack and placed at the upper left corner of the desk (Figure 30, left). For the Tablet condition, all tablets were set to the first page and also placed in a stack at the top left corner of the desk (Figure 30, middle). For the LHS condition, all windows were initialized to the first page and overlapped at the top left corner of the display, so that only one window was fully visible and the rest had their title bars showing (Figure 30, right). The rationale behind these initial arrangements was to ensure that users rearranged the

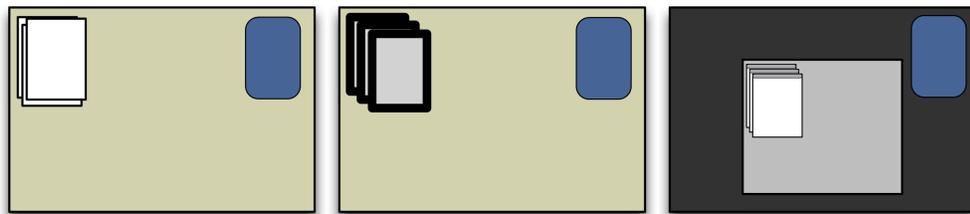


Figure 30 Starting configuration of the paper, tablet and LHS conditions.

materials at the start of the experiment rather than relying on the status quo. For a similar reason, the laptop was placed on the far right corner of the desk.

7.5.3 Data Collection

All users were filmed as they were completing the experimental task. Following each condition, participants were given a set of rating scales on which they were to place a mark describing how easy or difficult it was to perform various reading activities. The activities we asked users to rate were: planning and arranging materials, finding information, comparing information across multiple pages, marking up and annotating, and interleaving reading and writing. At the end of all three conditions, we asked participants to rank their preference for devices with respect to different aspects of the aforementioned reading activities. In the ranking of preferences, we additionally had users rank comfort, perceived speed, and overall preference—features that are difficult to judge in isolation. We then asked users to explain why they provided the preference orderings they did. Finally, we concluded with a brief interview that included questions about a participant’s background, the types of reading the participant performed most frequently, and whether they thought the experimental tasks reflected real-world reading tasks. Following the experiment, we scrubbed through the video and annotated points in the video in which the user changed the layout of the surfaces in their workspace, adjusted surfaces, wrote, lifted materials off the work surface, and turned pages.

7.5.4 Participants

A total of 13 participants (8 females and 5 males) took part in our experiment. All of our participants were right handed (although we did not specifically recruit for right-handed users). We discarded the data of one female participant because her command of the English language was considerably weaker than that of the other participants; we were concerned that she struggled during the reading tasks, possibly affecting her ability to

effectively judge the condition being tested. Furthermore, the participant had trouble communicating her subjective impressions of each system at the end of the experiment.

7.5.5 Analysis

Our analysis examined two sources of data. The first source was video footage captured during the experiment. To analyze the video, locations in the video stream where participants started a task, completed a task, or adjusted the layout of materials on the desk (or screen) were noted and video frames from these times were extracted so that these layouts could be compared with each other.

The second piece of data comes from the preference rankings and information gleaned from interviews conducted at the end of the experiment. The user preferences had the nice property that all users provided identical amounts of data and could be directly compared. The reasons for the preferences were qualified during the user interviews that followed users' assignment of rankings.

The analysis started with the more objective data from the video frame grabs. From the frame grabs commonalities between layouts or observed behaviors were identified. Then, to help explain why users were performing the observed behavior, the subjective user ratings and interview responses were examined.

7.6 Results

To help the reader visualize what users were doing, especially the reading surface layouts that users created, we have included the extracted frame grabs from the experiment video for each of the three tasks. The frame grabs are ordered on a 2-D plane based on the layout of the reading surfaces: increasing horizontal spread along the X-axis and increasing vertical spread along the Y-axis.

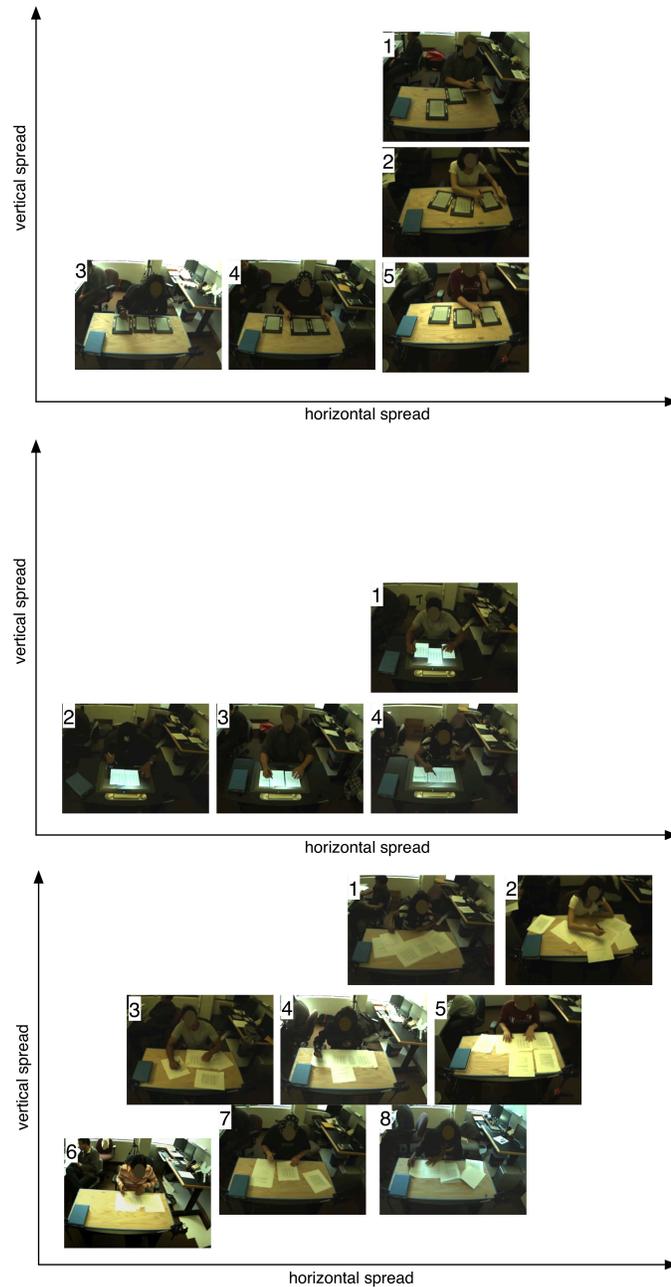


Figure 31 Range of layouts users produced in the linear reading task (Task 1). Conditions from top to bottom: Tablet, LHS, Paper.

7.6.1 Task 1: Linear Reading

Users started the linear reading task (Figure 31) by moving the pile of materials from the starting position and then accessing the material to read. On paper, users either prepared their workspace immediately, laying out the answer and question sheet next to the restaurant review bundle, or just took out the review bundle and started reading. For those

that prepared their workspace, some kept the pages for the other experimental tasks underneath the pertinent pages, while others created a separate “unused” pile. In the Tablet condition, users moved one, two or all three tablets in front of them to start the task (5 users left one tablet inactive, 2 users left two tablets inactive). On the LHS, all participants moved at least two of the windows, with some (P1, 8, 9) moving all three surfaces.

Layout Characteristics

Through the course of the task, participants performed minor adjustments to the surfaces. In all conditions, users would shift the reading and writing surfaces closer to their bodies in performing the reading and writing tasks required in the task. Another adjustment that the devices had in common was moving the reading surface away from other surfaces to isolate it. Tablet and LHS users frequently did this by shifting surfaces to create more space.

Tablet and Paper also lifted the paper off the plane of the desk. Six participants spent time reading while holding the papers above the desk (P1, 5, 6, 7, 10, 11), and 4 participants held tablets in hand (P1, 2, 7, 10). Naturally, the windows on the LHS could not be picked up. Participants did not explicitly comment on the ability to hold up the paper documents but did about the tablets. All participants who held up the tablets did so for reasons of comfort. Two users (P5, 9) who did not hold up the tablets during the task mentioned during the interview that although they left the tablet flat on the desk, they were sure there would be situations where they would want to tilt or hold the tablet. One interaction that only occurred in Tablet and Paper conditions was adjusting the reading surfaces *during* the reading process. Tablet and Paper users shifted the position of their surfaces as they were reading so that they would not need to change the direction of their gaze. These results, taken in aggregate, suggest that holding up materials to read is a deeply ingrained reading behavior that results in enhanced comfort. As such, it stood out as a beneficial feature in the conditions with physical surfaces.

Writing

The writing activities during linear reading included marking up the article and responding to questions about the review. Some users annotated the restaurant review heavily with underlines and circling while others did not make annotations in any of the conditions. However, all users were required to provide written answers to the question about the restaurant review at the conclusion of the linear reading task. For the writing task, all users supplied answers on a separate surface that was distinct from the one displaying the reading materials, regardless of whether they consulted the article during the writing process. The fact that users prefer to set up a new dedicated writing surface suggests that users assign a role to a device, which they are reluctant to alter. This recalls the behavior of users in multi-monitor environments, who give monitors special roles [Grudin 2001]. Lastly, many users in the Paper and Tablet conditions rotated the surfaces while writing but rotation was not possible in our implementation of the LHS. Based on that, we believe that rotation would be an important feature for writing activities on the LHS in general.

7.6.2 Task 2: Cross-document Comparison

Layout Characteristics

In the electronic conditions, with the exception of P8 and P10, participants performed the comparison task by laying out flight schedules side by side in order to find the correct set of flights. Tablet users who did not have three tablets in front of them after Task 1 moved and arranged the tablets so that all three were available. Users additionally pushed the tablets together so that their edges were touching. LHS users also arranged their windows side-by-side; some users spent a great deal of time adjusting the windows so that all three would be visible within the area of the display.

There was more creativity and diversity in how participants laid out the materials for the task in the Paper condition (Figure 32). For instance, with Paper, we observed that users

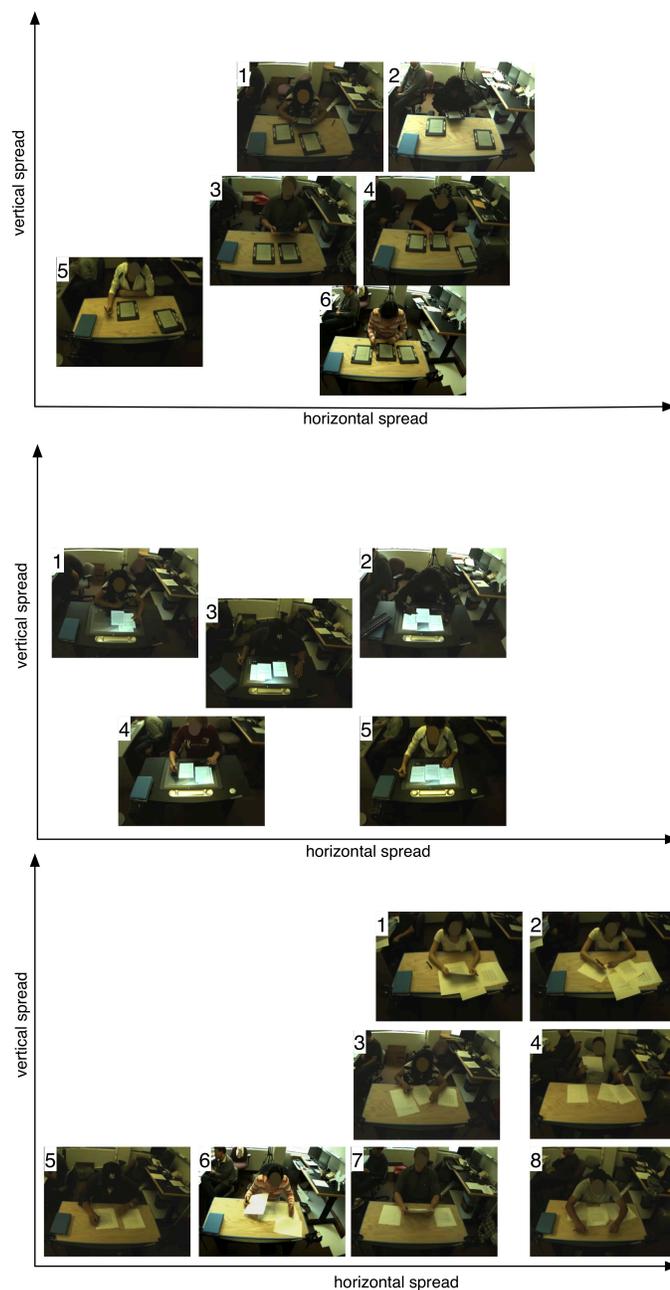


Figure 32 Range of layouts users produced in the comparison reading task (Task 2). Conditions from top to bottom: Tablet, LHS, Paper.

took advantage of the fact that paper had more reading surfaces than the electronic conditions and used them accordingly. Many paper users did not feel the need to un-staple the packet of flight schedules. Instead, these users chose to rapidly flip between pages. One user mentioned it was very easy to “stick three fingers into the sheets and jump around quickly.” Oftentimes users took advantage of the fact that they are able to look at multiple

pages at once by folding a portion of one sheet of paper out of the way. In general, users who did not un-staple pages had the flight schedules in hand and we observed users lifting the packet of schedules closer in order to study the schedules.

The advantages of paper having many surfaces also extended to those users who removed the staples, however. P2 summarizes the situation well during the interview when she says, “Paper was the best for comparing because I would have 4 screens [and this was useful] because I needed all three schedules to answer the question.” P2’s comment indicates that the ability to have more surfaces (than the three we used for the experiment) would be useful in certain situations (See Figure 32 Paper #1, 2, 4, 8).

One behavior that emerged that arose from the observed pattern of paper use was moving multiple sheets of paper at once. Moving pages simultaneously allowed things like shifting the entire workspace to a more comfortable position, or isolating two pages that the user was comparing away from other pages. Users used different strategies to move multiple sheets of paper simultaneously. Techniques included using both hands and many different fingers to move up to 4 sheets of paper together, or pushing down hard on overlapping areas of pages so that the friction was transferred to all the pages underneath the finger. We did see users move multiple tablets simultaneously for the same reasons as for paper, but to a lesser degree. This could have been because there were simply fewer tablets to move, or that they were more difficult to move. Simultaneous tablet movement was performed bimanually, but in one case a user took advantage of the physicality and thickness of the tablets to push one tablet against another, which in turn pushed against the third, in effect moving all three tablets together.

Surface Spacing

One major difference between the device configurations that affected comparison activities was the spacing between pages being viewed simultaneously, which was a function of the

page margin, device bezel, or window border and the total available space on which to layout the reading surfaces. Some users liked the fact that surfaces on the LHS could be placed very close together. P4 said, “With the single surface, all the files were there so [it was] better for looking at multiple things.” P12 expressed a similar sentiment, mentioning that “there’s less distance between the windows, things are closer, so it’s easier to look.” Others felt that the space was too cramped, with P9 going so far as to close one of the windows permanently at the end of the comparison task.

Tablets, with their thicker bezels gave users the impression that they were more spread apart. Some participants thought the spread out nature was a disadvantage, forcing them to look around for the desired tablet (P4, 6, 10). Others liked the fact that tablets had better defined separation (P7, 8, 9). Paper had the interesting property of having margins that also caused the content to be more spaced out than in the LHS. However, as we explain in the next section, certain users found that overlapping sheets of paper mitigated the issues of spacing to some degree.

Overlap

Overlap was another key difference between paper, tablets and LHS that was apparent during the layout process for comparison reading. Both Paper and LHS windows could overlap each other, while users generally avoided overlapping tablets given their thickness. A large number of positional adjustments for Paper and LHS (40% and 55%, respectively) were for resolving overlap. Users were able to overlap papers to reduce spacing between sheets of paper (Figure 32, Paper #2, #3). Overlap could not get rid of paper spacing entirely, however, given that at least one of the margins was always present. One issue users experienced on the LHS was arranging windows in such a way that the overlap was initially not an issue, but then running into problems when the ordering of the windows changed (this occurred when writing on a window at the bottom the stack). In this situation, content that was previously visible would get covered up, necessitating more re-arranging.

User comments confirmed that this behavior was problematic. Both P8 and P9 both mentioned that the overlap and the resulting obscuring of information were sometimes unpredictable and confusing during the interview. In contrast, overlap was better received in the Paper condition, where users thought that overlapping sheets of paper were good because it allowed several pages to be laid out in a compact space.

At the end of the comparison task, we observed that many paper users had papers scattered across the entire desk (Figure 32, Paper #2, #5 are good examples). For the comparison task, users needed to work with up to 5 pages of content (3 flight schedules, question sheet, answer sheet). Working with this many pages simultaneously proved to be a challenge for some in the Paper condition who mentioned that paper was difficult to manage: “messy”, as one participant put it. For that reason, many users preferred the layouts produced in the electronic conditions.

7.6.3 Task 3: Cross-device Reading

Layout Characteristics

In the cross-device reading task (Figure 33), users cleared away papers, or re-arranged tablets and windows to prepare for the introduction of the laptop computer into the workspace. One clear disadvantage of the LHS was that a great number of users mentioned the fact the single surface restricted laptop positioning (P2, 4, 5, 6, 7, 8). These users all said that they treated the entire surface as a single workspace, and were averse to placing another device on the surface, even over empty areas without any content. Users noted that their decision to not place the laptop on the surface was not due to any carefully calculated reasons. Instead, users said their instinct was simply against placing the laptop over a display.

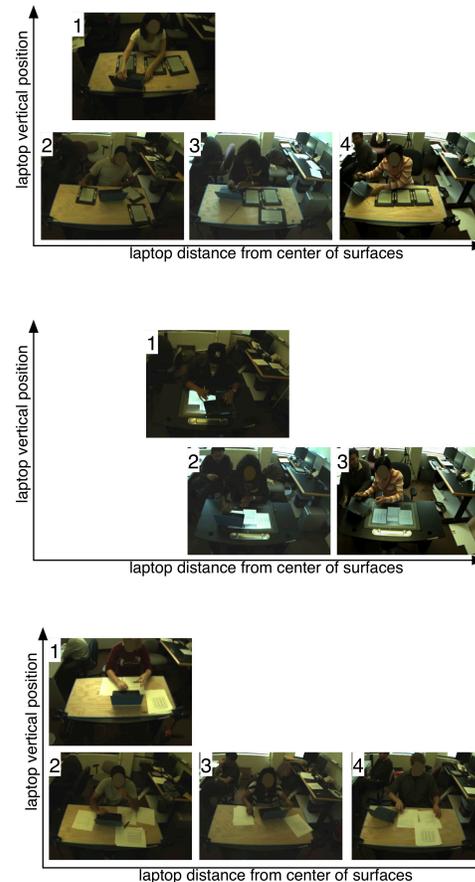


Figure 33 Range of layouts users produced in the cross-device reading task (Task 3). Conditions from top to bottom: Tablet, LHS, Paper.

Placing the laptop to the side was generally viewed to be disadvantageous. Users reported that the laptop was in a suboptimal location (either the laptop ended being too far away or off to the side) to read comfortably (P2, 7, 9, 10). In contrast, benefits of having the laptop in close proximity to the tablet and paper surfaces included the ability to arrange the workspace to support a work sequence and increased ease of interleaving reading and writing when transcribing information from the laptop to the answer sheet. P9 liked being able to “pick one of them (tablets) up, look at it and then look at the (laptop) screen” P7 thought that tablets were really good for establishing a “sequence” of operations by putting the laptop in between the question and response sheet. In this way, he could simply proceed down the line of devices to complete the task.

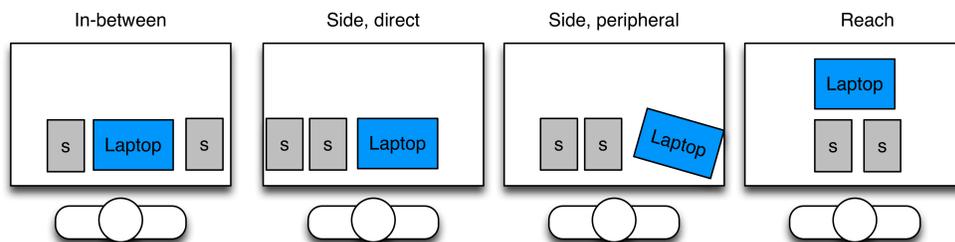


Figure 34 Diagram of layouts users created to accommodate the laptop. Rectangles labeled “S” denote a surface (window, tablet or paper).

The layouts users created corroborate the users’ interview responses. The layouts in conjunction with the laptop fell into four categories: in-between surfaces, to the side of surfaces, put in front of the user, to the side of the surfaces and off in the periphery of the user, or in front of the user but behind the reading and writing surfaces (Figure 34).

Plotting the frequencies of each laptop arrangement (Figure 35) shows that LHS users overwhelmingly placed the laptop off to the side of the main display of the tabletop, as they reported in the interview. Tablet users tended to adopt an even mix of the techniques, with the exception of the reach layout. We surmise that users’ reluctance to place tablets in the

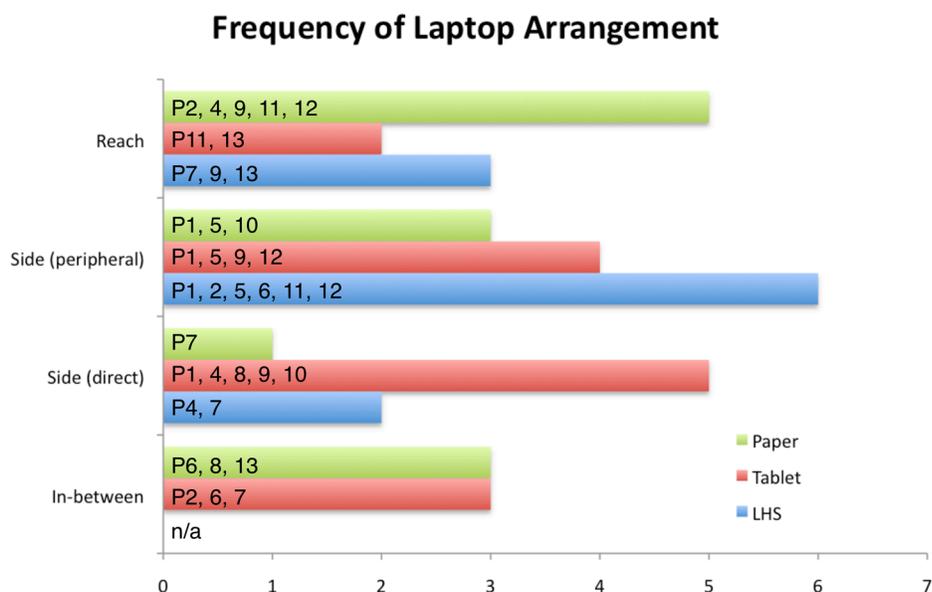


Figure 35 Frequency of each type of laptop arrangements.

“Reach” configuration is due to the fact that they were bulkier than paper. P12 described it as, “The bulkiness of the tablets plus the size of the bezel made the laptop seem far away.” P13 added, “The tablets were just kind of awkward—it was awkward leaning over the tablets.”

Adjustments While Using the Laptop

Although the layout characteristics capture the general setup of the surfaces, it inadequately illustrates the marked differences between conditions of how frequently users adjusted the reading surfaces during the task. Users moved tablets and paper around much more than they moved windows (25 adjustments for Paper, 14 adjustments for Tablet, but only 4 adjustments for LHS). The Paper and Tablet conditions saw many instances of users moving the laptop and surfaces together into more comfortable positions. This behavior may not be too surprising given the fact that Paper and Tablet support laptop arrangements where the laptop is more tightly integrated with the tablet or paper surfaces, rather than acting as a peripheral object.

Tablet and Paper users also shifted papers and tablets closer to the user’s body to write and then shifting the surface back to be next to the laptop. One particularly interesting adjustment technique involved a user pushing a tablet against a corner of the laptop in order to use that corner as a fulcrum to help rotate the tablet. Even if we exclude the rotation events users performed for writing, the number of adjustments in the Paper and Tablet conditions still exceeded those for LHS (11 and 11 vs. 4). Most notable, however, was that many of these adjustments were performed without the user looking at the surface being moved; generally users were reading from the laptop. The ability to move surfaces without visual feedback also occurred in the other tasks but it was especially prevalent in the third task, which emphasized an interleaving of reading and writing across different displays. Users did not perform *any* adjustments without looking at the surface in the LHS

condition. We believe that the tactile feedback of a physical surface factors strongly into this observation.

Reflexive Motions

Finally, one curious behavior we observed for both Tablet and Paper appears to come from the cyclic structure of the third task, which required searching on the laptop and then writing an answer. We observed users shifting surfaces closer to write, but reflexively pushing the surfaces back to their original position following the writing process. Prior tasks that had only a single writing segment did not produce this behavior. Two users employed the strategy of using the base of the hand holding the pen to push the page away in a smooth motion at the conclusion of writing. The reason we found the behavior interesting was because the subtle adjustments in the position of the surface on which the participant was writing had no obvious effect on the surfaces from which participant was reading. We suspect that the automatic return of the materials to the starting position for no apparent reason may be related to task chunking [Buxton 1986], with the replacing of the materials bracketing the end of the task.

7.6.4 Summary of Layout Behaviors

Arranging the reading surfaces was an essential part of all the tasks and users had a number of general comments about the process of laying out the reading surfaces for each experimental condition. Most users found windows to be the easiest to move mainly because they were not physical objects. Users also found paper easy to move, with “familiarity” being the main reason given. Tablets were generally judged to be difficult to move because they were heavier and thicker than the other surfaces. However, one user (P1) found that the tangible conditions (Paper and Tablets), by virtue of being graspable, simplified layout. Along the same lines, Paper and Tablets were seen to be more predictable in their positioning (P6, 7). One interesting comment revealed that the extra

effort to move surfaces resulted in some negative consequences beyond layout activities.

When describing the experience of acquiring a surface on which to begin writing, P2 explained that because the tablets were difficult to move, “I had to go out of my way to move [the tablet] in order to write and there was a risk I could lose my train of thought.”

Another common theme regarding layout was that the different device configurations provided different sets of implicit layout constraints. Users mentioned that Paper and Tablets, being physical devices, constrained the set of layout possibilities. The limited active area in the LHS condition’s display served an equivalent purpose in limiting the space of layouts.

The final topic on which most users commented during the interview was that the writing experience on paper was vastly superior to that in the electronic conditions. The fact that users preferred writing on paper was not terribly surprising given their familiarity with paper. What was surprising was that almost every user preferred Tablet to the LHS even though the Tablets and LHS used similar digitizer technology and pen tips. One of the reasons users gave for preferring Tablet to the LHS for writing activities include the surface being graspable while writing (P10) and the LHS pen being too sensitive. The former response is consistent with existing research which points to the importance of being able to continuously shift the writing surface with the non-preferred hand while writing with the preferred hand [Guiard 1987].

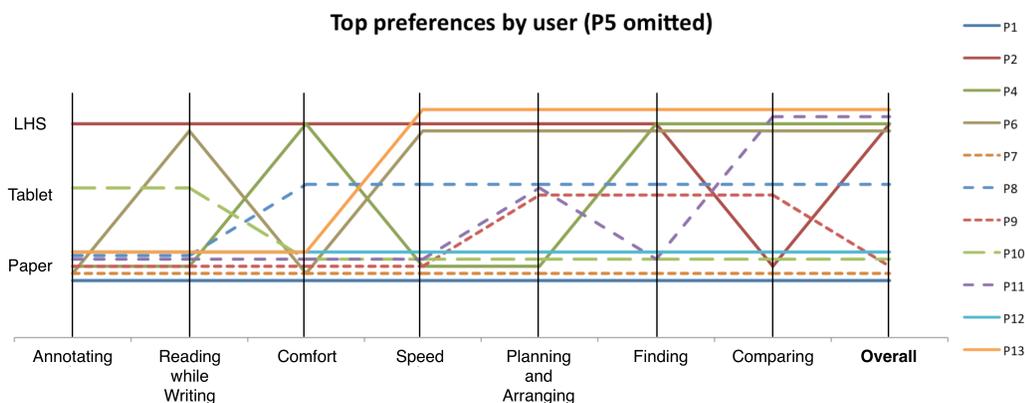
7.6.5 Overall User Preferences

Having covered the dynamics of the individual tasks in the experiment, we now discuss the overall device preferences of the participants. With the exception of P5, who felt none of the three configurations was best, and P8, who could not decide whether Paper or Tablet was best, the rest of the users were evenly split between preferring Paper or the LHS (5 users for each).

Figure 36 shows that for most users, their top overall preference is consistent with the majority of their preferences for specific reading activities. Furthermore, during the interviews users confirmed that there were no outstanding factors contributing to their overall preferences. However, a few users' overall preferences (P4, P11, in particular) appear to diverge from the preferences they give for specific activities and warrant some clarification here. P4 explained that overall she preferred the LHS because she found the advantages of the LHS to be compelling. However, given the limited amount of experience she had with the LHS configuration during the experiment, paper felt more familiar for most tasks. P11 explained that many of her issues with operating the LHS were rooted in the LHS's extremely poor writing experience, which carried over to activities that did not specifically deal with writing. As a result, her overall preference for the LHS was qualified with the comment that "with a better pen it (the LHS condition) would be perfect."

7.6.6 User Sub-populations

Since the overall preferences aggregate users' individual scores for devices across a range of reading activities, they somewhat oversimplify the true preferences of the users. One



Notes: P5, who did not have a top overall preference, was omitted. Ties were resolved by selecting the point that coincides with overall top preference. The preferences are sorted with activities that skewed towards paper first. The sorting makes it more obvious that Annotating along with Reading and Writing was judged to be overwhelmingly better with paper. However, Finding and Comparing activities were split between Paper and LHS.

Figure 36 Most preferred devices for each reading activity, by user.

pattern that we noticed while interviewing users was that there appeared to be different “camps” of users that tended to express similar preferences and gave similar reasons for those preferences. To investigate the matter further, we tried to determine whether there were such sub-groups of users sharing similar opinions. To perform the analysis, we aggregated user preferences across all reading activities into per-user heat maps that show how users tended to view a particular device configuration with respect to the others (Figure 37). We aggregated user preferences across all users by assigning the top preference a score of 3, a middle preference a score of 2, and a low preference a score of 1. Since we allowed users to rank devices “equal”, to resolve these ties, we used the following process: if all three conditions were equivalent, all received most preferred status. If two were equally better than a third, than those two received most preferred points and the third received

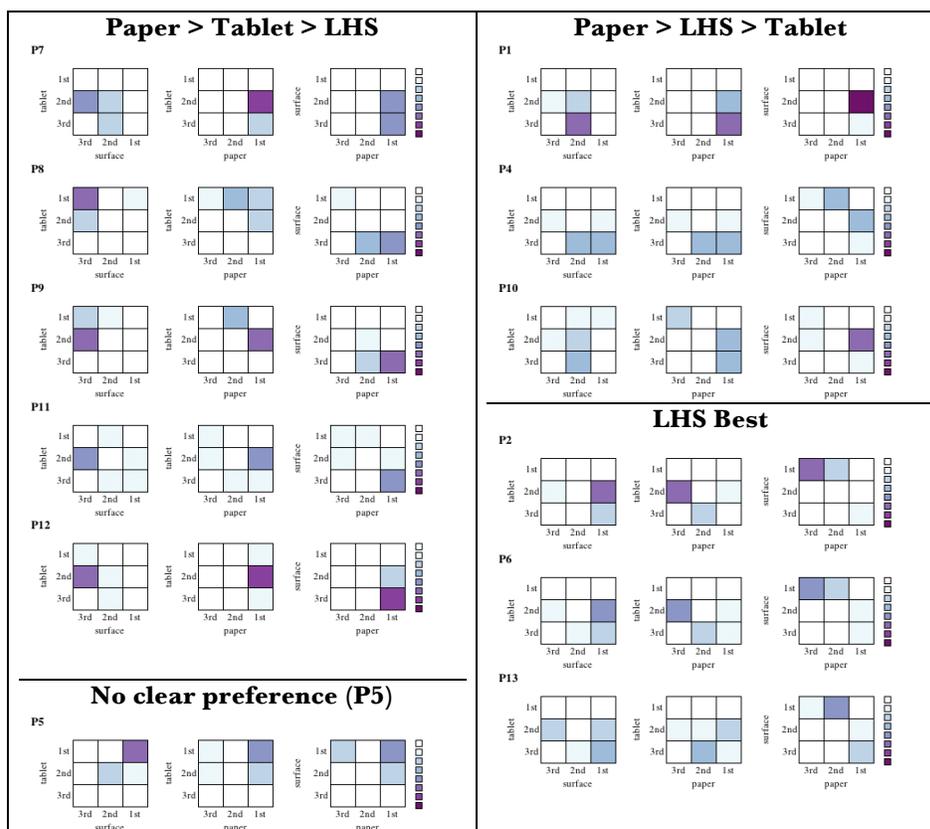


Figure 37. Distinct subpopulations based on user preference. (Note: Y axis labels are tablet, tablet, LHS; X axis labels are LHS, paper, paper).

least preferred status. If two conditions were equally preferred, but both were less preferred than the third condition, then the third condition would get most preferred status while the two would get middle preference.

We identified four main clusters of users. First, there were users who preferred Paper most, but preferred Tablet to LHS (P7, 8, 9, 11, 12). Next, there were users who preferred Paper but preferred LHS to Tablet (P1, 4, 10). Then, there were users who greatly preferred the LHS to the other devices (P2, 6, 13). Finally, there was the individual (P5) who did not have an overall preference because she thought that the different devices excelled at different tasks. Having split the users into these subgroups, the distribution of ratings within each subgroup reveals insights that were lost when all of the users are muddled together.

The main result from the analysis is that the LHS best group tended to prefer the Tablets to Paper in the areas of “Finding” and “Comparing”. We are able to connect this result with the comments we report earlier in the results section. Recall that the main advantage of the electronic conditions versus paper was due to having a neater workspace in which less flipping around was necessary to locate content. From this we infer that the LHS-best group tends to assign high importance to problems of layout.

For those that preferred Paper to Tablet to LHS, the preference largely stems from the fact that Paper and Tablets offered a greater flexibility and simplicity when laying out the surfaces. Users cited their increased satisfaction with the positioning of the laptop, as well as the simplicity and predictability of working with physical objects.

Users that preferred Paper to LHS to Tablets found LHS superior to both Paper and Tablets in terms of speed and the ability to find things easily. Furthermore, the user comments from users in this group focused on the fact that Tablets felt sluggish and that they were too spread out. Both of these comments are almost certainly due to the design of

our tablet prototypes. The sluggishness can be attributed to the electronic paper display in our tablets, and the separation stems from the width of the bezel in our prototypes.

It appears that all of our groups thought that Paper and Tablets were generally superior for writing-intensive tasks like annotation and writing. The LHS appears to be the generally preferred device for comparison tasks, likely owing to its characteristic of keeping information in a compact area.

Could Tablets Be A Good Compromise Between Paper and LHS?

Because only a single user found Tablets to be the best overall, we wondered if this indicated that Tablets were simply a poor choice overall, or whether Tablets represented a compromise choice where user opinions were split.

So, one final analysis we performed was to check how user preferences might change if paper was removed from consideration. The reason for doing this stems from the fact that Tablets were more similar in terms of physicality to Paper. Consequently, there may be a number of users who valued the paper-like benefits of Tablets but rated Paper higher because of the unsurprising fact that Paper was better at providing paper-like qualities.

Returning to the overall preferences presented in Figure 36, of the 5 users who rated Paper best overall (not including P5, who rated all conditions equally good), 4 found Tablets second-best overall. These users added to the ones that originally found Tablets to be preferable (P5 and P8) give a total of 6 users who preferred the Tablets, matching the number of users preferring LHS. These numbers indicate that Tablets might be the better compromise solution, appealing to those that prefer paper in situations where they are unable to use plain paper given its lack of digital functionality. The slight discrepancy between these numbers and those found in the sub-population analysis (Figure 37) can be attributed to the fact that the heat maps were generated by weighting individual reading activities equally, which users probably did not do when arriving at their overall ratings.

7.7 Discussion

7.7.1 General Efficacy of Multiple Reading/Writing Surfaces

First and foremost, we believe that our results confirm the notion that having multiple reading surfaces available is essential for the completion of a number of common active reading tasks. Despite exhibiting a wide range of processes when completing the reading tasks, all users, at some point took advantage of having multiple reading and writing surfaces. This confirms prior research, which highlighted the difficulty of performing active reading tasks with only a single display surface [O'Hara and Sellen 1997].

The way participants employed multiple surfaces simultaneously mirrored how multiple sheets of paper are used to support active reading [O'Hara and Sellen 1997; Adler et al. 1998]. Every participant adjusted the spatial layout of the surfaces multiple times during the experiment. Nearly all participants placed surfaces side by side to perform the document comparison in the flight-scheduling task. Lastly, multiple surfaces were used when interleaving reading and writing, such as when participants referred to the question page while writing on the answer page.

The fact that these behaviors were observed for all conditions suggests that users were able to use multiple reading and writing surfaces in a manner similar (but not identical) to paper. Therefore, we believe that task specifics and user preferences will be all the more important in deciding which of these conditions is most appropriate. To help researchers and designers make that decision, we dedicate the remainder of the discussion to the implications of the differences between different approaches to supporting multiple reading and writing surfaces.

7.7.2 Making Sense of the Differences Between Conditions

Our results provided a large amount of data about what participants considered to be the beneficial features of each device condition. We distill the benefits that participants cited into the diagram in Figure 38. The intersections in the Venn diagram represent features in common between pairs of devices. For example, Paper and Tablet both provide tangible reading surfaces. Ignoring the boldfaces and colors for the moment, each bullet point lists a beneficial aspect that users cited. The bullet points are grouped by the surface traits that give rise to these benefits.

Figure 38 serves an additional purpose beyond providing a summary of the subjective data. The intersections that do not include a specific condition (e.g., “Tangible Surfaces” when looking at the LHS) can also be interpreted as the features users missed about a specific condition. For example, the benefits in the “Tangible Surfaces” intersection show that LHS currently lacks support for laptop layouts. Identifying these shortcomings provides information about how each condition can be improved.

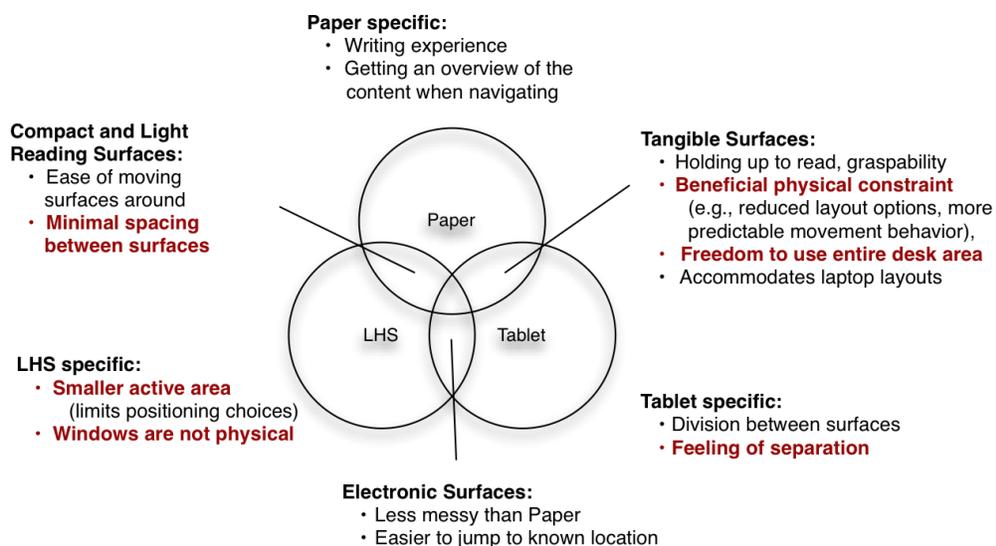


Figure 38 A diagram of the beneficial properties identified in each device condition. Benefits that conflict with each other are bolded and highlighted in red.

7.7.3 Functionality Cuts Both Ways

One interesting observation to notice in Figure 38 is that certain benefits appear to contradict others. These are boldfaced and displayed in red. For instance, one of the advantages cited for Tablets was that surfaces are separate. However, this benefit seems to contradict that the comment that having very little spacing between surfaces in the Paper and LHS conditions was an advantage. What these conflicting opinions suggest is that specific device characteristics should not be interpreted as being definitively good and bad. A combination of personal preferences and task requirements determine whether an affordance is beneficial, detrimental, or not of consequence.

Although it may seem odd at first that characteristics that constrain a user's capabilities (e.g., limited layout space for the LHS, physical surfaces in the Paper and Tablet conditions must respect laws of motion) can be perceived to be benefits, the observation is consistent with the literature. For instance, the fact that some of our users believed that the physical limitations associated with the Paper and Tablet conditions helped to simplify the space of layout possibilities brings to mind Terrenghi et al.'s [2007] sorting task, where users took advantage of the physical constraints in the tangible condition. Similarly, Grudin [2001] noticed that the physical space between monitors, which restricted window placement, helped to reinforce the separate roles of the display area. Our observations that limited layout area and physical restrictions turn out to be helpful are additional examples of these beneficial constraints.

7.7.4 Suggested Improvements for the Tablet

The capabilities that the Tablet condition lacked in comparison to the LHS and Paper conditions stemmed mostly from the physical bulk of the tablet, despite the prototype devices we used being close to the limits of weight and thickness possible with contemporary technology. Clearly, hardware developments like smaller electronics will

result in lighter, thinner devices with less obtrusive bezels. However, an important realization is that, no matter the technical innovation, tablets will always be thicker and heavier than a virtual window. Consequently, interactions that substitute for physical manipulation should generally be desirable for Tablets.

For example, the interactions to virtually “throw” and “retrieve” content from one device to another (similar to what has been done for large tabletops [Reetz et al. 2006]) can help reduce the situations where tablets need to be moved around. A simple to implement method of moving information would be to use lightweight data transfer techniques such as Pick and Drop [Rekimoto 1997] or Stitching [Hinckley et al. 2004]. With the appropriate tablet localization support, these techniques could be further enhanced to use real-time positioning information about the tablets. For instance, rather than explicitly interacting with different devices, a “throwing” operation could be made in the direction of the destination device.

Viewing and working with the contents of tablets in the environment at a distance may also be a good way to reduce physical rearrangement. For instance, a tablet may provide an interactive “portal” onto a different tablet. This approach could be extended to juxtapose the contents of multiple tablets on a single tablet, which draws inspiration from the comments of one of our participants who suggested that tablets provide multiple windows. Keeping the full contents of the page visible on surrounding tablets mitigates some of the problems associated with seeing less of documents when they are inside windows.

7.7.5 Suggested Improvements for the LHS

The capabilities that the LHS condition lacked in comparison to the Tablet and Paper conditions stemmed from the lack of physicality and to some degree layout flexibility. Increasing the active area on which the user can lay out surfaces will also increase layout flexibility. In addition, it seems possible that with a large active region, the user may

become less aware of the fact that the LHS is a separate device. However, it would appear that preserving some functionality to allow the user to define small regions where layout is slightly more constrained would be prudent.

Conveying some of the benefits of physicality in a LHS system can be done in a variety of ways. One way is to use classic UI techniques like snapping (e.g., snap-dragging [Bier and Stone 1986]) to impose helpful constraints. Although snapping does reduce layout flexibility, it also reduces the cognitive load of performing these layout operations.

Alternatively, giving virtual objects the appearance of obeying the laws of physics may also recapture the advantages gained from physicality. Recent work like BumpTop [Agarawala and Balakrishnan 2006], ShapeTouch [Cao et al. 2008], and Surface Physics [Wilson et al. 2008] offer a guide. These techniques can provide more predictable physical behavior for virtual reading surfaces as well as allow users to interact with virtual surfaces in a more natural and efficient fashion. These techniques would allow users move multiple surfaces at once and not restrict interactions to the just the finger.

Finally, modern UIs tend to rely heavily on visual techniques to substitute for physical interactions on other media. Therefore, improving tactile feedback on virtual surfaces using techniques like TeslaTouch [Bau et al. 2010], or through the use of physical widgets [Weiss et al. 2009] are attractive because they aid users' interactions with virtual surfaces without placing additional demands on the visual system. For the same reason, the use of auditory feedback to support eyes-free interactions, using techniques similar to those presented in EarPod [Zhao et al. 2007] and Shoogle [Williamson, Murray-Smith, and Hughes 2007], would be helpful as well.

7.7.6 Other Important Improvements

The two other weaknesses that were common to both our electronic conditions were that our electronic conditions did not provide the same navigational capabilities as paper, nor did they have a comparable writing experience.

For navigation, users mentioned that even though electronic conditions provided random access, it was difficult to know where things were. Navigation interfaces like Space Filling Thumbnails [Cockburn, Gutwin, and Alexander 2006], could have addressed this, but we opted not to use them in this study to minimize differences between conditions. With faster displays, additional solutions are possible. Techniques for conveying more information about the pages include showing dynamic previews of the page and zooming interfaces.

For writing, there do not appear to be any technologies on the horizon that will suitably address the feel of writing on an electronic screen. As such, it would appear that digital paper technologies like Anoto¹³ have an important role in reading activities that require large amounts of writing.

7.7.7 Tablets Occupy the “Middle Ground”

Our results suggest that tablets offer an in-between solution between paper and LHS systems. Our experiment showed that tablets were successful in incorporating a number of beneficial properties of paper. These include: (1) more easily accommodating a cross-device reading environment, (2) allowing users the ability to pick up reading materials for increased reading comfort and to aid comparison and (3) providing a tangible object to manipulate, which assisted in layout planning and also facilitated subtle adjustments that benefited the writing experience. Therefore, it seems that Tablets are the most direct way to recapture paper affordances in an electronic system. So, even though tablets were not preferred outright, we believe that by occupying this middle ground between Paper and

¹³ <http://www.anoto.com>

LHS, they are the most promising platform for which to develop improvements. Improving the Tablet condition would likely produce a solution that appeals to both users who prefer the LHS as well as those who prefer Paper.

7.8 Summary

We ran a controlled study to determine how different methods of implementing multiple reading and writing surfaces would affect the reading experience. Specifically, we focused our attention on three device configurations: Paper, a single Large Horizontal Screens (LHS) and a set of three interlinked tablet devices. For all of these approaches, users were able to spatially arrange the surfaces and interweave writing tasks. Although these three conditions differ in many ways, the key difference between the reading surfaces that we investigated was the impact of different ways of implementing physicality/tangibility. We report on how users performed the tasks as well as how users judged each configuration. LHS and Paper were preferred overall, but for distinct reasons. Paper was preferred mainly because it was more familiar, comfortable to use, and easy to write on, and the LHS system for its organized presentation of materials. Observations revealed that tangibility in both the Paper and Tablet conditions serves a number of useful roles that include tactile feedback, and improved layout options, especially when using a traditional computing device like a laptop. On the other hand, the fact that virtual windows on the LHS had no physical constraints was also appreciated. This revelation illustrates an important general theme of our results, which was that different users were often conflicted about deciding whether specific features of each approach were positive or negative. Even aspects of the devices that seemed to constrain user's options could be seen as beneficial. Having identified some of the distinct benefits for each condition (and those that are missing) we discussed methods by which researchers and designers could seek to compensate for these missing features.

8 Real World Deployment of the United Slates System

Active reading activities often occur within larger workflows into which new tools must successfully integrate. This means that controlled studies in the lab must be supplemented with observations of use with real-world reading tasks to get an accurate picture of whether a system is actually effective in supporting reading activities. This chapter describes an *in-situ* study in which the United Slates system was deployed with 7 doctoral students in the humanities who each used the system for their reading tasks over a 3-4 week period.

8.1 Introduction

The United Slates system was designed to support the highly complex reading activities encountered in classroom and workplace environments. Therefore, the true test of United Slates system and its associated concepts would be a real-world deployment into the very environments for which it was designed.

Several essential data points can only be gleaned from evaluating the use of reading systems in the real world. First, and perhaps most obviously, is whether the system meets real reading needs, rather than that of an experimental task in a controlled study. Real reading tasks often take longer than what can be accomplished during a laboratory study. Also, laboratory studies cannot represent a full cross section of reading activities. Second, real-world use reveals how a novel system fits into a user's existing ecosystem of reading, which consists of a mix of paper and electronic resources, software tools, and individual reading habits. Third, the introduction of a new technology can produce shifts in reading practice, as users adjust their workflow or use the technology in unanticipated ways.

Despite the crucial data in-situ evaluation provides, there is currently a gap in the literature when it comes to evaluating electronic reading systems in the real world. Research systems designed to support academic reading and writing activities have rarely been deployed

outside of the lab, much less for an extended period of time. Although previous studies of electronic reading technology in the classroom have examined the use of commercial devices like PDAs [Marshall and Ruotolo 2002], and e-book readers [Behler 2009; Young 2009; The Trustees of Princeton University 2010; Thayer et al. 2011] “in the wild”, these devices were found to be of limited utility for active reading. As a result, the information they provide about how active reading processes might change as a result of using an electronic device is restricted.

We set out to contribute additional knowledge about the use of electronic reading systems for academic use. To do accomplish this, we conducted an extended deployment of a version of the United Slates system. To fully exercise the system, we selected a demanding set of users for whom reading is a primary activity: Ph.D. students in the humanities.

Our findings from the study include all three types of insights we hoped to glean from an in-situ deployment. We learned a great deal about our users' reading needs and the extent to which the United Slates catered to those needs. Our users described their experiences using the slates for general reading of books and articles, dissertation writing, teaching, among other activities. Although some of these activities involved just the slates, a sizable percentage involved PCs, loose-leaf printouts, books and even commercially available reading devices.

Our reports detail the broad spectrum of tasks that an electronic reading device would be used for in an academic environment. With these tasks as a backdrop, we detail the key aspects of the United Slates that users found beneficial. These include inking capabilities, micro-mobility, task partitioning, and different multi-display configurations. Some features proved helpful in unexpected ways. For instance, automatic synchronization of documents and annotations simplified the process of digitally archiving content in addition to its primary goal of keeping data consistent across devices. Areas where the system fell short,

like certain types of navigation, availability of content, and speed were informative, and are included in our report.

Our study primarily contributes a more in-depth look at the role of electronic reading technology in the realm of academia. However, from a broader sense, the deployment helps shed light on the overall relationship between new electronic technologies and the existing electronic and paper tools they hope to supplement or replace.

8.2 Study Goals and Setup

For the deployment, the primary goal was to observe the use of a multi-slate system once the overhead of working across slates was removed. Although many features could potentially have been implemented, for purposes of simplicity the focus was on testing the core ideas of this dissertation (e.g., multiple screens, interactivity between screens, integration with PC). Previous experiences with the technology probe also suggested that simplicity would be an important quality because it would allow users to get started with the system more quickly. Also, there were limits to what types of functions could be successfully supported on the e-paper screens.

The version of the United Slates we deployed with users was a slightly modified version of the system described in Chapter 4 through Chapter 6. Here, we provide a walkthrough describing the details of features in the version of the system provided to participants during the study.

8.2.1 Hardware Configuration

For maximum reliability we rely solely on the tactile switches in the input modules of each slate to operate the devices we deployed. The upper rectangular buttons are used to move backwards through the document, the lower buttons were used for moving forwards and the center circular button is the command button. One problem we identified during

preliminary testing was that the magnetic sensors for sensing device proximity were somewhat unpredictable. First, the proper alignment needed to trigger a contact event differed between slates. Also, in the course of arranging slates in space, it was possible to accidentally trigger the proximity links by placing devices too close to each other. For these reasons, we disabled the proximity link aspects for the deployment.

8.2.2 Reading Interface

The two main views of an active document are the active page view and the document overview (Figure 39). The biggest difference in the deployed system with the United States system we tested in the lab was that the deployed system used gesture-based marking for command selection rather than the pop-up context menus. We thought that with more extensive training and extended use, users would attain expert-level performance and benefit from the faster gesture-based approach. We specifically use the Scriboli [Hinckley et al. 2005a] command selection technique. Using the technique, the user first selects items

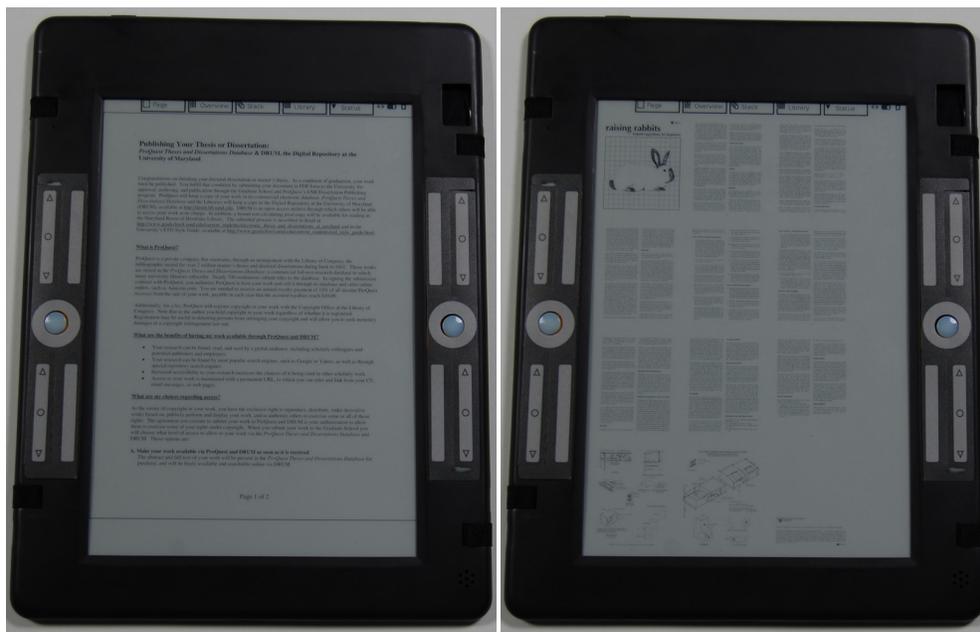


Figure 39 Left: Full page view of a document. **Right:** Space-filling thumbnail view of a document.

to establish the scope, creates a self-intersecting “pigtail” to call a directional marking menu, and marks in one of four directions to select a command (Figure 40).

Like other versions of United Slates, ink marks can be applied directly onto any page. These marks are stored within the PDF documents using native PDF annotation capabilities. We had to disable the ability to make margin callouts since the PDF standard did not support ink annotations within a popup window (Figure 41). The compromise we provided to enable more writing to fit in the margin was to provide a floating writing window that allows the user to write larger. Upon closing the window, the ink is scaled down by 75% and placed onto the page. The writing window is invoked using a marking gesture. We use the beginning of the gesture stroke to allow the user to specify the anchor point of the scaled annotation (Figure 40).

For extra security, we removed ability to erase annotations using the slates. Stroke deletion had to be performed using an external PDF editing tool. We were concerned that inadvertent contact with the eraser end of the stylus could accidentally remove strokes. We therefore felt a more prudent approach was to disable any possibility that the system might

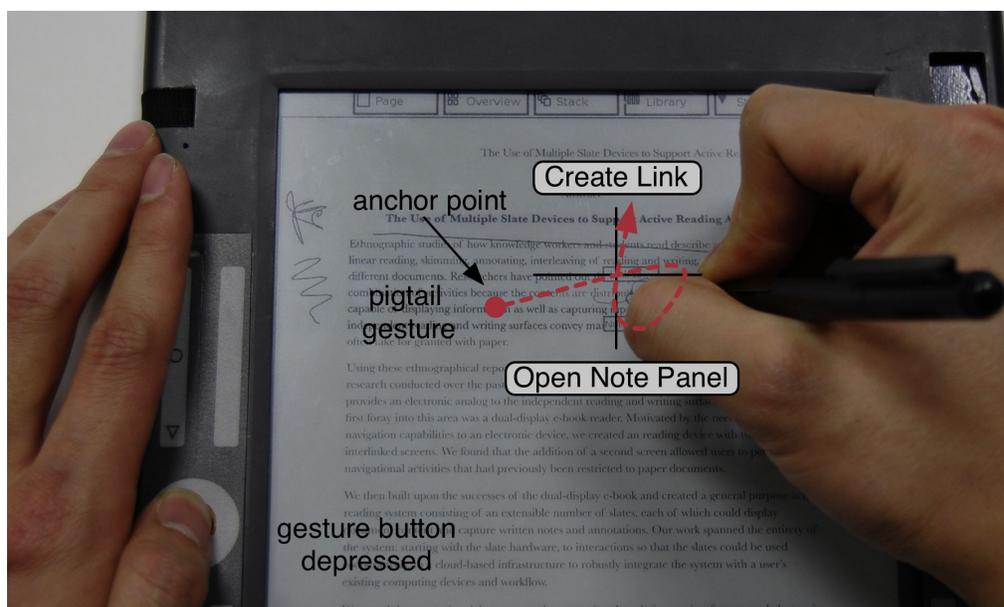


Figure 40 Scriboli-style radial menu for selecting commands in the deployed system.



Figure 41 Left: Writing in a note panel (middle) is scaled by 75% and placed at the panel anchor point (bottom). **Right:** User-created hyperlinks are indicated by a hollow square. Annotations can be placed beside the hyperlink as a reminder about the destination.

delete important data. Stroke permanence made the writing experience of the slates more similar to that of using an ink pen.

For the document overview we use Space-Filling Thumbnails (SFT) to show every page in a document together on a single screen. Larger documents resulted in smaller thumbnails.

In practice, we found that a practical limit for the SFT that enables pages to be resolved from each other was approximately 100 pages (less if documents are primarily block text).

To address this problem, we gave participants the option to send us particularly lengthy documents, which we would help break down into several smaller sections. 3 users (Cathy, Jane, and Anne) provided documents to be split. Larger documents were often usable in the overview if users made identifying ink marks on the pages.

8.2.3 Managing Documents

The deployed system used the active stack and library components characteristic of the United Slates system. We restricted the size of the active stack to 10 documents, which was done to ensure that all documents could be comfortably listed in one screen of data. Users

links between arbitrary documents. We use Stitching as the mechanism for performing link creation. To create a link between documents the user draws the directional gesture corresponding to the link command and then taps on the slate displaying the desired target of the link. The link is inserted at the starting point of the gesture and is denoted with a square (Figure 41). Based on prior feedback that it was equally easy to control slates directly rather than proxy interactions through slates, we removed the slate remote control feature from the system we deployed.

8.2.5 PC Connectivity

To support interactions between the slates and PC, we created PC and Macintosh versions of an application called “Slates” (Figure 43). The Slates application had three primary functions. First, Slates allowed the user to see the active document on each connected slate and also enabled the user to turn pages from the PC. The Slates application also provided a fast way to call up any of the active documents on the PC, which is done by double clicking one of the slate visualizations. The Slates application launches the copy of the document in the user’s Dropbox folder, which matches the content of what the user is viewing on the slate device. Finally, Slates is used to implement the distributed clipboard [Miller and Myers 1999] functionality by connecting the PC’s clipboard with the United Slates



Figure 43 Macintosh version of the Slates application. The application shows the active document on the slates in the environment. Users turn pages using the arrow keys on the keyboard and launch a document on the PC by clicking.

clipboard in the cloud.

8.2.6 User Feedback and Journal

To help users track their usage and to provide a method for users to submit immediate feedback and bug reports, we included an integrated feedback and usage tracking utility on the deployed system (Figure 44, top). Using the feedback tool, users could classify the nature of the feedback and write themselves a short reminder note that would be attached to the submission. The user can later supply additional information using the Web interface described below.

8.2.7 Web Interface

In order to keep the slate interface simple, functions that we decided were not central to reading activities were offloaded to a study website. Three key functions were performed

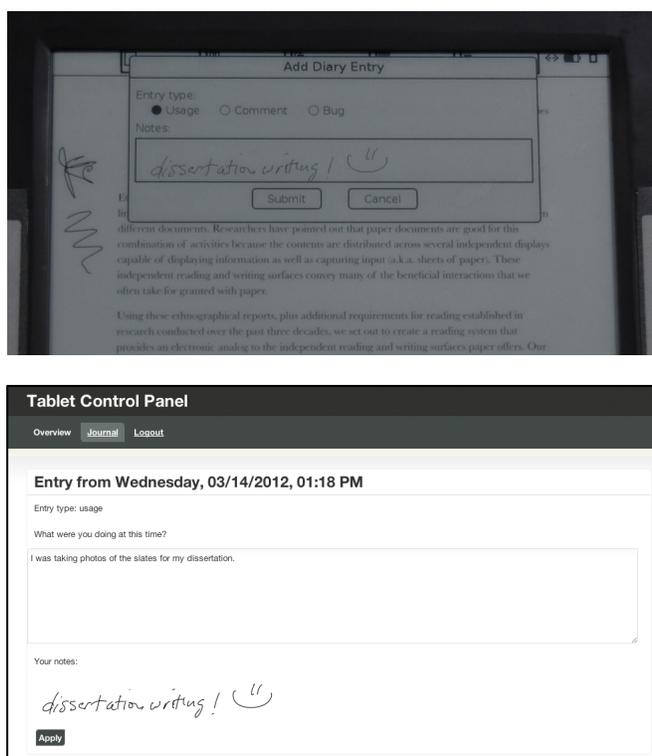


Figure 44 Top: Diary entry dialog on the slate. **Bottom:** Corresponding entry in the web diary interface showing that a user's notes are transferred to the web interface.

using the Web interface. First, users use the interface to create, edit and remove attribute tags for the Library. Doing the editing over the web was helpful because the interface was more responsive and also the user had access to a keyboard to type. Second, the web interface allows users to add extended details to their feedback and journal entries (Figure 44, bottom). The use of a web interface to supplement quick entries submitted from a device is inspired by Sohn et al.'s [2008] mobile phone diary study, which employed a similar technique. Third, the web interface allows users to set flags that tell the system to pre-render problematic PDF documents that are slow to display on the slates. The majority of these slow documents are PDFs created from high resolution scans of books and papers.

8.3 Methodology

The study we conducted involved multiple phases that culminated in the extended deployment of the United Slates system with users. To gain a better understanding of how the multi-slate system fits in with, and changes existing practices we first conducted a photo-diary study to get a better sense of how users currently employed paper documents and existing electronic tools. Once devices were available to loan to a user, we conducted a weeklong training session to make sure that users were familiar with all aspects of operating the system. Finally, participants used the system for their reading activities for 30-40 hours.

8.3.1 Photo-Diary

We started by asking users to keep a photo diary to document the range of reading activities they performed on a day-to-day basis over a one-week (7-day) period. The goal of the diary was two-fold. First, we wanted to ensure that users provided an accurate and representative picture of the reading activities, environment and tools that formed their current practice. Second, the photo diary was useful for eliciting additional information about past experiences working with documents and future goals.

During the week when the photo diary portion of the study took place, we asked users to take a photo whenever they found themselves reading for work, no matter how insignificant the activity. We met with all users at least three times over the course of the week. We met users once to get them started, once halfway through the week, and once at the end of the week. The schedule was not strictly controlled, however. If users mentioned they expected to perform an activity not previously discussed after the week was over, we made arrangements to meet with the user once they had documented that activity.

In the meetings with our users, we reviewed photos together with our users. We asked users about: when and where the photo was taken; the reading task the user was trying to perform at the time; the documents and tools (e.g., computer, pen, notebook, etc.) in use; and whether there was any special organization significance to the materials. These meetings lasted approximately one hour.

Our original plan was to immediately start users on the slate reading system following the photo-diary meetings. However, we encountered technical problems when preparing the devices for deployment. As a result, we needed to delay the deployment. Once we had solved the problems and tested the slates internally in our lab, we started the training and deployment process.

8.3.2 Training Sessions

When we were ready to give users slates, we first spent one week training the user on the operation of the devices. The training allowed us to give the users a full presentation of the device's features, to identify potential problems, and to make sure the user was familiar with the system before the study proper. The training was broken up into four 30-45 minute meeting sessions. We aimed to meet with the users every other or every two days to give users time on their own to try out the concepts we covered but not so much time that the users would forget the ideas.

The first training session introduced the basic operation of a single slate and covered how to move documents onto the slates, how to open documents, how to navigate, and how to operate the web interface. We asked users to ensure that the slates could connect to their home and office wireless networks. The second training session covered the command system, the writing box, and the Slates application.

During the third meeting, we provided users with a second device and introduced the Conduit and Stitching inter-slate interactions presented in Section 5.3.3 (we described the technique as “Pick-and-Drop” to users because of users’ familiarity with the term “Drag-and-Drop”) for documents and links. For these more advanced features, we covered potential use cases where we thought the feature might be beneficial. For instance, we suggested that links could be used to support cross-referencing, and that Stitching may be a good way setup the environment to look at different parts of a document. We also gave users the mobile hotspot that would allow them to connect to the Internet in locations without Wi-Fi access points.

At the fourth session we gave users the remaining two slates for their deployment. With four slates, available, we revisited the thumbnail navigation and showed how multiple panels of thumbnails could be used to jump between pages of different documents.

Afterwards, we performed a recap of all of the features we covered and conducted a short interview asking about the expected applicability of each feature to their reading activities in the upcoming weeks. Although the interview was primarily for us to gather information about projected system use, a secondary purpose of the final meeting was to motivate users to employ the device in their reading activities. We believed that talking with the users about how they would use the system would at the very least encourage the user to attempt to use the system for a particular reading task. If the user later discovered that the system was ill suited for the task, we would receive feedback about that. The worst-case scenario,

which we tried to prevent, was non-use of the system. Finally, we explained to users how to use the in-device diary feature to conveniently track their usage, but told users they could track their usage in whatever format they most preferred.

8.3.3 Extended Deployment

Following the training sessions, we started the study proper. The study protocol specified the deployment would last as long as it took users to perform 30 to 40 hours of work using the United Slates system. We kept the definition of “work” very loose to encourage slate use for a variety of activities. Since many reading activities span multiple devices and mediums, we explained to participants that they should count the full duration of work that involved the slates into the target time, not just the time spent operating the slates. We also explained to users that the reason we provided 4 slates was not necessarily because we expected them to use all of them at the same time. Rather, the purpose was to give them the option of using more slates if the need arises.

For the deployment, we asked users to take photos of device use as they did for the photo diary study. Recognizing that reading tasks would often repeat, and to avoid overburdening users with responsibilities, we told users that they did not have to capture every instance of reading, but instead should photograph unique activities. We did not set up formal meetings during the deployment but did end up meeting with users for technical support issues.

One issue encountered with the first set of users was that these users were not particularly fastidious about tracking their usage and taking photos, owing to the extra effort to perform these activities. For subsequent users, we periodically emailed users questions about their usage habits and reading activities, and reminded users to track their usage.

The time required for users to get to the prescribed 30-40 hours of slate use took significantly longer than expected. Initial estimates were that it would take users about one

week to reach the prescribed amount of use. These calculations were based on a 6-8 hour workday over 5 days. In practice, participants did not read continuously through out the day and also had other commitments, like travel, that occupied their time. So, even though most users used the United Slates system regularly, for a large portion of their reading tasks, it took participants about 3 weeks on average to reach the amount of usage we specified.

At the end of the extended usage, we setup a final wrap-up meeting and debriefing session. To prepare for the meeting, we aggregated users' predicted use of the system (from the end of the training), diary entries, email status updates and photographs. This information was used to prepare personalized questions about the participants' activities with the United Slates. In addition to these personalized questions, we had general questions about the slate functionality, and also overall impressions about the system. The final wrap-up interview covered a large amount of material and lasted between 2-3 hours.

8.3.4 Analysis

The analysis of the data focused primarily on the interviews conducted at the end of the deployment. However, the preparation for the final interview required a review of the participants' photo diary data as well as post-training feedback. These previously collected pieces of information were important for developing questions asking users to compare reading from the slates against previous practice. Questions about differences between projected use from the post-training meeting and actual use during the deployment were also prepared at this stage.

The analysis of the final interview first involved going through interview notes to enumerate the ways each participant used the United Slates. With a comprehensive record of how the slates were used, behaviors stemming from the unique characteristics of the United Slates were highlighted. Examples of these include multi-slate use, new ways of

working as a result of digital functionality, and differences users noted between reading with the United Slates system and with paper documents. Having identified the key themes from the interviews, relevant portions of the recorded interview audio was transcribed in order to check the accuracy of notes, and to obtain exact quotations.

8.4 User Profiles and Activities

7 users participated in our study. We specifically recruited for graduate students who were engaged in research, so the majority of our users (6 out of 7) had advanced to candidacy and all were engaged in activities related to their dissertation research. The participants were paid \$750 for their participation, which was computed based on a rate of \$15 per hour for the approximately 50 hours of total time commitment that was expected over the full duration of the study.

The highly specialized work of our graduate student participants meant that our users were engaged in the different types of research, with reading materials varying greatly between users. In addition to research, some of our users had teaching and grading responsibilities. The participants in our study provided many sample points in the different reading tasks graduate students in the humanities perform in their work. To provide some context about our participants, we will give some background about each user's stage in their graduate work and typical reading activities they were engaged in. Participant names have been changed.

Hannah - Hannah was a student in the History department and was approximately 6 months away from finishing her studies. Hannah primarily used the slates in support of completing her dissertation. She would mainly use one or two slates alongside her computer while she was writing. Hannah's work consisted largely of revisiting documents that she had collected on past research trips, which were archived electronically on her computer. With these documents, Hannah marked up important sections and then either

made notes about the documents, or referred to them for specific details as she was writing. Prior to receiving the slates, Hannah needed to print these documents out in order to read and mark them up. The marked up printouts were then scanned back into her computer so that they could be more easily accessed. Another strategy Hannah employed before receiving the using two laptops side by side to simplify the process of reading and entering notes into Zotero¹⁴, a reference and notes management software package. Hannah would use one laptop for viewing documents and the second for taking notes, preferring to use two devices rather than flipping between windows (Figure 45).

With the slates, Hannah continued working on her dissertation. However, instead of using printouts or her second laptop, Hannah read documents from the slates. Hannah also used the slates to take notes and annotate papers for History department colloquia, during which students and professors read and then discuss each other's papers. Finally, Hannah used the slates for grading student papers and for job applications.

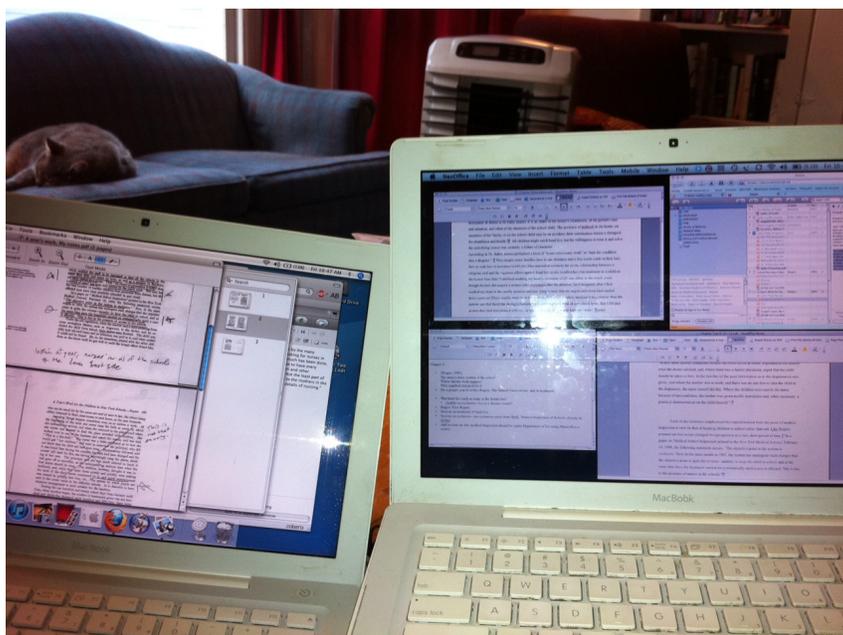


Figure 45 Hannah's two laptop reading environment. One laptop was for reading and the other for taking notes. Hannah preferred this arrangement to switching between windows.

¹⁴ <http://www.zotero.org>

Justin - Justin was a student in the English department who was preparing to defend his dissertation. Justin's other major responsibility during the semester was teaching two Freshman Writing Seminar courses. For his dissertation, Justin's activities skewed towards editing and polishing owing to the fact he was so close to defending the dissertation.

With the slates, Justin would output PDF versions of his dissertation and use the slates to review, proofread, and edit different sections of his dissertation in lieu of printing them out. For Justin's teaching responsibilities, Justin placed reading materials that he assigned to students on the slates and brought the slates to the classroom for his lecture and discussion sessions. Justin additionally graded student papers from the class on the reading system. Both of these tasks were done with paper documents prior to the deployment. Beyond these two core responsibilities, Justin regularly followed a handful of academic journals. As these journals came out with new issues, Justin would log on and download articles of interest onto the slates. Normally, his standard practice would be to read the journal articles on the PC, and occasionally read the ones that arrived in the mail in print form. Over the course of the deployment Justin made several trips to academic conferences and brought slates to the conferences. Justin also documented his use of the reading system for some of the aforementioned tasks while he was in transit between the University and the conference venue.

Barry - Barry was a student in English who had passed his examination to advance to candidacy several weeks prior to receiving the reading system. Barry's reading immediately prior to receiving the slates consisted of readings in preparation for his A-exam and readings relevant to his dissertation. Barry described a general reading strategy for most of the materials he encounters. The first pass involved determining the significance of the text and producing a summary. Subsequent passes involves a closer reading of the material in which he circles specific words and phrases that the author uses to modulate the flow of the

text. For Barry, the marks made during the close reading forms a useful artifact that helps to visualize the writing style of the author.

Barry's reading activities after receiving the slates focused on readings related to his dissertation. The materials included: a book central to his dissertation that he had been primarily reading on the iPad, segments of a historical manuscript, and a very large book that was only available in Amazon's proprietary .mobi e-book format. Since the reading was for his dissertation, Barry made extensive annotations on these materials, the historical manuscript in particular, on account of it being a primary source. For the final book, Barry had no choice but to read using his Amazon Kindle device, but employed one of the slates alongside the Kindle to make notes. In addition to the reading he was doing for his dissertation, Barry downloaded and read journal on the slates. Normally, Barry would travel to the library, obtain a stack of recent journals and peruse the bound paper versions of the journal. Lastly, Barry sat in on several lectures while he had the slate and took notes on the slates during those lectures.

Cathy - Cathy was a student in History and the participant who was least far along in her graduate studies. Cathy was our only participant that had not yet advanced to candidacy. Prior to and during the time she had slates, Cathy was simultaneously preparing a research prospectus, and tackling the list of readings that she needed to be familiar with for her A-exam. For the research component, Cathy used the slates for one long book and several historical manuscripts. Normally, she would have read these materials on the computer, and taken notes on a paper notebook on the side. When reading with the slates, Cathy mixed note taking in the paper notebook with direct annotations on the slate. Cathy only placed one item from the reading list on the slate, owing to the difficulty of finding electronic editions of the items on the list. For this particular title, Cathy also had a copy of the paper book and alternated between using the paper book and slates to read. Slates were

also used beside the PC for writing a paper in a course Cathy was taking. For writing, the slates were used for an initial read-through, as well as for reference as Cathy was writing. Slates were also used to read and participate in a History colloquium and to read and comment on a student paper from the class where Cathy was a teaching assistant. These activities would have otherwise been performed with paper printouts.

Jane - Jane was a student in Architectural History who had completed her examinations to advance to candidacy approximately two months prior to the deployment. During the deployment, however, Jane's reading activities on the slates focused on readings for the course she was teaching that semester. The vast majority of these readings were chapters of books and journal articles. Jane read and made notes on the documents she planned to use in class. Normally, Jane would have printed out these documents and written on the printouts. Although for teaching Jane preferred paper documents, for other purposes, Jane mentioned she had transitioned to using an iPad to read most electronic documents. However, the fact the iPad lacked a stylus meant she often had to take notes in a separate

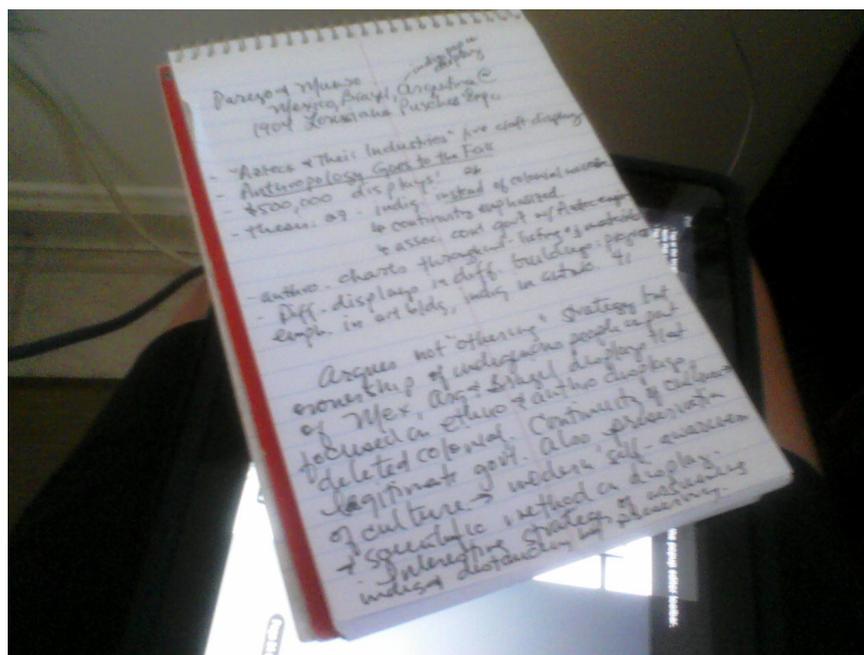


Figure 46 Jane's notebook for taking written notes when reading with the iPad.

notebook (Figure 46). In class, Jane used the slates to teach and to lead discussions. Since the time Jane had her slates was close to the end of the semester, she also used the slates to assemble a comprehensive outline reviewing the topics and documents covered in the class. For the review, Jane used all four slates we provided together, taking turns using each slate to view materials that had been covered over the course of the semester. Jane's use of four slates was notable and unique among our users. Outside of teaching, Jane used the slates for multiple History colloquiums, each time employing a different number of slates. Finally, Jane incorporated slates into the process of making final changes to a journal article that she was about to submit.

Anne - Anne was a student in History who had advanced to candidacy about 6 months prior to receiving the slates. Anne was at a stage where she was gathering documents for her dissertation, reading the documents, and placing notes about these documents into the software tools she used for notes management and writing (DevonThink Database and Scrivener for Mac OS X, respectively) (Figure 47). Anne emphasized that her goal was to extract out all of the relevant information from the documents she was reading so that she would not need to return to the original document. The materials Anne used ranged from digitized books, to shorter article-length documents. Normally, Anne's work process would

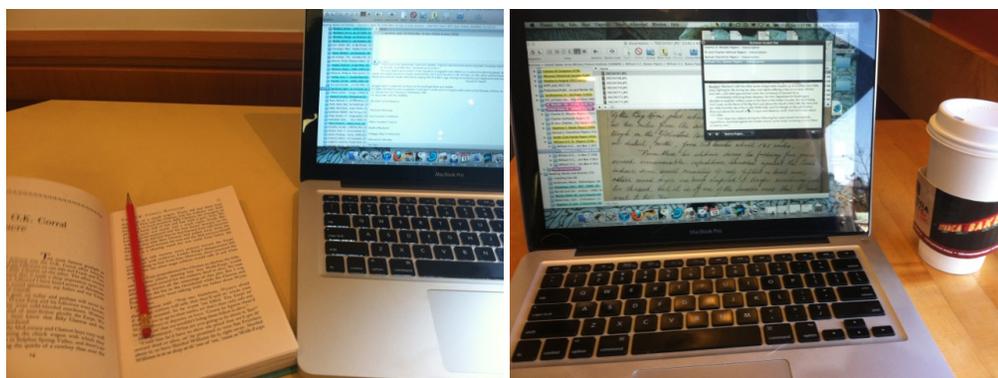


Figure 47 Left: Anne would often read and annotate on a paper book. **Right:** Notes from paper books and digital documents would be transferred into the Scrivener software package, which provided tools for writing, organizing ideas, and research management.

be to read on paper, mark up the documents, and then transfer the annotations and relevant excerpts from the document into the software tools on her computer.

When provided the slates, Anne did roughly the same thing, but did not rely as much on manual transcription. During the deployment, Anne was also preparing the syllabus for a course she was going to teach the following semester. For that task, she went through many iterations of the document, each time making small changes. Anne used the slates and paper almost interchangeably, as versions were paper printouts while others were documents on the slate.

Gina - Gina was a student in the English department who had completed her dissertation but not yet defended. The majority of reading Gina performed was from the various journals she followed. Normally, Gina would read these journal articles on her Kindle. Therefore, Gina mostly used the slates for reading articles from the roughly 10 journals she follows, along with additional articles her colleagues referred her to. Gina's reading activities also included texts from the 19th century that Gina was planning to use for future writing projects and for group discussion. Most of these books came in electronic form. To read on the slates Gina would print these out on paper documents to read and to mark up. Although Gina was not officially enrolled in any courses during the semester, she regularly sat in on courses in the department and read the materials associated with the courses and brought the slates to class discussions. The readings for these courses tended to be book chapters that she would have printed out on paper. One other activity that Gina did with the slates was reviewing and commenting on a section of a professor's unpublished manuscript that was circulated prior to a brown-bag discussion.

8.5 Findings and Key Themes

The duration of our deployment coupled with the individual differences between our participants meant that we encountered a wide range of reading activities over the entirety of the study. Our users employed a variety of materials and used the reading system in a wide range of venues. Despite these differences between users, there were several recurring themes in terms of how users used and responded to the system we provided.

8.5.1 Writing

We start with examining the role that writing played for our users. Users as a whole were dissatisfied that their computer or e-book readers lacked writing capabilities. In fact, the opportunity to try out a reading system that allowed paper-like annotation was a big reason our participants decided to take part in the study. Hannah explained that the ability to write on digital documents was a feature she had “been waiting for my entire life.” All of the other participants similarly mentioned that they had been intrigued by an electronic reading platform that offered the capability to mark up documents.

Writing served a number of important roles in our participants’ reading activities. The ability to make notes that were superimposed over the text was important. For her teaching duties, Jane mentioned that “I have an iPad but it doesn’t make sense to use it for this task. With an iPad I would use a notebook on the side to take notes. With the slates the notes were not split between two different locations.” However, another important role writing served was to help our users stay engaged with the material. Jane remarked, “Most of the writing is to help me stay engaged. The way I read actively is through taking notes. Even when I read the *New Yorker* I use a pen.” Jane’s comments are representative of the other users. For example, Justin recounts how he “enjoyed having the increased engagement of reading on paper back...I found myself highlighting, drawing a star, even if it wasn’t a journal article I expected to keep; these helped me stay engaged.” The benefits from

writing confirm Marshall's [1997] recommendation that electronic reading devices provide support for both notes and non-interpretive markings.

In previous studies where we employed identical slate hardware, users were often negative about the writing experience on the slates compared to paper. Therefore, we were a bit surprised to find that most users were satisfied with the writing experience. Although the stylus did not perform at the level of, say, a favorite pen—Justin mentioned some issues about the calibration and Gina said she had to write at bigger sizes (Figure 48)—most users actually said that after getting used to the stylus, they did not feel the writing experience held them back compared to writing on paper. This suggests that the writing experience on the slates was good enough for most uses. Barry, who regularly used an iPad with a capacitive stylus found the writing experience on the slates to be far better. Given these results, we believe the implication here is that writing with a stylus is better than not having one at all, but that users do appreciate attempts to improve the quality of the

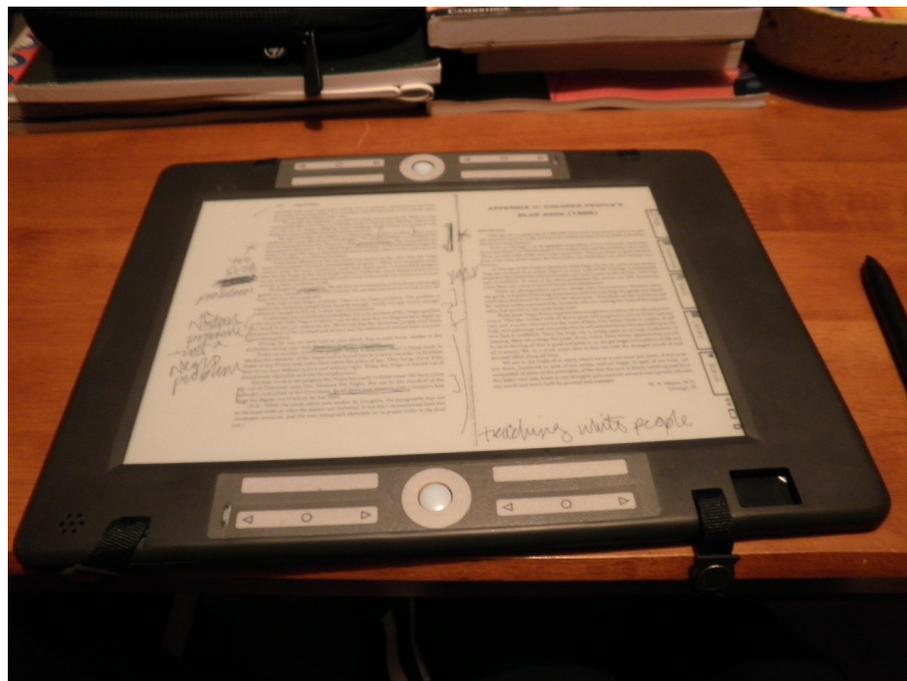


Figure 48 Annotations on one of Gina's document. Gina reported having to write larger than she was used when using the slates.

interaction. The freedom to place marks anywhere and ease and fluidity of writing afforded with a stylus to be the qualities most important for supporting reading.

Although no users mentioned being hindered by the lack of ability to erase strokes on the deployed system, it was consistently mentioned as a desirable feature. Interestingly the reason users wanted to remove strokes had more to do with wanting to send “clean” versions of documents to colleagues and students than correcting mistakes. Hannah and Gina additionally mentioned that the ability to distinguish between different reading passes on paper documents (using things like ink color) was a helpful feature on paper that was absent on the slates. The lessons here are that users want strokes to persist on their documents as they serve as useful reminders of what they had read previously, perhaps through the creation of layers of annotations. Thus, the ability to selectively hide marks, especially when sharing a document, appears to be the best way to support this behavior.

One common problem users encountered on the slates was the inability to access additional writing space. The problem seemed to show up predominantly during grading activities. Justin said that “I thought it was going to be great...with paper, I really liked being able to underline and circle...but I found that it actually gave me all of the problems of working with paper, but removed some of the advantages—with paper I could turn the page over and write on the back.” Other users like Hannah, and Cathy who used the slates to grade student papers, mentioned similar difficulties. Our users ended up having to type comments in a separate email rather than returning the marked up PDF back to their students, which they would have preferred. Suggested improvements our users gave included adding the ability to type out particularly long comments, add blank pages, and attach virtual Post-its.

Summary

Since writing played such a major role for all of our users, users cited the ability to write on top of a document to be a major improvement over alternative reading technologies like PCs, iPads, and Kindles. Based on this, it seems that support for stylus-based writing seems particularly essential. The way that people annotated (e.g., drawing, underlining, bracketing) would have been difficult to support otherwise. Although there is room for improvement, the general sentiments of our users overwhelmingly support Jane's characterization of the writing capabilities as being "the best thing about the slate" and "a game changer."

8.5.2 Ubiquitous Access to Data

The availability of documents and annotations when reading with the slates was most frequently cited as the feature that made slates a superior reading platform to paper. First, the ability to carry around considerably more material effectively increased the portability of the system relative to paper. Second, the seamless capture and synchronization of written annotations, and the new functionality it enabled, actually helped users realize that there were significant annoyances and deficiencies associated with paper documents.

Availability of Documents

Most of the documents our users used existed in electronic form. Users were acutely aware of the difficulties working with paper and many told stories of the difficulties they had when working with large quantities of paper. Gina recalled filling up an entire large filing cabinet with marked up printouts after just one year. Hannah, who was temporarily living abroad earlier in her graduate studies, mentioned, "If I had these during my A Exams, it would have changed my life; carrying around papers without a place to live is miserable."

The constant availability of documents meant that users could do work in a wider variety of locations. Some users ended up having significant travel obligations during their time using

the slates. Additionally, changing reading venues was a strategy to maintain concentration over an extended period of time for many users. Justin's experiences provide some good examples of how accessibility to documents improved when and where users could work. In Justin's case, he was freed from having to carry around the 35 student papers he needed to grade. With the slates, whenever he had free time, he had the option of grading some of them. Also, the fact that the documents were wirelessly synchronized with all devices meant that Justin was able to keep one device at his carrel in the library. As a result, Justin was excited to report that if he was going to read in the library, he could sometimes get by without carrying any devices at all. Given Tashman and Edwards' [2011a] findings about needing to support reading in different venues, we believe our finding is interesting because it highlights that factors beyond the device hardware have an influence on macro mobility. The availability of documents seemed to promote increased peace of mind as well. The slates freed users from having to plan ahead of time about what materials to bring around. Anne, who did quite a bit of traveling during the deployment, describes the problem:

“When I travel I bring one book and I worry that I'm going to run out and so sometimes I pack two books...[with the slates] there were plenty of options. Sometimes I don't know what kind of work I want to do and you don't want to run out of work, so I tend to bring way more with me than I need to do.”

Justin and Gina, both of whom also travelled heavily during the deployment, also expressed similar sentiments about the feeling of security they gained from knowing they had the reading materials ready if they needed to get work done. We believe that removing the pressures of having to assemble resources ahead of time also increases and encourages portability in a non-obvious way.

Availability of Annotations

In the United Slates, annotations were synchronized between slates to ensure a consistent multi-slate experience. It turned out that synchronized annotations had implications that went far beyond ensuring that a document on different slates looked the same. The most-

liked aspect about synchronized annotations was the fact that annotations made on any slate became instantly available on the PC. This was valuable for several reasons. For one, users needed to expend no additional effort to have a digital version of the marked up printout available on their computer. In contrast, when working with paper, users would have to go through the laborious process of scanning a document back. The ease with which annotations were now captured meant that users could preserve more of their work. Justin was pleased that the notes he made would automatically be available in the future, explaining that “in the past I never kept good teaching notes, now I have these marked up PDFs to use for the next time I teach.” Jane was also very keen about how the slates preserved her work, saying, “An immediate archive of the work you’ve done is something I’ll miss about the slates. The notes are available anytime I open it up again. It’s amazing how a small mark can remind you of an entire argument.”

The fact that annotations were captured and stored on the computer (and the cloud) was perceived to be more reliable than paper. Justin really liked the slates for teaching for that reason. He explained, “I didn’t have to worry about losing a sheet of paper, or fumble through papers.” Cathy echoed this view when she said,

“I still prefer to write on paper, but I also lose paper. I’m less likely to lose stuff on here – that surprised me. It’s more secure knowing my writing is on the computer. Or in a weird way it’s all in one place instead of 5 different notebooks and random pieces of papers and folders.”

Cathy’s statement also points to another major benefit associated with annotating on the slates, which is that there is always just one “master” copy of a document. Hannah, when describing the time she would work exclusively with paper, said, “sometimes I had two copies printed out but couldn’t store one or find one—it was a pain.” Having multiple printouts meant that Hannah would eventually need to find and re-integrate notes across several different versions of a document. Furthermore, the knowledge that their work would not be lost if a slate ran out of power, or malfunctioned, was incredibly important for

giving users the confidence to shift more of their activities to United Slates system. We believe that inspiring this kind of confidence, while not necessarily a reading requirement, is key to the practical adoption of a reading system.

Once annotations were on the computer they could easily be called up. Once, Hannah went to a cafe to work only to discover that the slate she brought refused to start up.

Hannah was pleased that she could still work off the documents on of her computer, which contained the annotations. As she recalled, “of all of the features, the one that surprised me the most was when I didn’t have slates and being able to pull up my notes on a computer.”

Jane would use the marked up documents on her PC, finding that the annotations helped her navigate while scrolling. Barry also liked that his notes were available when he had his PC. In fact, Barry was so impressed with this aspect of the United Slates system he made plans to scan in his paper documents following the deployment so that he could reproduce the workflow he enjoyed with the slates.

Summary

The far-reaching effects of having annotations propagate across devices were unexpected.

An important insight from our results is that data synchronization conferred benefits beyond ease of file exchange. The fully automatic capture and archiving of operations performed on paper offered clear advantages over working with paper documents.

On the other hand, the added peace of mind, portability, and ease of retrieval that synchronization provided were interesting because they helped recapture qualities that are traditionally associated with paper [Sellen and Harper 2003]. In fact, this finding updates the results of previous research. When the XLibris system was shown to legal scholars [Marshall et al. 2001], a curious result was that the scholars felt an inexplicable need to print out documents on paper. In contrast, our users welcomed the fact that the slates dramatically reduced their need to print documents out, with no mention of a desire to also

produce a hard copy. Why is this? We believe that the difference in responses can be traced to the theme of increased accessibility of documents. We suspect that with an isolated reading device, there are well-founded concerns that documents placed on the device will be stranded. Therefore, having a paper copy that provided universal access independent of device technology was attractive. With the slates, it appears that the connectivity between devices has reached to point of providing an effective substitute for that function of paper.

8.5.3 Extra Displays

Providing extra display real estate through the use of additional screens underpins much of the design of the reading system. Although the majority of the reading activities our users engaged in involved a single slate, the availability and option of having a second screen proved to essential in select situations for many users. Perhaps the most striking observation was the dramatic change in user opinions about the additional slate before and after the deployment. When starting out, all users were somewhat wary of the multi-slate configuration. Hannah for instance, said “I didn’t think I’d use two but it turned out I did...It was strange, it seemed that I would not need more than one because I’ve already got a laptop and iPhone.” But in the end, nearly all users changed their mind about the utility of having additional displays.

The most prevalent use case for multiple screens involved working with two parts of the same document simultaneously. One scenario where the ability was valuable was to follow up on endnotes. Justin, Hannah and Jane specifically mentioned instances where books and journals made heavy use of endnotes. Moreover, many of these endnotes were discursive, meaning they were meant to elaborate further on something mentioned in the text. For these notes, it was essential to be able to view the endnote and return back to the main text. Existing tools like the PDF reader on the PC made this difficult. Justin described that when he would read journal articles on the PC, he found himself disregarding endnotes because

it was too much work to scroll back and forth. Even users like Barry and Cathy, whose reading activities during the deployment did not require the use of multiple slates, mentioned that referring to endnotes and bibliographies were areas where additional displays would be useful.

Having additional screens of content was also used for working with several documents at once. Hannah, who frequently used one or two slates beside her PC when writing, said that “I didn’t have to go back and forth; having two meant that all the information was there...I would have normally read on the computer and it’s a real pain to read across two windows. I can never layout things the way I want.” At the upper end of the spectrum of multi-screen use was Jane, who made use of 4 slates when preparing a final review document for her class. Working with more than one slate helped Jane better see the connections between the materials she was summarizing. Jane also opened documents in advance on the slates, which provided visual reminders about materials she still needed to include into her review. Jane’s use of her slates resonates with Malone’s [1983] finding that paper documents in the workspace act as markers of tasks that need to be completed. It would appear that Jane’s synthesis of course materials, which required seeing many things at once (Jane referred to it as “free-association”), uniquely benefited from the multi-display configuration.

The ability to work with multiple documents simultaneously had the added effect of supporting the interleaving of reading and writing [O’Hara and Sellen 1997]. Barry and Jane cited specific cases where they used the notebook alongside a document of interest. Barry would take notes on a slate while reading on his Kindle. Jane used a slate as a notepad in two of the three colloquiums she attended. Working with the notebook on a separate device enabled her to take notes during the colloquium without having to look at the notebook, thus freeing her to follow along on the main document, Jane explained.

The other ways that users said they could have made use of additional displays (but did not do so during the deployment) are informative as well. Justin described a reading activity from earlier in the year that involved comparing multiple documents. To perform the activity at the time, Justin connected his laptop to the HDTV in his living room to see the documents together. Justin speculated that having multiple slates might have been useful in that situation. Anne discussed at length about how it had simply not occurred to her to use more slates during her syllabus revision task. In retrospect, Anne thought it would have been extremely helpful for getting a “big picture” of what was happening, and to compare specific weeks of the syllabus—activities for which she used paper printouts.

One area that could potentially limit the applicability of multiple screens (beyond the fact that many reading activities do not require the use of multiple slates) has to do with the difficulties associated with transporting multiple slates. Our participants’ experiences provide some indication about the quantity of hardware users could realistically transport. Users generally thought that a single slate was superior than carrying paper documents and Hannah, Justin and Jane had little trouble with transporting two slates. Hannah mentioned that two slates were comparable to carrying around a book. No users carried three or four slates. Hannah remarked that having to remember to pack and carry around the *accessories* for the slates, like their power adapters and the wireless hotspot, rather than the slates themselves, turned out to be the bigger nuisance of the multi-slate configuration. Hannah’s comment highlights the unexpected result that the real barriers in portability may not be from the quantity of slates at all.

Another interesting finding was that the need for multiple devices was sometimes seen as superfluous, and that contending with this perception may actually be the bigger obstacle that a multi-slate system faces. Barry and Jane mentioned that they needed to have used the system to realize the utility of having additional screens. Jane remarked “It’s hard to

imagine how to convince people how an additional slate would benefit their reading. Two was definitely the most useful configuration but that wasn't obvious until I started using them." In a similar vein, Anne said, "there's something ridiculous to me about having four of them. Maybe it's that I work in the humanities and no one is ever going to give me four tablets of anything. I'll be lucky to have one." Increased ownership and reduced cost of tablet devices may automatically alter this perception of multi-slate environments. Multi-slate environments may slowly gain traction in the same way that multi-monitor computing shifted from being something of an extravagance to something quite standard.

Summary

The multi-slate aspect of our reading system appears to be effective in supporting reading activities traditionally performed on paper. Users made use of multiple slates to take notes, compare documents, and view different documents, echoing the sorts of activities that multiple paper surfaces support [Adler et al. 1998]. Thus, the deployment supports the notion that an electronic reading system that provides multiple displays can more effectively support active reading activities, confirming Adler et al.'s design recommendations.

A more subtle insight comes when considering the full range of reading activities our participants engaged in. We found that for many activities a single slate was sufficient, and extremely desirable given its favorable physical characteristics compared to paper. Similarly, Adler et al. found that multi-surface reading activities make up approximately 50% of all reading activities. Although this is significant, it also means the other half of reading activities do not require the additional surfaces. The fact that not all activities require multiple screens suggests that the most practical solution is to provide a compact slate with a single display that can be extended with additional devices, as the United States does. This approach is preferable to a device with several screens that are permanently

attached, or one extremely large screen, because it supports the most common use case with minimal overhead while accommodating more demanding reading tasks.

8.5.4 Increased Micro-Mobility

A crucial element that distinguishes the slates from alternative implementations of multiple reading and writing surfaces (such as the use of several windows on one large screen, or fixed multiple monitors) is the considerably higher amounts of micro-mobility that thin, lightweight slates afford. By this, we mean that slates allow for the small, fine-grained manipulations that one often sees when dealing with paper documents: the small shifts in orientation of paper, and the stacking up, moving around and repositioning of paper documents that occur throughout the course of reading and writing [Luff and Heath 1998]. Like paper, the increased micro-mobility of the slate form factor enabled users to work more effectively in a wide variety of reading configurations.

With a single slate, high micro-mobility helped users read in different positions and with greater comfort. For instance, Hannah talked about how at times she would take the slate and crawl into the couch to read with increased focus and concentration. Hannah's experience was mirrored by a number of other users (Barry, Cathy, Anne, and Gina), who also described moving the slates to read on the couch. The additional concentration afforded through holding the device matches Marshall's [Marshall and Ruotolo 2002; Marshall 2005] findings about the importance of physicality when reading. Similar behavior was observed in Morris, Brush and Meyers' [2007] laboratory study, in which users picked up a TabletPC and moved it to read on a sofa in the room. The ability to hold up the slate likely made it easier to see the text as well. To illustrate, one of the cafes where Jane did a lot of reading had dim lighting, making text harder to read. Jane explains that in that cafe "the slates were more pleasant to use than paper. Since the slates were really light, I could hold them closer."

The benefits that micro-mobility conferred extended beyond the single-slate configuration though. The fact that the slates provided increased flexibility or orientation made it that much easier to use the slate alongside the PC. Thus, the slates not only provided extra display space, they provided a more *tailored* use of space.

For Hannah, who used the slates with her computer for dissertation work and job applications, the ability to orient the slates in 3-D space was important. Hannah found herself with slates on her lap, nestled between her body and the laptop (Figure 49), and alongside the laptop on a desk. For instance, while working on her job applications, Hannah needed to make sure her cover letter matched the points specified in the job description. “With the description [flat] on the side, it’s there to refer to, it’s a light screen and it’s much more natural than looking at two upright screens” Recall also that Hannah had previously employed two laptops for her research work. Comparing the experience of using a slate against the extra laptop, Hannah says “[a slate] is more comfortable than two laptops, it’s so light on my lap.”

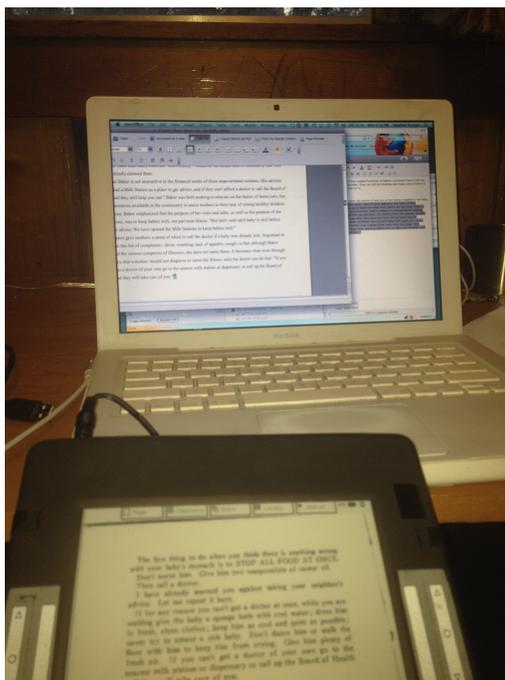


Figure 49 Hannah would rest a slate on her lap to more easily read while writing on the PC.

Increased mobility also allowed the slates to more quickly adapt to changing task requirements. Jane used slates beside her computer as she was revising her journal article. For the task, Jane was referencing documents for the purpose of adding footnotes to the main article. Jane discussed how she set up the slate to the right of her laptop as she was typing (Figure 50, left), allowing her to glance over. Occasionally, Jane would need to read the slates more carefully because, as she describes, “some of these were articles I haven’t read in the past year.” When these situations occurred, the slates accommodated the changing requirements by being able to be picked up (Figure 50, right).

Finally, Justin’s experiences while using the multiple slates as a passenger in a car highlights how micro-mobility helped to make multi-slate use more viable in unconventional reading venues:

“At one point, I had two slates in the car to revise my introduction. I had one showing the Intro and the other with Chapter 2 to check what I wrote in the Intro matched what I said in Chapter 2. I would hold the slates one on top of the other most of the time. When I needed to check something, I would fan the two out to see both at the same time.”

Justin’s description was interesting on a number of levels. For one, his use of multiple slates in the car was unique; most users tended to use multiple slates on flat surfaces like a kitchen table or office desk. More interesting, however, was the fact that the stacking and fanning interaction Justin employed was a technique that was not anticipated.

Summary

Our users’ experiences suggest that the micro-mobility of slates plays an important role in reading. The additional mobility helped to increase engagement, extended where and how the slates could be used (thus, also supported their “macro-mobility”), and helped users to work with the laptop more effectively by allowing the extra display area to be positioned freely in space. The results, especially the benefits of mobility when working with the laptop, corroborate the findings from our controlled study comparing the physicality of reading surfaces. The importance of having at components within a reading system that

exhibits high mobility reinforces the idea that it is desirable to have an environment that consists of at least one highly mobile slate.

8.5.5 Partitioning Information Spaces

The slates' full-screen presentation of materials and support for multiple devices and displays naturally forces users to divide their reading across physically separate devices. Grudin, when studying the use of multi-monitor desktop environments, observed that one of the benefits that arising from physical divisions between displays was "Space with a dedicated purpose, accessible at a glance" [Grudin 2001]. From our users' reports, it appears that the physical divisions between individual slates, and between the slates and the computer served in a similar capacity.

Many users (Justin, Gina, Jane, Hannah, and Anne) felt that the slates provided a reprieve from the distractions of the laptop. Justin and Gina both specifically mentioned that the slates allowed them a chance to get work done away from the computer. For instance, Justin welcomed the fact that the slates allowed him to turn off the computer: "The PC is distracting, I have lots of stuff open on the PC..." Gina mentioned that so much of her other work occurs on the computer that reading away from it was a relief. Hannah, Jane, and Anne, who used the slates with the PC, saw the problem from a slightly different angle. For them, the fact that the slates provided additional display space that *wasn't* a window on



Figure 50 **Left:** Jane would keep a slate on the side of the PC for reference. **Right:** Jane would periodically hold the slate to read in detail.

the computer desktop was beneficial in reducing the amount of distracting task switching they needed to do. Nevertheless, they were in agreement that the separateness of the slates helped them focus on their work. Hannah's explanation about the implications of separation was particularly descriptive:

“When I'm writing, the fewer things I have open, the better. Things like Zotero, email, being open are a pain. The slate was a separate space. And I can stay focused on the writing. It's how you focus your space to prevent distractions. A lot of writing is [about] routine and focus. Maintaining that focus is much more important.”

The physical differences between the laptop and slate screens also influenced the reading experience beyond focus and concentration. Justin remarked that viewing documents with the entire page visible was important to getting a sense of the page as a whole. In particular, proofreading became easier when seeing the full page. Justin, describing the process of editing his dissertation with the slates said, “In general, slates allowed for global editing, like how does the document look as an object?” Moreover, by shifting the editing task away from the PC, made it easier to catch errors. Justin explains, “From my experience, when you edit on the computer, you have a channel between the brain and the document on-screen that does not match the reality of what is actually [written] in the document.” Justin's final assessment of the slates was that “for editing purposes, the slates were not as clarifying as paper, but they had the same effect and got the job done.”

Finally, we also saw some examples where individual slates took on specific roles. Cathy had one slate in the living room for reading and one slate on her writing desk for writing activities. Although she did not make use of two slates together, she found that giving each slate a separate role to be helpful for organizing and dividing up her work. Cathy further added that looking ahead to her dissertation and future work, she could see herself using separate slates to keep track of different projects.

Summary

Our findings suggest that there are clear benefits from partitioning data across physically distinct slates. Among these benefits was reducing one's susceptibility to distractions from electronic devices, providing a direct association between physical objects to a specific role, and presenting information into more logical chunks.

8.5.6 Navigation Support

As part of supporting a variety of reading strategies, the version of the United States that was deployed provided several different ways to help users navigate. The two primary modes of navigation users performed were within-document navigation and between-document navigation.

Within-Document Navigation

For within-document navigation, users had a choice between turning pages manually or switching to the Space-Filling Thumbnail (SFT) overview of the document. The SFT view proved successful in helping users to jump to specific locations in the document without having to repeatedly turn pages. The main situations for which users employed SFT were for getting to a specific point in a document for discussions in the classroom or in colloquiums (Justin, Hannah, Jane), for resuming reading activities (all users), and for skipping to endnotes or references in documents (Justin, Anne).

The situations where SFT proved useful were interesting. Users consistently mentioned that for larger (more than 50 pages) documents, it was hard to see the content of text-heavy documents in the SFT. Therefore, for general random access, SFT was not always useful. However, if the user applied custom annotations to the page, the SFT overview became far more conducive for helping users find locations of interest.

In addition to navigation, the SFT view served other purposes as well. Jonathan, Barry, and Cathy all mentioned using the SFT view to get a sense of how much of a document

remained. Justin had an additional use for SFT that we had not expected, which was to quickly check how frequently long quotes came up in his dissertation. Although these uses for SFT may not be navigation activities per se, they suggest that the ability to assess a document at a global scale is important for reading, especially when the tangible cues of paper documents are missing.

The main area where within-document navigation was insufficient was in speed. When Cathy was quickly reviewing materials for her A-exam¹⁵, she noted that the slates were unable to support skimming activities as well as paper. For similar reasons, Jane and Anne found it easier to go through very large documents on the PC. Navigation speed was particularly crucial in activities where users have little advance idea about what locations they will need to visit. Reading for discussion fell into this category because users could not easily predict what topics the speaker would bring up. Jane and Hannah recalled difficulties when leading class discussion, and attending colloquiums, respectively. Our users did attempt to devise novel ways to try and compensate for the speed, however. For instance, Hannah would have different pages of the document open on additional slates to reduce page-turning time. Hannah's strategy inspires the exploration of new navigation interfaces that would perhaps leverage multiple displays to show the pages that had just been viewed, or anticipates relevant pages to display based on other users' devices.

We also discovered that page numbers played a more important role than we had expected. Page numbers were frequently present in notes and marginalia as the means to record a specific document location. Thus, many users (Hannah, Barry, Jane) suggested that the SFT overlay page numbers of the thumbnails to assist in tasks where the user relies on page numbers as an index into the document.

¹⁵ The A-exam at Cornell University is the examination doctoral students must take to advance to candidacy.

Summary

From the above, we believe that immediate gains to within-document navigation on electronic reading systems can be achieved in a relatively straightforward manner. Faster processors and displays will enable more interactive interfaces that can address the shortcomings we uncovered. Reading systems will benefit from more responsive page turning compared to the 0.5-second to 1-second delay on the slates. The ability to smoothly zoom in and out of thumbnails would likely solve the problem of the thumbnails being difficult to resolve. Many users had expressed that they would have wanted to use pinch-to-zoom gestures on the SFT (as well as the main document view). Lastly the ability to toggle informative overlays such as page numbers or section headings over pages in the SFT would allow the system to alternate between being a visual overview, for preview and proofing activities, or a more structured navigation interface. These solutions, which primarily rely on improvements to visual interface elements, are the easiest to integrate into existing electronic devices. However, more radical approaches that support navigation and awareness about a document in a non-visual manner may ultimately be more important. Paper documents support non-visual interactions like assessing document length by feel and physical place holding for navigation [O'Hara and Sellen 1997; Mangen 2008; Thayer et al. 2011]. Providing comparable non-visual interactions in an electronic system will be useful because it would allow these auxiliary interactions to be performed without conflicting with the primary activity of viewing the content on the screen.

Between-Document Navigation

When working across documents, multiple slates were greatly preferred. Hannah and Jane, who were the primary users of multiple documents together, eschewed the active stack and preferred to load documents on different slates. Most other users did not need to consult more than one document at a time. Users who did use the active stack, like Cathy and Gina, mostly did so to resume reading activities.

However, using multiple slates only addressed one class of multi-document use. Multiple slates helps when the reading task involves repeated navigation between a small set of documents. Information triage activities [Marshall and Shipman 1997], on the other hand require the ability to quickly skim through a large set of documents to determine their relevance. One feature that was missing on the slates was the ability to tell, at a glance, what information a document contains. Users wanted to know this information when they needed to revisit documents used previously. For example, Cathy mentioned that after having the slates for a while, she started having many groups of documents on the slates. With just a listing of document titles, there was no way to tell which of these documents were relevant for the task she was trying to accomplish.

Looking back, it is interesting to note that we had actually seen users employing strategies supporting triage interactions on the PC. Justin had described renaming his document files to include the title, as well as a summary of key points of that document. Hannah and Anne's use of inline notes alongside the document file in research management software probably catered to the same need. We had initially interpreted that these notes were to be used as a summary of the document to be used in lieu of the full text. However, we now believe that these notes also play a helpful role in triage operations.

Another type of navigation was the process of moving from one document to a related one. Document tagging in the Library worked to accelerate this process. Grouping documents with tags was most useful for users like Jane and Justin, who needed to switch between several different types of reading activities. Users like Hannah and Cathy, who had large books that were split up into chapters also made use of the groups to keep chapters of a book together in one place. The other users felt that document organization would only become important if they were using the reading system for a longer period of time. For the time they were using the system, these users thought that finding documents from the full

list to be quick enough. As a whole, accessing documents electronically was preferred over paper. For instance, Hannah explained that once all of her documents had been transferred to the slate, working across different documents was far simpler because she did not need to find and take out paper documents from her binders.

Summary

For simultaneous use of documents, the speed and directness of working with multiple screens was far preferred to multiplexing multiple documents on a single device. However, it is prudent that facilities for switching documents on a single device be made faster and smoother regardless. For one, as our findings in Section 8.5.3 indicated, users sometimes prefer, or are limited to, a single screen. But more generally, there is always the possibility that the number of documents one requires exceeds the number of reading surfaces available. We believe that removing the need to go to a different view to change documents, such as through document switching gestures, and presenting graphical previews of documents rather than a textual description are both good candidates for accelerating the switching process.

As for supporting information triage, we see several possible ways to give users a sense about the content of their documents. Based on user suggestions, the ability to attach pages of notes to a document and to browse those would achieve something akin to the inline notes in reference managers. Also, the ability to skim through a collection of documents rapidly rather than having to open each one separately to view them would enhance the triage process. The skimming could involve document content, annotations, or specific pages of interest. However, what seems to be most important is the ability to blur the divisions between documents for certain cross-document navigation tasks.

8.5.7 Inter-Slate Interactions

Page Linking

The linking system was provided as a means for users to accelerate cross-document navigation. In particular, we believed the system would allow multiple slates to better support cross-document navigation. The linking system turned out to be more applicable for users like Hannah and Jane who would use the slates during group discussions. When engaged in these discussions, Hannah and Jane made use of the notebook, on which they could make notes before and during the discussion about the paper being discussed. Jane, explains her thinking behind the links:

“At first, I would take notes in the notebook like I did with paper, which is to write down the page number. However, I wanted to experiment with the links; I thought this could be more useful. In my notebook, my notes became a link to the page, along with a quick note, which would be a little note to myself. I thought I could use one slate and click”

In practice however, users thought the implementation of the links needed to be improved to for them to be broadly useful. For one, the process of opening links suffered from delays, particularly when the system had to open and render a brand new PDF document.

Hannah explained that she wished the links were smoother in order to keep up with the speaker.

Users additionally mentioned that the presentation of the links could also have been better. Barry explained that the main problem with the links for him was that they did not give an indication about what was at the other end of the link. Jane thought the links should be double-ended, to maximize the chance of her rediscovering associated content should she come across a link. Jane also felt that for her purposes (e.g., substituting for handwritten page numbers) options to align or otherwise tidy up links would have made the linking interface better.

Although Linking received limited use, it appears that the capabilities are useful based on what participants said after the training period. Justin, had been trying out the feature and

mentioned it would be useful for jumping back and forth between notes when only a single slate was available. Gina thought it would be a good way to add structure to documents; she specifically mentioned it would be useful for teaching or assembling a syllabus. Gina did not have a chance to use the slates in that capacity, however, because she was not teaching during the deployment.

Stitching Content Across Slates

Our users reported few instances of using the Stitching/Conduit interactions for transferring documents and opening thumbnails across slates beyond trying the feature out a few times at the start of the deployment. All users said that either they did not encounter situations that necessitated inter-slate navigation, or that directly navigating seemed just as fast.

With Stitching, it was possible to open a document visible on one slate on another slate. But, with only slightly more effort, users manually open up the document using the Library interface they were more familiar with. However, since most of the navigation time was dominated by the computationally intensive task of opening a document and displaying a page, the perceived speed advantage of using Stitching was diminished. The limitations based on speed seem quite significant especially when considering that after training the main reasons participants found Stitching compelling had to do with being able to navigate more quickly. For instance, Hannah did mentioned that the ability to Stitch thumbnails reminded her of being able to save a location while flipping through other pages, like using a finger to save a spot. Anne thought Stitching could be useful for quickly following up on footnotes and endnotes on a second slate. The speed of the slates interfered with both of those intended uses.

Another important factor was that these interactions were only needed sporadically. Users spent most of their time reading, rather than performing Stitching operations to initialize

the slates. As a result, these interactions did not become ingrained in the reading process. For example, Hannah specifically mentioned that when she first received the slates, she was familiar with the inter-slate interactions, but they were not applicable. At one point during a group discussion Hannah thought the Stitching would have been useful, but she could not immediately remember how to invoke the command, resulting in a missed opportunity to use the interaction.

A comment we received came from Hannah corroborates this interpretation. Hannah mentioned that “To me, it was like: wow you can write on a PDF and it’s going to be there on the page and I didn’t have to print it out...that was all I wanted to do. All the other stuff was like “I have to learn other stuff?” Hannah’s comments suggest that unless users repeatedly encounter a glaring obstacle for which the more advanced interactions address, it might not be worth the overhead of learning to use a new feature.

Summary

At first glance, the lack of use of inter-slate interactions like Linking and Stitching could be interpreted as these features being extraneous. However, that particular view might be excessively negative. For one, these interactions were not in the critical path of users’ reading goals. As a result, users had little incentive to rely on these interactions. Also, the general perception that these features were complex and required extra effort to learn suggests that there were likely shortcomings in the current implementation that need to be corrected. Possible directions for improvement include implementing the design changes users suggested, providing better visual feedback about the operation and incorporating more direct manipulation (versus the gesture-based approach we employed) to aid discoverability of these interactions. Once obvious implementation issues have been fixed, the utility of inter-slate interactions should be revisited.

Another way where these interactions may eventually prove themselves useful would be if the application software on the slates provided more opportunities for data transfer. For instance, some users mentioned wanting to extract text from one slate to the notebook on another (a feature we had to disable because the experience was not satisfactory with the current slate hardware). Also, the current deployment involved reading materials geared toward the humanities; other styles of reading and more complex content (e.g., those containing interactive elements) may benefit more strongly from Stitching content from one slate to another. For example, activities that have a more significant navigational component, like traversing documents with hyperlinks, may have more opportunity to take advantage of Linking and Stitching. Finally, multi-user environments may derive benefits from inter-slate interactions. For instance, Stitching would be useful when the target slate belongs to a different user and is not a mirror of the source slate.

8.5.8 Connecting with the PC

Most of our users performed activities where they used the slates alongside the PC.

Although we used some of these activities to illustrate the benefits that arose from slate micro-mobility or the extra display real estate, these activities leveraged the electronic connection between slates and PC in the form of text extraction and remote control.

The connectivity between slates and PC was most applicable to activities where the PC occupied a central role in the activity. These activities were generally writing and editing tasks that involved a significant typing component. Most of the interactions between slates and PC consisted of transferring information from documents on the slate to the PC. Anne and Justin used the text selection and distributed clipboard functions on the system most extensively. Anne used the feature heavily to transfer her notes and excerpts into the notes management software on her PC after she was finished with a book or article. Anne preferred to copy text from the slates, rather than a document on her computer, because

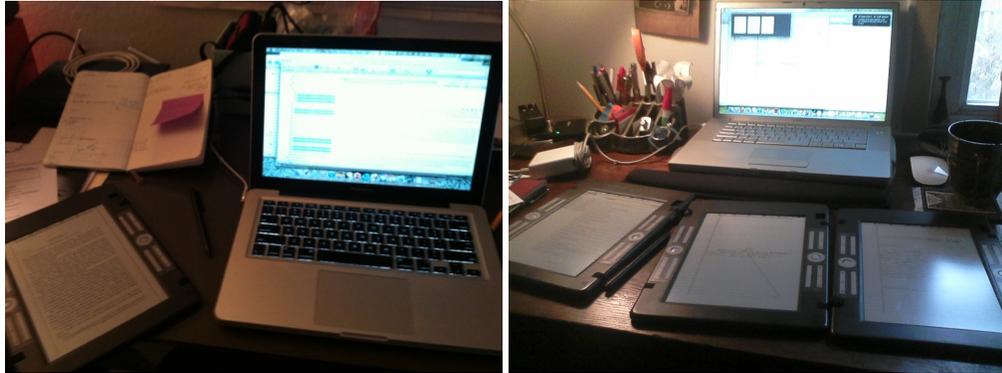


Figure 51 Left: Anne preferred to copy and paste from her slates directly into her notes on the computer. **Right:** Jane would use the Slates application (note application running on PC) to open documents she was viewing on the slates in order to access the digital content.

she found that the selection gestures on the slates offered more precision. Another reason why Anne preferred to work with the slates was that the PDF parsing engine on the slates correctly extracted the text more reliably (Figure 51, left).

Justin made use of the distributed clipboard in a completely different way while editing his dissertation. Using the distributed clipboard, Justin adopted an entirely new editing technique that was only possible with the slate functionality. As he was revising his dissertation, Justin identified portions of the text he wanted to remove or otherwise change. As he marked these segments up with the stylus, he also selected them and copied them into a “scrap pile document” on his PC. Later, when he went to implement the changes, he could make the changes directly, secure in knowing that the original version was safe in the scrap pile. Justin remarked that this new editing process was more streamlined; he could proceed down the list of changes on the slates and implement them, without having to flip between windows, or otherwise interrupt the process.

Jane employed yet another technique to get the text on a slate over onto the PC. Jane’s strategy was to call up documents from the Slates application (Figure 51, right). Jane found copying and pasting using the PDF viewer on her computer to be easier than through the gestural interface.

Summary

In general, all of our users thought that the ability to move text from the slates to the PC would be applicable at some point to their work. The main reasons why the feature did not get used more frequently was mostly due to users not being involved in activities that required that functionality (Gina, Cathy, and Barry). Hannah did need to transfer text between slates and PC, but since most of her documents were scanned they did not provide easy access to the digitized text. Besides Jane, the Slate application went largely unused by the other users. No users appeared to use the Slates application to control the slates from the PC either. When asked about why they did not make use of the Slates application for controlling the slates, users mostly said that it was easy enough to operate the slates directly. Users also mentioned that it was difficult to resolve the contents of the page using the thumbnail provided in the Slates application.

Users thought that even tighter integration with the PC might have been beneficial. Hannah and Anne, in particular wanted better integration with PC tools like reference managers, databases and writing tools. A lot of the organization they had already established on the PC did not transfer over to the slates. For instance, Hannah explained that moving documents from Zotero to the slates was difficult. Hannah had the following to say concerning the difference between the experience of using her slates to grade papers compared to working with her dissertation research: “[Grading papers] is an example of the simple projects where it was much easier to use the slates. The process didn’t need to integrate with the system I already had. To start fresh was important because the inertia problem was bad.” With Anne’s heavy reliance on software tools on the computer, we were slightly surprised she was generally positive about how the slates fit into her workflow. Even so, Anne thought that the ability to more quickly and directly view the content in her tools on the slates would have been beneficial.

Previous reading systems have generally not taken into account the relationship between the reading device and existing computing devices in the environment. For example, when researchers created a version of the XLibris device for legal scholars, a big question surrounding the device was its role relative to laptop computers. Researchers in that study concluded that since laptop functionality was important, the appropriate solution would be to combine the two into a single device to create a “document laptop.” We believe, in light of our results, that a promising alternative is a tightly integrated multi-device configuration. Pairing PCs with slates unifies the capabilities of both of devices, while enabling each device to retain its unique functionality. Given the variety of ways participants used slates and the PC together, it is quite clear that the PC complements the slates in reading activities.

8.6 Discussion

8.6.1 Comparing United Slates with Other Technologies

Table 5 on the next page compares the United Slates to the current state of affairs of using reading technologies. The table details how United Slates supports the themes identified and summarizes the feedback from users.

	Purpose	Current State of Affairs	United Slates Support	Feedback from Users
Writing	Annotations, note-taking, staying engaged with a text	Paper currently the gold standard. TabletPC offers pen-input capabilities	Free-from inking on documents	Not as good as pen on paper but adequate for short notes and markings. “A game changer” compared to alternative e-reading devices.
Access to Data	Reading in different venues, travelling, macro-mobility (portability)	PC holds much of the content, but PC not always available. Paper is cumbersome, have to remember to print out and can be lost.	Interchangeability of devices, annotations applied to the copy of the document on the PC automatically.	Increased peace of mind, ability to archive notes that would normally be lost (“one of the features I will miss most”), more portable than paper, no need to print.
Extra Display Area	Synthesizing information, checking endnotes, comparing documents, reading while writing	Paper well suited for this. Electronic strategies are large displays attached to the PC, or flipping between multiple windows	Ability to use multiple slates for added display area.	Recaptures navigation possible with paper. Superior to flipping between windows. Can serve as a persistent reminder / placeholder in workspace.
Micro-mobility	Staying engaged with material, reading comfort, use alongside PC	Paper thinnest and lightest. Kindles and iPad offer high micro-mobility for electronic devices.	Thin and lightweight, wireless slates that can be positioned independently.	Increased mobility helpful for working alongside the PC, can be used in variety of poses (e.g., reading on sofa, placed next to PC, held up)
Partitioning	Eliminating distractions, seeing a document as a single unit	Paper best for this. Working on the PC fraught with distractions.	Physical separation between slates and PCs.	“Clarifying” like paper. Provides a separate space that guards against distractions from the computer,
Navigation Support	Resuming reading activities, keeping up with discussions, and getting a sense of reading progress.	Paper provides kinesthetic cues about progress, allows for rapid skimming. PC allows rapid scrolling.	SFT document overview. Stack interface to switch between recently used “hot” documents. Document groups in the library.	SFT difficult to use for long, text-heavy docs. However, useful if pages are annotated. Not fast enough for skimming or unpredictable discussions. Often easier to use multiple screens rather than perform explicit navigation.
Inter-Slate Interactions	Enable multiple slates to be used in a cooperative fashion.	Prior systems did not offer inter-slate interaction support.	Conduit and Stitching allowing thumbnail navigation across screens, link creation and visitation.	Linking interface if improved, would be more applicable. Other interactions did not save much time compared to doing things manually. Not so critical for getting work done.
Connecting with PC	Obtaining documents, writing papers, creating electronic notes, sending docs to collaborators/students	Manual process to copy content from paper. Switching between multiple windows on the PC is annoying.	Seamless synchronization with PC version of document, Slates application for remote control and shared clipboard.	Shared clipboard considered very useful. Allowed for easier movement of data to use on the PC. Allowed mixing of digital and paper functionality (e.g., Justin editing)

Table 5 United Slates functionality and feedback placed in context with existing technologies.

Slates' Relationship with Paper

The second column from Table 5 shows that paper has traditionally had the best support for many of the tasks that our users were engaged in. Therefore, a crucial metric for evaluating the effectiveness of the deployed system is to examine how the system compared against paper. For certain aspects of reading, the United Slates offered a reading experience much closer to paper than existing electronic devices. For things like writing, viewing several things at once, and micro-mobility, users commented how the slates recaptured elements of paper and offered an acceptable substitute. However, there were also areas where the slates were seen as surpassing paper. Compared to paper, participants found slates to be more secure, easier to manage, and better integrated with the PC.

One key property of paper that the United Slates shares with paper is a high degree of versatility. Most participants transitioned a large portion of their overall reading activities to the slates. We believe the versatility of the United Slates comes from the combination of the extensibility of the multi-slate configuration, micro-mobility of individual slates, and freeform nature of the writing support. There are obvious parallels that can be drawn between these features and the characteristics of paper.

There are two important points related to versatility. First, versatility is an important dimension that has been under-explored in the design space. A versatile system is important to ensure that users are not overburdened from having too many devices that each fill too narrow of a niche. Therefore, versatility should be considered when evaluating the merits of a system. Second, it appears that paper may be more versatile than the slates at this point in time. Users' familiarity with paper and its capabilities means that paper remains an important tool for increasing versatility in situations when the electronic system does not provide adequate support for a task (e.g., flipping over for more writing space). Therefore, it may be prudent to consider using paper to augment electronic devices. For that task, it is possible that techniques for augmenting paper like PADD [Guimbretière

2003] and PapierCraft [Liao et al. 2008] can help to harmonize between the paper and digital worlds.

One last finding that is pertinent was that some users expressed some concerns that indicated there are some elements of paper may not be compatible with an electronic device at all. For instance, Cathy mentioned worries she had about not being able to encounter texts in the format in which they were published, which make sense given her background as a historian. Although, these concerns were not prevalent, we believe they indicate that paper will continue to play a role in reading for a long time to come.

Slates' Relationship With the PC

The slates' role relative to the PC is complex. On the one hand, users expressed a desire to work away from the PC, given their enthusiasm about the slates providing a sanctuary from the distractions of the PC. But, in both the photo-diary and deployment phases of our study, the presence of the PC in reading activities was constant. The PC not only served as a reading platform, it also served as a tool for writing, for note-taking, for archiving, for conducting queries at libraries, for corresponding with collaborators, and many other tasks that occur in support of reading. It is highly unlikely that reading slates will replace the PC for all of these tasks. Therefore, one challenge of all reading systems is satisfy the seemingly contradictory goals of integrating with PC use, while insulating the user from the PC.

Interestingly, providing superior connectivity with the PC appears to solve both problems. For instance, making data easily available on the slates made it possible to use slates more closely with the PC. The mobility of the slates actually made them preferable to a second monitor. Also, extending the operations that users expect to be able to perform across windows to the slates added new interaction possibilities. But, the connectivity also saved users from having to expend extra effort to move their work on the slates back onto the PC. As a result, users were also able to more successfully work away from the PC. For these

reasons, we believe the loose interaction links but tight data synchronization between the slates and PC provided in the United Slates is ideal.

8.6.2 New Interactions Resulting from the United Slates

The new ways users approached and used documents on the United Slates during the deployment illustrate the ways that electronic reading technology can provide useful functionality not available on paper. From our findings, we identify four areas where we saw glimpses of users taking advantage of electronic functionality to do new things. We believe these areas are ripe for further investigation.

First, an electronic reading system offers a more streamlined approach to working with documents. Since an increasing number of the documents are created, edited and distributed electronically, the digital functionality the United Slates provides results in a much more tightly integrated workflow between reading and document creation. Case in point, users like Justin and Anne took advantage of the ease of moving documents they were writing onto the slates so that they could quickly shift to a revision phase. Justin additionally demonstrated how having a document in digital form could enable a new, more streamlined way of editing a document.

Second, a digital reading system can capture a far more comprehensive archive of reading activities. Having the archive of reading activities was deemed useful for teaching, and when a user needs to revisit materials. Oftentimes, these pieces of their work that would normally have been lost because the cost of converting it to a storable format was too great.

Third, digital reading solutions offer increased security and piece of mind. This observation may seem unintuitive, given that people often print out hard copies for similar reasons.

However, we believe this is an example where adding the correct digital functionality can result in a situation that upends the status quo. The peace of mind that comes with a digital reading solution extends beyond storage and archiving. The physical properties of paper

documents force users to make tradeoffs about what materials they will have available. Electronic reading devices allow the user to delay the decision about what materials to work with by virtue of letting the user bring around much more content than will be used. As a result, a user can be much more flexible about the work they can do on the go.

Lastly, the multi-slate setup enables documents to have useful physical properties that are not available with paper. These properties allow electronic documents to fill niches better than their paper counterparts. A prime example was how users were able view the main text with endnotes on the multi-slate system without having to flip between pages. This type of interaction, which we believe to be superior, would not be possible with a bound paper document. Users do sometimes remove the bindings of documents so that they can view disjoint locations without flipping, but the additional effort as well as resulting tradeoffs limits the practicality of this strategy. With multiple electronic slates, users are freed from making these types of costly tradeoffs.

8.6.3 Improving the United Slates

The main complaints users had about the United Slates system centered on technical issues. These issues were largely due to the fact that the processing power of the slates is one or two generations behind contemporary tablet hardware. The slates were also limited by the fact that they were still research prototypes. Fortunately, these problems can be easily addressed using non-exotic means.

The previous sections about navigation discussed the potential for improved interactions given a more responsive display and more powerful processing. Those are well within reach. Another easily addressable area is power management. Our slates could run for several hours, but users still felt compelled to bring along the power adapter, reducing the portability of the system. Since the power management support on the system was not fully implemented, the slates had to turn off in between sessions to conserve battery capacity.

Though most users got used to the longer startup of cold booting—they would go and occupy themselves with other tasks while waiting—the startup costs did remove opportunities to perform some spur-of-the-moment reading activities.

The effort required to ensure network connectivity also annoyed users. The main problem was that it forced users to have to remember to carry around an additional electronic device (the mobile hotspot) when not at home or on campus. The reliance users had on the connectivity between slates and the PC made network connectivity that much more important. As several users mentioned, the devices in the system worked together so well that it was particularly frustrating when it stopped working. The solution here is to take the approach that many tablet appliances have adopted: integrate a cellular modem directly into the device. This approach would completely abstract the details of network connectivity.

In that light, we believe that a system based on the one we deployed, but correcting the most glaring problems is completely feasible to construct. In fact, doing so may not even require creating new hardware. The widespread availability of tablet appliances matching many of the above criteria means that there are a number of good candidate platforms on top of which a similar system could be created. Our efforts to move the functionality of the United Slates to the iPad are a step in this direction.

8.7 Summary

We deployed a version of the United Slates reading system to Ph.D. students in the History and English departments to understand how a multi-slate reading system would be used in real world reading scenarios. Using a photo-diary technique, we first gathered information about these individuals' reading needs and how these individuals currently satisfy these needs with existing tools. Then, we provided these students with a set of 4 slate devices and asked the students to track how they used the devices for their work. Our participants

ended up using the system for a wide range of reading activities, from which we were able to draw out a number of important lessons. Key findings include the central role of writing in reading activities, the importance of micro-mobility for enhancing comfort and versatility, the benefits from having multiple displays, and the transformative role of having annotations and documents simultaneously available. Although there were aspects of the system that users thought could be improved, the general consensus was that the deployed system offered functionality beyond existing electronic devices. In some scenarios, the digital affordances of the system, such as the effortless capture of annotations, and ability to transport a vast quantity of reading materials offered capabilities exceeding that of paper documents. Overall, we believe the general success the system had in supporting a diverse range of real world reading activities demonstrates that the ideas underpinning the United Slates system have merit and ecological validity.

9 The Big Picture

Testing whether a multi-slate reading system satisfies the claims laid out in the thesis of this dissertation involved multiple stages of designing, building, and testing, as detailed in the preceding chapters. Collectively, these stages provide information about the overall efficacy of using multiple slates for active reading. This chapter ties together the work presented in the previous chapters to offer a comprehensive view into the successes, shortcomings and implications of the multi-slate reading configuration.

9.1 Supporting Reading Using Multiple Slates

The central goal for this dissertation research was to create a reading system that supported active reading better than existing electronic devices and paper. The thesis of the dissertation was that linked electronic slates, as exemplified in the United Slates system, satisfy that requirement. Given the generally positive reception of the United Slates and its precursors, we now discuss what aspects of these systems contributed to their success.

9.1.1 Definite Advantages of Multiple Slates

The aspect in common between the systems tested in this dissertation was the use of multiple electronic reading and writing surfaces. Although the use of paper surfaces has been shown to be beneficial, our findings suggest that these benefits extend to electronic slate surfaces as well. The sections below highlight the different ways multiple slates were found to benefit reading.

Seeing More at Once

Having multiple slates allows users to see more at once. The ability to see more resulted in benefits when working with single as well as multiple documents. Experiences with the Dual-Display E-book demonstrated the extra space from an additional display enhanced local navigation activities. In the controlled study of multiple surface reading, the vast

majority of participants viewed documents side by side to perform the comparison task rather than flipping between pages. Finally during the deployment, the extra display area allowed different locations of the document to be viewed at once, which was particularly helpful for endnotes. The problems this capability addresses are real and significant. Users in the deployment described how reading electronically suffered because things like visiting endnotes were too cumbersome to perform on the PC.

The advantages that come out of being able to see more mainly come from reducing the amount of interaction required to access additional information. Glancing between screens was faster and required less work than alternatives. For instance, when working with multiple documents, users in the deployment found that distributing the documents on slates to be superior to the overhead and potential distraction of flipping between windows. Finally, it bears noting that users derived benefit from the extra content that was visible, even if they were not actively using it. Like paper documents scattered in the workspace, the deployment saw an example where information on slates served to provide helpful reminders about tasks the users still needed to attend to.

The above evidence suggests that the multiple screens of the United Slates offers clear advantages over devices that only present a single page's worth of content, like an e-book or iPad. Even if it is possible to multiplex pages on a screen, having the information at a glance, by virtue of having multiple displays, is desirable. This need for more display area shows why slates, large monitors, or tabletops are uniquely useful for reading. However, display area alone is not the only advantage of the multi-slate configuration. The following sections detail why splitting information across physically separate slates contributes benefits that are unique to the multi-slate configuration.

Versatility In Numbers

The ability for the multi-slate hardware to adapt to changing needs makes the multi-slate configuration somewhat unique. The adaptability of the multi-slate configuration can be understood as an example of a tool offering a low threshold but high ceiling [Myers, Hudson, and Pausch 2000]. The multi-slate configuration sets a low threshold in that the simple case of a single slate by itself supports many reading activities. However, the overall system realizes a high ceiling by making it possible to extend the environment with additional slates with few restrictions. As such, the multi-slate environment is suitable for basic reading activities but can quickly adapt to more complex reading requirements that benefit from having more reading surfaces (such as those observed in O'Hara et al. [2002]). Therefore, in an interesting twist, one advantage of the multi-slate configuration is that it gives the user the freedom to use only a single slate, if desired.

The importance of versatility in numbers was most prominent during real-world use in both the technology probe and the final deployment. We observed that for many reading activities a single slate sufficed. Participants often preferred the single slate given the increased ease of transporting one slate. As some users mentioned, a single slate was generally more portable than the majority of paper materials.

Designing hardware devices that have this quality of having a low threshold and high ceiling will be increasingly important. These types of devices support the most common use cases without burdening users with extraneous functionality, while ensuring the same devices can be employed for more involved tasks.

Micro-Mobility and Physicality

Multiple slates could be used in a more physical manner than devices such as laptops or large tabletops due to their increased micro-mobility. The added micro-mobility stems from the fact that by distributing information and functionality across multiple devices, each device did not need to be as complex. The increased micro-mobility and physicality

enhanced reading comfort, added versatility, and made it easier to work with other computing devices.

In the controlled studies (Section 5.5 and Chapter 7), participants picked up slates to read and commented on how that made the reading experience using the slates preferable.

During the extended deployment, almost all users took advantage of the fact that individual slates were highly mobile to read in different locations, like sofas, cars, kitchens, and offices.

Paper, by virtue of its thinness and low weight, is generally seen as having desirable physical properties. But, in certain cases, the physical characteristics of the slates were more conducive to reading than paper. For example, in the study comparing physicality (Chapter 7), users liked that slates were neater to work with. In deployments (Section 5.2 and Chapter 8), users mentioned that slates would lay flat in situations where books did not. Also, slate surfaces supported single-handed operation.

The high mobility of individual slates also allowed the multi-slate reading environment to quickly transition between supporting different styles of reading. For example, a slate acting in support of work on the PC can be picked up easily and moved into the role of primary reading device. The physicality of the slates also resulted in other benefits. For instance, the tangibility of the slates provided non-visual feedback regarding the spatial orientation of a slate on the desktop. This made it possible for users to write without having to dedicate constant visual attention to the slate. This ability supported interleaving of reading and writing, and also made the slates conducive for discussion-type activities where visual attention is not necessarily focused on the slate.

There were, however, instances where physicality was not necessarily desirable. For example, when using the Dual-Display E-book, participants preferred turning pages with a button rather than using physical interactions for navigation. In the preliminary evaluation of the United Slates, participants could not easily hold on to more than a single slate a time.

Finally, in the study comparing reading surfaces, users found that virtual surfaces were easier to move and had fewer layout restrictions. That extra physicality is not always better is an added reason why electronic systems are promising for reading. Electronic systems allow elements of physicality to be selectively introduced. In the United Slates, physicality was modulated in various ways: Stitching operations substituted for physical transfers and the document stack metaphor replaced having to contend with actual stacks of physical objects. More opportunities to perform this type of modulation might come from physical slates cohabitating with virtual surfaces, which might be realized using slates over an interactive tabletop or by integrating projection technology into the slates. One useful capability of a mixed-physicality configuration is to more easily allow users to switch between the physical or virtual instantiations of a document, depending on which is more convenient.

Separation

Finally, it appears that an equally important feature of the multi-slate configuration is how it reinforces separation. The need to balance interaction unification and separation may appear to be a contradiction but this is not the case. The idea that even linked ensembles of devices benefit from separation recalls Weiser's recommendation for the creation of seamful composition of computing devices [Weiser 1994; Chalmers and Galani 2004].

The United Slates system supported the idea of having well-defined divisions between interaction spaces using several techniques. These techniques included the physical separation between devices, the document stack model that allowed information to be separately assigned to each space, and the ability for each slate to operate independently of the rest of the system.

The separateness of spaces resulted in many advantages. One crucial function that emerged was the fact that separation removed distractions and helped to increase users'

ability to focus on the reading task. The availability of a dedicated space reduced the effort to bring up relevant content and reduced the chance that unrelated materials would distract the user. Another important benefit arising from respecting the separation between slates was that it allowed individual slates to work equally well by themselves or together with other slates. As a result, the United Slates could exhibit complex functionality when it was required, but did not saddle slates being used individually with unnecessary complexity.

9.1.2 Other Beneficial Features

Our studies also found benefits arising from the slate hardware that were not due to having multiple slates. These benefits are important as they illustrate features that would be useful to incorporate into reading devices in general.

Writing

The lack of support for writing was a major problem in existing electronic devices. During the deployment, every participant was acutely aware of the limitations that resulted from not being able to write when reading on the PC. For the participants who employed devices like E-books and iPads, the lack of writing support was always mentioned as a clear shortcoming.

By addressing the lack of writing support, the slates were deemed to be a “game changer” (Section 8.5.1). Participants’ complaints about the writing experience in the controlled studies (Section 7.6.4) indicate there is room for improvement. However, these complaints can also be interpreted to mean that writing was important and central enough for these participants to notice and care.

The reason writing was so important was because it was integral to the reading process. The availability of writing in the United Slates resulted in practical benefits like increased engagement, which users explicitly cited when discussing why the slates were superior to

reading on the PC. The importance of writing plus the fact that there does not appear to be input technologies that support the combination of speed, fluidity and range of information that can be capture via writing suggests that systems aiming to support active reading activities almost certainly need to support writing as well.

It bears mentioning that one of the features in the United Slates that dramatically amplified the usefulness of writing was the fact that user-generated marks were automatically captured and synchronized across all of a users' slates and PC. We detailed these added benefits arising from synchronization in Section 8.5.2.

Synchronization

Synchronization across devices was initially provided during the deployment of the United Slates as an added convenience so that users could avoid manually copying files to every slate. However, synchronization turned out to be far more important than we had expected. The ability to sync is, of course, not a new development; utilities have long enabled devices to sync their data against each other (e.g., Palm HotSync, rsync [Tridgell 1999], Dropbox, etc.). The important lessons from the deployment are the unexpected ancillary benefits that came from having a tightly synchronized system.

First, the increased peace of mind associated with having documents and annotations available on all devices reduced the need to produce hard copies of documents. Second, the automatic capture and propagation of user-created information enabled new interactions that would not have been possible with paper. For instance, the ability to effortlessly track and archive annotations applied over the course of a reading session was deemed to be helpful. Third, device interchangeability afforded by synchronization helped to sidestep practical concerns about the portability of a multi-device environment by decoupling the users' data from the hardware used to access that data. Overall, the combination of these three features formed a major reason why the United Slates offered a usable alternative to

paper, by offering a similar degree of security and reliability as paper documents, while surpassing paper functionality by automatically propagating changes across copies of documents and making vastly more data available.

Integration Into Electronic Workflows

The technology probe (Section 5.2) and deployment of the United Slates corroborates previous observations that modern workflows revolve around electronic documents and other electronic tools more so than ever. Morris, Brush and Myers [2007] had suggested that composition activities would likely occur on the PC and during the deployment we observed the PC being used for composition, and other, activities. Recognizing the important role the PC played, the United Slates was designed from the ground up to be used in conjunction with the PC, which produced several benefits.

First, enhanced connectivity between slates and surrounding devices allowed users to design new workflows that spanned slate and PC. For example, the shared clipboard between slates and PC streamlined text transcription and editing tasks. Second, the fact that the majority of documents used were electronic in origin meant the slates were a welcome tool; the integration of the slates with the PC made it so that users no longer needed to print. Third, the ability to feed the results of activities traditionally associated with paper directly back into the digital world was enthusiastically received. Despite problems during actual use, users in the deployment found the idea they could provide comments to students and colleagues without printing anything out incredibly attractive.

9.1.3 Limitations of the System

Despite the above successes, the present implementation of the United Slates still leaves room for additional improvement and investigation. First, there remain questions about the overall importance of interactional unification. Second, the deployment of the United Slates was limited in scope to graduate students in the humanities. Third, the slates suffered

from some key hardware limitations that restricted how participants used the slates. Lastly, the slates lacked certain digital functionality that would have potentially made them more powerful.

Actual Role of Interaction Unification

Interaction unification, the ability for interactions to occur across devices, was a major part of the United Slates system. However, its significance was inconclusive when considering the results from our evaluations. In the preliminary evaluation of the United Slates system, users were quite positive and enthusiastic about the functionality made possible through interaction unification. Users in the deployment had also felt that the features from interaction unification would be applicable to their work. However, during the actual deployment of the United Slates, these cross-slate interactions were not used extensively. The discussion in Chapter 8 explained some of the possible reasons for this, which included the problems with the design of the interactions and supporting interface, insufficient time for users to incorporate the interactions into their work, or the lack of applicability for the specific reading tasks participants were engaged during the deployment.

Restricted Deployment Domain

Another limitation of our findings is that they are largely restricted to the reading activities associated with students in the humanities. These students' reading activities likely involve longer stretches of linear reading than those of office professionals or students in technical fields. From our personal experiences in an engineering-centric field, quick skimming of technical papers and datasheets seem to be more of the norm. We believe different reading domains will shift how people use the United Slates. For instance, activities that make increased use of skimming and cross-referencing may result in the Stitching and hyperlinking functions becoming more essential (or make shortcomings in their implementations more obvious).

Hardware Issues

Although we produced a novel hardware platform that addressed many issues with existing devices, a combination of limited parts availability, problems with prototype hardware, and bounded engineering resources left some room where the slate hardware could be improved. Three important areas where hardware improvements would produce gains are in visual feedback, control schemes, and power management.

First, our slates, even with fast page flipping algorithms enabled, were considerably less responsive than devices using traditional LCDs. This limited the types of interfaces that could be implemented on the slates. For instance, multiple displays would likely have complemented the content extraction capabilities of InkSeine [Hinckley et al. 2007].

However, it was not possible to create a version InkSeine on the slates given their hardware limitations. Also, increased visual feedback can better convey electronic affordances to encourage users to use the more advanced features the United Slates.

Second, the touch controls on the slates did not play as big a role as we had hoped they would. The touch sensing buttons were intended to provide added input expressivity to aid users in performing random-access navigation and command selection. However, repeated unpredictability with the controls forced us to bypass using the touch controls in our studies. Nevertheless, we continue to believe richer controls offer a great deal of promise in reducing the amount of interaction required to accomplish reading tasks. In future devices, we hope to leverage touch, perhaps in conjunction with haptic feedback, to enhance reading.

Third, the power management of the slates was considerably less robust than that of commercial devices. Reducing the startup time and increasing the battery life of the slates would greatly expand where and when people could read. Although participants were still

able to read in mobile settings during the deployment, we believe improvements in this area would have made opportunistic reading much more prevalent.

Missing Features

Finally, there were features that would have likely been useful for end users that we did not implement. An example of such a feature was text search. Although text search would undoubtedly have been useful for reading, we omitted it after judging that the research insights it would provide might not be as immediately relevant to the central goals of this dissertation. Other interesting digital functions, such as the Reader's Notebook [Schilit, Price, and Golovchinsky 1998], back-links [Marshall et al. 2001], and auto-generated hyperlinks to related content [Price, Golovchinsky, and Schilit 1998] were similarly omitted. Incorporating the full range of electronic functionality possible on a computing device and testing it would likely have made the United Slates an even more compelling alternative to paper. These features could also be further adapted to further benefit from the multi-slate aspects of the United Slates system.

The limitations identified above suggest that research into multi-slate reading is not yet complete; further investigation is possible and necessary. Follow up studies can take the form of improved devices, adjustments to the interface and interactions, the use of the United Slates system for reading activities in other domains (such as those associated with office work or technical tasks), or the inclusion of additional digital functionality that would take advantage of the multi-slate interactions.

9.1.4 United Slates In Context

Table 6 revisits the table of active reading requirements presented in Chapter 2 (Section 2.5, page 31). In Table 6, we populate an extra column for the United Slates. Below, we provide justification for these choices based on the results of the different evaluations presented in this dissertation.

	United Slates	PC	Tablet	Tabletop	Tablet Appliance	E-book Appliance	Paper
KEY							
	○ low support						
	● moderate support						
	● () good support () denote target experience						
Hardware / Physical Requirements							
Mobility; thin, lightweight, graspable	●	○	●	○	●	●	(●)
Readability	●	●	●	○	●	●	(●)
Support for writing	●	○	●	●	○	○	(●)
Page Level Requirements							
Linear reading	●	●	●	●	●	●	(●)
Illustrations	○	●	●	●	●	○	(●)
Superimposed annotations	●	○	●	○	○	○	(●)
Text search	○	●	●	●	●	●	○
Glancing back to re-read	●	●	○	●	○	○	(●)
Kinesthetic cues	○	○	○	○	○	○	(●)
Document Level Requirements							
Non-sequential navigation	●	●	●	●	●	○	(●)
Building cognitive map of content	●	●	●	●	○	○	(●)
Skimming to get overview	○	●	●	●	●	○	(●)
Discovery of topical knowledge	●	●	●	●	●	○	(●)
Switching between reading styles	●	●	●	○	○	○	(●)
Workspace Level Requirements							
Sorting documents	●	●	●	●	○	○	(●)
Spatial layout	●	○	○	●	○	○	(●)
Extracting information	●	(●)	●	●	●	○	○
Reading from multiple documents	●	●	●	●	○	○	(●)
Integrating with PC workflows	●	●	●	●	○	○	○
Macro-Mobility Requirements							
Reading in different venues	●	(●)	●	○	●	●	○
Filing and archiving	●	(●)	●	●	○	○	●
Restoring reading workspace, resuming reading activities	●	○	○	○	○	○	○

Table 6 Updated comparison of how well different electronic reading devices support active reading activities.

The use of a custom hardware slate with explicit design goals helped the United Slates satisfy the hardware and physical requirements for reading. Participants were pleased about the micro-mobility, readability, and particularly the writing support on the slates.

At the page level, the United Slates supported basic reading activities as well as existing electronic devices. However, by virtue of more display real estate and proximity links, glancing back to re-read was easier. The limitations of the black and white e-paper screen limited the types of illustrations the United Slates supported, but the use of alternative display technologies such as electrowetting [Hayes and Feenstra 2003] (e.g., Liquavista) or interferometric modulation [Qualcomm 2011] (e.g., Mirasol) can potentially solve that

problem while preserving excellent readability characteristics. We did not implement text search in the United Slates system but there appear to be no barriers to doing so. At the moment, United Slates still lacks the helpful non-visual cues that paper provides, and addressing that issue should be a priority for future work.

At the document level, the ability to skip around using Space-Filling thumbnails and also being able to view disjoint locations of a document side-by-side dramatically was a clear advantage of the United Slates. Occasionally, support for non-sequential navigation surpassed that of PCs and paper. The micro-mobility of each slate contributed to the United Slates' ability to quickly switch between different styles of reading. The ability to obtain a cognitive map of the content was not something that was explicitly tested, but participants' use of the thumbnail overview to get a global picture of the document appeared to be in that direction. The slower screen of the slates limited the United Slates' utility for skimming and knowledge discovery, but these were not fundamental limitations.

The workspace level is where the United Slates begins to exhibit its unique strengths. The United Slates shined in its integration with other electronic devices. Users in fact preferred reading on the United Slates even while they were working primarily with the PC. The ability to work with many documents at once was also cited as an advantage of the system. Also, the United Slates system contributes the finding that maintaining focus and insulating the reader from the distractions presented by other devices is an important feature of using physically separate reading slates. Although the United Slates document model provided a mechanism for sorting, we did not perform evaluations to confirm its utility.

Finally, the area of macro-mobility is where the United Slates differentiated itself most from paper. The United Slates' ability to make reading materials available, and then also archive activities on these materials were greatly appreciated by participants in the deployment.

These qualities were the ones that our users missed the most at the conclusion of the deployment.

Overall, this more complete support of active reading requirements, which we hypothesized at the outset of the dissertation, and confirmed through testing, illustrates that a multi-device approach can preserve the unique affordances of its constituent devices while increasing the overall capabilities of a system—something that might not have been possible with a monolithic hybrid. The success of the United Slates also confirms the integral role that multiple reading and writing surfaces, regardless of whether paper or electronic, play in supporting active reading activities.

9.2 Comparing to Augmented Paper Approaches

Chapter 2 of this dissertation drew a distinction between augmented paper approaches and systems that drew inspiration from paper. Both of these approaches share a common goal of striking a balance between adding digital functionality while retaining desirable paper affordances. Therefore, it is not entirely surprising there is overlap between the features of the multi-slate system and augmented paper systems. For example, augmented paper systems like PapierCraft [Liao et al. 2008] also provide connectivity with the PC, support multiple reading and writing surfaces, text extraction and the capture of written markings. Portable augmented paper systems that provide output such as PenLight [Song et al. 2009] and MouseLight [Song et al. 2010] can overlay extra content and add dynamism to printed content while retaining the mobility of paper.

Even so, these approaches are not completely interchangeable because electronic reading surfaces such as slates offer certain capabilities that paper surfaces cannot easily match. First, since augmented paper approaches depend on users reading on paper to convey paper-like affordances, situations where reading on paper is undesirable also afflict augmented paper approaches. For instance, given the difficulty of transporting and

working with large quantities of paper, the United States' ability to call up and work with a near-limitless number of documents cannot be replicated. During the deployment this ability was useful, as it enabled opportunistic reading activities that would not have been possible with paper (e.g., Justin's ability to grade student papers when he had some downtime).

The second big difference between slates and paper stems from the nature of the dynamic display. Barring dramatic improvements in projector technology, the type of information that can be reproduced will be different. Augmented paper approaches impose limits on how much the content can change before a new paper version must be regenerated. For PDF documents, which are relatively static, slates and augmented paper approaches do not differ significantly. However, the limitations of augmented paper are more apparent when documents can change in more dramatic ways. For instance, users of augmented paper approaches cannot easily adjust the text size while reading. Larger-scale changes to documents that can be edited, such as word processor documents and text files, also cannot be represented when using an augmented paper system.

Finally, the combination of the display characteristics and limited mobility of augmented paper means it places limits on the types of documents that can be used. For example, augmented paper approaches would be ill suited for reading documents containing hyperlinks. With augmented paper approaches, the destinations of the links would necessarily have to be viewed electronically, eliminating any advantages gained from reading on paper. Moreover, visiting and viewing multiple destinations simultaneously would be quite difficult.

The ability to view a wide variety of electronic documents will likely grow in importance. Electronic document technologies offer a wealth capabilities such as hyperlinks, interactivity, and kinetic typography [Lee, Forlizzi, and Hudson 2002]. Therefore, as

mentioned earlier in Section 2.8, one forward-looking reason to prefer electronic reading surfaces over augmented paper surfaces is that they open the door to new forms of reading that have yet to be popularized.

9.3 Recommended Approach for Building and Testing HCI Systems

Most studies in HCI use off-the-shelf hardware and controlled laboratory testing. The combination of a) hardware being a factor in the reading experience, b) complexity and differences in how people read, led us to believe this approach would lead to inaccurate or incomplete findings. Therefore, the research strategy of this dissertation employs a somewhat unconventional approach in the literature that required creating custom hardware and accommodating real world challenges.

Looking back at the evaluation results, the benefits arising from this research approach included the ability to test a hardware configuration that did not yet exist, and also the opportunity to observe a more complete cross section of reading activities. The usage during the deployment was more reflective of what happens in the real world, consisted of more complex activities, and produced instances of unexpected use. In the past, creating deployable hardware was the exclusive domain of industry given the large set up costs and labor required. Fortunately for researchers, these barriers are much lower given the availability of rapid prototyping technologies and ease of access to cloud resources.

For situations that demand custom hardware, rapid prototyping technologies enable the creation of hardware prototypes that are small, lightweight and mobile, supporting activities that require these properties. Furthermore, these technologies allow the creation of devices that can stand up to the rigors of real-world use, enabling their use in extended deployments, which allows the benefits from extended deployments to be realized.

The ability to create deployment-ready systems in the cloud without much additional effort is another significant development. The cloud-based approach offers advantages of a

potentially simpler architecture, increased robustness, and scalability. Moreover, the cloud approach is cross platform, meaning it is appropriate for a broader population of users and is friendly to both custom hardware as well as commercial devices. This makes it easy to transition between research prototypes and more readily available commercial devices (which may be required for reasons of cost, volume, or robustness). For example, one way of making the ideas behind the United Slates available to a wider audience is port it to a commercial tablet appliance like an iPad. The cloud architecture that the United Slates employs accommodates this hardware change with minimal modification.

In summary, recent developments in rapid prototyping and cloud services offer significant enhancements to the process of creating and testing HCI systems. First, these developments free researchers from the compromises and limitations of using commercially available devices in roles for which they were not designed. Second, researchers can easily create systems that are robust and ready for deployment at the outset rather than having to duplicate work later on in the evaluation cycle.

9.4 Summary

This chapter provided a synthesis of the findings from the studies presented in previous chapters of this dissertation. In this synthesis, we presented several key advantages that can be attributed to the unique, multi-slate nature of the United Slates system. By showing that distinct advantages arise from the multi-slate configuration, these findings reconnect to ethnographers' hypotheses that providing independent reading and writing surfaces would better support active reading activities, thereby extending theories in that area. For designers of reading systems in general, this chapter also highlights useful features of the United Slates that were not tied to the multi-slate configuration but were helpful nonetheless.

The evaluation of the United Slates system also provided key pieces of evidence with which to compare the relative merits of the United Slates versus alternative systems that seek to bridge the paper and digital worlds. In particular, the evaluation produced concrete examples of how multiple slates offer useful functionality that is not easily matched using augmented paper systems—an area about which we were only able to speculate previously. The successes deploying the United Slates also resulted in lessons about the value of and recommendations for performing real-world deployments, which were also presented in this chapter. We believe these lessons will be helpful for researchers hoping to embark on a similar program of research.

Finally, this chapter details the limitations of the United Slates and notes areas for future investigation. Although users expressed enthusiasm for United Slates system, many features of the United Slates were not employed extensively in the deployment. Possible reasons for this range from reading tasks and reading materials that do not necessarily leverage these capabilities, to issues with the present user interface design, to hardware limitations.

Further studies are needed to more conclusively confirm or reject the hypothesis that inter-slate interactions are crucial to a multi-slate reading system.

10 Conclusions and Future Work

Paper supports reading activities well because it has desirable characteristics that allow it to satisfy not just the requirements associated with active reading, but also reading in general. This dissertation showed that a reading system consisting of several lightweight and interconnected slate devices provides a similar range of functionality and presents a compelling alternative to paper. In lab studies, and in the extended deployment of the system in the real world, we found that the multi-slate reading system could be used for reading activities for which existing electronic devices were not satisfactory. These activities included lengthy linear reading, reading that required writing or annotations, and cross-document reading. Moreover, the digital functionality the system provided, particularly in the areas of increasing portability, availability, and interoperability with the PC, proved to be beneficial.

One of the key reasons behind the success of the United Slates system is that it combines the positive aspects of a variety of computing devices: the mobility and readability of e-book readers, the pen input of TabletPCs, the general computing abilities of PCs, and the spatial layout capabilities of tabletops. We managed to pool these qualities by adopting the approach of linking together several independent devices. This approach allowed each device to retain its unique affordances, like high mobility, while allowing them to operate synergistically in support of complex tasks. In effect, this dissertation provides key evidence that with appropriate interaction techniques, a computing environment consisting of an ensemble of devices can offer capabilities that would be difficult to realize on a monolithic system.

10.1 Contributions

In this section, we revisit the contributions presented in the beginning of the dissertation and summarize how our work makes these contributions.

An analysis of the requirements for work-related and academic reading and the degree to which existing computing devices support these requirements.

By aggregating the literature about the process of reading, we provided a more concrete understanding of what it means to support reading activities. Table 1 (page 32) shows that reading requirements extend beyond what one can do at the page-level, or even the document-level. Adequately supporting the needs of active reading means a system must also consider how multiple documents interact with each other, and also ensure macro-mobility needs like the ability to read in different venues.

With a better idea of what comprises active reading, we were able to more effectively determine why people often prefer reading on paper to reading on existing electronic devices. Our analysis showed that existing electronic reading solutions each provided targeted functionality that was comparable (or sometimes superior) to paper. However, no single device provided as complete of coverage of reading requirements as paper, even taking into paper's shortcomings in certain areas like portability and inability to work with existing computing devices. From this analysis, we developed the insight that an effective alternative to paper would need to address reading needs in a comprehensive manner akin to paper.

Design and implementation of a multi-slate reading system

Starting with the Dual-Display E-book, we explored the benefits that would arise from a reading system with multiple independent reading and writing surfaces. The Dual-Display E-book validated hypotheses about the navigational benefits of adding a second display and

also proved instructive in establishing that the physical properties of a device contribute to the overall reading experience.

We then tested whether a multi-slate device configuration would support a broader set of reading activities. Given the lack of an appropriate hardware platform on which we could implement a robust multi-slate reading system, we created our own. Since building custom hardware to fill an available niche is now realistic for even small research groups, we believe that the implementation details we provide offers guidance to others who wish to do the same.

We then developed and tested a set of interactions for multi-slate reading to leverage the new affordances of the multi-slate configuration. The interactions allow users to leverage the combined display area, streamline operations that span multiple devices, simplify management of documents across devices, and more tightly connect the slate devices with existing computer tools. An evaluation with 12 users validated the interactions we developed.

Last, we showed how the functionality of the multi-slate system could be implemented in a way that would survive the rigors of a real world deployment. Our solution relies on shifting a significant portion of the system to cloud-based services. In creating the cloud-based infrastructure, we discovered that the maturity of cloud platforms now makes it a compelling component to incorporate into research systems from the very start. Doing so adds simplicity, robustness, and makes a system deployment ready with roughly the same amount of work as a non-cloud-based solution. We believe this insight will allow research systems to be tested in the wild with increased ease and frequency.

Evaluation of multi-slate reading systems

We conducted two studies to better understand the effects of using a multi-slate system for reading activities.

The first, a lab study with 12 participants, explored the consequences of using slates for reading activities vis-à-vis alternative technologies like paper and tabletop computers. These three systems offer vastly different physical characteristics. Our experiment was unique in exploring the role of physicality by employing three distinct reading tasks. Our study revealed that the tangibility of the slate condition served a useful role. However, we also made the interesting finding that a beneficial characteristic like tangibility can also be a disadvantage depending on context. Therefore, our recommendations center on providing functionality that imparts the illusion of physicality to virtual reading surfaces, or provides non-physical interaction alternatives for tangible surfaces.

The second was an *in situ* evaluation where the United Slates reading system was evaluated with 7 doctoral students at Cornell University. Through this study, we were able to report on how a multi-slate electronic reading system integrates into and shapes real world reading activities, something that would not have been possible with a laboratory study. Participants used the system for 3 to 4 weeks on average. Students successfully used the United Slates to read books and articles for courses and research, to teach, and to compose and edit papers. The enthusiastic response that the students had toward the multi-slate reading system, plus the wide range of reading activities for which the system was employed confirmed our hypothesis that such a system would serve an effective electronic alternative to paper documents. Further, the use of cloud-based technologies during the deployment turned out to be a serendipitous choice with unexpected benefits. Users in particular appreciated the superior integration of the United Slates with their PCs, and automatic synchronization of documents and annotations, which took the place of and often went beyond paper functionality. The deployment also identified possible areas for improvement and future research, particularly in the area of cross-slate interactions.

As a whole, the findings from these two evaluations allow us to provide concrete recommendations for how to support reading activities electronically. The findings also confirm longstanding hypotheses regarding the importance of multiple reading and writing surfaces for active reading. Finally, these findings demonstrated that the United Slates system provides an effective, and sometimes superior, alternative to existing reading technologies for active reading.

10.2 Future Work

The first area of future work involves providing the capabilities demonstrated in this dissertation with off-the-shelf hardware. The enthusiastic response of those who used the United Slates strongly encourages us to make a system with comparable functionality available to a wider audience. The current limitation we have comes from having a limited number of our custom slate devices. However, present-day hardware is superior in many respects to the devices we employed. For instance, the newly available third-generation iPad with its ultra-high resolution screen, long battery life, and superior networking capabilities is a good candidate for broadening the set of users we are able to reach. Moreover, the cloud-based infrastructure that underpins the operation of the system is hardware agnostic and could easily accommodate these alternative devices. As commercial devices in a slate form factor drop in price and become a regular staple of people's computing environments, we believe that the multi-slate reading configuration will see more widespread adoption. The increased adoption of such a system will in turn enable longitudinal reports of impacts of new reading technology at the multi-year or multi-decade time scale [Marchionini 2000].

The second area of future work is to expand the system into supporting multi-user scenarios. The research presented in the dissertation has focused on a single user, multi-device environment. However, the users to whom we have shown the device have

consistently mentioned the potential of a multi-slate reading system for collaborative activities. The Ph.D. students in our deployment described rich group interactions in teaching, colloquia, reading groups, and writing seminars. The multi-slate configuration seems naturally suited for these scenarios as the physical independence of devices automatically establishes boundaries between private and public space. Furthermore, the cross-device interactions we described in this dissertation would have utility in ensuring that users working together can easily exchange data and use their devices cooperatively. However, difficult problems surrounding user identification, device ownership, and effective multi-user interfaces must be solved in order to make a multi-user version of the system viable. There are many pointers outlining the way to proceed in this area, however. For instance, Hinckley et al.'s work on Stitching [Hinckley et al. 2004], from which this dissertation draws much inspiration, also introduces the idea that it is possible to incorporate the principles of proxemics in multi-user scenarios. Also, the pioneering work on collaborative editing in MATE [Hardock, Kurtenbach, and Buxton 1993] provides inspiring ideas about what an effective multi-user interface might encompass. In fact MATE seems quite amenable to being re-implemented using multiple slate devices.

The third area of future work is to push the boundaries surrounding the implication that new capabilities can be derived from the composition of multiple electronic devices. The increasing diversity of devices motivates exploration of new types of mixed device environments. So far, we have primarily examined environments that consist of slate devices and personal computers. Other combinations of devices offer exciting new possibilities.

For instance, in our studies comparing the physicality of reading surfaces, we showed that the virtual reading surfaces and also paper reading surfaces offer unique advantages that slates do not provide. Therefore, systems that combine slates with paper and slates with

tabletops will offer new possibilities. Projects like Bonfire [Kane et al. 2009] and LightSpace [Wilson and Benko 2010], which mix tangible components with interactive surfaces, are delving into this space. A more systematic exploration of the different ways distinct computing platforms can be used cooperatively should reveal additional synergies. An even more radical proposition is the possibility of designing versatile, highly mobile hardware devices that can be used together to form a customized environment of slates, laptops, and large interactive surfaces. Such an arrangement will address growing issues of device overload without sacrificing the functionality gained from device specialization.

Successfully addressing the areas of future work identified above will extend the work of this dissertation beyond the realm of active reading. In addition to reading, these avenues of future work will likely shape how people collaborate and use computers. This dissertation and the work it spawns contribute to a vision of the future where it is much easier to access, peruse and share information, anytime and anywhere.

Appendix A: Dual-Display E-Book Interview Questions

These interview questions were used in the evaluation of the Dual-Display E-book presented in Chapter 3.

Observations During Usage

Did they make use of thumbnails?

Did they use the wheel?

Did they use the button?

Did they use fanning?

Did they use flipping?

Questions for Both Single Display and Dual Display Conditions

How does this device compare to finding articles in a regular printed media (magazine, newspaper)?

How does this device compare to finding articles when reading online (web)?

Did you have difficulty reading the content? How would you improve the reading experience on the device?

Did you have any problems with the system responsiveness when you were reading?

Did you have any problems with the system responsiveness when you were trying to find an article to read?

Was the thumbnail feature useful for you while looking for new content to read? What were the pros and cons of this feature?

Evaluate the size of the thumbnails. Were they too small? Could they be reduced and still be useful?

Did you have any issues using the controls? For example, control location or responsiveness? Which controls had problems?

Wheel Control

Did this control behave responsively?

Did this control work reliably?

Did this control behave the way you would expect it to behave? If you got to decide how the control works, would it do the same thing?

Would you say this is a crucial control?

Was this control good for moving one page at a time?

Was this control good for moving multiple pages at a time?

How much effort did it require to use the control?

Button Control

Did this control behave responsively?

Did this control work reliably?

Did this control behave the way you would expect it to behave? If you got to decide how the control works, would it do the same thing?

Would you say this is a crucial control?

Was this control good for moving one page at a time?

Was this control good for moving multiple pages at a time?

How much effort did it require to use the control?

Fanning Control (Only for Dual-Display)

Did this control behave responsively?

Did this control work reliably?

Did this control behave the way you would expect it to behave? If you got to decide how the control works, would it do the same thing?

Would you say this is a crucial control?

Was this control good for moving one page at a time?

Was this control good for moving multiple pages at a time?

How much effort did it require to use the control?

Flipping Control (Only for Dual-Display)

Did this control behave responsively?

Did this control work reliably?

Did this control behave the way you would expect it to behave? If you got to decide how the control works, would it do the same thing?

Would you say this is a crucial control?

Was this control good for moving one page at a time?

Was this control good for moving multiple pages at a time?

How much effort did it require to use the control?

Device Characteristics

Was the device too light or too heavy?

Was the device too thin or too thick?

Was the device too small or too big?

Wrap Up Questions

Can you discuss the pros and cons of using one and two screens for reading?

Can you discuss the pros and cons of using one and two screens for navigating?

Appendix B: United States Preliminary Evaluation Materials

These materials were used for the lab study presented in Chapter 5. The materials consist of a questionnaire that participants completed and interview questions and procedural notes for the researcher.

How confusing was it to “send” a link?

1 2 3 4 5 6 7
Completely straightforward Extremely confusing

How useful/applicable do you think this feature would be for your reading?

1 2 3 4 5 6 7
Not useful at all Extremely useful

Screen Teleport

How did using the screen teleport function compare to reaching?

-3 -2 -1 0 1 2 3
Much Worse No Difference Much Better

How confusing was the process of using one device to access another?

1 2 3 4 5 6 7
Completely straightforward Extremely confusing

How useful/applicable do you think this feature would be for your reading?

1 2 3 4 5 6 7
Not useful at all Extremely useful

Moving text to the PC

How confusing was the process of copying and pasting to the PC?

1 2 3 4 5 6 7
Completely straightforward Extremely confusing

How useful/applicable do you think this feature would be for your reading?

1 2 3 4 5 6 7
Not useful at all Extremely useful

PC Control

How did controlling the devices from the PC compare to operating them directly?

-3 -2 -1 0 1 2 3
Much Worse No Difference Much Better

How confusing was the process of using the PC to control a device?

USER # _____

Background Interview About Reading

I'd like to start out by finding out a little about the reading that you do. I'm primarily interested in the reading you need to do for class or for work, rather than for pleasure.

If you think to your classes / work, can you describe the materials you've had to use in those classes?

Are any of these things electronic in format?

How often do you find yourself printing things out? Are there things that you tend to print out?

What types of reading do you do on electronic devices (e.g. laptop, phone, e-book reader)?

Do you have a preference about the format of your readings?

Where do you do most of your reading? Do you read in different locations? Where?

When was the last time you had to do a significant amount of composition (i.e. writing a paper, report, etc.)?

Do you ever find yourself using multiple documents, books at the same time?

Do you ever find yourself reading across different mediums? (For example, paper books along with the computer)?

When you're doing that, is it a pretty smooth process? Or do you wish certain things might be improved?

Basic Annotation and Navigation

Scenario setup

One of the main differences between this device and others

Things to cover:

Writing

How to turn pages

How to move around (Stacks, Page, Overview)

Command button – important for invoking commands

Margin notes

The stack is just a collection of documents we'll be using for the experiment.

Try moving around, scribbling, switching documents.

Introducing More Devices

Scenario setup:

What I am trying to test is whether providing extra screens can assist the reading process. With the functionality I just showed you, you could look at different things at once. However, I want to test features that might make that process simpler.

Growing Margin

One idea is to use the extra device as an

Basic:

Create notes on several pages

Make multiple notes

Browse through notes

Advanced:

Hook up devices

Make notes using 2nd device

Browse through notes

Having used one screen and two screens

Are there any other things in your experience where something like this might be useful?

What do you think are the advantages of this connection approach?

What do you think are disadvantages? What did you think was confusing?

Do you have any suggestions about how we might improve this?

Thumbnails

Scenario:

One of the things we think people will need to do is to switch between several documents. Having two screens can help because you can keep two things visible at once. And to jump to a specific page you can use the overview.

However, with more than one device, we thought that jumping with multiple overviews available is even faster. We think this might be a good option because you tend to only read one thing at a time. To do this you can “**grab**” a thumbnail.

Reason it’s called grab is because I’m pulling it to the screen I’m holding.

Basic:

Using the standard process, move between the Fruits and Animals using a single device

Advanced:

Do the same but with two overview panels open and jump between the devices

Are there any other things in your experience where something like this might be useful?

What do you think are the advantages of this overview-based navigation?

What do you think are disadvantages? What did you think was confusing?

Do you have any suggestions about how we might improve this?

Hyperlinks

The Scenario:

Documents frequently link to each other. We might have footnotes and endnotes, which are links from one part of the document to another, and then we have

references, which are links to other documents. Computers make automatically following links and coming back much easier.

With one screen, you can explore links like you would on a web browser. Links are shown with a rectangle. Tapping on the rectangle takes you to that location.

With two screens, we think that certain linking tasks might be made easier because you do not have to backtrack as much.

Command button to open link options, release goes away

Grab and Send

Grab a lot like thumbnails.

Reverse is possible. Show how to initiate a “send” link operation. If you want to open a link on a different device. This allows you to keep the thing you’re looking at

1. The paper cites 3 papers. Try using one device at a time to navigate to each of the references.
2. Now do the same by sending links
 1. Suppose you want to check out the references to see if they are useful. With one device, visit each one, read the abstract and go back.
 2. Now, try it with two devices, one with the list of references, and grab the links to the device

Are there any other things in your experience where something like this might be useful? How might you take advantage of this feature?

Do both grabbing and sending make sense?

What do you think are disadvantages? What did you think was confusing?

Do you have any suggestions about how we might improve how the linking system works?

Screen Teleporting

One of the great things about spreading documents around on multiple tablets is that the tablets can be spread out on the desk, like paper. The thing about paper is that to work with something that's far away because you can stretch, or you can pick up the paper.

With tablets, picking things up is a bit harder. I'm going to have you try that in a just a bit.

Instead we have this feature called screen teleporting. By holding the command button and then pressing the command button of another device, you can send the screen over and start using that device.

I'm going to have you

1. Change docs to animals
2. Navigate to an animal of your choice
3. Color in the animal

Do it with reaching, picking it up, and by teleport.

Compare working with the device by reaching to picking it up, to teleporting it.

Does the teleporting make sense? Did it give the illusion that it sent that device over to you?

Where do you think this might apply to your reading activities?

How would you make this better?

Moving Text to the PC

The research is to work with multiple screens, not just e-book screens but also that of your computer. We think this is important because a lot of reading activities are in support of other activities on the PC. For instance, like writing a paper.

One thing for you to try is to copy and paste text to the PC.
Hold down the command button and circle or underline text to select it
Tap on copy to move it to the clipboard.

Does this make sense?

Do you think this will be useful? How will it apply to your reading activities?

What do you think are disadvantages? What is difficult or awkward?

How would you improve?

Controlling from PC

So we showed how we can move data from the reading device to the PC when using the device. Another situation may be that we are using the PC and need to change the document or page on one of our devices.

For instance, I want to look at Figure 5 in the Earpod paper while writing. One can obviously pick up the pen and turn pages manually. What I'd like you to try instead is a PC remote control interface

Works like this:

Ctrl F2 pops open a window with your devices

Click on device to select it.

Navigate, change the page, etc.

Click outside of the virtual device to send it back

Start typing a sentence. "the quick brown fox jumps..."

Now, change page to go Pineapple

Go to figure 4

Type another sentence. "over the lazy dog"

Change page of Animal to go to SemTouch figure 4

Try that again using remote control (Ctrl F2)

Did you find one easier to do than the other? Why?

Did one way feel clumsier than the other?

Do you think this sort of remote control would be useful for you?

What do you think are problems with this approach?

How would you improve it?

Stacks and the Document Library

At the beginning of the experiment I said that stacks were just a set of documents you could switch between. The reason we want to just have a subset of documents at any given time is because you may potentially have hundreds of documents – as many as you might have paper. And you may need some at some point, and others at another point.

To help simplify, we have this idea of stacks, which is a set of documents you pick that are important. Additional things you look at get added to the stack.

How stacks get created

How to add documents to a stack

How to replace a stack

Sending and grabbing stacks

How to look for a stack you were using before

How to save a stack for later

Try out:

Add a couple of documents that start with A to your stack. Check how that goes on top of the experiment docs you've been using.

Find some books you like add them to a new group called "Faves"

Go to another device and load up "Faves", notice how the stack updates when you replace a stack.

Power 1st device off.

Recover the stack on the 1st device on a different device.

Is this concept confusing?

What were the differences you noticed using the stack system versus manually opening documents?

How might it apply to your reading activities?

How would you improve the stack system?

Appendix C: Physicality of Reading Surfaces Questionnaire

The following materials were used for the lab study presented in Chapter 7. They include a questionnaire that participants completed after each condition, a global ranking questionnaire that participants completed after all conditions, and interview questions for the interview at the end of the experiment.

How easy was it **to layout or arrange** the documents you needed during the course of the task?

|_____||_____||_____||_____||_____||_____||

I didn't do this

Very
easy

Very
difficult

How easy was it **to find** the documents you needed during the course of the task?

|_____||_____||_____||_____||_____||_____||

I didn't do this

Very
easy

Very
difficult

How easy was it to **mark up or annotate** the documents as you were reading them?

|_____||_____||_____||_____||_____||_____||

I didn't do this

Very
easy

Very
difficult

How easy was it to **read across multiple documents** to compare or contrast the information on them?

|_____||_____||_____||_____||_____||_____||

I didn't do this

Very
easy

Very
difficult

How easy was it to **remember what you had read or marked up** on other documents as you worked?

|_____||_____||_____||_____||_____||_____||

I didn't do this

Very
easy

Very
difficult

How easy was it to **write on a separate document** to note-take or fill in answers as you did the task?

|_____||_____||_____||_____||_____||_____||

I didn't do this

Very
easy

Very
difficult

Please rank the following conditions:

Which condition did you find easiest in terms of **planning and arranging the documents** to do the task? (Assign a 1, 2 or 3 to each condition where 1 is the best (or easiest) and 3 is the worst (or most difficult):

- ___ Paper
- ___ Single Surface
- ___ Separate Tablets

Which condition did you find easiest **to find** the documents you needed to do the task? (Assign a 1, 2 or 3 to each condition where 1 is the best (or easiest) and 3 is the worst (or most difficult):

- ___ Paper
- ___ Single Surface
- ___ Separate Tablets

Which condition did you find easiest **to mark up** the documents as you did the task? (Assign a 1, 2 or 3 to each condition where 1 is the best (or easiest) and 3 is the worst (or most difficult):

- ___ Paper
- ___ Single Surface
- ___ Separate Tablets

Which condition did you find easiest **to write as you were reading** during the task? (Assign a 1, 2 or 3 to each condition where 1 is the best (or easiest) and 3 is the worst (or most difficult):

- ___ Paper
- ___ Single Surface
- ___ Separate Tablets

Which condition did you find easiest **to read and compare across documents** during the task? (Assign a 1, 2 or 3 to each condition where 1 is the best (or easiest) and 3 is the worst (or most difficult):

- ___ Paper
- ___ Single Surface
- ___ Separate Tablets

Which condition did you find most **comfortable to use** during the task? (Assign a 1, 2 or 3 to each condition where 1 is the best (or easiest) and 3 is the worst (or most difficult):

- ___ Paper
- ___ Single Surface
- ___ Separate Tablets

On which condition do you think you were able to complete the task the fastest? (Assign a 1, 2 or 3 to each condition where 1 is fastest and 3 is the slowest:

- ___ Paper

- Single Surface
- Separate Tablets

Overall, which condition **did you prefer** for doing these tasks?

- Paper
- Single Surface
- Separate Tablets

Interview Questions

Are you left handed or right handed? Did you find anything awkward or difficult to reach?

What sorts of reading do you normally do?

While these tasks aren't exactly the same as your "everyday" tasks are they similar to the sort of reading tasks you've had to do? For instance, do you have to do search on the computer and then write? Comparing things side by side?

How often do you find yourself printing things out? Are there things that you tend to print out? Emails? Webpages? Course readings?

What types of reading do you do on electronic devices (e.g. laptop, phone, e-book reader)?

About the tasks themselves...

Were there any that you found particularly difficult to perform?

What was your overall strategy for performing the task? Did it change for the different devices? Why?

I was wondering why you

- Arranged things the way you did
- Moved this over here
- Chose to do X
- Didn't do something in one condition but did it in another condition

Did any of the tasks or devices confuse you?

Appendix D: Deployment Final Interview Questions

These interview questions were used during the final debriefing interview for the deployment in Chapter 8. The questions listed here are the common set of questions we asked all users. The first part of the interview varied between users as it focused on the specific activities each user engaged in during the deployment.

Experiment Debrief Interview

Activities Specific to the User

Discuss all the different tasks you completed on the devices over the past 2 weeks

For each task [user dependent (e.g., Marking up documents while writing, teaching, grading papers, reading journals, conducting research, attending conferences, etc.)]:

- Describe the process, what you were doing, where you were doing this (desk, sofa, coffee shop)
- How many Slates were in use? What were you doing with each one?
- How long did it last?

Features of the System

For each:

- Useful? Not useful? If not, when would it be useful?
- Were there any problems?
- IF you were to improve it, what would you change to make it better?

Multiple Screen Use

Links

Writing

Text Extraction

Groups and Stacks

Thumbnail Overview

Using with PC

PC Remote Control

Pick and Drop (Stitching)

Overall Impressions

What features would have made the system more usable?

How could the system have been more useful for you?

In what situations would you find additional screens useful?

What do you see as the biggest weakness in these devices?

Did you have any expectations coming into the study? What aspects were fulfilled? What aspects should I improve?

Bibliography

- Adler, A., Gujar, A., Harrison, B.L., O'Hara, K., and Sellen, A. 1998. A diary study of work-related reading: design implications for digital reading devices. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, 241–248. CHI '98. ACM Press/Addison-Wesley Publishing Co., Los Angeles, California, United States.
- Adler, M.J. and van Doren, C. 1972. *How to Read a Book*. Simon and Schuster, New York, NY.
- Agarawala, A. and Balakrishnan, R. 2006. Keepin' it real: pushing the desktop metaphor with physics, piles and the pen. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, 1283–1292. CHI '06. ACM, New York, NY, USA.
- Alexander, J., Cockburn, A., Fitchett, S., Gutwin, C., and Greenberg, S. 2009. Revisiting read wear. In *Proceedings of the 27th international conference on Human factors in computing systems*, 1665. CHI '09. Boston, MA, USA.
- Back, M., Cohen, J., Gold, R., Harrison, S., and Minneman, S. 2001. Listen reader: an electronically augmented paper-based book. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, 23–29. CHI '01. ACM, New York, NY, USA.
- Bartlett, J.F. 2000. Rock 'n' Scroll Is Here to Stay. *Computer Graphics and Applications, IEEE* 20(3) (May), 40–45.
- Bau, O., Poupyrev, I., Israr, A., and Harrison, C. 2010. TeslaTouch: electrovibration for touch surfaces. In *Proceedings of the 23rd annual ACM symposium on User interface software and technology*, 283–292. UIST '10. ACM, New York, NY, USA.
- Behler, A. 2009. E-readers in Action. *American Libraries*, October.
- Bier, E.A. and Stone, M.C. 1986. Snap-dragging. In *Proceedings of the 13th annual conference on Computer graphics and interactive techniques*, 233–240. SIGGRAPH '86. ACM, New York, NY, USA.
- Block, R. iriver's E-BOOK reader. *Engadget*.
- Bondarenko, O. and Janssen, R. 2005. Documents at Hand: Learning from Paper to Improve Digital Technologies. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, 121–130. CHI '05. ACM, Portland, Oregon, USA.
- Broxton, M., Lifton, J., and Paradiso, J.A. 2006. Localization on the pushpin computing sensor network using spectral graph drawing and mesh relaxation. *SIGMOBILE Mob. Comput. Commun. Rev.* 10(1) (January), 1–12.
- de Bruijn, O. and Spence, R. 2000. Rapid serial visual presentation: a space-time trade-off in information presentation. In *Proceedings of the working conference on Advanced visual interfaces*, 189–192. AVI '00. ACM, New York, NY, USA.

- Buxton, W. and Myers, B. 1986. A study in two-handed input. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, 321–326. CHI '86. April.
- Buxton, W.A.S. 1986. Chunking and Phrasing and the Design of Human-Computer Dialogues. In *Proceedings of the IFIP World Computer Congress*, 475–480. Dublin, Ireland.
- Cao, X., Wilson, A., Balakrishnan, R., Hinckley, K., and Hudson, S. 2008. ShapeTouch: Leveraging contact shape on interactive surfaces. In *3rd IEEE International Workshop on Horizontal Interactive Human Computer Systems*, 129–136. TABLETOP 2008. IEEE, Amsterdam, October.
- Card, S.K., Hong, L., Mackinlay, J.D., and Chi, E.H. 2004. 3Book: a 3D electronic smart book. In *Proceedings of the working conference on Advanced visual interfaces*, 303–307. AVI '04. ACM, Gallipoli, Italy.
- Chalmers, M. and Galani, A. 2004. Seamful interweaving: heterogeneity in the theory and design of interactive systems. In *Proceedings of the 5th conference on Designing interactive systems: processes, practices, methods, and techniques*, 243–252. DIS '04. ACM, New York, NY, USA.
- Chen, N., Guimbretière, F., Dixon, M., Lewis, C., and Agrawala, M. 2008. Navigation techniques for dual-display e-book readers. In *Proceedings of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, 1779–1788. CHI '08. ACM, Florence, Italy.
- Clanchy, M.T. 2007. Parchment and Paper: Manuscript Culture 1100-1500. In *A Companion to The History of the Book*, S. Eliot and J. Rose, Eds. , 194–206. Wiley-Blackwell, Hoboken, NJ.
- Cockburn, A., Gutwin, C., and Alexander, J. 2006. Faster document navigation with space-filling thumbnails. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, 1–10. CHI '06. ACM, Montréal, Québec, Canada.
- Conroy, K., Levin, D., and Guimbretière, F. 2004. *ProofRite: A Paper-Augmented Word Processor*. Tech Report. HCIL-2004-22, CS-TR-4652. University of Maryland.
- Dietz, P. and Leigh, D. 2001. DiamondTouch. In *Proceedings of the 14th annual ACM symposium on User interface software and technology*, 219–226. UIST '01. ACM, Orlando, Florida.
- Dillon, A. 1992. Reading from paper versus screens: a critical review of the empirical literature. *Ergonomics* 35(10), 1297–1326.
- Dixon, M., Guimbretière, F., and Chen, N. 2008. Optimal parameters for efficient crossing-based dialog boxes. In *Proceedings of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, 1623–1632. CHI '08. ACM, Florence, Italy.
- Dourish, P., Edwards, W.K., LaMarca, A., and Salisbury, M. 1999. Presto: an experimental architecture for fluid interactive document spaces. *ACM Trans. Comput.-Hum. Interact.* 6(2), 133–161.
- Eisenstein, E. 1979. *The printing press as an agent of change : communications and cultural transformations in early modern Europe*. Cambridge University Press, Cambridge, England.
- Fails, J.A., Druin, A., and Guha, M.L. 2011. Content splitting & space sharing: collaboratively reading & sharing children's stories on mobile devices. In *Proceedings of the 13th International*

- Conference on Human Computer Interaction with Mobile Devices and Services*, 361–370. MobileHCI '11. ACM, New York, NY, USA.
- Feng, G., Gormish, M.J., Chemishkian, S., Gudan, K.F., and Barrus, J. 2008. High Performance Pen Tracking on Electronic Paper Displays. *Proceedings of IDW* 15(1), 245–248.
- Fielding, R.T. 2000. Architectural styles and the design of network-based software architectures. Doctoral Thesis, University of California, Irvine.
- Fishkin, K.P., Gujar, A., Harrison, B.L., Moran, T.P., and Want, R. 2000. Embodied user interfaces for really direct manipulation. *Commun. ACM* 43(9), 74–80.
- Geller, T. 2011. The promise of flexible displays. *Commun. ACM* 54(6) (June), 16–18.
- Greenberg, S. and Marwood, D. 1994. Real time groupware as a distributed system. In *Proceedings of the 1994 ACM conference on Computer supported cooperative work*, 207–217. CSCW '94. Chapel Hill, North Carolina, United States.
- Grudin, J. 2001. Partitioning digital worlds: focal and peripheral awareness in multiple monitor use. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, 458–465. CHI '01. ACM, Seattle, Washington, United States.
- Guiard, Y. 1987. Asymmetric division of labor in human skilled bimanual action: the kinematic chain as a model. *Journal of Motor Behavior* 19(4) (December), 486–517.
- Guimbretière, F. 2003. Paper augmented digital documents. In *Proceedings of the 16th annual ACM symposium on User interface software and technology*, 51–60. UIST '03. ACM, Vancouver, Canada.
- Han, J.Y. 2005. Low-cost multi-touch sensing through frustrated total internal reflection. In *Proceedings of the 18th annual ACM symposium on User interface software and technology*, 115–118. UIST '05. Seattle, WA, USA.
- Hardock, G., Kurtenbach, G., and Buxton, W. 1993. A marking based interface for collaborative writing. In *Proceedings of the 6th annual ACM symposium on User interface software and technology*, 259–266. UIST '93. ACM, New York, NY, USA.
- Harrison, B.L., Fishkin, K.P., Gujar, A., Mochon, C., and Want, R. 1998. Squeeze me, hold me, tilt me! An exploration of manipulative user interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, 17–24. CHI '98. ACM Press/Addison-Wesley Publishing Co., Los Angeles, California, United States.
- Hayes, R.A. and Feenstra, B.J. 2003. Video-speed electronic paper based on electrowetting. *Nature* 425(6956) (September 25), 383–385.
- Hill, W.C., Hollan, J.D., Wroblewski, D., and McCandless, T. 1992. Edit wear and read wear. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, 3–9. CHI '92. ACM, New York, NY, USA.
- Hinckley, K. 2003. Synchronous Gestures for Multiple Persons and Computers. In *Proceedings of the 16th annual ACM symposium on User interface software and technology*, 149–158. UIST '03. ACM, Vancouver, Canada.

- Hinckley, K., Baudisch, P., Ramos, G., and Guimbretière, F. 2005a. Design and analysis of delimiters for selection-action pen gesture phrases in scriboli. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, 451–460. CHI '05. ACM, Portland, Oregon, USA.
- Hinckley, K., Dixon, M., Sarin, R., Guimbretière, F., and Balakrishnan, R. 2009. Codex: A Dual Screen Tablet Computer. In *Proceedings of the 27th international conference on Human factors in computing systems*. CHI '09. ACM, Boston, Massachusetts, USA.
- Hinckley, K., Pierce, J., Horvitz, E., and Sinclair, M. 2005b. Foreground and background interaction with sensor-enhanced mobile devices. *ACM Trans. Comput.-Hum. Interact.* 12(1), 31–52.
- Hinckley, K., Ramos, G., Guimbretière, F., Baudisch, P., and Smith, M. 2004. Stitching: pen gestures that span multiple displays. In *Proceedings of the working conference on Advanced visual interfaces*, 23–31. AVI '04. ACM, Gallipoli, Italy.
- Hinckley, K., Zhao, S., Sarin, R., et al. 2007. InkSeine: In Situ search for active note taking. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, 251–260. CHI '07. ACM, San Jose, California, USA.
- Hoffmann, R., Baudisch, P., and Weld, D.S. 2008. Evaluating visual cues for window switching on large screens. In *Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, 929–938. CHI '08. ACM, Florence, Italy.
- Holman, D., Vertegaal, R., Altosaar, M., Troje, N., and Johns, D. 2005. Paper windows: interaction techniques for digital paper. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, 591–599. CHI '05. ACM, Portland, Oregon, USA.
- Holzinger, A., Baerenthaler, M., Pammer, W., Katz, H., Bjelic-Radisic, V., and Ziefle, M. 2011. Investigating paper vs. screen in real-life hospital workflows: Performance contradicts perceived superiority of paper in the user experience. *Int. J. Hum.-Comput. Stud.* 69(9) (August), 563–570.
- Igarashi, T. and Hinckley, K. 2000. Speed-dependent automatic zooming for browsing large documents. In *Proceedings of the 13th annual ACM symposium on User interface software and technology*, 139–148. UIST '00. ACM, New York, NY, USA.
- Johnson, B.S. 1969. *The Unfortunates*. Panther, London,.
- Kane, S.K., Avrahami, D., Wobbrock, J.O., et al. 2009. Bonfire: a nomadic system for hybrid laptop-tabletop interaction. In *Proceedings of the 22nd annual ACM symposium on User interface software and technology*, 129–138. UIST '09. ACM, New York, NY, USA.
- Karlson, A., Bederson, B., and Contreras-Vidal, J. 2008. Understanding One-Handed Use of Mobile Devices. In *Handbook of Research on User Interface Design and Evaluation for Mobile Technology*, 86–101. Information Science Reference.
- Kay, A. and Goldberg, A. 1977. Personal Dynamic Media. *Computer* 10(3), 31–41.

- Klemmer, S.R., Graham, J., Wolff, G.J., and Landay, J.A. 2003. Books with voices: paper transcripts as a physical interface to oral histories. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, 89–96. CHI '03. ACM, New York, NY, USA.
- Lee, J.C., Forlizzi, J., and Hudson, S.E. 2002. The kinetic typography engine: an extensible system for animating expressive text. In *Proceedings of the 15th annual ACM symposium on User interface software and technology*, 81–90. UIST '02. ACM, New York, NY, USA.
- Leitner, J., Powell, J., Brandl, P., et al. 2009. Flux. In *Proceedings of the 27th international conference extended abstracts on Human factors in computing systems - CHI EA '09*, 3211. Boston, MA, USA.
- Liao, C., Guimbretière, F., Hinckley, K., and Hollan, J. 2008. Papiercraft: A gesture-based command system for interactive paper. *ACM Trans. Comput.-Hum. Interact.* 14(4), 1–27.
- Luff, P. and Heath, C. 1998. Mobility in collaboration. In *Proceedings of the 1998 ACM conference on Computer supported cooperative work*, 305–314. CSCW '98. ACM, New York, NY, USA.
- Luff, P., Heath, C., Norrie, M., Signer, B., and Herdman, P. 2004. Only touching the surface: creating affinities between digital content and paper. In *Proceedings of the 2004 ACM conference on Computer supported cooperative work*, 523–532. CSCW '04. ACM, New York, NY, USA.
- Mackay, W.E., Pagani, D.S., Faber, L., et al. 1995. Ariel: augmenting paper engineering drawings. In *Conference companion on Human factors in computing systems*, 421–422. CHI '95. ACM, New York, NY, USA.
- Mackay, W.E., Pothier, G., Letondal, C., Bøegh, K., and Sørensen, H.E. 2002. The missing link: augmenting biology laboratory notebooks. In *Proceedings of the 15th annual ACM symposium on User interface software and technology*, 41–50. UIST '02. ACM, Paris, France.
- Malone, T.W. 1983. How do people organize their desks?: Implications for the design of office information systems. *ACM Trans. Inf. Syst.* 1(1), 99–112.
- Mangen, A. 2008. Hypertext fiction reading: haptics and immersion. *Journal of Research in Reading* 31(4) (November), 404–419.
- Manguel, A. 1996. *A History of Reading*. Viking, New York.
- Marchionini, G. 2000. Evaluating Digital Libraries: A Longitudinal and Multifaceted View. *Library Trends* 49, 304–333.
- Marshall, C. 1997. Annotation: from paper books to the digital library. In *Proceedings of the second ACM international conference on Digital libraries*, 131–140. DL '97. ACM Press.
- Marshall, C. 2005. Reading and Interactivity in the Digital Library: Creating an Experience that Transcends Paper. In *Digital Library Development: The View from Kanazawa*, D. Marcum and G. George, Eds. Libraries Unlimited, Westport, CT.
- Marshall, C.C. and Bly, S. 2005. Turning the page on navigation. In *Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries*, 225–234. JCDL '05. ACM, Denver, CO, USA.

- Marshall, C.C., Price, M.N., Golovchinsky, G., and Schilit, B.N. 1999. Introducing a digital library reading appliance into a reading group. In *Proceedings of the fourth ACM conference on Digital libraries*, 77–84. DL '99. ACM, Berkeley, California, United States.
- Marshall, C.C., Price, M.N., Golovchinsky, G., and Schilit, B.N. 2001. Designing e-books for legal research. In *Proceedings of the 1st ACM/IEEE-CS joint conference on Digital libraries*, 41–48. JCDL '01. ACM, Roanoke, Virginia, United States.
- Marshall, C.C. and Ruotolo, C. 2002. Reading-in-the-small: a study of reading on small form factor devices. In *Proceedings of the 2nd ACM/IEEE-CS joint conference on Digital libraries*, 56–64. JCDL '02. ACM, Portland, Oregon, USA.
- Marshall, C.C. and Shipman, F.M. 1997. Spatial hypertext and the practice of information triage. In *Proceedings of the eighth ACM conference on Hypertext*, 124–133. HYPERTEXT '97. ACM, Southampton, United Kingdom.
- Martin, J.A. 1998. EveryBook unveils first true ebook. *CNN Interactive*. December 18.
- Merrill, D., Kalanithi, J., and Maes, P. 2007. Siftables: towards sensor network user interfaces. In *Proceedings of the 1st international conference on Tangible and embedded interaction*, 75–78. TEI '07. ACM, New York, NY, USA.
- Miller, R.C. and Myers, B.A. 1999. Synchronizing clipboards of multiple computers. In *Proceedings of the 12th annual ACM symposium on User interface software and technology*, 65–66. UIST '99. ACM, Asheville, North Carolina, United States.
- Mills, C.B. and Weldon, L.J. 1987. Reading text from computer screens. *ACM Comput. Surv.* 19(4) (December), 329–357.
- Morris, M.R., Brush, A.J.B., and Meyers, B.R. 2007. Reading Revisited: Evaluating the Usability of Digital Display Surfaces for Active Reading Tasks. In *Proceedings of the Second Annual IEEE International Workshop on Horizontal Interactive Human-Computer Systems*, 79–86. TABLETOP '07.
- Muter, P. and Maurutto, P. 1991. Reading and skimming from computer screens and books: the paperless office revisited? *Behaviour & Information Technology* 10(4), 257–266.
- Myers, B., Hudson, S.E., and Pausch, R. 2000. Past, present, and future of user interface software tools. *ACM Trans. Comput.-Hum. Interact.* 7(1) (March), 3–28.
- O'Hara, K. and Sellen, A. 1997. A comparison of reading paper and on-line documents. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, 335–342. CHI '97. ACM, Atlanta, Georgia, United States.
- O'Hara, K., Sellen, A., and Bentley, R. 1999. Supporting memory for spatial location while reading from small displays. In *CHI '99 extended abstracts on Human factors in computing systems*, 220–221. CHI EA '99. ACM, Pittsburgh, Pennsylvania.
- O'Hara, K., Smith, F., Newman, W., and Sellen, A. 1998. Student readers' use of library documents: implications for library technologies. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, 233–240. CHI '98. ACM Press/Addison-Wesley Publishing Co., Los Angeles, California, United States.

- O'Hara, K., Taylor, A.S., Newman, W.M., and Sellen, A. 2002. Understanding the materiality of writing from multiple sources. *Int. J. Hum.-Comput. Stud.* 56(3), 269–305.
- Oulasvirta, A. and Sumari, L. 2007. Mobile kits and laptop trays: managing multiple devices in mobile information work. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, 1127–1136. CHI '07. ACM, San Jose, California, USA.
- Pearson, J. and Buchanan, G. 2011. CloudBooks: An Infrastructure for Reading on Multiple Devices. In *Proceedings of the 15th international conference on Theory and practice of digital libraries: research and advanced technology for digital libraries*, 488–492. TPD L '11.
- Pearson, J., Buchanan, G., and Thimbleby, H. 2009. Improving annotations in digital documents. In *Proceedings of the 13th European conference on Research and advanced technology for digital libraries*, 429–432. ECDL '09. Springer-Verlag, Berlin, Heidelberg.
- Peng, C., Shen, G., Zhang, Y., Li, Y., and Tan, K. 2007. BeepBeep: a high accuracy acoustic ranging system using COTS mobile devices. In *Proceedings of the 5th international conference on Embedded networked sensor systems*, 1–14. SenSys '07. ACM, Sydney, Australia.
- Pierce, J.S. and Nichols, J. 2008. An infrastructure for extending applications' user experiences across multiple personal devices. In *Proceedings of the 21st annual ACM symposium on User interface software and technology*, 101–110. UIST '08. ACM, Monterey, CA, USA.
- Price, M.N., Golovchinsky, G., and Schilit, B.N. 1998. Linking by inking. In *Proceedings of the ninth ACM conference on Hypertext and hypermedia*, 30–39. HYPERTEXT '98. Pittsburgh, Pennsylvania, United States.
- Pugh, A.K. 1978. *Silent reading: an introduction to its study and teaching*. Heinemann Educational.
- Qualcomm. 2011. IMOD Technology Overview. Whitepaper. June 6.
- Reenskaug, T. 1979. *Models - Views - Controllers*. Xerox PARC Technical Note December 1979. Xerox PARC.
- Reetz, A., Gutwin, C., Stach, T., Nacenta, M., and Subramanian, S. 2006. Superflick: a natural and efficient technique for long-distance object placement on digital tables. In *Proceedings of Graphics Interface 2006*, 163–170. GI '06. Canadian Information Processing Society, Quebec, Canada.
- Rekimoto, J. 1997. Pick-and-drop: a direct manipulation technique for multiple computer environments. In *Proceedings of the 10th annual ACM symposium on User interface software and technology*, 31–39. UIST '97. ACM, Banff, Alberta, Canada.
- Rekimoto, J. 2004. SyncTap: synchronous user operation for spontaneous network connection. *Personal Ubiquitous Comput.* 8(2), 126–134.
- Rekimoto, J. and Saitoh, M. 1999. Augmented surfaces: a spatially continuous work space for hybrid computing environments. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, 378–385. CHI '99. ACM, Pittsburgh, Pennsylvania, United States.
- Rheinfrank, J. 1995. A conversation with Don Norman. *interactions* 2(2), 47–55.

- Rhodes, B. 2009. Fast Updating of Bistable Displays For Pen Tracking and Page Flipping. *Proceedings of IDW* 16(3).
- Saint-Andre, P. 2004a. *Extensible Messaging and Presence Protocol (XMPP): Core*. RFC 3920.
- Saint-Andre, P. 2004b. *Extensible Messaging and Presence Protocol (XMPP): Instant Messaging and Presence*. RFC 3921. institution.
- Schilit, B.N., Price, M.N., and Golovchinsky, G. 1998. Digital library information appliances. In *Proceedings of the third ACM conference on Digital libraries*, 217–226. DL '98. ACM, Pittsburgh, Pennsylvania, United States.
- Schilit, B.N., Price, M.N., Golovchinsky, G., Tanaka, K., and Marshall, C.C. 1999. As We May Read: The Reading Appliance Revolution. *Computer* 32(1), 65–73.
- Schilit, B.N. and Sengupta, U. 2004. Device ensembles. *Computer* 37(12) (December), 56–64.
- Sellen, A.J. and Harper, R.H.R. 2003. *The Myth of the Paperless Office*. MIT Press.
- Sohn, T., Li, K.A., Griswold, W.G., and Hollan, J.D. 2008. A diary study of mobile information needs. In *Proceedings of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, 433–442. CHI '08. ACM, New York, NY, USA.
- Song, H., Grossman, T., Fitzmaurice, G., et al. 2009. PenLight: combining a mobile projector and a digital pen for dynamic visual overlay. In *Proceedings of the 27th international conference on Human factors in computing systems*, 143–152. CHI '09. ACM, New York, NY, USA.
- Song, H., Guimbretière, F., Grossman, T., and Fitzmaurice, G. 2010. MouseLight: bimanual interactions on digital paper using a pen and a spatially-aware mobile projector. In *Proceedings of the 28th international conference on Human factors in computing systems*, 2451–2460. CHI '10. ACM, New York, NY, USA.
- Sporkin, A. 2011. Popularity of Books in Digital Platforms Continues to Grow, According to AAP Publishers. April 14.
- Stifelman, L., Arons, B., and Schmandt, C. 2001. The audio notebook: paper and pen interaction with structured speech. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, 182–189. CHI '01. ACM, New York, NY, USA.
- Sun, L. and Guimbretière, F. 2005. Flipper: a new method of digital document navigation. In *CHI '05 extended abstracts on Human factors in computing systems*, 2001–2004. CHI EA '05. ACM, New York, NY, USA.
- Tabard, A., Mackay, W.E., and Eastmond, E. 2008. From individual to collaborative: the evolution of prism, a hybrid laboratory notebook. In *Proceedings of the 2008 ACM conference on Computer supported cooperative work*, 569–578. CSCW '08. ACM, San Diego, CA, USA.
- Tandler, P., Prante, T., Müller-Tomfelde, C., Streitz, N., and Steinmetz, R. 2001. Connectables: dynamic coupling of displays for the flexible creation of shared workspaces. In *Proceedings of the 14th annual ACM symposium on User interface software and technology*, 11–20. UIST '01. ACM, Orlando, Florida.

- Tashman, C.S. and Edwards, W.K. 2011a. Active reading and its discontents: the situations, problems and ideas of readers. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, 2927–2936. CHI '11. ACM, Vancouver, BC, Canada.
- Tashman, C.S. and Edwards, W.K. 2011b. LiquidText: a flexible, multitouch environment to support active reading. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, 3285–3294. CHI '11. ACM, Vancouver, BC, Canada.
- Terrenghi, L., Kirk, D., Sellen, A., and Izadi, S. 2007. Affordances for manipulation of physical versus digital media on interactive surfaces. In *Proceedings of CHI 2007*, 1157–1166. ACM, San Jose, California, USA.
- Thayer, A., Lee, C.P., Hwang, L.H., Sales, H., Sen, P., and Dalal, N. 2011. The imposition and superimposition of digital reading technology: the academic potential of e-readers. In *Proceedings of CHI '11*, 2917–2926. CHI '11. ACM, Vancouver, BC, Canada.
- The Economist. 2008. The paperless office: On its way, at last. *The Economist*, October 9.
- The Trustees of Princeton University. 2010. Princeton e-reader pilot, Final report. October 27.
- Tridgell, A. 1999. Efficient Algorithms for Sorting and Synchronization. Doctoral Thesis, Australian National University.
- Weibel, N., Ispas, A., Signer, B., and Norrie, M.C. 2008. Paperproof: a paper-digital proof-editing system. In *CHI '08 extended abstracts on Human factors in computing systems*, 2349–2354. CHI EA '08. ACM, New York, NY, USA.
- Weiser, M. 1994. Creating the invisible interface: (invited talk). In *Proceedings of the 7th annual ACM symposium on User interface software and technology*, 1–. UIST '94. ACM, New York, NY, USA.
- Weiser, M. 1995. The computer for the 21st century. In *Human-computer interaction: toward the year 2000*, 933–940. Morgan Kaufmann Publishers Inc.
- Weiss, M., Jennings, R., Khoshabeh, R., et al. 2009. SLAP widgets: bridging the gap between virtual and physical controls on tabletops. In *Proceedings of the 27th international conference on Human factors in computing systems*, 3229–3234. CHI '09. ACM, New York, NY, USA.
- Wellner, P. 1993. Interacting with paper on the DigitalDesk. *Commun. ACM* 36(7), 87–96.
- Williamson, J., Murray-Smith, R., and Hughes, S. 2007. Shoogle: excitatory multimodal interaction on mobile devices. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, 121–124. CHI '07. ACM, New York, NY, USA.
- Wilson, A.D. 2005. PlayAnywhere: a compact interactive tabletop projection-vision system. In *Proceedings of the 18th annual ACM symposium on User interface software and technology*, 83–92. UIST '05. Seattle, WA, USA.
- Wilson, A.D. and Benko, H. 2010. Combining multiple depth cameras and projectors for interactions on, above and between surfaces. In *Proceedings of the 23rd annual ACM symposium on User interface software and technology*, 273–282. UIST '10. ACM, New York, NY, USA.

Wilson, A.D., Izadi, S., Hilliges, O., Garcia-Mendoza, A., and Kirk, D. 2008. Bringing physics to the surface. In *Proceedings of the 21st annual ACM symposium on User interface software and technology*, 67–76. UIST '08. ACM, New York, NY, USA.

Wilson, R. and Landoni, M. 2003. Evaluating the usability of portable electronic books. In *Proceedings of the 2003 ACM symposium on Applied computing*, 564–568. SAC '03. ACM, Melbourne, Florida.

Wobbrock, J.O., Wilson, A.D., and Li, Y. 2007. Gestures without libraries, toolkits or training: a \$1 recognizer for user interface prototypes. In *Proceedings of the 20th annual ACM symposium on User interface software and technology*, 159–168. UIST '07. ACM, New York, NY, USA.

Wolin, A., Eoff, B., and Hammond, T. 2008. ShortStraw: A Simple and Effective Corner Finder for Polylines. In *EUROGRAPHICS Workshop on Sketch-Based Interfaces and Modeling*, 33–40.

Yeh, R., Liao, C., Klemmer, S., et al. 2006. ButterflyNet: a mobile capture and access system for field biology research. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, 571–580. CHI '06. ACM, New York, NY, USA.

Young, J.R. 2009. 6 Lessons One Campus Learned About E-Textbooks. *The Chronicle of Higher Education*, June 4, sec. Technology.

Zhao, S., Dragicevic, P., Chignell, M., Balakrishnan, R., and Baudisch, P. 2007. Earpod: eyes-free menu selection using touch input and reactive audio feedback. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, 1395–1404. CHI '07. ACM, New York, NY, USA.