



Swansea University  
Prifysgol Abertawe



## Swansea University E-Theses

---

# Investigating lightweight interaction for active reading in digital documents.

**Pearson, Jennifer Sarah**

### How to cite:

---

Pearson, Jennifer Sarah (2012) *Investigating lightweight interaction for active reading in digital documents..* thesis, Swansea University.

<http://cronfa.swan.ac.uk/Record/cronfa42856>

### Use policy:

---

This item is brought to you by Swansea University. Any person downloading material is agreeing to abide by the terms of the repository licence: copies of full text items may be used or reproduced in any format or medium, without prior permission for personal research or study, educational or non-commercial purposes only. The copyright for any work remains with the original author unless otherwise specified. The full-text must not be sold in any format or medium without the formal permission of the copyright holder. Permission for multiple reproductions should be obtained from the original author.

Authors are personally responsible for adhering to copyright and publisher restrictions when uploading content to the repository.


Please link to the metadata record in the Swansea University repository, Cronfa (link given in the citation reference above.)

<http://www.swansea.ac.uk/library/researchsupport/ris-support/>

# Investigating Lightweight Interaction for Active Reading in Digital Documents

Jennifer Sarah Pearson

2012



Submitted to Swansea University in fulfilment of the  
requirements for the Degree of Doctor of Philosophy



**Prifysgol Abertawe**  
**Swansea University**

.....

ProQuest Number: 10821246

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



ProQuest 10821246

Published by ProQuest LLC (2018). Copyright of the Dissertation is held by the Author.

All rights reserved.

This work is protected against unauthorized copying under Title 17, United States Code  
Microform Edition © ProQuest LLC.

ProQuest LLC.  
789 East Eisenhower Parkway  
P.O. Box 1346  
Ann Arbor, MI 48106 – 1346



# Declaration

This work has not been previously accepted in substance for any degree and is not being concurrently submitted in candidature for any degree.

Signed \_\_\_\_\_ (candidate)

Date 20/6/2012

# Statement 1

This Thesis is the result of my own investigations, except where otherwise stated. Other sources are acknowledged by footnotes giving explicit references. A bibliography is appended.

Signed \_\_\_\_\_ (candidate)

Date 20/6/2012

# Statement 2

I hereby give my consent for my Thesis, if accepted, to be available for photocopying and for inter-library loan, and for the title and summary to be made available to outside organisations.

Signed \_\_\_\_\_ (candidate)

Date 20/6/2012

*i Mamgu  
for Granny*

# Abstract

Reading is a complex human activity evolved, and co-evolved, with technology over thousands of years. Mass printing in the fifteenth century firmly established what we know as the modern book, with its physical format of covers and paper pages, and now-standard features such as page numbers. Today, electronic documents are enabling paperless reading supported by eReading technologies such as Kindles and Nooks, yet a high proportion of users still opt to print on paper before reading. This persistent ‘print to read’ mentality is one sign of the shortcomings of digital documents.

The physical properties of paper (for example, it is light, thin and flexible) contribute to the ease with which physical documents are manipulated and pose a completely different set of affordances to their digital ‘equivalents’. Paper can be folded, ripped or scribbled on almost subconsciously – activities that require significant cognitive attention in their digital form. The almost subliminal interaction that comes from years of learned behaviour with paper has been described as ‘lightweight interaction’ which is achieved when a person actively reads an article in a way that is so unselfconscious that they are not apt to remember their actions later.

This Thesis investigates the advantages of paper, how the affordances of paper can be reified in digital form, and what forms best support lightweight interaction for reading. It explores the reasoning behind reader behaviour and introduces several interfaces that implement the lightweight properties found on paper to produce new ways of reading and marking up digital texts. Each of the tools described have been evaluated on an appropriate population of users.

As a starting point, the concept of placeholders are investigated as this is a good example of an area where the digital equivalent of a common paper practice is lacking in usability. From this investigation, several notable possibilities for lightweight attributes were identified which then leads systematically on to the second investigation into annotation – another common activity that has been poorly migrated from paper to digital. After studying both paper and digital based annotation, the Thesis then moves on to note-taking, focusing on a desk-based implementation that makes use of amalgamated tools – just like paper. Finally, to investigate the possibility of solely digital lightweight properties, the concept of back-of-book indexing is investigated – an example of a task that is currently difficult to perform both physically and digitally.

After gaining ideas from each of these these implemented interactions and user studies, a set of rules for designing lightweight interactions on digital interactive texts is described. This list of lightweight attributes can then be used as guidelines to improve the usability of future digital reading implementations.

# Acknowledgements

I would like to express my thanks to many individuals who have supported me greatly throughout the course of my studies. Without you I am sure this Thesis would never have materialised.

First I would like to thank George Buchanan my mentor and friend for all his help and support throughout every stage of my PhD. I would not have got this far without you. I would also like to thank Harold Thimbleby and Matt Jones for their everyday help and advice, as well as a significant amount of patience through many stressful occasions. Vicky, I owe you thanks too – I think I'd probably still be waiting for a viva date if it weren't for you.

Thank you to my parents for your unconditional support while I spend yet another three years in education. I promise I am done now Mum :-)) Thank you Dad for reading over my Thesis - it's nice you are still willing to help me with my homework after all these years. I would like to thank Uncle David and Auntie Susie for looking after me so well and for so long when I decide I need a break from Swansea, and to Laura, Hannah and Thomas for always making me smile. Thanks to my boyfriend, Tom for all he has had to endure throughout the duration, and to Granny - my biggest fan.

In addition, I would also like to thank my friends in the FIT lab for their friendship and advice during my studies. Simon and Patrick, you are awesome. The Tetris and coffee tournaments were a very welcomed and pleasant distraction from the joys of Thesis writing. For the sake of this argument - Ben, you are also an honorary member of the FIT lab, thanks for the proof reading and for helping us get the top score on the itBox. We rock.

Finally I would like to express my gratitude to Microsoft Research Cambridge, particularly Richard Harper, and Swansea University's Computer Science Department for the financial support they have offered me throughout the duration of my studies.

Thanks Everyone!

Jen ☺



# Preface

My Mum always said that if I wasn't sure what I wanted to do, that I should carry on in education until I found something that I really enjoyed enough to make a career out of. I'm not sure if she expected me to get this far before I finally made up my mind, or if she imagined my decision would be rooted in academia. Either way, I value the advice she has given me which has eventually led me to producing this Thesis.

Although I am not what you would generally consider to be a big lover of books, I stand by my decision to dedicate three years of my life to the topic of digital reading. Despite my preference of video games and art over books and magazines, I still relish in the sporadic enjoyment of a good read – particularly if it involves learning new material. As a child I remember being mesmerised by an old book I found at my grandparents' house, named simply *How Stuff Works* - the content of which I could soon recite by memory. As the years went on I found myself reading more and more, not for fun per se, but to learn, an activity that commonly required the use of other materials and what I now know to be the in-depth task of active reading.

Now, as the concept of electronic books are becoming evermore fluidly integrated into our daily lives, we find ourselves reading more frequently on-screen. Whether it be to undertake active reading on a PC workstation or reading a novel on an eReader, the act of reading is shifting more and more towards the digital, making it an opportune time to investigate the on-screen reading process.

In the three years I have been working on this topic, I have implemented several lightweight active reading tools which I have consequently evaluated and published in peer reviewed conferences in both the human-computer-interaction and digital library domains. Doing so has allowed me to attend several international destinations and meet many interesting individuals.

I have thoroughly enjoyed the time I have spent researching this area, and hope that you have as much fun reading it, as I did writing it. Enjoy!

## Style

In terms of stylistic approach, I have followed the opinions of Strunk and White [127] and opted to take an active voice in the writing in this Thesis. I will, therefore, be taking a first-person singular narrative, a choice I have made based on personal preference for clear and concise scientific writing.

## Ethical Issues

The user studies documented in this Thesis involve human participants recruited from staff and students at Swansea University (see Sections 3.6, 4.4, 4.7, 5.6, 6.6). As always, I carefully considered the ethical issues associated with each of the experiments I have run by following the guidance of Field and Hole [42].

Prior to all studies, each participant were asked to complete a consent form and was informed of their right to terminate and withdraw from the study at any time, without penalty. Monetary incentives were given in specified cases in return for the participants' time.

The comfort and well-being of the participants was paramount to each study, as well as the privacy of the data collected. Where recording equipment was used, anonymity was strictly maintained by focusing cameras away from faces and all data collected were later stored securely in a locked cupboard.

These controlled experiments have all been conducted as laboratory tests, as opposed to field studies, and were conducted in the Future Interaction Technology Laboratory, Faraday Tower, Swansea University and were approved by the Computer Science Department's *Ethics and Risk Assessment Committee*.



# Refereed Publications

## Conference Papers

1. George Buchanan, **Jennifer Pearson** : Improving Placeholders in Digital Documents - *European Conference for Digital Libraries (ECDL) 2008*, Pages 1–12, Aarhus, Denmark - **WINNER OF BEST PAPER**.
2. **Jennifer Pearson**, George Buchanan, Harold Thimbleby : Creating Visualisations for Digital Document Indexing - *European Conference for Digital Libraries (ECDL) 2009*, Pages 87–93, Corfu, Greece.
3. **Jennifer Pearson**, George Buchanan, Harold Thimbleby : Improving Annotations in Digital Documents - *European Conference for Digital Libraries (ECDL) 2009*, Pages 429–432, Corfu, Greece.
4. George Buchanan, **Jennifer Pearson** : An Architecture for Supporting RFID-Enhanced Interactions in Digital Libraries - *European Conference for Digital Libraries (ECDL) 2010*, Pages 92–103, Glasgow, UK.
5. **Jennifer Pearson**, George Buchanan, Harold Thimbleby : The Reading Desk: Applying Physical Interactions to Digital Documents - *Computer Human Interaction (CHI) 2011*, Pages 3199–3202, Vancouver, British Columbia, Canada.
6. **Jennifer Pearson**, George Buchanan: CloudBooks : An Infrastructure for Reading on Multiple Devices - *Theory and Practice of Digital Libraries (TPDL) 2011*, Pages 488–492, Berlin, Germany.
7. **Jennifer Pearson**, George Buchanan, Harold Thimbleby : The Reading Desk: Supporting Lightweight Note-taking in Digital Documents - *Theory and Practice of Digital Libraries (TPDL) 2011*, Pages 438–441, Berlin, Germany.
8. **Jennifer Pearson**, Tom Owen, Harold Thimbleby, George Buchanan : Co-Reading: Investigating Collaborative Group Reading - *ACM/IEEE Joint Conference on Digital Libraries (JCDL) 2012*, Pages 325–334, Washington DC, USA. - **NOMINATED FOR THE VANNEVAR BUSH BEST PAPER AWARD**.
9. **Jennifer Pearson**, George Buchanan, Harold Thimbleby : Investigating Collaborative Annotation on Slate PCs - *International Conference on Human Computer Interaction with Mobile Devices & Services (Mobile HCI) 2012*, In Print, San Francisco, USA.

## Journal Articles

1. **Jennifer Pearson**, George Buchanan, Harold Thimbleby, Matt Jones : The Digital Reading Desk: A Lightweight Approach to digital Note-Taking - *Interacting with Computers*, In Print, Elsevier.

## Workshops

1. **Jennifer Pearson**, George Buchanan : Improving User Interaction in Digital Books - *BooksOnline '09: 2nd Workshop on Research Advances in Large Digital Book Collections*, Found at: <http://research.microsoft.com/en-us/um/cambridge/events/booksonline09/papers/p1.pdf>, Corfu, Greece.
2. **Jennifer Pearson**, George Buchanan : Real-Time Document Collaboration Using iPads - *BooksOnline '10: 3rd Workshop on Research Advances in Large Digital Book Collections*, Found at: <http://research.microsoft.com/en-us/events/booksonline10/papers.aspx>, Toronto, Ontario, Canada.
3. **Jennifer Pearson**, George Buchanan, Harold Thimbleby, : HCI Design Principles for eReaders - *BooksOnline '10: 3rd Workshop on Research Advances in Large Digital Book Collections*, Found at: <http://research.microsoft.com/en-us/events/booksonline10/papers.aspx>, Toronto, Ontario, Canada.

## Doctoral Consortia

1. **Jennifer Pearson** : Supporting Effective User Navigation in Digital Documents - Bulletin of IEEE Technical Committee on Digital Libraries: *European Conference for Digital Libraries (ECDL) 2009*, Found at: <http://www.ieee-tcdl.org/Bulletin/v6n1/Pearson/pearson.html>, Corfu, Greece.
2. **Jennifer Pearson** : Supporting Effective User Navigation in Digital Documents - CHI 2010 Extended Abstracts: *Computer Human Interaction (CHI) 2010*, Pages 2947-2950, Atlanta, Georgia, USA.

# Contents

<b>Abstract</b>	<b>i</b>
<b>Acknowledgements</b>	<b>iii</b>
<b>Preface</b>	<b>v</b>
<b>Refereed Publications</b>	<b>ix</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Overview . . . . .	1
1.2 Background . . . . .	2
1.3 Thesis Outline . . . . .	4
1.4 Scope . . . . .	5
<b>2 Background</b>	<b>7</b>
2.1 Overview . . . . .	7
2.2 Reading . . . . .	7
2.2.1 Active Reading . . . . .	7
2.2.2 Reading in Conjunction with Writing . . . . .	8
2.2.3 On-Screen Reading . . . . .	9
2.2.4 Reading and Cognition . . . . .	10
2.3 Lightweight Interaction . . . . .	11
2.3.1 Paper Versus Digital . . . . .	12
2.4 Cognition and the User . . . . .	13
2.4.1 Ready-to-hand and Present-at-hand . . . . .	13
2.4.2 Invisible Computers . . . . .	13
2.4.3 Flow . . . . .	14
Distraction . . . . .	14
2.4.4 Affordance . . . . .	15
2.4.5 Metaphors . . . . .	17
2.5 Digital Document Interaction . . . . .	17
2.5.1 The Visual Book Metaphor . . . . .	17
Page Display . . . . .	19
2.5.2 Affordances of Paper Project . . . . .	21
2.5.3 Augmented Reading Hardware . . . . .	22
Single Display . . . . .	22
Dual Display . . . . .	23

	Kinesthetic Augmentations . . . . .	24
2.5.4	Realistic Book Software . . . . .	25
2.5.5	Hybrid Systems . . . . .	27
2.6	Summary . . . . .	29
<b>3</b>	<b>Placeholders . . . . .</b>	<b>31</b>
3.1	Introduction . . . . .	31
3.1.1	Overview . . . . .	31
3.2	Background . . . . .	32
3.2.1	Document Revisitation . . . . .	32
3.2.2	Placeholder Use . . . . .	32
3.2.3	Physical Placeholders . . . . .	33
3.2.4	Digital Placeholders . . . . .	34
	Web-Bookmarking . . . . .	34
	Digital Document Bookmarks . . . . .	34
	Electronic Reading Device Bookmarks . . . . .	34
3.2.5	Aims . . . . .	36
3.3	Literature Review . . . . .	36
3.3.1	Web Favourites . . . . .	36
3.3.2	Digital Document Bookmarks . . . . .	37
3.4	System Design . . . . .	38
3.4.1	Tabs . . . . .	38
3.4.2	Ordering . . . . .	39
3.4.3	Form Layout and Navigation . . . . .	40
	Alternative Designs . . . . .	40
3.5	System Implementation . . . . .	41
3.5.1	Bookmarks . . . . .	42
3.5.2	Saving and Opening . . . . .	42
3.5.3	Mouse-overs . . . . .	43
3.6	User Study . . . . .	44
3.6.1	Study Design . . . . .	44
	Control Systems . . . . .	44
	Procedure . . . . .	44
	Study Tasks and Documents . . . . .	45
	Interview Procedure . . . . .	45
3.6.2	Hypotheses . . . . .	46
3.6.3	Study Results . . . . .	46
	Placeholder Usage . . . . .	46
	Placeholder Types . . . . .	48
	Subjective Reviews . . . . .	48
3.7	Conclusions . . . . .	49
3.7.1	Lightweight Properties . . . . .	50
3.7.2	Future Work . . . . .	51
3.7.3	Related Publications . . . . .	51

<b>4</b>	<b>Annotations</b>	<b>53</b>
4.1	Introduction	53
4.2	Background	53
4.2.1	Physical Annotations	54
4.2.2	Digital Annotations	55
4.2.3	Anatomy of Annotation	56
4.2.4	Marginalia	58
4.3	Literature Review	59
4.4	Paper Study	60
4.4.1	Study Design	61
	The Systems	61
	Documents and Apparatus	61
	Procedure	62
4.4.2	The Tasks	63
4.4.3	Hypotheses	63
4.4.4	Study Results	63
	Subjective Results	63
	Annotation Placement	64
4.4.5	Paper Study Conclusions	66
4.5	System Design	67
4.5.1	Aims	67
4.5.2	Additional Document Space	68
4.5.3	Storage	68
4.5.4	Additional Media	69
	Scenario	69
	Alternative Designs	70
4.6	System Implementation	70
4.6.1	Tools	70
	PDF Tools	70
	Margin Tools	72
	Miscellaneous Tools	72
4.6.2	Saving and Opening	73
4.7	User Study	73
4.7.1	Study Design	73
	Research Questions	73
	Control Systems	73
	Procedure	74
	Study Tasks	74
4.7.2	Study Results	75
	Research Question 1 Results	75
	Research Question 2 Results	76
	Research Question 3 Results	77
	Research Question 4	78
4.8	Conclusions	78
4.8.1	Lightweight Features	79
4.8.2	Related Publications	80



<b>5</b>	<b>Note-Taking</b>	<b>81</b>
5.1	Introduction	81
5.1.1	Overview	81
5.2	Background	82
5.2.1	Issues Associated with Digital Note-Taking	82
	Space	82
	Tool Overload	83
	Menu Navigation	83
5.2.2	Scenario	83
5.2.3	Aims	84
5.3	Literature Review	85
5.4	System Design	86
5.4.1	Static Workspace	87
5.4.2	Appropriation	87
5.4.3	Drag-and-Drop Manipulation	87
5.5	System Implementation	87
5.5.1	The Virtual Desk	87
5.5.2	The Unified Post-it Tool	87
5.5.3	Drag-and-Drop	89
5.5.4	Form Layout and Navigation	89
5.5.5	Additional Functionality	90
	Colour and Size	90
	Snapping	90
5.6	User Study	90
5.6.1	Study Design	90
	Control Systems	90
	Procedure	93
	The Tasks	93
	Study Documents	94
	Selection of Study Metrics	94
5.6.2	Hypotheses	95
5.6.3	Study Results	95
	Post-it Usage	95
	Subjective Findings	96
	General Questions (Q1 - Q5)	96
	Specific Tasks (T1 - T9)	98
	Patterns of Use and Behaviour	99
	Hypotheses	100
	Digital Versus Physical	102
5.7	Conclusions	103
5.7.1	Lightweight Properties	103
5.7.2	Future Work	105
5.7.3	Related Publications	105

<b>6</b>	<b>Visual Indexing</b>	<b>107</b>
6.1	Introduction . . . . .	107
6.1.1	Overview . . . . .	107
6.2	Background . . . . .	108
6.2.1	The Four-Phase Framework for Text Searches . . . . .	108
6.2.2	Current Methods . . . . .	108
	Printed . . . . .	108
	Digital . . . . .	110
6.2.3	Aims . . . . .	111
6.3	Literature Review . . . . .	111
6.4	System Design . . . . .	113
6.4.1	Index Builder . . . . .	113
	Page Clusters . . . . .	113
	Hyper-links . . . . .	113
	Tool Tips . . . . .	114
6.4.2	Visual Features . . . . .	114
	Colour . . . . .	114
	Size . . . . .	114
6.4.3	Form Layout Navigation . . . . .	115
	Alternative Designs . . . . .	115
6.5	System Implementation . . . . .	116
6.5.1	Page Relevance . . . . .	116
6.5.2	Visual Representations . . . . .	117
	Colour Tag . . . . .	118
	Tag Cloud . . . . .	118
	Graph . . . . .	118
6.6	User Study . . . . .	119
6.6.1	Study Design . . . . .	119
	Control Systems . . . . .	119
	Procedure . . . . .	120
	The Tasks . . . . .	120
	Selection of Study Metrics . . . . .	120
6.6.2	Hypotheses . . . . .	121
6.6.3	Study Results . . . . .	121
	Speed . . . . .	123
	Accuracy . . . . .	123
	Subjective User Ratings . . . . .	125
	User Behaviours . . . . .	125
6.7	Conclusions . . . . .	126
6.7.1	Lightweight Features . . . . .	126
6.7.2	Future Work . . . . .	128
6.7.3	Related Publications . . . . .	128

<b>7</b>	<b>Discussion and Conclusions</b>	<b>129</b>
7.1	Overview of Problem . . . . .	129
7.2	Proposed Solution . . . . .	131
7.3	Contributions . . . . .	131
7.3.1	Placeholders . . . . .	131
	Lightweight Properties . . . . .	132
7.3.2	Annotations . . . . .	132
	Lightweight Properties . . . . .	133
7.3.3	Note-Taking . . . . .	133
	Lightweight Properties . . . . .	134
7.3.4	Visual Indexing . . . . .	134
	Lightweight Properties . . . . .	134
7.4	Research Scope and Limitations . . . . .	135
7.4.1	Cognitive Load . . . . .	135
7.4.2	Study Duration . . . . .	135
7.4.3	Number of Participants . . . . .	136
7.4.4	Scale of Annotation . . . . .	136
7.5	Running Themes . . . . .	136
7.5.1	Space Beyond the Document . . . . .	136
7.5.2	Visual Language . . . . .	137
7.5.3	Immediacy of Access . . . . .	137
7.5.4	Digital Technologies . . . . .	138
7.5.5	Designing for Appropriation . . . . .	139
7.5.6	Completeness of Metaphor . . . . .	139
7.6	The Book Metaphor . . . . .	139
7.7	Electronic Documents . . . . .	140
7.8	Concluding Remarks . . . . .	142
<b>8</b>	<b>Future Work</b>	<b>145</b>
8.1	Introduction . . . . .	145
8.2	Lightweight Application . . . . .	146
8.2.1	Collaboration . . . . .	147
8.2.2	Related Publications . . . . .	149
<b>A</b>	<b>Research Techniques</b>	<b>153</b>
A.1	Statistical Tests . . . . .	153
A.2	Software Used . . . . .	154
A.2.1	Third Party PDF Rasterizers . . . . .	154
	JPedal . . . . .	155
	TallPDF . . . . .	155
	PDFTron . . . . .	155
<b>B</b>	<b>Technology Review</b>	<b>157</b>
B.1	Introduction . . . . .	157
	B.1.1 Overview . . . . .	157
B.2	Background . . . . .	158
	B.2.1 Reading Features . . . . .	158

Bookmarks . . . . .	158
Annotations . . . . .	158
Page Numbering . . . . .	158
Magnification . . . . .	159
Page Turning . . . . .	159
B.2.2 Consistency . . . . .	159
B.2.3 Completeness of Metaphor . . . . .	160
B.2.4 Ergonomics . . . . .	160
B.3 PRS-300 . . . . .	160
B.4 PRS-505 . . . . .	163
B.5 PRS-600 . . . . .	165
B.6 Kindle . . . . .	168
B.7 nook . . . . .	170
B.8 iPad . . . . .	173
B.9 Galaxy Tab . . . . .	176
B.10 Discussion and Conclusions . . . . .	177
B.10.1 Page Numbers . . . . .	177
B.10.2 Page Turning . . . . .	178
B.10.3 Bookmarks . . . . .	178
B.10.4 Magnification . . . . .	179
B.10.5 Organisation . . . . .	179
B.10.6 Consistency . . . . .	180
B.10.7 Completeness . . . . .	180
B.10.8 Ergonomics . . . . .	180
B.10.9 Problematic Areas and Suggestions for Improvement . . . . .	181
B.11 Related Publications . . . . .	182
<b>Bibliography</b>	<b>183</b>
<b>Index</b>	<b>195</b>



# Chapter 1

## Introduction

There is fun to be done!

---

Dr. Seuss

### 1.1 Overview

The advent of the personal computer brought predictions of the so called ‘paperless office’ [118], promising the elimination of paper from everyday office-related tasks. Twenty years on from this original prediction [20], and digital documents are becoming more and more portable with the introduction of eReaders and mobile multi-function devices such as iPads [137]. New release novels and articles which were once bound to paper are now available for download in an instant, direct from on-line retailers, revolutionising digital reading. Even in the relatively short time I have been working on this Thesis, the uptake of documents in digital form has significantly increased. In fact, the use of digital documents has increased so much that in July 2010, Amazon announced that for the first time it had sold more Kindle books than both paper and hard backs for the top 10, 100 and 1000 best-selling books on its US website<sup>1</sup> [95, 134]. It was also announced this year that the prestigious Man Booker Prize<sup>2</sup> would be digitised for the first time by giving judges eReaders instead of piles of physical books [96].

This substantial increase in digital document uptake means it is timely to investigate the issues surrounding reading in its electronic form. There are many examples where on-screen reading tools have proven themselves deficient compared to their physical equivalents [85, 107], an issue that if left unresolved, could hinder the continued growth of digital document use. This Thesis explores current tools designed to aid the on-screen reading process and improves upon them by incorporating lightweight techniques that pose a minimal intrusion on the primary reading task. By doing so, I aim to increase user satisfaction of the tools and reduce the cognitive attention required to use them. Furthermore, the evaluation of each of the improved implementations allows me to produce a list of guidelines to which future on-screen reading systems can adhere. It is my conjecture, that following these guidelines will aid in the creation of lightweight user-centred interfaces for digital reading.

---

<sup>1</sup>[www.amazon.com](http://www.amazon.com)

<sup>2</sup>[www.themanbookerprize.com](http://www.themanbookerprize.com)

## 1.2 Background

There have been numerous studies that investigate the concept of digital reading and its applications. Princeton University for example, has recently performed a pilot study on eReading in a classroom setting [104]. The voluntary project which spanned a University semester, studied 51 students' use of the Kindle DX to determine if eReaders could effectively reduce their paper use without harming the overall classroom experience. The results of this study conclude that (at least for the students participating in the study) the use of eReaders significantly reduces the amount of printing performed by students. In fact, the average difference in paper use for students using the Kindle versus the control group (no eReaders) was as much as 55%. Princeton University as a whole however, claim that since the inception of digital document delivery on campus, their printing use has actually *increased* [105], a reality that has also been documented by economists in recent years [39].

This evidence points to a significant increase in digital document availability coupled with a steady rise in paper usage. Although we might expect that a higher proportion of digital documents would mean a decrease in paper use, in actuality the opposite is true, which suggests that many users print their downloaded digital documents. Since digital print-outs are expendable and easily replaced, many users will print several copies to distribute amongst others or to make notes on, exacerbating the already increasing use of paper. Clearly then, this 'print to read' [81, 118] mentality is still a common occurrence despite the introduction of 'paper-like' reading screens, which causes scepticism [14, 43, 49] of an office free from the printed page.

The reasoning behind many users' insistence of printing digital documents demonstrates a clear failure of on-screen reading and the tools designed to aid it. In a perfect world, reading and manipulating digital material would be almost indistinguishable from reading and interacting with paper – an ideal that would certainly bring the notion of the paperless office closer to reality.

The key concept which is dependent upon the usability of paper-like software, and one that is central to the work presented in this Thesis, is the act of reading. Reading is a complex and multi-faceted activity that takes years to master. For the majority of the time reading is not passive, instead it is performed in conjunction with other activities such as writing and highlighting [107], or jumping non-linearly within a document [85, 106]. This concept can be referred to as 'attentive' or 'active' [3] reading. One of the early studies on the design of digital reading concluded that reading actually occurs in conjunction with writing over half of the time [2].

Performing active reading on-screen however, is fraught with problems [107], from hardware issues relating to screen size and resolution, to effective navigation and mark-up facilities. It is well documented that both reading and active reading are easier [2, 52, 139] and faster [93] to perform on paper than they are on-screen. The main reason for this is that paper offers several advantages over computer screens including ease of annotation, navigation and flexibility of spatial layout [107].

The tangible properties of paper (e.g., it is light, thin and flexible) afford many actions that are not possible on their digital equivalents [106, 117]. For example, paper can be folded, ripped and stacked, it can be flicked, highlighted and scribbled on, yet requires no batteries to function. There have been many researchers who have studied affordance in the context of paper [36, 107, 114, 117] both for investigative purposes and to aid in the production of more paper-like digital readers. The New Yorker magazine [47], in 2002 stated:

*"Digital documents, of course, have their own affordances. They can be easily searched, shared, stored, accessed remotely, and linked to other relevant material. But they lack the affordances that really matter to a group of people working together on a report."*

– The New Yorker [47].

The cognitive attention required to accomplish these tasks on paper is minimal, which often means that people will do it without thinking. For example, doodling on a piece of scrap paper then folding it inside a notebook can be done while in a lecture, an activity that does not usually break the user's attention out of the primary task of listening to the lecture content. In contrast, digital document manipulation is far less intuitive [52] and consequently makes active reading tasks significantly more difficult.

The nostalgia of paper may also play a role in the scepticism surrounding the use of digital documents. Physical books are familiar, and after years of use often 'tell a story' about their inanimate life. Back in 1999 Gass [43] wrote an article about the pleasures of reading printed books, describing how the pages of his beloved copy of *Treasure Island* still bore the jam-stains from his childhood – a memory that now provides him with increased engagement with the book itself. In contrast of course, digital books do not age, and therefore do not contain any of the personal quirks that physical books acquire over their lifespan. Although there has been research into representing wear and tear on digital documents [58], for the most part, digital reading software maintains the pristine aesthetics of an eDocument throughout its entire existence.

Previous work on digital document manipulation has repeatedly demonstrated that poor human-computer-interaction in digital reading inhibits user performance. For example, Marshall [85] and Sellen [107] have both diagnosed several problems with reading and interacting with documents on-screen, but do not posit any technical solutions. Other researchers such as Golovchinsky [114, 115] propose solutions to specific problems within the digital reading world, but do not give any generalised improvements, suggesting that there is a clear gap in knowledge.

One approach to creating electronic documents that a variety of researchers have utilised is to follow the traditional book metaphor [71]. Traditional paper-based books are familiar and there is speculation that maintaining the same model on-screen improves users' interaction to the same information. In this context, readers are presented with a digital representation of a document that can be browsed and manipulated in a similar way to a paper book. Clearly therefore, the physicality of paper is central [79, 117] to the way in which we interact with documents. The material value and 'heft' of the printed page undoubtedly contributes to its popularity with one of the main objections to the eDocument being the plaintive "*I could never curl up with a computer*" remark.

A useful way of encapsulating the physical properties of paper is by 'lightweight navigation', a term that describes actions that are performed with little conscious effort [85]. Although this term was originally coined to describe the affordances of paper, it is possible that it could also be applied to digital technology by probing what electronic properties can be performed with a minimal amount of cognitive attention.

Given the overwhelmingly large increase in digital document uptake in recent years [97, 134], it seems timely to investigate the problems associated with electronic document manipulation tools in order to better understand and better inform future designs. Ideally, digital document software would allow users to annotate, manipulate and interact seamlessly in a way that is minimally intrusive to the user, preferably by capturing the affordances of paper while at the same time transcending its limitations.

The term 'lightweight' and also its reverse: heavyweight, which describes actions that take a lot of conscious effort to perform, are central to the content of this Thesis. By paying close attention to the 'lightweight' properties seen in physical documents, digital document software can not only incorporate the physical affordances of paper, but also improve upon them by surpassing their limitations. With this in mind, my goal is to prove by example, that lightweight interaction is indeed possible on a digital level. From these examples it will then be possible for me to redefine the digital equivalent of the term to aid in future digital document designs.



## 1.3 Thesis Outline

This Thesis starts with a background investigation into several of the key concepts and literature on the topic. In this background, I explain several analogies that neatly describe what I am attempting to achieve and tie them together to form a general motivation and research plan. I also describe the general literature on the topic of digital document navigation leaving more focused and detailed literature investigations to be covered within later chapters.

Instead of attempting to prove all lightweight attributes in a single integrated system, I have taken the approach of implementing several individual systems each covering different, potentially lightweight features. Each of these implementations were engineered as solutions to problems identified with current digital document designs, and were all tested via individual user studies. I felt that constructing multiple systems, each offering its own contribution, would allow me to focus on specifics, and pin-point exact areas of lightweight interest.

Clearly, there are many potential areas which I could have explored to prove which aspects of digital design could be considered lightweight. Very broadly however, I have decided to focus on two main areas of investigation:

1. Tools that are currently lightweight on paper but are heavyweight digitally.
2. Tools that are currently heavyweight on paper and are also heavyweight digitally.

Both these areas investigate tools that are currently considered heavyweight digitally, as my mission is to improve upon them by incorporating lightweight navigation into their design. I will accomplish this by first introducing concepts that are currently seen as lightweight on paper but are heavyweight digitally (1). This will allow me to assess whether mimicking paper practices will improve the usability of digital tools.

Although this is a useful first step, it is crucial to note at this point that printed documents are *not* a panacea, they also contain problems of their own which, in some cases, cause their tools to be heavyweight. By incorporating digital techniques in these cases, I attempt to overcome the problems of paper by introducing potentially lightweight digital tools, in this case by investigating those that are heavyweight on paper and are also heavyweight digitally (2).

The first of the lightweight implementations which I discuss in Chapter 3 is concerned with placeholding – a very common activity within printed documents but one which is poorly translated digitally. This chapter then, is an example of a tool that is currently lightweight on paper but heavyweight digitally. I chose to investigate placeholders as a starting point for discovering lightweight attributes as it is one of a small set of functions that is commonly seen within digital reading systems, and one which I feel, is relatively heavyweight in its current form.

During my investigation into placeholding, I made several notable discoveries, including the possibility that the space surrounding the document provides a useful area for additional contextual information. In Chapter 4, I substantiate this theory, by undertaking an in-depth, formulative investigation into annotations in both their printed and digital forms. By first analysing how and where they are made on paper documents, I was able to implement a digital solution that mimics paper annotation practices in order to gain insight for Chapter 5, a system that incorporates both placeholding and annotation into a single tool.

Following on from the previous two chapters on placeholding and annotation, in Chapter 5, I investigate the possibility of appropriation [37], a concept that is rarely seen in the digital world but is extremely common physically. I also explore the idea of a static digital workspace comparable to

a desk or wall in the physical world, and also the notion of direct manipulation for the creation and deletion of common objects.

As well as researching topics that are considered lightweight physically, I also want to prove that some areas which are considered heavyweight on paper, can be improved on the digital level by incorporating electronic techniques. By doing so, I aim to determine some lightweight properties that are specific to digital document design and in turn prove that paper is not the perfect medium for active reading. In Chapter 6 therefore, I investigate back-of-book indexing, a feature that is not particularly lightweight on physical or digital documents. I chose to investigate the area of indexing, as it demonstrates a clear example of an activity that can be significantly improved by incorporating ideas from the physical world with the computational power of a machine.

All of the implementations discussed above have been thoroughly tested by means of individual user studies which has allowed me to gain insight into the specifics of each system's features. Using the information gathered from earlier studies, I was then better informed for the design of later interfaces which eventually enabled me to lay down several lightweight attributes in Chapter 7. This list of lightweight properties can then be used by designers as guidelines for the creation of new digital active reading software.

The penultimate chapter is a conclusion of the main findings and contributions of the Thesis, including the properties that can be considered digitally lightweight. I end with ideas for future work, and a short discussion of a lightweight interface implemented using the lightweight attributes I have discovered throughout this Thesis.

In addition to the main Thesis content, Appendix B includes an analytical review of several of the most popular portable reading devices on the market at the time of writing. This investigation provided me with an understanding of the strengths and weaknesses of these devices which in turn, aided in my investigation into lightweight design.

## 1.4 Scope

This is clearly a very large problem to which there are many possible avenues of investigation. Exactly what is digitally 'lightweight' can be found by probing a large number of areas. It would be unrealistic in the time available therefore, for me to investigate every possible problematic area. For this reason, I chose to investigate the topic by selecting several key areas that I felt covered the main points I intended to prove.

As well as the overall topics I have investigated, there are also several other areas where I could not cover every possible avenue:

- Digital reading software would not be possible without electronic versions of traditional reading material, known commonly as eDocuments. There is no single standard format of electronic document, and many of the devices I have evaluated in Appendix B can read multiple types. For simplicity, I will be using only Adobe PDF files in this Thesis and will often take the term PDF to mean any type of digital document. Obviously, there are many other types of eDocuments. Nevertheless, for the purposes of designing active reading software, the choice of format is arbitrary. I have therefore opted to select the most popular type of digital document for use within my implementations;
- I will only be focusing on single columned documents. This is primarily due to the complications that can arise from multi-columned text such as additional scrolling and the difficulties that come with reformatting for larger text sizes.

- Most of my studies involve standard 17 inch wide-screen displays (any exceptions will be noted), as this is a common type of desktop display. I could perform the same tasks on larger displays to ease reading tasks, but I felt this would not be representative of the majority of users day-to-day activities. Further to this point, I started investigating digital reading on desktop displays. However, the digital reading revolution has become far more substantiated in the short time I have been working on this Thesis. Digital documents are no longer limited to desktop machines and are now supported on a wide range of portable small screen devices such as mobile phones and eReaders.
- For the most part, I will be focusing on paginated documents as they are in-keeping with the book metaphor (see Section 2.5.1). There is one occasion (in the Visual Indexing System - Chapter 6) where I will use a continuous scrolling display (see Section 2.5.1) as it is more consistent with common digital software such as Adobe Acrobat Reader.

## Chapter 2

# Background

If you don't know where you are  
going, any road will get you there

---

Lewis Carroll

### 2.1 Overview

This chapter describes the core background to the Thesis and discusses the general literature on the topic of digital reading. The work reviewed here covers the general area of digital document design. More focused literature reviews will be carried out at the start of each specialised chapter.

### 2.2 Reading

Reading is a very broad topic that has been around for millennia and, inevitably, there is an extremely large amount of research on the topic. This subsection gives an introduction to the major literature on the topic of reading, paying close attention to the concept of *active reading* as well as on-screen reading, and also the cognitive issues associated with both.

#### 2.2.1 Active Reading

Reading is rarely passive. For many, the act of reading is usually accompanied by thinking and learning which often leads to other activities such as note-taking, highlighting and underlining [106]. This process is known as '*active reading*' and was first defined by Mortimer Adler in the 1940 edition of "How to Read a Book" [3]. Active reading is a common activity for those who engage in knowledge-based tasks [2], particularly when in the workplace. Even those who read books for fun may also engage in the active reading process, perhaps by making a list of characters within a novel or writing down words they wish to define later.

Active reading is a pervasive theme throughout this Thesis. The tools I implement in subsequent chapters have been designed specifically to aid in the active reading process, and do so in a way which is minimally disruptive to the primary task. In this context, the primary task will be the active reading task, that is, thinking and engaging with the text. Secondary tasks involve any other activity. For example, when marking-up a document, the primary task of the user will be reading the text

and thinking about how to annotate it; the secondary task will be physically picking up a pen and writing notes on the page. For a user to be fully engaged with the primary task, it is beneficial for the secondary tasks to be as minimally cognitively demanding as possible to ensure the maximum amount of attention is being left for the primary active reading task.

### 2.2.2 Reading in Conjunction with Writing

When considering the topic of active reading, it is worth investigating the topic of reading in conjunction with writing, an action that is often considered one of the main tasks of the active reading process. One of the early studies on the design of digital reading was performed by Alder et al. [2] in 1998. The diary study, which studied work-related reading habits, was conducted for the purpose of informing the design of what they refer to as “*digital reading devices*” and confirmed that in most cases, reading occurs in conjunction with writing. In the study, eight out of their 15 subjects performed reading with writing between 75% and 91% of the time. In fact, in all but one subject, reading was accompanied by writing more often than not. The study also confirmed that there are many different *purposes* for reading. Some examples of these purposes include reading to remind, reading to answer questions, reading to learn, reading for cross-reference and reading to support discussion.

This data concurs with the work of O’Hara [106] and confirms that there are many types of reading other than simply to read for information. Many of the purposes of reading involve reading *in conjunction with writing*. Clearly, the process of writing is a complex activity within the work-related reading process, and one which has many purposes. It is therefore vital to the success of any digital active reading interface, that the tools designed to aid in this process are as universal and easy to use as possible to ensure these motivations are effectively achieved.

In addition to the different reasons one might write while reading, Alder et al. [2] also describe a taxonomy for five different categories of writing used in their analysis:

1. **Creation:** Creating a new document or editing an existing document.
2. **Note-Taking:** The writing of abbreviated or unstructured text used primarily as a temporary way of jotting down ideas before the writing of a final finished document.
3. **Annotation:** Writing on an existing document about the text within it. Annotations usually contain markers to their surrounding document content.
4. **Form-Filling:** Filling in structured forms or writing in a pre-defined and prescribed manner, e.g., filling out a tax form.
5. **Updating:** Updating calendars or schedules.

Here, Adler et al. take the term *note-taking* to be the unstructured comments made before embarking on writing a larger document. In this context then, it is separated from the term *annotation*, which they take to mean the informal markings made upon pre-existing literature. In this Thesis, I too will be separating these terms, investigating the topic of annotation which discusses mark-ups made on top of pre-made documents, and note-taking which in my case refers to the process of making notes on and around documents using digital Post-its.

Although active reading covers a diverse range of activities, the process of reading in conjunction with writing forms a large portion of a user’s engagement with a document, and one which I will be following closely throughout several of the investigations I discuss in this Thesis. One topic that is central to all the work I present here however, is the concept of on-screen reading, an activity that is integral to any digital active reading system.

### 2.2.3 On-Screen Reading

The main focus of this Thesis is to improve the active reading process on digital documents, which of course relies heavily on reading from computer screens. There have been a host of studies that have investigated on-screen reading. In 1982, Muter et al. [93] performed a study asking 32 participants to read continuous text for two hours, half of them reading from a video screen and the other half reading from a printed book. Although they concluded that there was no significant difference between the comprehension or subjective measures of discomfort (i.e., dizziness, fatigue or eye-strain), the results showed that the participants read 28% more slowly on video screens than they did from physical books. A follow up study conducted in the same manner was undertaken by Krunk and Muter [69] in 1984 to investigate the reasoning behind this speed difference. The results from this study suggested that the formatting (i.e., the number of characters and lines per page) and interline spacing were the main reasons for the difference in speed between the two mediums as opposed to the contrast ratio or rendering time. Although these studies may seem somewhat archaic having been performed on monochrome CRT monitors, the principle of the study results remain valid, particularly when applied to small-screen devices such as eReaders where the formatting of text is periodically changed.

The view that text size and formatting hinder the on-screen reading process is also shared by Mills and Weldon [90]. After an extensive review of empirical studies regarding the readability of text from computer screens, the authors argue that:

*“Paper appears easier and faster to read than computer screens, but the size of the effect depends on the quality of both the paper and screen presentation.”*

– Mills and Weldon [90].

Hansen and Haas [52] conducted a series of experiments in an attempt to explain the differences in performance between the two media by identifying seven factors that affect it. The results of four experiments conducted to evaluate these differences, including page size, legibility, responsiveness and tangibility, all confirmed that paper was better than computers for reading in every condition.

In 1992, Dillon [36] produced an extensive review of the literature relating to on-screen reading, and evaluated the reading process in terms of speed, accuracy, fatigue, comprehension and performance. In terms of speed, Dillon concurs with previous works and suggests that there is a performance deficit of between 20% and 30% when reading from screen as opposed to reading on paper.

In terms of subjective preference, the majority of the literature (e.g., [2, 52, 107]) suggests that paper is the favoured medium for reading; a trait that is even more prominent when considering reading in conjunction with writing. There have been many studies that have focused on user behaviour while reading and writing on-screen. Adler et al. [2] for example, discussed the differences between paper and on-screen active reading and concluded that paper based reading and writing accounted for 85% of people’s total activity time whereas on-screen reading and writing accounted for only 13%. These results were even more surprising in this case, as the authors claim that the participants all had at least some essential data stored on their computers which required them to use them for at least some portion of their work. This suggests that paper is a considerably more popular alternative to reading and writing on-screen even in situations where a vital part of the task involves computers. The authors then justify the choice of paper over an on-screen alternative, at least in some situations because:

*“...paper supports their particular reading and writing tasks better, and that the on-line alternatives simply fail to provide the critical affordances of paper.”*

– Adler et al. [2].

The effect these paper affordances have on the active reading process is seen in many aspects of digital document interaction. O'Hara and Sellen [107] for example, state that:

*"Annotation on paper was relatively effortless and smoothly integrated with reading compared to on-line annotation which was cumbersome and detracted from the reading task."*

– O'Hara and Sellen [107].

The tangibility that paper affords is a common theme in the reasoning behind the poor performance of on-screen reading and writing, and results in a high proportion of users printing digital documents to mark them up. Marshall [84] commented that there were occasions when even environmentally mindful users would print documents:

*"Even the most environmentally conscientious reader turns to the printer when asked to review a journal submission, to proofread a document, or to refer to the document in situations that a laptop (or even a portable reading device) would be awkward."*

– Marshall [84].

Sellen and Harper [117] suggest that there are "good reasons" for the continuing use of paper in organisational life, and attempt to justify the low uptake of on-line reading and writing as well as the 'print to read' mentality of many users by stating that:

*"The critical differences have to do with the major advantages that paper offers in supporting annotation while reading, quick navigation, and flexibility of spatial layout. We found that these, in turn, allow readers to deepen their comprehension of the text, extract a sense of its structure, create a plan writing, cross-refer to other documents, and interleave reading and writing."*

– Sellen and Harper [117].

#### 2.2.4 Reading and Cognition

Mayes et al. [88] conducted experiments to determine if reading information on-screen resulted in poorer performance than that of paper. As with the work of Dillon [36], the categories used to determine performance within their studies were classified as: reading time, comprehension of the text and mental workload. The first of their two-study analysis concluded that although there were no significant differences between mental workload and comprehension, those who read on-screen took considerably longer to complete the reading tasks than those who were reading from paper. Realising this, the authors then conducted a second study in order to determine if an increase in demands on working memory was responsible for these performance detriments.

To investigate these issues, Mayes et al. introduce the term 'secondary tasks', (in this case, the memorisation of a list of letters) which can reveal when a user's limited working memory capacity is exceeded. They argue that if reading from a screen does indeed increase workload due to the burdening of additional cognitive demands of the task, then users will be less able to store secondary task information. Thus, in theory, those who are reading from a screen will have a lower performance on later recall of the secondary task information than those who read from paper.

Although their first study concurred with the work of Muter [93], that reading is faster to perform on paper than it is on screen, the results from their second study, where secondary tasks were introduced, indicated that people can in fact read from screens as fast as they can on paper. They suggest

however, that when the secondary task requires attention from the user, those who are reading from paper tended to recall information better than those who read from a screen. They conclude from this that performance, at least in terms of information comprehension, is negatively affected by on-screen presentation.

The suggestion that reading from a computer screen reduces the working memory capacity for reading has been further investigated by Wästlund et al. [140] who conducted studies that avoid any confounding page layout variables such as line length, fonts and kerning. The main findings of this experiment were that reading comprehension is more difficult and mental workload higher when performed on screen than when performed on paper. They therefore conclude that reading on-screen reduces the working memory capacity by reallocating cognitive resources to document navigation.

This evidence suggests that the process of actually interacting with the computer is the reasoning behind this performance deficit. In fact, in a later paper, Wästlund et al. [140] even describe the navigation of on-screen reading to be part of the reader's overall processing capacity:

*“Reading on a computer screen involves both the process of reading the presented text and handling the computer, thereby, while reading a document onscreen, the reader's processing capacity is being utilized not only for decoding but also for page navigation.”*

– Wästlund et al. [140].

It is clear from the literature, that the cognitive workload required to read on-screen is higher than that of paper, thus suggesting that the process of interacting and navigating with the document using the computer is hindering the on-screen reading process. In terms of active reading, this statistic is likely magnified. By introducing another cognitive process into the equation (i.e., writing, highlighting, underlining etc.), the user now has more actions to perform interactively using the computer, a process speculated to utilise additional processing capacity. It is vital to the success of any active reading software then, that the cognitive workload required to use the tools is as low as possible to leave more processing capacity for the primary active reading task.

## 2.3 Lightweight Interaction

In the context of documents, the term ‘lightweight interaction’ was defined by Marshall and Bly [85] and is a useful way of describing the affordances that paper offers over digital. The physical properties of paper afford many actions that are difficult to replicate on the digital spectrum. In their 2005 paper: ‘Turning the Page on Navigation’ [85], Marshall and Bly conducted an in-depth observational study into magazine users’ printed and digital navigation methods. After a thorough inspection of each participant’s reading sessions, as well as a series of interviews and talk-throughs of reading habits, the authors confirmed that the general patterns of within-document navigation were roughly the same in the ePeriodical and paper versions of the magazine. The most intriguing observation they made within the study however, is what they refer to as ‘lightweight navigation’:

*“...navigation that occurs either when people reach a particular page or when they move within an article in a way that is so unselfconscious that they aren't apt to remember it later.”*

– Marshall and Bly [85].



This definition of the term *lightweight navigation* was defined in the context of linear reading: specifically, the unselfconscious, seamless movements that users make out of the linear stream of text and then back again. If an action is lightweight then, the reader is seldom aware they are performing it.

Marshall et al. then go on to describe four separate lightweight navigation types:

- Narrowing or broadening focus by manipulating the physical magazine [article];
- Letting one's eyes stray to a page element out of the textual flow;
- Looking ahead in the text to preview or anticipate;
- Looking back to re-read for context.

These properties which are common activities on paper, are rarely seen in digital document navigation. For example, one of the major lightweight activities the authors observe within their study is the act of page turning. This is described as "*a complex combination of lightweight navigational activities*", a seamless interaction which sadly, is absent from even the most sophisticated of digital page-turning simulations [30, 74].

The lightweight properties described by Marshall and Bly therefore, are examples that are specific to the interactions of physical documents. They speculate however, that this concept of lightweight interaction can also be applied to digital technology, but do not give any concrete evidence to support it. There are many aspects of computerised technology that far exceed the capabilities of paper (i.e., searching, zooming etc.) and by paying closer attention to the possibility of 'lightweight' interaction, digital document software can not only incorporate the physical affordances of paper but also improve upon them by surpassing paper's limitations.

The term 'lightweight' is used frequently throughout this Thesis to describe a task or action that can be performed without a significant amount of cognitive attention. In addition to this, I will also be referring to its corollary: heavyweight - which can be used to describe a task or action that takes a lot of conscious effort to perform.

### 2.3.1 Paper Versus Digital

There are many properties of paper that can be considered lightweight. There are also a large number of physical characteristics that make paper documents so popular. For example, paper is cheap and familiar compared to electronic equipment, it can also be written on easily and can also contain useful meta-data that is not implicitly available on digital documents (e.g., the 'feel' and 'heft' of the book).

In comparison, there are also several aspects of digital documents that exceed the limitations of paper. For example, digital documents are quick and easy to edit, copy and search. They benefit from spell checkers, translators and other computationally complex activities, and can be magnified to reveal more detail. It is also possible to store many thousands of digital documents easily and relatively cheaply on light and portable storage devices.

It is clear from these examples that there are benefits and drawbacks to both paper and digital texts. It is my intention therefore, to combine some of the lightweight properties of paper with digital enhancements to improve the overall usability of electronic documents and hopefully identify what can be considered *digitally lightweight*.

## 2.4 Cognition and the User

Before embarking on a project that incorporates lightweight techniques, it is beneficial to investigate the cognitive demands of users while engaged in demanding tasks such as active reading.

### 2.4.1 Ready-to-hand and Present-at-hand

German philosopher Martin Heidegger originally coined two terms of being: ready-to-hand and present-at-hand [56], to explain the instinctive nature and attitudes towards things in the world. His ideas have greatly influenced the field of cognitive psychology and are also extremely relevant to the broad area of human-computer-interaction [148].

Ready-to-hand (*zuhanden*) describes a scenario where users are engaged with the world in a normal and involved way. Specifically, users are performing tasks without thinking about the mechanics or tools being used to facilitate them. For example, when using a mouse to manipulate menus on a computer, the user in this case is not thinking about the tool itself (i.e., the mouse), only the objects they are directing with it.

Present-at-hand (*vorhanden*) describes a situation where users are no longer thinking about the task at hand, but rather the tools that are facilitating the task. Users in this scenario are conscious of the physical object itself, for example, picking up a fork and thinking about its design as opposed to thinking about what they are eating. This state is extremely undesirable in HCI as it pulls the user's attention away from their primary task. Ideally, we would want users to be thinking about their primary task, *not* the tools that assist it.

When applied to the ideas proposed in this Thesis, it would be beneficial for users to be thinking about their primary task (ready-to-hand), i.e., the active reading task, as opposed to thinking about the tools they use to do so (present-at-hand), i.e., the writing implement. In fact, I believe that a truly intuitive tool should only be present-at-hand when it is in a broken state and users are therefore thinking about how to fix it, at all other times it should be *invisible*.

### 2.4.2 Invisible Computers

Another useful way of describing the ready-to-hand state is with 'invisible computers' [103]. In his 1988 book 'The Design of Everyday Things' [102], Donald Norman suggests that computers should be made invisible to lend more time to the task at hand:

*"When I use a direct manipulation system – whether for text editing, drawing pictures, or creating and playing games – I do think of myself not as using a computer but as doing the particular task. The computer is, in effect, invisible. The point cannot be overstressed: make the computer system invisible."*

– Norman [102].

The idea of invisible computers is similar to those laid down by Marshall et al. and Heidegger: essentially, that the tools employed to complete a task should be so intuitive that they are, in effect, *hidden* from the user:

*"You don't notice the computer because you think of yourself doing the task, not as using the computer."*

– Norman [102].

This concept has also been touched upon by Weiser who is considered by many as the founder of ubiquitous computing, who thinks that computers should “*get out of the way*” [143], or essentially, make technology disappear. Although this was originally defined for ubicomp, it also applies to the broader sense of human-computer-interaction and ties in well with the other views described above.

In summary therefore, computer systems should be designed to be as easy to use as possible, essentially rendering the tools themselves as *invisible*. If the technology is too complicated it can disrupt the user’s primary task of reading and understanding the text – an undesirable situation that can lead to the loss of *flow* with the active reading task.

### 2.4.3 Flow

Hungarian psychologist Mihaly Csikszentmihalyi first outlined the theory of ‘flow’, which suggests that people are most happy when they are in a state of complete concentration or absorption in a task. This level of attention is theorised as the point at which users reach the optimal state of ‘intrinsic motivation’, to the point at which they are so absorbed in the task at hand that all other considerations are forgotten.

In his influential work: *Flow: the Psychology of Optimal Experience* [34], Csikszentmihalyi defines the term as:

*“The state in which people are so involved in an activity that nothing else seems to matter; the experience itself is so enjoyable that people will do it even at great cost, for the sheer sake of doing it.”*

– Csikszentmihalyi [34].

This state of ‘flow’, which is also known colloquially as ‘being in the zone’, ‘on the ball’ or ‘in the groove’ is well illustrated in terms of challenge and skill level by the famous graph [35] shown in Figure 2.1. From this chart, we can clearly see that in order to achieve the desired state of ‘flow’, there must be a balance between skill and challenge levels, that is, the skill level of the user and the challenge level of the task must both be high. If a task is too easy, or the skill level of the individual is too small, then the state of flow cannot be achieved.

When applied to the active reading process, it is beneficial for a system to encourage a state of flow with the primary task by making the tools themselves invisible, avoiding the undesirable present-at-hand state. For example, a user may be in the state of flow writing notes on a document; in this state they are not thinking about the tool they are using to make the notes (i.e., the pen). If however, the pen runs out of ink during the process, their attention then shifts from the active reading task to the secondary task of repairing or replacing the pen, and the state of flow is consequently lost.

### Distraction

Pace [109] discusses the possibility of distraction and how in some cases, attention may be drawn involuntarily to stimuli outside the attentive focus [40].

*“Web users tend to ignore minor distractions during a flow experience because their attention is focused on the task at hand. But a distraction that has sufficient intensity, frequency or importance to cause a shift in the user’s attention will terminate a flow experience.”*

-Pace [109].

As Pace notes, attention is a vulnerable process that can be easily diverted via a host of distractions. For example, environmental factors such as loud noises, physiological distractions such as hunger or fatigue and, of course, computer-related distractions such as bugs, software error messages or unresponsive programs. Poor user interface design can also contribute to the loss of attention during a computer related task. Assuming therefore, that the user has already achieved the desired state of flow with their primary active reading task, poorly designed tools may consequently distract the user enough to break them out of their task flow experience.

*“A poorly designed interface can disrupt a flow experience by demanding an excessive amount of attention.”*

-Pace [109].

This point is also highlighted by Marshall [81] in an earlier paper on digital annotation.

*“switching midstream ... is distracting. How much attention is...[a user]...expending to switch from ‘annotate in the margin mode’ ...Will he still be engaged with the text after he has interacted with it?”*

-Marshall [81].

Hence, reducing the amount of attention required to use the tools required for active reading will consequently reduce distraction and therefore lessen the possibility of losing a flow state. This property can serve as an alternative definition or understanding of the term lightweight. Thus, making the tools used for active reading more lightweight, will decrease the likelihood of losing an state of flow with the primary task.

The concept of flow, as it stands is not a core topic of research in this Thesis. However, it is relevant to demonstrate how the use of lightweight design can benefit the usability of active reading systems. That is, the lightweight tools I create in this Thesis do not necessarily aid in the creation of a flow state, but are intended to decrease the possibility of distracting the user away from an already achieved flow experience.

#### 2.4.4 Affordance

There has been a significant amount of research into the concept of *affordance*, a term that the Oxford English Dictionary describes as:

**aff-or-dance** (*noun*)

1. A characteristic of an object, esp. relating to its potential utility, which can be inferred from visual or other perceptual signals; (more widely) a quality or utility which is readily apparent or available.

Although the term affordance was actually invented by perceptual psychologist J. J. Gibson [45, 46], in the context of human-computer-interaction it was first coined by Don Norman in his 1988 book ‘The Psychology of Everyday Things’ [102]. What Norman is actually referring to here is *perceived affordance*: specifically, whether or not the user perceives that some action is possible. Culturally then, we are taught to expect certain things. For example, it is widely recognised that knobs are for turning, buttons are for pushing and slots are for inserting things into. As Norman [103] commented:

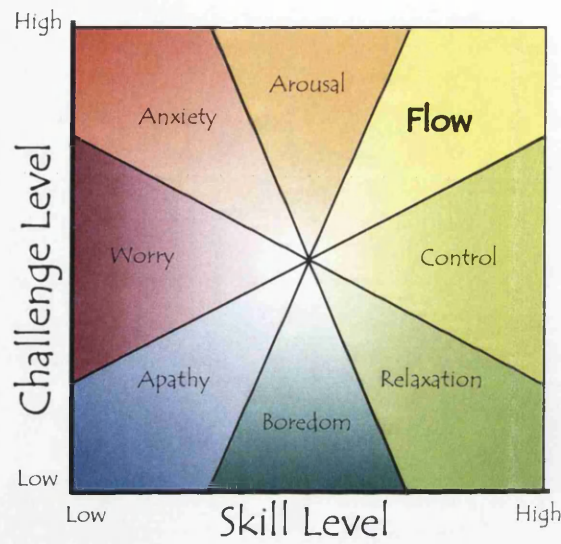


Figure 2.1: Csikszentmihalyi's [35] mental state graph showing challenge versus skill level

*“An affordance is not a property, it is a relationship that holds between the organism that is acting on the object. The same object may have different affordances for different individuals.”*

– Norman [103].

An adult for example, will usually look at everyday objects such as pens or keys in relation to the tasks they were designed to achieve, whereas a child will almost certainly regard them as tasty treats. The idea behind these affordances then, is that, if performed correctly, the user will know instinctively how to perform them and possibly even do it without thinking, making them primarily *ready-to-hand*.

In contrast, if apparent affordances have no actual function, we have *false affordance*, a problem that Gaver describes as people mistakenly try to act [44]. For example, a door with a handle that actually needs to be pushed exhibits false affordance as it indicates to the user that the handle should be used to pull the door.

Sellen and Harper [117] have investigated how the properties of paper support document-related tasks. Although they agree paper offers many useful affordances, they stress that the digital equivalents do not need to mimic the properties of paper. Instead, they suggest that designers take into account these affordances and attempt to provide them in other ways. Tashman et al. [132] also stressed a similar point:

*“Purely mimicking the affordances of paper in a computer-based system may not address all of the opportunities that digital technology can potentially provide.”*

– Tashman and Edwards [132].

Clearly therefore, slavishly replicating every physical interaction in the design of digital document tools is not the optimal way of enhancing them. However, by incorporating some of the more appropriate lightweight physical aspects into digital document design, we can not only reproduce some of the benefits of paper, but also surpass them by incorporating purely electronic techniques.

### 2.4.5 Metaphors

One concept that follows on from the notion of affordance is the visual metaphor.

**met-a-phor** (*noun*)

1. A figure of speech in which a word or phrase is applied to an object or action to which it is not literally applicable;
2. A thing regarded as representative or symbolic of something else, esp. something abstract.

In the context of HCI, the term metaphor can be used to describe the visual relationship users have between a computer interface and real-world objects. For example, the common raised button seen in many modern graphical user interfaces, being slightly three-dimensional in appearance, resembles a physical button and therefore implicitly suggests it should be pushed. This type of metaphorical representation of a real-world object takes advantage of users' intrinsic knowledge of the world and applies a new set of affordances to an electronic depiction of a physical interaction.

When designing a new metaphor, it is important to carefully consider how best to represent an interaction that is currently prominent on the physical plane. Some experts believe that if a digital interaction does not fully comply to the manner consistent with the looks and associations of the real world, then it will confuse users to the point at which their performance will suffer. Heim [57] for example, stated:

*"A metaphor's function must be consistent with real-world expectations...metaphor's should be employed to facilitate learning; they should not contradict the user's previous knowledge."*

– Heim [57].

Although this seems reasonable, there are other researchers who argue that by following this rationale, in some cases, designers can fall into the trap of attempting to create a virtual metaphor of a physical object that is in itself, badly designed. Rogers, Sharp and Preece [119] give the example of an old type of virtual calculator which was, at the time, designed to look and behave like a real-world calculator – a device that was poorly designed even in its physical form.

Before transferring a metaphor from the physical world to the digital domain then, it is essential to the overall usability of the end interface, that the real-world metaphor is carefully chosen. In short, metaphors should not be forced [57]. Rather, only the more intuitive and familiar real-world metaphors should be translated to the digital level.

## 2.5 Digital Document Interaction

### 2.5.1 The Visual Book Metaphor

There are several mediums that have made the transition from the physical to the digital world. When we think about how smoothly vinyls shifted to CDs and eventually to MP3s, or physical photographs to JPEGs or GIFs, it seems reasonable to question the migration of books and documents to the digital plane. One way in which documents can be represented digitally, is by using the 'visual book metaphor', a way of visualising and interacting with digital text using typically physical techniques.

There has been a significant amount of research in the area of the visual book metaphor and its effect on digital document design. Landoni and Gibb [71] for example, like many other researchers in the area, feel that a visual association with familiar concepts significantly aids the learning process. Their 2000 paper, which focuses on the role of metaphors in digital document design, includes the 'Visual Book Experiment', a study that investigates the importance of the visual book metaphor when presenting information on-screen.

The investigation uses software that displays a double-page spread to compare the effect of real books to the the presentation of the same book electronically. The results of the study concluded that users "*found the representation to be consistent with their mental model of a book*" and appreciated the enhanced functionalities the software provided.

In a later paper, Wilson and Landoni [147] make several recommendations for effective eBook design including the adherence to certain aspects of the book metaphor (e.g., covers, indication of book progress and easy to use bookmarking and annotation) and the incorporation of hypertext (e.g., for table of contents and indexes).

Though there are several researchers (e.g., [33]) who believe, like Landoni, that embedding book metaphors such as indexes, table of contents, bookmarks and so on, enhances digital versions of documents, there are also some who think that these are not necessary for ease of use. Nielsen [99] for example, suggests that although readers are more likely to assimilate information faster when they encounter a familiar format (in this case, a paper-like book representation), future designs should in fact avoid implementing the book metaphor due to the limits it enforces upon conceptual models of search and non-linear navigation. In a later AlertBox post [100], Nielsen commented that

*"...the book is too strong a metaphor."*

– Nielsen [100].

as it tends to lead designers away from the potential of new media capabilities, thus meeting with the functions of paper but never actually surpassing them.

The idea of exceeding the limitations of paper by applying digital techniques was also touched upon by Shneiderman [125] in his book *Designing the User Interface* where he discusses the concept of on-line manuals. He states here that:

*"...the designers will be most effective if they can redesign the manuals to fit the electronic medium and to take advantage of multiple windows, text highlighting, color..."*

– Shneiderman [125].

and:

*"...a close match between printed and online manuals can be useful."*

– Shneiderman [125].

This suggests that combining elements from the physical world with digital enhancements will improve the interaction with digital reading systems. The work presented in this Thesis focuses very much on the concept of the book metaphor by combining lightweight techniques seen in the physical world with purely electronic enhancements in an attempt to improve the overall interaction design of digital reading systems.

## Page Display

One important aspect of digital document representation is the way in which the text is displayed - particularly when considering the visual book metaphor. Pagination is the process of dividing documents into discrete pages for display.

### **pag-i-na-tion** (*noun*)

1. A sequence of numbers or signatures assigned to pages in a book, periodical etc. Also: The action of marking the pages of a book or other written text with such a sequence.

Before the age of computers and digital texts, the main purposes of pagination was to indicate the correct order of content within a document and to facilitate easy referencing to specific points within the text. Page breaks on such physical documents occur naturally either when the end of a page is reached or when the editor feels a new page is appropriate (e.g., at the end of a chapter, or to avoid orphaned lines for example). With the exception of ancient scroll parchments where pagination was not necessary, all modern physical books include pages.

Digital documents however, are not bound by the physical page and can therefore be split anywhere depending upon where and how they are being displayed. Many web-based applications (e.g., search page results or forum threads) will display a specific number of lines before a break, whereas smaller screen reading devices alter the page breaks depending upon the size of the text font. Alternately, one may choose not to paginate text at all, leaving reams of continuous text on the same page, forcing the use of scroll-bars for navigation as opposed to the common 'next' and 'previous' links.

During this research, I have made use of several types of pagination, each chosen depending upon the style of interface I was implementing. Some of the more common types of pagination used in digital document readers are described below:

**Single Page Display:** Only *one* page is displayed on the screen at a time. The 'next' and 'prev' (or single click scroll-bar) changes the view to the next or previous page. Example shown in Figure 2.2a.

**Single Page Continuous Display:** Displays every page of the document one directly after another on the screen. Controlled via the scroll-bar for continuous navigation, or the 'next' and 'prev' buttons to skip to the start of a page. This option facilitates un-interrupted reading as it can display the top of page at the same time as displaying the bottom of the previous page. Example shown in Figure 2.2b.

**Double Page Display:** A *double page spread* (two adjacent document pages) are displayed on the screen at a time. as with the single page display, the 'next' and 'prev' buttons are used to change the currently open pages. Example shown in Figure 2.2c.

**Double Page Continuous Display:** Identical in behaviour to the single page continuous display, except instead of a single page, it displays a double page spread. Example shown in Figure 2.2d.

**No Pagination:** Text content is *not* paginated at all, scroll-bars are used to progress linearly through the document and hyper-links are utilised to skip to specific points within the text. Example shown in Figure 2.2e.

The applications described throughout this document make use of several of these pagination methods, depending upon the function of the system in which they are included. See individual chapters for the reasoning behind each design choice.






Figure 2.2: Examples of five popular ways of paginating digital documents

### 2.5.2 Affordances of Paper Project

The Affordances of Paper Project [106] was a venture conducted in the mid 1990s within Xerox under the researchers Abigail Sellen and Kenton O'Hara. The aim of the project was to understand the persistence of paper in the workplace by analysing how the physical properties of paper support human work and interaction.

Reading is not a stand-alone activity; there are a diverse range of reading strategies as well as navigational and manipulation issues that can be applied to a wide collection of reading activities. One of the key goals of the project was to characterise how reading strategies are influenced by the goals and motivations of the reader and to consequently develop a framework to understand the interactive properties of paper.

The heart of the project focuses on the concept of affordance, paying particular attention to the understanding of how paper's affordances support the task of reading and comprehending documents. This understanding was achieved by comparing these paper based affordances with the affordances of digital document readers for the same class of reading tasks.

The typology they present shows how the different reading strategies adopted by users enforce a set of demands upon the properties required of the reading medium. The demands they mention include:

- The ability to make notes on and annotate a document in a way which allows integration with the primary task of reading and comprehension;
- The ability to integrate reading with other ongoing activities;
- The ability to move freely through a document in a non-linear fashion without getting lost;
- The ability to access several pages at once for simultaneous viewing;
- The ability to support concurrent access to multiple documents;
- The ability to browse the text to get a feel for which parts may be interesting;
- The ability to rapidly scan for specific bits of information;
- The ability to provide context and other cues which support the development of a sense of text.

The physical properties of paper mean that it generally supports these demands without significant effort from the user. The authors suggest that although emerging digital technologies make some progress towards meeting these demands, they do still suffer problems.

One of the major points to take from this study is the demands required of the medium in which a document is presented: in my case, electronic. Tying this in with the notion of primary and secondary tasks and the active reading process, the first demand is vital to the work presented in this thesis:

*"...the ability to make notes on and annotate a document in a way which allows integration with the primary task of reading and comprehension..."*

– O'Hara [106].

This point confirms the importance of making digital mark-up tools lightweight and ready-to-hand to allow the maximum amount of cognitive attention to be paid to the primary task of reading and comprehension.

### 2.5.3 Augmented Reading Hardware

Another way in which the book metaphor has been explored by researchers over the last few years, is by augmented reading appliances. There are many ways of categorising these devices, one of which is how the hardware itself has been modified to resemble the physical book. With this in mind, I have divided this review into three sections: hardware with a single display, hardware with dual display, and hardware that also includes bespoke kinesthetic augmentations.

#### Single Display

One of the principal works on the topic of digital reading appliances is the XLibris project [86, 111, 114, 115], which was developed at the FX Palo Alto Laboratory under the principal investigator Gene Golovchinsky. The project, which focuses heavily on augmented hardware that imitates physical paper, is concerned with the process of active reading outside of the desktop PC and the impact of both physical form and interaction on the overall reading experience.

In the projects two 1998 papers [114, 115] and demonstration [111] on the topic, the authors introduce the 'active reading machine' – a pen based tablet display that uses the 'paper document metaphor' to support the analytic reading of digital documents. By using a modified tablet PC that includes pressure strips and animations for page turning and free form digital ink for annotations, the XLibris system is able to support paper-like tangibility from an entirely digital device. The XLibris project endeavours to combine the affordances of paper with computational augmentation and portability in an attempt to reduce the standard 'search and print' model of digital libraries. Specifically they describe their creation as:

*"...a digital library reading appliance that combines the mobility and affordances of paper with computational augmentation."*

–Marshall, Golovchinsky and Schilit [86].

In 1999, the XLibris research team published a paper outlining a user study designed to test the active reading machine [86]. The study, which involved introducing the XLibris prototype into an on-going reading group, included observation of subsequent group meetings, interviews with group members and analysis of their annotations on both paper-based material and the XLibris device. The results gathered from this study were then used to describe the implications for designing future digital library interfaces and appliances and have great value to the work presented in this Thesis. The key results obtained include the characterisation of participants' annotation practices, reference use and analytic reading.

One key observation that resulted from this study was that annotations are used for a multitude of reasons, but the meaning of which is often unclear after the fact, even to the annotation author. They suggest then, that designers should be cautious in interpreting user marks without guidance from the individuals who created them, an idea which I am not concerning myself with in this Thesis.

This study also concluded that reference pursuing is an important, yet relatively rare part of the active reading process. Despite confirming in the study that references were frequently checked and regularly marked, they also found that they were rarely pursued. They therefore suggest, that designers digitise and/or hyper-link citations.

In terms of analytic reading, the authors confirm that XLibris' paper document metaphor supported the complex analytic reading practices found on physical documents (e.g., self-interruption, re-reading, reference checking and annotation). Perhaps more importantly, their result logs confirmed that the participants spent a minimal amount of time explicitly interacting with the device, suggesting

that designers should focus on readability, document layout and physical comfort when implementing a digital reading system. This evidence concurs with my previous statements that the tools designed to facilitate active reading should be *ready-to-hand* or, to some extent, *invisible*.

Another, more recent, attempt at providing a ‘book-like’ hardware platform, this time making use of multi-touch displays as opposed to the single touch tablets used by XLibris, is LiquidText [130, 132]. LiquidText has been designed to support the digital active reading process by incorporating a flexible digital document representation that is controlled via gestures on a multi-touch screen. Unlike the work I present in this Thesis, where I aim to migrate some of the lightweight properties of paper to the digital level, the LiquidText interface does not attempt to replicate the affordances of paper. Rather, it endeavours to surpass their limitations by facilitating a more fluid digital solution with a high degree-of-freedom to annotate and reorganise documents.

The touch-screen nature of the LiquidText hardware allows users to interact with multiple objects at once to provide a rich set of gesture-based interaction techniques for the digital document representation. The interface itself offers a non-paginated display which is collapsible to allow users to simultaneously view multiple parts of the document at once, an issue associated with digital documents that has previously been described as problematic [106]. Clearly, the use of non-paginated displays differs considerably from the paginated structure of the physical book which can cause problems when attempting to replicate some of the lightweight properties of paper. For example, what we know as bookmarks would be obsolete in such a display as there are no ‘pages’ to mark, instead markers must be made on specific parts of the text. The work I present in this Thesis therefore, does not make use of non-paginated displays like LiquidText as my research is concerned with taking the physical book metaphor and refining and enhancing its features in order to create digitally lightweight interactions for active reading systems.

One of the major similarities of LiquidText to the work presented in this Thesis is the extended workspace adjacent to the main document area which is intended to provide space for comments and document excerpts made by the user. This additional space, which includes a ‘fisheye’ lens function for re-scalability, provides a pseudo-infinite continuous workspace that I later identify as being one of the principal lightweight attributes of digital document design.

Another recent development in single display reading hardware is by Chen et al. [26] who have also utilised the slate PC form-factor in their implementation, PageSpark. As with the LiquidText system, the hardware used for the PageSpark interface is off-the-shelf (in this case an iPad) as opposed to the modified technology used by XLibris. The main goal of PageSpark is to enhance traditional static magazine content by creating interactive multimedia applications that run on devices that are close in size and resolution to physical books. One of the most relevant features of the PageSpark system is the page elements interaction that allows independent single column scrolling within multi-columned documents. This technique is a good example of how designers are making better use of digital technologies to enhance digital reading, as opposed to blindly following the conventions set by the physical page, a concept that I follow closely in all the systems I describe in this Thesis.

## Dual Display

Adler et al. [2] in their 1998 paper on work-related reading habits, investigated concurrent reading activities that occur across multiple display surfaces – a technique that is often utilised in cross-referencing. They define a display surface as an independent, tangible surface such as a computer screen or piece of paper. To classify surfaces as being used concurrently, they needed to be used side-by-side for the same activity. The authors here decided that ‘flipping’ between pages within the same paper document and holding open multiple windows on a single computer screen were *not* classed as

examples of parallel concurrent use.

The results of the study concludes that within their participant set, 48% of the time users used at least two concurrent display surfaces. This cross-document activity can be further categorised into three distinct uses:

1. Independent reading and writing displays: One screen is used predominantly to write, while any additional displays are used to read and cross-reference;
2. Multiple surfaces for reading only: No writing occurs; two or more displays were read with reference to each other;
3. Multiple surfaces for reading plus writing: More than one reading display with either an additional display for writing or writing on a display also used for reading.

The use of concurrent surfaces forms an important part of the work presented in this Thesis as the process of active reading is often accompanied by other activities, many of which will require the use of another work surface. For example, note-taking on a separate piece of paper or journal is an instance where multiple surfaces are required for reading and writing, whereas cross-referencing is an example where two surfaces are used to read two documents simultaneously.

One research programme which makes use of multiple surfaces for reading, is Codex – a dual-display e-book reader [27, 59] that facilitates embodied interactions such as folding, flipping and fanning which they describe as examples of lightweight navigation.

### **Kinesthetic Augmentations**

In addition to simply creating software for active reading on portable reading devices, there have also been instances where researchers have created reading hardware with bespoke kinesthetic augmentations to control active reading activities.

Watanabe et al. [141] for example, have developed a system known as ‘Bookisheet’, a bendable interface that consists of two thin acrylic plastic sheets augmented with bend sensors. Using Bookisheet, users are able to manipulate digital texts via a tangible interaction based on the physical book metaphor. The interface, which facilitates book-like page turning through bendable hardware, allows users to control the speed of scrolling by providing a similar elastic tangibility to real books. The two prototype models which consisted of both single and dual screen displays, included features such as audio and tactile scrolling feedback, and a micro-switching bookmark function.

This idea of modifying digital reading hardware with flexible augmentations to preserve the physical affordances of paper, has also been observed by Wightman et al. [146]. In this paper, the authors introduce TouchMark, a document navigation interface that consists of a standard touch-screen tablet augmented with two physical touch sensor tabs. By touching, flicking and squeezing this tab, users are able to navigate and bookmark eDocuments in a flexible way that mimics the affordances of paper.

Another example where flexible hardware has been created to better imitate the lightweight properties of paper, is PaperPhone [70, 98]. Named according to its interactional properties, the PaperPhone is a fully functioning smart-phone inspired device that incorporates a thin and flexible eInk display that resembles the tactile feel of physical paper. The device includes an eBook reader (among other things) and makes use of specific bend gestures to perform functional interactions in its software. For example, bending the corner will cause a sequential shift in page numbers in the reading application.




Figure 2.3: Screen shot from the British Library's 'Turning the Pages' project

Although these examples are primarily concerned with translating the affordances of paper into more intuitive lightweight hardware, they do provide clear examples of how researchers in this area are mimicking the affordances of paper in order to provide more lightweight electronic reading functionality, a concept which is core to this Thesis.

#### 2.5.4 Realistic Book Software

Given the widely recognised advantages of printed media, some researchers have sought to reproduce the experience of physical books as literally as possible when using digital texts, which allows the direct transfer of behaviours between the two mediums. Different approaches have been taken within this general paradigm that attempt to bridge this platform gap, often by incorporating graphical representations of physical book behaviour into digital document reader design. These types of realistic document implementations have a strong relationship with the visual book metaphor.

One of the earlier projects on the topic of realistic three-dimensional book representations was back in 1997, when the British Library introduced its *Turning the Pages* project [74], a virtual software environment for reading scanned documents in a realistic paper-like manner. The software, which has been designed for use on a touch-screen interface, shows a double page spread of what resembles a physical book sitting on a desk (see Figure 2.3). The book itself appears to be three-dimensional and pages can be turned by swiping a finger across the touch screen.

The Turning The Pages project has also inspired early work on realistic digital page turning [30]. The book metaphor and graphical page modelling described in Chu et al.'s 2003 paper [30], was a foundation for what they later describe as "realistic electronic books" [29, 76, 77], a concept that allows users to interact with realistic physical book visualisations of textual-based documents within a digital library. The design of the realistic book system imitates paper in several ways previously un-explored by digital document designers, including: a double page spread, pages curling as the user navigates between them, bookmarks protruding the closed page leaves and applying ageing processes to heavily viewed pages.

To test the effectiveness of three-dimensional realistic book implementations, Liesaputra et al. [75] undertook a user study designed to compare their performance to conventional digital formats (HTML




Figure 2.4: Page turning on the iPad iBooks app

and PDF) as well as with physical books. The results from this study proved that users not only liked the realistic book interface, but also performed information finding significantly quicker and equally as accurately as traditional document formats such as HTML and PDF.

Other researchers have also attempted to create graphical three-dimensional representations of physical books. Card et al. [24, 25] implemented a system called 3Book which was designed with the intention of rendering indefinitely large books at interactive speed levels. The interface, which is similar to the 'Realistic Book' [30] implementations, supports activities such as bookmarking, text extraction and on-the-fly indexes.

In a later paper on the topic [62], the creators of 3Book describe how they have adapted their system to allow annotations, specifically highlights and scribbles, to be added to their three-dimensional representation of the book. As with other attempts at realistic three-dimensional book representations, the 3Book interface uses a polygonal mesh with painted textures to provide physical page-turning animations, a function that the authors describe as complicated to implement when the user annotations are complex. Although the problems associated with three-dimensional graphics are not applicable to this Thesis, the work conducted by Card et al. does offer several key comparisons to the work I have conducted. For example, the annotations made on the 3Book by users are intentionally stored separately to the pages themselves to provide the flexibility to access unannotated pages or annotations alone, a property that I found to be a digitally lightweight feature in my evaluation of several of my end implementations.

Even portable computing devices such as tablets and smart phones are now attempting to replicate the book metaphor in their software. Applications such as iBooks and Stanza provide 'paper-like' page curling/turning by means of gestured finger swipes across the touch screen displays of the hardware they adopt (see Figure 2.4). Despite the familiarity of these types of interfaces however, they still suffer from some serious usability issues which could hinder the digital active reading process,

suggesting that there is still work to be done in the area.

Although there has been successful research into the concept of realistic books, there are however, limitations to these literal book metaphor approaches. The turning of the page, whilst aesthetically pleasing, is not necessarily a positive contribution to usability. The response speed of programs such as these, due to the computational cost of rendering page turning in detail is often poor. Furthermore, the ability to slowly 'curl' over digital pages is typically incomplete as in the majority of cases, only one page can be turned at a time. In contrast, on paper users can curl multiple pages at once, allowing fast flicking through a document as well as the ability to quickly switch between two non-adjacent pages.

Prior research conducted on realistic electronic books presents a grounding for the work presented in this Thesis. Although for the most part my implementations do not religiously follow the book metaphor, that is, all but one of my interfaces use single sided display as opposed to double page spread, I do make a direct comparison to a physical book in my largest investigation into note-taking. Instead of slavishly duplicating every aspect of the physical book in my design however, I cautiously trim what I feel are inappropriate migrations from the physical interaction. For example, as I mentioned above, the visual curling of pages within an electronic document, despite sparking familiarity, does not translate well to the digital plane due to restrictions placed upon its functionality, as well as problems resulting from page rendering [62].

My aim here then, is not to force physical metaphors into digital book design, but rather to pick out useful aspects that transfer in a meaningful and practical way. It is important to note, that just because a property can be considered lightweight physically, does not necessarily mean its digital metaphor will also prove lightweight [57].

### 2.5.5 Hybrid Systems

In addition to purely digital solutions, many researchers have opted to combine the features from both the physical and computerised worlds into single amalgamated systems. By exploiting or capturing some of the affordances of paper, they are in effect 'bridging the gap' between the physical and digital domains in an aim to enhance the digital reading process. This section describes several systems whose goals are to *integrate* rather than *remove* paper from our electronic lives [64].

The notion of merging physical interactions with digital enhancements can take many forms. One way in which researchers have approached the issue, is by creating physical desk spaces which are then enhanced using electronic modifications. Wellner et al. [144] for example, introduced the DigitalDesk, a tangible physical desk with additional electronic interaction. In later papers [94, 145], Wellner et al. go on to describe the system, which provides an augmented environment that enhances the physical paper world with computerised characteristics that overcome some of paper's limitations. The authors focus on performing actions without thinking, using fingers, arms, 3D vision and kinaesthetic memory to manipulate objects simultaneously. The DigitalDesk therefore, is essentially a physical desk where users can lay out papers, highlighters, Post-its, coffee mugs and so on, but also includes some digital features. As well as the physical items on the desk, there is also a projected computer display, as well as a series of video cameras that are used to read the paper documents placed upon it and to listen for user interaction using pens or fingers.

Another early example of this type of interaction, is that of Arai et al., who developed InteractiveDESK [8, 9], a computerised desk with reality awareness. The system uses a large CRT display with a keyboard and pen-input and responds electronically to physical operations performed on the physical workspace. For example, when a user puts a physical object down on the InteractiveDESK, the system recognises it and links it to a list of electronic files on the computer.



Yet another example of how physical desks have been used to aid in the digital active reading process is Koike et al. [67, 68] who have produced an augmented desk system called EnhancedDesk. The EnhancedDesk provides an intelligent environment that retrieves and projects digital information directly onto objects on the desk. The interaction of the system is similar but more advanced than that of Wellner's DigitalDesk as it allows users to manipulate the projection information in real-time using their fingers or by means of infra-red cameras, pattern matching and pan-tilted cameras, an interaction style that is similar to picking up physical objects.

A more recent example of how physical desks have been integrated into a digital working environment, is WikiTui [150, 151], a system which combines traditional paper based reading with innovative technology by providing an overhead projection of digital media on top of a physical book and table surface. The system, which gives access to multimedia information tagged to specific locations within physical books, is useful for exchanging information collaboratively and also captures gestures via a video camera.

In addition to interactive physical desks, there has also been work in the area of how to connect paper pages with digital media. One such example of this, is the Intelligent Paper interface [38]. Dymetman et al., the creators of Intelligent Paper, argue that the relationship between paper and digital documents is not, as many assume, *competition*, but rather *complementarity*, and aim to combine the two mediums to enhance the reading process. The design for Intelligent Paper system therefore, proposes collaboration between physical and digital documents by providing communication between a printed page and a digital peripheral. This interaction is accomplished by means of a printed page-id on the physical sheet and a pointer input device that records coordinates when pressed to paper. Using a local network connection, a communication link is established between the paper/input device and an output peripheral positioned close to the user.

Liao et al. [72, 73] have also created a gesture-based command interface that facilitates digital document manipulation using paper print-outs as proxies. Their system, known as PapierCraft, uses a digital 'pen' that allows users to annotate and create command gestures (such as operation indications like copying and pasting areas) on printed documents, and synchronise them to a customised digital viewer. This type of pen-based interaction is useful for many users due to the familiarity of performing common gestures, while simultaneously providing additional support for digital enhancements.

Some commercial systems have also realised the benefits of augmenting paper with digital media. One example of such a device is the 'LeapPad'<sup>1</sup>: a interactive electronic children's book platform that features a touch sensitive tablet with a paper book overlay. This educational technology allows physical books to be read to the user via an interactive 'magic pen' that works like a hand-held mouse pointer. Despite the physical constraints of the hardware, specifically, the device must sit in a plastic housing unit, the device proved extremely popular amongst consumers with reported sales of \$680 million in 2003 alone [89]. In 2008, the 'LeapPad' models were superseded by the 'Tag Reading System', a newer model that no longer requires a docking base plate. Instead, the 'Tag' system uses optical recognition to read uniquely identified tiny dot patterns surrounding letters in specifically made books. By using this technology, children are able to run their 'tag' pen across a physical book while listening to the audio book being read digitally.

A more recent example of a hybrid reading system is that of Back et al. who have developed the Listen Reader [10], a system that aims to keep the affordances of paper while enhancing them by adding electronic augmentation. The backbone of the Listen Reader system is a real, physical bound book with paper pages and printed images for the look and feel of traditional books. Additional electronic augmentations are subsequently added to this physical prototype in order to enhance the

---

<sup>1</sup>Produced by LeapFrog Enterprises ([www.leapfrog.com](http://www.leapfrog.com)) from 1999 to 2008

user experience of reading. Some of the extra computerised features of the system include, a multi-layered interactive soundtrack, proximity sensors to detect hand positioning and page specific RFID tags to facilitate fast robust page identification.

The clear advantage of systems that combine paper with electronic augmentations, is that users need not completely convert to a fully digital reading implementation, but rather keep the familiarity of physical books while at the same time enhancing them with electronic functionality. Conversely however, many of these cross-bred implementations require specialised hardware in addition to a large amount of physical space to function correctly, thus making them impractical for general use, particularly when on-the-go. Despite their limitations, this type of hybrid interaction may well provide a stop-gap between the physical world and a fully digitalised solution.

Although the focus of this Thesis is to create fully digital reading implementations, the work conducted in this area does strengthen my argument that the physicality of paper is still an important part of many users' reading practices and one which must be addressed when creating digital alternatives. The work in this area therefore, has provided me with several key suggestions for the improvement of electronic reading interfaces, for example, the use of desk space surrounding the central document.

## 2.6 Summary

This chapter has described a set of concepts and literature that are central to the work described in this Thesis. My investigations into both paper based and on-screen reading have given me a clearer picture on the shortcomings of digital reading and the psychology behind it.

It is clear from the majority of the literature [36, 93, 107, 139] that paper is a faster and more popular medium than computer screens for the attentive reading of a document. The reasoning behind this preference is linked to the affordances offered by physical documents and tools. The lightweight properties of paper enable often sub-conscious actions that divert very little cognitive attention from the primary active reading task.

Paper boasts many intuitive and tangible assets that enable smooth and unperturbed interactions from its users. One way to mimic parts of these interactions is by a transfer of metaphors from the physical plane – a concept I follow closely throughout my design decisions within this Thesis.

One common way researchers in the area have made use of metaphors, is by mimicking the interactions with physical books. When a metaphor is created, it provides its function with a certain set of affordances. The book metaphor then, supplies particular affordances for functions such as page turning, placeholder, annotation and so on. The majority of the literature seen in the area of digital document design focuses strongly on the physical book metaphor by using it in some way when creating and improving digital reading interfaces. For instance, there has been a significant body of work that focuses on augmented reading hardware by transferring eDocuments to portable book-shaped reading devices and enhancing them with paper-like tangibility [59, 86, 132, 141].

Other researchers have opted for the solely software route by closely mimicking the visual interaction styles of paper in realistic electronic books [25, 30, 74]. Other common areas of investigation are concerned with the amalgamation of physical documents with digital enhancements [10, 72, 144], which typically involves keeping the physical tangibility of the real-world book while improving it with electronic features.

The digital tools I implement to aid in the digital active reading process will often be designed around lightweight physical attributes by transferring metaphors from paper. In some situations, I will also be utilising solely electronic properties to create the same digitally lightweight properties. In both cases however, the tools I design should strive to be lightweight, ready-to-hand and invisible to ensure

the maximum amount of cognitive attention can be paid to the primary task as opposed to the tools themselves. This will then reduce the possibility that the tools will distract the user out of a flow state with the primary active reading task. With this in mind, the main motivations for this Thesis are:

1. Prove by example that lightweight interaction is possible in digital documents;
2. Create systems that use lightweight interaction in their design to improve the active reading process;
3. Produce a list of attributes that can be considered digitally lightweight to aid in future designs of active reading software.

To establish whether or not an attribute can be considered lightweight, I have performed a series of systematic evaluations using a set of human participants as a base. However, determining exactly which properties fall under the lightweight category will not be as easy as simply monitoring the speed or efficiency of a user undertaking an active reading task.

The overall aim of the systems I have tested is to improve the active reading process by incorporating lightweight techniques into their design. Lightweight interaction is categorised as being subliminal or subconscious actions that do not force the user out of the flow of their main task. It can be reasonably assumed then, that lightweight actions are not dependent upon speed or efficiency, but rather a low cognitive demand which leads to minimal intrusion on the primary reading task.


Although improving the speed at which a function is performed may slightly improve the time spent on the main reading activity, it is unlikely that it will significantly reduce the cognitive load required for the function itself. In fact, in some cases, measuring the speed at which a task is performed may well give false-positive results. More specifically, it is likely that the more tacit the tools, the more the user can focus on the main task and perhaps as a consequence improve their performance, which may well increase the time they spend on the task.

It is for these reasons that I have not used time and/or efficiency as metrics in the majority of my studies. Instead, I have chosen to focus more on subjective opinions from participants, coupled with observed behaviour during open tasks using my implementations.

Clearly there are occasions where speed and efficiency *are* factors in a user's ability to perform tasks. One such example of this, is in Chapter 6 where I compare automatic index creation to the slower and less efficient text search function. This is an example of a study where I use both speed and accuracy as metrics in my study evaluation.

## Chapter 3

# Placeholders



Experience is simply the name we  
give our mistakes

---

Oscar Wilde

### 3.1 Introduction

This chapter focuses on an area of digital document manipulation that is currently heavyweight in its digital form, but can be considered lightweight on paper. The majority of free digital reading interfaces used today (e.g., Adobe's Acrobat Reader or Apple's Preview) contain only a handful of functions, one of which is the placeholder. Yet research in the area of digital reading (e.g., [85]) has demonstrated that this function is rated unfavourably compared to its physical equivalents. Given that placeholdering is considered important enough for inclusion in digital reading software with such a small set of functions, I felt it a good starting point for my investigation into lightweight attributes.

#### 3.1.1 Overview

Placeholders in physical documents are a long-established method of locating information and provide crucial support for readers in remembering important places in the text. Placeholders can take many different forms. Scrap paper, dog-eared corners, bookmarks or even fingers that are used as location aids [36] require minimal effort, making them a perfect example of lightweight navigation. Unfortunately however, the equivalent tools on digital documents are far less intuitive and are consequently a more time consuming and cumbersome affair. In fact, much of the literature on this topic confirms that digital placeholders in their current form are extremely poorly used despite recommendations from experts that they should be simple to achieve [147]. I have chosen to pursue this topic as it is a good example of an area which is 'lightweight' on paper yet 'heavyweight' electronically.

This chapter introduces a digital placeholdering system that has been designed to act more like bookmarks in physical books. The Visual Bookmarking System incorporates a system of coloured 'tabs' which replace the typical list structure adopted by most electronic reading systems. After a discussion on the design and implementation of the new system, I move on to describe a comparative user study conducted to gain subjective reviews from the participants.

## 3.2 Background

### 3.2.1 Document Revisitation

Reading is rarely entirely linear. There are many occasions, even when reading novels, when the reader may decide to revisit a previous section to clarify or refresh their perception of the text.

Several studies have touched on the importance of document revisitation. Marshall and Bly for example [85], documented magazine readers' usage of placeholders for re-reading, refreshing or double checking, whereas O'Hara [106] states the importance of re-reading in learning situations. There has also been research into the frequency of revisitations on digitised media. Greenberg and Witten [51] for example, studied the reuse of Unix command lines and found that three quarters of them actually exist in the history list, that is, that three out of four are in fact revisitations as opposed to new commands.

Alexander et al. [7] have also studied the use of within-document revisitation by analysing the logs of 14 participants' navigational actions in Microsoft Word and Adobe Reader over a period of 120 days. By defining a revisit as: "*returning to anywhere within the bounds of an earlier visit and remaining there for more than two seconds*", they conclude that an average of 10 items in a frequency or recency list of bookmarks covers approximately 80% of revisited document locations.

It is clear from the literature then, that the revisitation of pages within documents is a key reading activity, and one that should be carefully considered when making the transition from paper to digital reading.

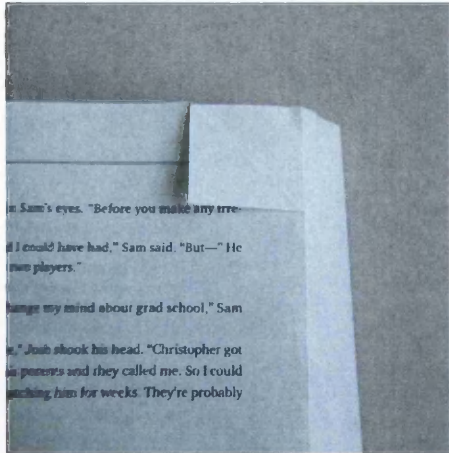
### 3.2.2 Placeholder Use

To aid in document revisitation, many users make use of placeholders which are essentially a means of marking a place within a document and can take many different forms. For example, ribbons or bookmarks in physical books or page markers in digital PDFs. All these different types of placeholder, can be considered to be either *temporary* or *permanent*:

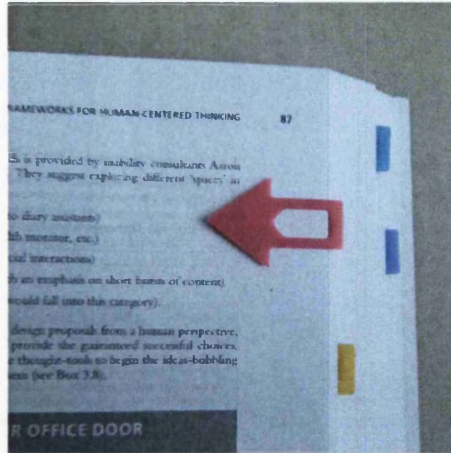
**Temporary Placeholders:** For example, a bookmark or a finger [36, 107] that is used temporarily to mark a place, then typically removed and placed elsewhere (i.e. when reading a novel, you may slip in a bookmark when you go and get a snack, then remove it again when you commence reading). A document is likely to have only one temporary placeholder within it which is placed solely to remember a particular place within the text.

**Permanent Placeholders:** For example, sticky post-its that can be used to mark pages that are looked up a lot (i.e., a chapter in a text book that contains an important piece of information). A document will often contain multiple permanent placeholders which are designed for long-term re-visitation of commonly used sections.


There are also techniques that could be considered to be *semi*-permanent placeholders. For example, folding down the corner of a page, or 'dog-earring' [112] as it is often referred, is a method of placeholdering that requires no external materials. It does however, cause slight damage to the book itself making it an undesirable long-term placeholder.




(a) 'Dog-eared' corner



(b) Arrow markers that can also point to lines in the text



(c) Book with built in ribbon marker



(d) See-through sticky note tab markers

Figure 3.1: Examples of physical placeholders

### 3.2.3 Physical Placeholders

There are many ways to hold a particular page within a paper document. Some examples of placeholders in physical documents can be seen in Figure 3.1.

Placeholders are extremely common in physical documents and are frequently employed without a significant amount of conscious effort from the user. As well as marking a place within a text however, physical placeholders can also provide other implicit information to the user, for example, sticking a Post-it half way down a particular page could indicate that the first half of the page has already been read. As Heim [57] stated:

*“Paper documents afford the user of place holders and allow us to make determinations of text length and location within the text.”*

– Heim [57]

### 3.2.4 Digital Placeholders

Digital placeholders are essentially links to specific points within text and typically take the form of a list. The most common forms of digital placeholders are: web-bookmarking (Figure 3.2a), digital document bookmarks (Figure 3.2b) and electronic reading device bookmarks (Figure 3.2c).

#### Web-Bookmarking

More commonly known as ‘favourites’, web bookmarks generally take the form of an un-ordered menu, that is, they usually exist within a drop down menu structure and are ordered according to the date in which they were added (see Figure 3.2a for an example of the Firefox web browser’s bookmark structure). Web bookmarks are not used to place marks within documents, but rather to keep an overall record of entire documents. Known problems with this method of placeholdering include: long, un-manageable bookmark lists and the inability to mark specific points within pages.

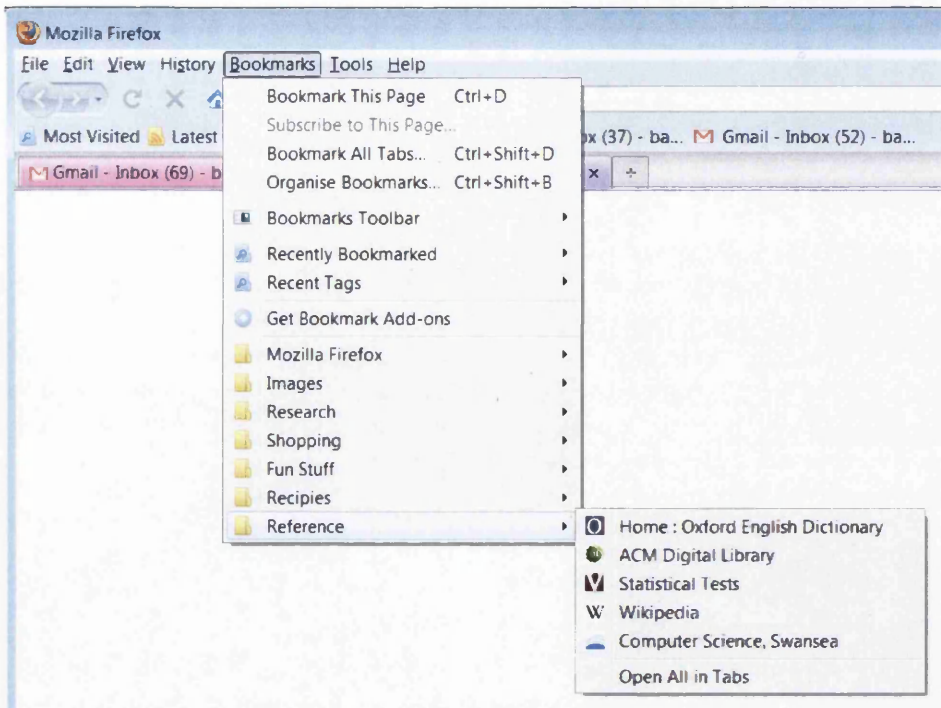
#### Digital Document Bookmarks

Placeholders within digital document readers usually take the form of an ordered list, that is, they usually exist in a tree list structure on the side of the interface and are ordered by page number (see Figure 3.2b for an example from Adobe Reader’s bookmark structure). One major problem with this method is the fact that many digital document readers do not allow user created bookmarks, only author created bookmarks that exist when the PDF is created, for example, chapter headings and so on. Adobe Acrobat is an example of such software; in the pay-to-use edition (Adobe Professional), users can create their own placeholders, yet in the free version (Adobe Reader) only author created bookmarks are available.

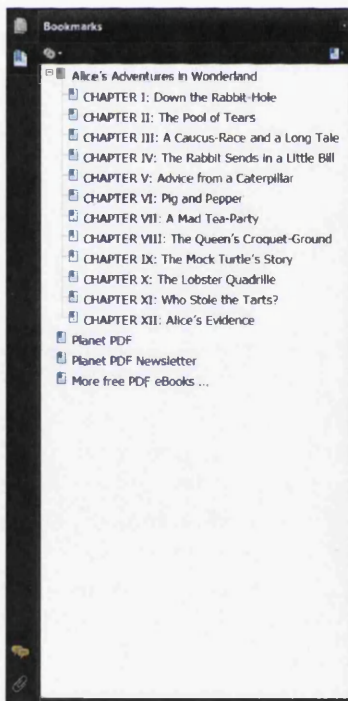
The way in which these structured lists behave is also drastically different to the way placeholders are used on paper resulting in the loss of many intuitive functions. For example, on paper, it is easy to see where a bookmark is in relation to the current page by simply observing which ones stick out of pages before and after it. Some digital document readers provide a function for this, (e.g., Apple’s Preview) highlights the appropriate bookmark when a page is changed, whereas others (e.g., Adobe’s Acrobat), only highlight the current bookmark when a user explicitly clicks in the bookmark list. Other useful information that has not migrated from paper to digital bookmarking, is the ability to see the spread of where bookmarks exist within the document. For example, on paper, it would be easy to see if a large document had all its bookmarks concentrated in the first 20 pages, whereas the majority of digital readers bookmark lists do not contain this information.

#### Electronic Reading Device Bookmarks

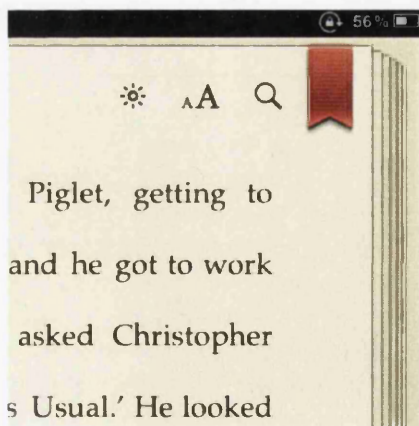
As well as desktop document readers, it is also important to consider the bookmark functions included within dedicated reading devices such as eReaders. The majority of eReader devices indicate bookmarked pages by displaying a small ‘dog-ear’ or ribbon in the top right corner of the reading area (Figure 3.2c). Even reading applications designed for portable computing devices (e.g., iPhones or iPads) use the ‘dog-eared’ or ribbon method for distinguishing between bookmarked pages. The major drawback of this method however, is that you must be on a page to know that it is bookmarked; there is no other way to tell where bookmarks exist whilst reading a document. In order to view all bookmarks within a document on one of these devices, the user would usually have to navigate away from the document itself, and then find the bookmarks list. These bookmark lists are almost identical to those found on traditional desktop document readers and therefore give little overview of where



(a) The Firefox web browser bookmark menu



(b) Adobe Reader's bookmark list



(c) The bookmark feature on the iPad's iBook app

Figure 3.2: Examples of digital placeholders



each bookmark is in relation to others in the document. These types of placeholder can be considered permanent as the reader must explicitly tell the device if they wish to remove one. In addition to these dedicated bookmark functions, many eReader devices also record the last place the reader was at before exiting, a function that can be characterised as a temporary placeholder.

### 3.2.5 Aims

The motivation for this research was to take something that is currently lightweight on physical documents but heavyweight digitally, and improve the digital version by imitating the way it is performed on paper. I felt this was a good starting point for my investigation.

It would have been virtually impossible for me to replicate the exact affordances of paper on a digital desktop solution. Therefore, I had to think strategically about which aspects of the system were easily replicated as well as those that were potentially lightweight. With this in mind, I laid down four aims to which the end system should be adhered to:

**Aim 1:** To create a visual system that looks like bookmarks are sticking out of the side of the document;

**Aim 2:** Bookmarks should always be visible no matter what page of the document the user is reading;

**Aim 3:** Users should be able to see, at a glance, where bookmarks are in relation to the current page. (i.e., which ones come before it and which ones come after it);

**Aim 4:** Bookmarks should provide ‘one click’ navigation to the pages they are referencing;

## 3.3 Literature Review

The following section reviews the major literature on bookmarks in two types of digital placeholders: web favourites and digital document bookmarking. Although I am not directly focusing on the web format of digital placeholdering in this section, I have briefly discussed the major literature on the topic to aid in the production of an improved digital bookmarking system.

### 3.3.1 Web Favourites

Research into the observed usage of web page bookmarks or ‘favourites’ as they are often referred, has been a popular topic since early on in their development. It has been well documented, that web bookmarking is more commonly used as a long-term archival aid as opposed to a method of relocating well-known material. Abrams et al. [1] for example, discovered that frequency of use was not a factor in determining whether or not a page was bookmarked, with users commonly choosing to memorise URLs or use search engines to locate commonly visited sites.

The infrequent use of web bookmarks for revisitation was also investigated by Tauschær and Greenberg [133] in 1997. The pair, who focus on *how* people revisit web pages, define the ‘recurrence rate’ of a page as the probability that a visited page is a repeat of a previous page visit (expressed as a percentage). With a revisitation result of 58% from their user study, they concluded that although the web is a ‘recurrent system’, the use of bookmarks as a revisitation tool was very limited.

Cockburn et al. [31] concur with these findings after a study into the Netscape history and bookmark files of a set of users. They discovered that 81% of their observed users’ page visits were revisitations and that most people have large web bookmark lists. They suggested however, that very

few users make constant use of the bookmarks they have, and furthermore, that the lists tend to grow over time but are rarely maintained, leading to sometimes overwhelmingly large and unmanageable bookmark lists. Realising this problem, some researchers have investigated means to provide suitable organisation for web favourites. WebTagger™ [66] for example provides a web bookmarking service for individuals and groups that provides a customisable means to store, access and rate web resources.

More recently, Weinreich et al. [142] and Tabard et al. [129] have studied the area again and reaffirmed the conclusion that web-bookmarks are poorly used and do not sufficiently support web page revisitation.

The key result of over a decade of research in this area therefore, concludes that the majority of web bookmarks are used for archival purposes, (i.e., when users want a permanent record of a particular page) as opposed to a general form of revisitation.

### 3.3.2 Digital Document Bookmarks

A web page can be considered as a non-paginated document, or a set of linked documents. In either case however, ‘bookmarking’ a web page will result in a placeholder to the document *as a whole* as opposed to a specific portion. Clearly then, the concept of web bookmarking does not provide the same functionality as classic placeholders do in physical books. In contrast to web favourites, paginated document readers allow bookmarks to be made on specific portions of a document, typically pages, facilitating, in effect, sub-document bookmarking, the type found in physical books.

Surprisingly, given their popularity on the physical plane, within-document bookmarking on paginated digital texts has been little studied. Marshall et al. [85] touched upon this topic briefly while making a comparison between user behaviours on physical and digital texts. The results from this portion of the study concluded that users rate digital bookmarking tools harshly compared to their physical equivalents and also make less use of them. Other studies of electronic document usage have also touched upon the lack of bookmarks within digital texts, but often with less direct evidence [149].

Several research groups have recognised that current electronic bookmarking tools bear no resemblance to their physical counterparts, and have attempted to make digital bookmarks more paper-like in order to bridge the gap and increase user uptake of the tools available. For example, the work conducted on ‘realistic electronic books’ [77] contain bookmarks similar to paper, that is, tabs that stick out of the side of pages. In an earlier paper [76], the authors performed a study on their 3D book visualisation comparing it to more conventional book formats, in this case, HTML, PDF (Adobe Reader in this case) and also physical documents. The bookmark portion of this particular study discovered that many users did not realise that a bookmark view existed in Adobe, as it was ‘hidden’, whereas the bookmark ‘tabs’ within their system were easy to see. Furthermore, they confirmed that many users did not make use of table of contents links within the realistic system as the bookmark tabs did it for them. Other implementations that make literal use of the book metaphor which also include physical style bookmark ‘tabs’ include the ones described by Card et al. [24, 25].

Alexander et al. [7] have created the Footprints Scrollbar – a system that aims to support within-document revisitation using marks in the scrollbar. Their implementation makes use of six inter-related techniques to aid in the process of revisiting sections of a document, including: coloured marks, animated scroll transitions, thumbnail image mouse-overs, back and forward functions, jumps to closest marks and go-to numbered markers. These within-document marks are not user-created bookmarks per se, as they are made automatically by the program if a scrolled region remains static for more than two seconds. Despite this however, it is a clear example of how the space surrounding the document area (in this case the scrollbar) can be modified to give static access to marker information

independent of the current position of the document.

Steimle et al. [126] go even further than simply mimicking paper bookmarks on the digital plane by incorporating a combination of physical adhesive stickers with electronic processing. Their prototype, 'Digital Paper Bookmarks' aims to combine the intuitiveness of paper bookmark interaction with the sharing and computational abilities of a synchronized electronic counterpart.

Other researchers have also focused on kinesthetic controls for digital bookmarking. Yoon et al. [152] for example, make use of a capacitive-type multi-touch display that allows users to hold 'touch points' – a bookmark interaction that is designed to mimic the temporary nature of fingers in real books. That is, when a touch point is held with a finger while a page is turned, the pages are 'stacked' on the marked page for revisitation later. If however, the touch is released, the marked page disappears, thus providing a temporary bookmarking facility which they claim does not impede the reading flow.

LiquidText [130, 132] also offers a simple transient bookmark function within its interface which allows user-defined touch-screen finger gestures to temporarily mark locations within its non-paginated document. In a similar way to Yoon et al. [152], LiquidText's ephemeral bookmarking function mimics the way fingers are used on paper and only holds a bookmark in place for as long as the user's finger is in place.

The Booksheet [141] interface is another research programme that offers a finger bookmark function, in this case by making use of two micro switches located near the forefinger positions. To temporarily mark a page using this finger function, the user must continually press a switch while at the same time scrolling through remaining pages. Releasing the switch is essentially like removing a finger placeholder from within the pages of a real book. TouchMark [146] also facilitates bookmarking via modifications to existing tablet hardware, in this case by allowing users to squeeze a bespoke navigation tab on the side of the device.

## 3.4 System Design

To mimic the way standard placeholders are used in books, I have designed the Visual Bookmarking System, an interface that uses coloured 'tabs' that stick out of the sides of a document. These 'tabs' are positioned in the virtual space located on either the right or left hand side of the PDF viewing area. The design of this representation has been modelled in the style of a telephone directory, that is, bookmarks that occur before the current page appear on the left of the PDF area, whereas bookmarks on the current or later pages appear on the right (see figure 3.3). When the current page is changed, these tabs then 'flip' from one side to the other depending upon where they exist in relation to the open page.

To make the relocation process easier for the user, the tabs are also ordered by page number, so the later the page, the further down the display the tab appears. Thus, a bookmark on page 1 will always appear at the top left corner of the display, whereas a bookmark on the last page will be seen on the bottom right. Figure 3.4 shows the design for the Visual Bookmarking System.

### 3.4.1 Tabs

To make the system more paper-like, I eradicated the standard list bookmark format and replaced it with a visual 'tab' system. These bookmark 'tabs' were designed to replicate bookmarks sticking out of the side of physical books. Due to the placement of these tabs (i.e., outside the PDF area), they are always visible independent of the currently open page. This ensures that all bookmarks can be




Figure 3.3: Example of telephone directory tabs

seen at all times, as opposed to the ‘dog ear’ type interfaces that are incorporated into devices such as eReaders.

Although it is not vital, it would be useful for users to be able to visually distinguish between each of the bookmarks at a glance. One option for this would be to make the tabs square and add text to each of them (see Figure 3.5). The longer the text is however, the less intuitive the system becomes; either the tabs need to become bigger, the font size needs to be reduced, or the text length needs to be reduced – all undesirable solutions.

To eliminate these problems, I opted not to add text to the bookmark tabs, but instead replace them with mouse-overs that give the same information. By doing so it was possible to change the shape of the tabs – circular for a less cluttered and aesthetically pleasing interface. To complement the pop-ups that appear when a user moves their mouse over a tab, the bookmarks can also be assigned a colour which can be used to distinguish between different pages (i.e., meaningfully, for example, red for important pages, or arbitrarily as the user chooses).

### 3.4.2 Ordering

As I have discussed earlier, the format of web bookmarks take the general form of an un-ordered list, that is, the bookmarks appear in the list in the order in which they were added. Although this format is acceptable for a bookmarking a collection of unrelated web pages, it would be undesirable in a document ordered by page number as it would give little contextual overview of where they exist in relation to the current page, as well as in relation to each other. In contrast, paginated document readers, being sequential, order their bookmarks in a list according to where they appear in the document, that is, a list ordered by page number.

These lists of page ordered bookmarks, despite differing from their paper equivalents, allow users to see, at a glance, which ones come before others in the document. What these lists do *not* show

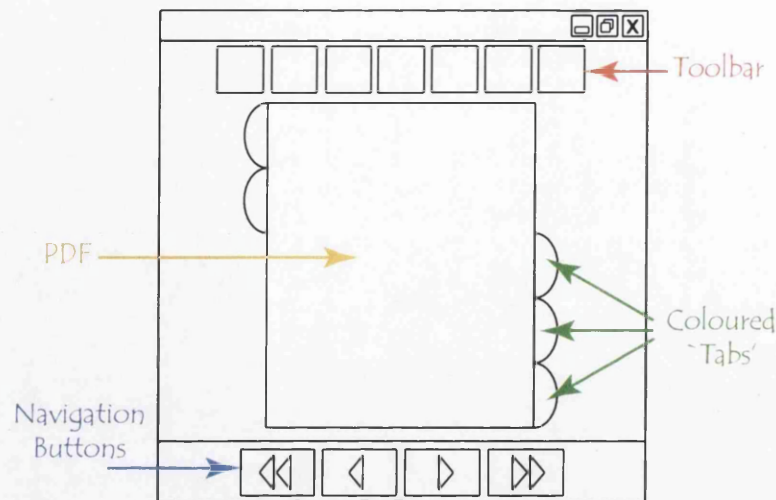






Figure 3.4: Design for the Visual Bookmarking System

however, is where they exist in relation to the currently open page. When reading a physical document, it is easy to see which placeholders occur on pages *before* the page you are reading as well as which ones occur on pages *after* it.

To include this functionality into the new system, each page ordered bookmark 'tab' will 'flip' with the page it is related to: that is, if a bookmark is on a currently open page or any pages *before* the current page it will be 'stuck' to the **top left** side, but if the bookmark is on any pages *after* the current page it will 'flip' to the **bottom right** side. This functionality ensures that users can tell which bookmarks come before and after the page they are reading, as well as the sequential order in which they exist. Figure 3.4 shows the design for the bookmark 'tabs' on the new system.

### 3.4.3 Form Layout and Navigation

Since placeholders are traditionally bound by *pages* and not by paragraphs or sentences, I have chosen *single page display* pagination (see Chapter 2.5.1) for the Visual Bookmarking System. This type of page display will ensure that whole pages will always be visible when a bookmark is clicked. Continuous displays were not a suitable option for this system as it would be difficult to illustrate where bookmarks that exist on previous and future pages. Since it was my intention to make a more paper-like bookmark system, it would have been less beneficial to include a pagination style that exists only on electronic documents as it may confuse users.

Page to page navigation within the system is accomplished by clicking on the 'next' , 'prev' , 'first'  and 'last'  arrows. The basic layout of the Visual Bookmark System is shown in Figure 3.4.

#### Alternative Designs

There are viable alternatives to the design that I finally implemented. One obvious approach would have been a derivative of that of Byrd [21] using small markers on the scrollbar to identify the location of bookmarks, where he instead had marked search terms. However, the use of small dots (no more than 2x2 pixels) seemed too small a cue to the user. A further enhancement of that approach would




Figure 3.5: An alternative to the round, blank tabs used in the design of the Visual Bookmarking System

have been to use a horizontal line to indicate each place - but this would result in two bookmarks in close proximity to each other having little or no distance between them. As a result, discriminating between the two lines when clicking the mouse would be difficult and distracting. My simple metaphor avoids this particular difficulty. The Footprints scrollbar of Alexander [7] takes a similar approach, using small thumbnails, but this metaphor emerged a year after I implemented this system.

Given that I was building on some observational work on users' behaviours on paper, I naturally wished to follow a more paper-like metaphor. Like Byrd, Liesaputra et al. [78] sought to indicate the location of search results in a document. Their Realistic Books tool used tabs that jutted out from the side of the document display. My design is similar, but provides larger tab sizes and maximises the size of each tab, considering both the display size and number of tabs. In contrast, the Realistic Books tool consistently uses small tabs that require greater accuracy with mouse clicks. Like Liesaputra, I could permit several columns of tabs when the number of bookmarks becomes large, but this was not necessary in the scope of my study.

Finally, I could have followed the metaphor of dog-earing pages. This, however, would potentially suffer from the mouse click accuracy problem, as a lot of information would be communicated in a very small part of the display, and the user would be required to click within a high precision. This would itself increase the demands on their attention, taking it away from the main text.

### 3.5 System Implementation

Figure 3.6 shows a screen shot from the Visual Bookmarking interface and illustrates the round colourful 'tabs' designed to mimic physical bookmarks. The physical size of each of these bookmark 'tabs' is dependent upon how many of them there are; the height of each tab is calculated by dividing the height of the PDF viewing area by the number of bookmarks in the document. Creating a new bookmark and similarly deleting an existing bookmark will resize and reorder every tab on the screen.

The aim of the visual bookmarking system is to allow users to quickly and easily relocate information in a document. This interaction has been achieved by a visual representation that mimics

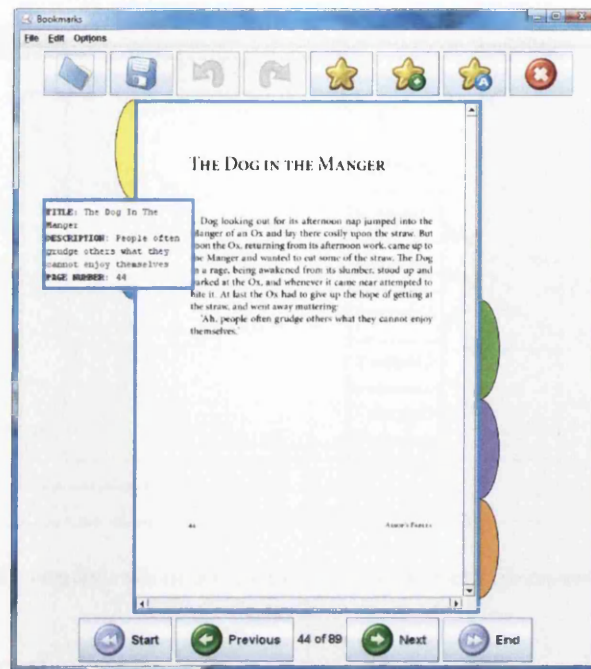


Figure 3.6: The Visual Bookmarking System

methods used in the physical domain (specifically utilising the ‘virtual’ space surrounding the PDF itself). Each bookmark in the system can be distinguished by its colour as well as its position in the display and clicking on a tab will take the user to the bookmarked page.

The implementation was written in Java using the open source third party software JPedal (see Section A.2.1) and is deployable on Windows, Macintosh and Linux machines. The following subsections describe the features and implementation methods of the Visual Bookmarking System.

### 3.5.1 Bookmarks

Each bookmark is created as a separate Java object that has several attributes associated with it and are stored in an array list which is used by the program to automatically create and position the bookmark ‘tabs’ on the form. Refreshing of these bookmark tabs is performed whenever a bookmark is created, deleted or edited, when the current page is changed or when an undo or redo action is performed.

When a new bookmark is created, the user has the option of editing each of these attributes (see Figure 3.7a) using the New Bookmark menu. Alternatively, the ‘Auto Bookmark’ function can be used which enters generic ‘Untitled’ text into the title and description and selects an arbitrary unused bookmark colour. Existing bookmarks can be edited in the ‘View Bookmarks’ menu (see Figure 3.7b) that also displays an ordered list of every bookmark in the document.

### 3.5.2 Saving and Opening

The program has a saving feature that allows users to save their bookmarks. This bookmark data is recorded independently of the PDF in a single file which saves every tab along with the location of the PDF on the hard drive and also the file size (in bytes) of the PDF. This allows the program to make educated choices about multiple copies of the same file. For example, if the user attempts to

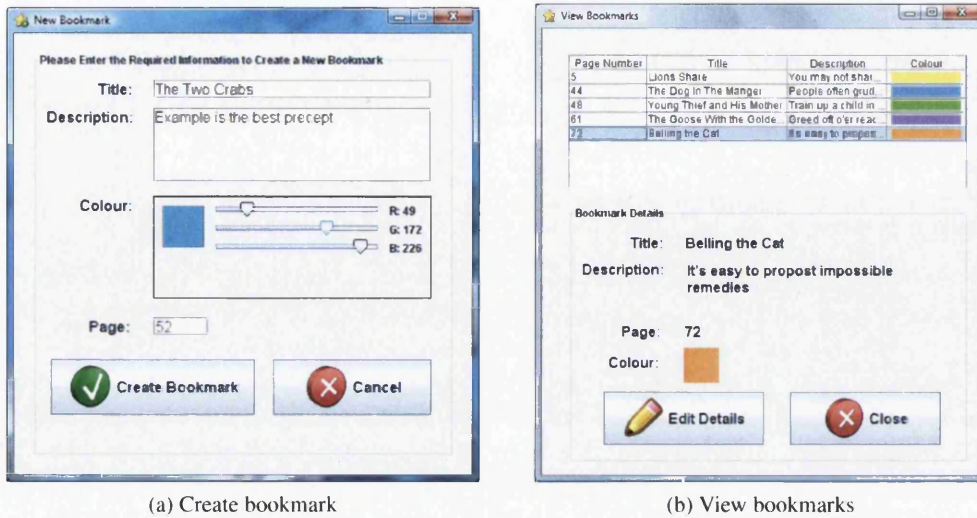


Figure 3.7: The create and view bookmark forms

open a copy of a PDF from a removable flash drive, the program will first check if bookmark data from this exact file and from this exact location has been saved previously. Failing that, the program automatically compares the local filename and file size against every entry in the Library file, if a match is found it then informs the user and invites them to link the PDF with the bookmark data.

### 3.5.3 Mouse-overs

It would be clumsy to allow text to be added to the bookmarks themselves, as word-wrapping etc., would cause them to look cluttered. Instead, to give an overview of the attributes associated with a particular bookmark, I have added mouse-over pop-ups that display the title, description, page number and appropriate border colour of each placeholder (see Figure 3.8). These pop-ups appear when the mouse is hovered over a tab and provide the user with quick and easy access to each placeholders information.

While mouse-over pop-ups have limitations, for example, they obscure content that lies behind them, this fleeting occlusion is more desirable than a permanent obscuration of the main text. An alternative approach would be to present the annotative text in a separate, static location, much like hovering over a link in a web-browser triggers the link's URL in the status bar. However, such an approach would itself suffer the problem that the location of the content would likely be some distance from the user's current visual focus of attention. Such a distraction is undesirable as my intention is to reduce distraction away from the primary task.

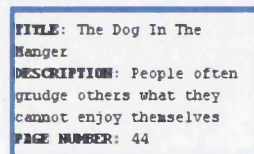


Figure 3.8: An example of a bookmark mouse-over



## 3.6 User Study

### 3.6.1 Study Design

To investigate the usability and overall value of the Visual Bookmarking System, a small comparative user study was performed to gain subjective reviews from a range of users.

#### Control Systems

To aid in the evaluation of the bookmark tabs, I implemented two control systems based on common methods of placeholdering available to date. These additional interfaces were then used to gain comparative subjective ratings from the participants in the post-study interview.

**Un-ordered Menu:** Based on the interaction traditionally found in web browsers where the bookmarks are displayed in a drop-down menu in the order in which they were added to the list (see Figure 3.9a for the control system used in the study and Figure 3.2a for an example of the real-life version it was based on).

**Ordered List:** Based on the current methods of placeholdering found in paginated document readers such as Adobe Acrobat. In this method, bookmarks are ordered by page number and are located in a side panel on the left hand side of the document display (see Figure 3.9b for the control system used in the study and Figure 3.2b for an example of the real-life version it was based on).

The three systems used for the comparison study contain the same basic interface and functionality including navigation, highlighting, saving, opening, undo and redo. The differences between the three systems were restricted to the bookmarking methods only and were identical in all other respects. To reduce unnecessary outside variables such as the differences between the participants' motor/spatial abilities, the navigation of both control systems have taken the single page paginated view to mimic the original design.

I chose not to compare the Visual Bookmarking System to commercial software such as Firefox or Adobe as I felt this would introduce many additional and unwanted variables into the experiment. For example, I did not want users to have any prior experience in using one system over the other (i.e., causing adverse learning effects), I wanted to keep the core functionalities of each system the same for consistency and I wanted to minimise the impact of different underlying technical issues (e.g., rendering speed or quality).

#### Procedure

The participant base for this study consisted of staff and students in the Computer Science department of Swansea University. The 13 participants I chose for the study were all aged between 20 and 35 and had a general or advanced level of computing ability. No incentive was given to the participants to take part in the experiment. Each study took on average between 30 and 60 minutes to complete and all comprised of the same format:

1. A short training tutorial on each of the three systems was provided to the participants to provide them with a basic understanding of each system's function and operation;
2. A semi-structured interview which I undertook to gain information on each participant's current web, printed and electronic bookmarking habits;

3. A set of tasks on the three systems which were undertaken based on a Latin-square design to minimise ordering effects. Each participant performed nine separate tasks on a pair of PDFs on each of the three implemented systems. The tasks themselves consisted of creating, editing, deleting and navigation of bookmarks;
4. Finally, I conducted a second interview in an attempt to differentiate between the three digital placeholding methods by gaining subjective reviews from the participants.

### Study Tasks and Documents

The tasks each user was asked to perform were intentionally short and simple, designed specifically to give the participants a chance to interact with each of the systems. Each participant was asked to perform two sets of five tasks on each interface. These tasks, which were spread between two PDFs, were of a closed nature and were specific to the PDF used.

The documents used in the study consisted of six novels (Treasure Island, Alice's Adventures in Wonderland, Jungle Book, A Christmas Carol, Aesop's Fables and Heart of Darkness) of varying length (average 149 pages) but all retained the same general format, resolution and font size.

These six PDFs were used to allow rotation amongst the systems and ensure the experiment was a product of the interface rather than the tasks. To avoid learning effects, I also varied the order in which the interfaces were presented to the participants.

The tasks consisted of: creating, deleting and editing bookmarks, and recalling bookmarks by both their name and the page number. These needed to be made specific to each document due to varying content and length. Example tasks given to the users in the study included: **'Go to the Bookmark entitled *Mad Tea Party*'**, **'Create a Bookmark on page 130'** and **'Create a new Bookmark with the title *My Bookmark on a non Bookmarked Page*'**.

For each of the three interfaces, participants were asked to rate the design out of ten against three criteria immediately after using it: ease of use, easy to learn to use, and the usefulness of its features. All participants were also asked to state which system they preferred in the post-study interview.

The tasks given to the participants were designed for the sole purpose of acquainting the participants with the system. Since the data used for analysis was solely subjective (e.g., ratings on the Visual Bookmarking System compared to the control systems) I collected no further data from the study tasks.

### Interview Procedure

I elicited the participants about their current use of electronic and paper bookmarks in the pre-study interview. I asked how often they used web bookmarks, and what strategies they used to manage them. I also asked how frequently they used bookmarks on paper, what medium they used (e.g. envelopes, post-its, etc.) and why they made that choice. The frequency with which they used PDF bookmarks was also addressed. They were finally asked to describe their experiences of the differences in bookmarking in the three media (web, paper and PDF), and which format they preferred to use.

The post-study interview was focussed on the participant's overall preferences between the three interfaces, and to elicit comments and explanations for that choice. The opportunity was also taken to obtain any further information to clarify their responses to the pre-study interview.

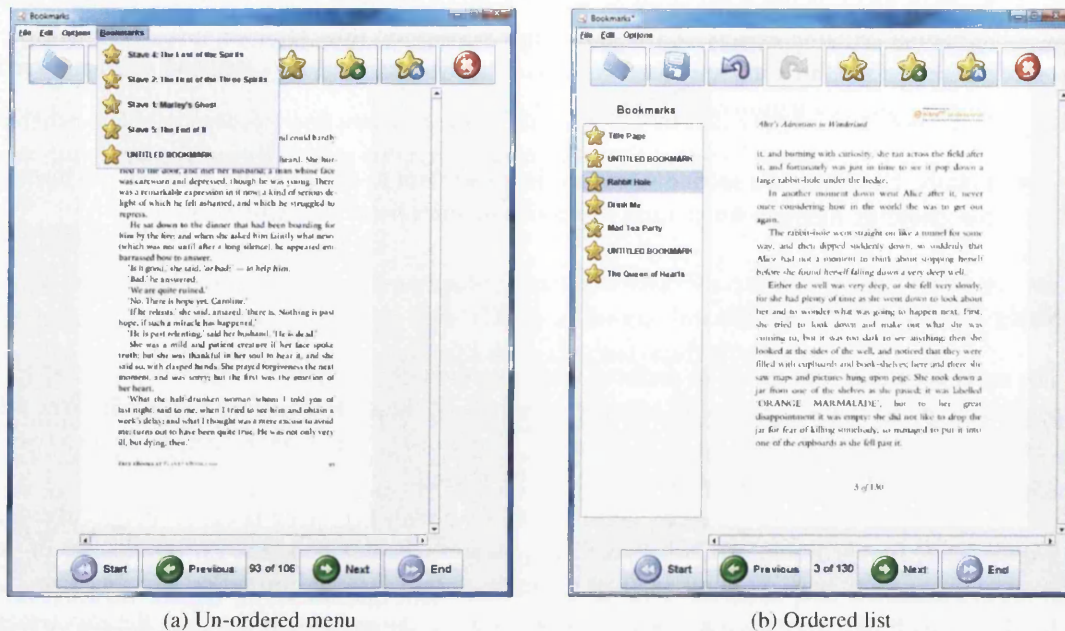


Figure 3.9: The two control systems created for the user study

### 3.6.2 Hypotheses

**Hypothesis 1:** The visual tab system will be the preferred method of placeholder than the more traditional control systems and will therefore yield higher subjective scores;

**Hypothesis 2:** The un-ordered menu will yield the lowest subjective user scores based on its low visibility (hidden within a menu) as well as the difficulties faced with ordering.

### 3.6.3 Study Results

This section describes the results from the study I conducted to assess the lightweight properties of the Visual Bookmarking System.

#### Placeholder Usage

How often people use bookmarks is often dependent on what type of document they are using. On average, the results for placeholder use on printed documents take the form shown in Figure 3.10 with over two thirds of the participants using placeholders either frequently (about half the time) or very frequently (almost all documents). However, several of the participants noted that their placeholder behaviour differed significantly between novels and reference books. Novels for example, typically favour a single temporary placeholder (i.e., a bookmark which is removed and replaced in a different spot every time the book is read) whereas reference books tend to include multiple more long term ones (i.e., fluorescent tabs to mark useful sections). These differences in placeholder use often dictate how often they are used, one participant stated *“I never use bookmarks in novels and hardly ever in reference books but I use them in my bible all the time”*, while another said *“I usually remember the page number in novels, but I do use them in text books quite a lot”*. For the purposes of this study, the

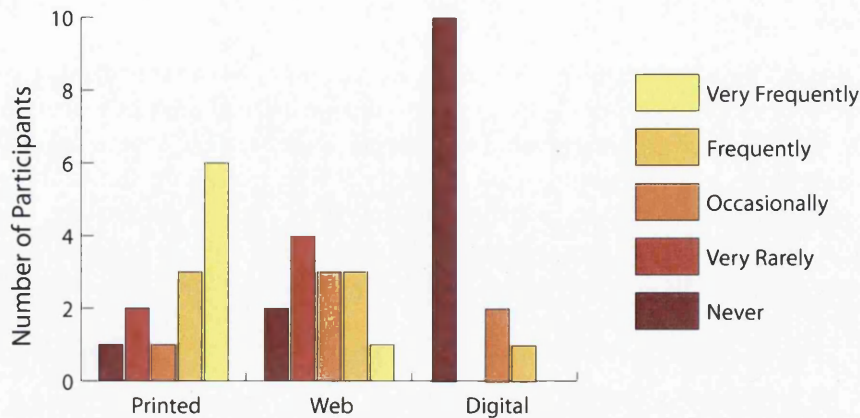


Figure 3.10: The frequency of placeholder use on printed documents, web documents and digital documents

participants were asked to give an approximate average of how often they use placeholders in general which yielded the results shown in Figure 3.10. This information was gathered during the pre-study interview where I asked users to select one option from a list of five possibilities: never, very rarely (less than one in 20 documents), occasionally (less than one in 10 documents), Frequently (about half the time) very frequently (almost all documents).

As Figure 3.10 shows, the differences between placeholder usage for printed and digital documents are substantial. The study concluded that six of the participants make very frequent (almost all documents) use of printed placeholders. In contrast, not a single participant makes use of digital placeholders this frequently. In fact, ten of the participants said that they have never used a placeholder in a digital document compared to only one on printed documents.

The results for how often the participants made use of web-bookmarks in their general lives produced a mixed set of reactions. Unlike the paper and electronic usage results where the majority vote was for an extreme case (i.e., very frequently in the case of paper and never in the case of digital), the web results yielded three quarters of its values within the middle range (i.e., frequently, occasionally and very rarely).

This data confirms that of the three placeholder methods tested, paper is the most frequently used, with web bookmarking coming in second and digital placeholder trailing with over three quarters of the participants never using them. Further questioning of the participants into why they prefer paper placeholders over digital versions uncovered some interesting opinions. One participant noted that *“paper bookmarks are easier and quicker to put in and move about”* whereas another commented *“books are more tangible which makes them more approachable”*.

One reason for the slow uptake of placeholders in digital documents is the possibility that they are not needed, with one participant pointing out that *“paper documents don’t have a search facility so you rely more on bookmarks”*. Although this particular user is noting a considerable advantage of digital documents (i.e., search) the flexibility and tangibility of paper documents remained dominant in the results with one participant commenting: *“paper is more physical and you can scribble your own notes on it”*.

### Placeholder Types

Part of the study was to assess how the participants currently make use of physical placeholders. To obtain this information, I posed a question in the semi-structured pre-study interview which asked “**What do you use to mark your place in a printed document?**”. The participants were allowed to select multiple answers from a set of four examples (loose paper, dog-eared corners, sticky Post-it notes and bookmark) as well as an option to select and specify others. The results from this portion of the study are shown in Figure 3.11 and conclude that loose paper is the preferred method with 7 out of the 13 participants selecting it as one of their commonly used placeholder techniques.




Figure 3.11: How the participant’s mark their place in a printed document

### Subjective Reviews

After completing a set of tasks on each of the three implementations, each participant was asked to rate (out of 10 where 1 is the lowest and 10 is the highest) each system for: how easy it was to use, how easy it was to *learn* to use, and the usefulness of its features. To assess statistical significance, I conducted a Kruskal Wallis test on the subjective data results.

The results of this portion of the study are illustrated in the graph in Figure 3.12. This qualitative feedback confirms Hypothesis 1: The visual tab system will be the preferred method of placeholdering than the more traditional control systems and as a result, yielded higher subjective scores. This is confirmed by the highly significant differences in scores for the **ease of use** ( $p < 0.0001$ ,  $H = 18.13$ ,  $df = 2$ ) and **usefulness of features** ( $p < 0.0001$ ,  $H = 23.43$ ,  $df = 2$ ) questions. Unsurprisingly given the fact that the two control systems were modelled on common bookmarking techniques, the **how easy each system was to learn to use** did not yield significant subjective results.

These participant ratings also aided me in confirming Hypothesis 2: The un-ordered menu will yield the lowest subjective user scores based on its low visibility (hidden within a menu) as well as the difficulties faced with ordering. This is clear from the results with the un-ordered menu achieving the

lowest subjective scores across the board. In fact, both the control systems performed poorly in the user evaluations with participants describing them as “confusing”, “clumsy”, “un-functional” and “rubbish”.

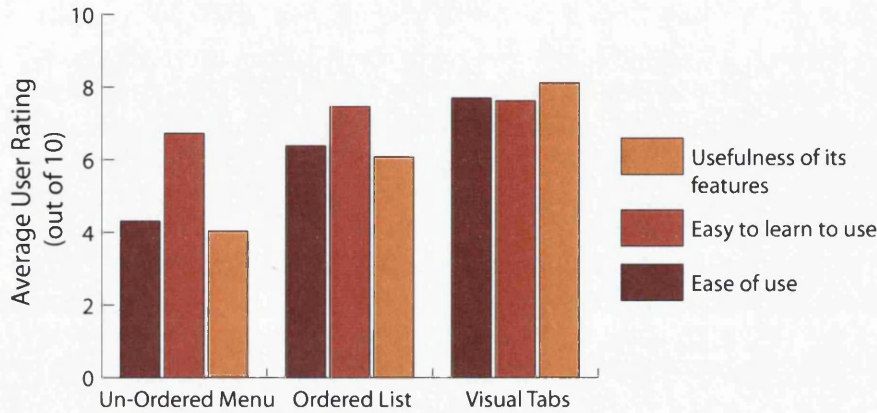


Figure 3.12: Average user ratings (out of 10)

As well as giving usability ratings for each implementation, the participants were asked which system they would prefer to use given the choice. Figure 3.13 shows the results of this investigation. As well as the three user study systems (un-ordered menu, ordered list and visual tabs) users were given the option to vote for a hybrid system which consists of the visual tabs combined with an ordered list. It is clear from this graph that the Visual Bookmarking System is the overall preferred choice amongst the participants with the hybrid option (which also includes the visual tabs) as the second. This is backed up by participant comments such as “I am a visual person so seeing rather than reading is good”, “it’s easy to find bookmarks because I can see them sticking out” and “this is a book metaphor that actually works”.



Figure 3.13: Users' preferred system

### 3.7 Conclusions

In summary, placeholder on physical documents can be considered very lightweight; adding and removing bookmarks to paper documents can be so subconscious that users are unlikely to think about how they are doing it. In contrast however, current digital placeholders are significantly less intuitive and are consequently rated unfavourably against their physical equivalents. Although some systems (e.g., eReaders) make use of ‘dog eared’ visualisations to signify bookmarked pages, the classical

bookmark ‘list’ structure found in the majority of digital document readers is vastly different from paper, making lightweight actions difficult.

To bridge the gap between the physical and digital forms of placeholdering, I implemented a system that constructs digital bookmark ‘tabs’ that mimic the way bookmarks are presented in physical books. These ‘tabs’ were designed to give a visual representation of placeholders within a document. This feature ensures that the placeholders are *always* visible, as opposed to only being seen while on a bookmarked page as on interfaces such as eReaders (see Section 3.2.4). This functionality could not have been achieved without additional space surrounding the document area – an idea that is little used within the digital library world.

The results of the user study conducted on the system produced promising subjective results from the participants. The cohort of users chosen for the study openly admit to making little to no use of web and digital placeholders and rate the visual system highly in comparison (Figure 3.12). As well as confirming the effectiveness of the visual system over more traditional placeholdering methods, the study also proved that the visual system or a hybrid system including the visual tabs and ordered list were the preferred methods of placeholdering for all but one of the participants.

### 3.7.1 Lightweight Properties

After implementing the system and conducting a user study to evaluate it, several lightweight properties have come to light. Firstly, it is clear from the subjective user study results that the visual tabs are the preferred way of displaying placeholders digitally. One key reason for this popularity as described by several of the participants, is the visual aspect of the design. That is, the shape, position and colour of the bookmark tabs.

Being placed outside the PDF area, the tabs are *always* visible, ensuring they can be seen regardless of the currently open page. Unlike many placeholder tools that display their bookmarks in drop down lists (e.g., web browsers) or those that indicate bookmarked pages by icons on the pages themselves (e.g., eReader devices), the tabs in the Visual Bookmarking System are static and are visible at all times.

Lists are not always the best way of displaying information digitally - particularly if they are hidden within menus. The results from the study concluded that the un-ordered menu system performed poorly compared to the ordered list system suggesting that:

- Hiding contextual information such as bookmarks within menus is a poor design choice and can be considered ‘heavyweight’ if the function is frequently accessed;
- The area surrounding the PDF is potentially useful for displaying information. In this case, the static area surrounding the document was used to display the bookmark tabs that are visible at all times independently of the currently open page.

In addition to this, instead of locating information via clicks, the Visual Bookmarking System utilises a mouse-over pop-ups to automatically display important information. So:

- Mouse-over pop-ups that allow important information to be viewed easily by simply hovering the cursor over certain areas of the screen provide lightweight information at a glance.

Another area of digital document placeholdering that can be considered ‘heavyweight’ if performed incorrectly, is the ability to see where bookmarks exist in relation to each other as well as in relation to the currently open page – something that is achieved painlessly on paper. The ‘tabbed’ nature of

the bookmarks in the system ensures the user can see exactly how many there are within a document as well as providing an easy method of navigation. They have been specifically designed in telephone directory style to clearly show their location in relation to each other, i.e., a bookmark at the top is further towards the beginning of the document than one at the bottom, as well as their relation to the current page, i.e., bookmarks on the left occur on pages before the current page and ones on the right occur after the current page (see Figure 3.3). Therefore:

- A visual approach, in this case using the *position* of the tabs to indicate where in the document they are positioned is a good way of illustrating important data. This also gives a neat overview of where each bookmark exists in relation to each other as well as in relation to the currently open page.

Using colour also seems to be a useful way of distinguishing between items in a list, for example, using red to indicate important pages. Therefore:

- The use of colour may be a lightweight property if used effectively.

In summary, my evaluation of the Visual Bookmarking System has identified several potentially lightweight properties which I investigate in greater depth in later chapters.

### 3.7.2 Future Work

One current problem with the bookmark paradigm is overcrowding; bookmark lists, particularly within web-browsers, tend to grow over time, but rarely get shorter as users are reluctant to organise and delete them. This could potentially be a problem with bookmarks in general. If users have such a long list that they cannot easily find what they are looking for, they may not bother using it at all. Sometimes then, overwhelming bookmark lists are almost as bad as having no bookmark list at all [31]. A useful area of future work in this area could be to study how people manage their bookmarks, paying close attention to when and why they delete them. This information could then be used to implement a bookmark management system to aid in the organisation and archive of user placeholders.

The bookmark tabs do not give an overview of how far they are into the book. That is, a document with five bookmarks could have them clustered on the first few pages, or spread over the entire document. A useful addition to the program would be to give some context to the length of the book and indeed the bookmarks within it.

### 3.7.3 Related Publications

1. George Buchanan, Jennifer Pearson : **Improving Placeholders in Digital Documents** - *European Conference for Digital Libraries (ECDL) 2008, Aarhus, Denmark* - **WINNER OF BEST PAPER.**





# Chapter 4

## Annotations



From the ashes of disaster grow the  
roses of success

---

Lionel Jeffries

### 4.1 Introduction

This chapter investigates the role of marginal space in improving the support of active reading in digital documents. The user study I performed on the Visual Bookmarking System concluded that the space surrounding the PDF was a useful area to display information about the document (in this case bookmarks). This lightweight property could also apply to annotations, an area that is currently described by many as being unintuitive [85, 107, 118] on digital documents.

The work presented in both this chapter, and the previous chapter, are used as stepping stones for the next which describes a system that makes use of both placeholding and annotation in a single tool. Before embarking on such a large project, I decided to first study how and *where* annotations are performed on paper by conducting a short user study. The results from this were promising, concluding that the margins perform an integral role in the physical mark-up process as they allow notes to be made close to the original text.

Using the results of this study as a guide, I then created a digital implementation that gives users additional marginal space surrounding the PDF area for annotations to be made. Finally, I end by describing a second user study conducted on this implementation to determine whether or not additional margin space can be considered a digitally lightweight property.

### 4.2 Background

Annotation and note-taking are very common active reading activities. In 1998, Adler et al. [2] reported on a study in which 15 people were logged for five days performing their daily document activities. The results of this study concluded that when users are working on documents, nearly half of their time is spent annotating or note-taking. Clearly then, the value of mark-up is high, yet the uptake on digital media is low [107, 118].

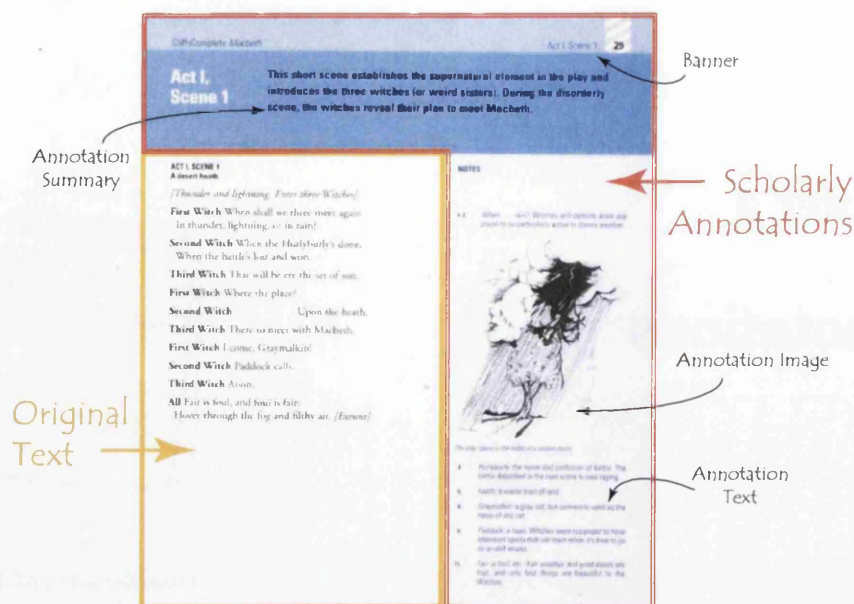


Figure 4.1: An extract from Cliffs Complete annotated Macbeth

### 4.2.1 Physical Annotations

Annotations on physical documents can take two general forms: the informal markings users make on a page while reading, and professional scholarly annotations found in literary work [84]. The latter here refers to critical edition books such as the extract from Cliffs Complete Annotated Macbeth shown in Figure 4.1. These professionally annotated versions of popular literature are intended to aid readers in the comprehension of classic and scholarly material and do so without altering the format of the original text. The example in Figure 4.1 shows in the bottom left the original Shakespeare text and in the other areas the annotations put in by the publisher. You can see here that the original text, in this case Act I, Scene I of Macbeth has not been altered, instead the annotations have been included in the margins.


A more common type of annotation made on documents are informal markings made by the user while reading and making sense of the text. These markings, which are commonly referred to as notes, annotations or scribbles, can include (but are not limited to): highlights, underlining, circling, doodling and text. Figure 4.2 shows four examples of informal user markings that show a diverse range of physical annotation possibilities. Figure 4.2a for example, shows a simple highlight, Figure 4.2b depicts the use of personalised symbols within the margin, Figure 4.2c uses a connector (in this case an arrow) to link the source material to an annotation, and Figure 4.2d shows an example of a key coding system that separates out different character's speech dialogues.

One issue regarding the annotation of pre-existing literature that has been little investigated thus far, is *where* the marks are made. Ideally the notes would be made close to the text to which they relate, for example, writing a definition next to a word in a text book would keep it in context. The problem in this instance however, is that it is not always possible to make notes close to the text, writing over the document itself may obscure the original text whereas writing in the margins provides finite space and could prove difficult to make a connection between the two.


Another potential problem with both these solutions is that writing over the text or even in the

Piglet!" And as Piglet looked sorrowfully round, Eeyore picked the balloon up with his teeth, and placed it carefully in the pot; picked it out and put it on the ground; and then picked it up again and put it carefully back.

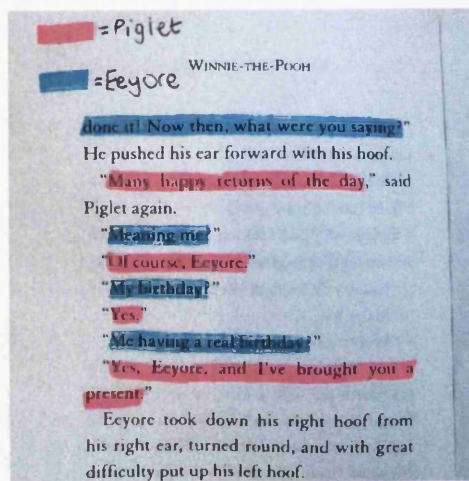
(a) Highlighting



(b) Margin marks



(c) Connections between source text and user annotations



(d) Key system

Figure 4.2: Examples of user annotation

margins damages the original document – a concept that has mixed opinions from readers [63]. A potential solution to this problem is to write the notes on a separate medium (e.g., spare paper, Post-its, notebook etc.). Although this method does give potentially unlimited space for writing notes, it now means that it can be difficult to reference specific locations within the original document and could easily get separated from the text to which it relates. Another solution to the problem of damaging the original document is to buy specialist note-taking equipment. The note tabs shown in Figure 3.1d in Chapter 3 for example, are sticky see-through notes that act as placeholders and allow users to make notes and highlights on a document without damaging the document itself.

### 4.2.2 Digital Annotations

Annotations within electronic documents are often considered unintuitive with some researchers even describing them as “cumbersome” [107]. This section discusses several of the most common annotation tools within digital document readers. For the sake of this discussion, I will not be pursuing other

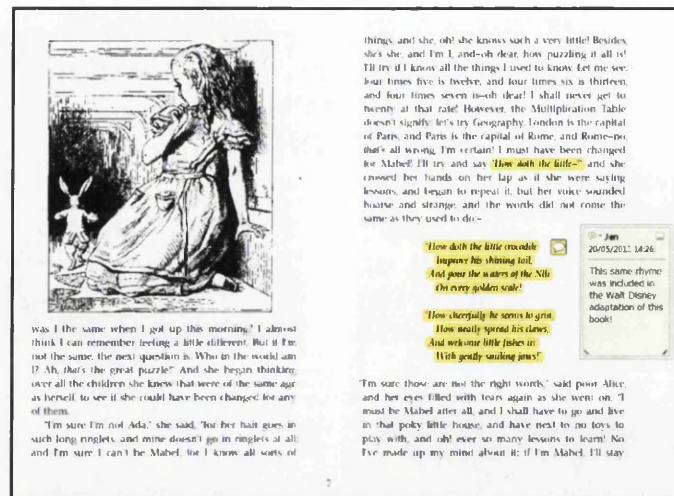


Figure 4.3: The comment features in Adobe Acrobat Reader X

types of digital annotation such as annotations seen in programming interfaces or web browsing.

One of the most popular digital reading applications on the market today is Adobe Acrobat Reader which is free for download and is available on Windows, Mac and Linux platforms. Unlike older versions, the newest version of the free Acrobat software (Adobe Acrobat X or version 10.0) also includes a small number of annotation features which were previously only available in the Professional edition. These features include text highlighting and sticky notes and can be seen in Figure 4.3. The non-free Adobe Acrobat Professional edition of the software, includes additional annotation tools such as call outs, shapes, text underlining/strike throughs etc.

Apple's Preview application also allows simple annotations to be made to existing PDF documents in a similar style to Acrobat. See Figure 4.4. The sticky note feature in this application forces notes to be made in the left hand margin of the screen and gives little flexibility to reorder and organise the notes within this area.

In its latest versions, Microsoft Word has also included a commenting feature which allows text to be highlighted and then comments to be added. This feature, which is aimed at collaborative editing, again forcibly stores comments in the right hand margin of the screen. See Figure 4.5.

### 4.2.3 Anatomy of Annotation

Researchers in the field of digital libraries have endeavoured to understand the different styles of annotation. For example, Agosti et al. [6] study historic practices of annotation, and the development of a digital library infrastructure to capture marginalia and other marks on mediaeval documents. By providing a consistent terminology for the different parts of digital annotation (e.g., [5]), it is possible to create a universal representation that facilitates the creation of usable annotations in digital documents. With this in mind, it seems sensible to first distinguish between the different parts of annotations in order to better understand how to create a complete digital solution. Marshall [84] describes three separate elements to annotations in her article "Reading and Writing the Electronic Book":

- *Body*: Any content added to the original document by the reader (e.g., a note or star written in the margin).

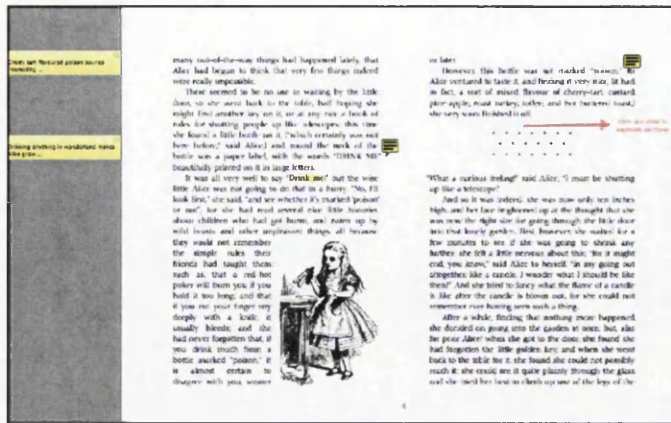


Figure 4.4: The annotations feature in Apple's Preview version 5

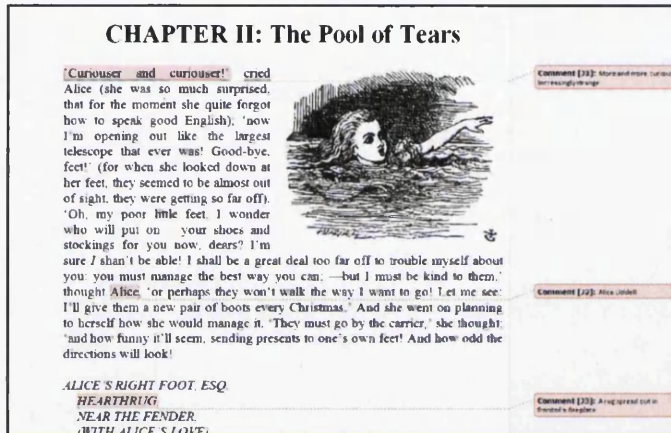


Figure 4.5: The comment feature in Microsoft Word 2007

- *Anchor*: The portion of the original text to which the annotation relates.
- *Marker*: How the anchor should be rendered when displayed, i.e., the type and colour of writing implement used.

To take an example, in Figure 4.2c, the *body* is the handwritten text “Ashdown Forest, Sussex. 100 Acre Wood”, its *anchor* is the source text “middle of the forest”, and its *marker* is a black rectangle. Similarly, Figure 4.2a, has an empty *body* (i.e., there is no additional content added by the user), its *anchor* is the text “Eeyore picked the balloon up with his teeth”, and its *marker* is a yellow highlight.

Digital implementations offer several other useful annotation elements that would be tricky and time consuming on paper. For example: the *author* which records who created the annotation, and the *time stamp* which keeps a record of exactly when the annotation was made. While it is unlikely that such a complex and complete annotation will be seen on paper, it is useful to have a set of common terms that will facilitate the discussion and comparison of annotations independent of implementation or platform.

It is clear from the examples in Figure 4.2b, that annotations are made up of several elements and can take many different forms. For example, highlighting a word then using a pen to draw a

connection to the margin where a note is made. Thus, providing just one tool to support such a feature, as many digital systems do, is not a complete solution. To create a complete tool set for a digital implementation then, these different annotation elements must be studied in order to provide the appropriate tool set for the jobs at hand. I will therefore, be using this anatomy of annotation as a grounding for the design of the tool set in my improved digital solution, as well as a base for evaluating if it is indeed the *minimal* complete set of tools.

#### 4.2.4 Marginalia

One of the key concepts relating to this chapter is the idea of marginal notes, or *marginalia*:

**mar-gi-na-li-a** (*plural noun*)

1. Notes, commentary, and similar material written or printed in the margin of a book or manuscript. Also (in extended use): notes, comments, etc., which are incidental or additional to the main topic.

There are many famous figures who are known for their marginalia (e.g., Samuel Taylor Coleridge [32]); in this case of course, a readers' notes can significantly increase the overall value of an otherwise 'damaged' document. The act of writing notes in the margin of a printed document has been around, probably for as long as there have been margins to write in, providing a useful space to make notes to aid memory and communication. Jackson [63] has investigated the concept of marginalia in depth, from casual scribbles to lengthy arguments, by studying the use of thousands of annotated books from a time span of over three centuries. After referring to readers' annotations as "*a familiar but unexplained phenomenon*", Jackson then reflects on the cultural and historical value of marginalia and presents a detailed review of some of the more interesting and famous examples in history.

The use of margins to place notes is a common occurrence in the physical world, but there have been occasions when researchers have encountered complaints from study participants with regards to the lack of marginal space on digital documents. Marshall et al. [86] for example, observed several user criticisms over the XLibris system's (see Section 2.5.3) lack of marginal space:

*"One reader complained that the margins, which were reduced when we scanned in the paper, were too small for marginalia. We have observed this complaint in other uses of XLibris as well. This emphasizes the importance of large margins for encouraging annotation, but it also introduces a design trade-off. Assuming a fixed size display, and scanned or pre-formatted documents (e.g. Postscript or PDF), increasing margin size ultimately decreases readability."*

– Marshall, Golovchinsky and Schilit [86].

This evidence suggests that marginalia can form an important part of the active reading process on both physical and digital documents. To investigate this issue further and to conclude exactly where the most common and popular areas for annotation are located, I first conducted a comparison study on paper based material. The results of this study were then used as a grounding for the design of a digital annotation system which I evaluate later in this Chapter.

### 4.3 Literature Review

The topic of annotation has been studied in great depth by a number of researchers over the last few decades. One such research programme, conducted by O'Hara et al. [108], is the study of what they call 'information recording' activities, that is, how readers extract and record information from a source text. They investigate the document related activities of 25 library users from the research sector. This set of intensive readers, were asked to record a complete diary of work-related and research based activities they perform in the library in a working day. Using the results of this study, they were able to create a model of document related activities which shows that reading can lead directly to information review, or first go through the phases of either annotation or note-taking. They observe in the study that the majority of annotations made (on photocopied documents) were made in the margin area and note that in most cases, the participants only wrote short notes or brief, incomplete phrases rather than large complete sentences. They observed from this that the meaning of these annotations was highly dependent upon the textual context to which they relate, and comment that they are not independently interpretable, rather that they are distinguishable only with the original text to which they are linked. This result strongly suggests that the positioning of annotations is an important factor in the active reading process, as poorly positioned notes may be hard to link back to the original content to which it relates.

Marshall [82] also conducted an investigation into work-related annotation by studying the marks made on text books in a University book store. The particular store provided a 'buy back' scheme which allowed students to sell their previously used books, along with any annotations, back to the store after use. Marshall uses the observations from this study to categorise several ways in which users associate annotations with the printed document elements. She states that users making marginal notes generally connect them to the source text in one of three ways: using arrows, using a bracket, brace, or some other mark and sometimes even relying on proximity alone to connect their own marginal (or interline) jottings with the text. Any attempt to recreate the physical annotation process digitally therefore, must ensure a complete set of tools is incorporated to allow all the types of anchoring mentioned by Marshall and others.

The creators of XLibris, the active reading machine, have also focused strongly on the concept of annotation [114]. Although the XLibris interface does not offer *additional* space surrounding the document area for making notes, the creators do make very clear that their free-form annotations can be made anywhere on the document itself because:

*"We believe that the ability to make unstructured, idiosyncratic marks is a key aspect of a paper-like user interface."*

– Schilit et al. [114].

This type of un-restricted annotation is key to creating a complete paper-document metaphor for use on digital documents and one which I intended to replicate in the design of a digital annotation system. Another resemblance of the work of XLibris to the work presented in this chapter is their use of margins, in this case to provide serendipitous links to related material during the reading process. Later in the Chapter, and in the remainder of this Thesis, I discuss the benefits offered by an extended workspace in the active reading process.

Another body of work which several researchers have investigated, is the concept of *re-flowing* annotations, specifically, content added to documents where the underlying content is modified. In such cases, each annotation must be transformed to fit in with the new layout of its context. Cadiz et al. [22] reported on a study where they observed around 450 people making use of shared annotations



within Microsoft Office 2000's web annotation facility and discovered that the biggest problem identified by users was orphaned annotations. They describe annotation orphaning (i.e., annotations that no longer correspond to a place within the original text content) as being "*understandably frustrating*", arguing that annotations mean little without the original content to which they relate.

Brush et al. [15] have also investigated user expectations of annotations on documents where the underlying document is subsequently changed and present a framework for building robust annotations that alter their position with the original text. Barger et al. [12] implemented a system that facilitates re-flowing annotations, this time using free-form ink on augmented reading hardware. The Xlibris project has also studied the digital ink repositioning algorithms for modifiable documents in their paper, Moving Markup: repositioning freeform annotations [48].

Dealing with annotations on documents where the original content is changed is a large topic and one which I decided not to pursue. Although digital documents do have the potential to be reformatted, for example, eReaders reformat pages when the font size is altered, the work presented in this chapter and indeed the Thesis as a whole, is not concerned with modifiable documents as all the software I implement assumes the original text remains unchanged. Concerns over orphaned annotations or lost bookmarks are therefore, not an issue.

In addition to re-flowing annotations, there is also an assortment of literature that focuses on digital annotation in the form of pen-based input, allowing the user to effectively 'draw' on electronic media much like they would on paper. Plimmer and Mason's paper [110] for example, focuses on annotations used within a classroom environment by exploring the paperless method of marking assignments. The main aim of the project, which focuses predominantly on marking up programming assignments is to allow mechanisms for recording and recognising scores and support rapid transition from assignment to assignment. The end product which uses the Tablet PC API to convert hand written comments to text, simultaneously annotates while recording the grades for assignments. When marking is finished a PDF file is automatically generated in order to provide feedback for the student.

There have also been investigations into the differences between paper and online annotations. Kawase et al. [65] for example, present a comparison of paper based versus online annotations by conducting a series of laboratory and field studies on users in a workplace. The results of their study concur with previous findings [107] that on paper, annotations support the learning process. The results also suggest however, that digital annotations provide an extra cognitive burden due to the lack of suitable tools available and problems relating to reading and interacting with documents on-screen. They suggest that online annotations are shorter and of less intrinsic value than paper based ones and suggest that digital annotation developers make better use of annotations digitally (e.g., visual overviews, enhanced finding techniques etc.) rather than attempting to mimic paper based interactions.

## 4.4 Paper Study

There have been a variety of studies that have focused on how users annotate on printed documents. However, a topic that has thus far been overlooked is the concept of space – specifically where these annotations are positioned. Superficially, this issue is straightforward: annotations will appear near to the material to which they relate. For example, a circle around a word within a paragraph and an arrow pointing to a definition somewhere in the margin. The problem faced with annotations on printed documents, is that space is finite and can potentially force additional mark-ups to be made outside the desired area. For example, an annotation about a paragraph in a document that has been made in a separate notebook. This scenario consequently presents the problem of referencing the

separate annotation back to the original text. As I have discovered from the literature [108], the concept of space has value within active reading systems, yet so far there has been little literature on the topic of exactly where is the best place to position notes.

On physical documents, there are 3 distinct locations where annotations can be placed:

1. Over the document itself;
2. In the margins of the document;
3. On a separate medium.

All of these locations suffer from potential problems however. For example, annotating over the document itself obscures the original text, whereas writing on a separate medium makes it hard to reference related material within the document. Which of these options is preferred by the majority of users is currently unclear as surprisingly little research in the area of annotation location has been undertaken. It would be useful therefore, to conduct a paper based study that investigates annotation placement to discover the most popular and most practical mark-up locations. A study such as this raised several research questions for me, such as: where is the *ideal* annotation location? and where is the *most common* annotation location? These questions also posed further questions such as: are the margins used for annotation simply because there is no other free space? and is the potentially infinite space provided by a separate medium being outweighed by the possibility of losing or dereferencing the additional material?

I chose these questions specifically to gain a better understanding of how and where annotations are made on paper, in order to better inform the design of a digital solution. To answer these questions, I undertook a paper based study and post-study interview with a set of participants who regularly undertake annotation tasks. The following section describes the design and results of this study.

#### 4.4.1 Study Design

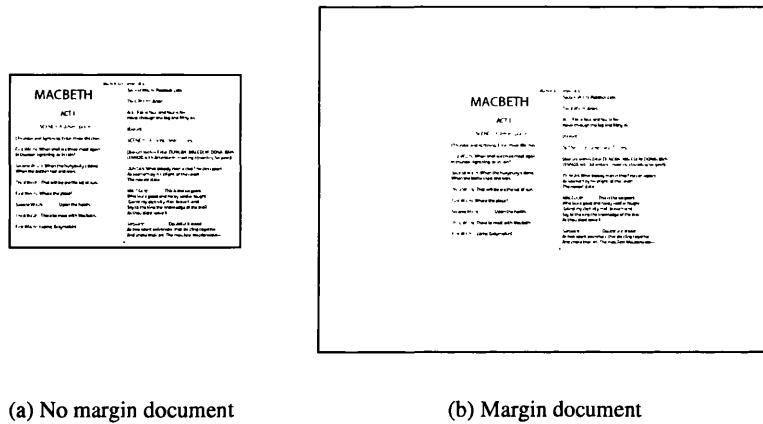
To best determine *where* annotations are made on paper documents, it seemed wise to ask users to perform simple mark-up tasks and observe the locations of the notes produced. The following sections describe the design of the study carried out to observe users' paper annotation practices and the post-study interview conducted to probe their ideal annotation location.

##### The Systems

I anticipated that the size of the margin would strongly affect the placement of annotations on a document. Therefore, to observe the effects of marginal space on the mark-up process I decided to perform a comparison test using two types of printed document. The first, which was used to observe user behaviour when there is little space on the original document to make notes, is shown in Figure 4.6a and is known as a 'no margin' document. In contrast to this, and to gain insight into user behaviour when there is ample space surrounding the text to make notes, I have also included a 'margin' document which is illustrated in Figure 4.6b.

##### Documents and Apparatus

Each participant was provided with two printed documents during the study: a margin document and a non-margin document (Figure 4.6b). In both printouts the size of the original text remained constant, but the margin document contained an additional 4cm uniform margin.



(a) No margin document

(b) Margin document

Figure 4.6: Examples of the two types of document used in the study

The choice of documents used in the study and consequently the tasks that were given, varied from user to user to remove any possible bias and ensure a fair and balanced comparison throughout the systems being tested. In total, I used four different PDFs in the study which were rotated amongst system and participants to ensure that no two participants received the same combination of system and document. The PDFs used were: Macbeth, Alice in Wonderland, Heart of Darkness and A Christmas Carol.

In addition to the printed documents, each participant was also provided with ample mark-up tools which included (but was not limited to<sup>1</sup>) coloured pens, pencils, highlighters, spare paper, Post-its and sticky tabs.

### Procedure

The participants I selected for the study were chosen from a group of post-graduate students and academic staff of Swansea University and consisted of 9 men and 1 woman between the ages of 21 and 35. I selected this user base intentionally due to their increased likelihood of annotation practices in everyday life.

In order to maximise the amount of information gathered in the study, I used video recording equipment to record the participant's actions throughout the testing process. Consent was obtained from each participant prior to the study taking place and anonymity was maintained by focusing the camera away from their faces. I kept the resulting artefacts produced by the study participants for analysis and further study.

The structure of each study comprised of:

1. A set of seven tasks to complete on each of the two document types (the non-margin and the margin systems);
2. A post-study questionnaire specific to the exact document used in the tasks. Given the variation of printed PDFs amongst the participants, each set of tasks and consequently the subsequent

<sup>1</sup>The participants were asked prior to the study taking place if there was any equipment they would like to be added. No participants requested any additional materials

written questionnaire varied slightly from user to user but adhered to the same general theme preserving the ratio of open and closed tasks;

3. A short series of semi-structured interview questions.

#### 4.4.2 The Tasks

To ensure the PDFs chosen would not affect the study, they were interchanged between each participant and the margin and non-margin documents. The tasks given to each participant varied depending upon the PDF used but all contained seven tasks per system (14 total). The format of these seven tasks included two closed and five open-ended assignments.

The closed tasks were designed to provide the user with a precise instruction for annotation that should vary only slightly between participants. An example of a closed task used within the study was: **'Highlight all of Alice's dialogue'**.

The open tasks however, were intentionally designed to give the user more freedom in their approach to annotation in order to tease out personal mark-up preferences. Two examples of the open tasks used within the study were **'Use the tools provided to create a link between the word 'heath' and the note you just made'** and **'Mark Page six as important'**.

#### 4.4.3 Hypotheses

The primary goal of the paper study was to identify where users make annotations on physical documents. In addition, I also wanted to gain subjective and qualitative data regarding users preferred annotation positions. My hypotheses for the results of the paper study were as follows:

**Hypothesis 1:** The margins would be a more popular area for annotation than over the document itself;

**Hypothesis 2:** The non-margin document would be rated poorly compared to the margin document for obscuring the original text;

**Hypothesis 3:** Additional mediums would be a useful alternative where margins are not available (i.e., the non-margin document).

#### 4.4.4 Study Results

##### Subjective Results

As part of the post-study questionnaire participants were asked to rate each of the document types (margin and non-margin) out of 10 (1 being the lowest 10 being the highest), for several attributes. The results of this portion of the study are shown in Table 4.1. To assess the statistical significance of the data I obtained, I conducted a Wilcoxon Signed-Rank test on the subjective scores.

It is clear from the results that the biggest divide in opinion between the margin and non-margin document types is seen in how well the original document can be seen after annotations have been made (Q8). This yielded average subjective scores of 3.2 and 6.2 for the non-margin and margin documents respectively and a significant Wilcoxon result of  $p < 0.011$  ( $W+ = 2$ ,  $W- = 43$ ,  $N = 9$ ). This confirms hypothesis 2 – that is, that the non-margin document would be rated poorly compared to the margin document for obscuring the original text. This result, however obvious, strongly suggests that the margins play an important role in the annotation of physical documents, as without them, the

	Non-Margin		Margin		Wilcoxon (p)
	Average	SD	Average	SD	
Q1: How easy it is to emphasise or highlight specific areas	9.0	0.942	9.0	0.942	=1*
Q2: How legible the notes you make are	7.1	2.025	8.8	0.919	=0.0313
Q3: The speed in which you can make notes	7.1	1.524	8.8	1.033	=0.0156
Q4: The effort it takes to make the notes	6.9	2.025	8.1	2.378	=0.0938*
Q5: How easy it is to make connections between sections	6.6	1.955	8.0	1.633	=0.0078
Q6: How easy it is to highlight an area	7.8	2.150	8.2	1.619	=0.4375*
Q7: How easy it is to indicate specific places	7.5	2.068	8.3	1.337	=0.125*
Q8: How well the original document can be seen afterwards	3.2	1.751	6.2	2.898	=0.0117

Table 4.1: Subjective ratings from the paper study (out of 10). \* indicates no significant result was found.




Figure 4.7: Users' annotation locations - ideal and actual

clarity of the original text can be obscured. This theory is also backed up by the results from Q2, Q3 and Q4, indicating that the participants felt their notes were more legible, speedier and easier to make on the margin document versus the non-margin document.

The results from the post-study interview substantiated these findings proving that the majority of participants (90%) agreed that having more space around the outside of the document is useful when marking up. In addition, I also confirmed that this additional space is not detrimental to the actual reading of the original document with all participants disagreeing with the statement "*Having bigger margins makes it more difficult to read the original document*". This evidence confirms hypothesis 1 – that the margins would be a more popular area for annotation than over the document itself.

### Annotation Placement

Figure 4.7 shows the results from the annotation placement portion of the study. Figure 4.7a illustrates the subjective results from the post-study interview, specifically the proportion of users who selected each of the three areas as their *ideal* position to place notes. Figure 4.7b shows the *actual* percentage of annotations made in each of the three positions, which was gathered from studying the resultant paper artefacts at the end of the study<sup>2</sup>. Statistical testing on these two sets of data would be in-feasible

<sup>2</sup>This data was obtained from the margin document types in the study as there was no option to place notes in the margin of a non-margin document. It also shows the distribution number of annotations made by all participants

due to differences in how the data was gathered. Specifically, in the ideal locations the users simply gave a number from one to three to indicate their preference for each location whereas the actual locations were identified via the number of annotations made by the participants during the study. Putting these sets of data side by side in graphs however, clearly illustrates that where users actually place their notes is very similar to their preferred placement of notes.

It is useful to note here however, that the documents used in the study were draft-like copies of hardback books, that is, they were simply ring-bound documents printed on standard office paper. I also made it clear to the participants at the start of the study that marking them with ink was the intention. When prompted about annotating their own purchased or borrowed books in the post-study interview, all participants admitted that they would never annotate the book itself for fear of damaging it. Conducting the same study on participants' own copies of books then, is likely to produce drastically different results for users' actual annotation locations. Indeed, when questioned about the topic every participant confirmed that when reading an actual book, they always use a separate medium for annotation or more surprisingly, opt not to make notes at all.

Despite this however, there is a mixed reaction to the subject of whether or not a separate medium is a resolution to the problem of annotation on documents with little or no margin. Although several users agreed that it is a potential solution, they all seemed to follow with a "but", suggesting that using a separate medium is an alternative to writing in the margins of a document as opposed to a comparable solution. This therefore, partially proves Hypothesis 3 – that additional media would be a useful alternative where margins are not available. The main problems the study highlighted in using a separate media to mark up documents were:

1. Their tendency to get lost or detached from the original text;
2. That separate notes mean little without the original text to refer to;
3. It is hard to reference a specific section in the original text on a separate medium;

The related subjective ratings from the participants reinforce this point, that is, using a separate medium is the lowest scoring place for annotation, with six users choosing it as their last resort, and only one participant rating it highest. Despite these reservations, observational data qualifies this low rating: eight participants used at least one separate medium during the study (either a piece of paper or a Post-it). Half of this group (4/10) used this method in both the minimal and extended margin documents. The other half used the separate medium *only* in the minimal margin presentation where there was less room to annotate; not a single participant used supplementary paper on the margin document only. The higher adoption of the separate medium when there is little or no margin available for annotation supports the theory that less margin space encourages the use of additional media.

Although a separate medium may not be the majority of users' ideal method of annotation, observational data suggests that many are forced to use it occasionally regardless. Despite the fact they were not particularly popular with the participants, there are some good attributes associated with separate media:

1. They offer potentially unlimited space;
2. They do not obscure the original document text;
3. They do not damage the original document.

Despite users' ideal and actual annotation placements (Figure 4.7), it is possible that a separate medium would be the ideal solution if the problems associated with them were eliminated, that is, if there was no chance of them getting lost and they are always linked with what they reference. On paper, this would be an impossible task, however, digital techniques can provide solutions to the problems associated with separate mediums.

One final useful observation I discovered during the study was the way in which certain participants completed the task 'Mark page X as important'. During this task, three out of the ten participants used a Post-it to indicate page importance by writing or drawing asterisks then proceeding to stick it to the side of the page, causing it to then act as a rudimentary bookmark. Although this may seem an obvious action to take on paper, the way in which a single physical object can perform more than one function does not extend to the digital world.

In fact, the positioning of Post-its that stick out of the page is not an action that is commonly seen in electronic documents at all, despite it being a common action to take physically. Many of the Post-its used within the study were used, not to write on, but to stick to the sides and act as independent bookmarks. This is an act that has clearly been anticipated by many manufactures, particularly those who make sticky tabs that are too small to write any significant notes on, yet perfectly sized and equipped to act as placeholders. These points form the basis of the next chapter on note-taking.

#### 4.4.5 Paper Study Conclusions

The results of the paper study indicate that people enjoy the freedom of physically writing on the document, whether it is on top of the original text or in the margins, as annotating close to the section of literature to which you are referring keeps the notes in context. This is backed up by participant comments such as "*I usually want to reference specific parts of the document and having the note near that section is very useful*" as well as the statistics – only 1 out of the 10 participants chose the additional media as their preferred method of annotation.

However, this study is an unrealistic test for annotation practices on *books* as opposed to printed pages as all of the participants agreed that they would only mark an actual book in extreme circumstances. They also confirmed that when making notes about a personal book, they either use a separate medium (e.g., scrap paper, Post-its, note-book etc.) or decide not to make notes at all. This is a problem that does not have to extend to digital documents. When designing a new digital annotation system, provisions can be made to ensure that the original document remains unaffected by saving the annotations in a separate file. This will overcome the problems associated with damaging the original text as notes can be toggled on and off, and allows multiple annotation files to be associated with a single copy of the document.

The study confirmed the popularity of margins for taking notes with 70% of the participants selecting it as their first choice for annotation compared with only 20% and 10% selecting over the document and separate medium respectively. It also uncovered some undesirable attributes associated with the separate medium method of annotation. Specifically, the fact that they can get lost easily, it is hard to link notes on a separate page to specific sections of the original text and the fact that they can be dereferenced (i.e., notes mean very little without the original text to link it to). Again, these are very real problems in the printed world but ones that need not carry over to the digital plane. This topic is discussed in more detail later in the Chapter.

After studying user behaviour during specific tasks, the study also suggests that the concept of appropriation, or making use of one tool for multiple purposes, may also be a potentially lightweight property if applied digitally. In the next chapter I investigate this point in detail by designing a digital tool that incorporates both placeholder and annotation into a single unified tool.




Figure 4.8: Design for the Margin Annotation System

## 4.5 System Design

By paying close attention to the results of the paper based study, I have designed a digital implementation that utilises the space outside the main document area for annotation. To better understand any potential routes for further investigation, I decided to create the digital system in a way that allowed me to make a formulative analysis of the comparative user study described later in the Chapter.

### 4.5.1 Aims

The main aim for this part of my investigation, was to design a digital reading system with an expandable margin area. In doing so, I hope to prove that the space surrounding the document is a lightweight property that can be applied to many different scenarios within the digital document domain.

- Aim 1:** Create a system that includes additional expandable marginal space surrounding the PDF which can be used as space for mark-ups;
- Aim 2:** To ensure the original document is not 'damaged' and can easily be seen minus the annotations, ensure the annotation files are separated from the original file;
- Aim 3:** Give the user a complete set of annotation tools to mark up the PDF itself, and the space surrounding it;
- Aim 4:** Create a tool that replaces the separate medium by incorporating its advantages (i.e., its potentially unlimited space, the fact it does not obscure or damage the original document) and overcomes its problems (i.e., it should always be referenced, there should be no chance of it getting lost; it should be hidden from view unless explicitly told to appear).



### 4.5.2 Additional Document Space

As I discovered from the paper study as well as my investigation into placeholders, the space surrounding the main document can provide a useful area for displaying and storing information. This was backed up by the results of the paper based comparison study described previously in the Chapter. The difference between the additional space used in the placeholder system and the space used in the comparison study are described below:

**Margins:** Document margins give each individual page additional space surrounding the document text to make notes. Therefore, each page's margin belongs to the page itself and turning to a new page gives a fresh margin. Notes in this form can be made on the document directly, or by sticking or inserting a separate medium (e.g., Post-it) to them.

**Static Space:** Other types of additional space where notes can be made about a document include static space which stays constant independent of the open page. For example, placing Post-its on the desk or the wall close to the space where the user is working. This type of additional space ensures that the content of the annotation can always be seen, despite the currently open page.

### 4.5.3 Storage

As I discovered while conducting the comparison study, users are reluctant to annotate their own books as they do not want to damage them. This preference is something that can easily be overcome in the digital world by simply making a copy of the original document before marking it up, ensuring that they always have an 'un-damaged' copy of the original file. A more efficient method of achieving the same effect however, would be to save all document annotations separately to the main document file. This would ensure the original document remains unchanged while at the same time facilitating toggleable annotations. It has been noted in previous investigations [107], that annotations are often perceived as:

*"a separate layer of the document."*

– O'Hara and Sellen [107].

Therefore, separating them from the original text seems a logical step to take.

Another advantage of separating the notes from the original document content is to allow multiple annotation files to be associated with a single PDF, allowing many versions of annotations to be viewed while only saving a single copy of the original text.

In the Margin Annotation system after a set of notes and/or annotations has been created, the program will allow them to be saved to a file that is independent of the PDF itself – this will be known as an annotation file. A PDF can have several of these files associated with it, each of which will be given a different name and saved separately. A library file containing information on all the PDFs and their respective annotations is used to connect these files and a menu within the main program will give access to a list of all the annotation files linked to the open PDF. The result of this is a standard format of annotation that can be created, viewed and edited easily without altering the PDF itself. Consequently, these files form a template which can be shared among computers without the necessity of having multiple copies of the same source text.

Sometimes it is beneficial to purchase a professionally annotated version of a classic book to aid in the understanding of the text. Many companies produce scholarly editions of popular books but to

view them all, the reader must in turn buy several copies of the original text. Similarly, if someone makes notes on an electronic document and wants to share it with others, they must not only send the annotations, but also the original document along with it. With the system proposed above however, users can open any number of annotation files while only owning one copy of the original text. See Section 4.5.4 for a scenario.

#### 4.5.4 Additional Media

Although the results of the paper study highlighted the unpopularity of the separate media annotation approach, it did suggest that it could prove useful if the problems associated with it were solved, specifically, if there is no chance of them getting lost or dereferenced. Incorporating a feature that overcomes these problems as well as utilising its benefits, could prove to be a valuable asset to the system. It is for this reason, that I have incorporated a 'Note' feature into the design.

Notes within the system are designed to be connected to either highlights or pen marks. This is solely due to the way in which they are displayed, that is, they are hidden from view unless the mouse passes over the object to which they are connected. When this action occurs, the program then paints a mouse-over pop-up containing the note text on the screen. This note feature then, retains the original benefits of the separate medium approach and, as long as the library file remains un-tampered with, it also solves the problems associated with the separate medium approach (the notes cannot get lost or dereferenced).

#### Scenario

It has been documented by experts [4] in the area of annotation, that there could be potential advantages to the layering of annotations to enrich their semantics and expressive power:

*"It is possible to have different layers of annotations on the same document: a private layer of annotations, which can be accessed only by the annotations author himself; a collective layer of annotations, shared by some users who are working on a document; finally a public layer of annotations, accessible to all users of the digital library."*

– Agosti and Ferro [4].

This description discusses the role of annotations in the context of the digital library and reinforces my decision to split mark-ups from their underlying documents. The following scenario describes a typical situation where the above statement becomes a reality:

*Joe is a student who is writing an essay on 'Alice in Wonderland' for a school assignment. He has an electronic version of the book which is stored on his computer in his local documents file. He opens and reads the document with the Margin Annotation System then proceeds to mark-it up with his own notes using the tools provided. Joe gets to a tricky piece of the text which he has trouble interpreting, so uses the program to switch to the professionally created annotation file he previously downloaded from the internet to help him decipher the material. After finishing the assignment, Joe opens his classmates' annotation files to compare answers before submission. He then sends his completed assignment, along with his annotation file to his teacher. The teacher is then able to open and view all 30 of her pupils annotation files on her computer while only storing a single copy of the original document.*

In this scenario, the Margin Annotation System, is allowing Joe to open and easy toggle between multiple annotation files (i.e., his own copy, the professional scholarly copy, and his classmates' copies) while only physically storing one version of the original document.

### Alternative Designs

In contrast to the previous chapter, my intention with this design was to elicit user requirements. The paper study revealed that it was beneficial to provide space for annotation outside the document content itself. In considering alternative designs, I will therefore focus on other designs that would help gain insight into user requirements for this wider annotation area. One possibility is the provision of an entirely separate note-taking space, detached from the document (like a notebook), mirroring paper. However, electronic media can be scrollable and need not have a paginated form.

Tashman et al. [132] also investigate the role of margins in supporting reading. However, they use only a right-hand margin and a scrolling display of the document pages. Given my focus on transferring behaviours from paper to digital media, it was important to match the physical world as much possible by using a paginated display and uniform margins on all sides of the page. Investigating the role of a margin in a scrollable display is a worthwhile research goal, but as Tashman et al. have now addressed this, there is little advantage in me simply reproducing their design and work.

XLibris [111] provided an optional separate notebook for capturing document content and adding the user's personal notes. That approach clearly mirrors the behaviours found when users work with paper, and would be very much worthwhile investigating as a further option when studying this topic.

Shipman et al. [120] examined the use of a dedicated, two-dimensional note-taking space versus Word in the context of recording and interpreting across a collection of documents. Their evaluation demonstrated several weaknesses in the continuous linear sequences of notes most readily supported by Word, and in Word's lack of connection to the collection the participants worked with. While XLibris' note-taking space remedies some of the minor problems reported by Shipman et al., it seemed more judicious to follow an approach that both mirrored paper and seemed more likely to avoid the shortcomings Shipman's study revealed.

It would, of course, be possible to encompass or adjust the behaviour of commonly used annotation tools, such as the pop-out note, to extend the space available to the reader. In the case of pop-up notes, they also occlude the text permanently (by the marker indicating their presence), and when activated also hide the underlying text. It is to adjust this behaviour to minimise such issues, but this would often result in a similar design to either a larger margin or a separate note-book document.

## 4.6 System Implementation

### 4.6.1 Tools

For simplicity, the tools within the original system have been split into several categories: PDF tools, margin tools and miscellaneous tools. The tool sets used within the user study altered slightly from this original design to test the different aspects of the system. Figure 4.10 shows a labelled screen shot of the the task bar.

#### PDF Tools

The following section describes the tools that I have designed for use specifically over the document itself:

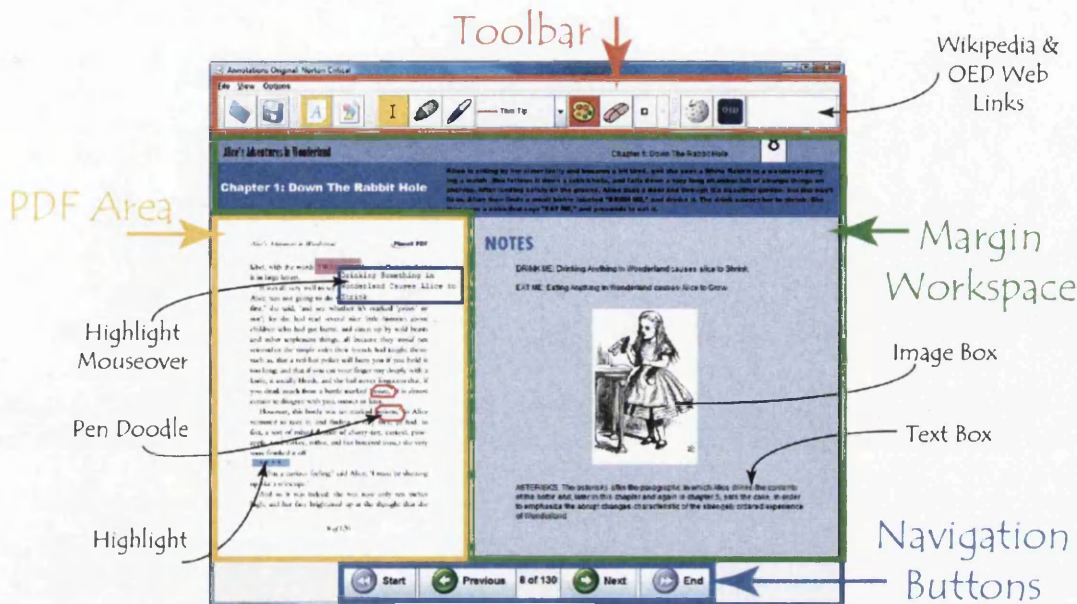


Figure 4.9: The Margin Annotation System

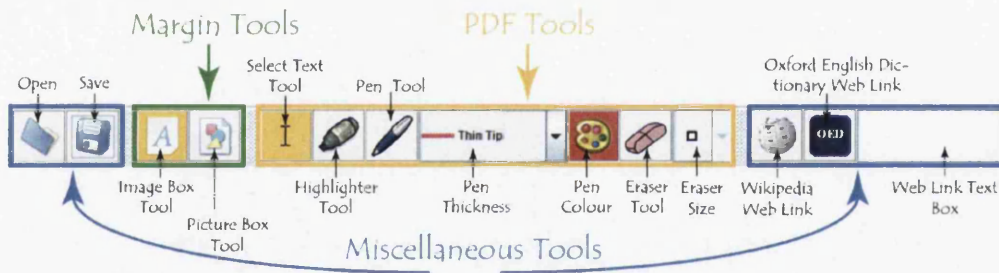


Figure 4.10: The Margin Annotation System's toolbar

**Highlighter:** A tool designed to highlight the main PDF. Highlights can be changed to any colour and can be resized to cover text, images or a combination of both. Notes can also be added to highlights by double clicking them (see below).

**Pen:** A tool controlled by mouse movements that allows free-hand lines to be drawn over the top of the original PDF. The colour and thickness of the pen tool can be changed in the toolbar at the top of the form. Notes can also be added to pen marks by double clicking them (see below).

**Note:** The tool created to act as an improvement to the separate medium. Notes can be added to highlights or pen marks by simply double clicking on them. Once a note has been created, it is hidden from view until the mouse is run over the object it is connected to; at this point, a pop-up window is displayed containing the contents of the text note.

**Eraser:** A simple tool that erases annotations made with the pen tool in a raster graphic fashion. Users can change the size of the eraser by means of the toolbar at the top of the form.

**Select Text:** A simple tool added for completeness to allow users to select and copy text from the main document into annotations. This tool is used in conjunction with the web links tool described below.

### Margin Tools

When given additional space, some of the tools that are used to mark-up text become less useful, for example, a highlighter pen would be redundant unless there was something (e.g., text) to highlight. I have decided therefore, to give users a separate and more appropriate set of tools to use on the margin area of the system:

**Image Box:** A tool that creates resizable image boxes that can be positioned in the margins of the document. The image displayed in the box can be changed by double clicking the box area.

**Text Box:** A tool that creates customisable text boxes for making notes outside the document area. Text boxes are very versatile and can be used to mimic scholarly editions as shown in Figure 4.1. The background colour, border colour and thickness, text size, font and colour and size of the text box can be modified for flexibility with the tools.

### Miscellaneous Tools

In addition to the PDF and margin based tools, I have also included several miscellaneous tools to increase the intuitive nature of the system and to probe their usefulness as part of my initial formulative design.

**Web Links:** Highlighting a portion of the original text using the 'select text' tool causes that text to be automatically copied into the web links box at the top of the form. The user then has the option to click Wikipedia<sup>3</sup> or Oxford English dictionary<sup>4</sup> buttons for web definitions of the selected text. User specified text can also be searched by manually typing directly into the web links box. I anticipated that this may be an innovative way of quickly looking up words and phrases from the document.

**Docking:** The PDF itself can be 'docked' at different locations around the background space to provide a higher degree of freedom with the notes being made.

**Toggling:** As Aim 1 states, I decided to save the annotation files independently of the PDF to ensure a) the original document is not 'damaged' by user created annotations and b) to allow multiple annotation files to be associated with a single PDF. To switch between different files then, the user can make use of the toggling feature which allows quick changing between different annotations as well as allowing the original document to be viewed alone.

---

<sup>3</sup>[www.wikipedia.org](http://www.wikipedia.org)

<sup>4</sup>[www.oed.com](http://www.oed.com)

## 4.6.2 Saving and Opening

Annotation files created using the program are saved independently of the PDF (i.e., the PDF itself remains unchanged) therefore it is possible to have multiple annotation files associated with a single document. The Margin Annotation System therefore, needs to save each annotation file separately as well as having a library file that connects them all together. Each computer running the system will only require one library file which will contain information about all the PDFs opened with the system in order to group the specific annotations associated with it. This library file is used when a PDF is opened to make a list of all the different annotation files that have been made on this particular document.

In addition to these two types of file, the system also uses PNG image files to save user 'drawings' made using the pen tool. All these files are stored in the same directory in a folder called 'Library'.

## 4.7 User Study

To confirm that additional marginal space is a lightweight property digitally, I conducted a comparative user study on the Margin Annotation System. This study was designed to test several aspects of the system as a whole including how useful the margin is for marking-up and if I can expand and improve on the separate media idea digitally.

### 4.7.1 Study Design

#### Research Questions

Since this study is mainly a formulative investigation, I have laid down research questions as opposed to hypotheses. I refer back to these questions in the results section. It is worth noting here that these questions are being answered purely on the subjective responses from the participants in the study – no usage behaviour was used to determine the success of any of the tools within the system.

**Research Question 1:** Is the margin a useful area for making notes?

**Research Question 2:** Is the mouse-over note function a suitable replacement for the separate medium seen in paper documents?

**Research Question 3:** What is the complete set of annotation tools?

**Research Question 4:** Are the miscellaneous tools useful? (e.g., OED links, toggleable annotation files etc.)

#### Control Systems

To test the effectiveness of the margin area and its tools, I implemented two additional systems: one with expandable margins and one without:

**Margin System:** The layout of the margin and the PDF remains the same, as do all the miscellaneous tools. In this system however, all the margin and PDF tools (described above) are available on the *margin only*.

**PDF System:** As with the Margin System, the layout and miscellaneous tools remain constant, but this time all margin and PDF tools are available on the *PDF only*.

By directly comparing these two implementations, I was able to assess which tools best suited the margin, and which best suited the PDF by gaining subjective scores on each from the participants following the study (see Procedure below). In total then, the study contained three interfaces:

- The Original System (set of tools for the PDF, and a set of tools for the margin);
- The Margin System (all tools on the margin only);
- The PDF System (all tools on the PDF Only);

### Procedure

The 16 participants I recruited for the study were selected randomly from a cohort of academic and research staff in the Physical Sciences Department of Swansea University. As with the paper study, participants with a research based background were chosen intentionally due to their increased likelihood of annotation practices in daily life. The age of the participants at the time of study ranged from 20 to 49 and the study itself lasted on average around 45 minutes.

To maximise the amount of information I could gather from the studies, I used screen capture software to record the specifics of each participant's actions during the tasks, and stored the resultant artefacts for later use. Each participant was given a £5 gift voucher as an incentive for participation. The structure of all 16 studies were carried out in the same manner, as described below:

### Study Tasks

The first set of four tasks given to the participants was of a closed nature to gain subjective responses during the interview questions that proceeded them. These tasks were undertaken to allow each of the tools to be used on each of the two interfaces (PDF and Margin systems). The first task in each set used the text box tool. Each text box task required the participant to find an item of information in the document and create a text box note with a summary of the discovered information. One example text box task was **"Turn to Page 4. What was the name of the firm? Create a text box on the PDF and enter this information in it"**. The second part of this set of tasks used the image box tool and asked the participants to add a specific image to a location in the document. A typical task from this portion of the study was **"Turn to Page 8: Add an image of Marley's Ghost on the PDF"**. The third task assessed the pen tool and required the user to draw a particular annotation on the document; a sample question from this was **"Turn to page 7: Change the pen colour to red and its thickness to medium and use it to put a star in the margin next to all references to Serpents"**. The last item of this task set examined the highlighter tool, with the participant highlighting some content and annotating it; one question used for this task was **"Turn to page 7: Use the Highlighter tool to emphasise the word *morose* and change its annotation"**.

The second set of 19 tasks was a mixture of 7 open and 12 closed tasks and was included in the study as a method of getting the users acquainted with the system. The closed tasks again included the creation of text boxes, image boxes and highlights. Some of the closed tasks in this part of the study included: **"Add an Image Box to the page, open the Picture *Tea Party* and resize it to fit in the Notes section"** and **"Turn to page 5: Add a Text box on this page, change its background colour, border font and font size and add in it the Text *Alice at the Mad Tea Party*"**. Some of the

more open tasks in the set were designed to give more flexibility to the participants in using the tools and included tasks such as: “**Mark page 8 as important**” and “**Use the tools provided to identify all the characters that appear on page 4**”.

1. A brief introduction to the study where the participants were introduced to each system and their tools;
2. A set of closed tasks on the PDF and Margin systems: 4 sets which each consisted of 2 closed tasks (1 on each system). These tasks were designed specifically to gain subjective ratings from the users regarding each of the tools and their ideal placement (either on the PDF or in the margin);
3. These closed tasks were immediately followed by a series of interview questions to find out where each tool is best suited;
4. 19 mixed open/closed tasks on the original system (Margin Annotation System) designed specifically to give the participants time to get to grips with each of the tools within the system, as well as its overall functionality.
5. A semi-structured interview based on the 19 tasks including a section of subjective ratings.

To reduce potential bias resulting from the ordering of systems or the content of documents, the order in which each participant used the interfaces was rotated, and the combination of documents to interface for each participant was unique.

#### 4.7.2 Study Results

##### Research Question 1: Is the margin a useful area for making notes?




Figure 4.11: Popularity of the PDF and the margin for each tool (percent)

To investigate whether the margin is a useful area for making notes, I asked the participants to use each tool on both control interfaces (the Margin System and the PDF System). They therefore used each tool on the margin *only* followed by the PDF *only*. Each tool was tested using different PDFs



with similar tasks and followed immediately with a question asking which location each tool was best suited to, (i.e., the PDF or the margin). The results from this section are shown in Figure 4.11.

At first glance, the graph showing these results seems to strongly favour the PDF as the most suitable area to use the tools. However, it should be noted here, that part of the tool set I implemented was designed specifically for use on the document itself and was not expected to perform well when used within a blank open margin space. For example, the pen function, a tool that was used by every participant in the paper study, is a poor choice to mark up the margins of an electronic document as the text box tool provides an easier and neater way of doing so. Similarly, since there is no text in the margins, the highlighter is a rather obsolete tool to use here. I anticipated that both the pen and the highlighter would perform poorly on the margins and was therefore more concerned with the user preference for the text box and image box tools.

Clearly then, the results of this experiment varied dramatically depending upon which tool was being tested, suggesting that each location has an optimum set of tools. 75% of the participants concurred that the Text Box tool was better suited to the margin as opposed to the PDF due to the difficulties of text occlusion on the PDF and the lack of free space available to place the object.

This concern over lack of space extends to the Image Box tool with 69% of the participants agreeing that the margin is the most favourable place to position images. The remaining 31% however, despite agreeing that space is an issue, chose to opt for the PDF as the optimum solution because they felt that they were more in context on the PDF itself.

As far as the Pen and Highlighter tools go, the decision was unanimous – 100% of participants agreed that the PDF is the best place to use these tools. The main reason for this as given by 88% of the participants, is that these tools can mark up the original text without obscuring it and, more importantly keeping it in context.

These results strongly imply that using the margin of the document as an annotation area is not an *alternative* to marking the PDF, but rather a useful *addition* to it.

#### **Research Question 2: Is the mouse-over note function a suitable replacement for the separate medium seen in paper documents?**

Although the problems associated with using a separate medium on paper resulted in only a small proportion of users making use of them in the paper study, the idea behind it is a useful one – you are able to make notes with a potentially unlimited amount of space without obscuring the original text. The highlighter tool has been designed with this in mind with the intention of replicating the positive aspect of marking up using a separate medium (i.e., that it doesn't take up space on the original document) and at the same time eliminating the negative points mentioned above. The end result is a tool that uses semi-transparent 'boxes' to highlight areas of the PDF while also allowing annotations to be 'tagged' to them appearing as a pop-up when a mouse pointer is passed over the highlight itself. This not only eliminates the problem of orphan annotations by ensuring that each one is referenced to their relative highlights but also eliminates the possibility of them getting lost or detached from the original text.

The highly desirable feature of not obscuring the original text, was a major factor in the mouse-over pop-ups popularity, with 10 out of the 16 study participants selecting the highlighter as their favourite tool. In summary then, the mouse-over note function is a popular and more importantly, non-intrusive method of marking up digital documents, suggesting that it is a potentially advantageous replacement to the separate medium used so often in paper documents.

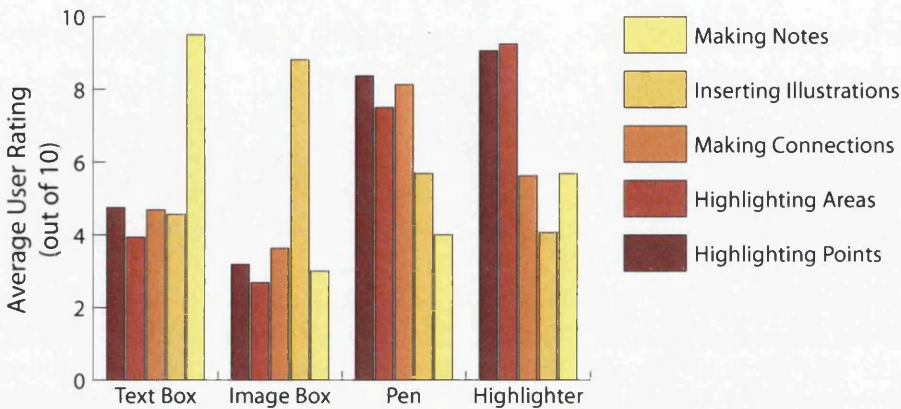
**Research Question 3: What is the complete set of annotation tools?**

Figure 4.12: Average user ratings for each tool (out of 10)

Since this issue is concerned with the tools themselves, I first wanted to ensure that the participants had ample opportunity to make use of each one individually. This was accomplished via the 19 tasks seen in step 4 of the user study procedure. To find out which annotation tasks each tool best supported, I asked the participants to give a score out of 10 (10 being the highest), for each tool based on five types of annotation:

- Making notes;
- Inserting illustrations;
- Making connections between points;
- Highlighting specific areas;
- Highlighting specific points.

I have chosen these five types of annotation as they coincide with the elements described in the anatomy of annotation in Section 4.2.3. The results of this are shown in Figure 4.12. I performed Kruskal Wallis tests on the data shown in the graph and concluded that all sets of ratings were statistically significant. As the graph shows, no single tool is best suited to every attribute, but together the four annotation tools allow users to perform all of the different types of task to a high standard.

The user study concluded that highlighting specific points is best performed by the Highlighter Tool with an average score of 9.00 out of 10 ( $p < 0.0001$ ,  $H = 32.43$ ,  $df = 3$ ) and similarly confirmed that this was also the best tool for highlighting specific areas with an average score of 9.25 out of 10 ( $p < 0.0001$ ,  $H = 33.89$ ,  $df = 3$ ). The Pen was identified as being the most popular tool for making connections between points with an average rating of 8.12 out of 10 ( $p = 0.0047$ ,  $H = 12.95$ ,  $df = 3$ ) and the Image Box was confirmed as the most suitable for illustrating something with an average score of 8.81 out of 10 ( $p = 0.0002$ ,  $H = 19.29$ ,  $df = 3$ ). Finally with an average score of 9.50 out of 10 ( $p < 0.0001$ ,  $H = 31.02$ ,  $df = 3$ ), the Text Box tool has been chosen as the most suitable tool for making notes.

The results of this experiment show that each tool has its optimum usage or set of usages and also confirms the necessity of having a complete set. For example, the absence of the Text Box tool would

leave the highlighter as the optimum method of making notes which, as indicated by its low rating (5.68/10) is not the perfect tool for this task.

**Research Question 4: Are the miscellaneous tools useful? (e.g., OED links, toggleable annotation files etc.)**

This research question was primarily included as a test to determine if there are any miscellaneous functions that are worth further investigation in later studies. To address some of these issues, I set a series of semi-structured interview questions to the participants to gain subjective data on each tool.

The first of these miscellaneous functions was the ability to associate multiple annotation files with the same primary document. To determine if the participants found it useful to have multiple annotation files associated with the same PDF, I included a comprehension style task into the study that makes use of several pre-made annotation files. The answers to these tasks could be found spread across the annotation files provided, resulting in the need to switch between annotations to locate the answers. These specific comprehension tasks allowed each participant to get a taste of the benefits and drawbacks of having independently saved annotations. The questions that follow these tasks reflect the opinions of the users after making use of this facility.

After the interview questions, it became clear that the popularity of having multiple annotation files was high, with the majority of users giving it a practicality score in excess of 7 out of 10. Some of the reasons for this popularity as defined by the participants included having different layouts, more space and straightforward collaboration. More commonly however, the justification for this fondness was the ability to cater towards several different audiences and/or create annotation files with different points of view.

The second of the additional functionalities I tested was the web links box where users are able to select text from the main document or annotations and simply click once to search either Wikipedia or the Oxford English Dictionary. The response from this portion of the study was promising – all participants stated it was a welcome addition to the system, with 3 even selecting it as their favourite tool overall.

## 4.8 Conclusions

In the physical world, there is a lot of space surrounding documents to work with (e.g., margins, wall, desk etc.) but electronically, notes are normally limited to the document area itself. The results of the study on placeholders suggested that additional space surrounding the document area could prove to be a lightweight addition to the active reading process, as the placeholder ‘tabs’ existed statically outside the confines of the PDF. To investigate this issue, I first conducted a paper study to observe and record how and where users make notes on physical documents. This comparison study consisted of two parts in which users were asked to perform annotation tasks on documents with large margins and documents with no margins.

The results of this paper based study indicated that placing notes close to the original material they relate to is important when making physical annotations. As I anticipated, users felt that writing these notes over the PDF itself hindered the reading of the original text, whereas writing them on a separate medium could lead to notes getting lost or dereferenced. The results of this investigation confirmed the importance of margins in the physical annotation process with over two thirds of the participants selecting the margins as their ideal annotation location. This study also suggested that a separate medium could be a useful avenue of investigation if the problems with it were eliminated.

More specifically, a separate medium could prove potentially useful as it does not obscure or damage the original text and offers potentially unlimited space, but suffers from problems such as getting lost or dereferenced.

The results from this paper based evaluation laid the ground work for the design of an improved digital annotation interface. The second phase in my investigation therefore, involved the implementation of a digital reading system with expandable margins. Unlike the placeholders system where the area outside the PDF was static and remained constant despite the currently open page, the Margin Annotation System makes use of *margins* that are linked to specific pages. Therefore, each page has its own separate margin area, the content of which is only visible on the screen when that particular page is open.

In addition, to further investigate the possibility of a separate medium being a useful and lightweight feature, I also incorporated a 'Note' function into the system design. This note feature can 'link' to other PDF tools within the system (highlights and pen marks) and is designed to assess the usefulness of a mouse-over method of displaying document notes. The results of this portion of the study strongly indicated that the mouse-over pop-up action is a lightweight property of active reading design.

### 4.8.1 Lightweight Features

The combined results including observed behaviour and subjective reviews from both the paper based and digitised studies have concluded that:

- Document *margins* are a valuable and popular asset in the annotation process. The results of the digital study strongly suggests that the margins are a useful addition to marking up the PDF as opposed to a straightforward alternative.

In addition to this, the digital study has also strengthened the finding from the previous chapter on placeholder confirming that:

- Mouse-over pop-ups are a lightweight way of displaying information that is not constantly needed. In the case of the Margin Annotation System, it is used as an alternative to the separate medium approach seen in physical annotation practices.

Although the second study confirmed that the Margin Annotation system included a minimum set of complete tools (i.e., every tool was the best at one particular annotation attribute), it may be possible to amalgamate two tools into one to reduce the number further. For example, perhaps the pen tool and the highlighter tool need not be separate, as long as it is possible to change the thickness and opacity, they are essentially the same entity and could therefore be merged into a single tool. Conversely, it may also be possible to create a Post-it style note that also doubles up as a placeholder, much like the physical example in Figure 3.1d. The observations I made in the paper study back up this theory, with several participants using a separate medium as a placeholder as well as to make notes. The action of performing more than one function with a single tool can be loosely described as appropriation [37], a concept that I discuss in depth in Chapter 5.

In summary, the results from this section suggest that:

- The concept of appropriation may be a lightweight feature that can be transferred to the digital level.

Research Question 4 was concerned with the miscellaneous tools within the system that I incorporated in an attempt to probe potentially useful areas of lightweight design. One of these features, was the structure of the saved annotation files. Specifically, user annotations are saved separately to the original document which not only ensures it does not get ‘damaged’, but also allows it to open multiple annotation files. A small addition to the design included a toggleable dialog box that allows easy transition between the different annotation files associated with the open PDF. The subjective results from the participants on this portion of the study were promising as the majority of the participants agreed that it offered several potential benefits to the overall active reading experience.

It is clear from this evidence that saving annotations independently of the PDF ensures that the original document remains unchanged and also allows more than one set of annotation files to be associated with a single copy of the original document. Therefore:

- Saving any user changes to a document separately, ensures the original does not get ‘damaged’ while allowing multiple annotation files to be used and shared easily while only physically storing one copy of the original document.

#### 4.8.2 Related Publications

1. Jennifer Pearson, George Buchanan, Harold Thimbleby : **Improving Annotations in Digital Documents** - *European Conference for Digital Libraries (ECDL) 2009*, Corfu, Greece.

## Chapter 5

# Note-Taking



Be yourself. No one can ever tell  
you you're doing it wrong

---

James Leo Herlihy

### 5.1 Introduction

This chapter builds upon the previous two on placeholdering and annotation. Where in Chapter 4 I investigated the use of *marginal* space surrounding the document area, here I further investigate the use of *static* space seen in Chapter 3: i.e., the space surrounding the document that stays constant despite the currently open page. By studying this area further, I aimed to conclusively prove that the area surrounding digital documents is a lightweight feature integral to the design of active reading software.

In addition to confirming the lightweight-ness of the static document space, this chapter also focuses on another potentially lightweight feature – the concept of appropriation; that is, using some technology in a way that was not originally intended by the designer [37]. It also investigates the idea of direct manipulation [125] to reduce the on-screen menu system.

#### 5.1.1 Overview

When undertaking active reading using physical books, it is common for users to make notes on spare paper, Post-its or other material. This additional media can also be used to mark positions within the book for relocation later, that is, to act as a placeholder. In this situation, the separate medium is dual functioning, that is, it is acting as a placeholder as well as a note-taking facility. On paper this makes sense, for example, it would be unlikely for someone to make notes using a pile of Post-its and a separate pile of bookmarks when one of these could easily do both jobs.

Digitally however, note-taking and placeholdering functions are typically split into two distinct tools which, unfortunately, tend to suffer from a range of usability issues. Digital placeholders for example, are a world away from what we are used to in physical books. As Chapter 3 explains in detail, the way

in which digital bookmarks are presented seriously affects the way in which they are used. In fact, the literature on this topic confirms their low usage within digital document readers as well as web browsers. Furthermore, the digital annotation tools which I discussed in Chapter 4 are also lacking in some areas. For example, the lack of marginal space on the majority of digital document annotation tools hinders the annotation process.

This chapter describes a system that combines placeholding and annotation into a single unified tool as well as incorporating additional static space for notes and a drag-and-drop interface for creation and deletion.

## 5.2 Background

When performing active reading on paper, it is likely that the central document will be used in conjunction with a multitude of other tools and equipment, for example, pens, notebooks, bookmarks etc. This behaviour has been observed by many researchers in the past; Luff et al. [79] for instance, stated:

*“Paper is rarely used in isolation from other artifacts. Most often it is used as part of a collection of various paper documents, writing devices, and other information sources, even digital displays.”*

– Luff et al [79].

Clearly then, the active reading process requires more than a mere document. The use of additional equipment and the availability of a suitable workspace to arrange material, are common features in a paper based reading environment, but something that is all-together different when digitised. On paper, the ability to lay information out in a physical space is extremely important for the active reading process [118], yet digitally, little consideration is given to the provision of an additional work area to make notes.

One of the major aims of this chapter was to provide users with a static workspace that enables them to “*integrate reading with other on-going activities*” [106] and facilitate the spatial organisation that comes naturally when working on paper. By applying physical interactions such as this into the design of an improved digital active reading system, I hoped to provide familiar interactions which will in turn increase user satisfaction with the system.

### 5.2.1 Issues Associated with Digital Note-Taking

#### Space

Figure 5.1, which shows an example of how note-taking is commonly performed on paper, demonstrates how the desk area surrounding the book is being used to keep notes that remain in place regardless of the currently open page. Although this concept of marginal space is seen in some digital systems (e.g., Adobe’s Illustrator has space surrounding its canvas area to place objects), it is not generally incorporated into document annotation software. Since there is no comparable workspace in most digital note-taking interfaces, users are thus forced to make any digital notes within the borders of the PDF itself. I believe this to be one potential contributory factor to the poor uptake of digital annotation tools and one that, if suitably solved, could be considered a lightweight attribute.

### Tool Overload

On paper, it is common for one tool to have multiple functions. For example, a ruler can be used to measure as well as to draw straight lines. This property is known as *appropriation* [37] and can be neatly described as improvisation, adaption and adoption of technology in a way that designers never envisaged. Although the concept of appropriation is a common occurrence in the physical world (e.g., I could use my heavy text book to press flowers or to prop a door open in addition to reading it), it is rarely seen within the digital document domain.

As Figure 5.1 illustrates, Physical Post-its can be used in a manner of different ways:

1. Being stuck to a page within the document, which is useful for making notes about a specific paragraph;
2. Being stuck completely outside the document, in this case, on the desk, which is useful for providing static access to notes independent of the currently open page, e.g. as a book summary or similar;
3. Being stuck to the edge of a page, which causes it to act as a bookmark.

Thus, one tool (in this case, a Post-it) is, in effect, performing two distinct functions – acting as a note-taking facility as well as a placeholder tool, something that in the digital world, is typically achieved with two separate, and often badly implemented functions.

### Menu Navigation

As I mentioned earlier, when a user is engaged in an active reading task, it is vital that they devote as much time as possible to the main task and not to the tools that are being used to facilitate it. Hiding important tools within click menu structures for example, may cause unnecessary attention to be paid to the tools themselves and not the primary active reading task. To help achieve a state of flow with the main task then, it is important to achieve direct manipulation [125] wherever possible by reducing the on-screen menu system.

#### 5.2.2 Scenario

To better illustrate common paper-based work practices, I have given a short scenario that observes the typical workspace of a young student who regularly performs homework tasks.

*Thomas, a 12 year old, is studying a range of subjects in school and comes home every night with homework from multiple classes. In his bedroom he has his own desk where he spreads out his books and assignments ready to complete his homework for the evening. Despite the untidiness of his work-space, Thomas has his own system for remembering and completing homework assignments that can spread over several semesters. For example, on the wall directly above his desk sits a cork pin board that contains numerous items including to-do lists, permission forms and time-tables. On the desk itself sits piles of books and paperwork along with a variety of Post-its that are stuck to drawers, lamps and pen pots. Looking closer at the text books on the desk, you can see scrap paper and Post-its sticking out – usually with hand-written notes and doodles drawn on them. Within text books you can also see Post-its stuck to individual pages with notes made on them, for example, in a history book there is a note on page 72 that says “\*learn 4 test fri\*”.*






Figure 5.1: An example of Post-its being used in a physical book

Thomas' behaviour is typical to many users of physical documents. By utilising the static space surrounding his reading material (i.e., walls, desk, lamp etc.), he creates an environment where he can view important information at a glance. For example, when he is reading a text book and wants to quickly find out when the homework is due, he can simply look up to his timetable on the pin board above him – an action that can be considered significantly lightweight. Alternatively, he may have this information recorded somewhere in a note-book or planner, which would require the extra step of first finding the item and then physically navigating to the correct page – an action that could easily break him out of a state of flow.

When applied to digital documents, this problem persists. For example, when reading a digital document, you are likely to have to switch between programs and/or windows to locate information about your time-table, which is again an action that could easily break you out of your train of thought, particularly, if the information is hidden within menus.

Thomas' methods for making notes and bookmarking pages within text books is also interesting. Specifically, to avoid marking or damaging what is clearly a borrowed school book, he makes notes on scrap paper and Post-its and fixes them close to the area to which they relate. In many cases he also uses these notes as rudimentary page holders, making regular use of Post-its to stick to the sides of important pages.

### 5.2.3 Aims

The main goal of this system was to encapsulate some of the lightweight properties described in the scenario above. Specifically, the static space surrounding the document (e.g., the desk and walls) and

the appropriation example of using scrap paper and Post-its to make notes and act as placeholders. Although screen size is an issue with computerised systems, the space surrounding the main document area can be very useful asset and should be utilised to provide the optimum amount of information to the user. The following lists the main aims of the Digital Reading Desk system:

**Aim 1:** Create an additional margin ‘desk’ space that will allow static notes to be viewed at all times whilst reading;

**Aim 2:** Combine both placeholder and annotation into a single unified tool;

**Aim 3:** Make the PDF look and act like a book; double page spread with notes that ‘flip’ according to the currently open page;

**Aim 4:** Facilitate quick drag-and-drop creation and deletion of notes;

### 5.3 Literature Review

A lot of the literature relating to note-taking is also closely connected to both placeholders and annotation and as such, can be found within the literature review sections of Chapters 3 and 4.

Several of the studies into realistic books have made use of tabbed bookmarks in a similar way to the work presented in this chapter. Chu et al.’s [29] implementation (see Section 2.5.4 for a full review), as well as providing page turning visualisations, also includes bookmarks that appear to project beyond the physical page. However, the position and size of these are standardized and computer-controlled, as opposed to the free-form, user-controlled approach I facilitate via the tools within the Digital Reading Desk implementation. These tabs do however, still provide additional contextual information to the user regarding their position within the document [75]:

*“Participants said they always knew where they were in the realistic books by looking at the left and right page edges. They could see where the current section ended by looking at the bookmark tabs.”*

– Liesaputra and Witten [75].

Card et al.’s [25] 3Book interface (see Section 2.5.4 for more) also includes protruding bookmark tabs in its design. In their case, bookmark tabs are created automatically when a highlight is made by the user, the tabs then appear in the centre of the highlighted space thus restricting where tabs can actually be made.

Some researchers have realised the advantages of physical Post-its and have therefore attempted to produce hybrid sticky notes, that is, systems that allow notes to be made on paper Post-its while offering additional digital enhancements. Steimle et al. [126] for example, produced a system known as ‘Digital Paper Bookmarks’, an interface designed to span both the physical and digital worlds. On the basic level, Digital Paper Bookmarks are simply adhesive Post-its of different colours which can be placed arbitrarily on a physical document to mark pages. These Post-its then synchronise with the digital system via identification of the page number and position of the note, as well as any text added to it. The digital visualisation of the physical notes consequently facilitates collaboration as well as providing customised index systems.

The closest resemblance of Steimle et al.’s work to the work presented in this chapter is the way in which the Post-its are deleted from the system. Due to the physical tangibility of the Digital Paper




Figure 5.2: Design for the Digital Reading Desk

Bookmarks, they are deleted by simply removing them from the pages and drawing a cross on them as opposed to the digital method of menu deletion. It is the lightweight action of discarding notes by removing them from the document that is one of the primary focuses of this chapter.

Mistry and Maes [91,92] have also implemented a system that employs the ubiquitous Post-it note in their system known as 'Quickies'. Using a series of techniques, from RFID readers and handwriting recognition, to natural language processing and artificial intelligence, the Quickies system has been designed to "act as an I/O interface to the digital information world" by bridging the gap between the two mediums. Users of the Quickie interface are able to write notes on special RFID encoded paper with a commercially available digital pen which automatically detects movements on the surface. The text from these handwritten notes is then processed by the system along with contextual information such as the whereabouts of the user. The data gathered then allows the system to make calendar appointments, create new phone contacts, amend or create new task-lists and even send reminders from other users utilising the system.

Although these hybrid systems are not directly associated with the work presented here, the fact that other researchers believe physical Post-its to have intrinsic advantages over any digital equivalents, gives justification for the need to better represent a digital version of the Post-it note.

## 5.4 System Design

Here I present the design of the Digital Reading Desk, a system which aims to provide a series of interactions similar to what can be achieved on a physical desk. Although I am aware that reading on paper does not provide a panacea that I can simply copy to improve a digital interaction, it does offer a proven, effective contrast which should increase the lightweight properties of the design.

### 5.4.1 Static Workspace

The design of the new system includes an additional workspace where notes can be placed outside the main document area while staying constant regardless of the currently open page. The position of this area is shown in Figure 5.2. The incorporation of this static workspace is my solution to the issue of space I discussed in Section 5.2.1.

### 5.4.2 Appropriation

One potentially lightweight digital property could be the idea of appropriation, or more specifically, the designing of tools that can perform more than one function. In the previous study on Annotations, I discovered that many of the participants making paper notes using Post-its, not only used them for writing on, but also stuck them to the sides of the pages to act as bookmarks, two functions that are typically separated digitally. To test this theory, I have designed the Digital Reading Desk system to include a single tool that will perform both placeholding and annotation. This idea of providing a multi-functionality tool is my solution to the problem of tool overload discussed in Section 5.2.1.

### 5.4.3 Drag-and-Drop Manipulation

To promote direct manipulation [125] with the tools in the system, the design of the features does not involve a menu system but rather a drag-and-drop infrastructure. Following closely to Neilson's usability heuristics [101], specifically the "aesthetic and minimalist design", I have ensured that no functions are hidden within menus, choosing instead to create a visual representation of Post-its that are draggable to and from Post-it 'piles' to facilitate easy, paper-like creation and deletion.

## 5.5 System Implementation

### 5.5.1 The Virtual Desk

The virtual desk area can be used to place static notes and has been designed to mimic the additional space surrounding physical books, such as the walls and desks described in the scenario in Section 5.2.2. This 'virtual desk' space, which can be seen in the screen shot in Figure 5.3, is a useful area for placing any notes that relate to the document as a whole (e.g., summaries or character lists), or notes that would be useful to see at all times (e.g., to do lists or timetables). It is useful to note that the desk area belongs to the document itself; therefore opening a new document will give the user a new desk. This feature fulfils Aim 1 (see Section 5.2.3).

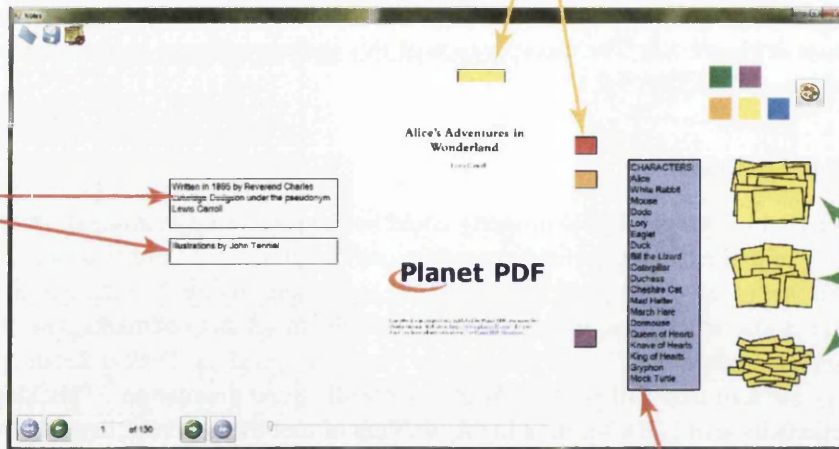
### 5.5.2 The Unified Post-it Tool

The Unified Post-it tool follows closely on from the bookmark 'tabs' seen in the Visual Bookmarking System, that is, coloured markers that protrude from the side of the document like tabs in physical books. The way in which the unified Post-it tool accomplishes two distinct tasks then, is dependent upon *where* they are positioned. Specifically, if they are placed on the document itself they are just notes whereas if they are stuck to the sides of a document, they become bookmark notes.

To encourage more flexibility with the tools in the system, I wanted to ensure that there are no constraints when it comes to where the Post-its can be positioned. Consequently, the Post-its can be placed either:

Post-its on desk behind the book (these will not be visible when the book is opened)

Post-its acting as bookmarks on pages after current page



Post-it Piles

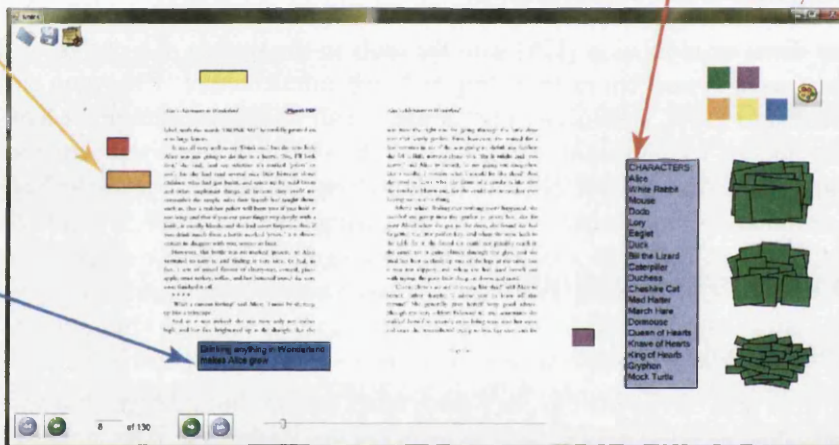
Post-its on desk (these will always be visible)

Post-it bookmark on current page

Post-it on the book itself with a text note added

Post-its acting as bookmarks on pages before current page

Post-it on the book itself that has been shrunk to reveal the text beneath



Post-its on desk behind the book (this is only visible when the book is closed)

Colour palette which changes the colour of the Post-it piles

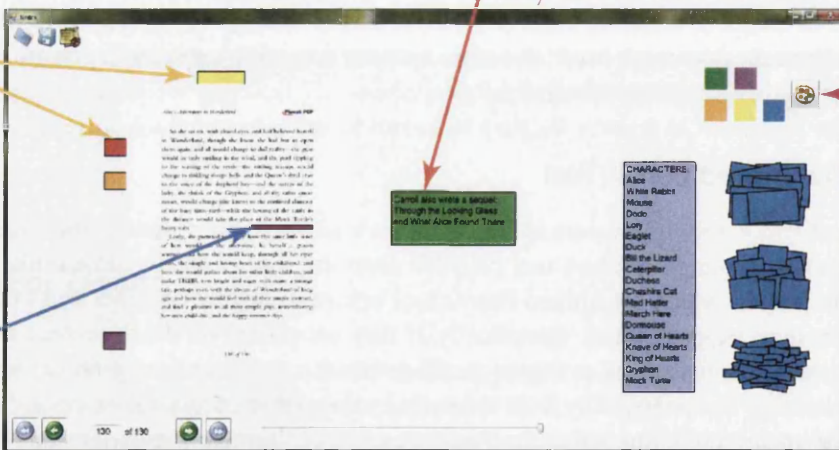


Figure 5.3: Screen shots from the Digital Reading Desk

1. Completely on the document;
2. On the desk next to the document, or if the book is closed, in the area behind the document;
3. Protruding from the document which will cause it to act as a bookmark.

This functionality essentially means that one tool now performs three separate functions: to make notes on specific pages of a document (Point 1), to make notes about the book as a whole (Point 2) and to make notes that also act as placeholders (Point 3). Post-its that also act as placeholders not only navigate to the correct page when clicked, but also ‘flip’ from one side of the book to the other depending upon which page is open. That is, Post-its that are bookmarking pages that are sequentially before the current page, are on the left of the book, whereas those that are on pages that are sequentially after the current page are on the right (much like the Visual Bookmarking System). See Figure 5.3 for screen shots of the system. This is a potential solution to the problem of tool overload which was discussed in Section 5.2.1, and fulfils Aim 2.





### 5.5.3 Drag-and-Drop

To reduce on-screen menu clutter, I use a drag-and-drop style interaction for the creation and deletion of Post-its. On the right of the virtual desk are three inexhaustible Post-it piles that change colour using the palette at the top (see below for details). To create a Post-it, the user drags from one of the piles onto the document, on to the side of the document or on to the virtual desk (Figure 5.3). Notes are removed in the same manner by dragging them back onto the pile which removes the extra heavyweight step of right clicking and selecting delete. As well as the addition of text to the Post-its, they can be moved, resized or ‘lifted up’ (to reveal text underneath), borrowing from and extending the behaviour of physical notes. All interactions are performed without menus: to add text is a double click, ‘lifting up’ is a single click, etc. Removing the need for menus promotes direct manipulation [125] and reduces the attention required for the tools themselves, leaving more cognitive resource for the primary active reading task. This feature fulfils Aim 4.

### 5.5.4 Form Layout and Navigation

My main motivation in designing this system was to improve on the shortcomings listed in Section 5.2.1 by creating an environment that applies physical interactions to digital documents, following closely to the work conducted on the visual book metaphor (see Section 2.5.1). One of the features of this design is a static workspace on which a digital book will be placed. As I was attempting to make a digital version of a physical book therefore, it seemed sensible to ensure the book itself turns and behaves in the same way as you would expect. To mimic paper books as closely as possible, the PDF document that sits on the desk will be a double page display (see Section 2.5.1) that opens in the same way as a physical book, that is, the display shows a double page spread when the book is opened and a single page (either the front or back covers) when it is closed (see Figure 5.3). I anticipated that visualizing document spreads in this manner, would encourage users to treat it more like a physical book and consequently make better and more frequent use of the tools provided. Unlike many of the commercially available reading systems on the market today (see Appendix B), I opted not to implement any page turning visualisations in the design of the Digital Reading Desk as I felt that it may pose an unnecessary distraction to users undertaking active reading tasks.

Clearly this functionality means that there are always (with the exception of the front cover and possibly the last page) *two* pages visible on the screen at any one time – like a double sided printed

book. The navigation system on the Digital Reading Desk system includes ‘next’ , ‘prev’ , ‘first’  and ‘last’  arrows which in actuality skip *two* pages at a time. In addition to this, there is also a scroll bar which allows easy ‘flicking’ through pages in the document and a text entry box that allows ‘jumping’ to specific page numbers. This functionality fulfils Aim 3.

### 5.5.5 Additional Functionality

#### Colour and Size

For lightweight creation of Post-its, the system has been designed so they can be dragged from ‘piles’ on the right hand side of the desk. The three piles represent three different sizes of Post-it that can be quickly dragged on to the desk or document for fast creation. To create Post-its of a different size to the three defaults, the user must first place one of a default size then use the resize tool (dragging the corners of the Post-it) to change its shape. To change the default colour of new Post-its, the user needs only to select a colour from the palette at the top which alters the colour of all the Post-it piles. This consequently changes the default colour of any newly created Post-its.

#### Snapping

To make the addition of bookmarks to a document slightly easier, I have implemented a ‘snapping’ function that pulls a Post-it to the edge of the page if it is within a certain distance of it (i.e., if the user positions the centre of a Post-it within ten pixels of the side of a page), it will ‘snap’ and centre it. This ensures that any bookmarks added can be easily lined up which makes Post-its easier to place and the interface look neater.

## 5.6 User Study

In order to test whether the features (i.e., the static desk, the unified Post-it tool and the reduced menu system) I incorporated into the Digital Reading Desk system could be considered digitally lightweight, I conducted an in-depth comparison study that tests my implementation against two control systems. Before undertaking the study described in this section, I first undertook a small (3 participant) pilot study in order to refine both the system and study designs. The study reported here, was the final evaluative study performed based on the results from this initial pilot.

### 5.6.1 Study Design

#### Control Systems

To make a direct comparison with the Digital Reading Desk system, I implemented two control systems as a basis to test against. These systems have been carefully designed to facilitate direct comparisons between traditional methods of bookmarking and placeholder as well as testing the availability of the static ‘desk’ space.

For consistency, both control systems use the same set of ‘traditional’ tools as shown in Figure 5.4. This traditional tool set has been designed to act in a similar manner to popular digital note taking and bookmarking functions available within commercially available document reader software. For example, the traditional bookmark function shown in Figure 5.4a, mimics the common tree list structure of digital bookmarks found in software such as Adobe Reader (see Figure 3.2b in Chapter 3 for a screen shot of the Adobe Reader’s tree-based bookmark list). The traditional notes function (Figure

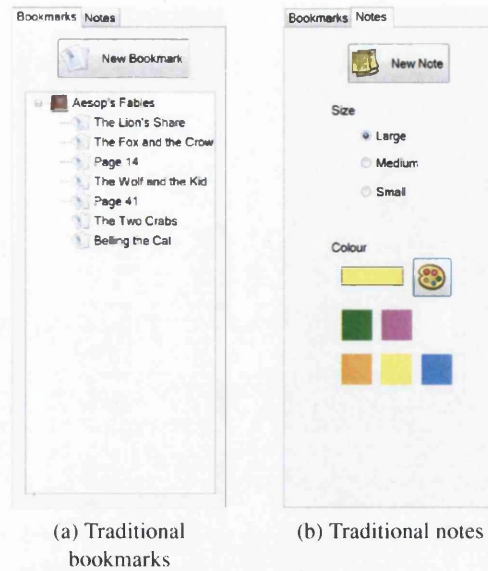


Figure 5.4: The bookmark and notes interfaces on the Traditional PDF and Traditional Desk control systems

5.4b) has also been built on traditional digital note taking methods.

**Traditional PDF:** The first of the control systems, which I named ‘Traditional PDF’ (see Figure 5.5), is a simple system modelled on conventional PDF readers such as Adobe Acrobat or Apple Preview. Post-its can be used on this system but only as notes (i.e., they do *not* double up as bookmarks) and are created and deleted in a more traditional manner, using dialog panes. For example, to create a note the user must first select its size and colour from options on the right of the window then click the ‘New Note’ button to add it to the PDF (see Figure 5.4b). To delete a note the user must right click over the note and then select ‘Delete’. Unlike the Digital Reading Desk, there is no desk area surrounding each document and therefore notes cannot be placed outside the confines of the PDF. Once notes have been added to the PDF however, they contain the same basic functionality as those on the Digital Reading Desk system (i.e., they can be moved around, collapsed, resized etc.). Bookmarks in the Traditional PDF system are a completely different entity to notes and are stored in the same tree structure that is used in most digital document reader software (see Figure 5.4a). To add a bookmark, the user must navigate to a page then click ‘New Bookmark’, and to delete or rename, they must use the right mouse button to access a short content menu. Note that the bookmarks available in this system are entirely user created and do not include any pre-defined bookmarks that may have been generated by the author or publisher (e.g., chapter headings, table of contents).

**Traditional Desk:** The second control system used in the comparison I called the ‘Traditional Desk’ (see Figure 5.6) and has been included to act as an intermediate between the Traditional PDF and Digital Reading Desk systems. Essentially, the Traditional Desk system is the Traditional PDF system with the addition of a desk area. Notes and bookmarks are in the same format as the Traditional PDF system (i.e., it uses the traditional tool set in Figure 5.4) and the only difference is that now notes can be added to the desk as well as to the PDF itself.





Figure 5.5: The Traditional PDF system

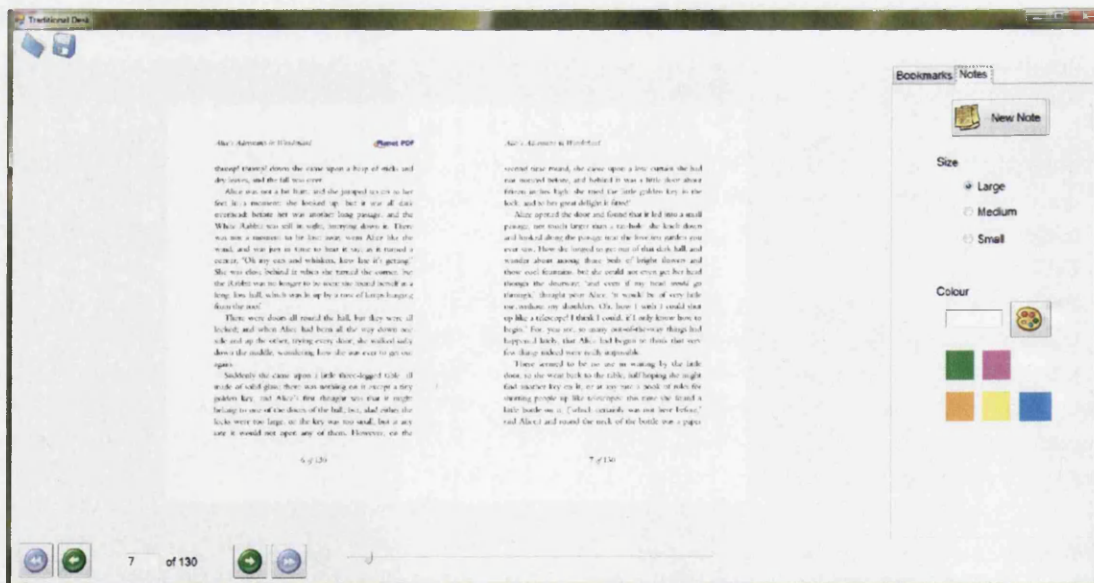


Figure 5.6: The Traditional Desk system

### Procedure

I recruited 16 participants (12 male and 4 female) between the ages of 21 and 41 all with a high skill level (post-graduate degrees and above) to take part in the study. Each study was conducted in a quiet, well lit room with sufficient natural lighting and took on average, an hour to complete. The equipment used in the study consisted of a standard mouse and keyboard input with a 17inch wide-screen display. No video or audio recordings were taken during the study, although I did observe and take notes of user behaviour during the tasks. Screen capture was not possible, as this proved incompatible with the PDF rendering software I used to implement the system. The final annotated files of each session were anonymously stored for later use. Each participant was given a £5 gift voucher as an incentive for participation and all 16 studies were carried out in the same manner, as described below:

1. Participants were given a brief overview of the experimental procedure and a short demonstration of each of the systems;
2. Following this, each participant undertook 3 sets (one for each system) of nine closed tasks (T1 – T9) which were designed to help them familiarise themselves with the tools and systems.
3. In order to observe user behaviours while undertaking less restrictive closed tasks, each participant completed an open-ended assignment on each interface;
4. After completing the study tasks, each participant filled out a short questionnaire which probed their current note-taking behaviour as well as providing subjective ratings on each system's features: five overall questions (numbered Q1 – Q5) and questions regarding each of the nine closed tasks (numbered T1 – T9 to link directly to the nine tasks). Each rating used a five point Likert scale from strongly disagree (1) to strongly agree (5);
5. Finally, a semi-structured interview was conducted that questioned each participant on every aspect of the systems including their experiences with other note-taking techniques on paper and digital.

To reduce any possible learning effects or bias, the order in which the systems were presented to the participants as well as the order of the task sets were rotated between users.

### The Tasks

As I mentioned above, the study comprised of two main sections: closed tasks (T1 – T9) and a larger open task. These tasks were performed on each of the three interfaces: Traditional PDF, Traditional Desk and Digital Reading Desk systems.

The closed tasks were designed as learning aids and covered all the tools included in each system, permitting a detailed measurement of specific low-level interactions. First I asked each participant to perform nine closed tasks on each system. These tasks posed varied slightly depending upon the exact PDF used, but contained the same basic function tasks. These tasks have been specifically designed to correspond to the questions T1 – T9 shown in Table 5.1. A typical closed task for a system within the study was: **T4: Create a note for page 6 that says “The name ‘Scrooge’ has entered the English vocabulary as a synonym for a miser”.**

To obtain better information on ‘natural’ user behaviour, the second part of a user's session with each interface was a larger open task. The open task encouraged the participant to mark up the documents in the way that they found the most appropriate for the job at hand. I anticipated that

any genuine difference in the interaction styles would result in distinct behaviours and artefacts from a participant's use of the system. The open tasks used three separate PDFs, rotated between the interfaces for balance. Each PDF document was near the same length (around 8 pages) and the tasks had a common structure. A sample task from a document entitled "Bio-fuels: Implications for food and agriculture" consisted of: **Read the article carefully, finding strong points for or against bio-fuels, as well as any open questions that occur to you or any significant statistical data. Mark up the article with notes to help you explain your interpretation of the text to another person. Summarise the article with a brief paragraph that would explain your impression of it.** These open tasks were designed to give the user more freedom in their approach to note-taking using the systems and to aid with their subjective responses. They were also used to observe each participant's natural note-taking behaviour and provide many of the patterns of use results I describe later.

### Study Documents

The PDF documents used in the study were carefully chosen to be of the appropriate length for the task at hand, as well as being of a suitable topic for discussion. In total, I used six different documents in the study. The first three were for use with the closed task set, contained the same font type and size, and averaged at 108 pages in length. These closed task PDFs were of classical children's novels (A Christmas Carol, Alice's Adventures in Wonderland and Aesop's Fables), chosen due to their relative lengths and the ease with which the content can be read.

The final three PDFs used in the study were designed and created by myself, for use specifically on the open task set. The topic of these three documents were varied within non-computer related topics and were chosen intentionally to include multi-opinion issues. The titles of the three PDFs used in the open tasks were: **Bio-fuels: Implications for food and agriculture**, **Video games and their effect on childhood obesity** and **Tanning Beds: is bronzed skin worth the risks?**. These PDFs were written in a clear and concise manner and were all exactly 8 pages in length. The reason for the short length of document in the open tasks was solely due to the type of task being performed. Specifically, each participant was asked to read each of these documents carefully and make notes about them. I felt that giving documents any longer than this would be counter productive to the main task of note-taking.

### Selection of Study Metrics

One major challenge I faced in designing a study such as this, is that it is often complex to precisely assess the benefits of a new interactive paradigm. Ultimately, my aim was that my new interaction would encourage a change in the pattern of user behaviour, underpinned by lower mental effort, rather than make substantial time savings (see Section 2.6). In advance of the study, pilot studies using standard HCI literature on direct manipulation demonstrated that time savings, whilst possible, were relatively small and, more importantly, unlikely to significantly change the patterns of user behaviour that have been seen as shortcomings of digital reading, for example, low rates of use, rather than problems centred on speed. In fact, if my interaction does indeed promote note-taking, I may actually see the opposite effect - that the time taken to produce notes will increase due to an increasing number of annotations being made. If the prior research correctly ascribes higher levels of annotation to more attentive reading, a more 'successful' design will encourage readers to spend more time writing their notes.

Mental effort is also difficult to assess objectively, so I hoped ultimately to identify differences through the artefacts of open tasks. By studying *how* and *where* the participants make use of notes as

well as the number of notes made, I was able to assess the usefulness of the different aspects of the system, e.g., the virtual desk area.

Using artefacts as an evaluative method is problematic however. In the case of this study for example, would a higher level of note-taking activity reflect a superior interface? Alternatively would more note-taking in fact indicate that reading is being interrupted and users are calling more on notes to support a reading task that has been made more complex? In order to obtain a principled model for my study, I have drawn strongly upon the techniques used by Shipman et al. in their analysis of VKB [121], who faced similar problems where the correct interpretation of artefacts and metrics for the task was imperfectly understood. Two ways in which this work has impacted the design of my study is the method in which subjective scores were obtained from the participants, and the way in which I present the results. Specifically, I split the subjective scoring into two parts, the five general questions (Q1 - Q5) regarding each system, followed by nine task related scores (T1 - T9) which focus on specific aspects of each interface. Like the work of Shipman et al., I have presented these results in a single table showing the average scores and statistical p values.

### 5.6.2 Hypotheses

Each aspect of the system and its controls have been carefully designed to prove specific assumptions made about features within the system. I have discussed these hypotheses below:

**Hypothesis 1:** The tools and interface of the Digital Reading Desk system will be more popular than those of the more traditional control interfaces (Traditional PDF and Traditional Desk systems);

**Hypothesis 2:** Drag-and-drop bookmarking (Digital Reading Desk system) is easier to use than traditional tree-listing of bookmarks (Traditional PDF and Traditional Desk systems);

**Hypothesis 3:** Incorporating both notes and bookmarks into one unified tool (Digital Reading Desk system) will make the interface easier to use than two separate ones (Traditional PDF and Traditional Desk systems);

**Hypothesis 4:** Allowing users to place notes on a desk (Traditional Desk and Digital Reading Desk systems) will make it easier to make notes about the book as a whole (Traditional PDF system).

**Hypothesis 5:** A visual approach to creating and deleting notes (Digital Reading Desk system) will improve their ease of use.

### 5.6.3 Study Results

The following section reports the findings from the comparative study described above.

#### Post-it Usage

The first part of the post-study questionnaire is concerned with each participant's current note-taking practices, in particular, their use of Post-it notes and scrap paper. The first part of this investigation involved the frequency in which the participants make use of Post-its and/or scrap paper in their everyday lives. As Figure 5.7 shows, the participants I recruited for the study represent a diverse range of Post-it/scrap paper users, with the majority of participants using them frequently (i.e., at least a couple of times a week). This Figure also indicates that the desk or wall is a common place to position notes – of all the participants that make frequent use of notes (i.e., at least a couple of times a week), 81%




Figure 5.7: The frequency of Post-its and scrap paper used by the participants in everyday life

place them on their desk or wall, compared with 50% who use them in books as notes or bookmarks. Such a high proportion of users placing notes on the area surrounding the document gives further justification for the addition of a virtual desk area into my design and poses the question as to why current digital document systems offer no comparable space to place notes. This is also backed up by subjective comments made by the participants during the post-study interviews: *“I often use Post-its to make quick calculations. I leave them on the desk then refer to them later”* and *“I use my notebook quite often to write down page numbers and paragraphs of interesting stuff I’m reading. Obviously I’m not going to write this on the actual book. It’s nice anyway to have a list of all the books I’m reading and all the pages that are useful on one sheet”*.

### Subjective Findings

Table 5.1 reports the participants average subjective ratings for both the general (Q1 – Q5) and specific task (T1 – T10) questions. To assess the statistical significance of the results obtained, I performed a Kruskal Wallis test on the data collected. I chose this test over the more traditional ANOVA due to the relatively small number of discrete values available, that is, the five point Likert values. The following subsections report these results in detail.

#### General Questions (Q1 - Q5)

The first general question I posed to the participants was if they found the interface as a whole, easy to use. The results of Q1 yielded average Likert scores of: 4.75 for the Digital Reading Desk compared to 3.56 and 3.94 for the PDF and Desk systems. Kruskal Wallis test results concluded this result was highly significant:  $p=0.0003$  ( $H=16.28$ ,  $df=2$ ), proving that the Digital Reading Desk interface is easier to use than the PDF and desk alternatives.

Similarly, Q2 and Q3 both produced significant Kruskal Wallis results of  $p=0.0037$  ( $H=11.2$ ,  $df=2$ ) and  $p=0.042$  ( $H=6.33$ ,  $df=2$ ) respectively, concluding that the Digital Reading Desk’s individual tool set is also superior to those included in the control systems.

	Traditional PDF		Traditional Desk		Reading Desk		Kruskal Wallis (p)
	Average	SD	Average	SD	Average	SD	
<b>Q1:</b> I found the interface easy to use	3.56	0.89	3.94	0.57	4.75	0.45	<0.0003
<b>Q2:</b> I found the tools easy to use	3.50	0.89	3.81	0.75	4.56	0.63	<0.004
<b>Q3:</b> I found the tools easy to <i>learn</i> to use	4.25	0.77	4.25	0.77	4.81	0.40	<0.05
<b>Q4:</b> I would use this system out of choice	2.50	1.03	2.81	1.05	4.06	1.18	<0.002
<b>Q5:</b> I think this system mimics paper well	2.44	1.03	3.06	0.93	4.31	0.79	<0.0001
<b>T1:</b> I found it easy to create new bookmarks	4.00	0.97	3.94	0.93	4.63	0.62	<0.05
<b>T2:</b> I found it easy to create new notes for a <i>specific page</i>	3.88	1.15	4.06	0.99	4.75	0.45	<0.05
<b>T3:</b> I found it easy to look up old bookmarks	4.31	0.95	4.31	0.95	4.31	0.60	=0.8694*
<b>T4:</b> I found it easy to create new notes for the <i>book as a whole</i>	2.81	1.42	4.38	0.89	4.63	0.62	<0.002
<b>T5:</b> I found it easy to look up old notes	2.63	0.96	2.63	0.96	3.63	1.26	<0.04
<b>T6:</b> I found it easy to collapse and expand notes	4.00	1.09	4.00	1.09	4.00	1.09	=1*
<b>T7:</b> I found amending old notes easy	3.94	1.06	3.94	1.06	3.94	1.06	=1*
<b>T8:</b> I found it easy to delete notes	4.25	0.93	4.31	0.87	4.63	0.80	=0.3679*
<b>T9:</b> I found it easy to delete bookmarks	4.19	0.98	4.25	0.93	4.56	0.63	=0.5975*

Table 5.1: Average subjective ratings (five point Likert). \* indicates no significant result was found

As well as proving that the Digital Reading Desk prototype was more popular than the alternative control implementations (see Section 5.6.1), I also wanted to determine if the participants felt it could be put to everyday use. Q4 produced promising Likert values of 4.06 for the Digital Reading Desk compared with only 2.5 and 2.81 for the PDF and Desk systems. Unsurprisingly, these results were proved significant in the Kruskal Wallis statistical test:  $p=0.0013$  ( $H=13.36$ ,  $df=2$ ). This result is backed up by user comments such as “*the last one [Digital Reading Desk] is pretty awesome, it’s a big improvement over Preview for the Mac*” and “*it [Digital Reading Desk] is cool - it makes logical sense, like real books*”.

The last of the general questions, is concerned with how well the Digital Reading Desk prototype mimics the behaviours of paper. Although imitating the properties of paper may not be a perfect solution to the problems faced with digital mark-up tools, the results show it may enhance usability in some areas. Q5 then, yielded highly significant results:  $p<0.0001$  ( $H=19.57$ ,  $df=2$ ) and Likert values of 4.36 for the Digital Reading Desk compared with only 2.44 and 3.06 for the PDF and Desk systems respectively.

### Specific Tasks (T1 - T9)

The first task, T1 confirmed that users find it easier to create new bookmarks with the Digital Reading Desk’s drag-and-drop method rather than the tree bookmark method with average subjective Likert scores of 4.63 for the Digital Reading Desk interface compared with 4 for the Traditional PDF and 3.94 for the Traditional Desk. Comments made by the participants during the post-study interviews strengthened my findings here: “*I prefer being able to drag and drop – it relates more to real life*”, “*it’s more intuitive, you don’t need to teach it because I already know how*”, “*dragging is much better, I’d rather drag than click - it’s far more fun*”.

After conducting pilot studies on the prototype, it became clear that users tend to make two types of note: those that are specific to a page (e.g., a note about a paragraph), and those that are generally for the document as a whole (e.g., a book summary). Where users place notes is usually an indication of their type. For example, notes about specific pages are usually placed on the page they relate to, whereas notes about the document as a whole are often placed either on the front cover, or on the desk or wall to enable persistent access to the material. With this in mind, I decided to split the tasks probing new note creation into two distinct questions: T2 and T4. I anticipated that the latter (T4) would elicit useful feedback on the desk portion of the system – the observed behaviours and artefact placement produced by these tasks are discussed later in this section. Subjectively, user ratings for the creation of notes for a *specific page* produced an advantage for the Digital Reading Desk interface, despite high Likert scores for all systems. The Kruskal Wallis test produced a statistically significant result of  $p=0.05$  ( $H=5.18$ ,  $df=2$ ).

For creating new notes for a *book as a whole*, Kruskal Wallis showed a clear distinction ( $p=0.0011$ ,  $H=13.62$ ,  $df=2$ ) between the two desk implementations (Digital Reading Desk and Traditional Desk) compared to the Traditional PDF, with Likert values of 4.63 and 4.38 for the Digital Reading Desk and Traditional Desk systems compared with only 2.81 for the Traditional PDF. This result strongly suggests that the desk feature provides a useful area for positioning notes about a document as a whole.

Another aspect of creating notes which tasks T2 and T4 explore is the method of creation. Specifically, the three user study systems exploit two separate interactions for note creation: the two traditional systems (Traditional PDF and Traditional Desk), that use a click button interface (seen in Figure 5.4b), and the Digital Reading Desk system that employs drag-and-drop. Both tasks T2 and T4 show differences between these two interactions and are backed up by user responses which included,

“drag-and-drop was much easier to create new notes in different colours and sizes [...] you could see what it would look like before you place it” and “dragging and dropping notes is more intuitive, why have menus when a picture [piles of Post-its] is better?”.

Despite the differences in interaction between the bookmarks on the systems, task T3 resulted in identical ratings for all interfaces and therefore did not yield a significant Kruskal Wallis result. I conclude from this that there is no significant performance problem with the look up feature of my new bookmark interaction as the average Likert rating for all three systems was 4.31 out of 5.

In the two control systems (Traditional PDF and Traditional Desk) the bookmark and note features were split into two distinct tools, while in the Digital Reading Desk interaction they were merged. With this in mind, it seemed appropriate to ask in addition to looking up old bookmarks (T3), how easy users found it to look up old notes (T5). The results of this task produced a significant result of  $p=0.0319$  ( $H=6.89$ ,  $df=2$ ), and Likert results of 2.63 for the Traditional PDF and Traditional Desk systems compared with 3.63 for the Digital Reading Desk. Although several participants suggested that an improved note look up system (e.g., an additional togglable list) would prove beneficial, these results clearly illustrate the popularity of the drag-and-drop interaction (Digital Reading Desk) compared to the more traditional implementations (Traditional PDF and Traditional Desk).

Tasks T6 and T7 produced the same average Likert score for each system and thus had no significant results from the Kruskal Wallis test. This data however, was the result of tasks where functionality was the same across all three systems, and similar results are both to be expected and confirm participant neutrality.

The final tests, T8 and T9 yielded no significant Kruskal Wallis results with  $p$  values of 0.36 and 0.59 respectively. These task scores however, resulted in high Likert values across the board (all above 4 out of 5). I can thus conclude from this, that the drag-and-drop method of deleting notes and bookmarks is of a similar (and relatively low) complexity, compared to the more traditional ‘menu-and-click’ method.

### Patterns of Use and Behaviour

As well as the post-study interviews and questionnaires, I also evaluated the participants’ use of the three systems through observations of their interaction in the open tasks and analysis of the final marked-up documents. From this data, I have identified a range of interesting mark-up behaviours.

Firstly, the use of bookmarks differed significantly between the systems: 75% of participants used them on the Digital Reading Desk system, while only 25% and 19% used them on the Traditional PDF and Desk systems respectively. This large increase in bookmark use strongly suggests that the unified Post-it tool actually encourages the creation of bookmarks. I conducted a Chi-squared test on the bookmark usage data and concluded its statistical significance:  $\chi^2 = 7.684$  ( $df=2$ ,  $p=0.02145$ ).

Secondly, the uptake of the desk area was promising, with many participants making use of it when available (i.e., on the Digital Reading Desk and Traditional Desk systems). The behavioural data from users placement of document summaries is one example of this. Table 5.2 reports the results from task T4 (I found it easy to make notes about the *book as a whole*) as well as the summarising behaviour I observed in the open tasks. To assess the statistical significance of this data, I performed a Chi-Squared test with Yates’ corrections (primarily due to the low frequency of ‘back page’ option). Not surprisingly, the difference between the Digital Reading Desk and Traditional Desk systems is not-significant as the ‘desk’ area is identical on both systems. However the results of these systems versus the PDF system produced a highly significant result of  $\chi^2 = 26.93$  ( $df=4$ ,  $p=0.00000565$ ). It is clear from these results, that the desk is the preferred position for document summaries. There is further support from the post-study interviews: “I always make summaries separately so being able



		Traditional PDF	Traditional Desk	Digital Reading Desk
Desk	Front Page	75%	19%	31%
	Back Page	16%	3%	3%
	Page Two	6%	0%	0%
	Other	3%	0%	0%
Book	Desk (always visible)	N/A	69%	57%
	Desk (behind first page)	N/A	9%	6%
	Desk (behind last page)	N/A	0%	3%
	<b>Desk (Combined)</b>	<b>N/A</b>	<b>78%</b>	<b>66%</b>
	<b>Book (Combined)</b>	<b>100%</b>	<b>22%</b>	<b>34%</b>

Table 5.2: Summary placement (% of Participants)

to put them on the desk is nice as you can always see it and make notes as you go” and “the desk is good because it is always in arms reach”. Some circumstances in which a desk area was said to be useful included: storing summaries of whole books rather than specific pages, notes to guide reading (e.g., ‘read this first’, or a task reminder), for common definitions and character biographies.

### Hypotheses

**Hypothesis 1: The tools and interface of the Digital Reading Desk system will be more popular than those of the more traditional control interfaces (Traditional PDF and Traditional Desk systems)**

The results from the general subjective questions (Q1 - Q5) confirm Hypothesis H1 and comments such as “the Digital Reading Desk system was so intuitive, it acts as real world experience would lead you to expect” and “it fits in with the book style and look - I think even my mother could use this” strengthened the statistical findings.

**Hypothesis 2: Drag-and-drop bookmarking (Digital Reading Desk system) is easier to use than traditional, tree-listing of bookmarks (Traditional PDF and Traditional Desk systems)**

Figure 5.8 reports participant preferences for bookmark presentation: traditional (tree-list used in the Traditional PDF and Traditional Desk systems) versus drag-and-drop (unified Post-it tool used in the Digital Reading Desk system). While 69% of the participants rated the unified Post-it tool to be slightly or much better than the traditional tree list, this result was not statistically conclusive ( $p=0.12$ ). The result for Hypothesis 2 is therefore supportive, rather than conclusive.

While tab style bookmarks are more popular overall, participants reported useful properties of the tree bookmark list, particularly when looking up notes. One commented “creating and reading should be drag-and-drop, and browsing should be a list”, indicating that the overview of a tree-list is a better alternative for reviewing notes than the visual tab interface, while the Digital Reading Desk interface is better for creating them. It is likely that an improvement to my interface can be made by incorporating a hybrid system that includes draggable notes spliced with a togglable tree-list to view bookmarks. In fact, I identified the potential for a hybrid bookmark implementation while interviewing a participant during the pilot study, which consequently prompted me to add a question regarding

it into the post-study interviews of the main 16 participant comparison. The results of this question concluded that 100% of the participants would prefer to have a system which included both drag-and-drop notes coupled with a list for ease of navigation.

**Hypothesis 3: Incorporating both notes and bookmarks into one unified tool (Digital Reading Desk system) will make the interface easier to use than two separate ones (Traditional PDF and Traditional Desk systems)**

Current digital note-taking interfaces typically make use of two distinct tools for placeholding and annotation, whereas the Digital Reading Desk combines them into a single unified tool. As the results show, the functionality of both the annotation and placeholder functions have not been significantly compromised by this change. My goal now however, is to identify whether or not users like the idea of a single tool which performs two functions and indeed whether they agree with Hypothesis 3.

Comments made by participants on this topic included: *“it [the unified tool] was smooth, it simplifies the interface without sacrificing functionality”*, *“it’s easier to act instinctively with the one tool as you are always reaching for the same things”* and *“it [the unified tool] makes life easier. In real life you wouldn’t have a pile of Post-its and a pile of bookmarks, you would just use a Post-it for both”*. Participants also observed that having one tool for two jobs makes it easier to turn notes that have already been created into bookmarks and vice versa, thus confirming Hypothesis 3.

**Hypothesis 4: Allowing users to place notes on a desk (Traditional Desk and Digital Reading Desk systems) will make it easier to make notes about the book as a whole (Traditional PDF system)**

The results from tasks T2 and T4, plus the data from participants’ summary placement (Table 5.2) confirms Hypothesis 4, that placing notes on a desk will be easier for making notes about a whole book.

**Hypothesis 5: A visual approach to creating and deleting notes (Digital Reading Desk system) will improve their ease of use**

Although the final two tasks (T8 and T9) produced no significant results, there is evidence that the drag-and-drop deletion of Post-its is more popular than the traditional menu based deletion method. While participants found both methods easy to use, they had a preference. Figure 5.8 presents a graph of users’ preference for menu based compared to drag-and-drop annotation (menu create and delete versus ‘drag-and-drop’ create and delete). The results from this portion of the study show that the majority (81%) of the participants felt that the drag and drop method of annotation (i.e., creating and deleting notes) was much better than menu based methods, (i.e. right click - delete). This is backed up by the highly significant result obtained in a Chi-squared test with Yates’ corrections:  $\chi^2 = 35.875$  (df=4, p=0.0000045). This result is also supported by comments such as: *“the drag-and-drop system was much the same as paper - chucking a note back on the pile is like screwing it up and throwing it away”* and *“dragging notes back onto the pile is much simpler and more intuitive than fiddling with menus as this breaks me out of my train of thought”*. This evidence, coupled with the results from tasks T2 and T4, proves Hypothesis 5.



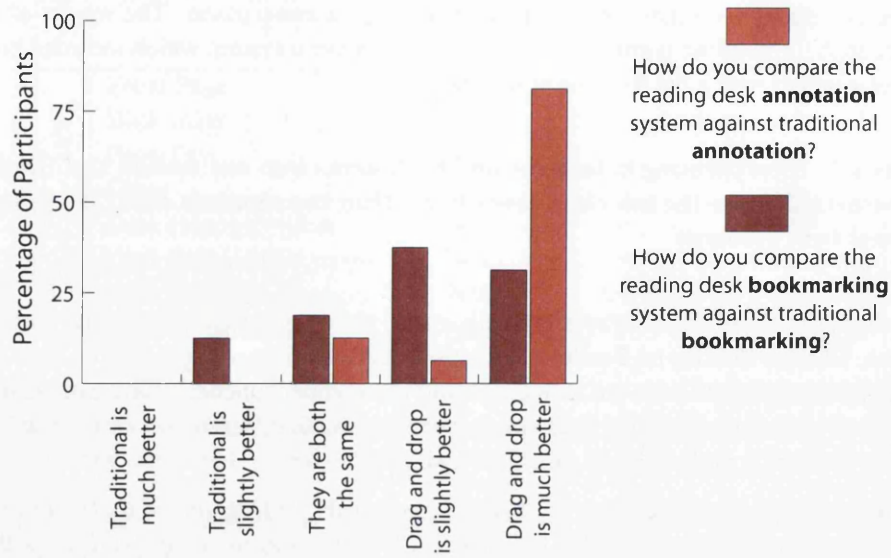


Figure 5.8: Subjective reviews of the annotation and bookmark portions of the systems

### Digital Versus Physical

Although I am not attempting to blindly reproduce the complete interaction of physical documents in my design, I do feel that incorporating certain paper based features with digital techniques will increase user satisfaction of the tools. How well the participants felt the system replicated the lightweight properties of paper is a research question that has been answered through evaluation of the qualitative results obtained via the subjective ratings and post-study interviews.

The results from Q5 of the subjective ratings given in Table 5.1, shows a strong preference for the Digital Reading Desk system with an average score of 4.3 out of 5 compared with only 2.4 and 3.1 for the Traditional PDF and Traditional Desk systems respectively. This highly significant result ( $p < 0.0001$ ,  $H = 19.51$ ,  $df = 2$ ) confirms that the additional design features I introduced into the Digital Reading Desk system (i.e., the virtual desk area, the unified Post-it tool and the drag-and-drop creation and deletion of notes) are fulfilling the book metaphor more completely than the two control systems. To further investigate each participant's views on paper versus digital mark-up, I also asked a specific question in the post-study interview: **How do you think the open tasks would compare with the Digital Reading Desk system if they were performed on paper?**

This question resulted in split views from the participants. Some could see the benefits of the digital solution, providing comments such as: *"It's easier to annotate digitally - real notes fall off, get lost and there's sometimes no room"*, *"Digitally is better because I won't lose my Post-its"*, *"Typing a note is a much better option than writing it by hand for people with bad handwriting"*, *"It's quicker digitally as I type faster than I can write"* and *"I usually use a notepad to make notes about a document as I'm reading it, but this [Digital Reading Desk] is better because you can make notes on the desk as you go"*. One participant even went as far as to say that there were no differences between the Digital Reading Desk system and a paper solution: *"Paper is no different to the last system [Digital Reading Desk], as notes can be attached to sides of pages"*. Other participants however, still prefer the physicality of paper when marking up: *"Computers are always more reusable but paper is nicer"* and *"It's probably faster on paper"*.

## 5.7 Conclusions

This chapter has described the design and implementation of the Digital Reading Desk – a system that builds on the previous two chapters on placeholding and annotation. By following the book metaphor and by taking ideas and thoughts from previous investigations, I have created a system that incorporates both placeholding and annotation into a single unified tool. The Digital Reading Desk system also incorporates a large workspace comparable to a physical desk or wall, for the addition of static notes, and a direct manipulation interaction style to reduce the time and effort required to create and delete Post-its.

Although the system mirrors some of the properties of paper, it also avoids its natural limitations: for example, the Post-it piles are unlimited and can be resized and changed in colour. Participants' comments on the topic of how the Digital Reading Desk compared to paper, included "*The Digital Reading Desk system is better than paper because the notes always stick the sides of the pages and don't damage the book.*" and "*You can't lose electronic notes and you never have to buy more because the piles are infinite.*"

To test the effectiveness of the different features of the system, I conducted a user study to compare it to two baseline implementations. The data resulting from the comparative study has provided evidence that there are several advantages of the Digital Reading Desk system, illustrated by differences in user behaviour during the open tasks.

Firstly, the observed behaviour in the study, and analysis of the subjective feedback, confirmed the utility of the desk area and consequently proved Hypothesis 4. The study also identified several key advantages of the unified tool, including its ability to support bookmarking and note-taking simultaneously. Participant behaviour and feedback both supported and confirmed Hypothesis 3. Participant comments included, "*I liked it [the unified Post-it tool], it's more similar to reality and it's less confusing to have one thing*". Participants also expressed their preference for the bookmark 'tab' system, an enhancement of the placeholding implementation in Chapter 3. As with the Visual Bookmarking System, users can track pages of interest by means of coloured tabs that stick out of the sides of the document. Unlike the earlier design however, which used page order to determine the visual position of a bookmark (i.e., bookmarks at the top are on lower page numbers than those on the bottom), the Digital Reading Desk system, by allowing users the freedom to drag notes anywhere, loses this sense of order. This may be a contributory factor for the 100% participant view that a hybrid bookmark interface with both drag-and-drop tabs and togglable tree-list features would be beneficial.

I did not expect significant differences in the time performance of adding notes in the various interfaces, nor did I anticipate that a drag-and-drop interaction would be dramatically easier to use than a menu based method. Rather, that the different interaction styles would result in a greater use of notes and placeholders in the unified tool, particularly in open tasks. This difference was observed in the study. The evidence gathered is not, I believe, conclusive proof that a single tool for annotation and placeholding is the optimal design, but there is a clear improvement over having two separate tools. It is my conjecture that a single tool will become more familiar to the user and through this will require less cognitive attention. It will require further experimentation to prove this conclusively.

### 5.7.1 Lightweight Properties

One of the primary goals of this chapter, was to conclusively determine whether the space surrounding the document area can be considered a lightweight property. The Visual Bookmarking System first implied this possibility after a study into the usefulness of the static bookmark 'tabs' proved to support lightweight interaction. To further investigate whether the space surrounding a document, in this case

the *marginal* space, was a positive attribute, I conducted an investigation into marginal annotations. This study concluded that margin space is a useful addition to annotating over the document itself as opposed to a straightforward alternative, confirming its lightweight status.

To confirm the suggested *static* workspace property is also a lightweight property then, the Digital Reading Desk implementation included an extended ‘desk’ area that allows users to make notes around the open document that stay constant independently of the currently open page. The results of this investigation confirmed the popularity of the desk area as well as positive differences in the patterns of use observed by participants in the open tasks, therefore:

- The *static* workspace surrounding the document is a useful area to display constant information. For example, notes and summaries about the book as a whole, or in the case of the Visual Bookmarking System, placeholder tabs.

The second major goal of this chapter was to investigate the concept of appropriation [37] and if designing for it can be considered a lightweight attribute. To achieve this, I opted to take two distinct digital functions (in this case, placeholdering and annotation) and amalgamate them into a single tool to mimic the way they behave on paper. The subjective results from the post-study interviews coupled with the patterns of use observed during the open tasks, strongly suggest that the unified Post-it tool is an improvement over more traditional digital placeholdering and note-taking techniques. Therefore:

- Designing for appropriation, if conducted in a way that is familiar to the user (e.g., functions that are currently appropriated on paper) can be considered to be a lightweight property.

As well as these features, the Digital Reading Desk system also incorporates the idea of direct manipulation [125] by creating a drag-and-drop style interface for the creation and deletion of the unified Post-it notes. The idea behind this feature lies within the Post-it ‘piles’ which I have designed to facilitate easy and familiar interaction from the users. In the physical world, users can just pick up a Post-it from a pile on their desk and place it on the book they are reading. Similarly, when the note becomes obsolete, the user can simply discard it by putting it back on the pile. To create a comparable interaction digitally then, I have removed the need for any ‘right click’ actions by allowing Post-its to be dragged directly to and from their piles. The effectiveness of the drag-and-drop creation and deletion of Post-its within the Digital Reading Desk system was confirmed in the comparison study. Therefore:

- Removing the need for menus wherever possible promotes direct manipulation, and is therefore less likely to distract from the primary active reading task or interrupt any state of flow that may have already been achieved.

Finally, the visual overview that I discovered in the placeholders implementation was also detected, albeit to a slightly lesser extent, in the Digital Reading Desk interface. Although the Post-it bookmarks have lost a portion of their ordering when compared with the tabs from the Visual Bookmarking System (i.e., their vertical position no longer indicates where they are in relation to other bookmarks), they do still provide an overview of where each one was positioned with regards to the currently open page (i.e., tabs on the left are on pages sequentially *before* the current page, whereas tabs on the right are on pages sequentially *after* the current page). This overview was rated favourably by the participants in the study. Therefore:

- The visual Post-it bookmark tabs are a useful way of giving an overview of where each bookmark is placed in relation to the currently open page.

### 5.7.2 Future Work

The information I gathered during the post-study interview has highlighted areas for improving the current design of the Digital Reading Desk system. For example, several users commented that the two-dimensional display of the text - particularly the edge of the book, made it hard to see where the Post-it bookmarks were positioned within the document (i.e., which ones were close to the front/back). To remedy this shortcoming, a visual cue could give a better indication of where each bookmark is placed in relation to others in the document. In the physical world, a number of subtle cues combine to give such information, for example, the differences in shadowing of Post-its with changing depth, and touch. However, literal reproduction of this digitally may be either ineffective or impossible. Potential digital solutions to this issue include adding dots to the navigation bar, similar to the interface presented by Byrd [21], or a graph visualization as suggested by Harper et al. [54].

It would also be extremely beneficial to conduct a more in-depth study of the system, when using larger documents and over a longer period of time. This would obtain a more naturalistic view of how users make use of the tools in their everyday mark-up tasks. A diary study over several days or weeks, or a longitudinal log analysis could also contribute to such a study.

Taking into consideration that when engaged in a reading task, many users will work with multiple documents, it may be worth considering the availability of the desk area. At present, the desk is connected with the PDF that is open, meaning that opening a new PDF will give you a fresh desk. Although this could be a useful feature, it may also be worth considering the need for a persistent 'meta-desk' that would mimic physical practices more closely, as some tasks may require the need for notes to be kept across multiple documents [83].


### 5.7.3 Related Publications

1. Jennifer Pearson, George Buchanan : **Improving User Interaction in Digital Books** - *Books Online 2009*, Corfu, Greece.
2. Jennifer Pearson, George Buchanan, Harold Thimbleby : **The Reading Desk: Applying Physical Interactions to Digital Documents** - *Computer Human Interaction (CHI) 2011*, Vancouver, British Columbia, Canada.
3. Jennifer Pearson, George Buchanan, Harold Thimbleby : **The Reading Desk: Supporting Lightweight Note-taking in Digital Documents** - *Theory and Practice of Digital Libraries (TPDL) 2011*, Berlin, Germany.
4. Jennifer Pearson, George Buchanan, Harold Thimbleby, Matt Jones : **The Digital Reading Desk: A Lightweight Approach to digital Note-Taking** - *Interacting with Computers*, Elsevier.



## Chapter 6

# Visual Indexing



Shoot for the moon. Even if you miss, you'll land among the stars

Les Brown

### 6.1 Introduction

This chapter focuses on the concept of indexing, an area of active reading that is relatively heavy-weight in both its physical and digital forms. To improve on the digital indexing process, I have implemented the Visual Indexing System, an interface that incorporates techniques from both the physical and digital domains in an attempt to make the overall electronic indexing process more lightweight. My aim here then, was to illustrate that paper is not a panacea; there are problems with physical documents that can be overcome in the electronic world by applying lightweight digital techniques.

#### 6.1.1 Overview

A common activity when undertaking research, and indeed within everyday life, is the process of locating relevant information within documents. Whether it be searching a text book for a particular topic or looking up a product within a store catalogue, searching for information is a tiresome but necessary part of the information retrieval process. In physical books, the process of recovering relevant information is supported greatly by the ubiquitous back-of-book index system – a classic, well-known structure that neatly lists key terms for easy document navigation. When books are digitised, this feature can be enhanced by means of hyper-links, but often this is not the case. The limitations of this type of document navigation are clear: they are time consuming to physically navigate to, and are also author-created, therefore restricting them to static terms that exist when the book is created.

In the digital world, tools can be used to aid in the recovery of useful information. For example, the text search feature can be used to quickly navigate to instances of a particular keyword within a document. However, this function can be extremely time consuming if the resulting list is long and gives no visual overview of where each keyword occurs over an entire document.

This chapter describes the design, implementation and evaluation of a system that combines the visual overview of the classic index, with the speed and convenience of digital searching.



## 6.2 Background

### 6.2.1 The Four-Phase Framework for Text Searches

There are several steps one must take in order to find a particular piece of information. Schneiderman et al. [123, 124] define these steps in their four-phase framework for text searches:

1. Formulation (what happens before a search takes place);
2. Action (starting the search);
3. Review of results (the results from the initial search);
4. Refinement (what happens after the review of results, before the user goes back to formulation with the same information need)

In this model then, the user must first decide on their information need, that is, what it is they want to find out, as well as where the best place is to search for it. This is known as the **formulation** phase and is considered by Schneiderman et al. to be the most complex as it involves multiple levels of cognitive processing. Following completion of this phase, users will then move on to the **action** phase which involves physically starting the search. This stage is typically achieved in recent years by typing in a keyword/phrase into an internet search engine and clicking 'Search'. Next, the **review of results** phase is where users search within each resultant document to assess its relevance in relation to their original information need. Finally the **refinement** stage is conducted, often by means of relevance feedback, to offer support for successive additional queries.

This chapter is concerned with the **review of results** phase. Specifically, the query in relation to *individual retrieved documents* (i.e., within-document searching) as opposed to the *retrieved documents in relation to each other* (i.e., multi-document searching).

### 6.2.2 Current Methods

Before I discuss the design and implementation of a new digital within-document information retrieval method, it is important to review the methods that are currently being used to locate relevant information in documents in both printed and digital media.

#### Printed

To locate specific information in printed books it is traditional to consult the back-of-book index. The following scenario describes a common situation where the index is used to locate relevant information within a large document (please refer to Figure 6.1 for specifics on the index used within this scenario).

*Wende wants to buy a new backpack, and decides to choose one from a popular product catalogue. To locate the product she requires, she first turns to the back-of-book index under 'B' to find the entry marked 'Backpacks', which contains an instruction to look under the more general 'Bags' section. Luckily, in this case the 'Bags' entry appears immediately after 'Backpacks' and is therefore easy for her to find. Once she has located the sub-entry for 'Backpacks' she then starts the process of looking at the index entry to find all the pages in the catalogue that contain backpacks. Thinking strategically, Wende realises that although the first entry is on page 467, the*

*next entry is clustered (i.e., 1004-1009) and is therefore more likely to be the backpacks section, rather than a single occurrence within a different product segment. She starts her search by locating page 1004 and browses sequentially to page 1009. While she is performing this action however, she keeps one finger in the index page which acts as a temporary placeholder and allows her to easily flick back and forth between the index and catalogue content.*

It is clear from this scenario that there are benefits and drawbacks to printed indexes. The advantages of the traditional index structure include:

- An alphabetised list of topics;
- A good overview of where areas of interest are located within a document;
- The clustering of pages (e.g. 3-10) shows relevant ‘hot-spots’ within the document.

Despite these advantages however, the printed index also suffers from many potential problems:

- It takes time to physically look up references;
- The possibility of un-indexed indexes;
- If an entry contains multiple page references, the user must flip back and forth from the index until the correct information is found;
- Circular references that never lead to page numbers, for example, **canine: see dog**, **dog: see canine**;
- References from a single topic distributed over several meanings, for example, **canine: 12, 35-42**, **dog: 45, 56-59, 73**;
- Inappropriate inversions, for example, **Woofers, Sub;**
- Indexing from section headings as a result of computer created indexes, for example, ‘Working with Threads in Java’ being indexed under ‘W’ as opposed to ‘T’ for Threads or ‘J’ for Java.

In the early 2000s, professional indexers Macmillan Computer Publishing conducted a usability study to determine how well indexes meet the informational needs of their users [113]. There are many known issues with indexes, that make creating them a difficult process. Despite triple posting entries, providing cross references or code concepts, Macmillan still identified problems during their observational study. Among other things, they were surprised at what users were looking up, that is, that many users were attempting to look up entries that publishers never would have thought of indexing. Macmillan concluded from this that different users take very different approaches to finding information in indexes – a major problem during creation, as anticipating the needs of everyone is virtually impossible. Clearly then, despite their benefits, statically created indexes still suffer from problems that could be overcome by allowing users the ability to create their own user-created, custom indexes, a theme which is central to the work presented in this chapter.

Athletics Supports.....	980	Backpacks.....	<i>see bags</i>
Audio		Bags	
• Accessories.....	1385	• Backpacks.....	467, 1004-09, 1611
• Audio Speakers.....	1463	• Camera/Camcorder.....	1204-5, 1213-18
• Home Cinema Systems.....	1408-11, 1357	• Children's Luggage/Bags.....	1020-1
• Storage.....	528-32, 1466	• Holdalls.....	1002-09, 934-36, 1146
Audio and Video Tapes.....	1385, 1466	• Ladies'.....	1014, 1006/7, 1146/7
		• Luggage.....	1013-21, 1003, 1009
		• Rucksacks.....	891
		• Sports Bags.....	1004-09, 1146/7, 1611
		Bakeware.....	712
		Bangles and Bracelets.....	42-53, 22, 31-33, 65, 93-95
		Bar Stools.....	241
		Barbecues.....	866
		Baseball and Basketball.....	938/9

## B

Baby and Nursery Equipment.....	1489-1524, 881
Baby Gifts/Clothing.....	22, 1145
Baby Monitors.....	1505-09, 1516

Figure 6.1: Snippet from a popular catalogue index

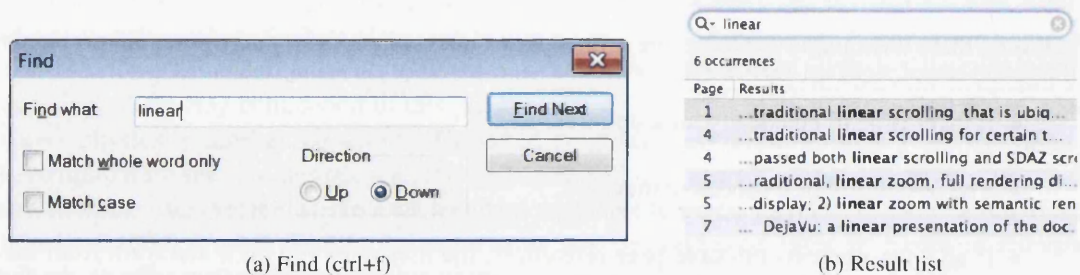


Figure 6.2: Current methods of digital searching

## Digital

When books are digitised, occasionally the index entries within it are hyper-linked, allowing easy navigation to indexed pages. Unfortunately however, even hyper-linked indexes suffer from the problem of flicking back and forth between index and book content. In fact, this problem is even more heavy-weight on digital media as fingers cannot be used to hold pages, and the 'flicking' action is performed with the scroll bar.

The most well known method of locating information within digital documents is to use the text search or 'Find' (Ctrl+f) feature (see Figure 6.2a) which linearly scrolls through every instance of the search term until no more occurrences can be found. Despite being fast in execution, this method can be extremely time consuming if the result list is long, as every instance of the search term is systematically accessed one mouse click or keyboard shortcut at a time. This traditional text search method also gives no overview of where within a document each search term occurs making it difficult for the user to pinpoint dense areas of interest.

An alternative to the constant clicking of the find feature, is a partial result list like the one used by Apple's Preview (see Figure 6.2b). Instead of linearly scrolling through every occurrence of the search term, Preview returns a single list of matches along with a snippet of the text surrounding the word itself.

### 6.2.3 Aims

The motivation for this system was to take something that is currently heavyweight on physical and digital documents, and improve it by incorporating features from both mediums resulting in a lightweight electronic solution. With this in mind, I constructed a set of aims for the Visual Indexing System, an interface designed to combine the speed and specificity of digital search with the visual overview of the printed index:

**Aim 1:** To allow users to create custom hyper-linked indexes from their own keyword or phrase;

**Aim 2:** To create a dedicated index panel to eliminate the need for flicking back and forth using the heavyweight scrollbar;

**Aim 3:** To give a visual overview of the number of text matches on a page/page cluster (i.e., the relevance of that page) that shows dense areas of keyword occurrences at a glance.

These aims have been carefully thought out to address specific problems of current index methods and form the basis of my hypotheses for the Visual Indexing System. Aim 1 for example, if properly implemented, will greatly reduce the time taken to look up entries, as well as overcoming the problem of un-indexed entries. Aim 2, in providing a dedicated space that is always visible on-screen will eliminate the need for flicking back and forth between the index and book content. Finally, Aim 3, will show dense areas of keyword occurrences at a glance, which should in turn increase the speed and accuracy of text searches.

## 6.3 Literature Review

Although there has been much research into general document retrieval, there has been relatively little exploration into within-document search visualisation. The following section describes several relevant bodies of research that focus on visualising within-document search results, focusing specifically on how they compare to the Visual Indexing System.

Harper et al. [53] created a system known as SmartSkim which is based on relevance profiling – a technique that generates a profile of showing where key terms exist across a whole document. The system uses histograms to give a within-document representation of relevant passages within a document based on a particular user query. Within each histogram, each bar corresponds to a fixed length section (known as a ‘tile’) of text in the document whereas the Visual Indexing System will split the document up by *pages*. The histogram tiles do make slight use of colour, but only as a means of indicating which bars had been visited, currently being visited and yet to be visited, *not* as I propose, to indicate the relevance of a particular section.

As a follow up to their 2002 paper, Harper et al. produced a slightly improved implementation known as ProfileSkim [54], which they went on to evaluate against a text search style baseline system called FindSkim. The results of this study confirmed that the visual approach to relevance profiling (ProfileSkim) is more efficient than the Ctrl+f (FindSkim) method and also suggests that Profile Skim improves precision.

Byrd [21] builds upon Shneiderman et al.’s [124] WInQuery system that displays relevance scores for multiple documents, by creating a scrollbar based visualisation for *within-document* information retrieval. The system covers the vertical scroll-bar with multi-coloured ‘dots’ that correspond to keyword occurrences within the document. When the user enters a new keyword into the search, the program assigns it a coloured dot and places it on top of the other terms within the scroll-bar.

A key of search terms and their assigned colours is then displayed on the bottom of the window. Although it can be difficult to distinguish between colours if there are a lot of search terms, it does allow users to identify clusters of a keyword occurrences by identifying areas of coloured dots in close proximity.

In 2006, Hoerber et al. first introduced the HotMaps system [60,61] which uses a visual representation to display query term occurrences using colour on a 'heat' scale. Although this system has been designed to aid users in general information retrieval, like my system, it uses visual communication to help the user target the most relevant material more quickly. In particular, it uses colour to indicate the number of times a word occurs within each document. More specifically, each search term is given a colour depending upon the number of times it occurs within a document, that is, multiple occurrences of a query term (a 'hot' result) will be dark red in colour whereas a small number of occurrences of a term (a 'cold' result) will be lighter in colour. What this system does not facilitate however, is the visualisation of search terms within a document; the Visual Index System builds on this idea of colour gradients to give users an overview of relevant sections within a document itself.

This concept of using shading to aid the information retrieval process by indicating the frequency of keywords within a document was also identified in Hearst's 1995 paper [55]. Her research describes the TileBars system which, like HotMaps, uses a gradient of colour to specify the number of occurrences of a query term in a list of retrieved documents. By using a series of coloured rectangles, she aims to simultaneously display the relative document length, keyword frequency and keyword distribution of a set of documents to help users decipher which ones are the most relevant. As with HotMaps, the visual information is presented to the user at the search result list level. The work I have done here studies search within a document, and also compares different methods for presenting this information.

Research has also been conducted into using word clouds to visualise the relevance of keywords within a document. Grotton [50] for example, created a within-document information retrieval system that uses TF/IDF<sup>1</sup> values combined with visualisation techniques borrowed from the web's tag clouds to enhance the perception of web documents. The word cloud system visualises what they describe as 'word importance' values within a document which are independent of user query. The study conducted on this system proved that participants found particular areas of interest 10% faster using the word cloud technique than by using ordinary text.

The Xlibris [115] system makes use of colour shading within its 'skimming' mode to indicate word importance within a document. Meaningful terms are identified within the document via standard TF/IDF scores and are assigned a shade of grey or black to indicate this information to the user. For example, terms that occur frequently in the current document but rarely in others (meaningful words) are coloured black, whereas those that occur in many documents (common words) are coloured light grey.

As well as within-document visualisation, there has been research into *how* to search within different types of digital document. Liesaputra et al. [78] conducted a study into searching for relevant information within different document mediums. They tested both back-of-book indexing and full text search on three different document types: realistic book format [77], HTML and PDF. The study investigates the influence of document format as well as search result presentation on the information search process. The results from the study showed that participants found information more efficiently and effectively with the subject index than full-text search, yet the post-study questionnaire revealed that the majority still preferred text search over the subject index. Despite the effectiveness of the index then, users of digital documents are more inclined to use text search. This is yet more evidence

---

<sup>1</sup>Term Frequency/Inverse Document Frequency

to suggest that a hybrid system allowing custom indexes to be built from user-defined search terms would be a useful method of within-document searching.

Chi et al. [28] have implemented a system that conceptually *reorganises* electronic document indexes based on the users' information need, but do not specifically *create* on-the-fly indexes like the Visual Indexing System. Their implementation, ScentIndex, is designed to narrow down the author-created indexes of electronic documents to the most relevant entries for display on a single page, as opposed to the complete index structure that commonly spreads over multiple pages. As with the Visual Indexing System, users start their search by entering their own keywords into the system. ScentIndex then uses a word-occurrence matrix to extract the most relevant indexes from the list based on the users original search term. A user study undertaken to evaluate the ScentIndex system concurred with the main findings of the user study I undertook on the Visual Indexing System described later in the chapter. Specifically, that tasks can be performed faster, more accurately and with higher user satisfaction using custom index systems than with more static, author-created back-of-book indexes.

## 6.4 System Design

The following section describes the design of the Visual Indexing System, an implementation that aims to replicate the lightweight overview from the physical back-of-book index, with the speed and efficiency of digital text search.

### 6.4.1 Index Builder

The foundation of the system is an index builder which allows users to create custom indexes from their own keyword or phrase. It then displays the relevant page numbers (i.e., the ones that have at least one occurrence of the keyword or phrase) in chronological order. I have not applied any stemming algorithms to the system hence the resultant indexes contain page results for exact text matches of the word or phrase in a document.

The following sub-sections describe the functional features included in the Visual Indexing System. These features include the clustering and hyper-linking of pages within the index structure, as well as the visual features I utilise in several of index representations I discuss in the next chapter.

#### Page Clusters

One useful feature of the program is its ability to 'cluster' page hits. For example, when looking for all references to sofas in a catalogue, the index may look something like this:

Sofa: 345, 467, **1067-1098**

That is, it will group major sections instead of listing each page separately. This not only simplifies the output but also allows users to see the sections of the document where there are the most occurrences of the keyword or phrase they entered.

#### Hyper-links

Each of the search results in the system is the form of a hyper-link which when clicked will take the user to the appropriate page and highlight all occurrences of the keyword or phrase on that page.

### Tool Tips

In addition to the visual cues (see next Section) that illustrate to the user the most relevant pages with respect to others in the list, the program also has the facility to view the exact number of text matches on a page by hovering the mouse over a hyper-link.

### 6.4.2 Visual Features

To complement the index builder software, I have applied several techniques to enhance the visibility of the word occurrence data.

#### Colour

Colour plays an important role in the field of information visualisation. There have been many examples of how colour has been used to illustrate differences within and between documents. Bier et al. [13] for example, use different shades of colour to visualise which documents in a book-plex collection have been read. Colin Ware [138] in his 2004 book on this topic, describes an application known as *nominal information coding or labelling*, which describes how best to classify different sets of elements using colour as distinct identifiers.

In his book 'Information Visualization: Perception for Design', Ware talks about colour-coding conventions in which he describes the meanings of certain colours in different cultures. For example, in western culture green is often associated with the word 'go' whereas white is considered to be the colour of purity. In relation to the design for the new system, I decided upon two distinct colours that have conventional meaning to act as maximum and minimums for word occurrences. A gradient of colours (see Figure 6.3) between these bounds is used to illustrate occurrence values that fall within these values. The colours chosen for the labels then, are based on a temperature gauge, that is, red for a 'hot' result (large number of occurrences of a query term) and blue for a 'cold' result (small number of occurrences of a term). Any results within this range will be coloured purple (i.e., a mixture of both red and blue), the shade of which is determined by how close they are to the maximum or minimum values. This type of gradual colour change based on 'hot' and 'cold' colours was also used in Alexander et al.'s Footprints Scrollbar [7].

It is important to note that this gradient of colours varies depending upon the maximum number of occurrences of the keyword or phrase but will be relative to that particular index term only. That is, a bright red link in one index term will be the maximum number of times it occurs but may well be significantly more or less than a bright red link in a different index term.

#### Size

Size can be a useful indicator to visualise the relative weights of different attributes, that is, something with a high weighting will be larger than something with a low weighting. In this case, it is desirable for users to see, at a glance, an overview of the most relevant pages within a document based on their keyword or phrase. It seems sensible then, to allocate pages with a small number of word occurrences



Figure 6.3: The colour gradient for the labels ranges from blue (cold) to red (hot)




Figure 6.4: Design for the Visual Index System

(a potentially irrelevant page) a small size whereas those with a high number of word occurrences (a relevant page) a big size.

### 6.4.3 Form Layout Navigation

One of the main problems with the common index structure in digital documents is the amount of scrolling required to navigate between the index and document content. If the index is not hyper-linked, to look up several references the user will either have to memorise all entries, or scroll to the first look-up, then back to the index, then to the second look-up, then back to the index and so on. On physical documents this process is relatively lightweight in that fingers can be used as temporary placeholders so pages can be ‘flicked’ back and forth while looking up several index entries. Digitally however, the process of scrolling is cumbersome and time consuming [107].

To reduce the amount of scrolling required in the program I opted to create a static index panel on the main interface that can always be seen while reading the document. This dedicated index panel is positioned to the right of the main document area and contains a set of user created indexes that can be easily deleted for a clutter free environment.

In addition to the index panel, and the main document area, there is a toolbar which will contain a set of radio buttons to control the type of visualisation that will be applied to the index (see Section 6.5.2 for details on the visual representations included in the program). The page layout of the document in the system is *single page continuous display* (see Chapter 2.5.1). This presentation has been included to allow quick and continuous scrolling across the entire document; a useful feature when scanning several pages for highlighted words. Figure 6.4 shows the design for the Visual Index System.

### Alternative Designs

In this chapter I have enumerated a number of different designs that either mirror or extend the traditional back-of-the-book index. One could take alternative approaches that dispense with that metaphor, and draw more directly from techniques such as information visualization. Marti Hearst’s



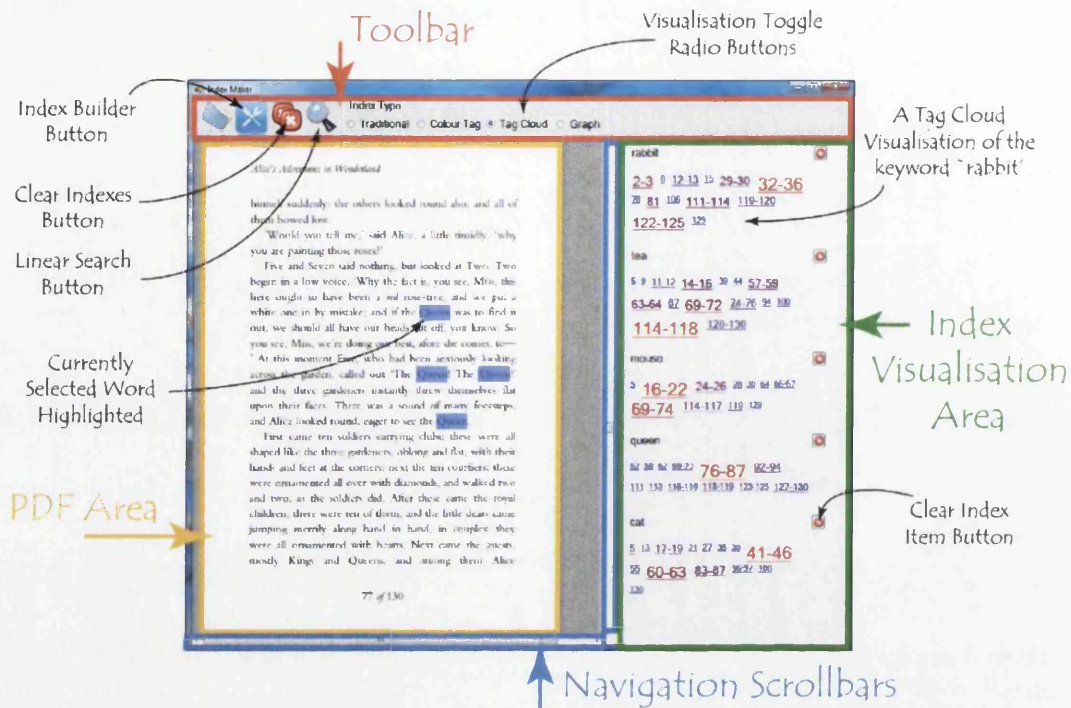


Figure 6.5: The Visual Index System showing the tag cloud visualisation

TileBars [55] would be one approach that breaks from the direct representation of the document body as text, and provides a more abstract overview. The general principles of the Footprints scrollbar [7] could also be explored as a metaphor for highlighting the parts of the document that best match the user's search.

At the more localised view, the presentation of the different formats could be varied. One simple example being the use of different colours, and a more complex area would be to explore different visual structures for the tag sets.

## 6.5 System Implementation

The following section describes the implementation of the Visual Indexing System, an interface designed to incorporate the speed of digital text search with the visual overview of the printed index.

### 6.5.1 Page Relevance

One of the main features of the Visual Indexing System is its ability to visualise page relevance. The relevance of any page within a document, is highly dependent upon the number of occurrences of a particular keyword or phrase within a particular page, relative to the rest of the pages within a document. That is, any page with high occurrences of the user-defined keyword is considered to be more relevant than pages with low occurrences of the keyword and are therefore visualised differently to indicate this difference.




Figure 6.6: Tool-tip that provides the user with exact world occurrence data for the label

To calculate the relevance of a page, the program first iterates through the entire document searching for exact matches to the search term, storing the occurrence values for each page in an array. After this is completed, it then runs through the array clustering any pages that are numerically adjacent and contain one or more occurrence of the search term. Next it calculates the maximum and minimum occurrence values for that particular index (Note: any pages that have no occurrences are not included in the array so occurrence values start at 1 not 0). The program then uses this information to calculate the size and colour (see Section 6.4.2) of the resulting hyper-link.

```
public Color CalculateColourOfLink(int Occurrences)
{
    int ActualRange = MaxNumberOfOccurrences;
    float ScaleValue = (float)255 / (float)ActualRange;
    float NewVal = ScaleValue * (float)Occurrences;
    // normalise occurrences so it lands between 0 and 255
    int ColourVal = (int)NewVal; // casts float value to int
    Color ReturnColour = Color.Black; // default colour
    ReturnColour = Color.FromArgb(ColourVal, 0, 255 - ColourVal);
    // use this normalised value to create a colour value between red and blue
    return ReturnColour; // return the new colour
}
```

Listing 6.1: The method that calculates the colour of a link using occurrence data

The code in Listing 6.1 shows the method for creating a colour from an occurrence number (passed in as a parameter). First it normalises the occurrence value so that it lands between 0 and 255, that is, the bounds for the RGB value of a colour. Then it uses the format  $RGB = (NormalisedValue, 0, 255 - NormalisedValue)$  to create a returnable colour. For example, a 'hot' link will be bright red and will have an RGB value of (255,0,0), whereas a 'cold' link will be blue and have an RGB value of (0,0,255). Anything in between, that is, a page that has an occurrence value between the maximum and the minimum will have a mixed value e.g.  $RGB = (100, 0, 155)$ .

There are also methods to calculate the size of both the tag cloud and graph visualisations which use a variation of the method shown in Listing 6.1.

## 6.5.2 Visual Representations

Once the occurrence data has been calculated for a particular keyword, the program builds the physical indexes which will be placed in the 'index visualisation bar' shown on the right in Figure 6.4. Items within the index visualisation bar are rendered in one of three ways (see Figure 6.7) and can be easily toggled by using the radio buttons on the toolbar. The purpose of these visual representations is to:

1. Display the pages/clusters that contain the user defined keyword or phrase, like a traditional index;

- Clearly illustrate the relevance of each page/cluster (i.e., how many occurrences of the keyword it contains) to give a clear representation of which pages are most relevant.

The basic structure of the index builder and the base of all the visualisations is a chronologically ordered list of hyper-linked page numbers/clusters. When the program builds an index, it first calculates the occurrence information (see Section 6.5.1), sifts out any pages that do not contain any instances of the keyword or phrase, then clusters adjacent pages. All links are then added to the index in numerical order and all instances of the keyword(s) are highlighted on the PDF. Each hyper-link is also equipped with a tool-tip feature that ‘pops up’ a bubble containing occurrence information when the mouse passes over it. Figure 6.6 shows an example tool-tip for the keyword ‘Mouse’, clearly illustrating the page cluster (16-22) and the number of occurrences of the word ‘Mouse’ between these pages (27).

Using the basic visual features I discussed in Section 6.4.2, I implemented three visual representations of this basic index builder design: Colour Tag, Tag Cloud and Graph.

### Colour Tag

The Colour Tag tool (Figure 6.7a) makes use of colour as a visual feature. Although the size of each hyper-link remains the same, the colour is changed according to its relative keyword occurrence value (see section 6.4.2 for details of the colour gradient used).

### Tag Cloud

The Tag Cloud tool makes use of both size and colour to aid in the visualisation of keyword occurrences. Tag clouds are traditionally chronologically ordered, size weighted hyper-links. In this system, each page or page cluster is ordered by page number as in the Colour Tag tool, but is now sized according to their relevance (i.e., the number of occurrences of the keyword relative to the other pages in the document). In addition to this, the colour gradient used in the Colour Tag tool is also utilised to further illustrate page relevance.

### Graph

As with the Tag Cloud interface, the graph visualisation also makes use of size and colour to indicate page relevance values. In this case however, each hyper-link is represented by a coloured bar in a graph – the larger the bar, the more times the keyword or phrase occurs within its page or cluster.

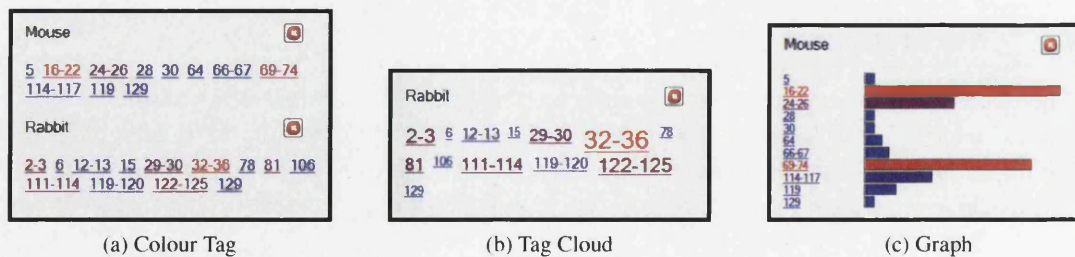


Figure 6.7: The visual index representations

## 6.6 User Study

To test the effectiveness of the Visual Indexing System, and to tease out any potentially lightweight attributes, I undertook a comparative user study on the three visual index interfaces. Unlike the user studies in my previous investigations however, I have used both speed and accuracy (among other things) as criteria for evaluating this system. I made this choice as I felt that unlike my previous studies, the act of locating information within a document is currently relatively slow if using digital text search, and sometimes quite inaccurate if using traditional indexes (see Section 6.2). It was also my conjecture, that the speed and accuracy would influence the amount of information eventually accessed by users making it a crucial element in the evaluation of the system.

### 6.6.1 Study Design

The main aim of the Visual Indexing System was to integrate the speed of digital search with the visual overview of keyword locations that comes with document indexing. I anticipated that this would increase the speed with which users can locate relevant information and improve the effectiveness of reading the text. My hypothesis was that a custom indexing solution would in itself be an improvement over traditional linear indexing, but furthermore, that a visual representation (colour and size) would prove to be a valuable asset in triage efficiency.

To investigate these issues, I performed a comparative user study which focused on qualitative data in the form of subjective user questionnaires, as well as quantitative data obtained from task timings and accuracy levels.

#### Control Systems

In order to test the effectiveness of the implemented visualisations, I implemented two control searching methods as a basis to test against. These two techniques were based on long-established methods of document searching:

**Linear Search:** The Linear Search tool implemented in the program is based on the standard Windows within-document text search feature that allows users to sequentially progress through a document one keyword at a time. This method of search depends highly upon the length of the document and the number of occurrences of a word in a document. I expect therefore, to have a wide range of timing results from this portion of the study tasks.

**Traditional Index:** The Traditional Index tool has been designed to resemble the classic index structure, that is, all entries are the same size and colour. It is not comparable to a traditional back-of-book index however, as it actually still contains the same general features of the other custom indexes within the system. That is, the Traditional Index tool is still a custom index builder that appears in the index visualisation area after the user enters keywords. The index values themselves are still hyper-linked and still contain mouse-over tool-tips. The only difference between the Traditional Indexing tool and the visual tools mentioned in the previous chapter, is that the index entries themselves are at a constant size and are always black (i.e., they do not use size or colour to visualise the relevance of the page).

### Procedure

The participants recruited for the study were selected from a set of Computer Science researchers due to their experienced knowledge of online and document searching. In total 14 participants were selected: 11 male and 3 female, all between the ages of 22 and 37 and all with normal colour perception<sup>2</sup>. All timings and questionnaires were stored anonymously, the study itself lasted on average around 30 minutes and the participants were given a £5 gift voucher in return for their time. No video or audio recording equipment was used within the study, but the timings of every task were recorded automatically by the program.

The structure of the study comprised of:

1. A pre-study questionnaire designed to gain general information about each participant's searching habits.
2. A short training session on each of the systems.
3. The tasks (described below).
4. A post-study questionnaire designed to obtain subjective views with regards to the tested systems.

### The Tasks

The tasks provided were designed to measure how long it takes to find the most relevant parts of a document based on a particular keyword or phrase. The relevance of a page within the document is calculated based upon the number of occurrences of the search term relative to the other pages within the document (see Section 6.5.1).

In total, each participant was given three PDFs and asked to perform five separate searches on each – one for each of the search tools. They were then asked to discover the part of a specific document that best (most) matched a given query. In order to accurately record timings for the tasks completed by the study participants, time data was logged automatically by the program.

Due to the nature of the study, there were several factors that could affect the time taken to complete the tasks. These include the length of the document and the number of occurrences of the keyword or phrase in the document. To minimise the effects of these factors, each search term was rotated to balance orderings and the logged time data across the 3 separate PDFs was averaged. To reduce any bias resulting from the different PDFs and any possible learning effects, the orderings of the documents and tasks were varied from participant to participant.

### Selection of Study Metrics

Before conducting the study, I laid down a set criteria to help assess how each of the study interfaces will be evaluated.

**Speed:** The speed of a look-up was measured in seconds from the moment a participant was given a task, to the moment they produced the results (i.e., the most relevant page/cluster of the document based on the keyword or phrase).

---

<sup>2</sup>Prior to the study, all participants confirmed that they were not colour blind and had no significant visual deficits. This step was taken as a precaution to reduce the possible bias resulting from outside variables relating to visual perception

**Accuracy:** Prior to the study, the optimum page relevance results were calculated and recorded as a basis to test the accuracy of user results. This was measured as a percentage according to how many of the relevant pages were discovered using each method (e.g., if the optimum result was pages 4, 17, 22 and 34 but the user only found 17 and 34, their accuracy result would be 50%).

**Subjective User Ratings:** The subjective ratings were recorded out of 10 (1 being the lowest and 10 being the highest), and tested five different aspects of the systems:

- How easy it was to use;
- How easy it was to *learn* to use;
- The clarity of the display;
- How easy it was to judge the number of times a word occurs in a page or set of pages;
- The speed in which they can find relevant sections of a document.

**User Behaviours:** As well as the criteria mentioned above, I also observed and made notes about the way in which users approached the tasks in order to better understand their searching habits.

As I have discussed previously (see Section 2.6), the studies I have undertaken in previous chapters have not used speed and accuracy as criteria for evaluating the effectiveness of my implemented systems. This was primarily due to my overall goals in designing lightweight active reading systems. That is, I felt that it was unlikely that improving the speed or accuracy of an interface would significantly reduce the cognitive load required for performing the functions.

In this case however, I have selected a topic of active reading which is currently heavyweight in both its physical and digital forms and suffers from problems associated with speed (e.g., text search is linear and slow) and accuracy (e.g., back-of-book indexes may have un-indexed entries). It is for this reason therefore, that I opted to use these criteria in addition to subjective user responses and observed behaviour to give a more concrete evaluation of the Visual Indexing System over more traditional methods of text searching.

### 6.6.2 Hypotheses

**Hypothesis 1:** The Traditional Index builder will be a faster and more accurate method of locating relevant sections in a document than text search;

**Hypothesis 2:** The use of colour and size will increase user satisfaction of the index builder and yield higher subjective scores from study participants;

### 6.6.3 Study Results

The initial questions posed to the participants in the pre-study questionnaire confirmed that the majority of users make use of physical (93% of participants) and digital (79% of participants) reference books at least once a week. To determine the most common methods of locating information in documents, the participants were asked to rate three separate methods according to how often they made use of them:

1. Using the index in a physical book;

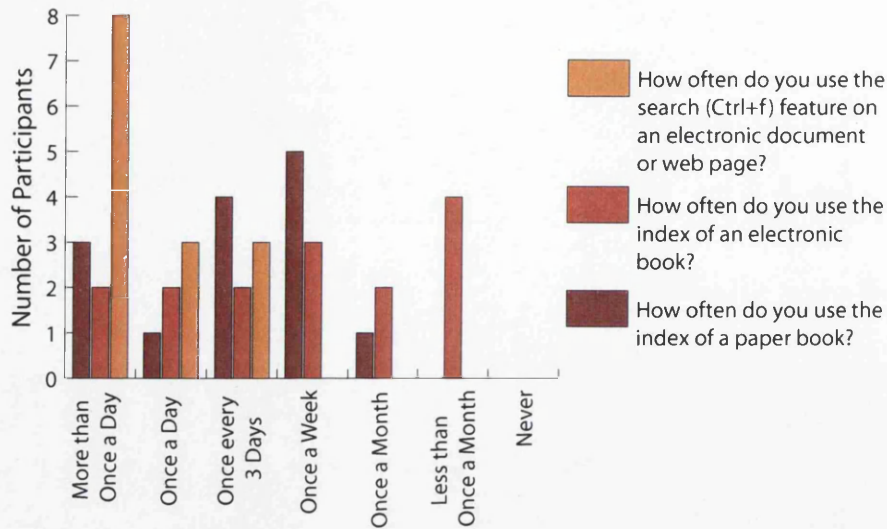


Figure 6.8: Usage statistics for digital search, digital indexes and physical indexes

	Linear Search	Traditional Index	Colour Tag	Tag Cloud	Graph
Average PDF 1	99.9	37.7	21.2	14.4	14.5
Average PDF 2	92.1	33.9	13.2	10.9	9.9
Average PDF 3	134.1	32.8	15.6	11.6	11.2
<b>Average All</b>	108.7	34.8	16.7	12.3	11.9
<b>SD All</b>	59.33	22.78	11.18	9.54	10.20

Table 6.1: Average timed data (seconds)

2. Using the index in a digital book;
3. Using the text search function in a digital document.

They then rated each one on a seven point Likert scale of the format: more than once a day, once a day, once every 3 days, once a week, once a month, less than once a month, never. It is clear from the results (see Figure 6.8) that the most common tool for locating information is text search (Ctrl+f) with all participants making use of it at least *once every 3 days* and 57% using it *more than once day*.

When asked which system they found more useful within *digital* documents (text search or indexes), participants were quite evenly divided with seven opting for text search, six choosing digital indexes and one selecting both. Although many participants commented on the usefulness of the overview that indexes provide, they rated them rather poorly in the pre-study questionnaire giving them averages of 5.43 out of 10 for ease of use and just 4.86 out of 10 for reliability. This evidence strengthens my argument that the traditional paper index, and indeed the imitated digital equivalents, are not the optimal methods of looking up information digitally.

### Speed

The results from the timed tasks (see Table 6.1 and Figure 6.9a) concluded that the Traditional Index is approximately three times faster than the standard Linear Search for locating relevant information in a document. In addition to this, the results also confirm that the use of visual cues offer further improvement to the index builder software with average results for the visual systems (Colour Tag, Tag Cloud and Graph) being at least twice as fast as the Traditional Index.

To assess the statistical significance of the timed test results I conducted a single-factor ANOVA test which makes simultaneous comparisons between the means of all five methods used in the study. This test produced a result of  $p < 0.0001$  ( $df = 4, 205, F = 81.915$  and  $F_{crit} = 2.416$ ), concluding that the resulting data was statistically significant. Due to the large differences in standard deviations between the Linear Search versus the other four systems (i.e., the Traditional Index and the three visual systems), I also performed a Welch's t-test with Bonferroni corrections upon the data to pin-point the areas of major significance.

The results of the Bonferroni t-test (see Table 6.2) confirm that the difference between linear search compared to the four indexing methods is significant, consistently yielding p results of less than 0.0001 and t values ranging from 4.87 to 15.22. Furthermore, comparing the traditional indexing method versus each of the visual indexing methods (Colour Tag, Tag Cloud and Graph) has also shows to have statistical significance yielding p results between 0.0039 and 0.0004 and t values ranging from 3.52 to 4.37.

	Traditional Index		Colour Tag		Tag Cloud		Graph	
	t	p	t	p	t	p	t	p
Linear Search	=4.87	=0.0001	=6.51	<0.00001	=15.18	<0.00001	=15.22	<0.00001
Traditional Index			=3.52	=0.0039	=4.37	=0.0004	=4.26	=0.0004
Colour Tag					=1.42*	=0.081*	=1.27*	=0.11*
Tag Cloud							=0.09*	=0.3*

Table 6.2: Statistical significance of timed data using Welch's t-test with Bonferroni corrections. \* indicates no significant result was found.

### Accuracy

For the purposes of this study, the most relevant section for any search term was defined to be the page/cluster that contain the most occurrences of the user defined keywords. To determine the precision of each document search, each participant was asked to report on their findings as to which page or cluster of pages *they* found to be the most relevant. This information was recorded immediately after each search.

After the study was completed, this data was compared to the 'answers', that is, my definition of the most relevant section, in order to determine if the participant had got it right or wrong. This yielded a set of five tables (one for each of the search methods) which included each of the 14 participants and



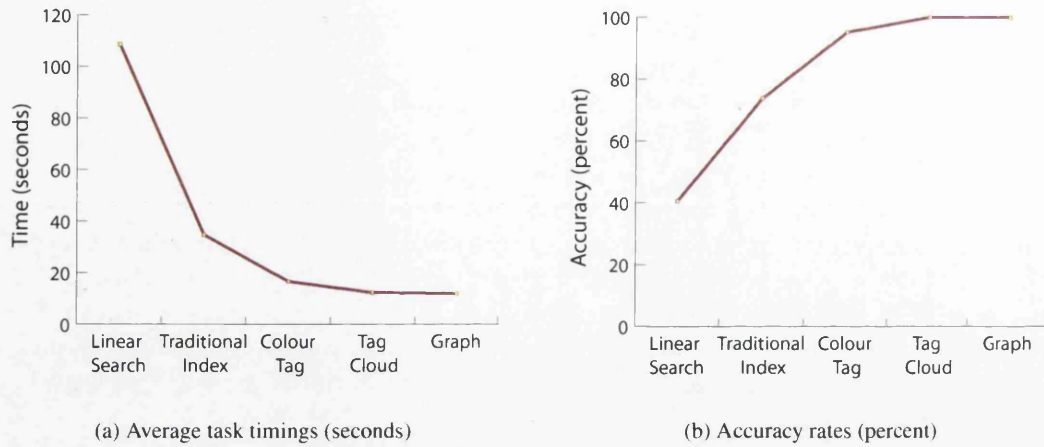


Figure 6.9: Speed and accuracy results




Figure 6.10: Average user ratings for each of the five systems

their results for all three PDFs, that is, 42 boolean results for each search method. The results shown in Figure 6.9b show the percentage of right answers amongst each 42 boolean set.

In the linear tasks, users were able to locate the most relevant section of a document only 40.47% of the time, whereas the traditional index method yielded a result of 73.81%. To assess the statistical significance of these two results, a Chi-squared test was performed yielding the results  $p = 0.002$  ( $\chi^2=9.528$ ) confirming that they are significant. The accuracy of the visual implementations resulted in even greater accuracy, with a result of 95.23% for the colour tag system and 100% for both the tag cloud and graph versions. As you can see from Figure 6.9, as the task timings reduced, the accuracy increased.

	Favourite System	Least Favourite System
Linear Search	0	12
Traditional Index	0	2
Colour Tag	0	0
Tag Cloud	4	0
Graph	10	0

Table 6.3: User System Preferences

### Subjective User Ratings

To gather subjective information about the functionality of the respective systems, the participants were asked to rate each one out of 10 (1 being the lowest 10 being the highest) for the following attributes:

1. How easy it was to use;
2. How easy it was to *learn* to use;
3. The clarity of its display;
4. How easy it was to judge the number of times a word occurs in a page or set of pages;
5. The speed with which they can find relevant sections.

The averaged user ratings for each of the five types of search systems are highlighted in Figure 6.10 and illustrate a clear trend in usability amongst the systems. Specifically, the visual systems (i.e. Colour Tag, Tag Cloud and Graph) score consistently higher ratings for every characteristic than the traditional approaches (linear search and traditional indexing). Furthermore, 93% of the participants concurred that the visual index systems are either a major (12 out of 14) or minor (1 out of 14) improvement over traditional indexing. Comments such as *“The relevance of pages can be shown easily in a graphical manner”* and *“These [the Visual Systems] are much faster and more dynamic - I don’t have to plough through pages of indexes to find the one I want”* further support this improvement.

### User Behaviours

The tasks presented to the participants instructed each user to locate the most relevant section of a particular document based on the keyword or phrase entered. To make an informed decision regarding this section using the linear search tool, users were required to first search through every term occurrence in the document to determine the most relevant section. The observed behaviour of the participants when performing the tasks however, suggests many users have a low inclination to spend a large amount of time searching one text systematically. For example, Participant seven read only the first third of their linear search document during the Linear Search task, a trend that was seen, albeit to a slightly lesser extent by four of the participants during the study.

It has previously been observed [17] that only in extreme circumstances will users be meticulous in their searching of a document while deciding on its relevance. I have identified this behaviour within the study – in the linear searching tasks, around a third of the participants did not span the entire document before making a decision about the most relevant section. In fact, it is this behaviour that caused many of them to select a page or cluster that did not match the highest frequency area of the keywords given.

## 6.7 Conclusions

This chapter has described a tool for document indexing; a common activity that is currently heavy-weight in both its digital and physical forms. For example, looking up entries in the index of a physical book is greatly supported by the use of fingers as temporary placeholders but still takes time to look up and also suffers from the potential problem of missing index entries. Indexes in digital documents are used in conjunction with buttons or scroll-bars, which can be difficult to manipulate quickly and still suffer from the missing index problem. An alternative to indexes in digital documents is text search, which linearly progresses through each instance of a word in the document. This method can also be considered heavyweight as it gives no overview of dense areas of interest and can be tedious and time consuming for words that occur many times. This method does overcome the problem of un-indexed entries however, as the keyword or phrase for the search is user-defined.

To solve these issues, I have designed and implemented a system that combines the visual overview of traditional back-of-book indexes with the speed of text search, while at the same time eliminating the need for flicking back and forth by including a dedicated index side panel. The system, which allows users to build their own custom indexes via their own keyword or phrase, incorporates size and colour to aid in the visualisation of keyword occurrences.

After discovering that the space surrounding the PDF area is useful (see Chapter 4), it seemed sensible to test this theory with something other than making notes. It was this realisation that led me to create the static side panel that displays the user created indexes. The main benefit of this dedicated panel for displaying indexes is that users no longer need to scroll between index and document content as it is always visible on the screen.

To test the implemented system and its three distinct word occurrence visualisations (Colour Tag, Tag cloud and Graph), I performed a comparison study against two control systems (Linear Search and Traditional Index). The results from this study showed significantly reduced timings and higher accuracy rates for the traditional indexing and visual systems compared with the linear search which confirms Hypothesis 1: The traditional index builder will be a faster and more accurate method of locating relevant sections in a document than text search.

The study also showed user preference for the visual features of the systems (i.e., colour and size) yielding higher subjective scores for the three visual implementations thus confirming Hypothesis 2: The use of colour and size will increase user satisfaction with the index builder and yield higher subjective scores from study participants.

### 6.7.1 Lightweight Features

The results from the user study concluded that participants performed tasks significantly faster and more accurately using the systems with the index builder tool (Traditional Index, Colour Tag, Tag Cloud and Graph) as opposed to the Linear Search system. The participants also expressed subjective preference for the index builder systems over the more common text search. These results therefore confirm:

- Giving an overview of where in the document word occurrences exist, significantly increases both speed and accuracy of a search task.

This overview of important information is also benefited by the clustering technique that groups consecutive pages together, as opposed to listing them separately. Participant comments on the topic of clustering included: *“It [page clustering] is nice - it’s not very cluttered this way”* and *“I didn’t even notice that - it just makes sense - like how books do it”*. This method of reducing the amount of

information, on screen aids users in the process of finding relevant clusters of information which is a useful asset when attempting to locate specific topic areas within books. I can therefore conclude from this that:

- Displaying information in a more concise way, in this case, by clustering results, can aid the visualisation of data to users.

Further to this, a comparison between the four index builder implementations (Traditional Index, Colour Tag, Tag Cloud and Graph) found the visual approaches (Colour Tag, Tag Cloud and Graph) to be even faster and more accurate than the Traditional Index system. These visual approaches, which make use of colour (Colour Tag) and colour and size (Tag Cloud and Graph), are also rated favourably by participants in their subjective review of the implementations. Since the only difference between the visual systems and the Traditional Index system was colour and size, I can therefore conclude:

- Colour is a useful way of visualising important information. In this case, the colour was used on a gradient heat scale to indicate to the user areas of the document with high occurrences of the search term;
- Size is also a useful way of visualising important information. In this case, the larger a link, the more word occurrences of the search term exist within the page(s).

One of the main problems associated with digital indexes is within non-hyper-linked documents where there is a need to scroll back and forth between the index and its entries. Unlike physical documents where fingers can be used as temporary placeholders to ease this process, digital indexes are usually used by simply scrolling to the first entry then back to the index then to the next entry then back to the index and so on. Although I did not investigate this issue during the study (as all of the index builder systems were hyper-linked), user comments such as: *"It's great having the index on the side like that, now I can see the index and the page at the same time"* and *"It [the side index] is nice because I don't need to scroll so much"* strongly indicate that:

- Providing static access to important information, in this case the index bar to the side of the document area, reduces the number of steps users take in retrieving the information, in this case, scrolling.

Finally, this chapter was designed to prove that paper is not a panacea; there are many aspects of digital document design that if properly implemented, can greatly exceed the limitations of printed texts. My intention in this case, was to design a system that incorporates useful features from both the physical and electronic domains to create a digital solution that out-performs current paper and digital systems. This was achieved in the Visual Indexing System by integrating the speed of digital search with the overview and access ability of the physical index. Therefore:

- Mimicking paper is not always the optimum solution to the heavyweight properties associated with digital documents. In this case, I have increased both the speed and accuracy of locating relevant information within a digital document by incorporating the uniquely electronic techniques of text search and hyper-links into a physical index design.

### 6.7.2 Future Work

A useful area of future work in this area would be to perform a user study comparing the Visual Indexing System to traditional indexes within physical books. A study such as this would confirm whether or not the implemented system is an improvement over the traditional paper based index method.

Another step which would greatly improve the usability of the interface itself, would be to apply a stemming algorithm to the search feature which will reduce a keyword down to its primary lexical unit (or root) to pick up additional search terms for the index.

### 6.7.3 Related Publications

1. Jennifer Pearson, George Buchanan, Harold Thimbleby : **Creating Visualisations for Digital Document Indexing** - *European Conference for Digital Libraries (ECDL) 2009*, Corfu, Greece.

## Chapter 7

# Discussion and Conclusions

If we knew what it was we were  
doing, it would not be called  
research, would it?

---

Albert Einstein

### 7.1 Overview of Problem

Reading is an age old activity that is increasingly being performed interactively on-screen. The soaring popularity of digital documents in recent years has been accelerated by an influx of electronic reading devices such as eReaders, tablets and mobile phones – hardware that is becoming evermore fluidly integrated into our daily lives [95, 97, 134, 137]. Despite the clear popularity of digital texts however, research indicates that paper usage is in fact increasing [39, 105], suggesting that people are obtaining digital documents, but are later printing them.

This underlying desire to print digital texts [81, 118] demonstrates a clear failure of on-screen reading and reading related activities. It stands to reason therefore, that if the motivation to print digital texts is due to the quality of the tools used for digital active reading, then improving them would decrease the need for printing. How best to improve the way in which users interact with such systems however, is a difficult task which, fundamentally, depends strongly upon the differences between paper documents and their electronic equivalents.

There are many key distinctions between electronic and paper documents. The tangible properties and affordances that paper offer are considerably different to those that can be performed electronically, resulting in a drastic variation in the interaction styles seen between the two media. For example, paper is thin and can therefore be folded and ripped, whereas current electronic screens, despite advances in technology [70] are still too bulky to manipulate in this way.

The way in which physical documents age and mature with help from their owners is another key feature which makes them different to digital texts. Over time, physical books will wear, they will get tatty and dirty, often providing a nostalgic reminder to previous occasions where one's own copy of a favourite book was stained during childhood [43]. As Marshall [84] commented:

*“The tattered cover is, in and of itself, a reminder of the book and all of the experiences surrounding it.”*

– Marshall [84].

The act of annotation is a record of a reader's previous interaction with a book; along with general wear and tear, physical documents 'tell a story' that is personal to the owners that have come into contact with it. It has also been observed that in some cases, the annotations of others provide value to future readers [81]. During a 1987 keynote presentation [136], Brown University Professor Andries van Dam stated:

*"I would always grab the dirtiest copy of a book in the library, rather than the cleanest one, because the dirtiest ones had the most marginalia, which I found very helpful."*

– van Dam [136].

The action of marking-up documents in this way can be considered not as simply reading, but as part of a more focused and cognitively demanding process known as *active reading* [3]. Active reading typically demands more attention than what is required to simply read a novel or newspaper, and frequently includes tasks such as highlighting, note-taking, searching and non-sequential navigation. It has been widely documented [85, 107, 117, 131] however, that paper, thanks to its familiar and lightweight [85] affordances, better supports active reading than on-screen implementations.

There have been many attempts to improve the usability of digital document tools by mimicking the interaction styles of paper. Liesaputra et al. [29, 76, 77] and Card et al. [24, 25] for example, have used the literal book metaphor [71] to create realistic visualisations of real books within their designs. Golovchinsky et al. [114, 115], Hinckley et al. [27, 59] and Tashman et al. [130, 132] have gone one step further, by enhancing augmented reading hardware to create 'book-like' reading appliances. There have also been attempts to bridge the gap between the physical and digital domains by amalgamating paper documents with digital enhancements (e.g., [10, 72, 144]).

Evidently then, the physical properties of paper are integral [79, 117] to the ease in which they are manipulated, and undoubtedly contribute to the high proportion of users opting to print digital documents. A useful way of describing the physical affordances of paper, is by what Marshall et al. [85] describe as lightweight interaction; that is:

*"...navigation that occurs either when people reach a particular page or when they move within an article in a way that is so unselfconscious that they aren't apt to remember it later."*

– Marshall and Bly [85].

The almost subliminal interaction of acts such as folding, highlighting and flicking, require little cognitive attention from the user – a desirable state that is neatly defined by Csikszentmihalyi's theory of flow [34]. The state of flow can be characterised as being so fully immersed in a task that all other considerations are forgotten. It is within this state of complete, focused motivation that users will achieve the maximum level of cognitive attention with the task they are performing.

To achieve the desired state of flow then, it is essential that the tools designed to aid the active reading process are made as intuitive as possible to ensure that the user is thinking about the task at hand, as opposed to the tools used to facilitate it. That is, the tools themselves must be ready-to-hand [56, 148] rather than present-at-hand. If the lightweight properties of paper hold (i.e., if the tools used to aid the active reading process are ready-to-hand), then a state of flow can be achieved with the primary active reading task.

On paper this ideal is often realised. For example, folding down the corner of a document to mark a page, or highlighting a word within a passage of text can be so unselfconscious that the user

is unlikely to think about doing it. It has been widely recognised [2, 85, 118] however, that digital active reading tools are less intuitive and consequently require more conscious effort to perform, often breaking users out of their state of flow.

## 7.2 Proposed Solution

The main motivation of this Thesis is to improve the digital active reading process by incorporating lightweight techniques into the design of the tools aimed to facilitate it. It is my conjecture, that doing so will make the tools themselves more intuitive, require less cognitive attention to perform and consequently help promote a state of flow with the primary active reading task, a situation that is often accomplished when using paper based tools.

Marshall and Bly's definition of the term lightweight [85] describes the properties of paper documents as opposed to eDocuments. They speculate however, that this concept of lightweight interaction can be also be applied to digital technology but do not give any concrete evidence to support it.

To gain a better understanding of what attributes can be considered digitally lightweight, I have implemented four systems, each of which have been carefully designed to investigate specific areas of potentially lightweight design. In some cases I have closely mirrored the interaction styles of physical documents in an attempt to create a digitally lightweight equivalent, and in others I have enhanced digital techniques to better utilise the availability of electronic functionality. In all cases, I have thoroughly tested each separate implementation by means of systematic evaluation on an appropriate set of users, helping me identify exactly which areas of my designs can fall under the lightweight category. To fully comprehend which properties of active reading tools can be deemed lightweight, these evaluations were based predominantly on user effort during the study tasks, coupled with subjective preference for my tools versus more traditional alternatives.

## 7.3 Contributions

The following section gives an overview of the main contributions I have made throughout this Thesis. I start by briefly describing each chapter's content, and move on to their individual contributions in the context of lightweight design. In the next Section, I bring together each of the running themes identified throughout my investigations, to produce a list of lightweight attributes for use within future active reading systems.

The definition of lightweight is highly dependent upon the cognitive workload of the individual performing the task, and not upon the time they take to achieve it. That is, the time taken to complete a secondary task is irrelevant if users are indeed performing it without thinking, yet *any* time taken performing an action that has stopped a state of flow with their primary task is problematic. Therefore, simply timing the effects of a change in interface design will not usually provide a meaningful representation of whether or not a property is indeed lightweight. It is for this reason, that the studies I undertook on each of the systems described below were evaluated predominantly using subjective user data to determine the lightweight-ness of properties within this Thesis.

### 7.3.1 Placeholders

In Chapter 3 I investigated the concept of placeholders; that is, a method of marking one's place within the pages of a document. In paper documents, placeholders are versatile and lightweight, for example, inserting a bookmark or piece of scrap paper between the pages, or 'dog-earring' the corners. Current



digital placeholders however, are significantly less intuitive and are also rated unfavourably by users. Placeholder within documents then, is a perfect example of a function that is currently lightweight on paper, but heavyweight digitally.

In an attempt to overcome the problems associated with digital placeholdering, I implemented the Visual Bookmarking System – a design that uses coloured ‘tabs’ to mimic the way placeholders are used in physical books. The comparative user study I conducted on the implementation indicated preference for the visual design over more traditional approaches such as drop down menus and tree lists.

### Lightweight Properties

The results from the study I conducted on the Visual Bookmarking System unearthed several useful possibilities for future exploration. The main points that resulted from my investigation into placeholders were:

- Hiding contextual information such as bookmarks within menus is a poor design choice and can be considered ‘heavyweight’ if the function is frequently accessed;
- The area surrounding the PDF is potentially useful for displaying information. In this case, the static area surrounding the document was used to display the bookmark tabs that are visible at all times independently of the currently open page;
- Mouse-over pop-ups that allow important information to be viewed easily by simply hovering the cursor over certain areas of the screen provide lightweight information at a glance;
- A visual approach, in this case using the *position* of the tabs to indicate where in the document they are positioned is a good way of illustrating important data. This also gives a neat overview of where each bookmark exists in relation to others as well as in relation to the currently open page;
- The use of colour may be a lightweight property if used effectively.

These findings provided me with a strong grounding for further investigation by contributing useful ideas that I subsequently incorporated into later implementations.

### 7.3.2 Annotations

In Chapter 4 I conducted several formulative investigations into annotation, on both physical and digital documents. Annotation is a major part of the active reading process, and a function that has received a considerable amount of attention from the digital library community. To gain a better understanding of annotation for future implementations, I first conducted a small study on paper documents to probe users’ use of annotation, specifically, where they choose to position them.

The results from this paper-based study indicated that the margins of a document form an important part of the physical annotation process, owing to the additional space they provide for notes and their proximity to the original text for easy referencing. Yet, digitally, users are very restricted in where they can physically place notes.

Following on from the results of the paper study, and to investigate how users perform annotation on digital documents, I implemented the Margin Annotation System – a digital annotation tool that includes an expandable margin area to make notes.

### Lightweight Properties

The studies I conducted on annotation were designed to both further investigate the potentially lightweight attributes I identified within the placeholder study, as well as probing additional areas of lightweight design. The main points I identified throughout my investigation into annotation were:

- Document *margins* are a valuable and popular asset in the annotation process. The results of the digital study strongly suggests that the margins are a useful addition to marking up the PDF as opposed to a straightforward alternative;
- Mouse-over pop-ups are a lightweight way of displaying information that is not constantly needed. In the case of the Margin Annotation System, it is used as an alternative to the separate medium approach seen in physical annotation practices;
- The concept of appropriation may be a lightweight feature that can be transferred to the digital level;
- Saving any user changes to a document separately ensures the original does not get ‘damaged’ while allowing multiple annotation files to be used and shared easily while only physically storing one copy of the original document.

This list confirms two of the lightweight properties identified within the placeholder study (i.e., the area surrounding the document and mouse-over pop-ups), and suggests two additional areas that could prove potentially beneficial (i.e., appropriation and independently saved file copies). I subsequently used these properties as a grounding for my next investigation into digital note-taking.

#### 7.3.3 Note-Taking

After learning several important lessons from the placeholder and annotation prototypes, I moved to the broader concept of note-taking. In this chapter, I identified three key deficits associated with note-taking on digital documents:

1. **Space:** On paper, notes can be made and placed in areas surrounding documents (e.g., desk or walls) to allow static access to information whilst reading. Digital documents however, rarely offer a comparable workspace to position notes while reading a document;
2. **Tool Overload:** On paper, one tool can often perform more than one function; for example, a Post-it note can be used to make notes as well as being used as a bookmark. In contrast, digital documents typically separate these functions, resulting in multiple single-purpose tools;
3. **Menu Navigation:** On paper, when an action needs to be performed, the user does not need to go looking through menu systems in order to find a way of doing so. For example, when a Post-it note is no longer required on a paper document, one can simply pick it up and throw it away. Digitally however, this function is usually achieved via a menu.

In order to overcome these problems and create a more lightweight digital document implementation, I designed the Digital Reading Desk – a system that incorporates the physical interactions of paper documents in its design. The interface of the Digital Reading Desk system contains an extended workspace comparable to a physical desk that allows users to make static notes that remain constant regardless of the currently open page. It also includes draggable Post-its that can act as both placeholders as well as notes.

### Lightweight Properties

The study I conducted on the Digital Reading Desk was carefully designed to evaluate three distinct areas of the interface: the static desk area, the unified Post-it tool and the draggable create and delete method. This extensive user study allowed me to identify several key properties:

- The *static* workspace surrounding the document is a useful area to display constant information. For example, notes and summaries about the book as a whole, or, in the case of the Visual Bookmarking System, placeholder tabs;
- Designing for appropriation, if conducted in a way that is familiar to the user (e.g., functions that are currently appropriated on paper) can be considered to be a lightweight property;
- Removing the need for menus wherever possible promotes direct manipulation, and is therefore less likely to distract from the primary active reading task or interrupt any state of flow that may have already been achieved. ;
- The visual Post-it style bookmark tabs are a useful way of giving an overview of where each bookmark exists in relation to the currently open page.

This list confirms several of the lightweight properties I identified in my previous investigations. Firstly, the concept of space surrounding the document, in this case *static* space as opposed to *marginal* space, was again verified to be a lightweight attribute. In addition, the concept of appropriation which was briefly touched upon in the study on annotations, was more firmly established to be a useful lightweight property. I also discovered that a reduced menu system could aid in the possibility of achieving a state of flow with the primary active reading task, and reconfirmed the use of visual cues for displaying information.

### 7.3.4 Visual Indexing

In Chapter 6, I implemented a system that creates custom visual indexes designed to aid the within-document information retrieval process. Indexes are an example of a feature which is relatively heavy-weight in both its physical and digital forms. For example, on paper, indexes are fixed, and cannot be customised to tailor to the specific needs of the user. It also takes time to physically navigate to pages within an index of a physical book. In contrast, digital indexes can be hyper-linked which improves the speed of look-up, but suffer from other problems such as the difficulties in scrolling between pages and index.

My intention in this chapter was to prove that paper is not the solution to every problem associated with digital documents. By combining techniques from the currently heavyweight physical and digital indexing methods, I aimed to arrive at a superior system that can then be considered lightweight. Specifically, I took the most useful aspects of each method, in this case, the visual overview seen in physical back-of-book indexes and the speed of digital text search, to create the Visual Indexing System.

### Lightweight Properties

The results of the study I conducted on the Visual Indexing System allowed me to identify a range of lightweight attributes, some of which coincide to results from previous studies, whereas others are new properties that have resulted from my investigations into solely digital interactions:

- Giving an overview of where in the document word occurrences exist, significantly increases both speed and accuracy of a search task;
- Displaying information in a more concise way, in this case by clustering results, can aid the visualisation of data to users;
- Colour is a useful way of visualising important information. In this case, the colour was used on a gradient heat scale to indicate to the user areas of the document with high occurrences of the search term;
- Size is also a useful way of visualising important information. In this case, the larger a link, the more word occurrences of the search term exist within the page(s);
- Providing static access to important information, in this case the index bar to the side of the document area, reduces the number of steps users take in retrieving the information, in this case, scrolling;
- Mimicking paper is not always the optimum solution to the heavyweight properties associated with digital documents. In this case, I have increased both the speed and accuracy of locating relevant information within a digital document by incorporating the uniquely electronic techniques of text search and hyper-links into a physical index design.

The predominant theme I have identified through this study is the visual language used to display useful information; in this case, colour, size and overview were used to better illustrate important contextual data. The work I conducted in this area also confirmed the value of solely digital techniques such as hyper-links, as well as reconfirming the lightweight-ness of the area surrounding the document.

## **7.4 Research Scope and Limitations**

### **7.4.1 Cognitive Load**

The existing work that motivated this research posited the hypothesis that cognitive load and the user's conflict between their primary reading task and the supporting work of annotation, and so on, was a problematic area of interaction design. Within my studies, a major difficulty was capturing the cognitive factors involved in the users' interactions. I primarily used subjective ratings and responses to elicit the participants' experience of the different interaction designs presented to them in the experiments. Previous research in the domain of digital libraries has used similar approaches (e.g., [86,121]), but a researcher with greater expertise in cognitive psychology could, no doubt, design a more detailed research method that would give specific and direct cognitive insights.

### **7.4.2 Study Duration**

My use of laboratory experiments allowed for close observation of participants' interaction, but cannot demonstrate the impact of long-term use of my designs. A longitudinal study would permit the evaluation of the long-term use of the systems, and indicate if the impact of the designs would persist beyond the initial period within which the interface is novel to the user. Furthermore, it may be that long-term use highlights further advantages or disadvantages of both my designs and the existing interactions that are in current use.

### 7.4.3 Number of Participants

Each of my experiments consisted of 10 to 16 participants. While this is consistent with the previous research in digital library interaction [121], other areas of research in HCI regularly use a larger cohort of participants. Given the power of the differences observed in most of my studies, it would be surprising if the outcomes would change significantly. It nonetheless should be considered that a larger number of participants may adjust the fine-detail of the results, and would also enhance the authority of the findings of my investigations.

### 7.4.4 Scale of Annotation

Given the short-term nature of the experiments I conducted, and the smaller number of participants involved, the scale of annotation undertaken across the studies is relatively limited. Readers who regularly engage in attentive reading on paper produce rich annotations on specific texts over many months [116, 135]. Within both digital library and information retrieval research, we know that the Zipfian distribution [11], also known as Bradford's law, is a fundamental pattern within information work [128]. This suggests that most documents used by an individual will have only limited annotation, and, conversely, only a few will be subject to intensive and prolonged mark-up. My findings will therefore likely hold for the majority of individual documents, but there are two points to be borne in mind: first, the cumulative complexity of many documents, each with a limited annotation, may produce different behaviours than a small corpus; second, the challenges for intense and detailed annotation may give rise to different phenomena in user behaviour and unveil different user requirements.

## 7.5 Running Themes

This section brings together the topics that have been prominent throughout much of my research. I use these long-running themes as a foundation to create a list of digitally lightweight properties for use within electronic active reading systems.

### 7.5.1 Space Beyond the Document

One of the main themes running throughout the systems I have implemented, is the use of additional space beyond the document. Superficially, this can take two forms: the marginal space surrounding document pages, and the static space surrounding the entire document (e.g., the desk or wall). These two approaches either give users extra page space for annotation/notes (the margin), or gives support for the user's management of the document as a whole (static desk or wall space).

This particular attribute has been seen in:

1. The Placeholders chapter where *static* space was used to display constant bookmark tabs;
2. The Annotations chapter where *marginal* space was used to make further notes and annotations on pages of a document;
3. The Note-Taking chapter where *static* space was utilised, this time, in the form of a 'desk' that allows notes to be made independently of the currently open page;
4. The Visual Indexing chapter where again, *static* space was used to display the index bar to reduce the need for scrolling between index and document content;

The evidence I have gathered from each of the individual user studies and my analytical review of current reading technology, has enabled me to confirm that the **space beyond the document** is a 'lightweight' property that should be used in future digital document designs.

### 7.5.2 Visual Language

The way in which information is displayed to a user is an important factor which can ultimately determine how effectively a system can be used. For example, if an important piece of information is hard to see or find, the user's attention will be shifted from their primary task to how best to access the information, breaking them out of their contextual flow.

There is a significant body of research that investigates the use of visual language such as colour, space and position to organise information (e.g., [87, 122]). This idea of using visual language to display relevant information is also one of the recurring themes in this Thesis. Specifically, these visual attributes have been identified in:

1. The Placeholders chapter where *position* was used to indicate the location of bookmarks in relation to the current page as well as to other bookmarks, providing an *overview* of where bookmarks exist in the document. In addition, the use of *colour* was also identified as a potentially useful attribute for distinguishing between bookmarks (e.g., red for important etc.);
2. The Note-Taking chapter where again, *position* was used to give an overview of where Post-it bookmarks are located in relation to the currently open page, which also gives a certain *overview* of the bookmarks within the document. The Post-its within this system also allowed their *size* and *colour* to be changed, which aids in the indication of page importance;
3. The Visual Indexing chapter where *colour* and *size* were used to illustrate word frequency within pages. The traditional index structure which includes clustered entries, gives an added *overview* of exactly where in the document important areas occur.

I can thus conclude from the evidence gathered from these investigations, that **the visual language** is an extremely important feature in displaying information digitally. Specifically, some of the lightweight features of the visual language I have identified are: **position, colour, size** and **overview**. Some of these attributes have been seen before in other work – Malone [80] for example, makes use of size, location (position) and colour to indicate priority in a group of documents.

### 7.5.3 Immediacy of Access

As I discuss in Section 2.4.3, the desirable state of flow [34] is achieved when an individual is so engrossed in a particular task, that nothing else around them seems to matter. To achieve a state of flow with the primary active reading task, it is essential that all the tools that are designed to aid it are made ready-to-hand [56, 148] which ensures that the maximum amount of cognitive attention is left for the main task, as opposed to the tools designed to facilitate it. This will then reduce the possibility that the user will be distracted by the tools themselves [109].

With this in mind, it seems reasonable to consider the immediacy of access of the tools within a digital system. That is, how soon can the tools be accessed and, more importantly, how cognitively demanding is it to use the tools designed to aid the active reading process.

One attribute that falls under this category and has been identified within two of my implementations is the idea of a reduced menu system. In theory, a system that has fewer menus will promote

direct manipulation and therefore allow functions to be performed faster and with less cognitive attention required from the user. The concept of a reduced menu system has been identified in:

1. The Placeholders chapter where I identified that hiding contextual information such as bookmarks within menus is a poor design choice and can be considered 'heavyweight' if the function is frequently accessed;
2. The Note-Taking chapter where a drag-and-drop interface for the creation and deletion of Post-its was implemented to promote direct manipulation of the tools.

The results from both these studies has given strong evidence to support the **reduced menu system** as a lightweight attribute of digital document design.

In addition to this, another useful function that has been identified as a lightweight way of displaying information that need not always be visible is the mouse-over pop-up, used in:

1. The Placeholders chapter when a mouse is run over a bookmark tab, a pop-up will appear giving the user additional information about the details of the bookmark (e.g., the page number, user defined descriptions etc);
2. The Annotations chapter where mouse-over pop-ups were used as an alternative to a separate medium seen in physical documents.

This function has proved useful for participants in viewing information that is not required to be visible at all times, but still needs to be regularly accessed. Allowing data such as this to be viewed easily by means of simple mouse movements as opposed to making button clicks or searching within menus, reduces the number of steps needed to access the information, and consequently the cognitive demand on the user. I can therefore confirm that **mouse-over pop-ups** are indeed a lightweight property of digital document design if used to visualise important but not constantly required information.

#### 7.5.4 Digital Technologies

There are many properties of digital documents that far exceed the limitations of the printed page. For example, digital documents can be searched, zoomed and edited easily, providing a new layer of interactive possibilities that are difficult to produce on paper. By optimising these advantages, it is possible to create solely digitally lightweight interactions without the need for mimicking paper. Some of these properties have been identified in:

1. The Annotations chapter<sup>1</sup> where saved annotations are made separately to the main document which ensures the original does not get 'damaged', while allowing multiple annotation files to be used and shared easily while only physically storing one copy of the original document. This idea of saving annotations separately to the main document is a concept that can only be achieved on paper via additional media loosely attached to the original content – a type of mark-up that proved unpopular amongst the participants in my paper based user study;

---

<sup>1</sup>This concept of saving mark-ups separately can actually been seen in all my implementations, but it is the Margin Annotation System where I make use of this property as a feature of the interface, i.e., users can toggle between different annotation files from a menu within the system.

2. The Visual Indexing chapter where I confirm that mimicking paper is not always the optimum solution to the heavyweight properties associated with digital documents. In this case, I increased both the speed and accuracy of locating relevant information within a digital document by incorporating the uniquely electronic techniques of text search and hyper-links into a physical index design.

I can conclude from this evidence that the use of digital technologies can be considered lightweight properties if used effectively. In these cases I have proved that the use of **independently saved file copies**, **text search** and **hyper-links** are all lightweight properties of digital active reading systems.

### 7.5.5 Designing for Appropriation

Appropriation is a method of interacting with an object or technology in a way that was not originally intended by the designer. Although it is near impossible to predict exactly how users may adapt a particular piece of technology, it is possible to design tools that support more than one function, mimicking the behaviours of paper. I have identified the concept of appropriation as a lightweight attribute in two of my investigations:

1. The Annotations chapter where I observed users in the paper study making use of Post-it notes as bookmarks;
2. The Note-Taking chapter where I created a digital version of Post-it notes that can be used as placeholders or as a method of annotation.

The results from the user evaluations I conducted on these systems have provided me with significant evidence that designing for **appropriation** is a lightweight property of digital active reading systems.

### 7.5.6 Completeness of Metaphor

Another key point which has been evident throughout my investigations is the completeness of metaphor. It is clear that the completeness of a metaphor is extremely important to the eventual success of any function that intends to support active reading; placing restrictions on what something means or where it can go it is a demonstration of bad interaction design. Specifically, the importance of the complete metaphor has been identified in:

1. The Note-Taking chapter where users were not restricted as to where they could position Post-it notes. Notes could either be placed: on the document itself, on the desk surrounding the document, or protruding the document which causes them to act as bookmarks;

I can therefore conclude that **completeness** is a lightweight property of digital active reading tools.

## 7.6 The Book Metaphor

There has been a significant amount of research that focuses on the visual book metaphor in an attempt to improve the usability of digital document tools. Although it is clear that mimicking every aspect of physical documents in digital designs is unfavourable, and in many cases, infeasible [118], it is also apparent that visualising documents using the physical book metaphor is a common design choice for consumers (e.g., in applications such as iBooks for the iPad) and researchers (e.g., [30, 115]) alike.



The reason why imitating physical books within digital designs proves so popular, is primarily due to familiarity. That is, the majority of people have been brought up reading physical books and will consequently find a digital text integrated with physical features profoundly recognisable. It is very possible therefore, that a new generation of children being brought up reading *only* from electronic texts, on eReaders for example, will feel very differently about the book metaphor being applied to digital document designs. As Nielsen [99] once commented:

*“...it is unnecessary to keep the book metaphor except for walk-up-and-use situations where immediate transfer of past skills is needed.”*

– Nielsen [99].

Features that are based on the book metaphor are not necessarily perfect in design and can easily be overtaken in functionality by applying solely digital techniques. However, the familiarity associated with such a well-used medium is overwhelmingly strong. For example, in the absence of any better proposals, the first motor cars were designed to look like carriages, which were, at the time, a very familiar sight. Over time however, these carriage shaped vehicles were eventually overtaken and cars now bear little resemblance to the once primitive-looking early automobiles. Yet, there are few complaints from consumers that modern motor vehicles should look more like carriages. Could it be possible then, that the integration of the book metaphor into digital document designs is merely a transition stage before arriving at a more intuitive, yet exclusively digital reading interaction?

This Thesis has focused strongly on the transfer of metaphors from the physical to electronic worlds. By mimicking some of the interactional properties of paper, I have, in the majority of cases, increased user satisfaction with the tools available for active reading. Although this suggests that users are still not ready to completely abandon the look and feel of the physical book, I have also acquired evidence for the contrary by proving in Chapter 6, that the physical book idiom can also be improved upon by using solely digital enhancements.

## 7.7 Electronic Documents

One of the major concerns users seem to have with the eDocument paradigm, is the fact that they do not actually own a physical copy of the book. Some of the studies I have conducted on electronic reading concluded that many users feel almost cheated by buying an electronic document, with comments like *“I’d rather have the [physical] book because it looks good on my shelf”*. This electronic ownership of a document also causes other problems, for example, attempting to sell a copy of a text book you no longer need or even lending a book to friend to read<sup>2</sup> for a short while can cause licensing and key problems for users.

Another problem I have identified from the technology I have reviewed, is that it is clear that the concept of ‘pages’ are not as concrete digitally as they are on paper. Although the familiarity of paper books causes ‘page-like’ breaks in digital media, possibly due to the fact they can be printed, in actuality, page breaks within digital documents can occur anywhere and alter frequently depending upon the zoom level of the document. This reformatting of documents, for magnification or otherwise, consequently causes not only the number of pages within the document to suddenly change, but also the page numbers of different sections of the document. For example, if Chapter Three of an eBook originally appears on page 40 of 100, changing the magnification to make the text twice as large will

<sup>2</sup>Some companies do cater for the transfer of electronic documents, the Barnes and Noble nook for example, has a virtual library system as well as a friend lending function.

cause the document to increase in page length to around 200 pages. What was page 40, now no longer contains the start of Chapter Three but perhaps the mid section of Chapter One; Chapter Three may now start on page 80. This type of inconsistency then has a knock-on effect on other functions within these devices, for example, in some devices, bookmarks only record page numbers, not where specific page numbers start. Therefore in the example above, bookmarking the start of Chapter Three in the original magnification will bookmark page 40; changing to the new magnification will not alter the bookmark to where Chapter Three now starts on page 80, but instead continues to mark page 40 - what is now in the middle of Chapter One.

Clearly then, using 'page numbers' in this context is not a consistent method of referencing within digital documents. As I have already discussed, this is not ammunition to completely abolish the use of pages from digital documents - far from it in fact. The notion of pages is still vital for the use of eDocuments as they are needed for functions such as printing as well as the potential need to reference physical pages to digital media. There is the need however, to more concretely design for specific locations within digital documents to ensure consistency when the document is reformatted for magnification and so on.

Other problems, such as the loss of paper-like tangibility within digital documents have also been identified. Marshall and Bly [85] for instance, comment on the the failure of digital systems in the context of flipping multiple pages within a document and the consequence of losing the serendipitous feature of simply opening a document at an interesting article. They also comment on the implicit metadata found within physical documents that often gives tactile feedback on information such as how much of the book is left to read.

Despite their current pitfalls, digital documents do have many useful qualities that, if correctly implemented, have the potential to overtake their physical equivalents. Digital enhancements such as text search, zooming and copying are all powerful features that significantly increase the usability of electronic text.

It is clear from several sources (e.g., [63]), that some users feel that marks within physical documents provide an further contextual layer of information in addition to the original content. For example, borrowing an old book from the library and finding it has already been marked up with highlights and notes from a previous borrower provides another persons's interpretation of the text and may even help identify worthwhile sections without the need to reading the entire book. In contrast however, I have encountered evidence from some study participants, that there are users who feel it is not at all appropriate to mark a physical book. In fact, some think it is a terrible, almost sinful action to take, even going as far as describing writing in books as "*a crime*" - suggesting that doing so will "*damage*" it [41]. Thinking back to my school days where a teacher once referred to my childish doodles as "*ruining a perfectly good book*", I realise that my mindset too prevents me from even 'dog-earring' the corners of physical books for fear of ruining them.

Digital documents however, need not suffer from the problems associated with additional marking. In fact, not only can they avoid these problems, they can also benefit from other advantages as well as potentially advancing in functionality by providing additional features. Although currently, digital annotation is deemed awkward [107] compared to its physical rivals, my work in the area of lightweight interaction aims to bridge the gap between the physical and digital domains and hopefully make digital reading as easy to perform as it is on paper.

Other useful advantages of digital documents that I have identified throughout my investigations, are the enhancements that can be made using the quick and accurate computational power of the eReading hardware. Chapter 6 was designed to prove that it is not only functionality that has migrated from the physical world that can be considered digitally lightweight. The results of this study proved that the common text search function, coupled with hyper-links and other useful techniques, can be

exploited to produce a digitally lightweight function that would not have been possible on paper.

## 7.8 Concluding Remarks

This Thesis has attempted to establish what properties of digital reader design can be considered digitally lightweight by creating and evaluating implementations to aid in the active reading process. There are many ways in which current digital reading implementations are deficient and similarly, many ways in which they can be improved. By incorporating lightweight techniques into the design of electronic reading software, the cognitive load required by users to perform document related activities can be reduced, helping users achieve a state of flow with the primary active reading task. As O'Hara and Sellen once commented, there is:

*"...the need to support quicker, more effortless navigation techniques."*

– O'Hara and Sellen [107].

This statement encapsulates the goal of lightweight design within digital active reading systems. The main aim of interfaces designed to aid this process is to provide tacit and effortless interaction that provides minimal intrusion on the main reading task.

As a general principle in human-computer-interaction, it is important to improve the user's effectiveness in their tasks. As previous research such as O'Hara and Sellen, and Marshall and Bly, has demonstrated, in the case of active reading, the main (and most important) task is engagement with the text. Increasing the amount of cognitive effort available for this task is therefore desirable. In this Thesis, I have sought to provide digital tools that require almost subliminal attention to control. Through this, more of the user's mental resources will be available for the primary active reading task. With this in mind, the main motivations of this Thesis, as I first described in Section 2, were:

1. Prove by example that lightweight interaction is possible in digital documents;
2. Create systems that use lightweight interaction in their design to improve the active reading process;
3. Produce a list of attributes that can be considered digitally lightweight to aid in future designs of active reading software.

To accomplish these goals, I first identified several active reading tools that can currently be considered 'heavyweight' in their digital forms. Specifically, these topics include: placeholder, annotation, note-taking and indexing. By paying close attention to how these activities are performed on paper, I was able to identify what properties can be considered physically lightweight, and consequently apply them to digital interactions. From this core idea, I then designed, implemented and evaluated several digital interactions that incorporate some of the lightweight properties of paper while at the same time adding electronic enhancements.

The results from these evaluations conclude that users prefer the new simplified systems over more commonplace equivalents (control systems). Not only did the participants generally rate my implementations higher than the baseline interfaces, but they also voluntarily offered corroborating judgements. For example, one participant in the study on note-taking commented that: *"I don't have to think too much about how to do it"* - a perfect example of what Marshall and Bly describe as lightweight navigation and proof that it is indeed possible digitally. This evidence also concludes that

the active reading process has been improved by the systems that include lightweight properties, a valuable contribution to the digital reading process – an activity that is becoming increasingly more commonplace [95, 97, 134, 137].

Finally, from extensively analysing both user behaviour and subjective responses from each of the user studies, I was able to compile a list of lightweight properties of digital active reading systems:

- Space Beyond the Document;
- Visual Language;
  - Position;
  - Colour;
  - Size;
  - Overview;
- Immediacy of Access;
  - Reduced Menus;
  - Mouse-over Pop-ups;
- Digital Technologies;
  - Independently Saved File Copies;
  - Text Search;
  - Hyper-links;
- Appropriation;
- Completeness of Metaphor.

This list of attributes is a valuable foundation for future active reading designs. To illustrate this point, I have used this list to design and implement an additional interface that makes use of the lightweight properties above. It is my conjecture that properly incorporating some of these lightweight properties into new digital reading implementations will reduce the cognitive attention required by users to make use of the tools, and leave more time and effort available for the main active reading task.

Current digital active reading software is far from being a replacement for paper, but my research in this area has moved towards bridging the gap between the physical and digital domains making interacting with digital documents significantly less cumbersome.



## Chapter 8

# Future Work

Don't cry because it's over, smile  
because it happened

---

Dr. Seuss

### 8.1 Introduction

This chapter describes some areas where the work presented in this Thesis can be extended to aid in the creation of future lightweight implementations. Chapters 3 to 6 each contain a short explanation about how their implementation or study can be extended. Now that I have laid down a set of lightweight guidelines, it is feasible to build upon them to improve the usability of digital active reading systems (Section 8.2 describes one such implementation I have produced by following the lightweight guidelines discovered in previous chapters), and also potentially adds to the list of digitally lightweight properties.

There are many areas in which the digital active reading process could be improved by incorporating lightweight techniques. In this Thesis I have investigated four (placeholdering, annotation, note-taking and indexing) different aspects of digital reading which I felt were lacking in certain areas, and improved upon them which led me to produce the list of lightweight attributes seen in Chapter 7. Some examples of other topics of investigation which I did not focus on in this Thesis, but are still areas which could benefit from the incorporation of lightweight interaction include:

- The action of 'flicking' between pages in a document is, on paper, a relatively lightweight activity; using temporary placeholders such as fingers to mark pages facilitates quick navigation between multiple locations within a document. Digitally however, it can be tricky to quickly flick between two or more non-sequential pages using existing methods such as scroll-bars or buttons. Some work has already been conducted to aid this process (e.g., [16, 18, 132]), but further investigation into more lightweight methods of flicking between non-sequential pages within a digital document would be beneficial.
- The investigation of lightweight active reading on portable devices such as eReaders or tablets would be an extremely beneficial area of future work. Clearly, the limitations of hardware such as these will impede the conversion of some of the lightweight properties I have discovered. For example, limited screen real-estate could hinder providing extra space surrounding the document and eInk devices are predominantly greyscale which leaves adding colour impossible.

Therefore, creating active reading software specifically tailored to such devices could provide additional lightweight properties specific to portable devices.

- Hybrid approaches to active reading that involve digital documents with more physical interactions could also be a useful area of future research (see Section 2.5.5 for a literature review on hybrid reading systems). For example, many people prefer to read on paper as opposed to reading on screen, therefore a potentially useful solution in this situation would be to provide a physical book incorporated with a digital note-taking facility. One step I have already taken to achieve this goal, is a system I implemented that makes use of a physical document augmented with RFID tags to link directly to a computer. Using the unique tag as an identifier, the computer system looks up the details of the book and provides the user with an electronic copy of the document, coupled with an area surrounding it to make notes [19].

In addition to these areas, the list of lightweight properties I have laid down need not be used merely for documents; it is also possible for these attributes to be applied to a wider range of electronic content, including:

- Use within web applications; there are several areas of browser design that could benefit significantly from making use of lightweight techniques. For example, the work I have conducted on placeholders could easily be modified to improve traditional web-favourite tools, making use of colour, position and space surrounding the main browser. The lightweight properties I discovered through my investigation into visual indexing can also be extended to web applications, either by providing users with relevant sections within a hierarchical web-page or by applying the techniques to a broader type of information retrieval such as search engines.
- Applying lightweight techniques to multimedia is also an area that could benefit from the work presented in this Thesis. For example, providing methods of placeholder or annotation within video or audio tracks may prove a useful feature, particularly if lightweight features such as colour, size and overview are used to clearly illustrate their positions. Digital images could also be greatly improved by incorporating lightweight properties. As with digital documents, the way in which pictures behave on paper is significantly different to the way in which they can be manipulated digitally, therefore, incorporating some lightweight features like I have done with documents in this Thesis, could increase user satisfaction with digital images.

## 8.2 Lightweight Application

This Thesis has focused strongly on creating implementations that utilise lightweight properties in order to produce more effective and intuitive active reading systems. Thus far, I have probed several areas of active reading that are currently seen as problematic digitally, and attempted to improve upon them by mirroring techniques used on paper, while incorporating digital enhancements. The results of these investigations have allowed me to lay down a set of lightweight attributes that can be used as guidelines for future active reading systems (see Section 7.5). To demonstrate the utility of the attributes I have discovered, I have included the following section that describes a specific area that represents a potentially fruitful area for new research. This section demonstrates how the list of attributes I have identified throughout my investigations into active reading can be applied to a new situation to create an improved interaction experience.

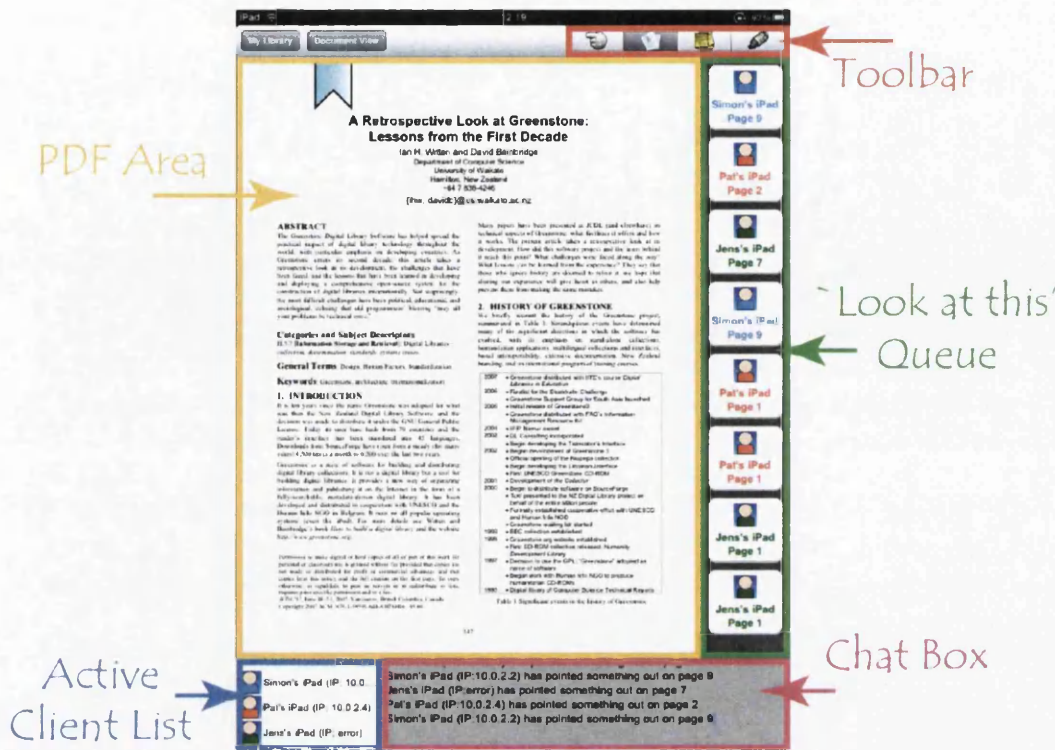


Figure 8.1: Screen shot from gsBooks

### 8.2.1 Collaboration

One area where the use of lightweight techniques could significantly improve the usability of the active reading tools, and one where I have already made a small amount of ground work, is the topic of collaborative reading. Reading is not often considered a social activity. In fact, the act of reading is typically solitary, for example, sitting alone quietly reading a novel or a newspaper. However, there are occasions where reading is performed interactively in groups. Some examples of these, where all members of the group are reading collectively from the same working document, are in revision/study groups, book clubs and classroom environments.

Despite the clear usage of collaborative reading, there is very limited literature studying how to design effective tools for group reading of the same working document. Although there are advantages to performing collaborative reading on paper (i.e., it is easy to scribble notes, flick between pages etc.), there are some areas where it can be considered considerably 'heavyweight'. For example, pointing out specific areas of a document to other members of the group as well as ensuring all notes made eventually end up on each user's personal copy of the document.

By applying some of the lightweight attributes I have discovered through my investigations into active reading, I was able to create gsBooks, an improved collaborative reading system that uses slate PCs to facilitate real-time updates of document mark-ups within a group environment. A screen shot from the gsBooks interface can be seen in Figure 8.1.

The basic architecture of gsBooks is a series of iPads connected to a central server. The backbone of this interaction, is the real-time communication between members of the working group. This con-



nection is used to provide each user with up-to-date mark-up information from other members and ensure that each client has access to the same working document. To achieve this, the gsBooks architecture uses a ‘cloud computing’ infrastructure that makes use of **independently saved annotation files**, a technique that I proved to be a lightweight property after my investigation into annotations (Chapter 4).

To use gsBooks, each client iPad must first connect to the host server sending it their IP address and user selected nickname. The server then allocates each user their own unique colour to aid in the distinction between group members. The use of **colour** within digital active reading systems is one of the lightweight attributes I have identified throughout the investigations described in this Thesis. It is for this reason therefore, that I opted to include it as an extra contextual element within the gsBooks interface. Using colour in this way, provides users with a visual aid to determine exactly which users created which mark-ups just like on paper where each member of the group can use a different coloured pen.

Once a client has successfully connected to the gsBooks service, it can then send and receive messages from the server. These messages are in XML format and are used to keep each client iPad up to date with any changes made to the working document by other members of the group.

Another useful feature of the gsBooks interface is its ability to ‘point out’ information to other users in a working group. Often in a group collaboration session, users will want to indicate items and locations to other group members. For example, “look at this figure on page 45”. When working on a single document or machine, this process is straightforward: the user can simply point out the section they are looking at with their finger. Assuming all members of the group have visible access to the shared document or screen, they will easily see the required information. When collaboration occurs over multiple copies of the same document however, this process becomes more troublesome.

To overcome this, I have incorporated a ‘look at this’ queue that allows users to ‘point out’ specific areas of the original document to other members of the group. The ‘look at this queue’ is designed to hold a list of temporary placeholders made by members of a working group to allow them to quickly and easily navigate to the exact area of the document indicated by the original creator. To ensure quick and easy access to this information then, it is essential that the list itself is always on display regardless of the currently open page. This idea of displaying information in the **static space surrounding the document** was identified as a lightweight property in several of my other investigations, including placeholders (Chapter 3), note-taking (Chapter 5) and indexing (Chapter 6).

When a ‘point out’ is made by a member of a group, it is subsequently added to the top of every members ‘look at this queue’ along with the colour and nickname of its creator. Receiving a ‘point out’ from another member of the group however, does not automatically navigate every user to the point at which it was made. Rather, each entry in the ‘look at this queue’ is hyper-linked, so users can choose a convenient time to access the information. The idea of **hyper-linked** information to reduce the overall navigation required by users was identified as a lightweight attribute in my investigation into visual indexing (Chapter 6).

For completeness, and to demonstrate the real-time properties of the system, the gsBooks system also includes two basic mark-up tools: annotations and bookmarks. Both the annotation and bookmark indicators are filled with their creators’ unique colour to give a quick **overview** of who created each one to other members of the group. Hovering the mouse over one of these icons provides the user with a **mouse-over pop-up** giving additional contextual information as to who created or edited each mark-up and when. Both visual **overview** and **mouse-over pop-ups** have been identified as lightweight properties in my previous investigations.

The following scenario provides a short description of a situation where gsBooks might be used:

*Kris, Joe, Laura and Hannah are studying for an exam. They each have a copy of the course text on their individual iPads and are discussing the material around a table. Joe notices that there is a section within one of the chapters that is not fully explained. He read a web article on this topic earlier in the day and decides to make a note clarifying the information for the benefit of his friends. Immediately after he makes the note, it automatically appears on the others' iPads along with a notification at the bottom stating that Joe has added a note. Later in the session, the group are discussing a past exam question and decide to look up the answer within the document. Kris finds the appropriate page and paragraph and 'points it out' to the other three via the tools within the system. Joe, Laura and Hannah's iPads instantly recognize that Kris is pointing something out and offers them the chance to automatically jump to the right page.*

This scenario illustrates how a lightweight reading interface has practical use within collaborative working groups. The real-time aspect of the gsBooks system ensures that each member of the group has immediate access to all notes made by other users, as well as facilitating an easy method of referencing specific sections.

In summary, this section has discussed the design and implementation of gsBooks, a system that facilitates real-time collaboration using iPads. The gsBooks interface makes use of several of the lightweight properties I have already identified throughout my other investigations to create a quick and initiative infrastructure for working groups. The gsBooks system supports real-time lightweight mark-ups between group members as well as providing a quick and easy method of pointing out specific areas of a document to other client machines. The slate PC form factor of the iPad provides a sleek interaction that mimics a physical clipboard to facilitate a familiar and closely integrated working environment.

The lightweight attributes I have used in the gsBooks interface are:

- The use of **colour** to give an **overview** of which annotations, bookmarks and 'point outs' are made by which members of the group;
- The use of **independently saved file copies** to create a cloud computing infrastructure and allow mark-ups and 'point outs' to be shared in real-time amongst group members;
- The utilisation of the **static space surrounding the document** to place the 'look at this' queue and 'chat box' which allows constant access to the changes and 'point outs' made by other group members;
- The use of **digital technologies** which make use of **hyper-links** for quick and easy access to modifications and also the real-time connection between iPads, a function that is not possible on paper equivalents;
- The use of **mouse-over pop-ups** to give further information about the user who created and edited mark-ups.

### 8.2.2 Related Publications

1. George Buchanan, Jennifer Pearson : **An Architecture for Supporting RFID-Enhanced Interactions in Digital Libraries** - *European Conference for Digital Libraries (ECDL) 2010*, Glasgow, UK.

2. Jennifer Pearson, George Buchanan : **Real-Time Document Collaboration Using iPads** - *BooksOnline '10: 3rd Workshop on Research Advances in Large Digital Book Collections*, Toronto, Ontario, Canada.
3. Jennifer Pearson, George Buchanan : **CloudBooks: An Infrastructure for Reading on Multiple Devices** - *Theory and Practice of Digital Libraries (TPDL) 2011*, Berlin, Germany.
4. Jennifer Pearson, Tom Owen, Harold Thimbleby, George Buchanan : **Co-Reading: Investigating Collaborative Group Reading** - *ACM/IEEE Joint Conference on Digital Libraries (JCDL) 2012*, Washington DC, USA. - **NOMINATED FOR THE VANNEVAR BUSH BEST PAPER AWARD.**
5. Jennifer Pearson, George Buchanan, Harold Thimbleby : **Investigating Collaborative Annotation on Slate PCs** - *International Conference on Human Computer Interaction with Mobile Devices & Services (Mobile HCI) 2012*, San Francisco, USA.

# Appendices

# Appendix A

## Research Techniques

I am learning all the time. The  
tombstone will be my diploma

---

Eartha Kitt

### A.1 Statistical Tests

When conducting studies that involve numerical data, it is important to ensure that this information is meaningful and properly represents the appropriate sample population. As Cairns and Cox once commented:

*“...we need to be sure that what we see in our numerical data is not just natural variation between people but variation due to the real differences between interfaces and their effect on people.”*

– Cairns and Cox [23]

To conclude whether the numerical data I gather throughout this Thesis is meaningful then, I conducted statistical significance tests. These tests will not provide un-deniable proof that the participants I have sampled are entirely representative of the complete population, but will, if proved significant, give strong evidence to support it.

There are a great many statistical tests available to make sense of numerical data. For simplicity, I am only describing the statistical tests I have used throughout this Thesis:

**Student t-test:** For comparing two sets of values when the data is parametric;

**Analysis of Variance (ANOVA):** For comparing more than two values when the data is parametric;

**Welch’s t-test:** For comparing two values with Gaussian populations where there is a large difference in standard deviation between the two samples;

**Kruskal-Wallis test:** For comparing more than two values where the data is non-parametric (i.e., where the results are not normally distributed or there are a relatively small number of discrete values rather than a scale; e.g., 5 point Likert);

**Willcoxon Signed Rank Test:** For comparing two paired groups with a non-Gaussian population;

**Chi Squared:** For comparing two or more values with a binomial outcome (e.g., on populations).

In addition to the standard tests discussed above, in some cases, I have also applied fixes for tests when their pre-conditions are violated:

**Bonferenni Corrections:** Allows the use of a t-test for more than two sets of values;

**Yates' Corrections:** Fix for Chi Squared test when the cell values are small (i.e., <5).

Therefore, there are also occasions throughout this Thesis where I will be making use of the following tests:

**Welch's t-test with Bonferroni corrections:** For comparing more than two values with largely different standard deviations;

**Chi Squared test with Yates' corrections:** For use on two or more values with binomial outcome where one or more of the cell values are small.

Studies that involve Likert scales, provide *ordinal* data and should not assume a normal distribution. Therefore, in these cases, I will use **non-parametric** statistical tests. I will also be taking a conservative approach to statistical testing in this Thesis and will consequently be considering ratings given out of 10 to also be non-parametric.

## A.2 Software Used

I have used a variety of software and programming languages to code the implementations described in this Thesis. Specifically, the languages I have used are:

- Java: which was used to code the Visual Bookmarking System (Chapter 3) and the Margin Annotation System (Chapter 4));
- C#: which was used to code the Visual Indexing System (Chapter 6) and the RFID Book Reader (Chapter 8.1);
- Objective C: which was used to code gsBooks (Chapter 8.1).

In addition, I have also made use of several IDEs and development tools including: the .NET Framework, XCode and Notepad++. The following subsection describes the third party rasterizers I have made use of within my implementations.

### A.2.1 Third Party PDF Rasterizers

This thesis is concerned with implementing innovative software to aid in the active reading process, not to demonstrate my knowledge of how PDFs are visualised. To avoid the extra burden of coding PDF rasterizers then, I have, in most cases, used third party software to handle the drawing of PDFs in my implementations. With the exception of the Future Work chapter, which makes use of a library within the XCode programming environment, I have tried and tested several third-party pieces of software that take over the rendering of PDFs.

**JPedal**

Java PDF Extraction Decoding Access Library is an open source library created by ID resolutions. The software, which has been designed for use with Java, suffers from several major problems including slow rendering speed and difficulties in extracting text. Despite these problems however, it was a good starting point for me to create software for multi-platform use and was consequently used in both the Visual Bookmarking System (Chapter 3) and the Margin Annotation System (Chapter 4).

**TallPDF**

After receiving comments from user study participants on the slow rendering speed of my first two implementations, I consequently decided to search out another third party PDF rasterizer. At this point I had also opted to switch from Java to C# to aid in my creation of graphical user interfaces. The software I decided upon for my next two implementations was TallPDF created by Tall Components Ltd, which was available on a trial basis from [www.tallcomponents.com](http://www.tallcomponents.com). This PDF rasterizer software was used in the Visual Indexing System (Chapter 6) and the RFID Book Reader (Chapter 8.1).

**PDFTron**

Although the TallPDF software was adequate, I felt that further exploration into PDF rasterisers would be beneficial before I making a purchase decision. With this in mind, I then moved on to PDFTron; a professional PDF rendering suite that can be cleanly integrated into the .NET framework. The PDFTron .NET SDK is available from [www.pdftron.com](http://www.pdftron.com), and although they do operate strict licensing conditions (i.e., unless you buy a multi-OEM licence, you are restricted to one machine), it proved to be smoothly compatible and easy to integrate with the tools I was implementing. I therefore used PDFTron in my largest implementation, The Digital Reading Desk (Chapter 5).





# Appendix B

## Technology Review

It's kind of fun to do the impossible

---

Walt Disney

### B.1 Introduction

To gain a better understanding of current digital reading features, I have undertaken a critical analysis of several popular reading appliances. By reviewing commonly used digital reading methods, I hoped to learn more about the problematic areas in order to better inform future designs.

#### B.1.1 Overview

Nowadays digital reading can take several forms. Users are no longer bound to read digital texts from immobile monochrome CRT displays [93] – instead they now have the option of reading from several more portable sources, such as laptops, tablets, mobile phones and eReaders. The fast delivery and large storage capacity that digital devices offer should make reading from bulky physical books seem obsolete, even nonsensical. Bistable eInk displays, the latest craze in digital reading, follows from the introduction of eInk which is illuminated, like print on paper, by reflected light and promises paper-like reading capabilities with added digital benefits.

The employment of digital reading has increased massively over recent years [134] prompting the production of dozens of portable reading devices into the consumer market. How well these devices support reading and active reading tasks however, is an issue that has been little investigated. There are clear advantages to digital reading devices over more traditional printed books. For example, someone going on a long holiday may prefer being able to carry many books without weight penalties or space constraints, but are then reliant on battery power and must take care not to damage the device by dropping it and so on. In contrast, conventional physical books rarely *fail*.

Traditional books are fixed, and can only serve one sort of use. For example, a large format book, which is easy to read for a user with vision defects, would be found large and unwieldy for individuals with good vision. Most portable reading devices allow the size of text to be changed under user control, so the *same* book can be read by a user with high acuity or with vision defects needing large text. Some devices can generate a synthesized voice, so they are suitable for blind users. Similarly, dynamic search of the text for a word or phrase is a distinct advantage to digital books. However, printed books behave (broadly) consistently, whereas digital interaction allows for a

greater variety and inconsistency. Clearly then, the design of the hardware and of the software within the devices is crucial to the overall reading experience of the user.

In this review, I discuss the usability issues associated with several popular electronic reading appliances. By analysing dedicated reading devices and applications, I hoped to gain insight into different features that can be considered digitally lightweight as well as determining what features suffer from interaction deficits .

Although reading can be performed on a variety of digital appliances (e.g., desktops, laptops, televisions etc.), I have chosen here to critically analyse portable reading devices only as I discuss features from common desktop reading software (e.g., Adobe's Acrobat and Apple's Preview) throughout the remainder of this Thesis.

## **B.2 Background**

Before embarking upon such a review, I first lay down a set of features to aid in the evaluation of each of the devices or applications. These features include standard reading tools such as bookmarking and annotation, relevant human-computer-interaction issues such as consistency and completeness of metaphor, and eDocument specific functions such as magnification.

### **B.2.1 Reading Features**

One of the main areas of portable reading I investigated, was the additional reading features that have close relation to the work presented in this Thesis. Specifically, the way in which bookmarks and annotations are performed on these devices, as well as the functions that they are dependent on (e.g., page numbering etc.), formed a large part of my understanding of current digital reading methods, and aided me in the creation of more lightweight alternatives.

#### **Bookmarks**

One of the more common features of reading is the simple and lightweight task of bookmarking. As I discuss in depth in Section 3, digital placeholders are significantly different from their physical counterparts and perform poorly in subjective user comparisons. The fundamental differences in the design of bookmarks for these two mediums largely contribute to the frequency they are used. How well these dedicated devices incorporate such a well used physical feature into their designs dictates the device's overall usability and consequently, how well it will be adopted by users.

#### **Annotations**

Annotations can be considered a by-product of the active reading process and can be easily achieved on paper documents using one of many lightweight methods. Assuming that eReader devices have catered for all reading audiences, as opposed to those who simply want to read a novel without making notes or placing bookmarks, how well do they incorporate active reading tools into their design? – A reading device that has poorly adopted annotation tools cannot be considered a complete working system as it hinders the physically lightweight process of making notes on books.

#### **Page Numbering**

The way in which documents are displayed can take several forms. When we are dealing with small screen devices however, the pagination options are limited. The biggest challenge when paginating

documents on small screens is that displaying the text in its original format is (usually) not viable as it would not fit entirely on the screen. To overcome this, digital documents can be reformatted to fit less text on the screen which in turn completely alters the page numbering system (see Magnification below). Reordering page numbers in this way can have a detrimental effect upon other functions in the device, for example, bookmarks which are highly dependent on pages.

### **Magnification**

One feature of electronic documents that cannot be accomplished on paper without the use of additional tools, is magnification. This magnification or 'zoom' function should ideally allow visually impaired users to read from the same device and read the same books as those with perfect vision. One side problem with this feature, is how these devices restructure the document to fit on the screen when the zoom level has altered. Fundamentally, this is not an issue, the device will simply put more or less words on the screen than the original document presented. Looking deeper here however, it is clear that this now alters what we know as page numbers. For example, if the original document had 100 pages and then we zoom in to double magnification, it will now contain around 200 pages. What was page one will now be spread over page one and page two. Clearly, this will now effect features such as bookmarks that rely on the page-referencing method of page numbers for relocation. How devices deal with this issue can effect a users' experience with reading, as well as the features that are bound by it.

Alternatively, a page can be zoomed to appear larger, without being re-flowed. This approach retains the original page format and page numbering, but now requires the user to pan across pages to read them. This panning action could disrupt the overall reading of text but is often the only option available for pictures or figures. This key design decision is typically forced by the format of the document being read: some documents cannot be re-flowed, so most devices must provide both mechanisms.

### **Page Turning**

Although page turning may seem trivial, how the device changes pages is a rather fundamental feature that if implemented incorrectly, could seriously hinder the reading process. When using the device for reading only, the most commonly used function will be the 'Next' button as it is used every time a user wants to change the page. The positioning of this next button therefore, is critical to the ease with which these devices are used. Placing the button in an awkward spot will cause unnatural hand movements which could be uncomfortable for users, and possibly lead to Repetitive Strain Injury (RSI).

### **B.2.2 Consistency**

When designing a digital interface, thought should be given to the consistency of buttons and functions within it to ensure smooth user interaction with the tools. Poor consistency leads to difficult interaction, and possibly low rates of use of the tools included in the system. It is vital therefore, that dedicated buttons should be well labelled and always perform the same function. For example, back buttons should *always* go back to the previous screen, zoom buttons should *always* zoom, and so on. Any variation in the functionality of these buttons could lead to confusion for the user.

### B.2.3 Completeness of Metaphor

Although there are certain affordances that cannot be replicated digitally, there are other actions and functions that could easily be incorporated into the software but are left out leaving the user with a strong sense of incompleteness. In a physical book for example, you can insert as many bookmarks as you please – this not the case however, in the majority of digital reading devices who often place restrictions on bookmarks allowing only one per page.

### B.2.4 Ergonomics









The overall ergonomic design of a digital reading device is crucial to its eventual success. Reading is an activity that can take place over prolonged periods of time, so, to optimise the reading experience, the design of these devices must provide an adequate and perhaps even invisible (see Section 2.4.2) platform for which to do so. More specifically, these devices need to be ready-to-hand (see Section 2.4.1) to ensure that users are not thinking about the devices themselves, but rather the material to which they are interacting. Occasionally, with large or heavy books, a user's attention may drift to how best the books can be positioned to make them more comfortable to hold, a scenario that can be easily overcome in digital document readers with a good ergonomic design.

Making these devices thin and portable enough to hold in one hand is only a portion of what should be considered before manufacture. For example, which functions should have dedicated buttons as opposed to on screen menus? – The incorporation of physical buttons may affect the aesthetic design of the device but may well aid its usability – a trade-off that must be carefully considered by the manufacturers. Further to this, how easy is it to access these buttons from the main fascia while holding the device as you would a book? And are frequently needed buttons or actions (e.g., for page turns) comfortable and flexible so the user is unlikely to get strain injuries? – In commercial computer-based activities there are protocols to follow, for example, the screen must be eye level, the chair must support the worker's back, the mouse must be within arm's reach and so on. Failing to provide these things may cause health problems for the worker after prolonged periods at the workstation (e.g., eye strain, neck or back problems, RSI etc.). Therefore, considering user comfort in the design of devices for digital reading – an activity that can take hours, is integral to user satisfaction of the end product.

Other potential ergonomic considerations include the re-assignment of buttons and actions so that the device can cater for both left and right handed users, the legibility of the display – which affects the magnification possibilities and the size of the physical buttons – smaller buttons can be difficult to use if the user has large hands or fingers and can be difficult to see for users with vision defects.

## B.3 Sony Reader Pocket Edition (PRS-300)

**Display:** 5 inch eInk (8-Level Greyscale)

**Physical Buttons/Devices:** Power Switch,  (Home),  (Return),  (Previous Page),  (Next Page),  (Next Page & Move Cursor),  (Previous Page & Move Cursor), <blank>(Enter),  (Bookmark),  (Size), Number Buttons X 10

**Resolution:** 800 x 600 pixels

**Size:** 159 x 108 x 10mm

**Weight:** 220g

**Operating System:** Linux

**Wireless Connectivity:** None

**Memory:** 512MB Internal




Figure B.1: The Sony Reader pocket edition (PRS-300)




**Expansion:** None

**Power:** Sealed Internal Lithium Ion Battery - around 6800 page turns on one charge

**eBook Formats Supported:** Adobe PDF, TXT, RTF, Microsoft Word, BBeB, EPUB

**RRP:** £199.99


## Bookmarks

A useful feature of the PRS-300 model that many other eReaders have neglected is its dedicated bookmark button . Clicking on this button within a document will cause a 'dog ear' to appear in the top right-hand corner indicating that the page has been bookmarked. Clicking the button on an already bookmarked page, deletes the current bookmark. This 'dog ear' visualisation however, does not give an overview of which pages are bookmarked as only the currently open page can be seen on screen. To view a list of all bookmarked pages, the user has one of two choices. The first option, that can be accessed through the main menu  (button) –  **All Bookmarks (on-screen menu)**, lists all bookmarks in all books stored on the device in one big list. The second and possibly more useful option, is to view only the bookmarks within a single book. This function however, must be accessed from the **Options** menu for that particular document which unfortunately, requires several button clicks to locate.

## Annotations

There is no facility to create annotations, notes or highlights on the PRS-300. The manual does confirm that these functions can be used within the eBook library on a paired PC, but the device would be unable to open them rendering the function rather useless.

### Magnification

Although there is no zoom mode within the device, there is a size  button that toggles between three different font sizes. This button however, does not change the font size of the home or options menus which seem to be set at a constant size throughout. Although the device does give visual feedback to state that the menus are not able to be zoomed, this is a serious inconsistency that could hinder visually impaired users.

### Page Numbering

The page numbering on the PRS-300 follows suit with most other Sony eReader models. The device does not keep the page numbers of the original document unless they are encoded into the document itself. If the document does include pre-encoded page numbers then, the pages stay constant throughout, that is, if you zoom in then the device does not reorder the pages, it simply splits each page into smaller pages and names them accordingly. For example, if an original document was ten pages and we zoom to 4x magnification, page one would now span over four pages – all of which would be called Page one. This method of page numbering is useful if users are coordinating with physical copies of the same book, for example, in a reading group when you are asked to turn to a particular page. It can be confusing however, to have multiple pages with the same page number. A more suitable solution to this problem when on a larger zoom level would be to keep the original page numbers but also include how many of them they are (i.e., Page X (Part Y of Z)).

The other, more common page numbering system used by Sony's eReaders, is when there are no page numbers encoded into the document which causes the device to calculate the number of pages on-the-fly depending upon the zoom level. Therefore, the more zoomed in you are, the more pages there are in the book. However, this method now leads to further problems, such as the shared reference problem encountered by those collaborating using the same document [84]. Previously, when two users were reading separate copies of the same document, they could say "hey look at page X, paragraph Y". Now to find the particular passage a user would also have to say "go to zoom level S then page X paragraph Y".

Another problem faced here is bookmarking. If a user places a bookmark on a page then zooms in, the old bookmark is 'lost' as only one of the multiple pages the original page now spans over is left bookmarked by the system. For example, if I were to bookmark page ten on the smallest zoom level then zoom in to 4x magnification, my original page ten now spans over four pages, but the system only bookmarks one of these four pages. Similarly, if a bookmark is made when the document is on a high zoom level, then it is zoomed out again, the bookmarked page is now far bigger than it originally was; how then, will the user know exactly which part they intended on bookmarking?



What the system does here is take the first line on the page being bookmarked and use it as the anchor point for the bookmark, ignoring any other text on the page. That is, whatever page the first word of the original bookmarked page is now on, is the only one which gets a bookmark on the new zoom level.

### Page Turning

The PRS-300 edition has several sets of buttons on its main fascia, including a circular directional pad which is used for page turning: ▼ and ► for Next and ▲ and ◀ for Previous. Presumably as a cosmetic feature, the buttons on the device have been arranged in a somewhat symmetrical pattern forcing these directional inputs to be placed in the middle of the device, directly below the display screen (see Figure B.1). Assuming the device will be held in one hand for reading, the position of






this directional pad is just about in range of the thumb of a right handed individual holding it on the bottom half. Left handed users however, may find it a strain to stretch their thumb to the next button after several page turns.

### Consistency

One of the major inconsistencies of the PRS-300 model is the functionality of its dedicated buttons. The return button  for example, does not always behave in the way you would expect, that is, it does not always return the user to the previously viewed screen. Under normal circumstances, for example, if you are in the Date and Time menu, clicking return , takes you back to the previous screen. However, in some instances it behaves as a menu button, for example, if you were on page five then use the number buttons to take you to page 55, the return button *should* take you back to page five. Similarly, if you were on the home screen then open a document, the return button *should* take you back to the home screen again. However, in these two examples, the return button actually takes the user to the options screen for a book instead of returning them to the previous screen, an inconsistency that can make viewing the options for a book time consuming and difficult.

## B.4 Sony Portable Reader System (PRS-505)

**Display:** 6 inch eInk (8-Level Greyscale)

**Physical Buttons/Devices:** Power Switch, Vol -, Vol +,  (Size),  <(Previous Page) X 2,  >(Next Page) X 2,  (Bookmark),  (Directional Buttons), <blank>(Enter), Menu, Number Buttons X 10

**Resolution:** 800 x 600 pixels

**Size:** 175 x 122 x 8 mm

**Weight:** 250g

**Operating System:** Linux

**Wireless Connectivity:** None

**Memory:** 250MB Internal




**Expansion:** Accepts Memory Stick Pro Duo (MS) up to 8GB, Secure Digital (SD) cards up to 2GB and Secure Digital High Capacity (SDHC) up to 32GB

**Power:** Sealed Internal Lithium Ion Battery - around 6800 page turns on one charge

**eBook Formats Supported:** Adobe PDF, TXT, RTF, Microsoft Word, BBeB, EPUB

**RRP:** £199

### Bookmarks

The bookmark function on the PRS-505 has the exact same functionality as the PRS-300; specifically there is a dedicated bookmark  button on the device which facilitates easy bookmark creation and deletion. The bookmark menus are also identical to those found on the PRS-300, that is, all bookmarks in all books can be accessed via the main 'Reader' menu by following the  All Bookmarks menu option and bookmarks within individual books can be found by following the 'Book' menu (Menu (button) –  Bookmarks) from within a book.

### Annotations

As with the PRS-300 there is no facility to create notes or annotations on the device.





Figure B.2: The Sony portable reader system (PRS-505)

### Magnification

The magnification on the PRS-505 follows suit with many other devices of this kind. For reflowable documents such as ePubs, the device allows users to alter the size of the text thus reformatting the entire document's page layout. This function is achieved on the device by pressing the  button and can display one of three font sizes. Sadly, as with all Sony eReader models, the magnification function does not extend to the menus of the device.


### Page Numbering

The page numbering on the PRS-505 is the same as on the PRS-300 (see Section B.3).

### Page Turning

The designers of the PRS-505 have opted to include two 'wheel' shaped directional pads on its main facia, one for page turning (where strangely, only two directions of the wheel actually work) and the other for controlling the on screen menus (see Figure B.2). In addition to these two directional pads, the device also includes an extra set of page turning buttons (next and prev) on the right hand side. Despite these buttons being perfectly positioned for the thumb of a right handed individual, they do seem rather unnecessary for a device that already contains a dedicated page turning wheel.

### Completeness

As I have already discussed, the PRS-505 contains not one, but two directional 'wheel' pads which perform slightly different functions. The left pad for example, only controls the two page turning  functions and does nothing when the user is not within a book. To integrate a page turning control into 'wheel' shape seems slightly unnecessary and confusing as only two of the many directions a wheel



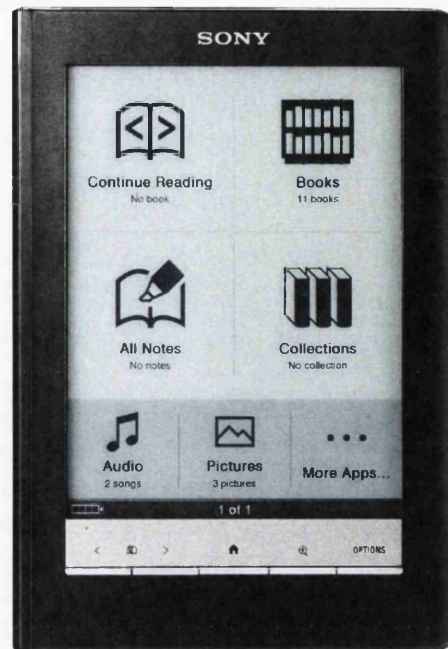


Figure B.3: The Sony Reader touch edition (PRS-600)

can perform are used in the function. A simple two-directional switch would have done the same job while at the same time saving space and improving the completeness of the device.

The second pad contains four directional buttons that are used to navigate within menus on the device and works as you would expect. The first experience of using the directional buttons while reading a book however, may seem very erratic and unpredictable. However, on closer inspection it becomes clear that the left ◀ and right ▶ directional buttons are in fact undo and redo features. There are no functions for the up or down directional buttons of the second wheel pad while reading a book, yet another example of an incomplete function on the device.

## B.5 Sony Reader Touch Edition (PRS-600)

**Display:** 6 inch touch-screen eInk (8-Level Greyscale)

**Physical Buttons/Devices:** Power Switch, Vol -, Vol +, ◀ (Previous Page), ▶ (Next Page), ⬆ (Home), ⊕ (Size), Options

**Resolution:** 800 x 600 pixels

**Size:** 175 x 122 x 10 mm

**Weight:** 286g

**Operating System:** Linux

**Wireless Connectivity:** None

**Memory:** 512MB Internal



**Expansion:** Accepts Memory Stick Pro Duo (MS) and Secure Digital (SD) cards up to 16GB



**Power:** Sealed Internal Lithium Ion Battery - around 7500 page turns on one charge

**eBook Formats Supported:** Adobe PDF, TXT, RTF, Microsoft Word, BBeB, EPUB


**RRP:** £279.99

## Bookmarks



One of the features of the PRS-600 model is a bookmarking system that allows users to mark multiple pages in a single document. There are two methods of creating a bookmark on the device but unfortunately, both demonstrate poor user interaction. The bookmark feature is classed by the device as a 'Note' and therefore exists within the Notes menu: **Options (button) –  Create/Edit Notes (on-screen menu) –  (on-screen icon)**. This is a rather long-winded method of performing a simple operation like bookmarking; a task that is so easily accomplished on paper. To get around this, the device has a short-cut bookmark option that requires double tapping in the top right corner. However, this feature is not supported by any kind of visual cue and therefore can only be discovered by close reading of the full user manual (which is not included in the box) or by random discovery.

To indicate to the user which pages are bookmarked, the device displays a small triangular 'dog ear' in the top right corner of the display which can be seen only when a bookmarked page is open. To view a list of all bookmarks within a document, the user must navigate through a series of menus: **Options (button) –  Create/Edit Notes (on-screen menu) –  Notes (on-screen menu)**. The 'Notes' menu is an ordered list of all annotations associated with a particular document, specifically, it merges every user-created mark-up item within a specific document into the same list (i.e. bookmarks, highlights, handwriting and comments). Each of these different annotation types are distinguished in the list by a different icon, and the lists can be ordered by page number, type or comment. One major flaw of this type of list-based interaction however, is the lack of overview of where bookmarks exist in the document. In paper books, the relative placement of bookmarks can be seen easily as well as how much of the document is left to read.



## Notes

The Touch edition Sony reader contains a within-document 'notes' system that can be accessed via a series of menus, specifically: **Options (button) –  Create/Edit Notes (on-screen menu)**. This then scales the document view very slightly to provide a small, clear space for the notes menu that sits along the top of the viewing window. Within this menu, users can either; add highlights to the book content, access free-hand mode using the stylus, delete using the eraser, add/remove bookmarks and access the notes list. In reality, notes and bookmarks mean different things to the majority of users. It seems odd, never mind longwinded therefore, that Sony has decided to combine them into a single entity and effectively bury one within the other. The following sub-sections describe the functions that Sony has grouped under the common name 'Notes'.

## Highlights

Existing text within the open document can be highlighted by first accessing the highlighter function hidden within the Notes menu: **Options (button) –  Create/Edit Notes (on-screen menu) –  (on screen icon)**. They can be removed by using the eraser tool also found within the Notes menu. As mentioned above, once a highlight has been created it is automatically added to the Notes list along with other types of annotations in the document.

## Handwriting

The device makes good use of its touch screen technology by incorporating a handwriting function into its design. By accessing the same Notes menu as other document annotation tools (**Options (button) –  Create/Edit Notes (on-screen menu) –  (on screen icon)**), users can 'draw' directly onto

the PDF surface with a stylus or finger. This is a useful function that may be better accompanied by variable pen thickness sizes to give better flexibility to the tool itself.

### **Comments**

Although you cannot simply add a text box to a page and type directly onto it, there is a tool that allows users to add hidden text and handwriting to both bookmarks and highlights (these are known as comments). Sadly however, this facility is so well hidden within the device that it can take weeks to find via random discovery. This situation is seriously exacerbated by the absence of a full manual with the Reader purchase. To add a comment, a previously created bookmark or highlight must be tapped once (double tapping accesses a different function) to bring up a menu that allows users to choose the type of comment they wish to anchor to the note: either handwriting or keyboard entry (Note: comment creation can only be accomplished when the Notes menu is off). Once a comment has been made and committed, it will disappear from view and can be accessed by tapping once on the icon adjacent to the bookmark or highlight. Not only does this functionality deny the creation of comments randomly on a page as each one must be anchored to an existing note, but each note (i.e., bookmarks or highlights) can only sustain one comment each.

### **Drawings and Text Memos**

As well as annotations on pages within a document, the Sony Touch Reader also allows handwritten drawings and text memos to be made independently of any documents. This feature allows the reader to be used as a separate notepad as well as a reading device.

### **Magnification**

The PRS-600 offers two types of document resizing. The first, which allows users to change to one of five different font sizes (S, M, L, XL and XXL, where S is the original size), is designed to re-render and display a newly formatted version of an open document page with the desired font size. This feature however, can be slightly temperamental on certain documents. For example, some tables, specifically those that span the entire width of a two columned document prohibit larger font sizes for that particular page (i.e. M, L, XL and XXL). In addition, some images, particularly embedded PDF images, are rendered incorrectly with the device reformatting any text within them along with the main body of the document. To get around any possible magnification problems, the device also offers a standard zoom function that does not reformat document pages, but zooms in on specific points and facilitates panning using four directional arrows on the touch screen.


One major downfall of the device is its lack of consistency of the zoom function. Specifically, one can zoom in on documents but not on menus, making the device extremely difficult to use for visually impaired users.

### **Page Numbering**

The page numbering on the PRS-600 is the same as on the PRS-300 and PRS-505 (see Section B.3).

### **Page Turning**

By adding a touch screen, Sony has been able to make the eInk display the main focus of the device by limiting the number of physical buttons on its fascia. They have opted to place 5 thin buttons on

the device and place the icons of their functionality just above them effectively reducing their target boundary. The positioning of these buttons however has not been thoroughly thought out and looks strikingly similar to the layout you would expect to find on a digital music player (see Figure B.3). The page turning  buttons (<) and (>) have been placed adjacent to each other on the left hand side of the device. This puts the most commonly used function, i.e., the next page button (>), in an extremely awkward position and one that would not suit the thumb or finger position of a left or right handed individual.



In addition to using dedicated hardware buttons, pages can be changed on the PRS-600 by making sweeping gestures on the touch screen. To move to the next or previous pages within the document, the screen must be dragged in the appropriate direction with a stylus or finger, and to repeatedly change pages (i.e., flick quickly through document), the screen must be held after the gesture has been made. The default page turning gesture directions have been designed to coincide with the buttons on the device, specifically, gesturing from right to left (<) takes you to the previous page and from left to right (>) takes you to the next page. These directions, however logical, contradict the learned behaviour of reading physical books. For example, when reading ahead in a paper book (i.e., going to the next page); one would take the **right** page and flip it to the **left**. Conversely, to go back to a previous page, one would take the **left** page and flip it to the **right**; completely opposite to the gestures provided in the device's hardware. Fortunately, the page turn gesture can be modified from its default direction by altering the device settings in the main menu.

### Consistency

As mentioned above, a major problem with the device is its inability to zoom menus making the device extremely inconsistent and a poor choice for visually impaired users. In addition to this, the main options menu unnecessarily spans two pages (i.e., you must click to change to the second half of the menu) despite the fact that there is sufficient room on the screen to place all items within it on one page. This adds a needless extra click to an already slow menu system.

## B.6 Amazon Kindle 2

**Display:** 6 inch eInk (16-Level Greyscale)

**Physical Buttons/Devices:** Power Switch, Vol -, Vol +, PREV PAGE, NEXT PAGE X 2, HOME, MENU, BACK, Full QWERTY Keyboard,  (Symbol),  (Text Size), 5-way Directional Controller

**Resolution:** 800 x 600 pixels

**Size:** 203 x 135 x 9 mm

**Weight:** 289g

**Operating System:** Linux

**Wireless Connectivity:** 3G + Wi-Fi

**Memory:** 2GB Internal

**Expansion:** None

**Power:** Sealed Internal Lithium Polymer Battery - up to one week (with wireless on) on one charge

**eBook Formats Supported:** Kindle (AZW), TXT, ADOBE PDF, HTML, MS WORD

**RRP:** £149.99

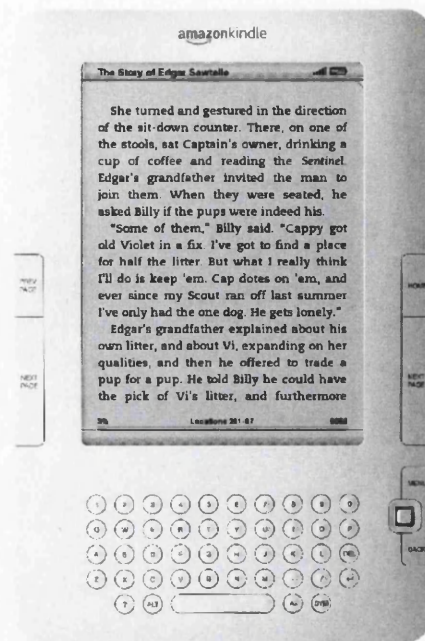


Figure B.4: The Amazon Kindle 2

## Bookmarks

Following suit with the Sony Touch's design, the Kindle does not have a dedicated button for bookmarking. Instead, to bookmark a page the user must click **Menu (button)** then scroll down six items and select **Add a Bookmark (on-screen menu)**. To bookmark a page on the Kindle then, requires eight button clicks; a rather large number for such a commonly used function. On closer inspection of the user manual, I discovered the Kindle does have an accelerator short-cut for this which makes bookmarking far less cumbersome: **ALT (button) + B (button)**.

When a page has been bookmarked, it renders a 'dog eared' corner in the top right of the screen. Surprisingly however, pages that have not been bookmarked, show a dashed dog ear in the top right corner; a confusing design that has no apparent function whatsoever.

## Annotation

Users can add notes to specific points in the text by first entering annotation mode, then moving the cursor to the place in the text they want to anchor the note. To create a note in a specific location, the user only needs to start typing using the full QWERTY keyboard included on the fascia of the device. They then select the 'save note' option using the five-way directional stick and a new note is created in the specified position.

## Highlights

To create a highlight the user must first be in annotation mode which can be accomplished either by moving the caret into the document area (simply by using the directional stick) or by selecting **Menu (button) – Add a Note or Highlight (on-screen menu)**. Next, the user must move the caret to the

beginning of the place they wish to highlight using the five-way directional controller and 'start' by clicking the five-way down. The caret then turns italic and will highlight text that is selected using the five-way directional buttons. Clicking the 5-way down again will finish the highlight which is subsequently indicated by a faint grey underline.

### Magnification

The Kindle 2 offers six font size variations, accessible through the size  $\text{A}$  button on the device. It also allows users to change the margin size of the text using the strangely worded 'Words Per Line' feature of: *fewest, fewer, default*. Following suit with the Sony's inconsistency, the Kindle 2 also sets its menus at a fixed size. Unlike the Sony's however, the Kindle gives no visual, audio or tactile feedback when attempting to change the font size within a menu.

### Page Numbers

The Kindle 2 has opted to do-away with conventional page number format and replaces them with 'locations' that correspond to specific places within the text. Instead of page numbers then, the Kindle measures the file in locations that are linked to specific positions of text within the document. Each location always corresponds to the same position in the text regardless of how the text is being split for the display, that is, how large or small the font is. This means that one screen can contain more than one location depending upon the zoom level and where the text splits.

Despite this location system however, the device still seems to have trouble with bookmarks. As with Sony devices, the Kindle 2 also takes the first word of the originally bookmarked page to be the anchor for the bookmark and ignores any text thereafter. This means, of course, that if the interesting part of the page you bookmarked happened to be at the bottom, you would lose its place if you changed to a different zoom setting.

### Page Turning


The Kindle designers have carefully engineered their hardware to cater for both left and right handed individuals. That is, there are *two* 'Next' buttons positioned on the sides of the device in such a way that they can be utilised by either the left or right thumb and more importantly, can be used without significant strain, with only one hand (see Figure B.4).

Since books are designed to be read sequentially from start to finish, the previous button is expected to be used less frequently which is reflected in its design position – a single small button above the 'Next Page' button on the left of the device.

## B.7 Barnes and Noble nook

**Display 1:** 6 inch eInk (16-Level Greyscale)

**Display 2:** 3.5 inch colour touch-screen LCD

**Physical Buttons/Devices:** Power Switch, 2x < (Previous Page), 2x > (Next Page), Menu Touch Button 

**Display 1 Resolution:** 800 x 600 pixels

**Display 2 Resolution:** 480 x 144 pixels

**Size:** 196 x 125 x 13 mm

**Weight:** 343g

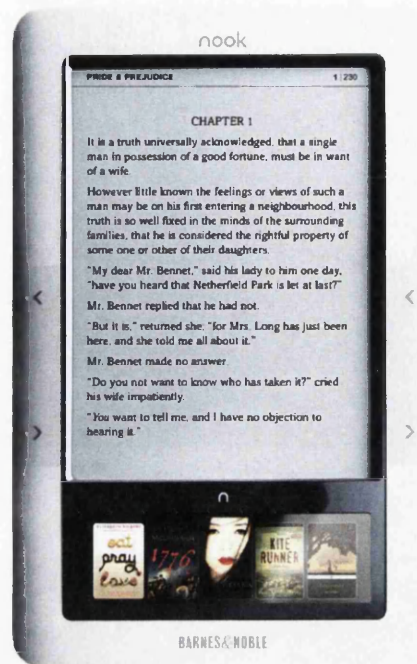


Figure B.5: Barnes and Noble nook

**Operating System:** Android

**Wireless Connectivity:** 3G + Wi-Fi

**Memory:** 2GB Internal


**Expansion:** Accepts microSD cards up to 16GB

**Power:** Replaceable Lithium Ion Polymer Battery

**eBook Formats Supported:** PDF, EPUB, PDB, BMP, TXT

**RRP:** \$199

## Bookmarks

The nook, like the Kindle 2 and the three Sony models, indicates bookmarked pages via icons on the pages themselves. In this case, the device displays a tiny, almost unidentifiable icon  in the top right corner of the screen. This icon can be hard to see at the best of times, particularly if it is placed on a page that is not blank in that area, for example, a black cover page will completely obscure the bookmark icon.

There is no dedicated bookmark button on the device. To view all of a document's bookmarks, you must first be within the document, either by selecting 'Reading Now' or an entry from the 'My Library' options found in the Home menu on the touch screen. Then, in the document options displayed on the colour touch screen, you can select: **Bookmarks (on-screen menu) – Go to bookmark (on-screen menu)**.

The bookmarks menu within the device is in the form un-ordered list, that is, bookmarks exist in the list in the order they were added. There is also no facility to change the name of the bookmarks; they are simply named after their number, for example, 'page x'. There are also similar bookmarking issues arising from page numbering as seen on the Sony models (see Section B.7).

As well as user-added bookmarks, the nook also has an automatic function that saves the last position a user was reading within each book. This information is then displayed under the title of each document within the library menu.

### Annotations

Annotations on documents on the nook can be made by first highlighting words within the text; notes cannot exist without a highlight. Unfortunately, creating highlights on the nook demonstrates extremely poor interaction. Due to the lack of physical buttons on the device, highlights must be made using the touch screen which displays a four-way 'D-pad' (see Figure B.6). The lack of feedback from their D-pad touch-display however, makes navigating to and selecting specific words very cumbersome. To add a highlight to a document, the user must first select: **Highlights and notes (on-screen menu) – Add highlight or note** to bring up the D-pad display in Figure B.6. The D-pad is then used to navigate to the place where the annotation starts, before the user selects the **Start selection** option. The D-pad is then used again to select the place where the highlight ends before selecting **End Selection** and finally **Submit**.




Figure B.6: The nook's 'D-Pad' display

### Magnification

As with every other eReader I have studied, changing the font size on the device does not change the size of the menus on the eInk screen or the colour touch screen.

### Organisation

One feature of the most recent nook (software version 1.5) which has not been included in devices such as the Kindle or the Sony's, is a 'shelf' organisation system. Readers can use the shelf system to organise their books into personal categories. Within the Library menu, it is possible to: go to a shelf, place or remove books from a shelf, create a new shelf and remove a shelf. Users can then view books either in one large list, or in a list of shelves.

### Page Numbering

Instead of reordering page numbers on-the-fly when a zoom level is changed, the nook maintains the original page numbers encoded into the document – therefore, resizing the text does not reorder the pages. Unfortunately, the way in which the device deals with this method of numbering is poor. For example, if there is more text on the screen than there was on the original page (i.e., if the text font is very small and can display more than one of the original pages on the screen), then it does not display both page numbers. Instead, it simply picks *one* of the original page numbers to displays. Conversely, if one original page no longer fit on the screen (i.e., if the text is very large) then the original page is



displayed on multiple pages, for example, Page four could now span over six pages – all of which will display ‘Page 4’ at the bottom of their screen.

## Page Turning

In comparison with Amazon, Barnes and Noble have opted to include two sets of page turning buttons < > (two on the left and right sides), presumably to cater to both left and right handed individuals (see Figure B.5).

The main eInk display on the nook is not touch sensitive. Below this screen however, exists a second display, in this case, a colour touch-screen LCD which is used primarily for displaying menus and as a touch pad for navigation and page turning.

## B.8 iPad

**Display:** 7.9 inch backlit multi-touch colour LCD

**Physical Buttons/Devices:** On/Off, Screen Rotation Button, Vol -, Vol +, Home Button, Accelerometer, Ambient Light Sensor, GPS

**Resolution:** 1024 x 768 pixels

**Size:** 243 x 190 x 13 mm

**Weight:** 730g

**Wireless Connectivity:** Wi-fi and 3G Connectivity

**Memory:** 64GB Internal

**Expansion:** None

**Power:** Sealed internal lithium Polymer battery - around 9 hours of web surfing (3G) on one charge

**RRP:** £699

## iBooks

### Bookmarks

The bookmark interface on the iBook application uses red ribbon-shaped markers which are positioned in the top right corner of the page. This interface makes it easy to see whether or not a currently open page is bookmarked, but cannot be used to give an overview of all bookmarks in the document. To view a list of every bookmark within the document, the user must navigate to the beginning and select the bookmark button.

### Notes and Highlights

The iBook application includes rudimentary highlighting of words and images within the document, but, due to the way it has been implemented, is tricky to control with the primary input device: a finger. In addition to this, the application also allows users to make notes on their documents but only in connection with words within the text. To add a note to a document, the user must first select a word or selection of words from the document and click the ‘Note’ option. This first highlights the word in the same manner as the highlight function, then brings up a Post-it style sticky note to which the user can add text. Once the text is completed, the application places a small Post-it icon in the margin close, but not always in line with the word to which it relates. These icons are distinguishable from each other by the date in which they were added; notes made on the same day produce identical





Figure B.7: The Apple iPad

icons. There is no facility to change the caption on these icons making it very difficult to tell them apart.

To retrieve a previously created note, the user must click on the note icon itself, *not* the highlighted word to which it relates. Clicking on the highlighted word only gives options to change the colour of the highlight or to *remove* the note attached to it. Connecting notes to their respective highlighted words then, is extremely difficult if not impossible if there are multiple notes on the same line or page.

This problem with accessing notes is further exacerbated by the way in which the note icons are positioned; in single file list in the margin of the document. In itself, this causes problems with referencing as mentioned above, but also with space as the interface provides no method of scrolling through these notes once the margin is full.

### Magnification

When using iBooks for the default document type (ePub), the conventional 'zoom' gesture (pinching two fingers) does not cause the document to be magnified as you would expect. In replacement to a zoom on this document format, the application allows users to adjust the text size and font by selecting the  icon and tapping the larger or smaller options (11 sizes total). This method does not however, change the size of any images within the document which remain constant throughout.

As with many of the eReader devices I discussed earlier, the menus within the iBooks application do not zoom with the document content.

## Organisation

Like many eReader devices, the organisation of the iBooks application is poor. Despite the attractive ‘book shelf’ style document list, which displays the title and cover image of a book, the iBooks application does not facilitate folder organisation or even user-defined tagging of books for easy referencing. Surprisingly, it simply throws all the books in the collection into one large bookshelf; an undesirable situation if the user owns a large collection of eDocuments.

If the user selects the ‘list’ view, they have the option of sorting the books within the device either: in the order in which they were added (Bookshelf view), in order of title (Titles view), in order of author surname (Authors view) or by Category (Categories view). All these options use information encoded into the eDocument upon creation and give no facility for users to tag or organise documents themselves.

## Page Numbering

As with many of the other devices I have studied, the iBooks application reformats the page numbers in a document when the font size is changed. A nice additional feature of the application is small text area telling the user how many pages there are left in the current chapter (e.g. five pages left in this chapter) – a useful feature that can be difficult to calculate manually.

## Page Turning

The application follows the book metaphor closely and has engineered a neat page turning interface that mimics the way paper pages behave. By using a finger, pages can be grabbed, turned slowly, or almost folded to reveal the page underneath. This, unfortunately is as far as the graphics can be manipulated. Despite the application showing multiple pages behind the open ones (like a real book), it is impossible to ‘grab’ these pages with a finger; only the page immediately preceding and following the current page can be grabbed at any one time. The interface would be greatly improved if users were able to grab more than one page either side of the open one as opposed to simply viewing an inactive graphic that depicts the book’s other pages.

As well as to the graphical page turning functions in the application, there is also a simpler method which involves tapping the right (to go forward) or the left (to go backwards) sides of the screen to change page. In addition to this, the application has also included a scroll-bar at the bottom of the screen that not only facilitates quick flicking through a document, but also gives a neat overview of how far through the book you are currently reading.

## Consistency

One of the main inconsistencies of the iBook application is the way in which documents are displayed. For aesthetic purposes, the book metaphor the designers have used includes a background graphic that depicts book pages behind the currently open ones. This design has been included to make the digital text look more physical, like a real book. Unfortunately however, this graphic does not behave in any way like a real book should. As I have mentioned above, the many pages behind the current ones cannot be grabbed and turned – a feature that, if implemented could allow users to flick easily through pages like they would in a real book. As well as its lack of functionality, the background graphic of the application does not change when the current page is either the first or last page of the document (double page spread view). For example, when the title page (i.e., page one) is open, it appears in the right hand side of the double page spread, like a real book should if it was closed, but



Figure B.8: The Library of the eBook Application on the Samsung Galaxy Tab

the background graphic has not altered, leaving the left hand side blank and still showing other pages behind it. Similarly, if you are on the last page which ends on the left, the graphic has not changed and the book still looks open.

Another inconsistency, this time in conjunction with the device itself, is the magnification function. The iPad hardware includes a sensitive multi-touch display that is used by other device applications to zoom by pinching together two fingers. The iBooks application however, does not make use of this gesture; a small inconsistency that could prove a nuisance to those well trained with iPad and iPhone software.

## B.9 Samsung Galaxy Tab

**Display:** 7 inch backlit multi-touch colour LCD display

**Physical Buttons/Devices:** On/Off, Vol -, Vol +, @TODO GPS, 3MP rear-facing camera, 1.3MP front-facing camera

**Resolution:** 1024 x 600 pixels

**Size:** 190 x 120 x 12 mm

**Weight:** 380g

**Operating System:** Android

**Wireless Connectivity:** Wi-fi, Bluetooth and 3G Connectivity

**Memory:** 16GB Internal

**Expansion:** Accepts MicroSD cards up to 32GB

**Power:** Removable Lithium battery - around 7 hours of video playback or 10 hours of talk time

**RRP:** £299.99

### eBook

#### Bookmarks

The bookmarks in the eBook application, as with the majority of other reading software I have evaluated, do not mark specific positions in text, but rather the page the text happens to be on. This again, creates problems with consistent referencing when the magnification level of the device is changed.

The application also remembers by default the last page the user was reading as a type of temporary bookmark, and visualises this information nicely in the library section of the interface (see Figure B.8). This library interface provides a clear overview of approximately where the reader is in each book, by inserting a small red ribbon into each book icon in the shelf. These ribbons change position

vertically down the book depending upon the last page the user was reading (Note: these are not user created bookmarks, they are simply a record of the last read page in the document).

### **Highlights**

The app also has the facility to add highlights to a document by simply pressing and holding a finger over a word. This then presents the user with a pop-up menu with several options such as: copy; search; highlight and dictionary. There is no facility to add text to books within the application.

### **Magnification**

The eBook application allows resizing of the text within the documents to one of seven pre-defined font sizes. As with all the other devices I have reviewed however, changing the size of the font is not consistent as all menus, headers and footers stay at a constant size.

### **Page Numbering**

Following suit with the majority of other electronic reading devices on the market, the eBook app reorders page numbers according to the size of the font which, again hinders the behaviour of the bookmarking function.

### **Page Turning**

The page turning on the application is achieved via sweeping or double tapping gestures on the touch screen and makes use of graphical page 'curling' visualisations. The implementation of this book metaphor is not complete however, as only one page can be 'grabbed' and curled at a time.

## **B.10 Discussion and Conclusions**

The following section brings together my findings from the critical analysis of the devices I have reviewed.

### **B.10.1 Page Numbers**

It is clear from experience that what we know as 'pages' mean little in the digital world. Unlike paper documents where we are bound by page numbers, digital texts are free from the constraints of physical pages and can be re-formatted easily; completely altering the numbering system. The main problem with this method, is that digital 'page numbers' do not correspond to anything concrete within a document as pages are resized depending upon the size of the text. For example, when a book is at its original size it may contain 100 pages. Altering the font of this document to a larger size will now cause the pages to be re-constructed and it may now span over 120 pages.

This reformatting can lead to a series of user issues ranging from placeholder problems to the difficulties faced in referencing physical books with their digital equivalents. It is also undesirable for users who remember important parts of a document based on page numbers and indeed in collaborative situations where two or more users are attempting to refer to the same page when on different zoom levels.

An alternative to page numbering in digital documents is to use specific document *locations*, as used in the Amazon Kindle models. These locations correspond with specific points in the text and

remain unchanged after each reformatting. So, instead of referencing text by page numbers like the majority of electronic reading devices, for example, Page 40 of 150, the Kindle uses locations, for example, Locations 132-145 (312) and specifies the percentage of the document already read. This method of referencing is consistent, and solves many of the problems mentioned previously. It can however, make remembering specific points tricky, as locations can quickly run into thousands in large documents.

Another possible solution to this problem could be to use locations as a method of defining exact points in the text and use them to define the numbering system seen by the user. In this scenario, the original page numbers will be kept in tact by splitting them into parts if the font size is changed (for example, Page Z (Part X of Y)). This method of page numbering ensures that the original page numbers correspond to those in their physical equivalents for easy referencing between the mediums as well as making other features such as bookmarking easier to accomplish.

### **B.10.2 Page Turning**

The way in which pages are turned on digital reading hardware varies from device to device. Some make use of dedicated buttons which are often poorly positioned causing possible discomfort to users after long periods of reading. For example, placing the button in the centre of the device means reaching to press it with the holding hand will cause an unnatural stretch of the thumb. Other designs utilise touch-screen technology to change pages, either by using taps or 'sweeping' gestures to move between single pages.

One common theme that runs through many portable reading systems is page turning animation – that is, a visual representation of the current page being pulled back to reveal the next page underneath. This visual approach to page turning attempts to make the feature more paper-like by mimicking the physical method allowing quick flicking between two adjacent pages within a document. Unfortunately however, the page turning animations in all the devices I have reviewed only extend to the next (or previous if in double page view) sequential page; subsequent pages cannot be grabbed using the feature. This means that although additional pages can usually be seen behind the next page, the software does not let the user pick up multiple pages and therefore cannot be used to 'flick' through the remainder of the document or switch easily between two non-adjacent pages. A feature such as this which would allow users to quickly skip through future pages to get a feel for what is coming up, would be extremely useful, particularly for non-sequential reading such as magazines [85].

### **B.10.3 Bookmarks**

The majority of the reading devices I evaluated focused on permanent placeholders, that is, putting in a new placeholder does not discard the old one. Some of them however, did incorporate a temporary 'bookmark' placeholder into their design, for example, a marker that is used simply to keep the place in which reading stopped in the last session. Like in physical documents, a bookmark will be kept until the user reads ahead in the book when it is then removed and replaced at the new end point. Although the hardware and applications I have reviewed do not explicitly call this feature a bookmark, the 'carry on reading' tool incorporated into many of these devices is a good example of a temporary placeholder.

Bookmarking is an extremely common feature of reading, yet many of the reading devices I evaluated did not include a dedicated bookmark button. For many of these devices, bookmarking a page is accomplished by navigating through an often lengthy series of menus. The usability of these

devices would be greatly improved if more commonly used features like bookmarking were given dedicated physical buttons on the device's fascia.

A more pressing usability issue of the placeholder functions within many of these devices is how they are displayed. In physical books, most placeholders can be seen sticking out of the pages. Unfortunately however, this lightweight feature of physical placeholders is not as easily achieved in the digital world. In fact, many of these devices make little attempt to even display a static list of bookmarks, never-mind the visual overview of their related positions. In the majority of the reader devices I evaluated, a bookmarked page could only be seen while reading if it was, itself, a bookmarked page. To view all bookmarks within a document, the user would need to navigate away from the document (using menus) to retrieve a list of placeholders and notes.

Another major downfall with the bookmarks on these devices is the fact they assume pages are constant when in fact they are not. They let you bookmark 'pages' which consequently get 'lost' once a new zoom level is selected. A more accurate method of marking points in the text would be to use the location method described above, where specific points in the text are referenced as opposed to the (almost random) points in which the text split (i.e., page breaks).

#### **B.10.4 Magnification**

A major advantage that digital texts have over traditional material is the fact they can be magnified easily to suit the needs of the user. In the physical world, users with visual impairments would normally have to use a magnifying glass or even buy a copy of the document with larger text. Digital documents however, can be easily manipulated, and the text changed to whatever size the user requires.

After examining these devices it is clear that there are two types of digital magnification:

1. **Zooming:** The font is enlarged when the viewing window zooms into a smaller part of the page. In this method directional buttons are required to pan around the page, but the original page numbers remain unchanged.
2. **Reformatting:** The size of the font is changed, which automatically alters the format of the pages. This method means there is no panning required but causes the original page numbers to be lost.

In general, document formats which are designed to be reflowed (e.g., ePubs), are magnified via reformatting, which causes problems with page numbering (see Section B.10.1). More constricted document formats such as PDFs which cannot easily be reformatted, make use of the zooming and panning option.


The main problem with the majority of these devices is that the magnification function does not extend to the menus which stay constant throughout. This serious inconsistency renders many of the device's functions useless for those with poor vision unless a third party performs all the navigation. These devices would be greatly improved if they a) included both reformatting and zoom magnification, and b) extended the magnification to include menus.

#### **B.10.5 Organisation**

In the physical world, those with a large number of books may want to separate them into genres for easy referencing, for example, keeping all the cookery books on one shelf, all the reference books on another and so on. Organisation for documents within digital reading devices however, is surprisingly poor with the majority forcing all books into one long list. Although many of them do sort by title,

author and/or date, most do not allow custom organisation into folders or even into genres, making browsing through a large collection extremely difficult.

### **B.10.6 Consistency**

The consistency of the software within devices is certainly something that can be improved upon. I have touched upon several cases within my evaluation of the chosen devices in which the consistency of certain functions does not follow platform conventions. Examples of these range from button inconsistencies as in the case of the Sony PRS-300's Return  button, to function inconsistencies, in the case of many zoom functions' inability to increase the size of menus. Even the page numbers within several of these devices are inconsistent (see B.10.1). Recalculating page numbers at every zoom level undermines the traditional connection between a word or sentence and a particular page (or location) identifier, and complicates the interaction for functions that are dependent upon them (e.g., bookmarks).

### **B.10.7 Completeness**

Digital devices that mimic physical ones have certain expectations that are not always fulfilled; consequently the sense of completeness within devices such as these, is rare. For example, in a paper book it is possible that in certain situations you would want to add more than one bookmark to a page; perhaps in the form of Post-its to mark interesting parts on the top and bottom sections. This function however, is impossible on the majority of devices; in all I sampled, a page could either be bookmarked or not bookmarked with no facility to add additional ones.

Another example of this would be flicking through a document: a process that is easy to accomplish on paper and useful for quick navigation. Achieving this interaction on digital reading devices however, is far more cumbersome. In the majority of eReaders I reviewed, there is no comparable interaction to simply 'flicking' through the document with a finger. The three Sony models for instance, had no tool for such an interaction. The PRS-300 and PRS-505 models did include dedicated number buttons for navigation to specific page numbers, but neither allowed for the free flicking through numerous pages within a document. In contrast, the Kindle 2 and the nook do allow scrolling through pages using additional hardware (the Kindle's five-way directional stick and the nook's additional colour screen), but do not provide a complete flicking interaction primarily due to the slow refresh speed of eInk displays. Another solution to this problem as used by more general devices such as tablets and phones, is to include touch screen scroll bars to flick quickly through pages.

### **B.10.8 Ergonomics**

In work-based computer tasks there are many protocols to follow; for example, the screen must be at eye level, the chair must keep your back straight and so on, otherwise the user could end up with eye strain, neck/back injuries or RSI. A major appeal of eInk devices is the 'paper-like' viewing of non-backlit displays, a feature that is designed to reduce eye-strain. Reading for a prolonged period of time in the same position may also trigger other problems however, particularly if a repetitive action is being performed; for example, pushing the 'next' button to change the page on an eReader. To reduce the possibility of strain injuries while using these devices then, it is essential that the positioning of commonly used features is carefully considered. During my investigation into digital reading devices, I found that several of the top-selling eReader products on the market have been designed with aesthetics as opposed to comfort in mind. The Sony Touch for example, has opted for the minimalist



design approach, putting only five thin strip buttons on its fascia. Unfortunately however, the most commonly used button on the device (i.e., the 'next' button) has been placed second from the left in the row of buttons – an extremely awkward position for users to press while holding the device with one hand (regardless of which hand they use).

### **B.10.9 Problematic Areas and Suggestions for Improvement**

There are several areas of digital reading design that could be significantly improved by paying attention to the basic HCI principles outlined previously.

- In the case of devices such as eReaders that do not offer a touch-screen interface, the button choice and placement on the device's fascia should be carefully considered. Commonly used functions such as next and bookmarks should have their own dedicated buttons and should be placed conveniently to avoid discomfort during long reading sessions. As I have discussed, the 'next' buttons on the Kindle are perfectly positioned to coincide with the thumb positions of users holding the device in their left or right hand, whereas the corresponding button on the Sony PRS-600 is poorly positioned for even ambidextrous users;
- Further to button positioning, device manufacturers must also ensure that the buttons and functions within the device are consistent. Users should not have to wonder what action the button or function will perform every time they use it; it should consistently perform the same action;
- Designers should also strive for completeness: ensuring that tools and actions within the device mimic the actions that can be performed on paper, unless this is demonstrably inefficient in a digital interaction. In doing so, the functions within the device will become more lightweight and improve the interaction between the user and device;
- There are certain principle conventions that should be followed in order to ensure good user interaction with the tools. For example, due to the technology used in eInk devices, the screen update speed is slow. Therefore, an obvious remedy to this problem would be to reduce the number of screen updates. Sadly in the examples I chose there were several instances where the devices failed to do this. In the Sony Touch for example, the main menu screen is not one big list; you must scroll down after several entries to view the second half despite there being ample room for the entire list on the screen at one time. This causes unnecessary button clicks and page refreshes that could easily be avoided if the menu was contained in a single list;
- The organisation of documents within portable reading devices is an issue that few designers have addressed; many simply dump all documents into a single list. To provide a fully integrated reading interface, provisions must be made to ensure users are able to reorder and organise their book collections into bespoke clusters for easy access later;
- Pages in the digital world are a forced metaphor that have migrated (badly) from physical books. The problems that arise from the reformatting of text for zooming and so on (e.g., difficulties in referencing certain parts), are directly associated with the numbering of 'pages' within digital texts. Designers should therefore be thinking of alternative ways of referencing specific areas of documents. This is not to say we should abolish the ubiquitous page from digital implementations, but rather, simply consider how these so-called pages hold-up against common digital functions such as zooming;

- The function of zooming, despite causing problems like the ones mentioned above, is also a key advantage that digital documents have over physical ones; reformatting the text of eDocuments, allows it to be made larger or smaller depending upon the visual capacity of the user. Unfortunately however, in most cases, zoom levels do not extend to menus which stay at a constant size, making additional functionality difficult for users with poor vision. Clearly, the solution to this problem is to strive for consistency in zooming by ensuring that all text displayed on screen adjusts in accordance with the current magnification applied to the device;
- The bookmarking functions on the majority of these devices leave much to be desired as they give no indication of where bookmarks exist within the book while it is currently being read. Instead, bookmarked pages are typically signified via some visual cue (e.g., a dog-eared corner or ribbon) on the bookmarked pages themselves, giving no overview of where others exist within the document. To view a list of all bookmarks, users will typically have to navigate away from the document they are reading to a list of entries. Although screen real-estate on these devices is in short supply, it would be highly beneficial if some indication of bookmark overview within a document could always be visible on screen while reading. One possible solution to this problem could be to take a similar approach to Byrd [21] and insert dots in the sidebar to signify the location of bookmarked pages within a document;
- Turning pages on these devices ranges in design from simple button clicks to sweeping gestures across touch-screen displays. Some designs then (usually eInk), have opted for a straightforward transition between one page and another, yet others have based their page turning function around the book metaphor; coding page ‘curling’ visualisations – usually in conjunction with gesture based controls. From my experience with these interactions however, it is clear that the majority have not completed the metaphor. For example, in most cases, the user can curl one page, but not grab multiple pages; it is also not possible to create dog-ears by curling pages in this manner. My recommendation here then is that page turning on digital documents need not be modelled on paper; if the designers choose to do so however, then it should behave *exactly* like paper to avoid the incompleteness associated with a half-implemented book metaphor.

## B.11 Related Publications

1. Jennifer Pearson, George Buchanan, Harold Thimbleby : **HCI Design Principles for eReaders** - *BooksOnline '10: 3rd Workshop on Research Advances in Large Digital Book Collections*, Toronto, Ontario, Canada.

# Bibliography

- [1] David Abrams, Ron Baecker, and Mark Chignell. Information Archiving with Bookmarks: Personal Web Space Construction and Organization. In *CHI '98: Proceedings of the 16th annual SIGCHI conference on Human Factors in Computing Systems*, pages 41–48, New York, NY, USA, 1998. ACM.
- [2] Annette Adler, Anuj Gujar, Beverly L. Harrison, Kenton O'Hara, and Abigail Sellen. A Diary Study of Work-Related Reading: Design Implications for Digital Reading Devices. In *CHI '98: Proceedings of the 16th annual SIGCHI conference on Human Factors in Computing Systems*, pages 241–248, New York, NY, USA, 1998. ACM.
- [3] M. J. Adler and C. Doren. *How to Read a Book: The Art of Getting a Liberal Education*. Simon and Schuster, 1972.
- [4] Maristella Agosti and Nicola Ferro. Annotations: Enriching a Digital Library. In *ECDL '03: Proceedings from the 7th European Conference on Research and Advanced Technology for Digital Libraries*, volume 2769 of *Lecture Notes in Computer Science*, pages 88–100. Springer Berlin / Heidelberg, 2003.
- [5] Maristella Agosti, Nicola Ferro, Ingo Frommholz, and Ulrich Thiel. Annotations in Digital Libraries and Collaboratories Facets, Models and Usage. In Rachel Heery and Liz Lyon, editors, *ECDL '04: Proceedings from the 8th European Conference on Research and Advanced Technology for Digital Libraries*, volume 3232 of *Lecture Notes in Computer Science*, pages 244–255. Springer Berlin / Heidelberg, 2004.
- [6] Maristella Agosti, Nicola Ferro, Emanuele Panizzi, and Rosa Trinchese. Annotation as a Support to User Interaction for Content Enhancement in Digital Libraries. In *AVI '06: Proceedings of the Working Conference on Advanced Visual Interfaces*, pages 151–154, New York, NY, USA, 2006. ACM.
- [7] Jason Alexander, Andy Cockburn, Stephen Fitchett, Carl Gutwin, and Saul Greenberg. Revisiting Read Wear: Analysis, Design, and Evaluation of a Footprints Scrollbar. In *CHI '09: Proceedings of the 27th annual SIGCHI conference on Human Factors in Computing Systems*, pages 1665–1674, New York, NY, USA, 2009. ACM.
- [8] Toshifumi Arai, Kimiyoshi Machii, and Soshiro Kusunuki. Retrieving Electronic Documents with Real-World Objects on InteractiveDESK. In *UIST '95: Proceedings of the 8th annual ACM Symposium on User interface and Software Technology*, pages 37–38, New York, NY, USA, 1995. ACM.

- [9] Toshifumi Arai, Kimiyoshi Machii, Soshiro Kuzunuki, and Hiroshi Shojima. InteractiveDESK: a Computer-Augmented Desk which Responds to Operations on Real Objects. In *CHI '95: Proceedings of the 13th annual SIGCHI conference on Human Factors in Computing Systems*, pages 141–142, New York, NY, USA, 1995. ACM.
- [10] Maribeth Back, Jonathan Cohen, Rich Gold, Steve Harrison, and Scott Minneman. Listen Reader: an Electronically Augmented Paper-Based Book. In *CHI '01: Proceedings of the 19th annual SIGCHI conference on Human Factors in Computing Systems*, pages 23–29, New York, NY, USA, 2001. ACM.
- [11] Ricardo Baeza-Yates and Berthier Ribeiro-Neto. *Modern Information Retrieval*. Addison-Wesley Publishing Company, USA, 2nd edition, 2008.
- [12] David Barger and Tomer Moscovich. Reflowing Digital Ink Annotations. In *CHI '03: Proceedings of the 21st annual SIGCHI conference on Human Factors in Computing Systems*, pages 385–393, New York, NY, USA, 2003. ACM.
- [13] Eric Bier, Lance Good, Kris Popat, and Alan Newberger. A Document Corpus Browser for in-depth Reading. In *JCDL '04: Proceedings of the 4th ACM/IEEE-CS Joint Conference on Digital libraries*, pages 87–96, New York, NY, USA, 2004. ACM.
- [14] Sven Birkerts. *The Gutenberg Elegies: The Fate of Reading in an Electronic Age*. Faber and Faber, Winchester, MA, USA, 1994.
- [15] A. J. Bernheim Brush, David Barger, Anoop Gupta, and J. J. Cadiz. Robust Annotation Positioning in Digital Documents. In *CHI '01: Proceedings of the 19th annual SIGCHI conference on Human Factors in Computing Systems*, pages 285–292, New York, NY, USA, 2001. ACM.
- [16] George Buchanan. Rapid document navigation for information triage support. In *Proceedings of the 7th ACM/IEEE-CS Joint Conference on Digital libraries*, JCDL '07, pages 503–503, New York, NY, USA, 2007. ACM.
- [17] George Buchanan and Fernando Loizides. Investigating Document Triage on Paper and Electronic Media. In *ECDL '07: Proceedings from the 11th European Conference on Research and Advanced Technology for Digital Libraries*, volume 4675 of *Lecture Notes in Computer Science*, pages 416–427. Springer Berlin / Heidelberg, 2007.
- [18] George Buchanan and Tom Owen. Improving navigation interaction in digital documents. In *Proceedings of the 8th ACM/IEEE-CS Joint Conference on Digital libraries*, JCDL '08, pages 389–392, New York, NY, USA, 2008. ACM.
- [19] George Buchanan and Jennifer Pearson. An Architecture for Supporting RFID-Enhanced Interactions in Digital Libraries. In *ECDL '10: Proceedings from the 14th European Conference on Research and Advanced Technology for Digital Libraries*, volume 6273 of *Lecture Notes in Computer Science*, pages 92–103. Springer Berlin / Heidelberg, 2010.
- [20] BusinessWeek. The office of the future. Found at: [http://www.businessweek.com/technology/content/may2008/tc20080526\\_547942.htm](http://www.businessweek.com/technology/content/may2008/tc20080526_547942.htm), 30 June 1975. Business Week.

- [21] Donald Byrd. A Scrollbar-Based Visualization for Document Navigation. In *DL '99: Proceedings of the 4th ACM International Conference on Digital libraries*, pages 122–129, New York, NY, USA, 1999. ACM.
- [22] J. J. Cadiz, Anop Gupta, and Jonathan Grudin. Using Web Annotations for Asynchronous Collaboration Around Documents. In *CSCW '00: Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work*, pages 309–318, New York, NY, USA, 2000. ACM.
- [23] Paul Cairns and Anna L. Cox. *Research Methods for Human-Computer Interaction*. Cambridge University Press, 2008.
- [24] Stuart K. Card, Lichan Hong, Jock D. Mackinlay, and Ed H. Chi. 3Book: A 3D Electronic Smart Book. In *AVI '04: Proceedings of the International Working Conference on Advanced Visual Interfaces*, pages 303–307. ACM Press, 2004.
- [25] Stuart K. Card, Lichan Hong, Jock D. Mackinlay, and Ed H. Chi. 3Book: A Scalable 3D Virtual Book. In *CHI EA '04: Extended Abstracts of the 22nd annual SIGCHI conference on Human Factors in Computing Systems*, pages 1095–1098, New York, NY, USA, 2004. ACM.
- [26] Jiajian Chen, Jun Xiao, Jian Fan, and Eamonn O'Brien-Strain. PageSpark: An e-Magazine Reader With Enhanced Reading Experiences on Handheld Devices. In *EICS '11: Proceedings of the 3rd ACM SIGCHI symposium on Engineering interactive computing systems*, pages 149–152, New York, NY, USA, 2011. ACM.
- [27] Nicholas Chen, Francois Guimbretiere, Morgan Dixon, Cassandra Lewis, and Maneesh Agrawala. Navigation Techniques for Dual-Display e-Book Readers. In *CHI '08: Proceedings of the 26th annual SIGCHI conference on Human Factors in Computing Systems*, pages 1779–1788, New York, NY, USA, 2008. ACM.
- [28] Ed H. Chi, Lichan Hong, Julie Heiser, and Stuart K. Card. eBooks with Indexes that Reorganize Conceptually. In *CHI EA '04: Extended Abstracts of the 22nd annual SIGCHI conference on Human factors in computing systems*, pages 1223–1226, New York, NY, USA, 2004. ACM.
- [29] Yi-Chun Chu, David Bainbridge, Matt Jones, and Ian H. Witten. Realistic Books: a Bizarre Homage to an Obsolete Medium? In *JCDL '04: Proceedings of the 4th ACM/IEEE-CS Joint Conference on Digital libraries*, pages 78–86, New York, NY, USA, 2004. ACM.
- [30] Yi-Chun Chu, Ian H. Witten, Richard Lobb, and David Bainbridge. How to Turn the Page. In *JCDL '03: Proceedings of the 3rd ACM/IEEE-CS Joint Conference on Digital libraries*, pages 186–188, Washington, DC, USA, 2003. IEEE Computer Society.
- [31] Andy Cockburn and Bruce McKenzie. What do Web Users do? An Empirical Analysis of Web Use. *International Journal of Human-Computer Studies*, 54(6):903–922, 2001.
- [32] Samuel T. Coleridge. *A Book I Value: Selected Marginalia*. Princeton University Press, Princeton, NJ, USA, 2003.
- [33] Fabio Crestani and Massimo Melucci. A Case Study of Automatic Authoring: From a Textbook to a Hyper-Textbook. *Data and Knowledge Engineering*, 27(1):1–30, 1998.
- [34] Mihaly Csikszentmihalyi. *Flow: The Psychology of Optimal Experience*. Harper and Row, 1990.

- [35] Mihaly Csikszentmihalyi. *Finding Flow: The Psychology of Engagement with Everyday Life*. Basic Books, 1997.
- [36] Andrew Dillon. Reading from Paper Versus Screens: a Critical Review of the Empirical Literature. *Ergonomics*, 35:1297–1326, October 1992.
- [37] Alan Dix. Designing for Appropriation. In *BCS-HCI '07: Proceedings of the 21st British HCI Group Annual Conference on People and Computers*, pages 27–30, Swinton, UK, UK, 2007. British Computer Society.
- [38] Marc Dymetman and Max Copperman. Intelligent Paper. In *EP '98/RIDT '98: Proceedings of the 7th International Conference on Electronic Publishing/4th International Conference on Raster Imaging and Digital Typography: Electronic Publishing, Artistic Imaging, and Digital Typography*, pages 392–406, London, UK, 1998. Springer-Verlag.
- [39] The Economist. Technological Comebacks: Not Dead, Just Resting. Found at: <http://www.economist.com/node/12381449>, 9 October 2008. The Economist.
- [40] Michael W. Eysenck and Mark T. Keane. *Cognitive psychology: a student's handbook*. Psychology Press, 2005.
- [41] Anne Fadiman. *Ex Libris: Confessions of a Common Reader*. Farrar, Straus and Giroux, 1998.
- [42] Andy Field and Graham Hole. *How to Design and Report Experiments*. Sage Publications Ltd, 2003.
- [43] William H. Gass. In Defense of the Book: On the Enduring Pleasures of Paper, Type, Page, and Ink. Found at: <http://www.harper.org/archive/1999/11/0060708>, November 1999. Harper's Magazine.
- [44] William W. Gaver. Technology Affordances. In *CHI '91: Proceedings of the 9th annual SIGCHI conference on Human Factors in Computing Systems*, pages 79–84, New York, NY, USA, 1991. ACM.
- [45] J. J. Gibson. *The Theory of Affordances*. Lawrence Erlbaum, 1977.
- [46] James J. Gibson. *The Ecological Approach To Visual Perception*. Lawrence Erlbaum Associates, September 1979.
- [47] Malcolm Gladwell. The Social Life of Paper. Found at: [http://www.newyorker.com/archive/2002/03/25/020325crbo\\_books](http://www.newyorker.com/archive/2002/03/25/020325crbo_books), 25 March 2002. The New Yorker.
- [48] Gene Golovchinsky and Laurent Denoue. Moving Markup: Repositioning Freeform Annotations. In *UIST '02: Proceedings of the 15th annual SIGCHI conference on Human Factors in Computing Systems*, pages 21–30, New York, NY, USA, 2002. ACM.
- [49] Jeff Gomez. *Print is Dead: Books in our Digital Age*. Palgrave Macmillan, 2009.
- [50] Thomas Gottron. Document Word Clouds: Visualising Web Documents as Tag Clouds to Aid Users in Relevance Decisions. In *ECDL '09: Proceedings from the 13th European Conference on Research and Advanced Technology for Digital Libraries*, volume 5714 of *Lecture Notes in Computer Science*, pages 94–105. Springer Berlin / Heidelberg, 2009.

- [51] Saul Greenberg and Ian H. Witten. Supporting Command Reuse: Empirical Foundations and Principles. *Int. J. Man-Mach. Stud.*, 39:353–390, September 1993.
- [52] Wilfred J. Hansen and Christina Haas. Reading and Writing with Computers: a Framework for Explaining Differences in Performance. *Communications of the ACM*, 31(9):1080–1089, 1988.
- [53] David J. Harper, Sara Coulthard, and Sun Yixing. A Language Modelling Approach to Relevance Profiling for Document Browsing. In *JCDL '02: Proceedings of the 2nd ACM/IEEE-CS Joint Conference on Digital Libraries*, pages 76–83, New York, NY, USA, 2002. ACM.
- [54] David J. Harper, Ivan Koychev, and Yixing Sun. Query-Based Document Skimming: A User-Centred Evaluation of Relevance Profiling. In *Advances in Information Retrieval*, volume 2633 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, 2003.
- [55] Marti A. Hearst. TileBars: Visualization of Term Distribution Information in Full Text Information Access. In *CHI '95: Proceedings of the 13th annual SIGCHI conference on Human Factors in Computing Systems*, pages 59–66, New York, NY, USA, 1995. ACM Press/Addison-Wesley Publishing Co.
- [56] Martin Heidegger. *Being and Time (Translated From the German Sein und Zeit)*. Blackwell Publishers, 1962.
- [57] Steven Heim. *The Resonant Interface: HCI Foundations for Interaction Design*. Pearson/Addison Wesley, 2007.
- [58] William C. Hill, James D. Hollan, Dave Wroblewski, and Tim McCandless. Edit Wear and Read Wear. In *CHI '92: Proceedings of the 10th annual SIGCHI conference on Human Factors in Computing Systems*, pages 3–9, New York, NY, USA, 1992. ACM.
- [59] Ken Hinckley, Morgan Dixon, Raman Sarin, Francois Guimbretiere, and Ravin Balakrishnan. Codex: a Dual Screen Tablet Computer. In *CHI '09: Proceedings of the 27th annual SIGCHI conference on Human Factors in Computing Systems*, pages 1933–1942, New York, NY, USA, 2009. ACM.
- [60] Orland Hoeber and Xue Dong Yang. A Comparative User Study of Web Search Interfaces: HotMap, Concept Highlighter, and Google. In *WI '06: Proceedings of the 2006 IEEE/WIC/ACM International Conference on Web Intelligence*, pages 866–874, Washington, DC, USA, 2006. IEEE Computer Society.
- [61] Orland Hoeber and Xue Dong Yang. The Visual Exploration of Web Search Results Using HotMap. In *IV '06: Proceedings of the Conference on Information Visualization*, pages 157–165, Washington, DC, USA, 2006. IEEE Computer Society.
- [62] Lichan Hong, Ed H. Chi, and Stuart K. Card. Annotating 3D Electronic Books. In *CHI EA '05: Extended Abstracts of the 23rd annual SIGCHI conference on Human Factors in Computing Systems*, pages 1463–1466, New York, NY, USA, 2005. ACM.
- [63] H. J. Jackson. *Marginalia: Readers Writing in Books*. Yale University Press, New Haven, CT, USA, 2002.

- [64] Walter Johnson, Herbert Jellinek, Leigh Klotz, Jr., Ramana Rao, and Stuart K. Card. Bridging the Paper and Electronic Worlds: the Paper User Interface. In *CHI '93: Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems*, pages 507–512, New York, NY, USA, 1993. ACM.
- [65] Ricardo Kawase, Eelco Herder, and Wolfgang Nejdl. A Comparison of Paper-Based and Online Annotations in the Workplace. In *EC-TEL '09: Proceedings of the 4th European Conference on Technology Enhanced Learning: Learning in the Synergy of Multiple Disciplines*, pages 240–253, Berlin, Heidelberg, 2009. Springer-Verlag.
- [66] Richard M. Keller, Shawn R. Wolfe, James R. Chen, Joshua L. Rabinowitz, and Nathalie Mathe. A Bookmarking Service for Organizing and Sharing URLs. *Computer Networks and ISDN Systems*, 29(8-13):1103–1114, 1997.
- [67] Hideki Koike and Motoki Kobayashi. EnhancedDesk: Integrating Paper Documents and Digital Documents. In *APCHI*, pages 57–62. IEEE Computer Society, 1998.
- [68] Hideki Koike, Yoichi Sato, and Yoshinori Kobayashi. Integrating Paper and Digital Information on EnhancedDesk: A Method for Realtime Finger Tracking on an Augmented Desk System. *ACM Transactions on Computer-Human Interaction*, 8(4):307–322, 2001.
- [69] Richard S. Krunk and Paul Muter. Reading of Continuous Text on Video Screens. *Human Factors*, 26:339–345, 1984.
- [70] Byron Lahey, Audrey Girouard, Winslow Bursleson, and Roel Vertegaal. PaperPhone: Understanding the Use of Bend Gestures in Mobile Devices with Flexible Electronic Paper Displays. In *CHI '11: Proceedings of the 2011 annual SIGCHI conference on Human Factors in Computing Systems*, pages 1303–1312, New York, NY, USA, 2011. ACM.
- [71] Monica Landoni and Forbes Gibb. The Role of Visual Rhetoric in the Design and Production of Electronic Books: The Visual Book. *The Electronic Library*, 18(3):190–201, 2000.
- [72] Chunyuan Liao, François Guimbretière, and Ken Hinckley. PapierCraft: a Command System for Interactive Paper. In *UIST '05: Proceedings of the 18th annual ACM Symposium on User Interface Software and Technology*, pages 241–244, New York, NY, USA, 2005. ACM.
- [73] Chunyuan Liao, François Guimbretière, Ken Hinckley, and Jim Hollan. Papiercraft: A Gesture-Based Command System for Interactive Paper. *ACM Transactions Computer-Human Interaction*, 14(4):1–27, 2008.
- [74] British Library. Turning the pages. Found at: <http://www.bl.uk/collections/treasures/digitisation1.html>.
- [75] Veronica Liesaputra and Witten H. Ian. Seeking Information in Realistic Books: a User Study. In *JCDL '08: Proceedings of the 8th ACM/IEEE-CS Joint Conference on Digital libraries*, pages 29–38, New York, NY, USA, 2008. ACM.
- [76] Veronica Liesaputra, Ian H. Witten, and David Bainbridge. Lightweight realistic books: The greenstone connection. In *JCDL '07: Proceedings of the 7th ACM/IEEE-CS Joint Conference on Digital libraries*, pages 502–502, New York, NY, USA, 2007. ACM.



- [77] Veronica Liesaputra, Ian H. Witten, and David Bainbridge. Creating and Reading Realistic Electronic Books. *Computer*, 42(2):72–81, 2009.
- [78] Veronica Liesaputra, Ian H. Witten, and David Bainbridge. Searching in a Book. In *ECDL '09: Proceedings from the 13th European Conference on Research and Advanced Technology for Digital Libraries*, volume 5714 of *Lecture Notes in Computer Science*, pages 442–446. Springer Berlin / Heidelberg, 2009.
- [79] Paul Luff, Guy Adams, Wolfgang Bock, Adam Drazin, David Frohlich, Christian Heath, Peter Herdman, Heather King, Nadja Linketscher, Rachel Murphy, Moira Norrie, Abigail Sellen, Beat Signer, Ella Tallyn, and Emil Zeller. The Disappearing Computer. *Augmented paper: Developing Relationships Between Digital Content and Paper*, pages 275–297, 2007.
- [80] Thomas W. Malone. How do People Organize Their Desks?: Implications for the Design of Office Information Systems. *ACM Trans. Inf. Syst.*, 1:99–112, January 1983.
- [81] Catherine C. Marshall. Annotation: From Paper Books to the Digital Library. In *DL '97: Proceedings of the 2nd ACM International Conference on Digital libraries*, pages 131–140, New York, NY, USA, 1997. ACM.
- [82] Catherine C. Marshall. Toward an Ecology of Hypertext Annotation. In *HYPertext '98: Proceedings of the 9TH ACM conference on Hypertext and Hypermedia: Links, Objects, Time and Space—Structure in Hypermedia Systems*, pages 40–49, New York, NY, USA, 1998. ACM.
- [83] Catherine C. Marshall. Collection-Level Analysis Tools for Books Online. In *BooksOnline '08: Proceeding of the 2008 ACM workshop on Research Advances in Large Digital Book Repositories*, pages 41–44, New York, NY, USA, 2008. ACM.
- [84] Catherine C. Marshall. *Reading and Writing the Electronic Book*. Synthesis Lectures on Information Concepts, Retrieval, and Services. Morgan & Claypool Publishers, 2009.
- [85] Catherine C. Marshall and Sara Bly. Turning the Page on Navigation. In *JCDL '05: Proceedings of the 5th ACM/IEEE-CS Joint Conference on Digital libraries*, pages 225–234, New York, NY, USA, 2005. ACM.
- [86] Catherine C. Marshall, Morgan N. Price, Gene Golovchinsky, and Bill N. Schilit. Introducing a Digital Library Reading Appliance into a Reading Group. In *DL '99: Proceedings of the 4TH ACM Conference on Digital libraries*, pages 77–84, New York, NY, USA, 1999. ACM.
- [87] Catherine C. Marshall and Frank M. Shipman, III. Spatial Hypertext: Designing for Change. *Commun. ACM*, 38(8):88–97, 1995.
- [88] Daniel K. Mayes, Valerie K. Sims, and Jefferson M. Koonce. Comprehension and Workload Differences for VDT and Paper-Based Reading. *International Journal of Industrial Ergonomics*, 28:367–378, April 2001.
- [89] Josh McHugh. LeapFrog's Wild Ride. Found at: <http://www.wired.com/wired/archive/13.11/leapfrog.html>, November 2005. Wired Magazine.
- [90] Carol Bergfeld Mills and Linda J. Weldon. Reading Text from Computer Screens. *ACM Comput. Surv.*, 19:329–357, December 1987.

- [91] Pranav Mistry. QUICKIES: the Future of Sticky Notes. In *CHI EA '09: Extended Abstracts of the 27th annual SIGCHI conference on Human Factors in Computing Systems*, pages 3529–3530, New York, NY, USA, 2009. ACM.
- [92] Pranav Mistry and Pattie Maes. Intelligent Sticky Notes that can be Searched, Located and can Send Reminders and Messages. In *IUI '08: Proceedings of the 13th International Conference on Intelligent user Interfaces*, pages 425–426, New York, NY, USA, 2008. ACM.
- [93] Paul Muter, Susane A. Latremouille, William C. Treurniet, and Paul Beam. Extended Reading of Continuous Text on Television Screens. *Human Factors*, 24:501–508, 1982.
- [94] William Newman and Pierre Wellner. A Desk Supporting Computer-Based Interaction with Paper Documents. In *CHI '92: Proceedings of the 10th annual SIGCHI conference on Human Factors in Computing Systems*, pages 587–592, New York, NY, USA, 1992. ACM.
- [95] BBC News. Amazon Kindle e-Book Downloads Outsell Paperbacks. Found at: <http://www.bbc.co.uk/news/business-12305015>, 28th January 2011.
- [96] BBC News. Booker Prize Judges Sent e-Readers. Found at: <http://www.bbc.co.uk/news/entertainment-arts-12306173>, 28th January 2011.
- [97] BBC News. UK General Title Digital Book Sales Soar to 16m in 2010. Found at: <http://www.bbc.co.uk/news/technology-13262842>, 3rd May 2011.
- [98] BBC News. Flexible Phone Made from Electronic Paper to Debut. Found at: <http://www.bbc.co.uk/news/technology-13308452>, 6th May 2011.
- [99] Jakob Nielsen. The Art of Navigating Through Hypertext. *Communications of the ACM*, 33:296–310, March 1990.
- [100] Jakob Nielsen. Electronic Books: A Bad Idea. Found at: <http://www.useit.com/alertbox/980726.html>, 26th July 1998.
- [101] Jakob Nielsen and Rolf Molich. Heuristic Evaluation of User Interfaces. In *CHI '90: Proceedings of the 8th annual SIGCHI conference on Human Factors in Computing Systems*, pages 249–256, New York, NY, USA, 1990. ACM.
- [102] Donald Norman. *The Design of Everyday Things*. Basic Books, 1988.
- [103] Donald A. Norman. *The Invisible Computer*. MIT Press, Cambridge, MA, USA, 1998.
- [104] The Trustees of Princeton University. The e-Reader Pilot at Princeton. Found at: <http://www.princeton.edu/ereaderpilot/eReaderFinalReportLong.pdf>, 2009.
- [105] The Trustees of Princeton University. Why we are Doing an e-Reader Pilot. Found at: <http://www.princeton.edu/ereaderpilot/rationale>, Taken on: 18th April 2011.
- [106] Kenton O'Hara. Towards a Typology of Reading Goals. In *Affordances of Paper Project Technical Report*, Rank Xerox Research Centre, Cambridge, UK, 1996.
- [107] Kenton O'Hara and Abigail Sellen. A Comparison of Reading Paper and on-line Documents. In *CHI '97: Proceedings of the 15th annual SIGCHI conference on Human Factors in Computing Systems*, pages 335–342, New York, NY, USA, 1997. ACM.

- [108] Kenton O'Hara, Fiona Smith, William Newman, and Abigail Sellen. Student Readers' Use of Library Documents: Implications for Library Technologies. In *CHI '98: Proceedings of the 16th annual SIGCHI conference on Human Factors in Computing Systems*, pages 233–240, New York, NY, USA, 1998. ACM Press/Addison-Wesley Publishing Co.
- [109] Steven Pace. A grounded theory of the flow experiences of web users. *International Journal of Human Computer Studies*, 60:327–363, March 2004.
- [110] Beryl Plimmer and Paul Mason. A Pen-Based Paperless Environment for Annotating and Marking Student Assignments. In *AUIC '06: Proceedings of the 7th Australasian User Interface Conference*, pages 37–44, Darlinghurst, Australia, Australia, 2006. Australian Computer Society, Inc.
- [111] Morgan N. Price, Bill N. Schilit, and Gene Golovchinsky. XLibris: The Active Reading Machine. In *CHI '98: Proceedings of the 16th annual SIGCHI conference on Human Factors in Computing Systems*, pages 22–23, New York, NY, USA, 1998. ACM.
- [112] Daniela K. Rosner, Lora Oehlberg, and Kimiko Ryokai. Studying Paper Use to Inform the Design of Personal and Portable Technology. In *CHI EA '08: Extended Abstracts of the 26th annual SIGCHI conference on Human Factors in Computing Systems*, pages 3405–3410, New York, NY, USA, 2008. ACM.
- [113] Christine Nelsen Ryan and Sandra Henselmeier. Usability testing at Macmillan. *Key Words Vol. 8*, pages 198–202, 2000.
- [114] Bill N. Schilit, Gene Golovchinsky, and Morgan N. Price. Beyond Paper: Supporting Active Reading with Free Form Digital Ink Annotations. In *CHI '98: Proceedings of the 16th annual SIGCHI conference on Human Factors in Computing Systems*, pages 249–256, New York, NY, USA, 1998. ACM Press/Addison-Wesley Publishing Co.
- [115] Bill N. Schilit, Morgan N. Price, and Gene Golovchinsky. Digital Library Information Appliances. In *DL '98: Proceedings of the 3rd ACM Conference on Digital Libraries*, pages 217–226, New York, NY, USA, 1998. ACM.
- [116] Susan Schreibman, Ray Siemens, and John Unsworth. *A Companion to Digital Humanities*. Wiley Publishing, 2008.
- [117] Abigail Sellen and Richard Harper. Paper as an Analytic Resource for the Design of New Technologies. In *CHI '97: Proceedings of the 15th annual SIGCHI conference on Human Factors in Computing Systems*, pages 319–326, New York, NY, USA, 1997. ACM.
- [118] Abigail J. Sellen and Richard H.R. Harper. *The Myth of the Paperless Office*. MIT Press, Cambridge, MA, USA, 2003.
- [119] Helen Sharp, Yvonne Rogers, and Jenny Preece. *Interaction Design: Beyond Human-Computer Interaction*. John Wiley & Sons Ltd, 3 edition, 2011.
- [120] Frank M. Shipman, Haowei Hsieh, J. Michael Moore, and Anna Zacchi. Supporting personal collections across digital libraries in spatial hypertext. In *Proceedings of the 4th ACM/IEEE-CS joint conference on Digital libraries, JCDL '04*, pages 358–367, New York, NY, USA, 2004. ACM.

- [121] Frank M. Shipman, Haowei Hsieh, J. Michael Moore, and Anna Zacchi. Supporting Personal Collections Across Digital Libraries in Spatial Hypertext. In *JCDL '04: Proceedings of the 4th ACM/IEEE-CS Joint Conference on Digital Libraries*, pages 358–367, New York, NY, USA, 2004. ACM.
- [122] Frank M. Shipman, III, Catherine C. Marshall, and Thomas P. Moran. Finding and Using Implicit Structure in Human-Organized Spatial Layouts of Information. In *CHI '95: Proceedings of the 13th annual SIGCHI conference on Human Factors in Computing Systems*, pages 346–353, New York, NY, USA, 1995. ACM Press/Addison-Wesley Publishing Co.
- [123] Ben Shneiderman, Donald Byrd, and William Bruce Croft. Sorting Out Searching: a User-Interface Framework for Ttext Searches. *Commun. ACM*, 41(4):95–98, 1998.
- [124] Ben. Shneiderman, Donald Byrd, and William Bruce Croft. Clarifying Search: A User-Interface Framework for Text Searches. *D-Lib Magazine*, January 1997.
- [125] Ben Shneiderman and Catherine Plaisant. *Designing the User Interface: Strategies for Effective Human-Computer Interaction (4th Edition)*. Pearson Addison Wesley, 2004.
- [126] Jürgen Steimle, Oliver Brdiczka, and Max Mühlhäuser. Digital Paper Bookmarks: Collaborative Structuring, Indexing and Tagging of Paper Documents. In *CHI EA '08: Extended Abstracts of the 26th annual SIGCHI conference on Human Factors in Computing Systems*, pages 2895–2900, New York, NY, USA, 2008. ACM.
- [127] W. Strunk Jr and E.B. White. *The Elements of Style*. Pearson Education Company, 1918.
- [128] Elaine Svenonius. *The Intellectual Foundation of Information Organization*. MIT Press, 2000.
- [129] Aurélien Tabard, Wendy Mackay, Nicolas Roussel, and Catherine Letondal. PageLinker: Integrating Contextual Bookmarks Within a Browser. In *CHI '07: Proceedings of the 25th annual SIGCHI conference on Human Factors in Computing Systems*, pages 337–346, New York, NY, USA, 2007. ACM.
- [130] Craig Tashman. LiquidText: Active Reading Through Multitouch Document Manipulation. In *CHI EA '10: Extended Abstracts of the 28th annual SIGCHI conference on Human Factors in Computing Systems*, pages 2959–2962, New York, NY, USA, 2010. ACM.
- [131] Craig S. Tashman and W. Keith Edwards. Active Reading and its Discontents: The Situations, Problems and Ideas of Readers. In *CHI '11: Proceedings of the 2011 annual SIGCHI conference on Human Factors in Computing Systems*, pages 2927–2936, New York, NY, USA, 2011. ACM.
- [132] Craig S. Tashman and W. Keith Edwards. LiquidText: A Flexible, Multitouch Environment to Support Active Reading. In *CHI '11: Proceedings of the 2011 annual SIGCHI conference on Human Factors in Computing Systems*, pages 3285–3294, New York, NY, USA, 2011. ACM.
- [133] Linda Tauscher and Saul Greenberg. How People Revisit Web Pages: Empirical Findings and Implications for the Design of History Systems. *International Journal on Human-Computer Studies*, 47(1):97–137, 1997.

- [134] The New York Times. E-Books Top Hardcover at Amazon. Found at: <http://www.nytimes.com/2010/07/20/technology/20kindle.html>, 19 July 2010. The New York Times.
- [135] E. G. Toms and H. L. O'Brien. Understanding the Information and Communication Technology Needs of the e-Humanist. *Journal of Documentation*, 64(1):102–130, 2008.
- [136] Andries van Dam. Hypertext '87 Keynote Address. *Communications of the ACM*, 31(7):887–895, 1988.
- [137] Phil Wahba. e-Readers Gain Traction, Spur Sales: Poll. Found at: <http://www.reuters.com/article/2010/09/22/us-books-ereaders-poll-idUSTRE68L3RG20100922>, 22nd September 2010. Reuters.
- [138] Colin Ware. *Information Visualization: Perception for Design*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2004.
- [139] Erik Wästlund, Torsten Norlander, and Trevor Archer. The Effect of Page Layout on Mental Workload: A Dual-Task Experiment. *Computers in Human Behavior*, 24:1229–1245, May 2008.
- [140] Erik Wästlund, Henrik Reinikka, Torsten Norlander, and Trevor Archer. Effects of VDT and Paper Presentation on Consumption and Production of Information: Psychological and Physiological Factors. *Computers in Human Behavior*, 21(2):377–394, 2005.
- [141] Jun-ichiro Watanabe, Arito Mochizuki, and Youichi Horry. Booksheet: Bendable Device for Browsing Content Using the Metaphor of Leafing Through the Pages. In *UbiComp '08: Proceedings of the 10th International Conference on Ubiquitous Computing*, pages 360–369, New York, NY, USA, 2008. ACM.
- [142] Harald Weinreich, Hartmut Obendorf, Eelco Herder, and Matthias Mayer. Off the Beaten Tracks: Exploring Three Aspects of Web Navigation. In *WWW '06: Proceedings of the 15th International Conference on World Wide Web*, pages 133–142, New York, NY, USA, 2006. ACM.
- [143] Mark Weiser. Some Computer Science Issues in Ubiquitous Computing. *Communications of the ACM*, 36(7):75–84, 1993.
- [144] Pierre Wellner. The DigitalDesk Calculator: Tangible Manipulation on a Desk Top Display. In *UIST '91: Proceedings of the 4th annual ACM Symposium on User Interface Software and Technology*, pages 27–33, New York, NY, USA, 1991. ACM.
- [145] Pierre Wellner. Interacting with Paper on the DigitalDesk. *Communications of the ACM*, 36(7):87–96, 1993.
- [146] Doug Wightman, Tim Ginn, and Roel Vertegaal. TouchMark: Flexible Document Navigation and Bookmarking Techniques for e-Book Readers. In *Proceedings of Graphics Interface 2010, GI '10*, pages 241–244, Toronto, ON, Canada, 2010. Canadian Information Processing Society.

- [147] Ruth Wilson and Monica Landoni. Evaluating the Usability of Portable Electronic Books. In *Proceedings of the 2003 ACM Symposium on Applied Computing, SAC '03*, pages 564–568, New York, NY, USA, 2003. ACM.
- [148] Terry Winograd and Fernando Flores. *Understanding computers and cognition*. Ablex Publishing Corp., Norwood, NJ, USA, 1985.
- [149] Hazel Woodward, Fytton Rowland, Cliff McKnight, Carolyn Pritchett, and Jack Meadows. Cafe Jus: an Electronic Journals User Survey. *Journal of Digital Information*, 1, 1998.
- [150] Chih-Sung Andy Wu, Susan J Robinson, and Alexandra Mazalek. WikiTUI: Leaving Digital Traces in Physical Books. In *ACE '07: Proceedings of the International Conference on Advances in Computer Entertainment Technology*, pages 264–265, New York, NY, USA, 2007. ACM.
- [151] Chih-Sung (Andy) Wu, Susan J. Robinson, and Ali Mazalek. Turning a Page on the Digital Annotation of Physical Books. In *TEI '08: Proceedings of the 2nd International Conference on Tangible and Embedded Interaction*, pages 109–116, New York, NY, USA, 2008. ACM.
- [152] Dongwook Yoon, Yongjun Cho, Kiwon Yeom, and Ji-Hyung Park. Touch-Bookmark: a Lightweight Navigation and Bookmarking Technique for e-Books. In *CHI EA '11: Extended Abstracts of the 29th annual SIGCHI conference on Human Factors in Computing Systems*, pages 1189–1194, New York, NY, USA, 2011. ACM.

# Index

- 3Book, **26**, 85, 130
- Acrobat, 31
- Active Reading, 2, 3, 5, **7–8**, 9–11, 13, 14, 21–24, 26, 28–30, 53, 58, 59, 61, 78, 79, 81–83, 89, 104, 107, 121, 129–132, 134, 136, 137, 139, 140, 142, 143, 146–148, 157, 158
- Adobe  
Acrobat, 31, 34, 44, 56, 91, 158
- Adobe Acrobat, 34, 44, 56, 91, 158
- Affordance, 2, 3, 9–12, **15–16**, 17, 21–25, 27–29, 33, 36, 129, 130
- Affordances of Paper Project, **21**
- Agrawala, Maneesh  
Codex, **24**, 130
- Alexander, Jason  
Footprints Scrollbar, 37, 41, 114, 116
- Amazon  
Kindle, 1, 2, **168–170**
- Apple  
iPad, 1, 26, 34, 147–149, **173–176**  
iPhone, 34  
Preview, 31, 34, 56, 57, 91, 98, 110, 158
- Arai, Toshifumi  
InteractiveDESK, 27
- Back, Meribeth  
Listen Reader, **28–29**, 130
- Bainbridge, David  
Realistic Electronic Books, **25–26**, 37, 41, 85, 112, 130
- Balakrishnan, Ravin  
Codex, **24**, 130
- Barnes and Noble  
nook, **170–173**
- Bly, Sara  
Lightweight, 1, 3–5, **11–12**, 12, 13, 16, 18, 21, 23–27, 29, 31, 36, 41, 46, 49–51, 53, 66, 67, 73, 78–82, 84, 86, 87, 90, 102–104, 107, 111, 113, 115, 119, 121, 126–127, 130–143, 145–149, 158, 179, 181
- Book Metaphor, *see* Metaphor
- Booksheet, 24, 38
- Bookmarks, 166, 169
- Brdiczka, Oliver  
Digital Paper Bookmarks, 38, 85
- British Library  
Turning The Pages, **25**
- Brown, Les, 107
- Burleson, Winslow  
PaperPhone, 24
- Byrd, Donald  
Four-Phase Framework for Text Searches, **108**  
Scrollbar-based Visualization, 40, 105, **111–112**
- Card, Stuart K.  
3Book, **26**, 85, 130  
ScentIndex, **113**
- Carroll, Lewis, 7
- Chen, Jiajian  
PageSpark, **23**
- Chen, Nicholas  
Codex, **24**, 130
- Chi, Ed H.  
3Book, **26**, 85, 130  
ScentIndex, **113**
- Chu, Yi-Chun  
Realistic Electronic Books, **25–26**, 37, 41, 85, 130
- Cockburn, Andy  
Footprints Scrollbar, 37, 41, 114, 116  
Codex, **24**, 130
- Cohen, Jonathan  
Listen Reader, **28–29**, 130
- Copperman, Max  
Intelligent Paper, **28**

- Coulthard, Sara  
SmartSkim, 105, **111**
- Croft, William Bruce  
Four-Phase Framework for Text Searches, **108**
- Csikszentmihalyi, Mihaly  
Flow, 14, **14–15**, 30, 38, 83, 84, 104, 130–131, 134, 137
- Ctrl+f, 110
- Digital Paper Bookmarks, 38, 85
- DigitalDesk, **27–28**, 130
- Direct Manipulation, 81, 83, 87, 89, 94, 103, 104
- Disney, Walt, 157
- Distraction, **14–15**
- Dixon, Morgan  
Codex, **24**, 130
- Dymetman, Marc  
Intelligent Paper, **28**
- Edwards, W. Keith  
LiquidText, **23**, 130
- Einstein, Albert, 129
- EnhancedDesk, **27–28**
- eReader, 1, 2, 6, 9, 34–36, 39, 49, 50, 129, 140, **157–182**
- Fan, Jian  
PageSpark, **23**
- Firefox, 34, 44
- Fitchett, Stephen  
Footprints Scrollbar, 37, 41, 114, 116
- Flow, 14, **14–15**, 30, 38, 83, 84, 104, 130–131, 134, 137
- Footprints Scrollbar, 37, 41, 114, 116
- Four-Phase Framework for Text Searches, **108**
- Ginn, Tim  
TouchMark, 24
- Girouard, Audrey  
PaperPhone, 24
- Gold, Rich  
Listen Reader, **28–29**, 130
- Golovchinsky, Gene  
XLibris, **22–23**, 58–59, 70, 112, 130
- Greenberg, Saul  
Footprints Scrollbar, 37, 41, 114, 116
- Guimbretière, François  
Codex, **24**, 130
- PapierCraft, **28**, 130
- Gutwin, Carl  
Footprints Scrollbar, 37, 41, 114, 116
- Harper, David, J.  
ProfileSkim, **111**  
SmartSkim, 105, **111**
- Harper, Richard  
Paperless Office, 1, 2
- Harrison, Steve  
Listen Reader, **28–29**, 130
- Hearst, Marti, A.  
TileBars, **112**, 115
- Heidegger, Martin  
Present-at-hand, **13**, 130  
Ready-to-hand, **13**, 13, 16, 21, 23, 29, 130, 137, 160
- Heiser, Julie  
ScentIndex, **113**
- Herlihy, James Leo, 81
- Hinckley, Ken  
Codex, **24**, 130  
PapierCraft, **28**, 130
- Hoeber, Orland  
HotMaps, **112**
- Hollan, Jim  
PapierCraft, **28**, 130
- Hong, Lichan  
3Book, **26**, 85, 130  
ScentIndex, **113**
- Horry, Youichi  
Booksheet, 24, 38
- HotMaps, **112**
- Hyper-link, 19, 22, 119, **107–127**, 134, 135, 139, 141
- Intelligent Paper, **28**
- InteractiveDESK, 27
- Invisible Computers, 13, **13–14**, 14, 23, 29, 160
- iPad, 1, 26, 34, 147–149, **173–176**
- iPhone, 34
- Jeffries, Lionel, 53
- Jones, Matt  
Realistic Electronic Books, **25–26**, 37, 41, 85, 130
- Kindle, 1, 2, **168–170**



- Kitt, Eartha, 153
- Kobayashi, Motoki  
EnhancedDesk, **27–28**
- Kobayashi, Yoshinori  
EnhancedDesk, **27–28**
- Koike, Hideki  
EnhancedDesk, **27–28**
- Koychev, Ivan  
ProfileSkim, **111**
- Kuzunuki, Soshiro  
InteractiveDESK, 27
- Lahey, Byron  
PaperPhone, 24
- LeapFrog  
LeapPad, **28**
- LeapFrog Enterprises  
LeapPad, 28  
Tag Reading System, 28
- LeapPad, **28, 28**
- Lewis, Cassandra  
Codex, **24**, 130
- Liao, Chunyuan  
PapierCraft, **28**, 130
- Liesaputra, Veronica  
Realistic Electronic Books, **25–26**, 37, 41, 85, 112, 130
- Lightweight, 1, 3–5, **11–12**, 12, 13, 16, 18, 21, 23–27, 29, 31, 36, 41, 46, 49–51, 53, 66, 67, 73, 78–82, 84, 86, 87, 90, 102–104, 107, 111, 113, 115, 119, 121, 126–127, 130–143, 145–149, 158, 179, 181
- LiquidText, **23**, 70, 130
- Listen Reader, **28–29**, 130
- Mühlhäuser, Max  
Digital Paper Bookmarks, 85
- Machii, Kimiyoshi  
InteractiveDESK, 27
- Mackinlay, Jock D.  
3Book, **26**, 85, 130
- Maes, Pattie  
Quickies, 86
- Man Booker Prize, 1
- Marshall, Catherine, C.  
Lightweight, 1, 3–5, **11–12**, 12, 13, 16, 18, 21, 23–27, 29, 31, 36, 41, 46, 49–51, 53, 66, 67, 73, 78–82, 84, 86, 87, 90, 102–104, 107, 111, 113, 115, 119, 121, 126–127, 130–143, 145–149, 158, 179, 181
- Mazalek, Alexandra  
WikiTUI, **28**
- Metaphor, 3, 6, **17–18**, 19, 22–27, 29, 37, 49, 59, 89, 102, 103, 130, **139–140**, **160**, 175, 177, 181, 182
- Minneman, Scott  
Listen Reader, **28–29**, 130
- Mistry, Pranav  
Quickies, 86
- Mochizuki, Arito  
Booksheet, 24, 38
- Newman, William  
DigitalDesk, **27–28**, 130
- nook, **170**, **173**
- Norman, Donald  
Affordance, **15–16**  
Invisible Computers, 13, **13–14**, 14, 23, 29, 160
- O'Brien-Strain, Eamonn  
PageSpark, **23**
- O'Hara, Kenton  
Affordances of Paper Project, **21**
- PageSpark, **23**
- Pagination, 19, **19**
- Paperless Office, 1, 2
- PaperPhone, 24
- PapierCraft, **28**, 130
- Plaisant, Catherine  
Direct Manipulation, 81, 83, 87, 89, 94, 103, 104
- Post-it, 32, 48, 55, 62, 66, 68, 79, **81–105**, 133, 134, 137, 138
- Present-at-hand, **13**, 130
- Preview, 31, 34, 56, 57, 91, 98, 110, 158
- ProfileSkim, **111**
- Quickies, 86
- Ready-to-hand, **13**, 13, 16, 21, 23, 29, 130, 137, 160
- Realistic Electronic Books, **25–26**, 37, 41, 85, 112, 130

- Robinson, Susan J.  
WikiTUI, **28**
- Samsung  
Galaxy Tab, **176–177**
- Sarin, Raman  
Codex, **24**, 130
- Sato, Yoichi  
EnhancedDesk, **27–28**
- ScentIndex, **113**
- Schilit, Bill  
XLibris, **22–23**, 58–59, 70, 112, 130
- Sellen, Abigail  
Affordances of Paper Project, **21**  
Paperless, 1, 2
- Seuss, Dr, 1, 145
- Shneiderman, Ben  
Direct Manipulation, 81, 83, 87, 89, 94, 103,  
104  
Four-Phase Framework for Text Searches, **108**
- SmartSkim, 105, **111**
- Sony Reader  
PRS-300, **160–163**  
PRS-505, **163–165**  
PRS-600, **165–168**
- Steimle, Jürgen  
Digital Paper Bookmarks, 38, 85
- Sun, Yixing  
ProfileSkim, **111**
- Tag Reading System, 28
- Tashman, Craig  
LiquidText, **23**, 70, 130
- TileBars, **112**, 115
- TouchMark, 24
- Turning The Pages, **25**
- Vertegaal, Roel  
PaperPhone, 24  
TouchMark, 24
- Watanabe, Jun-ichiro  
Booksheet, 24, 38
- Wellner, Pierre  
DigitalDesk, **27–28**, 130
- Wightman, Doug  
TouchMark, 24
- WikiTUI, **28**
- Wilde, Oscar, 31
- Witten, Ian H.  
Realistic Electronic Books, **25–26**, 37, 41,  
85, 112, 130
- Wu, Chih-Sung  
WikiTUI, **28**
- Xiao, Jun  
PageSpark, **23**
- XLibris, **22–23**, 58–59, 70, 112, 130
- Yang, Xue Dong  
HotMaps, **112**
- Yixing, Sun  
SmartSkim, 105, **111**
- Zipfian Distribution, 136

