

Citation for published version:

Alexander, J & Cockburn, A 2008, An Empirical Characterisation of Electronic Document Navigation. in *Proceedings of Graphics Interface 2008*. GI '08, Canadian Information Processing Society, pp. 123-130.

Publication date:
2008

Document Version
Peer reviewed version

[Link to publication](#)

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

An Empirical Characterisation of Electronic Document Navigation

Jason Alexander*

Andy Cockburn†

Department of Computer Science and Software Engineering
University of Canterbury
New Zealand

ABSTRACT

To establish an empirical foundation for analysis and redesign of document navigation tools, we implemented a system that logs all user actions within Microsoft Word and Adobe Reader. We then conducted a four month longitudinal study of fourteen users' document navigation activities.

The study found that approximately half of all documents manipulated are reopenings of previously used documents and that recent document lists are rarely used to return to a document. The two most used navigation tools (by distance moved) are the mousewheel and scrollbar thumb, accounting for 44% and 29% of Word movement and 17% and 31% of Reader navigation. Participants were grouped into stereotypical navigator categories based on the tools they used the most. Majority of the navigation actions observed were short, both in distance (less than one page) and in time (less than one second). We identified three types of within document hunting, with the scrollbar identified as the greatest contributor.

Keywords: Document navigation, document use, scrolling, event logging

Index Terms: H.5.2 [Information Interfaces and Presentation]: User Interfaces—Interaction Styles

1 INTRODUCTION

In many computer applications users need to work with documents that are larger than can be conveniently displayed within one window. Users therefore need tools such as scrollbars, the mousewheel, zooming, and so on, to navigate between document regions. Although these tools are supported by most user interfaces and are heavily used by almost all computer users, there is little empirical data characterising their use.

Facilities for moving the document region displayed within the window have been provided since the very first graphical user interface, with Sutherland's Sketchpad supporting dials for moving the x- and y-coordinates [25]. Since then, many improvements have been proposed. The scrollbar has been embellished with visualisations and bookmarking facilities [4, 5, 14, 20], new navigation behaviours based on panning and zooming have been proposed and evaluated (e.g. [3, 16]), and content-based navigation paths have been explored [17]. Performance analysis based on target acquisition has also been conducted [2, 15]. However, questions of how navigation facilities are used 'in the wild' remain largely unanswered.

A few studies, however, have reported high level observations of document navigation, noting aspects of poor user support. In their analysis of web navigation, Byrne et al. [6] observed that scrolling was an 'obvious case where widget design could make a difference'

(p. 550), and in comparing reading electronic documents with paper ones O'Hara and Sellen [23] stated that electronic navigation was 'irritatingly slow and distracting' and that users need 'quicker, more effortless navigation'. These observations, made in 1999 and 1997, may still be valid today, or, the frustration and inefficiency may have been ameliorated through improved system responsiveness and the many enhanced features such as zooming, bookmarks, advanced scroll-wheels, rate based scrolling, and so on, that are prevalent in current user interfaces. More recent studies such as Weinreich et al's [27] into web navigation (2006) noted that 'requiring users to both scroll and flip through pages seems to be inefficient'.

The goal of our research is to establish a firm empirical characterisation of what users currently do when navigating documents in the 'real world' by logging their naturalistic interaction patterns over a period of several months. This analysis will answer questions such as 'how frequently do users need to navigate' and 'to what extent are different navigation tools used?'. Empirical characterisations similar to these exist for web navigation (e.g. [9, 26]), but surprisingly they do not exist for the more common task of general document navigation.

The next section describes related work on document navigation systems, studies and analysis methods, followed by a brief description of our logging tool, called AppMonitor. We then describe the method and results of our four month longitudinal study. The results are complimented by inline discussion, providing suggestions as to how interface designers may use this data to improve navigation interface design.

2 RELATED WORK

There are three areas of related research that are important for this work, a background of document navigation systems, studies into document navigation, and methods for observing interaction.

2.1 Document Navigation Systems

Desktop document interaction is one of the foremost computer based activities. Many researchers and developers have tried to improve the efficiency of interactive controls for this task.

The scrollbar is the primary interface control for navigating within electronic documents. Embellishments to scrollbars have been proposed by Byrd [5], who added visualisations within the scrollbar to aid the search for query terms, by Hill and Hollan [14] who added edit and read volume visualisations to the scroll trough, and by Björk and Redström [4] who added change, readability and search indicators into the scroll troughs. Laakso et al. [20] added a bookmark area next to the scroll thumb, to allow the easy relocation of content. Ishak and Feiner [17] modified the linear behaviour of the scrollbar, to follow the author's pre-specified content paths, useful in multi-column documents.

Rate-control systems (e.g. dragging with the middle mouse button in Microsoft Word) give the user direct control over the scroll velocity rather than location. Several researchers have tried to improve the performance of rate-control devices and interfaces. Zhai et al. [28] demonstrated that web browsing tasks (including scrolling) were completed more quickly when using a mouse with

*e-mail: jason@cosc.canterbury.ac.nz

†e-mail: andy@cosc.canterbury.ac.nz



a rate-control isometric joystick than when using a standard scroll-wheel. In a subsequent study, Hinckley et al. [15] demonstrated a cross-over effect, with the scroll-wheel outperforming the joystick for short distance movement, but not for long distances.

Rather than researching the input device, several researchers have examined interfaces that modify the systems interactive response to the users control. Igarashi and Hinckley [16] proposed Speed-Dependent Automatic Zooming (SDAZ), a technique that zooms out from the document as the rate of scrolling increases, with the goal of reducing human perceptual problems associated with rapidly moving data. Cockburn et al. [10] empirically determined that SDAZ outperformed navigation using scrollbars and rate-based scrolling. Sun and Guimbretière [24] combined SDAZ with Rapid Serial Visual Presentation (RSVP) [11], to produce “Flipper”, finding that Flipper outperformed SDAZ. Finally, Appert and Fekete [3] introduced the OrthoZoom Scroller, a navigation technique that requires only the mouse to pan and zoom in a 1D space. They found that OrthoZoom was twice as fast as SDAZ.

The rate-based systems in the previous paragraph all seek to increase the speed that the user can move through the data. An alternative approach, taking thumbnail sidebars to their logical extent, was proposed by Cockburn et al. [8] with Space Filling Thumbnails (SFT). SFT replaces motion with direct page selection from a display showing miniaturised thumbnails of every single document page at once. Clicking on a thumbnail moves directly to the associated page. Their evaluation showed that SFT users quickly learned to apply their spatial memory to aid navigation, resulting in faster page acquisition than other interfaces.

In seeking a consistent methodology for conducting empirical comparisons of navigation techniques, Guiard et al. [13] proposed using Shakespeare’s complete works as a platform for evaluating multi-scale navigation techniques.

2.2 Document Navigation Studies

Several researchers have performed controlled studies to investigate document navigation and presented their theories to model interaction.

Hornbæk and Frøkjær [18] compared a standard scrolling document interface with fisheye and overview+detail interfaces for reading electronic documents. The overview+detail interface supplemented the standard scrolling interface with sidebar thumbnails. The fisheye interface diminished all but the first and last paragraphs of each section, allowing users to click to expand the diminished paragraphs. Their method involved asking participants to write essays and answer questions about a document’s contents, and results showed that answers were of the highest quality with the overview+detail interface, and that using the fisheye interface resulted in the shortest task completion times.

Attempts to form scrolling performance models have resulted in contention between researchers. Hinckley et al. [15] argued that scroll distance and target size were the key factors, demonstrating that Fitts’ law [12] holds for a bi-directional tapping task. Guiard et al.’s [13] evaluation platform is based around the assumption that navigation tasks reduce to target acquisition tasks that adhere to a Fitts’ law model. Anderson [2] proposed that scrolling time is limited by visual search ability and that it was in fact better modeled using a linear function of movement distance.

2.3 Document Revisitation

Document revisitation has most widely been studied in the context of the Internet. Numerous researchers [7, 9, 22, 26] have empirically characterised the revisitation patterns of web users. The most recent of these studies ([22], 2007) found that 43.7% of web page requests were revisitations of pages previously viewed. The usefulness of history mechanisms to aid revisitation has also been studied. Obendorf et al. [22] found the ‘back’ button was used for 31% of all

page revisits, while the combination of bookmarks, the homepage button, the history list and typed URLs were only responsible for 13.2% of page revisits. The web, however, is quite a different context to that of interest to this paper—desktop electronic document manipulation.

2.4 Client-side Log Analyses

Although relatively rare, client-side log analysis of interaction has a variety of advantages over other evaluation techniques such as observational field studies, survey methods, and controlled experiments. First, it is relatively easy to conduct a longitudinal study over weeks or months of system use. Second, it is arguably less susceptible to the Hawthorne effect [21] than other methods because users are engaged in their everyday work during the analysis, without the physical presence of observers (in the case of field studies) and without risk of responses that are post-hoc rationalisations of activities (in the case of survey techniques). Third, they can generate data at a fine granularity and they can capture rare events that may not occur frequently enough to be observed in a field or controlled study. Finally, because the logging process does not directly intrude on the user’s work, the method is easily scaled to high numbers of users—indeed, large software companies now use this technique to gain feedback on their software as part of ‘customer experience’ programs (see www.microsoft.com/products/ceip/ as an example).

The primary disadvantages of automated client-side logging studies are that they do not capture the context of the user’s actions (for example, if a user randomly fidgets with the scrollbar while drinking tea these actions will be logged as equally salient with all others), and that developing logging software has been prohibitively complex.

Until recently, the difficulty in implementing logging software has discouraged and impaired its use. For example, in analysing web navigation, participants have been asked to abandon their preferred proprietary web-browsers in favour of customised logging versions of open source systems (e.g. Tauscher and Greenberg [26]) and to use ‘roll-your-own’ systems (e.g. Kellar et al. [19]).

Screen recorders such as TechSmith’s Morae (<http://www.techsmith.com/morae>) offer similar capabilities to client-side logging tools, but they normally only record low-level events such as keyboard and mouse presses. They therefore fail to capture system-level semantics of the action that can be determined by client-side logging tools (e.g. discriminating the particular widget pressed, and the current system state).

3 APPMONITOR

AppMonitor is a Microsoft Windows based program that logs user actions in Microsoft Word and Adobe Reader. The system is described fully by [1], however, due to its integral part in this work, we briefly describe it here.

AppMonitor allows logging of all user actions in unmodified Microsoft Windows based programs. Once installed, AppMonitor requires no input from study participants. It automatically starts when a user logs into the computer and captures salient events whenever an application of interest is opened. Log files are uploaded to a web server when a local buffer size limit is reached.

AppMonitor can record low level events, such as mouse movement and key presses, as well as high level interaction such as menu selections and button presses. It also tracks changes in document state such as the current view and page length, scrollbar position(s), and zoom level.

The set of events logged is configurable by the researcher. Each event is entered into a log file with a date, timestamp, and window handle information, allowing individual actions to be linked to specific documents and applications, even when multiple documents are open simultaneously.



Importantly for the studies presented here, AppMonitor allows either full key logging (where every key press is logged) or shortcut key logging. AppMonitor defines shortcut key logging as any key press with a modifier (Ctrl or Alt), any of the arrow keys, number pad keys or Insert, Delete, Home, End, Page Up or Page Down, Enter, Tab or Space. When logging a shortcut key press AppMonitor also checks to determine whether the Shift key is depressed (as this can be used to differentiate, for example, between simply moving the cursor and highlighting a region).

The AppMonitor system was developed on the Microsoft Windows XP platform for logging actions in Microsoft Word 2003 and Adobe Reader 7. These software versions were used by all participants in our study. Interested readers should see [1] for using AppMonitor on other versions of these applications.

4 LONGITUDINAL STUDY

The longitudinal study monitored the navigation actions of 14 volunteer participants over a period of 120 days. The participants were Computer Science staff and postgraduate students, two of whom were female. They were asked to install AppMonitor and continue their everyday work as normal.

We took two actions to reduce participants' privacy concerns regarding the data that was to be logged. First, we disabled AppMonitor's ability to record all keystrokes. Without this action we would have been able to reconstruct all of the documents the users created during the study, raising clear privacy concerns. Only shortcut key logging was active during the study. Second, we showed all of the users how to display AppMonitor's logging window, which shows the events being recorded to the log files in real time. Several participants stated that they felt more comfortable running AppMonitor having seen that the logs were not capturing sensitive information. Finally, we also disabled mouse-movement logging to reduce both the size of the data files and the processing demands of running AppMonitor.

5 RESULTS AND DISCUSSION

This section presents the results and a discussion of the data recorded in the longitudinal study. Some of the analyses maintain the distinction between Microsoft Word and Adobe Reader. This is to allow a comparison between the tool sets available in each of these applications, not to try and rate one as "better" than the other.

The plain text log files for each participant ranged in size from 0.82MB to 19MB (mean 5.4MB, s.d. 4.9MB). Participants' use of the logged applications varied widely, some used these programs nearly every day, others less regularly, as Windows was installed only on their secondary machine.

5.1 Document Usage Sessions

To ease analysis, we introduce the concept of a Document Usage Session or DUS. A DUS is a period of time where the user is interacting with their document, or more accurately, a period of time where AppMonitor has recorded at least one event from that document. A session begins when an event for that particular document is registered. It then continues until a period of five minutes has passed without AppMonitor collecting any events. The document is then classified as *idle*, meaning it is assumed that the user is not using or interacting with the document in any manner. A DUS is terminated when the document is closed. A document always contains at least one DUS.

It should be noted, as stated earlier, that full keyboard and mouse logging was disabled for this study. Potentially, a user may spend a large amount of time reading a fixed part of the document (without moving its position) and after five minutes this interaction time will no longer be counted. A user could also spend a long period of time typing, without using any navigation keys (arrows, enter, tab, home, end etc). This situation is unlikely, as typing is an activity

that requires corrections and re-writing, and will inevitably require the system to automatically scroll to keep the cursor position on-screen. We believe any "missed" interaction time is balanced by the time where users move immediately away from the application while we continue to think they are interacting.

DUS's are useful when applying time analysis to the data recorded by AppMonitor. We noticed that users regularly leave their documents open for hours, overnight, or even days without interacting with them. DUS discards this idle time when a user is not interacting with a document.

5.2 Document Use and Re-use

Document use and re-use gives an indication as to the extent of application utilisation likely to be observed. Table 1 summarises overall document use and reuse, with means over all participants and daily statistics based on the number of days where AppMonitor recorded events for that application. All documents are treated with equal salinity—unsaved "scratch-pad" type documents are not differentiated from full documents with extensive interaction.

The average number of documents opened per day was four for both Word and Reader. This is approximately one every two hours of a working day, per application, or one an hour between both interfaces.

Analysis	Word (s.d.)	Reader (s.d.)
Mean docs. opened per day	3.7 (1.6)	4.1 (5.6)
Min/max docs. opened per day	1.8 / 7.3	1.3 / 23.1
Mean num. days application used	41.6 (21.6)	25.6 (15.1)
Min/max days applications used	16 / 81	4 / 53
Mean % docs. that are unique*	44.3% (18.2%)	60.4% (26.9%)
Min/max % docs. that are unique	14.2% / 75.6%	12.0% / 100%
Mean doc. usage dist. Zipfs R^2	0.91 (0.05)	0.82 (0.11)
Min/max usage dist. Zipfs R^2	0.8 / 0.99	0.63 / 0.98

Table 1: Document use and re-use statistics (*some titles/lengths of documents were unavailable, so these are omitted)

Approximately half of the documents opened were reopenings of those previously viewed by the user. The usage distributions (the number of times documents were opened), averaged over all users, followed a Zipfian distribution for both interfaces. Both interfaces had minimum uniqueness values below 15% (i.e. only 15% of the documents for at least one user were unique), indicating a high chance of the same document being opened multiple times for some participants. One user had a 100% uniqueness measure, although this is due to them only ever opening five documents in Reader.

5.2.1 Document History Systems

The large range of document uniqueness values observed indicates that interface history features, namely *recent document* lists, are likely to be of value to only some users. Document history systems within Microsoft Word and Adobe Reader are virtually identical and based purely on the last few documents (four by default in Word and five in Reader) that the user has opened¹. This is in contrast to web-browsers such as Internet Explorer (www.microsoft.com/ie) and Mozilla Firefox (www.mozilla.com/firefox) where history systems maintain listings of all of the web-pages viewed over the last period of time (often two weeks), regardless of the number of pages visited. We now investigate how our observations have implications for history mechanisms in document navigation systems.

¹ Adobe Acrobat Standard and Adobe Acrobat Professional (the commercial, feature rich versions of Adobe Reader) do contain a further history mechanism that allows the user to view the documents from: today, yesterday, the last seven, fourteen or thirty days or the last twelve months.



Table 2 documents the use of the history lists in the “File” menu of each application. In Adobe Reader we found only one participant used the history list, selecting the second history item twice. Greater use of the history list was observed in Microsoft Word, two thirds of participants used the history function at some point. This difference can partially be explained by the difference in uniqueness counts (44% in Word and 60% in Reader), however such little use of the history function in Reader can only be explained by user unwillingness to open documents in this manner. Use of this mechanism accounts for only a small percentage in Word and a minuscule percentage of the Reader document reopenings.

Analysis	Word	Reader
History selections, total count	68	2
% of reopenings accounted for	7.1%	0.26%
Number of participants utilising	9/14	1/14
Selection distribution (counts):		
1st History Item	27	0
2nd History Item	21	2
3rd History Item	11	0
4th History Item	9	0

Table 2: History menu item selections

There are two factors that influence the usefulness of the history systems utilised in these two applications: the length of time between reuse of a document and the number of other documents opened in between reuse of a document. Figure 1 shows the average length of time between the closing and re-opening of the same document.

Most documents followed a regular pattern of having a shorter period between closing and opening the more frequently it is used. The anomalies in this graph are interesting—the points where a user has opened the same document many times in the space of a day or two. These points may be explained by inadequacies in the document navigation software. One participant commented that part of his logs may look strange, as he had to reopen Reader every time he recompiled his document using L^AT_EX. The ability to reload the document instead of closing and reopening may ease this user’s navigation requirements.




Figure 1: Average length of time between closing and re-opening of the same document

The number of other documents used between the re-opening of a document effects whether the document in question is still in the *recent documents* list. Figure 2 reports the number of documents opened between reopenings of a particular document. Word documents almost always have less than 20 documents opened between reopenings, however, most counts are still above ten, rendering the

Analysis	Word (s.d.)	Reader (s.d.)
Mean doc. length (pages)*	6.3 (4.5)	38.2 (35.7)
Min/max doc. lengths (pages)	0 / 160	1 / 1743
Mean doc open time (mins)	402.4 (396.8)	356.3 (500.6)
Min/max means of doc open time (mins)	1.2 / 1245.4	1.23 / 1370.8
Mean doc. usage time, DUS idle=5mins	18.8 (8.5)	6.4 (3.3)
Min/max means of usage time, DUS idle=5mins	1.1 / 30.8	2.5 / 14.0
Mean doc. usage time, DUS idle=10mins	25.5 (11.2)	9.1 (5.1)
Min/max means of usage time, DUS idle=10mins	1.2 / 40.7	3.1 / 17.9

Table 3: Document properties (*some document lengths were unavailable, so are omitted from this analysis)

recent documents history mechanism unhelpful for quickly reopening documents. In Adobe Reader, this history mechanism is even less likely to be helpful, with majority having over 20 other documents used between re-openings.




Figure 2: Average number of documents used between closing and reopenings of the same document

The number of other documents between re-openings shown in Figure 2 and the lengths of time between reopenings shown in Figure 1 suggest that providing history mechanisms such as those in web browsers would be more useful to users than the current systems. A two week history list (14 days) would be adequate to cover virtually all reopening requirements, as described in Figure 1.

5.3 Document Properties

The document properties collected characterise the length and period of document interaction and are summarised in Table 3. The average length of documents opened in Word was 6.3 pages and in Reader was 38.2 pages. Figure 3 demonstrates further the large difference in distributions of document lengths for the two applications. The longest observed document in Word was 160 pages and Reader 1743 pages. Such extremities in document length mean that designers of new navigation systems must be fully aware of the size of the content expected to be used in their system. For example, when opening a large document, the scroll thumb on a standard scrollbar quickly reaches its minimum size and thereafter a one pixel movement in the thumb can result in a movement of several pages in the document, rendering it useless for fine scale adjustments. Designers should consider whether the default scrollbar is always the correct tool for the application, especially when large variances in the size of document are expected.






Figure 3: Distribution of mean document lengths per participant, by application

5.4 Navigation Mechanisms

Document navigation systems allow a user to move around a document that is too large to fit completely on the screen at one time. The applications in this study both contain a variety of navigation tools. The ten most used (by distance) navigation tools are shown in Table 4. The thumb and the mousewheel are the most commonly used navigation tools, followed by keys that allow paging through a document. On average, the mousewheel is used to scroll over six pages in Word and around three pages in Reader, for every document opened. The thumb is used to scroll 4.5 pages in Word and nearly six pages in Reader, per document.

The collected navigation tool use data is used to build up categories of stereotypical navigators for each participant, as shown in Table 5 for Microsoft Word and Table 6 for Adobe Reader. These categories are based on the tools used to move the greatest distance by each participant. Clustering was performed manually and was based on the two most used tools. The two top ranked tools, on average, accounted for 77.0% (s.d. 12.8%) of all scrolling actions, per participant in Word and 75.8% (s.d. 13.4%) for Reader. Only one Word participant and one Reader participant had a third tool of high usage, the remainder of participants had the rest of their navigation distributed between several other tools. In Word, the thumb/mousewheel combination is clearly the most commonly used navigation combination. Reader users had a greater diversity in the tools they applied for their navigation.

Navigator	Num part.	Mean % of dist. using tools (s.d.)
Thumb/Mousewheel	9	33.0% (17.8%) / 45.9% (21.1%)
Paging keys/Mousewheel	2	30.4% (0.8%) / 30.8% (7.3%)
Other combinations	3	–

Table 5: Microsoft Word navigator categories

A second category of tools important in document navigation are those used to adjust the document zoom. Microsoft Word documents had an average zoom over time of 94.2% (s.d. 28.3%), and Reader documents 99.7%, (s.d. 13.8%). Users spent 59% of their time with the zoom level of their document at 100%. 24% of their time was spent at zoom greater than 100% and the remaining 17% at a level below 100%.

Table 7 illustrates the tools used for zooming in each of the applications. On average, the zoom was changed 0.2 times per document in Word and 0.3 times per document in Reader. The Ctrl-Scrollwheel zooming technique made up 14% of Word and 67%

Navigator	Num part.	Mean % of dist. using tools (s.d.)
Thumb/Mousewheel	4	52.5% (18.4%) / 28.5% (17.6%)
Paging keys/Mousewheel	2	35.8% (10.7%) / 35.6% (14.0%)
Paging keys/Thumb	2	41.9% (7.3%) / 27.9% (4.5%)
Paging/One other tool	2	37.5% (21.3%) / 30.2% (0.2%)
Thumb/One other tool	2	46.1% (40.1%) / 45.3% (42.7%)
Other	2	–

Table 6: Adobe Reader navigator categories

of Reader zoom changes, indicating some of our study group are aware of the advanced zoom tools available in these interfaces. A post-study interview of eight of the study participants found that half could competently demonstrate the Ctrl-Scrollwheel technique, and half had never heard of it. Given both the empirical and contextual evidence, it is surprising that more participants did not transfer their knowledge of the Ctrl-Scrollwheel technique in Reader into Word. Advanced navigation features such as this should be better publicised by applications. Two participants never found the need to change the zoom level in Word. Everyone changed the zoom level at some point in Reader.

Word Tool	% of zoom actions	Reader Tool	% of zoom actions
Zoom combo-box	45.6%	Ctrl-Scrollwheel	66.7%
Adjust document layout or document view	28.7%	Zoom In/Zoom Out push buttons	28.8%
Ctrl-Scrollwheel	14%	Zoom In/Out tool	4.3%
Other	12%	Other	0.2%

Table 7: Use of zoom tools

5.5 Navigation Patterns

This section takes a deeper look at some of the navigation actions described in the previous section. A navigation pattern is the continuous use of the same navigation tool, *without* a pause of more than 1.5 seconds. Informal trials found that this pause period was a suitable length of time to assume the user had stopped to use (read) the document at that position. When a pause of greater than 1.5 seconds was observed, the next navigation action was classed as a new navigation pattern. Figure 5 shows the distribution of navigation actions, according to their period. The graph is truncated at 10 seconds for clarity. Over 50% of navigation actions take less than one second to complete. The figure shows an exponential reduction in the percent of navigation actions as the length of the navigation action increases. However, it still shows that a significant portion (46%) of actions take several seconds to complete. Developers of new navigation systems should be aiming to push as many actions as possible into the 0–1 second range.

The distribution of navigation actions according to the distance moved is shown in Figure 6. The graph differentiates between positive (forward) and negative (backward) movement in the document. Over 70% of navigation actions adjust the document by less than a single page. Many researchers are focusing their attention on de-




Figure 4: A vertical scrollbar



Microsoft Word Tool	% total dist. moved	Avg. dist. per doc. (pages)	Adobe Reader Tool	% total dist. moved	Avg. dist. per doc. (pages)
Mousewheel	44.0%	6.8	Scrollbar thumb	30.5%	5.8
Scrollbar thumb	29.1%	4.5	Mousewheel	16.6%	3.2
Paging keys	7.1%	1.1	Paging keys	15.1%	2.9
Scrollbar trough	3.21%	0.5	Bookmarks	12.0%	2.3
Extremity keys	2.3%	0.4	Hand tool	7.8%	1.5
Find	1.3%	0.2	Extremity keys	5.4%	1.0
Highlight-drag	1.3%	0.2	Thumbnails	3.4%	0.7
Arrow keys	1.1%	0.2	Find	3.0%	0.6
Scrollbar arrows	0.5%	0.1	Scrollbar trough	2.3%	0.4
Rate-based scrolling	0.1%	0.01	Next/previous push buttons	1.8%	0.3

Table 4: Navigation tool use, by distance

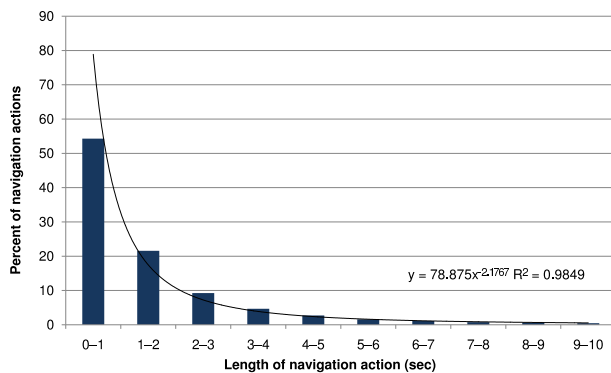


Figure 5: Period of navigation actions

signing navigation systems that are efficient at achieving targets a large distance from the starting point. This result demonstrates that users predominantly navigate small distances at a time. Whether this would change if more efficient long distance navigation tools were deployed remains to be seen.




Figure 6: Distance moved when navigating

5.5.1 Document Coverage

AppMonitor allows the monitoring of scrollbar movements as a document is navigated using tools and techniques such as the scrollbar and the mousewheel. This section utilises the scrollbar data to examine the amount of time spent in different areas of a document.

The scrollbar thumb (see Figure 4) not only provides feedback

on the current position in the document, its size is also directly proportional to amount of the document currently visible on-screen. A small thumb indicates a small percentage of the document is currently visible, while a large thumb indicates a large percentage is visible. AppMonitor records the size of the scrollbar trough, the thumb position and the thumb size. We use this data to calculate the currently visible “window” onto the document, along with the time spent in that window to form a measure of document coverage.

Figure 7 shows a summary of document coverage for all documents in the two applications of interest. For clarity, we have divided the document into 10% chunks—each representing a section of the document. As with other measures reported in this paper, we use a DUS timeout of five minutes.




Figure 7: Document coverage, by application

We observe two quite different patterns for the two applications. Adobe Reader documents, in general have less time spent in a particular section the further towards the end one progresses. In contrast, Microsoft Word documents follow an almost normal distribution, with the most time spent in the middle third of the document. These differences likely arise from the editing ability of Microsoft Word, with the middle of the document taking longer, or requiring more time than the extremities. In Reader this usage pattern is likely caused by a tendency for people to read the start of a document more slowly and thoroughly than later sections. This result highlights the importance of executive summaries or abstracts written at the start of an electronic document. They are more likely to be read than other areas of the document, possibly prompting further reading, printing or closing.

5.5.2 Document Hunting

We use the term *document hunting* to refer to the repeated backwards and forwards movement of a document in order to locate



or center an object of interest on-screen. Document hunting can be observed when utilising any multi-directional tool, such as the scrollbar thumb, the mousewheel and Adobe Reader's Hand-tool. It is important to have an empirical understanding of document hunting, as users can waste a large amount of time moving indirectly toward their target.

We determine document hunting to be occurring when the user navigates to a position in the document and immediately begins a navigation action in the opposite direction. This may then be repeated several times until the user settles on a position in the document. We restrict our recordings of document hunting to immediate changes in direction using the *same* tool.

There are three possible hunting scenarios, as illustrated in Figure 8. Each diagram represents the user's position in the document over time, with start and end points of the navigation indicated by a filled circle. Figure 8(a) illustrates closing in on a point, where the user gradually narrows down on the object of interest, after initially over-running the target. Figure 8(b) illustrates a hunting movement that returns the user to their starting position. Finally, Figure 8(c) illustrates indirect movement to a position elsewhere in the document. Each diagram indicates the *hunting time* and the *actual distance to target*.

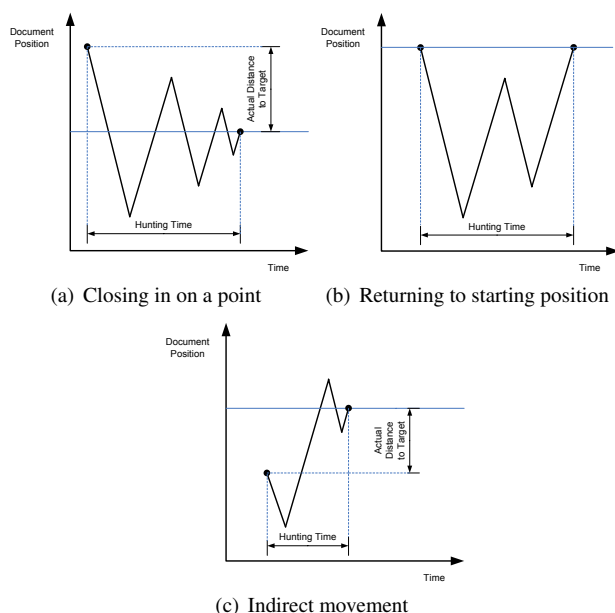


Figure 8: Document hunting scenarios

A summary of the hunting patterns aggregated over all of our participants is presented in Table 8. In total, document hunting accounts for approximately 16% of all navigation actions. Closing in on a point and indirect movement are the most commonly observed types of hunting. Users can spend a significant amount of time and move over a large number of extra pages due to the hunting effect.

Analysis	Close on point	Return to start	Indirect
% scroll actions	7.3%	1.0%	8.2%
Mean # direction changes	1.8	1.3	2.9
Mean hunting time (sec)	56.4	15.0	19.0
Mean doc. length hunting in	19.0	8.8	19.6
Mean extra pgs. moved over	4.2	5.9	6.9

Table 8: Document hunting statistics

A break-down of hunting, by tool, is displayed in Table 9. Clearly, the scrollbar thumb is the worst "offender" for all types of document hunting. Unfortunately, as per most of the empirical results presented here, the data cannot tell us whether these are voluntary actions, or interface deficiencies. Our future studies hope to reveal reasons for this observation. Each type of hunting raises possible areas of improvement for interaction designers. 'Closing in on a point' hunting could be reduced by improving target acquisition using the scrollbar, or by providing alternate tools for this task. 'Returning to the starting position' and 'indirect movement' are a product of the user not being aware of their final destination from the current position. Providing the user with improved document overview tools would aid the identification of, and direct movement to, the desired document position.

Tool	Close on point	Return to start	Indirect
Scrollbar thumb	72%	71%	80%
Mousewheel	18%	22%	9%
Hand-tool (Reader only)	5%	3%	8%
Other	5%	4%	5%

Table 9: Document hunting, by tool

6 LIMITATIONS

This paper characterised the document navigation habits of the 14 participants in our study. All participants were members of Computer Science departments, and classified themselves as "advanced" or "expert" users. There is no reason to believe these results do not generalise to any "expert" user. A larger, more broad field study, with participants who don't classify themselves as experts would be needed to determine whether navigation patterns change as proficiency increases.

One of the primary disadvantages of studies such as this, that utilise logging techniques, is their inability to provide contextual information. The data cannot provide information on the task the user was trying to achieve, or why a particular tool was used for a particular task. The work presented here forms part of a wider study aiming to understand the use of document navigation tools and the results used to suggest improvements. Many avenues of future work are being explored, including the use of think-aloud protocols during unconstrained observations of document navigation.

7 CONCLUSIONS

This paper has presented a longitudinal study characterising how users navigate through their documents using two of the prominent document preparation and review systems: Microsoft Word and Adobe Reader.

Approximately half of documents viewed are reopenings of ones previously used, however between individuals we observed a wide variation of reuse. There was minimal to no use of recent document history mechanisms to access these 'popular' documents, giving rise to the possibility of applying web-browser style history lists to aid frequent revisitation.

The mousewheel, scrollbar thumb and paging keys account for 80.2% and 62.2% of all navigated distance in Word and Reader respectively. Word users can be stereotyped into one of three 'navigator' categories based on their most common interaction techniques: scrollbar thumb/mousewheel users, paging keys/mousewheel users and others. Reader users were more diverse in their tool use, requiring six categories for complete classification. Most navigation actions are short, both in terms of time and distance. Over half take less than one second and over 70% move less than a single page.



Each of the applications were observed to have different document coverage patterns. Word documents followed an approximate normal distribution, with the middle third most viewed. Reader documents followed a more linear model, with a smaller amount of time spent in a region of the document, the closer it was to the end.

Finally, three types of document hunting were identified: Closing in on a point, returning to the starting position and indirect movement. Document hunting accounted for 16.5% of navigation actions. The scrollbar thumb is the tool that most encourage hunting.

This work forms only the beginning for other researcher to further characterise, model and build theories describing document navigation. The discussion provided with these results indicates possible areas of improvement for document navigation systems.

ACKNOWLEDGEMENTS

The authors wish to thank all of the participants who volunteered to be part of this study. This research is partly funded by New Zealand Marsden Grant number 07-UOC-013.

REFERENCES

- [1] J. Alexander, A. Cockburn, and R. Lobb. AppMonitor: A Tool for Recording User Actions in Unmodified Windows Applications. *Behavior Research Methods, In Press*, 2008.
- [2] T. H. Andersen. A Simple Movement Time Model for Scrolling. In *CHI '05: Extended Abstracts on Human Factors in Computing Systems*, pages 1180–1183, New York, NY, USA, 2005. ACM Press.
- [3] C. Appert and J.-D. Fekete. OrthoZoom Scroller: 1D Multi-Scale Navigation. In *CHI '06: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 21–30, New York, NY, USA, 2006. ACM Press.
- [4] S. Björk and J. Redström. Window Frames as Areas for Information Visualization. In *NordiCHI '02: Proceedings of the Second Nordic Conference on Human-computer Interaction*, pages 247–250, New York, NY, USA, 2002. ACM Press.
- [5] D. Byrd. A Scrollbar-Based Visualization for Document Navigation. In *DL '99: Proceedings of the Fourth ACM Conference on Digital Libraries*, pages 122–129, New York, NY, USA, 1999. ACM Press.
- [6] M. D. Byrne, B. E. John, N. S. Wehrle, and D. C. Crow. The Tangled Web We Wove: A Taskonomy of WWW Use. In *CHI '99: Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pages 544–551, New York, NY, USA, 1999. ACM Press.
- [7] L. D. Catledge and J. E. Pitkow. Characterizing Browsing Strategies in the World-Wide Web. In *Proceedings of the Third International World-Wide Web Conference on Technology, Tools and Applications*, pages 1065–1073, New York, NY, USA, 1995. Elsevier North-Holland, Inc.
- [8] A. Cockburn, C. Gutwin, and J. Alexander. Faster Document Navigation with Space-Filling Thumbnails. In *CHI '06: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1–10, New York, NY, USA, 2006. ACM Press.
- [9] A. Cockburn and B. McKenzie. What Do Web Users Do? An Empirical Analysis of Web Use. *International Journal of Human-Computer Studies*, 54(6):903–922, 2001.
- [10] A. Cockburn, J. Savage, and A. Wallace. Tuning and Testing Scrolling Interfaces that Automatically Zoom. In *CHI '05: Proceeding of the SIGCHI Conference on Human Factors in Computing Systems*, pages 71–80, New York, NY, USA, 2005. ACM Press.
- [11] O. de Bruijn and R. Spence. Rapid Serial Visual Presentation: A Space-Time Trade-Off in Information Presentation. In *AVI '00: Proceedings of the Working Conference on Advanced Visual Interfaces*, pages 189–192, New York, NY, USA, 2000. ACM Press.
- [12] P. M. Fitts. The Information Capacity of the Human Motor System in Controlling the Amplitude of Movement. *Journal of Experimental Psychology*, 47(6):381–391, 1954.
- [13] Y. Guiard, M. Beaudouin-Lafon, Y. Du, C. Appert, J.-D. Fekete, and O. Chapuis. Shakespeare's Complete Works as a Benchmark for Evaluating Multiscale Document Navigation Techniques. In *BELIV '06: Proceedings of the 2006 AVI Workshop on Beyond Time and Errors*, pages 1–6, New York, NY, USA, 2006. ACM Press.
- [14] W. C. Hill, J. D. Hollan, D. Wroblewski, and T. McCandless. Edit Wear and Read Wear. In *CHI '92: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 3–9, New York, NY, USA, 1992. ACM Press.
- [15] K. Hinckley, E. Cutrell, S. Bathiche, and T. Muss. Quantitative Analysis of Scrolling Techniques. In *CHI '02: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 65–72, New York, NY, USA, 2002. ACM Press.
- [16] T. Igarashi and K. Hinckley. Speed-Dependent Automatic Zooming for Browsing Large Documents. In *UIST '00: Proceedings of the 13th Annual ACM Symposium on User Interface Software and Technology*, pages 139–148, New York, NY, USA, 2000. ACM Press.
- [17] E. W. Ishak and S. K. Feiner. Content-Aware Scrolling. In *UIST '06: Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology*, pages 155–158, New York, NY, USA, 2006. ACM Press.
- [18] Kasper Hornbæk and Erik Frøkjær. Reading of Electronic Documents: the Usability of Linear, Fisheye, and Overview+Detail Interfaces. In *CHI '01: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 293–300, New York, NY, USA, 2001. ACM Press.
- [19] M. Kellar, K. Hawkey, K. M. Inkpen, and C. Watters. Challenges of Capturing Natural Web-based User Behaviours. *International Journal of Human-Computer Interaction, In Press*, 2007.
- [20] S. A. Laakso, K.-P. Laakso, and A. J. Saura. Improved Scroll Bars. In *CHI '00: Extended Abstracts on Human Factors in Computing Systems*, pages 97–98, New York, NY, USA, 2000. ACM Press.
- [21] E. Mayo. *The Human Problems of an Industrial Civilization*. Cambridge, MA: Harvard University Press, 1933.
- [22] H. Obendorf, H. Weinreich, E. Herder, and M. Mayer. Web Page Revisitation Revisited: Implications of a Long-Term Click-Stream Study of Browser Usage. In *CHI '07: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 597–606, New York, NY, USA, 2007. ACM Press.
- [23] K. O'Hara and A. Sellen. A Comparison of Reading Paper and On-Line Documents. In *CHI '97: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 335–342, New York, NY, USA, 1997. ACM Press.
- [24] L. Sun and F. Guimbretière. Flipper: A New Method of Digital Document Navigation. In *CHI '05: Extended Abstracts on Human Factors in Computing Systems*, pages 2001–2004, New York, NY, USA, 2005. ACM Press.
- [25] I. E. Sutherland. Sketch Pad A Man-Machine Graphical Communication System. In *DAC '64: Proceedings of the SHARE Design Automation Workshop*, pages 6.329–6.346, New York, NY, USA, 1964. ACM Press.
- [26] L. Tauscher and S. Greenberg. How People Revisit Web Pages: Empirical Findings and Implications for the Design of History Systems. *International Journal of Human-Computer Studies*, 47(1):97–137, 1997.
- [27] H. Weinreich, H. Obendorf, E. Herder, and M. 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 Press.
- [28] S. Zhai, B. A. Smith, and T. Selker. Improving Browsing Performance: A Study of Four Input Devices for Scrolling and Pointing Tasks. In *INTERACT '97: Proceedings of the IFIP TC13 International Conference on Human-Computer Interaction*, pages 286–293, London, UK, 1997. Chapman & Hall, Ltd.

