IMPROVING WEB PAGE REVISITATION: ANALYSIS, DESIGN, AND EVALUATION

ANDY COCKBURN
SAUL GREENBERG
STEVE JONES
BRUCE MCKENZIE
MICHAEL MOYLE

ABSTRACT

In this paper, we distill several years of our research on understanding and improving how people return to their previously visited web pages. Our motivation is that web page revisitation is one of the most frequent actions in computer use, and consequently any interface improvements in this area—no matter how small—can have a very large effect. We report our findings across five categories of revisitation research: characterisations of user behaviour; system models of navigation and their impact on the user's understanding; interface methods for increasing the efficiency of the Back button; alternative system models for navigation; and alternative methods for presenting web navigation histories. The behaviour characterisation shows that revisitation is a dominant activity, with an average of four out of five page visits being to previously seen pages. It also shows that the Back button is heavily used, but poorly understood. Three interface strategies for improving web page revisitation are described. The first, a gesture-based mechanism for issuing the frequent Back and Forward commands, addresses low-level interface issues, and is shown to be both popular and effective. The second, a 'temporal' behaviour for the Back and Forward buttons, aims to overcome the problems associated with poor understanding of the current behaviour of Back. Although the results do not conclusively show advantages for the temporal behaviour of Back, they strongly suggest that revisitation can be improved by providing temporally ordered lists of previously visited pages. The third interface scheme investigates how next-generation browsers could integrate the current tools for revisitation into a single utility, and how simple visualisation methods can be used to aid users in identifying target pages displayed in miniature.

Andy Cockburn is a senior lecturer in the Department of Computer Science at the University of Canterbury, Christchurch, New Zealand. He is the director of the Canterbury Human-Computer Interaction and Multi-Media Laboratory.

Saul Greenberg is a full professor in the Department of Computer Science at the University of Calgary, Canada. He is the director of the Grouplab Laboratory for HCI and CSCW.

Steve Jones is a senior lecturer in the Department of Computer Science at the University of Waikato, Hamilton, New Zealand.

Bruce McKenzie is a software engineer at the Jade Development Centre in Christchurch, New Zealand.

Michael Moyle is a software engineer at Allied Telesyn Research in Christchurch, New Zealand.

Every day, millions of people worldwide have problems returning to previously visited web pages. These problems often amount to little more than the minor annoyance of finding that a page has 'disappeared' off the set accessible with the Back button. Sometimes, however, extensive searching is necessary to return to a page: for example, when the user forgets to bookmark a 'valuable' page, and he or she cannot remember how it was originally retrieved. Although the impact of these problems on each individual may be small, it is clear that easing these problems can yield enormous benefits when multiplied across millions of users, and billions of page accesses.

Since 1994 we have examined the problems of revisiting web pages, and have designed and evaluated several systems aiming to improve web page revisitation. This paper provides an integration and synthesis of this work, and is presented in the following structure:

- *User behaviour*. We begin by reporting the recent results of a web use log-analysis, which show that revisiting pages is a dominant activity on the web.
- *System and user models of the current behaviour of the Back button*. We describe an easy to repeat experiment to demonstrate that many users misunderstand the rudimentary behaviour of the main interface tool for revisitation—the Back button. We also discuss why, despite this misunderstanding, Back is heavily used.
- *Improving the efficiency of the Back command*. We describe the efficiency limitations of the interface mechanisms used to issue the Back command, and present an evaluation of a gesture-based shortcut (similar to the scheme recently provided in the Opera web browser).
- *Improving understanding and efficiency of the Back model*. We describe an alternative 'temporal' behaviour for the Back and Forward buttons, and present the results of its evaluation.
- *Improving the presentation of revisitation tools*. In discussing the implications of the earlier findings, we demonstrate how next-generation web browsers could integrate and enhance the diverse tools for revisitation that are available in current browsers (Back/Forward, bookmarks, and history lists). This work is on-going, and preliminary results are encouraging.

USER BEHAVIOUR: WHAT DO WEB USERS DO?

Considering that web browsers are among the most widely used computer applications, there has been only modest research into how they are used. In this section, we briefly summarize prior analyses of browsing behaviour, and then describe results from our own recent study of what web users do as they navigate the web.

Prior analyses.

Some researchers detail the demographics of Internet users (e.g., gender percentages, age, occupations, educational attainment, etc.) as well as the demographics of their technologies (e.g., people's connection speed to the Internet and browser selection). The most well known of these is the 1994-1998 biannual WWW Surveys[1] developed by the Graphics, Visualization and Usability Center at the Georgia Institute of Technology (e.g., Kehoe and Pitkow (1996)).

Other researchers have analysed the tasks that people do as they navigate. For example, Byrne et al (1999) videotaped eight people as they used their browser over the course of their day and codified user behaviours. From this, they developed a task taxonomy of browsing, including the six general web tasks:

- Use information—a series of activities where people use information gathered from the web;
- Locate on page—searching for particular information on a page;
- Go to—the act of trying to get the browser to display a particular URL;
- Provide information—sending information to the browser (e.g., authentication, addresses, search terms);
- Configure browser—to change the configuration of the browser itself;
- React to environment—to supply information or deal with a problem on demand of the browser.

---

[1] www.cc.gatech.edu/user_surveys/

The authors then go onto sub-divide these general tasks into more specific ones, and to codify how often they occur. While these results only reflect a total of a few hours of use by a few people, it provides insight into the actual things that people do.

Still others have analysed traces (or logs) of users' actions to reveal statistics of use. Some use server-side logs. Since most web servers collect data indicating when a particular page has been accessed and by what IP address, this is an easy data source to mine. For example, it is relatively simple to analyse logs to expose the frequency of page hits on a web site. However, server-side logs are limited in that they often do not distinguish well between different users; they collect no data on actual browser use; and they are often missing crucial data. Pirolli, Pitkow, and Rao (1996), and Chi, Pirolli and Pitkow (2000) discuss the problems of extracting meaningful information from server-side logs. Another option is to instrument the client i.e., the browser itself (or a client-side browser proxy). The advantage is that client-side logging can record the exact history of the user's actions with his or her particular browser.

Perhaps the most well known client-side log analysis of web use are by Catledge and Pitkow (1995) and Tauscher and Greenberg (1997). These studies instrumented the then-popular XMosaic browser to record the pages that users visited and the interface mechanisms used to access them. The participants in both studies were primarily staff, faculty and students in University Computing Departments. Catledge and Pitkow logged three weeks of use by 107 users in 1995, while Tauscher and Greenberg analysed five to six weeks of use by 23 users in 1995. Catledge and Pitkow revealed that the dominant user interface techniques for visiting pages were clicking on hypertext anchors (52%) and on the Back button (41%). Navigating to pages by typing the URL, by clicking Forward, or by selecting from 'Bookmarks' were all lightly used, accounting for about 2% each. Tauscher and Greenberg confirmed that link selection and clicking Back are the dominant navigation mechanisms, accounting for approximately 50% and 30% of navigation acts. Tauscher and Greenberg also analyzed *the recurrence rate* of page visits: "the probability that any URL visited is a repeat of a previous visit, expressed as a percentage". They found that the recurrence rate for the subjects participating in their study was 58%, and by reanalysing the data from 55 of Catledge and Pitkow's subjects they found a recurrence rate of 61%. This result shows that users had previously seen approximately three out of five pages visited. Their analysis also reveals that people tend to revisit pages just visited a short while ago, that they access only a few pages frequently, that they browse in very small clusters of related pages, and that they generate only short sequences of repeated URL paths Tauscher and Greenberg (1997).

In these early studies, people rarely used bookmarks (less than 2% of user actions). However, a later 1996 survey by Abrams et al (1998) suggested that bookmark use was rising. 84% of his respondents had more than eleven bookmarks, indicating that people at least had the intention of returning to key pages. Indeed, Pitkow (1996) reported from a survey of 6619 users that "organizing retrieved information" is one of the top three usability problems of using the web, reported by 34% of participants.

Our Recent Client-Side Log Analysis of Web Use

While all excellent studies, we do not know if the findings of these early studies reflect current use of the web. Things have changed considerably since the mid-90s. Instead of early adapters and technologically savvy 'elite' users, people from virtually all demographics use the web (although not equally). Modern web browsers have highly polished interfaces with features far beyond those supported by XMosaic. Many web navigation aids such as search engines and web directories are now a fundamental part of web use, yet these tools were either in their infancy or did not exist at the time of the prior studies. The technology itself has changed, where broadband connections now gives nearly instantaneous response—a sharp contrast to the 'World Wide Wait' of the early days. The kinds of information accessible on the web are also completely unlike that of the prior decade. Once an academic repository, it is now a massive commercial and populist arena.

Method. To update and extend the findings of these previous studies, we analysed the web browsing activities of seventeen Computer Science staff and graduate students over a 119-day period from early October 1999 to late January 2000[2]. We gathered the data through the history and bookmark files that Netscape Navigator (versions

---

[2] The full details of this study are available in Cockburn and McKenzie (2001).

4.5-4.7) maintains. The history file keeps a list of the URLs the user has visited, the time of their last and first visit, the number of visits, and the title of each page. The bookmark file stores information about the user's bookmarks and their organisation into folders. We obtained copies of these files through incremental backups that are automatically created every night at our institution. To eliminate the chance that participants would modify their behaviour due to their awareness that their actions were being logged, we asked participants for permission to retrieve their backup files after the terminating date of the study. That is, we asked participants if we could mine their data retrospectively instead of asking if we could monitor their future web uses.

What is the average number of pages visited by a person per day? The participants made a total of 84841 page visits, spread across 17242 different URLs. On average, participants visited approximately 42 pages each per day. The actual number of pages visited by each user per active day would be substantially higher than this, because the average includes weekends and the Christmas/New Year vacation period during which few participants would have used their browsers at work. Even with this underestimation, these results suggest that people visit almost double the number of pages now than in the mid-90s. That is, our means of 42 pages per day mean compares to approximately fourteen (Catledge and Pitkow 1995) and twenty one (Tauscher and Greenberg 1997) pages per day in the earlier studies.

How many of a person's page visits are to ones they have visited previously? The premise behind the Back button, history systems, and bookmarks is that people frequently navigate to pages that they have seen before. As already mentioned, previous studies have found that the average revisitation rate is between 58-61%. Has this figure changed? To analyze revisitation, we first plotted the total pages visited over time by each participant (Figure 1a). The mean total number of page visits by each participant was 4991 (s.d. 6106), ranging from 281 to 24309. Participant 15's heavy web use statistics can be explained by his role as a web-master. We then plotted the *distinct* pages visited over time by each participant (Figure 1b). Here, the mean number of URLs visited was 1227 (s.d. 1086), with a range from 74 to 4251. As the figures indicate, the total number of page visits is strongly correlated with the growth of the number of distinct URLs visited. Linear regression over all subjects gives a slope of 5.1 and an R-squared value of 0.8837 ($F_{1,940}$=7140, p<0.0001). This slope reflects the revisitation rate for the subject pool: for each new URL added to the set of distinct URLs, four pages are revisited. This revisitation rate is confirmed using Tauscher and Greenberg's formula for revisitation rate:

Revisitation Rate $R = 100 * \dfrac{total\_visit\_count - total\_URL\_count}{total\_visit\_count}$

This calculates a revisitation rate of 81%, with individual participant revisitation rates ranging from 61% to 92%—substantially higher than the previously reported values of 58 and 61%.

Are bookmarks and other shortcuts to particular pages used? By probing our data, we noticed that one factor behind a high revisitation rate is that almost all participants had one or two pages that they visited far more often than any other. Participant 2, for instance, had visit counts of 4352, 384, 199, and 117 for his top four pages. In general, we found that the participants' top three pages accounted for 20% of all page accesses.

Given the high visit counts to particular pages, we would have thought these would be prime candidates for becoming shortcuts. Netscape Navigator supports a variety of shortcut techniques—including the configurable 'Home' button, bookmarks, and the 'personal toolbar'. We analyzed Netscape's history and bookmark files to see whether people created shortcuts to their frequently visited pages. What we found was that although most users had shortcuts to their top two pages, few had shortcuts for their third, fourth and fifth most frequently visited pages.
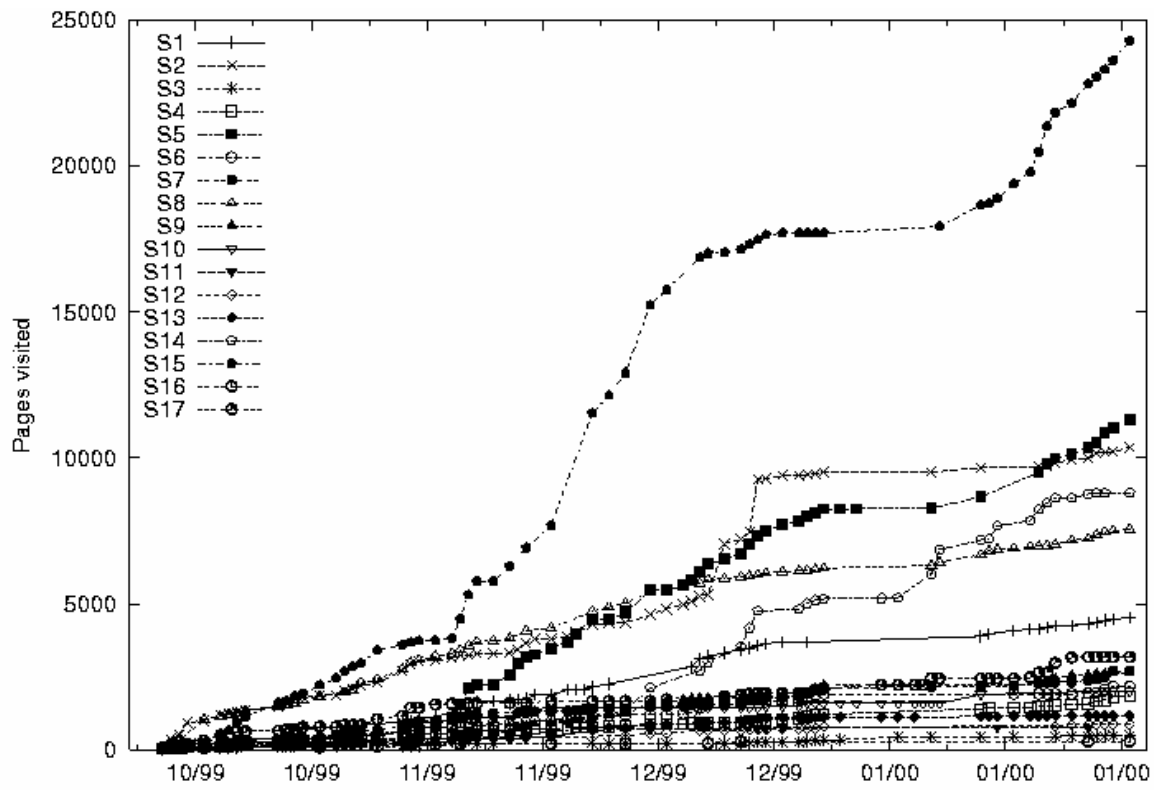
(a) Total pages visited
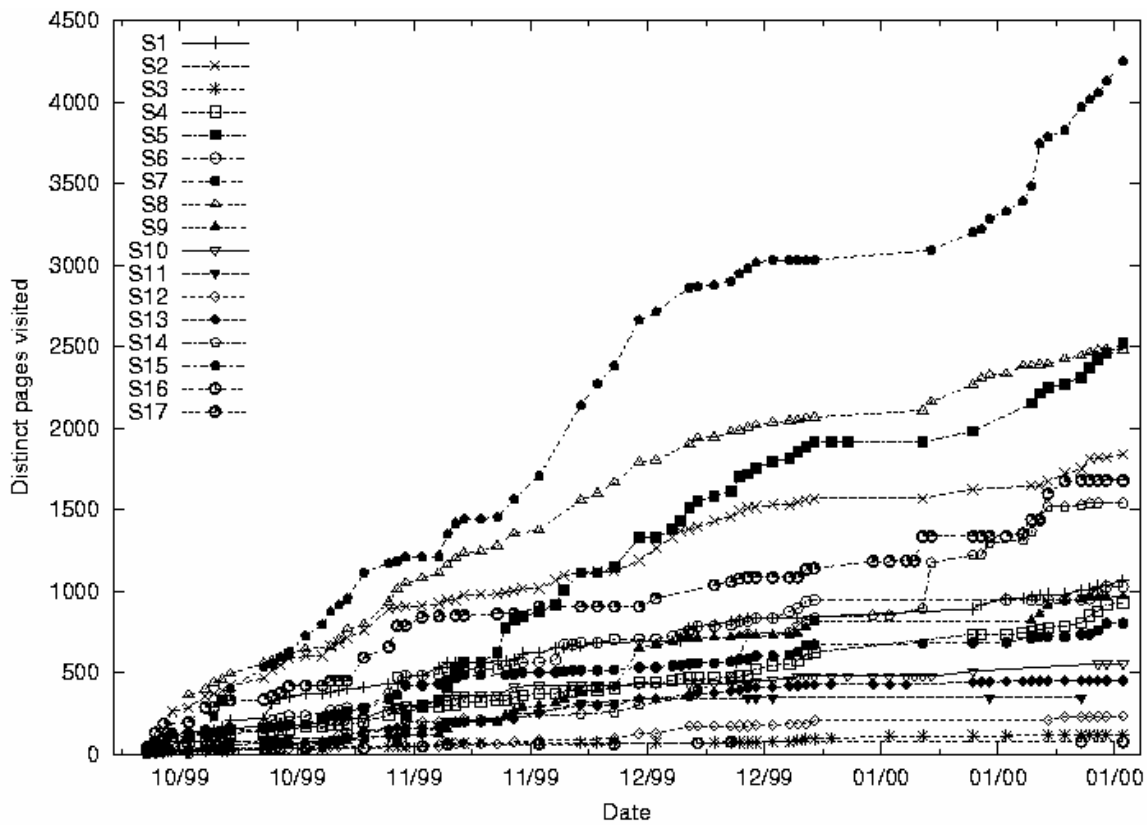
(b) Distinct pages visited.

Figure 1. Growth in the total number of pages visited (a) and distinct pages visited (b) over time (mm/yy) for each user.
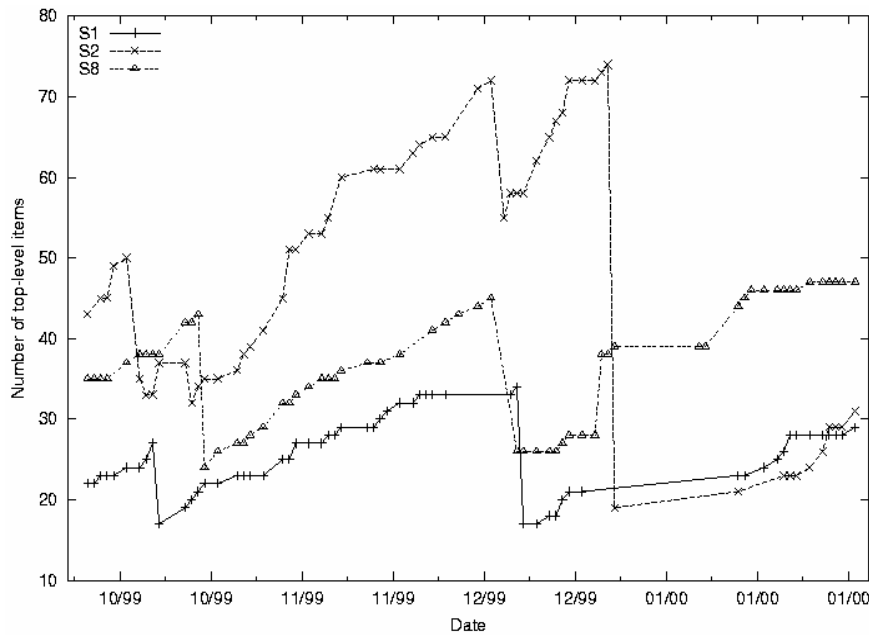
Figure 2. Number of URLs and folders in the top-level bookmark file of three users, plotted against time (mm/yy).

Further analysis of the bookmark file showed wide variation in bookmark use: participant 16 had none, while participants 2 and 8 had more than 500 each. The mean maximum size of the subjects' bookmarks collection was 184, but as suggested above the standard deviation was very high (s.d. = 166.15). Some people organized the bookmarks into folders, where the mean number of folders used to store bookmarks was 18.1 (although again the standard deviation is very high—s.d. = 16.5). We saw that the rate of bookmark addition heavily outweighed the rate of deletion, with averages of 27.6 additions and 3.7 deletions. This imbalance implies that users have, or will have, problems managing the size and organisation of their bookmark collections over time. For example, bookmarks in Netscape are normally selected via a pop-up cascading menu, the length of which depends on the number of "top level" items in the bookmark file. Participant 17 had 130 items in this top-level, which would produce a very cumbersome cascading menu.

Figure 2 plots the number of items in the top-level bookmark structure for three of the subjects over time. The obvious steps in the figure show that participants periodically tried to reorganise their bookmarks to overcome the problem of the menu growing too long; this effect was also noted by Abrams et al (1998). Rather than deleting items, we saw that subjects would typically relocate them to new folders (counted as a top-level item). It is worth noting that the interface mechanisms for managing bookmarks have improved since Netscape 4.5-4.7, possibly easing bookmark management.

The apparent reluctance to delete bookmarks is at odds with the relatively transient nature of Web sites and pages. Two months after collecting the bookmark data, we ran scripts that attempted to access each page in the participants' bookmark collections. Any page returning 404 "Not found", 301 "Moved Permanently", or 5xx (host unavailable) was deemed invalid. Approximately 25% of the pages were invalid. This indicates that over time, bookmark collections will become cluttered with useless items.

What proportion of pages do not have titles? Titles are used by Netscape and Microsoft Internet Explorer in a variety of ways, including labeling items on the Back pull down menu, default identification tags in the bookmark and history lists, and labeling the window-manager border. Missing, incorrect, and inconsistent titles can frustrate the user's ability to identify pages that they wish to return to (Cockburn and Greenberg 2000).

Five percent of the distinct URLs visited by the participants did not have an HTML "Title" tag associated with the page. Although there are alternative interface techniques that could be used to aid page identification (such as thumbnail images of the page (Cockburn et al 1999; Kaasten et al 2002; Robertson et al 1998)), it is likely that

text titles will remain an important page identification cue. There is little that any browser can do to ensure the presence and accuracy of the text titles, but web authoring software could promote careful consideration of page titles. Often page-titling facilities are 'hidden' under sub-menus or dialogue boxes, rather than prominently displayed for each newly authored page.

Summary

In summary, the main findings of the log analysis study are as follows. Web page revisitation is the dominant activity in web browsing. On average users have previously seen four out of every five pages visited. Most participants had one or two pages that they visited far more often than all other pages. Bookmark use was highly varied, with users either making light or heavy use of bookmark collections. Finally, there were two indications that users find bookmark management troublesome: first, 25% of the bookmarks did not refer to legitimate pages when we tried to access them; and second, some of the heavy bookmark users appeared to be forced into 'shuffling' bookmarks when their top-level menus became too long.

SYSTEM AND USER MODELS OF BACK

The experimental results summarised in the previous section show the importance of having efficient and effective mechanisms for returning to previously visited web pages. The Back button is one of the main interface components for returning to web pages (accounting for approximately 40% of navigational actions, according to Tauscher and Greenberg (1997)), yet an easy to repeat experiment, described below, shows that many users misunderstand its behaviour. Having described the experiment, we describe why the Back button is heavily used despite the common misunderstanding.

Understanding of the Back button

Through a simple mass-participation informal study, which takes only minutes to conduct, it is simple to demonstrate that many web users misunderstand the behaviour of the Back button.

To run the informal study in a class, use a web browser projected onto a large screen. Tell the audience that they will be asked to predict whether a page can be returned to using only the Back button (not using links or other history tools). Start with the browser showing a home page that is well known to the audience: we normally use the "Computer Science Homepage". Then slowly and clearly use a link to visit a page off the home page, for example, "People in Computer Science". Tell the audience that they will revisit this page shortly. Clearly demonstrate the use of the Back button to return to the homepage. Next, slowly and clearly use a different link to visit another page off the home page. With the browser displaying this final page, ask the students to write on a scrap of paper whether the first linked page ("People in Computer Science") can be returned to using the Back button and only the Back button (no links or history tools). Ask them to write the number of times the Back button will have to be clicked if the page is accessible with Back. Finally, collect the responses. Using paper responses encourages participation, and protects the participants from possible embarrassment in a show of hands.

The stack-based behaviour of Back means that the page is inaccessible with Back.

We first conducted a more elaborate version of this experiment in 1995, when the web was still relatively new (Cockburn and Jones 1996). The eleven participants were all Computer Scientists who used the web daily. Eight wrongly predicted that the page could be accessed with Back. Since then, we have repeated the experiment as described above many times in large Computer Science undergraduate classes. We estimate that approximately 50% of the students—all of who use the web in their studies—misunderstand whether pages are accessible with Back, with the most common error being to state that the "People in Computer Science" page can be returned to with two Back clicks.

Stack-based behaviour

All major commercial browsers use the same stack-based model for web page navigation. In this model, there are two ways of displaying pages in the browser, which we call 'load' and 'revisit'. Pages are loaded when the user clicks on a link, types a URL, selects a 'Favorite' page, and so on. The effect of 'load' is to add the page to the top of a stack of visited pages. Pages are revisited with the Back and Forward buttons, which move downward and upward through the stack of visited pages. The menus associated with Back and Forward allow users to directly

revisit pages on the stack. When the user loads a page after revisiting pages, the new page is added to the stack immediately above the current stack position, and all pages above that stack position are removed. This explains why the "People in Computer Science" page cannot be revisited in the previous example.

Using the symbol → to indicate loading a page (e.g. clicking on a link) and the symbol ⇐ to indicate revisiting a page with Back, if the user navigates through pages **a→b→c**, then the stack contains **{a, b, <u>c</u>}** with page **c** shown in the browser (indicated by underlining). If the user then navigates back to **a** by pressing the Back button twice (**c⇐b⇐a**) then the stack contains **{<u>a</u>, b, c}**. If the user then clicks on a link to page **h** (with **a→h**), then **b** and **c** are removed from the stack, giving **{a, <u>h</u>}**.

What's good about the Back button?

The two primary limitations of Back are that many users misunderstand its operation and that not all recently seen pages can be revisited. Yet, despite these limitations, Back is heavily used. There are several factors contributing to Back's success (Cockburn and Greenberg 2000):

- It can allow rapid return to recently visited pages (assuming they are still on the stack).
- It is robust. People can use it even with a naïve model of the way it works.
- It is cognitively undemanding. Users can backtrack through pages using a simple 'click until the desired page is recognised' strategy.
- It is 'ready-to-hand'. Unlike interface features such as the 'History list', which must be explicitly popped up when needed, there is little overhead in accessing Back because it is on constant display.
- It consumes minimal screen real estate.

These beneficial properties have led to the Back button being one of the most heavily used interface components in existence.

Summary

A large proportion of regular web users misunderstands the stack-based behaviour of the Back button. Despite this misunderstanding, they use it heavily because it has many desirable properties. This conflict between heavy use and misunderstanding clearly results in suboptimal use.

The following three sections describe the design and evaluation of three approaches for improving human performance with the Back button. The first investigates the efficiency of using a shortcut 'gesture' action for issuing the Back command. The second evaluates the potential cognitive and motor benefits of changing the stack-based Back model to a complete temporal list, overcoming the problems of pages disappearing off the set accessible with Back. The third investigates the potential benefits of integrating the diverse schemes for web revisitation currently used in browsers.

IMPROVING THE EFFICIENCY OF THE BACK COMMAND: GESTURE NAVIGATION

The mouse-driven cursor is the main input device for web navigation—to visit a link the user points the cursor to the link and to revisit a page the user moves the cursor to the Back button. Fitts' Law (Fitts 1954) predicts the time taken to move the cursor between targets (such as the Back button and the page links), but faster alternatives than mouse pointing are available.

In recognising the importance of providing efficient mechanisms for issuing the Back command, both Microsoft Internet Explorer and Netscape Navigator provide keyboard shortcuts, with Internet Explorer using backspace or 'Alt+left-arrow', and Netscape Navigator using 'Alt+left-arrow'. The main limitation of key-bindings for web browsing is the overhead in homing the hands between the mouse (used for links) and the keyboard. Homing is necessary for the 'Alt+left-arrow' keybinding because many keyboards are arranged so that the Alt and left-arrow keys cannot be simultaneously pressed with one hand. Similarly, for right handed users, the backspace keybinding is awkward because the user must either reach across the keyboard with their left hand, or take their right hand off the mouse. Another shortcut for Back, provided by both browsers, is the context menu that can be popped up by pressing the right mouse button. Accessing the Back menu item incurs overheads in waiting for the menu to be posted and in the Fitts' Law limitations of pointing to the menu item. Finally, specialised input devices such as the

Microsoft Intellimouse Explorer[3] mouse provide additional buttons for the Back and Foward commands, the these shortcuts are only available to the relatively small proportion of users who buy the devices. From our experience, studies, and subjects' comments, it appears that the Back command is seldom issued through any interface mechanism other than the Back button. The question, then, is how to improve the efficiency of issuing the Back command, without the low-level costs of target acquisition or of homing the hands between keyboard and mouse.

Inspired by work on gesture-based marking menus (Kurtenbach and Buxton 1991), in 2001 we designed, implemented and evaluated a gesture-based mechanism for issuing the Back and Forward commands. Similar features were simultaneously released in the Opera[4] commercial web browser, and slightly later by the Mozilla Optimoz project[5]. To our knowledge, neither of these commercial implementations has been formally evaluated.

With gesture navigation the user 'flicks' the mouse, with the left button held down, left (for Back) or right (for Forward). To support gesture navigation within unaltered commercial browsers in our evaluation, we constructed a web-site where each page contained a Javascript program instructing the browser to navigate Back or Forward whenever a gesture is recognised[6]. Several empirically determined values are used to recognise legal gestures. First, to prevent the script from recognising a simple mouse click (such as a link selection) as a gesture command, the mouse coordinates must change by at least 35 pixels between the button being pressed and released. Second, to distinguish between gestures and text selections (both completed by dragging with the left mouse button), the gesture must be completed within 250ms.

Method

The evaluation compared the effectiveness of gesture navigation with that of the normal Back button across two tasks. Both tasks involved following rehearsed sequences of nine web pages as quickly as possible, first using the normal browser, and then using gesture features. Multiple rehearsals of the web paths were used to minimise the impact of learning effects. Tasks were first completed using the Back button so that we could investigate any deterioration in subjective satisfaction with the Back button after using the gesture system.

Task one compared the effectiveness of the two interfaces in depth-first navigation. The path followed four links on subsequent pages then backtracked with four successive Back commands, giving a path **a→b→c→d→e⇐d⇐c⇐b⇐a** using the notation introduced earlier.This task represents a directed search style of web-use: for instance, searching for a faculty member's web page starting from a university's home page.

Task two examined the effectiveness of the two interfaces in breadth-first navigation, also called 'hub-and-spoke' navigation (Catledge and Pitkow 1995), which involves visiting a series of links (or 'spokes'), one at a time, off a central 'hub' page. Beginning at a 'start' page, the subjects followed a link to a main 'hub' page, then navigated to three 'spoke' links off that page, pressing Back to return to the hub each time. A final Back command returned to the start page, giving a complete path of **a→b→c⇐b→d⇐b→e⇐b⇐a**.

When using the Back button, the mouse-pointing requirements of task two are much higher than task one. After selecting each link to a spoke page the user must point to the Back button, and then point to the next page link. Fitts' Law predicts that the Back button will result in slower task performance. For this reason, a second factor for 'amplitude' was introduced in task two, allowing us to analyse the degradation of performance as the distance between the Back button and the links on the hub-page increased. The links on the hub-page were vertically aligned immediately above one another at one of three corners of the web page: top-left, bottom-left, and bottom-right to give three levels of amplitude 'low', 'medium' and 'high'. Mean distances between the Back button and the group of links were 8.5cm, 14cm, and 20cm for the three levels of amplitude.

---

[3] www.microsoft.com/hardware/mouse/ie_info.asp

[4] www.opera.com

[5] www.optimoz.org

[6] See www.cosc.canterbury.ac.nz/~andy/gestureSite/ for an example.

The twenty participants were all volunteer postgraduate Computer Science students familiar with web navigation. Each participant's training with the gesture system involved a brief (one or two minute) demonstration of the 'flick' gestures and explanation of the rules determining a valid gesture.

Results

In task one, the mean task completion times for the normal Back button and gesture systems were 6.1 (s.d. 1.1) seconds and 5.4 (s.d. 0.96) seconds. This is a significant difference (t_{18}=2.68, p<0.05), showing a reduction of 11% in the mean task time when using gestures.

The performance data in task two were analysed using a two-factor analysis of variance with repeated measures. The factors were 'interface type' with two levels (Back button and gesture system) and 'amplitude' with three levels (low, medium and high). The mean time for the Back button was 7.44 (s.d. 1.12) seconds compared to a mean of 6.09 (s.d. 1.19) seconds for the gesture system, giving a significant main effect: $F_{1,18}$=82.2, p<.001. The gesture system reduced the mean task time by 18%.

The mean times for the three levels of amplitude were not significantly different ($F_{2,36}$=1.17, p=0.32). This is unsurprising because amplitude has no effect on task completion time in the gesture system. As expected, the interaction between factors 'interface type' and 'amplitude' was significant: $F_{2,36}$=6.78, p<0.01. Figure 3 reveals the cause of the interaction—as the amplitude increases, the mean completion times for the Back button increase, but the task times for the gesture system remain relatively constant.

The participants' comments and subjective ratings dramatically amplified the effectiveness of the gesture system, with comments such as "Fantastic", "That's amazing", "Really really nice" and "Brilliant!" Eighteen of the twenty subjects rated the efficiency of the gesture system more highly than the Back button (5-point Likert scale responses). The mean responses for the Back button and gesture system were 2.42 (s.d. 1.1) and 4.26 (s.d. 0.8), giving a significant difference: Wilcoxon Signed Ranks, N=17, z=3.46, p<0.01.

Summary

Gesture based shortcuts for issuing the common Back and Forward commands are already being used in commercial web browsers. The evaluation summarised here shows that these shortcuts are both efficient and popular. The current 'Alt+left-arrow' keyboard shortcut for Back is inefficient because it requires the user to move a hand between the keyboard and the mouse because the keys are normally too far apart to be pressed with one hand. Similarly, the 'backspace' keyboard shortcut requires right-handed users to either remove their hand from the mouse or to reach across the keyboard with their left hand. Gesture solutions, in contrast, are efficient, popular, and available to all users.

Full details of this experiment are available in Moyle and Cockburn (2002).

IMPROVING UNDERSTANDING & EFFICIENCY OF BACK: A TEMPORAL MODEL FOR BACK & FORWARD

While the gesture interface appears to provide a more efficient motor mechanism for issuing the Back command, it does nothing to improve the poor understanding of the Back button, reported earlier.
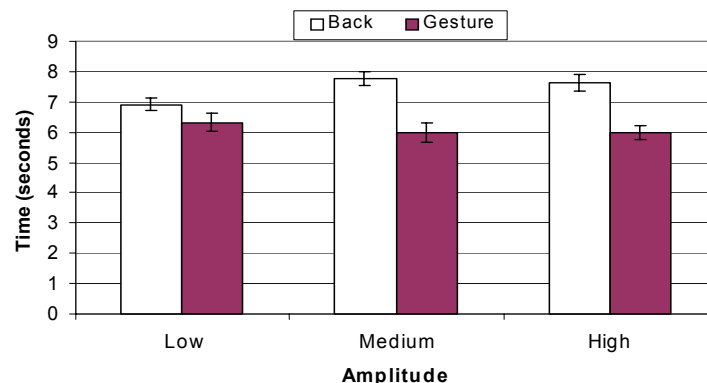

Figure 3. Mean times (and standard errors) in task two for the normal-Back and gesture interfaces.

In Greenberg and Cockburn (1999) we proposed an alternative behaviour for the Back button. This 'temporal scheme with duplicates removed' aimed to maintain the desirable properties of the stack-based behaviour, while overcoming its limitations of poor understanding and of providing an incomplete list of visited pages. The temporal model ensures that all previously visited pages are accessible with Back and Forward, and that Back navigates through pages in the order in which they were last displayed in the browser.

Like the stack-based technique, the temporal scheme maintains two semantics of page display (loading and revisiting) that are executed through a visually identical interface. Loading a page causes it to be added to the end of a 'main temporal list' (denoted **{a, b, c}**, as for the stack example earlier) that serves a function equivalent to the stack. When the Back and Forward buttons are used, each revisited page is added to the end of an internal secondary list (denoted **[a, b, c]**), which is not displayed to the user. After revisiting pages with Back and Forward, when the user next loads a page, the secondary list is concatenated to the end of the main temporal list, and any older duplicate pages are deleted. Finally, the new page is appended to the end of the main temporal list. So, for example, if the user navigates through pages **a→b→c⇐b⇐a**, then the main temporal list contains **{a, b, c}**, and the secondary list contains **[b, a]**. If the user then clicks on a link to page **h**, then the main temporal list will contain **{a, b, c, b, a, h}**. The older duplicates of pages **a** and **b**, shown crossed out at the start of the list, have been deleted.

We recently evaluated the temporal scheme, as summarised below. Full details of the study are presented in Cockburn and McKenzie (2002).

Method

The evaluation was based around a series of web-browsing tasks, with the participants using a 'new release of Netscape Navigator'. They were told that the new release 'may or may not have modified the behaviour of the Back and Forward buttons'. The visual appearance, mouse bindings and response time of the stack-based and temporal interfaces were identical to those of Netscape Navigator version 4.72. To construct the interfaces, unparented windows providing exact visual replicas of the Back and Forward buttons were placed over Netscape's actual buttons. When the subjects clicked the imitation Back or Forward buttons our program modified the state of its stack or temporal list and issued appropriate page display requests to the browser. Back and Forward menus that were visually indistinguishable from those normally produced by Netscape were also supported.

The thirty-four participants—all Computer Science students who used Netscape Navigator daily in their everyday course work—were randomly assigned to either the stack or the temporal interface. Each evaluation lasted approximately one hour. No training or instruction on the behaviour of the web browser was given.

The experiment consisted of repeating six tasks in three different web sites (one created specially for the experiment, one based on `www.boeing.com`, and one based on `www.cs.waikato.ac.nz`). Only two of the tasks are described in this paper—the *parent revisitation task* and the *cross-site distant revisitation task*.

The *parent revisitation task* examined the subjects' ability to return to a parent (or home) page from a relatively deeply nested page. In the Boeing site the task was: "You are at the page describing the 747-international aircraft. Return to Boeing's Homepage." Both the stack and the temporal system allowed the users to solve the task with a single selection from the Back menu, but otherwise the stack system would require three Back clicks, while the temporal system would need either four or six (depending on earlier actions). We therefore predicted that the stack interface would be more efficient.

The *cross-site distant revisitation* task examined the subjects' ability to revisit a temporally distant page. An example of the task is: "You are at the page describing Boeing's 747-freighter commercial aircraft. Return to the Waikato University page 'Introduction to Computer Science' that you were at approximately half an hour ago." In theory, the temporal system is more efficient than the stack-based system because it allows the page to be selected directly from the Back menu. With the stack, however, the page would have been popped off the stack, requiring users to switch from backtracking with Back to navigation using links. Finally, if users used the cognitively lightweight 'click Back until the desired page is seen' strategy, then the temporal scheme would succeed, but only after between 12 and 27 clicks! We were uncertain which interface would allow more efficient completion, but were interested to observe the solution strategies used.
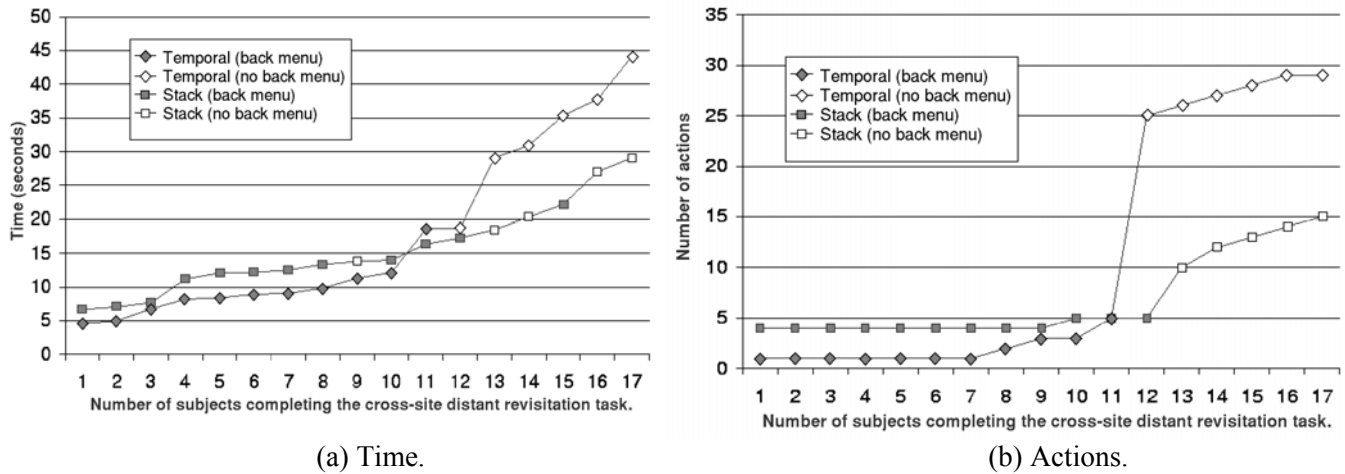
(a) Time.   (b) Actions.

Figure 4. Time and actions taken to complete the cross site distant revistation task using the stack and temporal interfaces. The x-axis is a sorted order of participants completing the task. Filled shapes show Back menu users; unfilled ones show users who did not use the Back menu.

Results

In the *parent revisitation task*, the mean task completion times with the stack and temporal interfaces were 6.2 (s.d. 2.2) and 10.6 (s.d. 6.2) seconds, providing a significant difference: unpaired T-Test, T_{100}=4.81, p<.01. Similarly, the mean action counts for the stack and temporal interfaces were significantly different at 2.4 (s.d. 1.3) and 4.1 (s.d. 2.8) actions: T_{100}=3.8, p<.01. As predicted, then, the longer paths with the temporal interface made it less efficient for returning to parent pages.

The *cross-site distant revisitation task* provided an interesting contrast between the interfaces. The results revealed highly varied solution strategies in both interfaces. Although the overall means were similar for both interfaces, some users solved the task quickly, while others struggled. The mean solution times for the stack and temporal interfaces were 15.5 (s.d. 6.5) and 16.2 (s.d. 12.0) seconds, and the corresponding action counts were 7.2 (s.d. 4.4) and 8.4 (s.d. 8.9). Neither of these comparisons provides a statistically significant difference. Analysing the ways in which participants solved the task provides interesting insights into opportunities for enhancing the efficiency of web browsers. The intention here is not to determine whether the two interfaces are statistically different from one another, but rather to observe patterns of behaviour that result in efficient and inefficient navigation.

Use of the Back menu (supported by both the normal stack and temporal interface) was strongly correlated with the efficiency of task completion (Spearman rank correlation between action count and menu use, r_s=0.83, p<.01). The mean number of moves used by the 23 participants who used the Back menu was 3.0 (s.d. 2.4), compared with 20.7 (s.d. 7.8) for the eleven participants that did not. Task times were similarly biased, with means of 11.0 seconds (s.d. 4.5) and 27.7 (s.d. 9.3) for menu and non-menu users. The correlation between action count and menu use holds for both the stack (r_s =0.89) and temporal (r_s =0.86) interfaces. Figures 4(a) and 4(b) show the time and number of moves (y-axis) used to complete the task plotted against a sorted order of participants (x-axis). It shows that the temporal system allowed most of the participants to solve the task rapidly (when compared to the stack users), but that the least efficient temporal users were dramatically slower the worst stack users. The figures also show that all eleven efficient users with the temporal system used the Back menu, and that none of the inefficient users did so.

One final problem, noted by many of the participants who used the Back menu in either system, was that web pages are often hard to identify from the text displayed in the Back menu. The text shown in the Back menu is extracted from the HTML <TITLE> tag associated with the page. As noted earlier, this information is often unavailable, or truncated, or poorly reflects the page contents. Other mechanisms for aiding page identification are discussed in the following section.
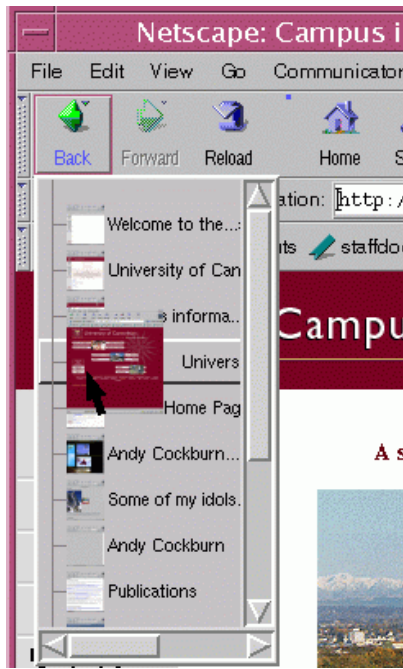
Figure 6. A version of the WebView prototype with zoomable thumbnails integrated into the Back/Forward menus.

tion task most rapidly when using the Back menu is
could be made more efficient if browsers encouraged
uch as the Back menu, rather than relying on multiple
lowing section.

VISITATION TOOLS

s to the temporal scheme that aim to improve the
k is motivated and informed by the findings of the

evisitation, inc... buttons,
hese tools ca... into one
of simple visu... web page

l WebView, t... epts. The
ows. First, th... rsions of
nplementation... complete
ociated with the... matically
scape responds to user actions in WebView. Second,
al list of previously visited pages, therefore integrating
vard buttons. Third, they provide zooming thumbnail
small thumbnails, which are automatically captured
/henever the user points the cursor over them. Fourth,

Figure 5. A "dogeared" thumbnail. The top-left dogear encodes visit counts, and the bottom-left dogear shows that the page is bookmarked.

the thumbnails are enhanced with additional identification cues including 'dogears' that encode information about the number of visits to a page and whether the page is bookmarked (see Figure 6).

Several systems have used thumbnails to depict web pages, including MosaicG (Ayers and Stasko 1995), PadPrints (Hightower et al 1998), and Glabster (Kaasten 2001). Recently Kaasten et al (2002) formally evaluated the effectiveness of various cues to web page identification, with results indicating that people are able to identify pages more accurately from thumbnail images than from titles or URLs. It seems reasonable to expect that automatically captured web page thumbnails, such as these, will soon appear in commercial web browsers.

Pilot study of WebView

To gain preliminary insights into WebView's usability, and to direct our development work, we performed a preliminary evaluation (Cockburn et al 1999). Seven volunteer participants, all graduate Computer Science students, received a very brief introduction to WebView's interface prior to completing a simple navigational task with WebView. They then repeated the task without WebView (using Netscape Navigator). The evaluated version of WebView was similar to that in shown in Figure 6, but displayed in a separate window rather than embedded within the Back menu.

To prepare for the task, the participants were asked to navigate from the "Computer Science Homepage" to a "Teaching Page", then back to the "Computer Science Homepage" and down through links to "People", then "Academic Staff", then "Wal's Page". Using the notation described earlier, this path is $a{\rightarrow}b{\Leftrightarrow}a{\rightarrow}c{\rightarrow}d{\rightarrow}e$, giving a stack content of {**a, c, d, e**}, and a temporal list content of {**b, a, c, d, e**}. Note that page **b** has been popped off the stack, but is displayed in WebView's temporal list. The timed task involved returning to the "Teaching Page" from "Wal's Page". With WebView, the page was immediately accessible on the temporal list, and could be revisited with a single click. With Netscape, however, the participants had to click Back three times to return to the "Computer Science Homepage" (or use the Back menu), and then select the link to the "Teaching Page".

All participants completed the task more rapidly with WebView, with mean times for Netscape and WebView of 9.6 seconds and 2.6 seconds respectively (paired T-Test, t_6=2.7, p<.05). Five of the seven participants were enthusiastic about the system, three of them extremely so, one stating "It'd be great to have a system like this." The two unenthusiastic participants were primarily concerned about the redundancy introduced by having two

different ways of navigating (the "Netscape way" and the "WebView way"). Integrating WebView's support into the browser, as shown in Figure 6, should overcome this concern.

We asked the participants to respond to the questions "Q1: If available, how likely would you be to use a system like WebView?" and "Q2: How useful were thumbnail images in recognising pages?", with responses rated on a five point scale from 1–"Not at all" to 5–"Very". The mean response to Q1 was relatively high at 3.9. We were surprised, however, that the mean response for the usefulness of the thumbnail images was low at 2.7. When asked to comment on the thumbnails, several participants stated that they would probably be more useful when navigating through unfamiliar sites. Participant 4 added that the thumbnails were not useful because the task was introduced verbally, using the same names for the pages as was displayed in the title text identifying each page.

One participant made an interesting observation that he "didn't like the thumbnails jumping about while navigating". When a new page is loaded, WebView redisplays the temporal list, causing all thumbnails in the display to move down one position. We had not anticipated this redisplay causing problems. Our current design solution is to use rapid animation to move items to their new location, rather than abrupt redisplaying them. Similar animated techniques are now common in many rapidly appearing/disappearing graphical user interface components, such as pop-up menus.

Summary

WebView demonstrates a variety of interface schemes aimed at rationalising and improving the diverse facilities for revisiting web pages currently used in browsers. The pilot study indicated that under favourable conditions these mechanisms could be successful, but extensive further work is necessary to determine how they scale to prolonged heavy use. This will be the focus of our further work.

OVERALL SUMMARY

This paper has provided an integration and synthesis of the findings from eight years of research on understanding and improving web browsers' support for revisiting web pages. The findings in five categories of revisitation research were presented: characterisations of user behaviour; system models of navigation and their impact on the user's understanding; interface methods for increasing the efficiency of the Back button; alternative system models for navigation; and alternative methods for presenting web navigation histories. Results and observations presented in the paper include the following:

- Web page revisitation is a dominant activity in web use, with more than 80% (on average) of web page visits being to pages previously visited by the user.
- Most users have one or two pages that they visit much more frequently than all other pages.
- Some users make very heavy use of bookmarks, and incur problems with managing them.
- Many users misunderstand the stack-based behaviour of the Back button, even though they use it regularly.
- Gesture navigation features in which users 'flick' the mouse to issue the common Back and Forward commands are both efficient and popular. They overcome the need to continually move the cursor away from the page links, or to home the hands between the mouse and the keyboard, when using Back. These features are now included in some commercial browsers.
- The evaluation of an alternative behaviour for the Back and Forward buttons suggests that revisitation can be improved by providing a complete temporally ordered list of previously visited pages. This observation has potentially important implications for the development of navigation features in the next generation of web browsers.
- The 'WebView' prototype system demonstrates how temporal lists could be used to integrate the diverse web revisitation tools in current browsers. It also demonstrates how the problems of page identification (prevalent in current browsers) can be eased through simple visualisation techniques.

REFERENCES

Abrams D, Baecker R, Chignell M, 1998 "Information Archiving with Bookmarks: Personal Web Space Construction and Organization", in *Proceedings of CHI'98 Conference on Human Factors in Computing Systems,* (Los Angeles: Los Angeles) pp 41--48

Ayers E, Stasko J, 1995 "Using Graphic History in Browsing the World Wide Web", in *Proceedings of the Fourth International World Wide Web Conference. 11--14 December.* (Boston: MA: Boston: MA)

Byrne M, John B, Wehrle N, Crow D, 1999 "The Tangled Web We Wove: A Taskonomy of WWW Use", in *Proceedings of CHI'99 Conference on Human Factors in Computing Systems* (Pittsburgh, May 15--20: Pittsburgh, May 15--20) pp 544--551

Catledge L, Pitkow J, 1995 "Characterizing Browsing Strategies in the World Wide Web", in *Computer Systems and ISDN Systems: Proceedings of the Third International World Wide Web Conference. 10--14 April.* (Darmstadt, Germany: Darmstadt, Germany) pp 1065--1073

Chi E, Pirolli P, Pitkow J, 2000 "The Scent of a Site: A System for Analyzing and Predicting Information Scent, Usage, and Usability of a Web Site", in *Proceedings of CHI'2000 Conference on Human Factors in Computing Systems, April 1--6* (The Hague, The Netherlands: The Hague, The Netherlands) pp 161-168

Cockburn A, Greenberg S, 2000 "Issues of Page Representation and Organisation in Web Browser's Revisitation Tools" *Australian Journal of Information Systems* **7** 120--127

Cockburn A, Greenberg S, McKenzie B, JasonSmith M, Kaasten S, 1999 "WebView: A Graphical Aid for Revisiting Web Pages", in *Proceedings of OzCHI'99: Computer Human Interaction Specialist Interest Group of the Ergonomics Society of Australia. November 28--30.* (Wagga Wagga.: Wagga Wagga.) pp 15--22

Cockburn A, Jones S, 1996 "Which way now? Analysing and Easing Inadequacies in WWW Navigation" *International Journal of Human-Computer Studies* **45.** 105--129

Cockburn A, McKenzie B, 2001 "What Do Web Users Do? An Empirical Analysis of Web Use" *International Journal of Human-Computer Studies* **54** 903--922

Cockburn A, McKenzie B, 2002 "Pushing Back: Evaluating a New Bahaviour for the Back and Forward Buttons in Web Browsers" *International Journal of Human-Computer Studies* **In Press**

Fitts P, 1954 "The Information Capacity of the Human Motor System in Controlling the Amplitude of Movement." *Journal of Experimental Psychology* **47** 381-391

Greenberg S, Cockburn A, 1999 "Getting Back to Back: Alternate Behaviors for a Web Browser's Back Button", in *5th Conference on Human Factors and the Web, June 3.* (Gaithersburg, Maryland.: Gaithersburg, Maryland.)

Hightower R, Ring L, Helfman J, Bederson B, Hollan J, 1998 "Graphical Multiscale Web Histories: A Study of PadPrints", in *Proceedings of the 1998 ACM Conference on Hypertext, June 20--24.* (Pittsburgh, Pennsylvania.: Pittsburgh, Pennsylvania.) pp 58--65

Kaasten S, 2001 *Integrating Back, History and Bookmarks in Web Browsers* Masters, University of Calgary, Calgary

Kaasten S, Greenberg S, Edwards C, 2002 "How People Recognize Previously Seen Web Pages from Titles, URLs and Thumbnails", in *People and Computers XVI (Proceedings of Human Computer Interaction 2002)* Eds J Faulkner, J Finlay and F Detienne pp 247-265

Kehoe C, Pitkow J, 1996 "Surveying the Territory: GVU's Five WWW User Surveys" *The World Wide Web Journal* **1** 77--84

Kurtenbach G, Buxton W, 1991 "Issues of Combining Marking and Direct Manipulation Techniques", in *Proceedings of ACM UIST'91* pp 137--144

Mayo E, 1933 *The Human Problems of an Industrial Civilization* (Harvard University Press: Cambrigde, MA.)

Moyle M, Cockburn A, 2002 "Gesture Navigation: An Alternative Back for the Future", in *Extended Abstracts of CHI'2002 Conference on Human Factors in Computing Systems, 20--25 April* (Minneapolis, Minnesota: Minneapolis, Minnesota)

Nielsen J, Levy J, 1993 "Measuring Usability: Preference versus Performance" *Communications of the ACM* **37** 66--75

Pirolli P, Pitkow J, Rao R, 1996 "Silk from a Sow's Ear: Extracting Usable Structures from the Web", in *Proceedings of CHI'96 Conference on Human Factors in Computing Systems, April 13--18* (Vancouver: Vancouver) pp 118--125

Pitkow J, 1996 "GVU's WWW User Surveys"

Robertson G, Czerwinski M, Larson K, Robbins D, Thiel D, Dantzich M v, 1998 "Data Mountain: Using Spatial Memory for Document Management", in *Proceedings of UIST'98 ACM Conference on User Interface*

Software and Technology, November 1--4. (San Francisco, California.: San Francisco, California.) pp 153--162

Tauscher L, Greenberg S, 1997 "How People Revisit Web Pages: Empirical Findings and Implications for the Design of History Systems" *International Journal of Human Computer Studies, Special issue on World Wide Web Usability* **47** 97--138