Beyond the 'Back' Button: Issues of Page Representation and Organisation in Graphical Web Navigation Tools

Andy Cockburn
Department of Computer Science
University of Canterbury
Christchurch, New Zealand
andy@cosc.canterbury.ac.nz

Saul Greenberg
Department of Computer Science
University of Calgary
Calgary, Alberta, Canada
saul@cpsc.ucalgary.ca

Abstract

Although the 'Back' button is good for revisiting very recently seen pages on the world-wide web, its recency and stack-based model makes it inefficient for navigating back to distant pages. The limitations of 'Back' have motivated researchers and developers to investigate graphical aids for web browsing. This paper examines the design and usability issues in two fundamental questions that all graphical tools for web-navigation must address: first, how can individual pages be represented to best support page identification?; and second, what display organisation schemes can be used to enhance the visualisation of large sets of previously visited pages? Our 'webView' graphical browsing system, which interacts with unaltered versions of Netscape Navigator, demonstrates new interface techniques for page representation and display organisation. WebView's page identification techniques included zoomable thumbnail images and a 'dogears' metaphor that offers a lightweight mechanism for bookmarking. Its display is organised using an integrated hybrid of three techniques: 'hub-and-spoke', which models the user's navigation within a site; 'site-maps', which model navigation between sites; and temporal organisation, which provides a recency ordered list of the visited sites.

Keywords:

World-Wide Web navigation, history and revisitation, graphical navigation maps.

1 Introduction

Web browsers have rapidly become an almost essential component of the desk-top computing environment. They provide access to local and global documents and news, digital libraries, powerful search facilities, and interactive pages let us conveniently shop from home. Navigating the web has become such a central paradigm for human-computer interaction that Microsoft has made its 'Internet Explorer' web browser the main interface for file management.

Despite this incredible popularity, the user interfaces to web browsers are far from ideal. Previous studies clearly show that users have a poor understanding of the most rudimentary navigational interface components (Cockburn & Jones 1996) such as the `Back` and `Forward` buttons. When considering that the `Back` button is almost certainly the most widely used graphical user interface component ever created (accounting over 30% of all web-browser user actions (Tauscher & Greenberg 1997)) it is clear that even the smallest inefficiency or confusion in its use will escalate to massive loss of productivity when multiplied across millions of users.

In recent related work (Greenberg & Cockburn 1999) we scrutinized the behaviour of linear navigation schemes that require only `Back` and `Forward` buttons: both the *actual* behaviour in browsers such as Netscape and Microsoft Internet Explorer, and the *possible* ways that the buttons could operate to improve revisitation. Although these simple interfaces have advantages they also have major limitations, including inefficiency in large data-sets, limited scalability, and poor support for orienting the user in the information space (Section 2.2). These limitations have motivated researchers to investigate graphical navigational schemes—see Cockburn & Jones (1997) for a review of systems supporting graphical web browsing.

In this paper we investigate user-interface considerations that influence the design and use of graphical mechanisms for web-browsing and page revisitation. In particular, we consider two questions that must be addressed by the designers of such systems. First, what interface mechanisms can be used to identify and distinguish each individual previously visited page?, and second, what structural arrangements can be used to enhance the visualisation of a large set of previously visited pages? The relative merits of each design alternative are discussed, and related systems are used to exemplify techniques. We also describe our graphical browsing system, webView, which we use to experiment with, and demonstrate, new design alternatives.

The structure of the paper is as follows. Section 2 provides background information on page revisitation on the web and precisely identifies limitations in linear revisitation schemes. Section 3 provides an introduction to the general problems associated with graphical mechanisms for web-browsing and page revisitation. Section 4 then addresses the question of how to graphically represent individual pages so that they can be readily identified and distinguished from other pages. Section 5 investigates the related question of how to organise the display of multiple pages. Our webView system is described in Section 6. Section 7 summarises the paper and identifies directions for our further work.

2 Revisitation on the web with linear navigation

In this section we discuss why page-revisitation is a fundamental activity in web navigation, and we examine the limitations of linear revisitation schemes.

2.1 Data on web-page revisitation

Tauscher & Greenberg (1997) analysed patterns of web revisitation by logging the user actions of twenty three users with a specially equipped version of XMosaic2.6. During a six week period they found that 58% of page visits were to pages that the user had previously seen. They also found a strong recency effect, where the most recently visited pages were most likely to be visited next. Figure 1 plots the cumulative percentage URL recurrence rate as a running sum against the history distance between pages visited. For example, there is a 39% chance that the next page will be within the six most recently visited pages. The study also indicates that Back is the dominant interface mechanism used for page revisitation, accounting for more than 30% of all navigational acts (other studies report that it accounts for 37%; Catledge & Pitkow (1995)). Other revisitation facilities were used infrequently: for example, less than 3% for bookmarks, and less than 1% for history systems.

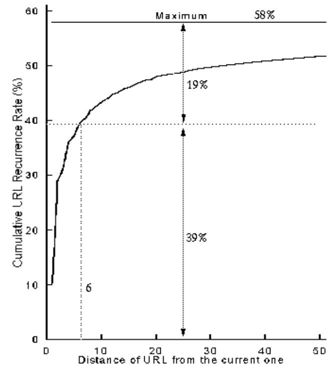

Figure 1: Cumulative URL recurrence rate against history distance between pages (extracted from Tauscher and Greenberg, 1997).

Tauscher & Greenberg's study provides powerful evidence that page revisitation is a fundamental requirement in web navigation. It also shows that while many revisitation paths are back to pages that have been seen recently, there is also a need to return to more distant pages. For example, Figure 1 shows that there is a 19% chance that the next page the user visits will be more than six pages away from the current page (measured in terms of recency).

2.2 Advantages and Limitations of linear navigation

A variety of linear navigation schemes can be implemented using only the web-browser's Back and

`Forward` buttons, as described in Greenberg & Cockburn (1999).

Although each of these techniques has relative merits and disadvantages with respect to its counterparts, in general there are two major advantages of page revisitation based on the `Back` and `Forward` buttons. First, it is cognitively undemanding in that there is little need for decision-making—users need only repeatedly click the `Back` button until either the desired page is displayed or until the start of the `Back` list is reached. Second, it is visually compact—the two buttons consume minimal screen real-estate, and for this reason it is reasonable for users to keep them on permanent display within their browsers; consequently they are always ready to hand.

The limitations of linear schemes can be generalised as follows.

Inefficiency in large data-sets. Revisiting distant pages is laborious and time consuming, requiring many clicks of the `Back` button and the associated redisplay of many pages in the browser[1].

Limited access. The stack-based schemes supported by current commercial browsers only allow a sub-set of pages to be revisited because of the branch-pruning that occurs when pages are popped off the top of the stack (Cockburn & Jones 1996).

Scalability. Current commercial browsers, such as Microsoft Internet Explorer, are incorporating file-management capabilities into their functionality. Consequently, the range of data sources accessed through 'web'-browsers is likely to dramatically increase, as will the temporal range of data elements that the user will want to revisit.

Context. Linear schemes give the user a single viewpoint—the current page—into the information space through which they are navigating. The user's sense of orientation within the information space is therefore entirely dependent on their memory of previously visited pages and the contents of the current page.

Honesty of representation. Linear lists model the user's actions in the web in one-dimension, yet the underlying information space is more honestly modeled in two or three dimensions.

[1] Web-browsers now support short-cut menus associated with the `Back` and `Forward` buttons, but we consider these techniques to be 'graphical' mechanisms, as discussed later in the paper.

3 Graphical schemes for page revisitation

Graphical overview maps have long been used to help users overcome the disorientation associated with navigating through hypertext (Nielsen 1990). In this paper we use a broad interpretation of 'graphical schemes' to include any technique that provides a representation of multiple pages, such as menus, visual lists, and 2D or 3D renderings of web-spaces.

Each of the limitations of linear navigation schemes can *potentially* be eased or resolved through graphical visualisations of web spaces. The ability of graphical displays to represent millions of individual data points addresses the issues of scalability and limited access. The problems of inefficiency in large data sets can be reduced (provided the user can identify pages in the graphical representation) because pages can be directly accessed from their representations. Finally, graphical overviews can reveal the structure and relationship between pages, helping users orient themselves in the information space and overcoming the problems of context and honesty of representation.

Even if this *potential* could be attained, there remains a fundamental problem confounding the use of graphical schemes for web navigation. Users make continual trade-offs between the screen real-estate consumed by applications and the value of the information provided. If users can successfully navigate some or most of the time without graphical aids, they are likely to remove, iconify, or bury the graphical window. Once the graphical aid is no longer visible it becomes overhead to redisplay it: the support it offers is no-longer ready to hand. Graphical browsing aids must therefore maximise the value of the information they present, and do so using minimal screen real-estate. The following subsections address two questions that focus on this issue. First, how can each page be best represented in a graphical display. Second, how can the display be organised to best reveal the relationships between pages.

4 Issues in graphical page representation

The ideal page representation in a graphical display would allow immediate and unique identification of the page while consuming only minimal screen real-estate. Page identification will be aided if it is displayed within the context of related neighbouring pages, but in this section we focus purely on techniques that can be used to identify individual pages.

Three categories of techniques for page identification are presented in the following sections: images of the rendered page, textual representations of the page, and abstract properties of the page. Many of the techniques discussed can be combined to reinforce page identification, as demonstrated by the webView system (Section 6).

4.1 Images of the rendered page

Miniaturized 'thumbnail' page images can be automatically captured as each page is rendered in the web-browser. At low scaling factors (large thumbnails), pages are readily identifiable from thumbnails, but image quality degrades as size is reduced. Figure 2 compares the effect with a page that contains several large distinctly coloured regions (top row), and with a page that is almost entirely text (bottom row).

The distinct coloured blocks in Figure 2 make the page identifiable at small sizes, but even at large sizes Figure 2 provides few visual aids to page identification. The problems are further exacerbated by the consistent page-styles (coloured backgrounds, frames, banners, and so on) used within web-sites. Figure 3 shows the poor distinction between four different pages within the Calgary Computer Science web-site with thumbnails that are 50 pixels wide. In reality, constraints on screen real-estate make icons of even this modest size unlikely to proliferate: for example, the icons used in file browsers such as Microsoft's Explorer are approximately 15 by 10 pixels.

Zooming techniques can be used to ease the problems of low distinction between pages at small sizes. MosaicG (Ayers & Stasko 1995) allows users to zoom in and out of its display of page thumbnails. PadPrints (Hightower, Ring, Helfman, Bederson & Hollan 1998) and webView (Section 6) support discrete sizes of thumbnails: the small images are magnified as the user moves the mouse over them. The Data Mountain (Robertson, Czerwinski, Larson, Robbins, Thiel & van Dantzich 1998) provides a 3D visualisation of a large set of web page thumbnails. When the user selects a page, an animated zooming technique magnifies it and moves it forward into the user's focus. The Data Mountain also uses a variety of animation cues to reinforce the user's awareness of the spatial organisation of thumbnail images (see Section 5). The pad++ (Bederson & Hollan 1994)

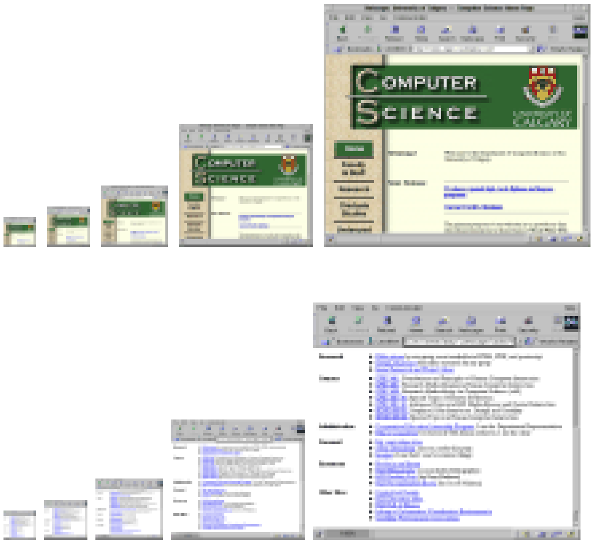

(f) 12. (g) 16. (h) 25. (i) 50 pixels wide. (j) 100 pixels wide.

Figure 2: The Calgary Computer Science Homepage (`www.cpsc.ucalgary.ca`) and a text page at varying pixel widths.

Figure 3: Four pages within the Calgary Computer Science site at 50 pixels wide.

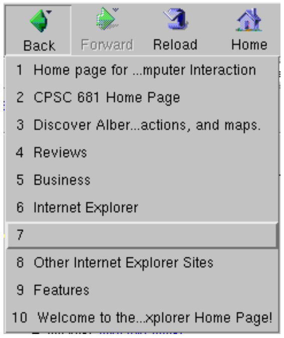

Figure 4: Netscape Navigator's Back menu, showing text extracted from <Title> tags.

zoomable web browser described in Bederson, Hollan, Steward, Rogers, Vick, Ring, Grose & Forsythe (1999) integrates browser functionality within each of the miniaturised page representations, allowing smooth transitions between thumbnail images and browsing actions. WebView automatically captures a zoomable thumbnail image of each page displayed in Netscape (Section 6.1).

4.2 Textual page representations

Several text-based mechanisms can be used to assist page identification. These include the page URL, the <Title> and <H1> text in the page's HTML source, samples of text from the body of the page, and explicit user annotations.

URLs are primarily machine-oriented. While some site addresses can be readily parsed and recognised by users (for example, `netscape.com`, `microsoft.com`, `acm.org`), not all sites are meaningful and lower level paths to specific pages are of limited use.

Titles, extracted from the <Title> tag in the page's source, are are often poor identifiers of page contents. Problems with titles affect a wide range of navigation aids including the 'Back' and 'Forward' menus supported by Netscape and Microsoft Internet Explorer (Figure 4), their history lists, and graphical overview maps. Titling problems include the following.

Erroneous and missing titles. Many pages provide little or no useful information about the contents of the page, even on pages that are are otherwise highly professional. For example, the blank entry '7' in Figure 4 resulted from an absent title in the page offering 'More Info' on Microsoft's Internet Explorer `www.microsoft.com/ie/info/default.asp`. It appears that many page authors are less concerned about the accuracy of title tags (which are normally 'hidden' in the title bar of the window) than they are about the presentation and information contents of the main window.

Excessive/insufficient title consistency. It is common to find a set of pages with identical or almost identical titles. Sites that use HTML frames are particularly problematical because title information remains constant while the user navigates through pages contained within the frame.

An absence of consistency in title format can also frustrate page identification. Entries '1' and '2' in Figure 4 demonstrate poor title-format consistency in undergraduate course home pages entitled 'CPSC 681 Home Page' and 'Home page for the course CPSC481: Foundations and Principle of Human Computer Interaction'. Although these titles accurately describe the page contents, the lack of consistency in their presentation complicates and slows page identification.

Long titles and truncation. Many graphical mechanisms constrain the number of characters that can be fit into the available display space. Truncation is used to shorten the title length, , but this can remove essential information for page identification. Entries '1', '3' and '10' in Figure 4 demonstrate truncation, and in entry '1' the major page identification cue 'CPSC 481' has been removed from the displayed title.

Other text sources. Heuristic techniques could attempt to extract text titles from a variety of sources other than the <Title> tag, such as from the major heading text for the page, or from the first few sentences of the body of the page. These techniques are also problematical, for instance <H1> is often intentionally omitted in pages that rely on graphical banners for page headings.

Explicit user annotations. Although explicit user-supplied annotations are likely to be extremely accurate and descriptive (for that particular user), it is unlikely that users will be willing to supply annotations for anything other than the most frequently

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.

WHAT WILL YOU BUILD? | **sales@docketalarm.com** | **1-866-77-FASTCASE**

fastcase®
Smarter legal research.