Cockburn, A. and Greenberg, S. (1999)
Beyond the 'Back' Button: Issues of Page Representation and Organisation in Graphical Web Navigation Tools. Report 1999-640-03, Dept of Computer Science, University of Calgary, Calgary, AB., Canada.

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.

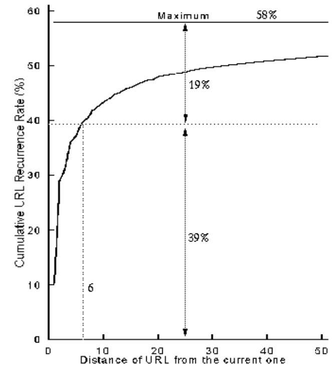

Figure 1: Cumulative URL recurrence rate against history distance between pages (extracted from Tauscher and Greenberg, 1997).

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)

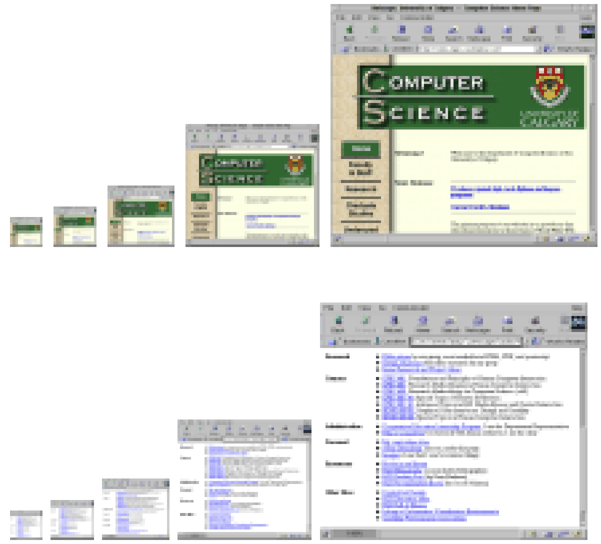

(f) 12. (g) 16. (h) 25. (i) 50 pixels wide. (j) 100 pixels wide.

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 visited pages.

WebView extracts a title from the `<Title>` tag when available, and uses the page URL when unavailable. Users can choose to have both the URL and the title text displayed simultaneously. The important issue, however, is that the text associated with the page representation is only one of several cues for page identification.

4.3 Abstract page properties

Abstract page properties can be used to reinforce page identification techniques. In particular, temporal information about the frequency and timing of previous visits may help users to identify pages. Netscape and Microsoft Internet Explorer both use temporal information in their history views of past visits.

Graphical representations of 'dogears' provide an appealing metaphor that can be used to capture several identification cues. Explicit dogears—based on a metaphor of folding down the corner of a book page—allow a natural means for marking pages. Implicit dogears—automatically created and modified—can be used to reveal the number of visits to a page. Further iconic mechanisms can be used to reveal the timing of first and most recent access to the page. Figure 5 shows a webView thumbnail page image that includes dogear and iconic identification clues. The red dogear at the top-left of the thumbnail is an explicit dogear that the user can add by clicking to 'bookmark' the page. The green shading in the dogear in the top-left of the thumbnail becomes denser with each visit to the page. The blue flashes on the right-hand side of the dogear indicate the time of the first and last access to the page, where the bottom of the thumbnail equates to the eldest page in the display and the top of the thumbnail equates to the current time. WebView's dogears are further described in Section 6.

Figure 5: WebView's dogear metaphor for capturing bookmarks, temporal properties of page visitation, and the visit count.

5 Issues in structural organisation of page display

This section discusses issues in organising the display of many pages. Four schemes are described: hub-and-spoke techniques that depict the user's navigational branching actions; spatial schemes that exploit the user's memory for the spatial location of objects; site-maps that model the physical storage location of pages; and temporal schemes. The relative merits of each organisational scheme are discussed, and systems demonstrating each technique are noted. Hybrid schemes can also be used, as demonstrated by the webView prototype described in the following section.

5.1 Hub-and-spoke dynamic trees

Hub-and-spoke dynamic trees are generated in response to the user's navigational acts. Each time a new page is visited it is added to the display, together with a link showing the connection between the previous page and the new one. Various graph drawing algorithms can be used to generate the display (with consideration of cyclic page links), resulting in tree-like displays (Figure 5), nested displays (Figure 6), and 'wheel-like' hub-and-spoke displays (Cockburn & Jones 1996). Tree and nesting displays clearly depict the hierarchical relationship between parent pages and sibling links.

The primary advantage of hub-and-spoke dynamic trees is that they cluster logically related pages, independent of their physical location. There are two main disadvantages of hub-and-spoke dynamic trees. First, because each branching navigational act is treated as a new spoke of the current page, there is restriction on the nesting level of the display. Second, hub-and-spoke browsing focuses on *re*-visitation, and offers little or no support for initial page access. WebView eases the first problem by creating a new top-level branch in the tree for each new site added to the display. The hub-and-spoke connection between pages at different sites is displayed using graphical links. WebView also partially eases the second problem by parsing the HTML contents of each visited page, allowing shortcut navigation to the spokes emanating from any current page. Robert & Lecolinet's system (1998) went one-stage further by pre-fetching pages to a pre-specified depth. Clearly there are performance costs incurred by pre-fetching large sets of pages.

Systems demonstrating hub-and-spoke dynamic trees include MosaicG (Ayers & Stasko 1995), webNet (Cockburn & Jones 1996), a pad++ based browser (Bederson et al. 1999) and padPrints (Hightower et al. 1998). WebView's implementation of hub-and-spoke representation is described in Section 6.

5.2 Spatial layouts

Spatial organisations aim to exploit people's memory for the spatial location of objects. Users explicitly position page representations (such as thumbnails) in a 2D or 3D space, and the location then serves as an additional cue to page identification.

Robertson et al. (1998) show that with a dataset of 100 pages, the Data Mountain—a 3D spatial organisation of pages—provides statistically significant improvements in the time and error rates of page access when compared with Internet Explorer's favourites utility.

The primary disadvantage of spatial techniques is that users must explicitly position pages within the information space, placing both a cognitive and manipulative burden on the user. First they must decide how to categorise the page, and consequently decide where the page should be located. Mackay's (1988) study of email usage suggests that while some users naturally categorize information, others do not. Having categorised the page and decided where it should be located, the user must move it to the desired location. The combination of these activities may be unacceptable to users who wish to simply glance at a page and go onto the next. If pages assume a default location to free the user from having to explicitly select a location, there is a risk that the default region may become excessively cluttered, reducing the user's ability to locate specific pages.

Systems demonstrating spatial organisation techniques include WebForager (Pitkow & Bharat 1996) and Data Mountain (Robertson et al. 1998).

5.3 Site maps

Site maps show a topology of the physical storage locations of pages. They are normally generated statically—as a batch process prior to the user's navigation to the site—and once generated, they can be made available to all site visitors. Page representations within site-maps are normally active, allowing users to navigate directly to pages by clicking.

The primary advantage of site-maps is their coverage. All pages at a site can be included in the site-map, aiding the user's orientation and navigation across the entire site. Statically generated site-maps can support initial access to pages as well as revisitation to pages previously seen.

Although the wide coverage of site-maps offers advantages, it also means that the representation is likely to show many more pages than the user needs, thus unnecessarily cluttering the display. There are other disadvantages of basing the site-map topology on the physical location of pages: in particular, the physical location of pages need not reflect semantic relationships, and pages on the same topic at different sites will be excluded from the display.

Systems supporting site-maps include the hyperbolic browser (Lamping, Rao & Pirolli 1995), the Navigational View Builder (Mukherjea & Foley 1995), WebViz (Pitkow & Bharat 1996). WebView uses partial site-maps to ease the problem of excessive nesting that occurs with hub-and-spoke organisation (Section 6).

5.4 Temporal organisation

Temporal organisation schemes exploit the user's memory for the timing of their actions. For recently accessed pages, a variety of linear `Back` and `Forward` revisitation schemes are efficient and easy to use (Greenberg & Cockburn 1999). For temporally distant page access, time segments can be used. For example, Microsoft Internet Explorer's history mechanism allows users to view the last few days, weeks, and months. Many pages are likely to be present within coarse-grained temporal segments such as 'last month', and it is therefore necessary to use secondary categorisation techniques to assist page identification.

WebView's display of visited pages includes several clues about temporal activity. The dogears associated with each page include clues about the frequency of page access (see Figure 5), and its display is organised to provide a temporal ordering of visited sites, assuring that the current site is always at a consistent location in the display.

5.5 Display organisation: summary

The capabilities and relative merits of the four display organisation techniques described are summarised in Table 1.

6 The webView prototype

Our investigation of graphical web navigation has been assisted by the development of prototype systems. The webView prototype shown in Figure 6 responds to and controls navigation in Netscape Navigator[2]. The following sections discuss webView's mechanisms for assisting individual page identification, for organising the display of many pages, and for supporting navigation shortcuts.

6.1 Page identification

Page identification in webView is aided by a combination of text-titles, zoomable thumbnails and dogears.

When a <Title> tag can be extracted from HTML, it is shown alongside the page thumbnail, otherwise the page URL is shown. The 'View' menu allows the user to select whether the title and/or the page URL are shown in the display.

WebView captures a thumbnail image of each new page once it is loaded into Netscape. The default thumbnail size is 25 pixels high, which in our experience is close to the minimum size for page identification. When the user moves the mouse over a thumbnail it is immediately magnified to a height of 64 pixels which provides a powerful identification cue. Figure 7(a) shows the zoomed 'Insects' thumbnail.

[2]Interaction with unaltered versions of Netscape is enabled by Netscape's remote control facilities `home.netscape.com/newsref/std/x-remote-proto.html`.

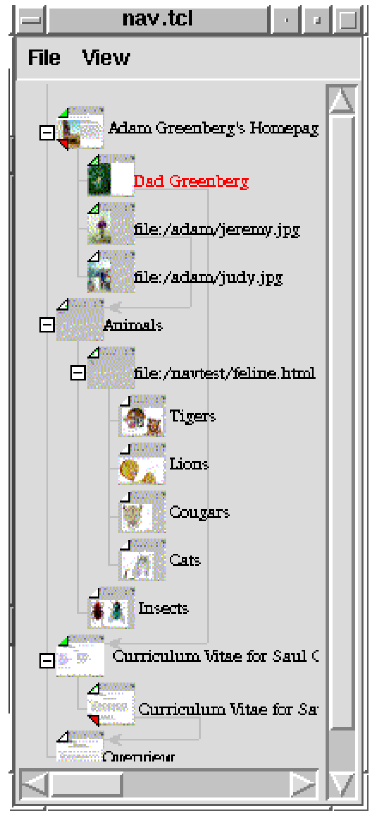

Figure 6: The webView prototype.

Explicit and implicit dogears further aid page identification. Thumbnail images can be explicitly marked with a red dogear in the bottom left corner by clicking with the middle mouse-button (for example, the 'Insects' thumbnail in Figure 7(a)). The dogear in the top-left corner of each thumbnail indicates the visitation count for each page, becoming progressively denser green with each revisitation to the page.

6.2 Page organisation

WebView's page display organisation is a hybrid of hub-and-spoke, site-map and temporal techniques. The tree representation is dynamically generated to reflect the user's navigational acts within the browser. The current page always appears within the top-most branch of the tree, and is identified by colouring the page's text red.

When the user navigates to a new page, an al-

Issue	Hub-and-spoke	Spatial	Site-map	Temporal
When generated?	Dynamically. Models user's actions.	Dynamically. Explicitly positioned by user.	Statically (or dynamically to a specified link-depth).	Dynamically.
Coverage	All visited pages.	All visited pages.	All site pages.	All visited pages.
Support for initial page access?	✗ ? (spoke shortcuts in webView)	✗	✔ (within site)	✗
Support for page revisitation?	✔	✔	✔ ✗ (site only)	✔
Display tailored to user's interest?	✔	✔	✗ (models site, not actions)	✔
Semantic groupings of related pages in display?	✔	✔ (explicit by user)	✔ ? (physical location need not model semantic relationships)	? ✔ (assuming user visits semantically related pages in sequence).
Primary advantage	Reveals browsing path within and across sites.	Location provides a cue to page identification.	Complete site coverage supports orientation within site.	Recency is a clue to page identification.
Primary disadvantage	Excessive display nesting with long browsing paths.	Explicit action necessary to position page representation.	Reveals a superset of 'interesting' pages.	Secondary structure necessary for temporally distant pages (e.g. 'last month').

Table 1: Summarising the capabilities and merits of display organisation techniques.

gorithm determines where it is inserted into the graphical tree. First, the site of the new page is compared with the site of the previous page. If they are the same then the new page is inserted as the first sibling of the previous page. If the sites are different then a check is made to see whether any page in the display is located at the same site as the new page. If so, the new page is inserted as the first sibling of the top-level page at that site. If not, the new page is inserted as a new top-level branch in the tree. The final action of the algorithm is to take the entire branch that contains the new page, and move it to the top of the graphical tree display to maintain a temporal-ordering of sites.

Through this algorithm, the graphical tree provides a clear depiction of the hub-and-spoke relationships between pages located at the same site, while avoiding the problem of continually increasing the nesting depth. In order to represent cross site hub-and-spoke relationships, a line is drawn between the previous and new pages when they reside at different sites. For example, the line and arrow in Figure 7(b)) shows that the user navigated to the 'Curriculum Vitae for Saul Greenberg' page from the 'Dad Greenberg' page. The movement of the site-branch containing the current page to the top of the tree maintains a temporal ordering of visited sites, and assures that the current page can always

(a) Zooming a thumbnail.

(b) Different site hub-and-spoke links.

be seen near the top of the display. Sites that the user has not navigated to for some time disappear off the bottom of the window, but they can be revisited by scrolling.

6.3 Navigation shortcuts

As well as responding to user actions that are initiated at the web-browser, webView can be used to control page display in Netscape, allowing several navigational shortcuts.

Clicking the text identifier of any page in webView causes Netscape to immediately navigate to that page, allowing revisitation to any page in the display. WebView also allows short-cut navigation to spoke links emanating from any page. Clicking any thumbnail (not only the current page) pops-up

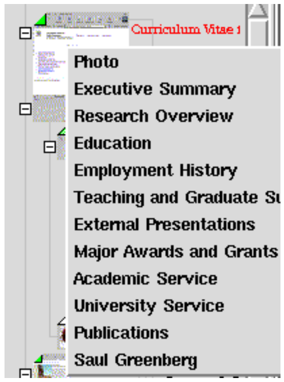

Figure 7: Thumbnail spoke menu.

a menu of the page's spoke links, and selecting one of the links causes Netscape to display that page. Figure 7 shows the spoke-menu for a Curriculum Vitae page. To avoid extremely long menus being generated for pages that contain many links, at most the first twenty page links are included in the menu.

7 Summary & Future Work

Interactive graphical tools have the potential to ease many user problems in navigating through the world wide web: they are scalable and can depict extremely large data-spaces, they can support revisitation to all previously seen pages, and they can aid the user's sense of orientation in the web-space. To realise this potential the tools must represent individual pages in a manner that aids rapid identification, and they must organise the display of large sets of pages in a manner that is meaningful to the user.

This paper has discussed the design and usability issues in providing identifiable representations of individual pages, and in organising the display of many pages. The webView prototype demonstrates novel techniques in both of these areas. Its zoomable thumbnail images and 'dogears' metaphor aid page identification, and its display organisation is implemented through a hybrid of hub-and-spoke, site-map, and temporal organisation schemes that promises to reduce many of the problems of the individual techniques.

Further work will focus on evaluating webView's page representation and organisation techniques.

References

Ayers, E. & Stasko, J. (1995), Using graphic history in browsing the world wide web, *in* 'Proceedings of the Fourth International WWW Conference. 11–14 December, Boston'.

Bederson, B. & Hollan, J. (1994), Pad++: A zooming graphical interface for exploring alternate interface physics, *in* 'Proceedings of UIST'94'.

Bederson, B., Hollan, J., Steward, J., Rogers, D., Vick, D., Ring, L., Grose, E. & Forsythe, C. (1999), A zooming web browser, *in* Ratner, Grose & Forsythe, eds, 'Human Factors in Web Development', p. To appear.

Catledge, L. & Pitkow, J. (1995), Characterizing browsing strategies in the world wide web, *in* 'Computer Systems and ISDN Systems: Proceedings of the Third International WWW Conference. 10–14 April, Darmstadt, Germany', Vol. 27, pp. 1065–1073.

Cockburn, A. & Jones, S. (1996), 'Which way now? Analysing and easing inadequacies in WWW navigation', *International Journal of Human-Computer Studies* **45.**(1), 105–129.

Cockburn, A. & Jones, S. (1997), Design issues for world wide web navigation visualisation tools, *in* 'Proceedings of RIAO'97: The 5th Conference on Computer-Assisted Research of Information. Montreal, June 1997', pp. 55–74.

Greenberg, S. & Cockburn, A. (1999), 'Getting back to back: Alternate behaviors for a web browser's back button', Submitted to the 5th Conference on Human Factors and the Web.

Hightower, R., Ring, L., Helfman, J., Bederson, B. & Hollan, J. (1998), Graphical multiscale web histories: A study of padprints, *in* 'Proceedings of Hypertext'98, 20–24 June. Pittsburgh.', ACM Press, pp. 58–65.

Lamping, J., Rao, R. & Pirolli, P. (1995), A focus+context technique based on hyperbolic geometry for visualising large hierarchies, *in* 'Proceedings of CHI'95 Denver, May 7–11', pp. 401–408.

Mackay, W. (1988), More than just a communication system: Diversity in the use of electronic mail, *in* 'Proceedings of CSCW'88, Portland, September 26–28', pp. 344–353.

Mukherjea, S. & Foley, J. (1995), Visualizing the world wide web with the navigational view builder, *in* 'Computer Systems and ISDN Systems: Proceedings of the Third International World Wide Web Conference. 10–14 April, Darmstadt, Germany', Vol. 27, pp. 1075–1087.

Nielsen, J. (1990), 'The art of navigating through hypertext: Lost in hyperspace.', *Communications of the ACM* **33**(3), 296–310.

Pitkow, J. & Bharat, K. (1996), Webviz: A tool for world-wide web access log analysis, *in* 'Proceedings of the First International WWW Conference'.

Robert, L. & Lecolinet, E. (1998), Browsing hyperdocuments with multiple focus+context views, *in* 'Proceedings of ACM Hypertext'98, 20–24 June. Pittsburgh.', ACM Press, pp. 293–294.

Robertson, G., Czerwinski, M., Larson, K., Robbins, D., Thiel, D. & van Dantzich, M. (1998), Data mountain: Using spatial memory for document management, *in* 'Proceedings of UIST'98, November 1–4. San Francisco.', ACM Press, 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* **47**(1), 97–138.