Design Issues for World Wide Web Navigation Visualisation Tools

Andy Cockburn[†]

Department of Computer Science
University of Canterbury
Christchurch, New Zealand
andy@cosc.canterbury.ac.nz
+64 3 364 2987 ext 7768

Steve Jones[†§]

Department of Computer Science
University of Waikato
Hamilton, New Zealand
stevej@cs.waikato.ac.nz
+64 7 838 4490

Abstract

The World Wide Web (WWW) is a successful hypermedia information space used by millions of people, yet it suffers from many deficiencies and problems in support for navigation around its vast information space. In this paper we identify the origins of these navigation problems, namely WWW browser design, WWW page design, and WWW page description languages. Regardless of their origins, these problems are eventually represented to the user at the browser's user interface. To help overcome these problems, many tools are being developed which allow users to visualise WWW subspaces. We identify five key issues in the design and functionality of these visualisation systems: characteristics of the visual representation, the scope of the subspace representation, the mechanisms for generating the visualisation, the degree of browser independence, and the navigation support facilities. We provide a critical review of the diverse range of WWW visualisation tools with respect to these issues.

Keywords

World Wide Web; hypermedia navigation; design issues; visualisation.

1 Introduction

With millions of users searching, browsing, and surfing the WWW, economies of scale are clearly relevant. Even a small inefficiency in user navigation within the WWW will result in enormous productivity losses if it is common to a fraction of WWW users. Our previous research has noted that users commonly have incorrect mental-models of even the most fundamental methods of WWW navigation, and has suggested interface mechanisms to ease the problems [CJ96].

Despite the importance of usability issues, there is a paucity of research on the underlying causes of problems in WWW navigation. Notable exceptions include the empirical works of

[†]Authors are listed alphabetically.
[§]Author for correspondence.

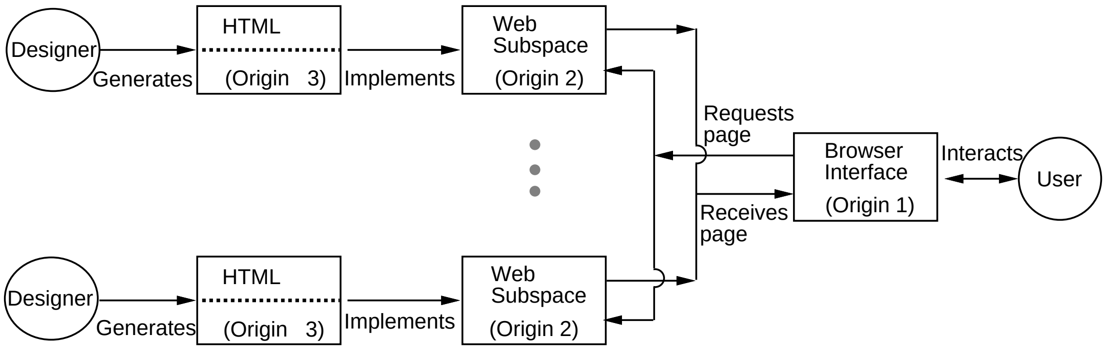

Figure 1: *Origin of user problems in WWW navigation*

Catledge and Pitkow [CP95] and of Tauscher and Greenberg [TG96] which infer patterns of interaction and user problems from extensive logs of browser use.

Although research on the *causes* of problems in WWW navigation is sparse, development of novel browsing techniques to assist navigation is prolific. This paper considers design issues in visualisation tools to support WWW navigation, integrated with a review of existing tools.

Section 2 identifies three levels of WWW design and use that can introduce difficulties for the end user, and briefly discusses potential solutions at each level. Section 3 focuses on one of these solutions, that of novel visual browsing support tools. Key issues in the design and functionality of these tools are identified. The diverse range of WWW visualisation tools are critically reviewed, including our own navigational assistant WebNet.

2 A Taxonomy of The Sources of User Problems in WWW Navigation

This section describes the origins of user problems when navigating the WWW, and briefly discusses potential solutions to these problems. We are primarily concerned with the user-centred problems of becoming "lost in hyperspace" [Nie90] rather than the problems deriving from issues such as network limitations, page-loading speeds, and network bandwidth. Naturally, some user-centred navigational difficulties are due to quirks in a particular user's mental model of the navigation facilities that their browser offers. Many misunderstandings, however, are encouraged by inadequacies in browser interfaces [CJ96], by poor WWW space design, and by the limitations forced on the page designer by the constraints of their hypertext description language. Figure 1 illustrates these origins of user problems when navigating the WWW. The following sections discuss the navigational difficulties arising at each of the three origins, and they briefly review some of the potential solutions to the problems. In each case, the problems are introduced and illustrated by a user scenario (in italics). Table 1 further summarises the problems arising from each level, noting the proportion of users affected and some of the potential solutions.

2.1 Origin 1: Browser User Interfaces

Jane uses Netscape Navigator[1] *about three times a week. She 'understands' that the 'Back' button lets her return to the pages she has just been to, but quite frequently she is surprised that the pages she has just visited 'disappear': she cannot find them using 'Back' and 'Forward'.*

[1]Netscape Navigator is a trademark of Netscape Communications Corporation.

Origin of Problem	Proportion of users affected	Example problems	Potential solutions
Browser Interfaces (Origin 1)	Very high (all WWW browsing is carried out through a handful of browsers).	Misunderstanding of client's facilities (for example, Back and Forward).	Improved system image to better communicate the system's facilities to the user.
		Range of client's facilities (for example, the absence of an interactive history list).	Extended facilities for user support.
Page Design (Origin 2)	Small for each site, but errors of page design are common across many sites.	Poor structure within a WWW site (promoting "lost in WWW space".	Tools to support structured WWW design (see section 2.2).
		Poor graphical design (consistency in representation, legibility and readability, visual appeal, etc).	Guidelines for page designers (for example, [NS95]).
Hypertext Markup Language (HTML) Features (Origin 3)	Very high (almost all pages are written using a few dialects of HTML).	Restricted range of expressible hypertext facilities. Inability to affect browser state (such as the history list).	Solutions to these problems are unlikely. There are strong conflicts between the need for standardisation, for advanced features, and for security.

Table 1: *Origins of user problems in WWW navigation, the proportion of users affected, example problems, and potential solutions.*

The scenario illustrates the navigation problems that arise when the user misinterprets the navigation facilities offered by their browser. Many users do not realise that the history mechanisms of most commercial browsers are stack-based, rather than being linear lists of visited pages [CJ96]. The 'Back' button descends into a stack of previously visited pages, and 'Forward' ascends towards the top of the stack. A link selection while within the stack removes all pages above the current stack position with the consequence that they cannot be re-accessed using 'Back' and 'Forward'.

There has been little research to investigate the efficacy of stack-based navigation, and what has been done indicates that other mechanisms would improve navigation. Tauscher and Greenberg [TG96], for instance, provide evidence, based on extensive logs of page revisitation patterns, that stack-based navigation mechanisms provide poor prediction of page revisitation. Their research indicates that simple recency-based mechanisms (with duplicate pages removed) would improve re-access to pages. Microsoft Internet Explorer goes some way towards providing a facility with a persistent date-ordered history mechanism. Tauscher and Greenberg also state that graphical mechanisms for page display could provide further improvements. Bieber and Wan [BW94] investigate multi-window backtracking in hypertext, and suggest that several alternative schemes for backtracking should be offered within a hypertext system. These implications are unsurprising within hypertext research communities which have extensive experience with rich navigation facilities to assist maintaining a sense of context within information spaces. Many current research projects are investigating the use of overview maps and graphical displays of WWW spaces. These novel mechanisms are reviewed in section 3.

2.2 Origin 2: WWW Subspace Design

Bob understands that Netscape Navigator uses a stack based history mechanism. When visiting WWW subspaces he finds that he bookmarks many pages by default because he is not sure whether he will be able to return to them via other links, or if they will disappear from his history list. He is frustrated that so few WWW pages indicate fundamental navigation information such as *their relationship to others in a WWW subspace, their links to the previous or next page in a collection, and visual clues regarding the WWW subspace structure. He is sure that he would spend less time managing temporary bookmarks if pages contained this information.*

Navigation problems can be forced on the user by inappropriate WWW subspace design, even when the user has perfect understanding of their browser's behaviour. The hypertext interconnections in WWW subspaces can be extremely rich, and designers would benefit from tools and guidelines to support and assist them. Several design methodologies for hypermedia design are detailed in [BI95]. Isakowitz, Stohr and Balasubramanian [ISB95] review the Relationship Management Methodology (RMM)—a diagramming technique for designing and constructing hypermedia applications, which provides a platform for collaborating designers to discuss, consider, and record the development of complex inter-related data-spaces. Nanard and Nanard [NN95] describe MacWeb, a prototyping tool that supports the construction and visualisation of hypertext structures. Thimbleby describes GenTL, a tool for systematic WWW subspace authoring and analysis [Thi96, Thi95].

Navigational problems in WWW subspace design go beyond the structure of page interconnection, and issues of appropriate graphical presentation of WWW pages can greatly affect the usability of WWW sites. Guidelines for graphical page design, such as those of Nielsen and Sano [NS95], can help designers to make appropriate decisions regarding the visual presentation of their pages.

2.3 Origin 3: WWW Subspace Description Language

Mary is hired to implement a WWW subspace for a large organisation. She has extensive experience in hypertext and multi-media design, but is relatively new to the WWW. She continually finds that the limitations of HTML constrains the facilities she can provide in her WWW subspace.

The range of hypertext facilities available to a WWW subspace designer is constrained by the language of WWW subspace description, HTML (HyperText Markup Language). Further constraints on the WWW subspace design are imposed by the differences between browsers which cause hypertext facilities that work on one browser to fail on another (or be presented in a different fashion).

HTML allows only one type of page-link, and authors who wish to provide "standard" hypertext facilities such as typed links or bi-directional links will find HTML's facilities constraining. Many other desirable hypertext facilities are unavailable within HTML. For example, a WWW space designer may wish to tailor the history state of the browser when the user enters a particular page—such tailoring would allow, for instance, forceful metaphors for entering and exiting WWW subspaces. Many current WWW spaces have an 'Exit' page to provide a sense of navigation closure, but having reached the exit page, the user must carefully search the browser's History list for the last page outside the WWW subspace if they are to avoid coincidentally re-entering the subspace through use of 'Back'.

It is unlikely that the problems originating from constraints on HTML will be resolved. There are strong conflicts between the need for standardised HTML, advanced features in HTML, and security in HTML.

2.4 Alleviating Navigational Problems: Discussion

To alleviate these navigational problems requires an integrated approach addressing each of the three levels concurrently. First, new WWW browsing applications could be developed to supersede current browsers. This approach has several drawbacks including the complexity of browser design, and the need to deal with the profusion of media types available on the WWW. Pragmatically, however, the marketplace dominance of Netscape Navigator and Microsoft Internet

Browser Tool	Visualisation Creation			Visualisation Dimension		Visualisation Style		View Filtering	
	Before Browsing	During Browsing	Both	2-D	3-D	Network	Tree	Visual	Semantic
WebMap		√		√		√		√	**x**
MosaicG		√		√			√	**x**	**x**
HyperSpace		√**(a)**			√	√		**x**	√**(b)**
Mosaic Enhancements (Gershon)		√		√			√	**x**	**x**
WebNet		√		√		√		√	√
Transducers/Associates	√	√	√**(c)**	√			√	**x**	**x**
Hyperbolic Views	√				√**(d)**	√		√**(e)**	**x**
HotSauce	√				√		√	√**(f)**	
WebCore	√			√**(g)**		**N/A**	**N/A**	**N/A**	**N/A**
Navigational View Builder	√			√**(h)**	√	√	√	√	√
WebViz	√			√		√		√**(i)**	
DeckScape		√		√**(j)**		**N/A**	**N/A**	**N/A**	**N/A**

(a) Hyperspace recreates the full visualisation rather than incrementally amend the current state
(b) HyperSpace positions objects in the visualisation using an algorithm for determining their relatedness. Closely related documents are positioned closely together.
(c) Provided by different tools within the architecture.
(d) VRML based
(e) Hyperbolic Views allows different layouts to be created.
(f) HotSauce allows the depth of lookahead in the hierarchy and the number of nodes on the screen at any one time to be user controlled.
(g) This is a textual visualisation.
(h) The description of Navigational View Builder suggests many potential visualisations.
(i) WebViz can provide alternative random visualisation layouts.
(j) DeckScape is a WWW browser but is included because it addresses some of the issues mentioned here.

Table 2: *Visualisation attributes specified for each visualisation tool.*

Explorer precludes radical browser advancements unless they originate from within Netscape and Microsoft.

Second, WWW subspace designers could be educated on how to support effective navigation through their information spaces. This could include directions on graphic design, on content presentation, and on the integration of user-centred navigation aids, such as overview maps, within WWW subspaces. Pragmatically, however, no set of design advice or guidelines will be adopted unilaterally across the Internet. Constraint-based WWW authoring tools such as GenTL [Thi95, Thi96], and automated HTML analysis software such as `weblint` [Bow96] can also help to produce effective navigation support within WWW subspaces, but as with design guidelines, only a small proportion of web designers are likely to adopt them.

Third, richer and more powerful languages for WWW subspace description would enable expanded facilities for describing navigations between pages. These could include hypertext facilities such as typed links [CB96] and bi-directional links [BA96]. A challenge for such languages or tools is to support these attributes across subspaces belonging to different authors where the level of content and navigation awareness and control is unpredictable. The final design and standard of such languages is beyond the sphere of influence of most WWW subspace designers and users.

So, considering these problems, and given that bringing about changes to WWW page description languages and page designers' perspectives is a very hard and likely long-term task, what alternative approaches remain?

We advocate augmenting the functionality and user interface of normative WWW browsers with graphical visualisations of web-subspaces. Our work with WebNet is investigating the development of browser-independent navigational assistants which map and adapt to the user's navigational acts. Users can initiate navigational acts either at the unaltered browser (such as Netscape, or Internet Explorer) or at WebNet's graphical overview diagram. The main aim of this work is to alleviate the deficiencies in navigation support of current WWW browsers, and to overcome problems derived from unsupportive page design. It does not directly address restrictions of HTML. Many other researchers are also developing systems, reviewed in the following section, with similar ambitions.

3 Augmenting Browser Navigation Support

The value of graphical overview diagrams in assisting user navigation through complex information spaces is well known within Hypertext research [Nie90, Con88, UY89, CB88]. Many researchers are now interested in providing overview diagrams to help overcome the WWW's navigational difficulties which were reviewed in the previous sections. Systems are rapidly being released, and the methods, means, and ambitions of these visualisation systems differs greatly across the diverse systems. To assist research in this area, this section critically reviews novel WWW browsing systems. The structure of this review is based around three key issues that these systems must address:

1. Characteristics of the visualisation, including issues such as the mechanisms used to generate the visualisation, its dimensionality, style, and scope.
2. Navigation support functions. The extent of the facilities that the WWW visualisation provides for users.
3. Browser independence. The degree to which a particular approach is tied to specific WWW browser.

These issues are addressed in turn below.

3.1 Visualisation Characteristics

WWW pages and the links between them can be represented through a variety of graphical means to help overcome the two primary problems of hypertext identified by Conklin [Con88]:

- *disorientation:* graphical representations aim to help users maintain a sense of context within an information space; and
- *cognitive overhead:* graphical representations can provide an external representation of the user's memory of their navigation session.

Additionally, graphical representations can provide an enriched interface for initiating navigational acts.

The use of visualisation techniques, however, begs the question of what is actually visualised? The novel WWW browsing systems developed to date differ greatly in their visualisation properties. Three critical issues that must be addressed, discussed below, are the mechanisms for creating the visualisation, the scope of the WWW space that can be visualised, and the style and dimension of the visual display. The user-interface facilities offered by the visualisation are discussed in section 3.2.

3.1.1 Visualisation Creation

A pivotal distinction between WWW visualisation tools is whether the representation of the subspace is generated *statically* before the user enters a site, or *dynamically* while the user navigates through a site. Table 2 summarises the approaches to visualisation creation across eleven WWW navigation visualisation tools.

Static systems typically provide users with a complete visualisation of a WWW subspace at the beginning of their navigation session. These views are generated prior to the user's arrival at a site, and consequently computationally expensive visualisations such as those provided by WebViz [PB94], Munzner and Buchard's hyperbolic views [MB94], and Navigational View Builder [MF95] can be created without affecting system response time. Their disadvantage is that they do not, or can only minimally, adapt to the user's actions: they are difficult to modify in real time. HyperSpace [WDBH95] supports statically generated views, but it attempts to provide dynamic modification during navigation. This requires the view to be recreated (rather than updated) with each navigational act—an inefficient solution.

A further difficulty with the static approach is that the representation becomes inconsistent with the web-subspace whenever the underlying pages are modified. Statically generated visualisations can be particularly valuable for WWW subspace designers and maintainers, and several tools for authoring HTML support visualisation of the specified subspace (see section 2.2).

With dynamic generation of the visual display, the visual representation of the WWW subspace adapts to the user's navigational acts. New pages are added to the display as the user encounters them, and the displays normally adapt to show the user's current context within previously browsed pages. The dynamic adaption of the display causes gradual degradation in system performance as the extent of the graphical representation increases. WebMap [Doe94], MosaicG [AS95], Gershon's Mosaic enhancements [GWL+95], and WebNet [CJ96] all provide dynamically updated views. A significant advantage of dynamically generating the the visualisation is that it reflects exactly the navigated subspace, and its integrity is ensured. It becomes a snapshot of the visited pages at the time of navigation.

3.1.2 Scope of the Subspace Visualisation

Two issues of the subspace scope that must be considered are its range (or size) and its temporal extent. Although both of these issues are strongly influenced by the mechanisms used to generate the subspace (discussed in the previous section) they are worthy of separate discussion because they capture the requirements of the visualisation rather than the facilities available in current systems.

With respect to the range of the subspace, the first approach is to create a visualisation of a finite, well defined subspace such as an intranet (normally collated through static generation). This approach is taken by systems such as WebViz, Munzner and Buchard's hyperbolic visualisations, WebCore [BA96], Navigational View Builder, and HotSauce [App96]. This has the benefit that a subspace can be completely and accurately mapped and contain representation of all pages and the relationships between them. The disadvantage is that these systems fail to support navigation outside the boundaries of the visualised space. Additionally the complete visualisations can swamp the user in information that is irrelevant to them, although it could be argued that users are unable to decide to exclude paths to pages until they have been shown.

The second approach, normally associated with dynamically generated web-spaces, is to provide flexible and unconstrained visualisations of whatever pages the user navigates. This has the advantages that the user is presented with relevant information (those that they chose to visit) and that the visualisation can display pages across subspace boundaries (and serve to make those boundaries as obvious or transparent as required). The user can use the visualisation to navigate around the subspace they defined during their browsing session. A disadvantage is that these visualisations do not show the *potential* destinations from visited locations. Example systems in this category are WebMap, MosaicG, and Gershon's Mosaic enhancements.

The key difference between the two approaches amounts to a distinction between the temporal nature of the navigation support provided. The first scheme primarily supports the user in deciding where to go to next, but provides little navigation history. The second approach allows users to view their navigation session and see relationships between previously visited pages, but provides little support for visualising potential destinations from those visited pages. A hybrid approach, adopted by HyperSpace and WebNet, provides both history and lookahead. WebNet, for example, dynamically updates its subspace representation as the user moves from page to page, but it also provides a lookahead facility to view the links contained within visited pages.

3.1.3 Visualisation Dimension

The graphical representation of pages and links can either be presented in two or three dimensions. Most visualisations are in two dimensions and are exemplified by systems including WebMap, the WWW 'associates' of Brooks *et al* [BMM95], WebCore [BA96], MosaicG [AS95], Gershon's Mosaic enhancements [GWL+95], WebViz [PB94] and WebNet [CJ96]. Examples of the visualisations of these systems can be seen in Figures 2, 3, 4 and 5. Conventionally these representations present pages as nodes (boxes or circles) containing page titles, URLs or unique identifiers. Links between pages are presented as edges (lines). In some systems the lines contain arrowheads indicating link directions. The benefits of two dimensional visualisations are that for reasonable numbers of nodes they are computationally inexpensive to update and rearrange and, it could be argued, match some users' mental model of the WWW's pages and links. A key disadvantage is that they quickly become overloaded with nodes and links making updates more computationally expensive and losing clarity of visualisation.

Three dimensional visualisations such as those exemplified in Munzner and Burchard's hyperbolic visualisations [MB94], HyperSpace [WDBH95], Navigational View Builder [MF95] and Hot Sauce (formerly Project X) [App96] (see Figures 2, 3, 4 and 5) vary significantly in representation of pages and links (discussed below). The advantage of three dimensional representations is that they can offer novel views of information spaces and notionally represent a larger information space than two dimensional systems within the same screen real estate. A distinct disadvantage is that the generation of such visualisations is computationally expensive and consequently their dynamic performance is poor on "standard" hardware.

3.1.4 Visualisation Representation

The structure of the WWW is a network. Several systems, however, represent web-subspaces as a hierarchy or tree like visualisation. These include Brooks *et al*'s WWW associates, Mosaic enhancements [GWL+95], MosaicG, and some suggestions for Navigational View Builder [MF95]. Implementation constraints, or the perceived simplicity of hierarchical displays may have motivated the use of non-network displays.

The majority of systems provide network (potentially cyclical) views. These include WebMap [Doe94], WebViz [Doe94] and WebNet [CJ96] in two dimensions, and hyperbolic visualisations [MB94] and HyperSpace [WDBH95] in three dimensions. Apple's recent tool HotSauce [App96] provides (at the time of writing) a three dimensional visualisation but in a strongly hierarchical manner, perhaps making it more suitable to file-store rather than WWW subspace representation.

Most of the systems make no provision for altering and adapting the visual representation of the pages and links. HotSauce, however, has highly dynamic visualisation properties, allowing users to 'fly' forwards and backwards and through its visualisations at varying velocities. Although this is essentially a zoom function, it is this ability which gives HotSauce its three dimensional 'feel'. Objects within the HotSauce view can be dynamically relocated in two dimensions using direct manipulation. Hyperbolic visualisations and HyperSpace use Virtual Reality Modelling Language (VRML) [Net96] to support manipulation of the view and movement around the information space in three dimensions. They do not support manipulation of objects within their views.

3.2 Navigation Support Functions

Passive visualisations, with no associated functionality, can assist navigation by contextualising the users navigational actions and by providing a surrogate for the user's short term memory. The WWW subspace representations can be greatly reinforced, however, by integrating functionality into the visualisation. Table 3 summarises the systems under consideration and the additional navigation functions they provide.

Browser Tool	Browser Independent	Shows links from visited pages	Saveable Session History	Dual Control	Subspace Management
WebMap	**x**	**x**	√	√	**x**
MosaicG	**x**	**x**	√	√	√**(a)**
HyperSpace	**x**	√	**?**	**x**	**x**
Mosaic Enhancements (Gershon)	**x**	**x**	**?**	√	**x**
WebNet	√	√	√	√	√
Transducers/Associates	**x**	**x(b)**	**?**	√	**x**
Hyperbolic Views	**x**	**N/A**	**x**	**x**	**x**
HotSauce	**x(c)**	**N/A**	**x**	**x**	**x**
WebCore	**x**	**N/A**	**N/A**	**N/A**	**N/A**
Navigational View Builder	**?**	**N/A**	**?**	**?**	**?**
WebViz	**N/A**	**N/A**	**N/A**	**N/A**	**N/A**
DeckScape	**N/A**	√	√	**N/A**	√

(a) MosaicG creates a new tree/hierarchy for each subspace.
(b) The links could be shown if the link tree associate was also used.
(c) HotSauce gives users the option of which of three browsers (Netscape Navigator, Internet Explorer and CyberDog) they wish to use to view pages, but is not browser independent.

Table 3: *Additional functional attributes specified for each visualisation tool.*

3.2.1 Dual Control

We use the term *dual control* to describe the ability of the WWW subspace visualisation to issue commands to an associated web-browser. Through dual control, the user's navigational acts can be initiated either at the browser or at the visual representation.

Several existing systems provide dual control: WebMap, MosaicG, Brooks' associates, Gershon's Mosaic enhancements, HotSauce[2] and WebNet. There are several benefits of such functionality. First, users do not have to move their focus between visualisation and browser to move from page to page. Second, users can return to previously visited pages though a single click on the appropriate node in the visual display. This avoids dependence on selection from potentially incomplete browser history lists [JC96] or multiple 'back' navigations to reach the desired page. Third, links to pages which are not currently visible in the browser can be followed.

3.2.2 Link Previews

Using standard browsers such as Netscape, users wishing to visit a series of links off a single page normally use a "hub and spoke" [CP95] browsing strategy in which the user repeatedly returns to the "hub" page to gain access to subsequent links. Graphical representations which support link previews (showing all the available links off a page) and dual control remove the need for these redundant traversals back to the hub. WebNet provides both navigation history and possible future destinations within its visualisation. From each node (page) in the visualisation, all links from that node can be displayed, and clicking on the nodes causes the (separate) browser to display the page.

3.2.3 Saveable Navigation Histories

Users often return to WWW subspaces or paths between subspaces to review information or to extend the scope of the navigated pages [TG96]. Standard browsers provide limited support for persistent records of visited pages through bookmarking facilities, but these fail to show complete histories or paths.

Statically generated visualisations provide persistent views of subspaces, but they do not represent a particular user's path through pages, nor do they differentiate between pages which have and have not been visited. In contrast, dynamically generated visualisations can allow

[2] The HotSauce visualisation is not updated as a result of activity within a browser.

users to save their navigation histories for use in other sessions (even perhaps by other users). WebMap, MosaicG, and WebNet allow navigation histories to be saved. Of course once such a history is created it can suffer similar problems to statically generated visualisation such as out-of-date links and pages, and a resultant incorrect mapping between visualisation and underlying subspace.

By combining these facilities for saveable histories of paths and dynamically generated visualisations, users are able to generate diverse navigational paths, possibly between diverse sites. This allows the visualisation to be split into 'chunks' which are manageable by the user, and keeps pages accessed in different phases of the navigation activity closely related to each other in the visualisation. DeckScape [BS95] (actually a browser rather than visualisation tool) and WebNet both support these facilities.

Browser Tool	**Problem Level Addressed**	**System Objectives and Comments**
WebMap	**2, 3**	Purpose: to support navigation in the WWW. Interesting features are the use of colour to differentiate between intra and inter-server links. A playback feature allows replay of navigation activity. Drawbacks are the lack of meaningful labels in visualisation nodes and the need to modify NCSA Mosaic code.
MosaicG	**2,3**	Purpose: to enhance history keeping and act as a navigation aid during WWW browsing. MosaicG uses thumbnail images as visualisation nodes rather than labels to support identification within the visualisation and allows `collapsing' of sections of the visualistion.
HyperSpace	**2,3**	Purpose: to support navigation in the WWW. Provides an initial random organisation of nodes(pages) in a 3D virtual space. The visualisation can be remapped to represent relatedness of node content.
Mosaic Enhancements (Gershon)	**2,3**	Purpose: to help users quickly locate where they are in hyperspace (particularly the WWW). Uses colour to impart information about link types and allows users to overlay further links on the actual WWW structure being visualised.
WebNet	**2,3**	Purpose: to support WWW navigation within and across sessions. Browser independent, with saveable navigation histories and subspace management. Has been extended to provide collaborative browsing.
Transducers/Associates (Brooks)	**2,3**	Purpose: to circumvent limitations in the usability of the WWW through the use of independent software entities (agents). Stream transducers intercept HTTP requests to facilitate browsing associate activity.
Hyperbolic Views	**2,3**	Purpose: to support WWW navigation by showing multiple documents and the links between them in an appropriate 3D representation. Uses hyperbolic space for visualisations to avoid clutter and provide movement.
HotSauce	**2,3**	Purpose: to support navigation around information spaces (including WWW subspaces). Can represent information spaces for which Meta Content Format (MCF) descriptions have been prepared. Allows highly dynamic movement around the visualisation both within WWW browsers and stand-alone applications.
WebCore	**1**	Purpose: to overcome the lack of reverse links from WWW pages (ie to show which pages contain links to the one currently being viewed). Doesn't use a graphical visualisation, the reverse links are shown in an HTML page. Requires analysis of a great many WWW pages.
Navigational View Builder	**2,3**	Purpose: to aid orientation and navigation in hypermedia systems (including the WWW) through the use of a variety of overview diagrams.
WebViz	**2**	Purpose: to provide graphical visualisations of WWW access logs. Although not directly concerned with navigation support the visualisation are similar to those adopted in other systems, and the use of such a tool may guide more appropriate WWW subspace design.
DeckScape	**2,3**	Purpose: to investigate new navigation methods through a new WWW browsing application. Uses a `deck' metaphor to create collections of WWW pages. It is actually a limited functionality browser with restricted visualisation but some aims and concepts relate to navigation support.

Table 4: *Additional functional attributes specified for each visualisation tool.*

3.2.4 Visualisation Filtering

Visualisation filters allow the user to control the amount, type, and representation style of information within the display. These facilities help in controlling clutter in the display, assist in highlighting important information. Mukherjea and Foley [MF95] identify three types of filtering operations: *content based* in which nodes with specific attributes are shown or hidden; *link based* in which certain types of links are shown or hidden; *structure-based* in which the topology of the visualisation are used as filtering criteria. We would add *interaction-based* filtering to this list, which would allow filtering based on the history of the user's navigational actions. In addition to controlling the content of the visualisation, the some systems also allow the user to control the style of the visualisation. For instance, WebMap users can choose between alternative graph layout algorithms to change the shape of the graph, and WebViz supports alternative random positioning of objects in the graph.

The Navigational View Builder supports content, link-based, and structure-based filtering. It also supports alternative styles of representation. WebMap supports link-based filtering, using different representations of line colour and style to differentiate between links within and across WWW servers. WebViz supports interaction-based filtering based on recency and frequency of access to pages within the displayed subspace, and HyperSpace used structure-based filtering to arrange the view according to a relatedness heuristic. WebNet supports structure-based filtering to include or exclude objects based on their logical distance (in terms of number of links) from the current page. It also supports interaction-based filtering based on the frequency and recency of the user's access to pages.

3.3 Browser Independence

Users are likely to have a particular browser preference, and new browsers may emerge in the future. It is there beneficial for the additional navigation support tools that we are advocating to be able to operate with any browser. It is important to note the difference between true browser independence and functioning with a large but finite set of browsers. Browser independence requires an underlying architecture that is free of browser specific protocols, while support for several browsers requires the incorporation of communication protocols for each supported browser.

Tools that augment browser functionality must establish a communication channel with the browser. Dual control (section 3.2.1) requires a two-way communication between the browser and the browser augmentation tool. NCSA Mosaic was the first browser to support a communication protocol with external applications, using its Common Client Interface (CCI). Consequently, by far the majority of browser augmentation tools developed to date are tied to Mosaic, including WebMap, MosaicG, HyperSpace, Brooks' associates, and Gershon's Mosaic Enhancements.

WebNet originally operated only with the tkWWW browser, but a browser independent architecture has been developed. In this architecture a restricted functionality proxy HTTP server (also termed WWW server) mediates the communication between WebNet and any browser. Page requests, which may originate at WebNet or the browser, are forwarded to the modified proxy server which returns details of the retrieved page to both the browser and to WebNet.

4 Summary

In this paper we have described three levels of problem which users may experience while navigating around the WWW. These were WWW page description languages, WWW page and subspace design and WWW browser design. We focused on one approach to ameliorating these problems: the augmentation of browsing applications with visualisations of WWW subspaces.

We critically reviewed novel WWW browsing applications across three key design issues: the characteristics of the visualisations; the functionality attached to the visualisations; and browser independence. Several novel systems that provide graphical representations of WWW subspaces were discussed with respect to these issues. The purpose and attributes of these systems are summarised in Table 4.

Our aim is to assist and motivate research in this area. Our own further work will continue developing WebNet, a browser independent navigational assistant that dynamically generates graphical overview maps of the user's history and potential navigational paths.

References

[App96] Apple Computer Inc. Hot sauce. World Wide Web Document, 1996. URL http://mcf.research.apple.com/.

[AS95] EZ Ayers and JT Stasko. Using graphic history in browsing the world wide web. In *Proceedings of the Fourth International World Wide Web Conference. 11–14 December, Boston*, 1995. Also available from: http://www.w3.org/pub/Conferences/WWW4/Papers2/270/.

[BA96] C Bacquero and JP Andrade. Webcore - forward and backward navigation on a multi-grain web representation. World Wide Web Document, 1996. URL http://shiva.di.uminho.pt/ cbm/htdoc/Overview.html.

[BI95] M Bieber and T (eds) Isakowitz. Designing hypermedia applications. *Communications of the ACM*, 38(3):27–29, 1995.

[BMM95] C Brooks, WS Meeks, and MS Mazer. An architecture for supporting quasi-agent entities in the www. CIKM'95 Intelligent Information Agents Workshop, Baltimore, December 1–2. World Wide Web Document, 1995. URL http://www.osf.org/www/waiba/papers/CIKM/CIKM.htm.

[Bow96] N Bowers. Weblint: Quality assurance for web pages. World Wide Web Document, 1996. URL http://www.khoral.com/staff/neilb/weblint.html.

[BS95] MH Brown and RA Shillner. Deckscape: An experimental web browser. *Computer Networks and ISDN Systems*, 27(6), 1995. Third International World-Wide Web Conference, Darmstadt. World Wide Web Document. URL http://www.igd.fhg.de/www/www95/papers/90/deckscape-final-v1/paper.html.

[BW94] M Bieber and J Wan. Backtracking in a multiple window hypertext environment. In *Proceedings of the Fifth ACM Conference on Hypermedia Technologies (ECHT'94)*, Edinburgh, September 1994, pages 158–166, September 1994.

[CB88] J Conklin and ML Begeman. gIBIS: A hypertext tool for exploratory policy discussion. *ACM Transactions on Office Information Systems*, 6(4):303–331, October 1988.

[CB96] C Cleary and R Bareiss. Practical methods for automatically generating typed links. In *Hypertext '96: the Seventh ACM Conference on Hypertext. Washington DC March 1996*, pages 31–41. ACM Press, 1996.

[CJ96] A Cockburn and S Jones. Which way now? analysing and easing inadeqacies in www navigation. *International Journal of Human-Computer Studies*, 45:105–129, 1996.

[Con88] J Conklin. Hypertext: An introduction and survey. In I Greif, editor, *Computer Supported Cooperative Work: A Book of Readings*. Morgan Kaufmann, 1988.

[CP95] LD Catledge and JE Pitkow. 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*, volume 27, pages 1065–1073, 1995.

[Doe94] P Doemel. Webmap - a graphical hypertext navigation tool. World Wide Web Document, 1994. URL http://www.tm.informatik.uni-frankfurt.de/Publications/Doemel/WWW94/www94.html.

[GWL+95] N Gershon, J Winstead, J LeVasseur, J Croall, A Pernick, and W Ruh. Visualizing world wide web information resources. World Wide Web Document, 1995. URL http://www.cs.umbc.edu/ cikm/1995/npiv/gershon/paper.html.

[ISB95] T Isakowitz, EA Stohr, and P Balasubramanian. Rmm: A methodology for structured hypermedia design. *Communications of the ACM*, 38(3):34–44, 1995.

[JC96] S Jones and A Cockburn. A study of navigational support provided by two world wide web browsing applications. In *Hypertext '96: the Seventh ACM Conference on Hypertext. Washington DC March 1996*, pages 161–169. ACM Press, 1996.

[MB94] T Munzner and B Burchard. Visualizing the structure of the world wide web in 3d hyperbolic space. World Wide Web Document, 1994. URL http://www.geom.umn.edu:80/docs/research/webviz/.

[MF95] S Mukherjea and JD Foley. Visualizing the world-wide web with the navigational view builder. *Computer Networks and ISDN Systems*, 27(6):1075–1087, 1995.

[Net96] Netscape Communications Corporation. An introduction to vrml. World Wide Web Document, 1996. URL http://partner.netscape.com/comprod/products/navigator/live3d/intro_vrml.html.

[Nie90] J Nielsen. The art of navigating through hypertext: Lost in hyperspace. *Communications of the ACM*, 33(3):296–310, 1990.

[NN95] J Nanard and M Nanard. Hypertext design environments and the hypertext design process. *Communications of the ACM*, 38(3):49–56, 1995.

[NS95] J Nielsen and D Sano. Interface design for sun's www site. WWW page: http://www.Sun.COM/sun-on-net/uidesign/, May 1995.

[PB94] J Pitkow and KA Bharat. Webviz: a tool for world-wide web access log analysis. World Wide Web Document, 1994. First International WWW Conference, Geneva. URL http://www.elsevier.nl/cgi-bin/WWW94link/31/overview.

[TG96] L Tauscher and S Greenberg. Patterns of revisitation in world wide web navigation. Technical Report Research report 96-587-07, University of Calgary, Canada, 1996.

[Thi95] H Thimbleby. Gentl manual. World Wide Web Document, 1995. URL http://www.cs.mdx.ac.uk/esrc/auto.html.

[Thi96] H Thimbleby. Systematic web authoring. World Wide Web Document, 1996. URL http://www.cs.mdx.ac.uk/harold/webpaper/.

[UY89] K Utting and N Yankelovich. Context and orientation in hypermedia networks. *ACM Transactions on Office Information Systems*, 7(1):58–84, 1989.

[WDBH95] A Wood, N Drew, R Beale, and B Hendley. Hyperspace: Web browsing with visualisation. World Wide Web Document, 1995. URL http://www.cs.bham.ac.uk/ amw/hyperspace/www95.

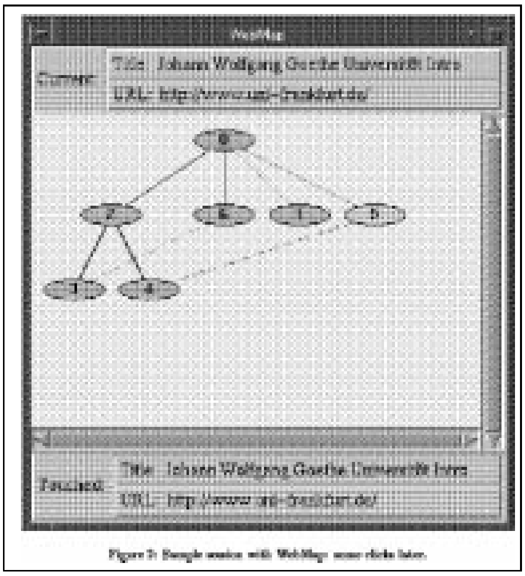

(a)

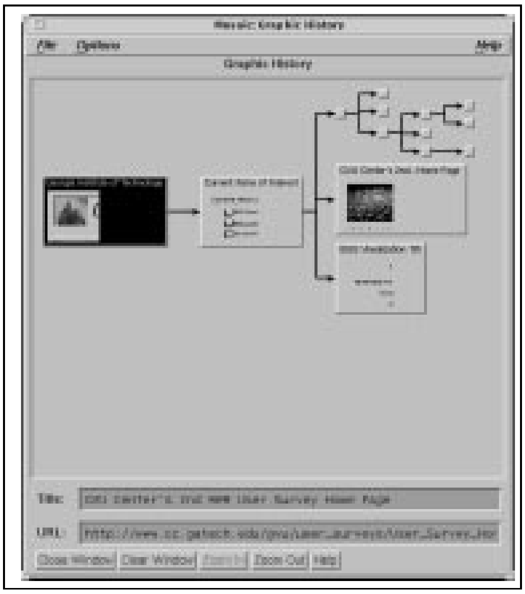

(b)

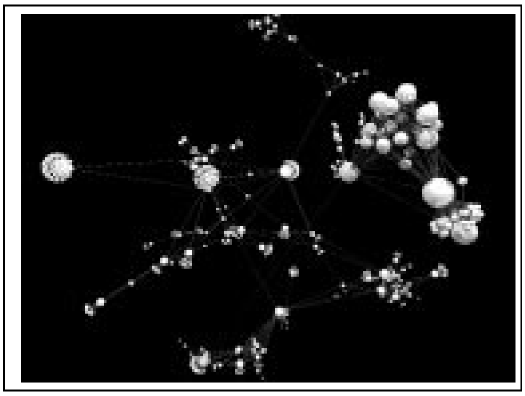

(c)

Figure 2: *Example visualisations from systems which create visualisations during the navigation process: (a) WebMap, (b) MosaicG, (c) HyperSpace.*

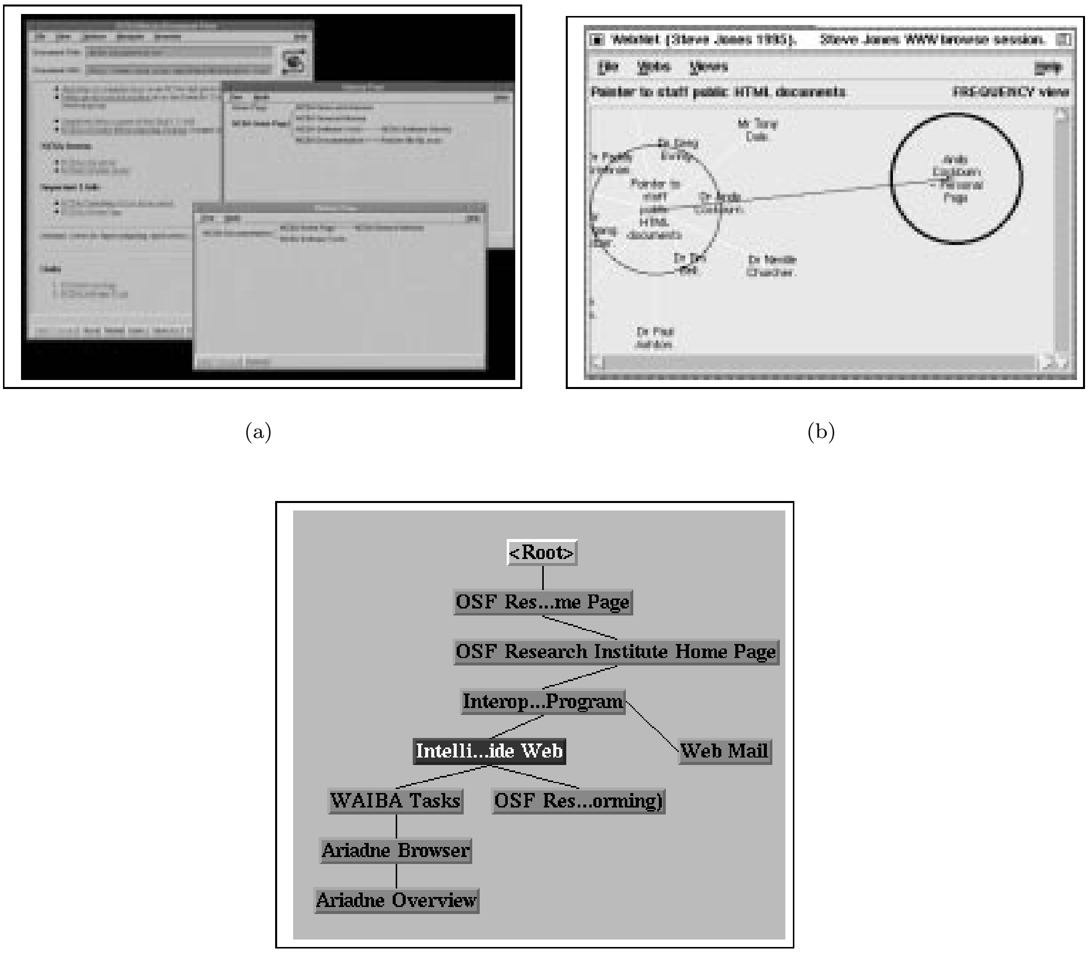

Figure 3: *Further example visualisations from systems which create visualisations during the navigation process: (a) Gershon's Mosaic enhancements, (b) WebNet and (c) one Brooks* et al*'s associates.*

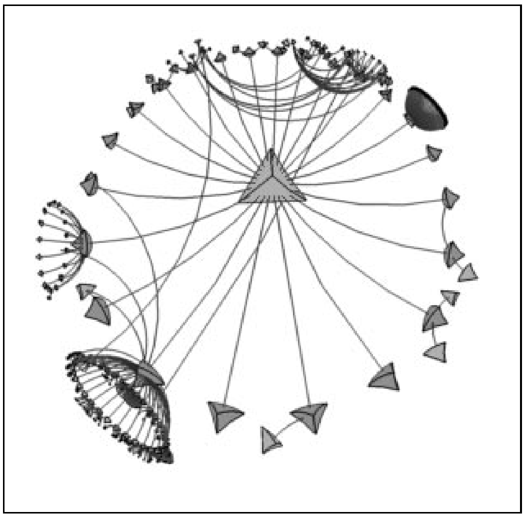

(a)

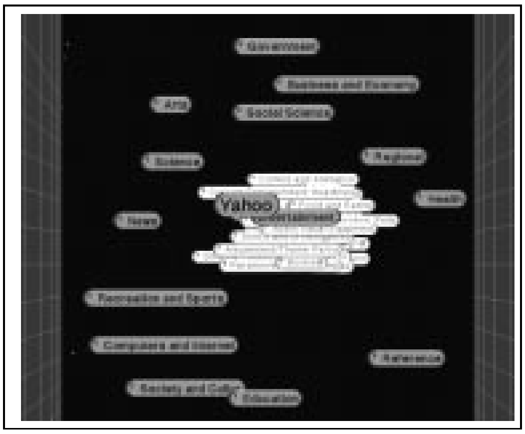

(b)

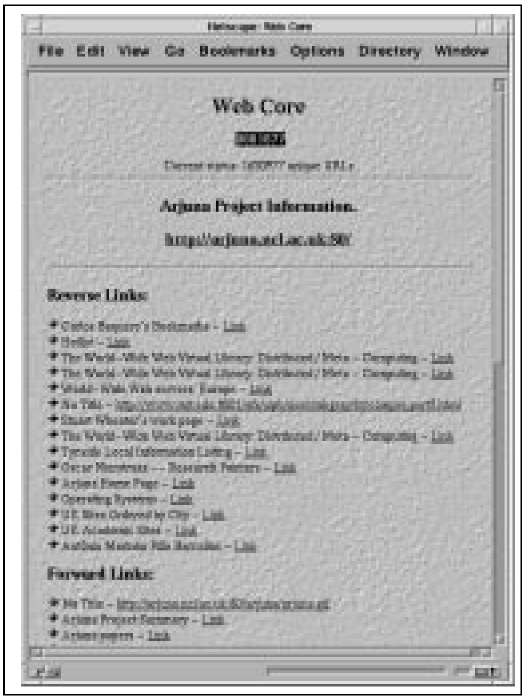

(c)

Figure 4: *Example visualisations from systems which create visualisations prior to the navigation process: (a) Munzner and Buchard's hyperbolic browser, (b) HotSauce, (c) WebCore.*

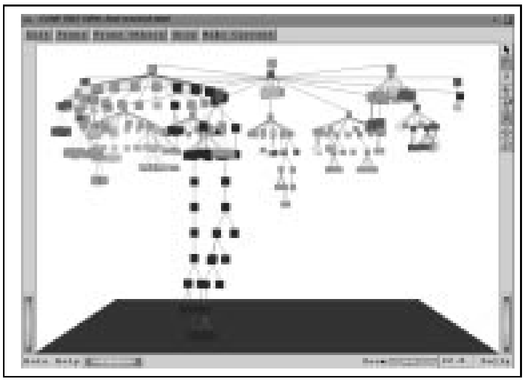

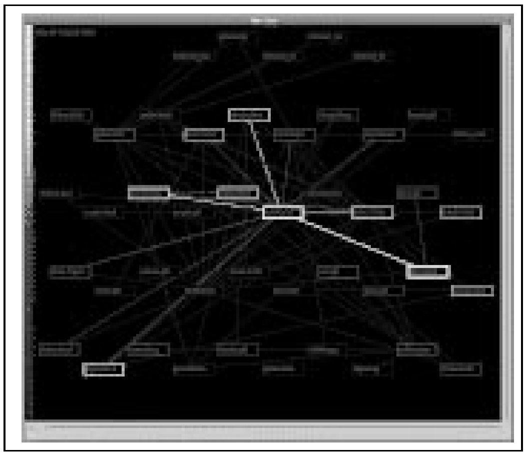

(a) (b)

(c)

Figure 5: *Further example visualisations from systems which create visualisations prior to the navigation process: (a) Navigational View Builder and (b) WebViz. (c) DeckScape is a WWW browser.*