Web Mobility Guidelines for Visually Impaired Surfers

Simon Harper, Carole Goble, Robert Stevens

Computer Science Department
University of Manchester
Manchester M13 9PL UK
simon.harper@man.ac.uk http://www.man.ac.uk/towel/

The 'Towel' project seeks to find solutions to problems encountered by visually impaired users when travelling in the World-Wide-Web (Web) by leveraging solutions found in real-world mobility. Visually impaired users find mobility on the Web particularly difficult because of the reliance of hypermedia on visual layout. Hypertext design and usability guidelines have traditionally concentrated upon navigation to facilitate this mobility; consequently other aspects of travel are neglected. This paper seeks to address these issues by extending current guidelines and design methods to include the real-world mobility concepts of orientation, memory, environment, preview and the purpose of the task at hand.

Keywords: Mobility, hypertext, visually impaired, guidelines, user agents

1. INTRODUCTION

Movement through and around complex hypermedia environments, of which the web is the most obvious example, has long been considered an important and major issue in the hypermedia design and usability field (Chen, 1997; Furuta, 1997). Indeed the commonly used slang phrase 'surfing the web' implies rapid and free movement, pointing to its importance among designers and users alike. It has also been long established (Brambring, 1984; Chieko, 1998) that this potentially complex and difficult movement is further complicated, and becomes neither rapid nor free, if the user happens to be visually impaired. This general term is used to encompass the World Health Organisation (WHO) definition of both profoundly blind and partially sighted individuals (Harper, 2000). This is because the richness of visual cues presented to a sighted user are not appropriate or accessible to a visually impaired user (Harper, 1999). Visually impaired users have a number of difficulties when interacting with this predominantly visual information. For example, a sighted user will be able to assimilate the page structure and visual cues on that page within a few seconds. This information is also continually present (on the page) for refreshing the memory of the user quickly when necessary. To fully realise the problems involved we suggest that our sighted readers start their browser and limit the window size to the top left fifth of the screen (Figure 1). Now browse a series of simple and complex web sites, and note the problems you have. We believe you'll find that the main problems encompass:

1. Can't get a feel for what's on the page;
2. Don't know how long the page is or where I am on it;

3. I become disoriented;
4. Frames, tables, spacer images, and large images become obstacles;
5. There is too much detail for my viewing area and it is too complex;
6. The whole movement and travel experience is neither satisfying nor enjoyable.

By trying to address these problems from a real world perspective we differentiate our contribution from previous ones. We assert that by appropriately applying our principles to both hypermedia and user agent design, the community can learn from our work and, by the application of this knowledge, the community can increase mobility for visually impaired web travellers.

Figure 1: Where am I?

In supporting these assertions, information from a number of interdisciplinary areas must be used. These are presented first and are split into a background section covering mobility, the Web Accessibility Initiative (WAI) guidelines, hypermedia design techniques, and web mobility and sensory translators (like screen readers). An explanation of mobility techniques and objects that those techniques are used upon will then be presented and similarities will be drawn between those present in both the real and virtual worlds. Finally, a study of these techniques and objects will enable a set of mobility guidelines to be derived and from these a number of solutions, and examples of how to implement them, presented.

2. BACKGROUND

Movement is synonymous with travel, and can be thought of as the whole experience of moving from one place to another, regardless of whether the destination is known at the start of travel or if the journey is initially aimless. Mobility is the ease at which travel can be accomplished, and also refers to items which are used to accomplish this travel. Navigation and Orientation (Farmer, 1979) are both used to achieve mobility. Orientation can be thought of as knowledge of the basic spatial relationships between objects within the environment (Bentzen, 1979). Navigation, in contrast, suggests an opportunity for movement within the local environment (Brambring, 1984). Mobility Techniques are used on mobility objects and inform both navigation and orientation. Mobility Objects are present within an environment and are used to increase mobility by having techniques performed upon them. The objects may by active (placed specifically to increase mobility e.g. signposts) or passive (placed for some other reason but still used as a mobility object e.g. pavement kerb).

Many organisations have formulated rules for web content accessibility, the American Foundation for the Blind (AFB), the Royal National Institute for the Blind (RNIB), the World Wide Web Consortium (W3C), and WebABLE, to name but four. These guidelines are however all comparable, with the W3C Web Content Accessibility Guidelines (WCAG) (WWW, 2000) covering the key points of all the others. The WCAG were formulated to make all web content accessible to

people with disabilities. These guidelines focus to a large extent on provisions for visually impaired people, as it is this user group who currently have most difficulty interacting with Web content. However, these guidelines are mainly concerned with making visual information gracefully transform to auditory information, but they do not address the issue of mobility around web sites and web pages (Goble, 2000). Although the ultimate purpose of interacting with a hypertext document is not to get to the right location but to read the location, the user is required to *reach* the information before it can be comprehended. While there has been extensive work on the usability of hypermedia (e.g. Garzotto, 1997), sometimes using the physical world as a model (Darken, 1996; Dillon, 1997; Nielsen, 1999; Spool, 1997), the focus is on sighted users. The systematic hypermedia evaluation methodology, SUE (Garzotto, 1997), includes *accessibility* (how easy it is for users to locate information) and *orientation* (a user's understanding of their current location and their own movements, and a user's grasp of their current navigation context) amongst its efficiency measures. However, accessibility concentrates on navigational richness, and link completeness rather than mobility, assuming that the user can easily travel within the web site or page. Conventionally these issues are addressed by text to speech translations that are performed by a 'screen-reader' which reads (audibly through a text-to-speech synthesiser) from the top left to the bottom right of the window or application in focus (Edwards, 1995). This is acceptable in most instances as the screen reader also translates the focused user event handler (such as a toolbar or menu system), and so the applications normal user event mechanisms can be used for interaction by the visually impaired user. This kind of translation is called 'screen scrapping' as it only produces superficial information about the text being translated, as opposed to examining the precise linguistic meaning of that text so that more complex meanings (associated with style, colour etc.) can be derived. However, when interacting with complex hyperlinked information, possibly containing multimedia resources, these screen-readers become inadequate for the purpose at hand, because of the reliance of hypermedia on context, linking, and the deeper document structure to convey information in a useful way.

Web browsers for visually impaired users have been created to access this deeper document structure, by directly examining the hypertext or the Document Object Model (DOM). These browsers, while good in certain respects, are flawed because they do not, and cannot make provision to relate mobility information about documents and clusters of documents (like web sites). This is because this information is not present or accurately derivable from an analysis of either the hypertext or the DOM. Therefore, any mobility information must be either explicitly stated within the document itself, or derivable by user agents based on a clear set of rules for analysing the Hypertext and DOM. Some technologies, such as WAB, CASTs Bobby, (WWW, 2000), do exist that process the Hypertext to aid understanding and to simplify the structure of complex documents. However these do not directly address the issue of mobility and any mobility gains are only derived as a bye-product and are therefore small and incomplete.

3. NAVIGATION AND ORIENTATION IS NOT MOBILITY

The WAI guidelines (WWW, 2000; Harper, 1998), and many other hypertext design methodologies (Garzotto, 1997; Spool, 1997), focus exclusively on navigation and orientation when addressing the issues of travel and mobility, and distinctions are not made between sighted and visually impaired users. This is important because if a distinction is not made, differences in a user's cognitive processes cannot be addressed, and these processes when forming mental maps are not the same across user groups (Harper, 1999; Dodds, 1982).

In many respects focusing on orientation and navigation leads to confusion because navigation and orientation are only parts of a picture of mobility and are by no means the whole story

(Brambring, 1984). In fact orientation and navigation are only high level classifiers to a set of techniques that represent how travel is accomplished (Bentzen, 1979; Farmer, 1979) and do not elucidate on the actual techniques themselves. Therefore, it is these techniques, and the objects they act upon that should be under investigation when working technical solutions need to be created.

3.1 Mobility Techniques

3.1.1 Preview and Probing
In real world blind mobility, a lack of preview of upcoming information is one of the major issues to be addressed. Consequently, this preview is considered to be a primary task for good mobility and is achieved by probing the environment. In a web-mobility context, the lack of previews of both upcoming hyperlinks and information relating to movement on the web page itself suggests that some degree of 'probing' must be implemented so that a limited preview can be obtained (Harper, 1999). Indeed if a user is observed traversing the web, they can be seen to select a hyperlink, preview the contents (by clicking or placing the caret over the link to see the destination) and return if the contents are not applicable. This probing is continued until each hyperlink is previewed, and interesting contents are found (Harper, 1999).

3.1.2 External Memory
Blind mobility solutions exist to accomplish obstacle avoidance and are based on both enhancing preview (as described above), planning to avoid obstacles through knowledge of the environment (orientation), and on navigating oneself around obstacles based a knowledge of one's orientation within that environment (Goble, 2000). Planning (in real world mobility) to avoid obstacles suggests a certain knowledge of an end goal to be achieved, while this is true in many cases it is not always known at the outset and related travel information may be used in transit as the goal becomes more evident. These problems are addressed in blind mobility by the provision of external memory aids like maps, guidebooks and route descriptions (Harper, 1999). This is also the case in web mobility, where a search for a specific goal may be instigated at the outset or where a user may choose to browse without much idea of a goal until well into browsing. Therefore, to find and avoid obstacles (like Feints – options that are not available can be thought of as obstacles) encountered 'on-the-fly', a web traveller, needs some form of preview. They also need to be supplied with fore knowledge of an area, or be supplied with it in-route, and have some knowledge of ones orientation within an environment. These obstacles like feints, graphics, and frames may also change with the context and task being performed; a graphic while an obstacle in the context of information searching, may be useful as a marker in the context of navigation.

3.1.3 Cueing
Orientation or 'where-ness' (detecting cyclic behaviour, direction and distance) is important in blind mobility as it enables travellers to navigate with some degree of accuracy. However, problems exist for visually impaired travellers, because they do not have the luxury of visual cues to base these judgements on. Therefore, the environment must be updated such that cues are provided in an appropriate manner, giving explicit orientation information such that navigational information can be detected (Goble, 2000). The similarities between real-world and web mobility for visually disabled people suggests that the provision of some form of explicit and appropriate orientation method (such as explicit cues) would be an advantage when travelling in the virtual web environment. This would mean that a user can make a choice as to whether they want to be at the current location and if not how to best attempt to get to their perceived destination (Goble, 2000).

3.2 Mobility Objects
To recap, preview is the ability to look ahead so that the physical or virtual environment can be detected and interacted with more easily. External memory is the ability to recognise environmental and spatial information from sources of learnt knowledge. Both preview and learnt knowledge exist in the environment as cues, and external to the environment as part of either internal or external memory such as maps and descriptions. Using this information it is therefore possible to group mobility objects within the virtual environment into a number of roles and sub-roles:

1. Cues – are objects or combinations of objects that a traveller actively users to facilitate their onward journey.
    1.1 Navigational Cues – answer a traveller's question 'Where can I?' and could be for example a signpost.
    1.2 Orientational Cues – answer a travellers question 'Where am I?' for example a unique combination of objects.
2. Obstacles – are objects that inhibit a users onward journey, however under certain conditions (such as familiarity with the object) an obstacle can change to a cue.
3. Memories – are either internal or external, and contain the knowledge that enables a traveller to decide if the object is a cue or an obstacle.
    3.1 Navigational Memory – answers a travellers question 'Where can I?', for example the results of a search on a help system.
    3.2 Orientational Memory - answers a travellers question 'Where am I?', for example a site map on the Web.
4. Out-of-view – addresses the concept of preview. Valid travel objects can be present but out-of-view. They are therefore not obstacles because they do not inhibit travel but they are not cues either because they do not facilitate travel until they come into view. Out-of-view objects are those that take over approx. 20 seconds to be spoken (in pilot studies, a sighted user can get an overview of a standard page within 20 seconds) or are below the current viewable area, and objects in the viewable area do not lead onto or suggest the presence of out of view objects.

3.3 Mobility Principles
A number of principles should also be followed when travelling the web. These principles are well known in real world mobility (Dodds, 1982; Edwards, 1995; Farmer, 1979; Green, 2000; Harper, 2000; Harper, 1998) and therefore should be used when travelling virtual worlds (like the web). They take the form of a series of questions:

• *Information Flow* – is feedback fast, appropriate and not too detailed, but detailed enough?
• *Granularity* – are there enough cues, are they close enough together and can they be found?
• *Egocentricity* – is feedback and guidance in terms of where a user is and their current focus?
• *Memory* – can I access memory appropriately and effortlessly at any point in a journey?
• *Regularity* – are travel objects deployed in a regular manner and can this be recognised and exploited?
• *Spatial* – can spatial metaphors be reformulated into a non-spatial representation?

3.4 Combining Objects, Techniques, and Principles
We can now see that objects, techniques, and principles fit together to form a coherent mobility model, that can be used by designers for web site analysis and modification, and can be encoded in a set of mobility guidelines. The relationships between these mobility items can be used such that

a cohesive view of the mobility information, tailored to the needs of a visually impaired user, can be supplied (see Figure 2).

Mobility objects such as cues, obstacles, out-of-view items, and memories, may or may not be present within a web page. These objects can be identified by using mobility techniques upon them (these techniques can be encoded within mobility instruments like software utilities or design guidelines) to decide their presence, type, and the course of action to take, when encountered. Finally, mobility techniques conform their actions to mobility principles to enable them to accurately relate relevant information to the visually impaired user.
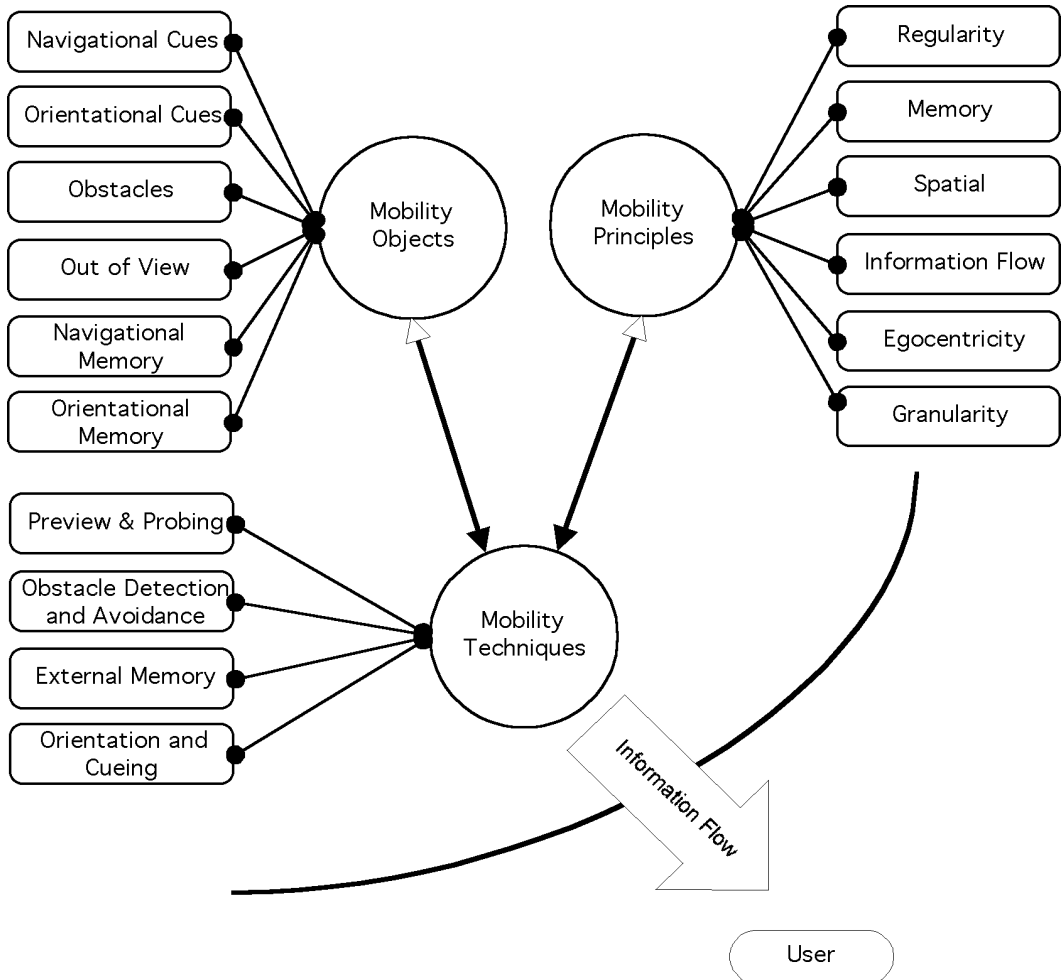

Figure 2: Combined techniques, objects, and principles Enhancing Mobility

4. ENHANCING MOBILITY

4.1 Fragmentation
Fragmentation of the hypermedia resource is key to facilitating good mobility for visually impaired people, because separating content into smaller units makes travelling through it more manageable

and meaningful (Harper, 1999). It also means that content can be reassembled in ways that are useful to the recipient and not necessarily dictated by the resource designer who is not familiar with a users individual needs. Fragmenting a page means that some type of re-joining method must be employed. This means that previews of the fragmented areas will be clustered towards the head of the page, to provide an access mechanism for each fragment. This previewing of the clustered fragment links therefore lowers the time taken for a visually impaired user to become familiar with the page content and structure.

Fragmentation is also useful in systems that have small viewable areas, like Personal Digital Assistants (PDAs) and mobile phones and communicators. This is because the problems of disorientation, the lack of good cognitive feedback, and a complex navigational system persistently faced by visually impaired users, are also faced by sighted users when interacting with these smaller devices. This is because screen area is limited and large amounts of information cannot be readily displayed as part of one whole page. The fragmentation of a hypermedia resource is made easier in Extensible Mark-up Language (XML) using the XML Fragmentation Recommendations (XFrag) (WWW, 2000). At this time, a technological solution is somewhat distant and so methods that allow fragmentation of the Hypertext Mark-up Language (HTML) are required.

4.2 Guidelines for Enhanced Mobility
Mobility within hypertext resources can therefore be enhanced by implementing good mobility techniques on good mobility objects within the hypermedia environment. In general these techniques and objects should make implicit mobility information explicit and easy to interact with, without compromising the design ethos of the hypermedia resource. These guidelines should be integrated into the WAI guidelines and are therefore couched in similar terms. Once the guidelines have been established, it will be possible to formulate technical solutions to address and investigate them by rapid prototyping.

4.2.1 Guidelines for Hypermedia Designers
- Give an overview of the content of the page so that a user can quickly re orientate themselves to the document when met again.
- Give an overview of the layout of the page so that a user can orientate themselves to the document structure therefore facilitating enhanced movement within that page.
- Provide a concise means of navigating the entire resource that does not interfere with the document content, and provide meaningful previews of linked content.
- Allow content sections of the document to be previewed at the beginning of the page so that the content of the page is easily reachable.
- Maintain the layout and regularity of the design so that user can re-orientate themselves to a familiar layout on each page.
- Make sure that mobility information feedback is appropriate (i.e. tailored to a users sensory and cognitive requirements) and useful for the cognitive processes of the intended recipient.
- Do not compromise the design ethos of the site.

4.2.2 Guidelines for User Agent Designers
- Implement systems to record and interact with journeys (in the form of pathways) undertaken by a user, both in this session and in previous sessions.
- Implement systems to enable users to easily exchange stored journeys so that users can gain knowledge from other travellers.

- Implement systems to access and interact with mobility content present in the hypermedia (placed there by hypermedia designers).
- Implement systems to derive implicit mobility content found in hypermedia.
- Make sure that mobility information feedback is appropriate (i.e. tailored to a users sensory and cognitive requirements) and useful for the cognitive processes of the intended recipient. This can most easily be accomplished by including mobility in hypermedia design methodologies.

5. IMPLEMENTING MOBILITY GUIDELINES – EXPERIMENTS BY RAPID PROTOTYPING

It is unlikely that hypermedia and user agent designers will implement mobility enhancements for visually impaired people if they either require a large amount of reworking of their implementations or if they have to overly compromise their design to include these features. Designers, being mostly sighted, also do not have a clear idea of the problems faced by visually impaired users, and so modifications to their design may seem overly complex or just not necessary (the WAI Guidelines are intended to address this last issue). Developments in user agent technology should lead in the effort to include facilities to enhance mobility in hypermedia, however, some additions to enhancing mobility can be made by hypermedia designers using current technologies.

In creating these additions our approach took into account the guidelines for hypermedia designers (previous section) but also our desire to build explicit mobility information into documents without compromising the visual design, using existing technologies (like HTML and JavaScript), and making them easily usable with screen reading devices. These additions are currently at an experimental stage and have been formulated by rapidly prototyping hypermedia documents to include explicit mobility information. Tests were then conducted to gauge the appropriateness of the experimental solution. The result of the solutions can be seen in Figure 3, and the methods of implementation can be found in Listing 1. As these solutions are experimental, a final all encompassing solution is not yet practical, however the objective of creating a starting point for further discussion has been achieved.

5.1 Proposed Additions

5.1.1 Visually Invisible Mobility Information

The key to the enhancement of hypermedia resources with mobility information is the non-interference of this type of information on the visual design itself. The way to reduce the impact of this mobility information is to make it invisible so that sighted users and designers are not preoccupied with the way this information looks or is presented. As mobility information must be placed at the top of the viewable document, so that screen readers will start to read it first, the visual effect of the information is very important. Mobility information will always take up some screen real-estate because it must have a screen presence (for rendering by a screen reader). This can, however, be minimised by using a very small font size (see Listing 1 Part Aii). While this reduces the visual effect it does not negate it and so a small, embedded Cascading Style Sheet (CSS) should be used to set the text colour to the background colour (see Listing 1 Part Ai). There are obvious problems, if using a background image; these are not now normally used on well-designed sites as they increase screen clutter. Visually invisible links can now be created to enhance mobility by displaying small amounts of explicit mobility information (see Listing 1 Part Aiii).

5.1.2 Explicit Document Structure

Firstly, a user needs to know the structure of the web page so that orientation to it becomes easier with both time and familiarity. This should be a short description of not more than 50–100 words,

```
<html>
<head>
<title>Towel - Real World Mobility on the Web</title>
<meta http-equiv="Content-Type" content="text/html; charset=iso-8859-1">
<SCRIPT LANGUAGE=JAVASCRIPT>
<!– – Hide
function docStructure() //Part Bi
{
ds=window.open(' ',"replace","toolbar=no,personalbar=no,directories=no,location=no,resizable=no,width=400,height=180")
ds.document.write("<html><head><title>Document  Structure  Window</title></head><body bgcolor='#00CCCC'
text='#000000' link='#FF8000' vlink='#000000' alink='#FFCCCC' background='bckg.jpg'><font face='Verdana, Arial,
Helvetica, sans-serif'>")
ds.document.write("<p>The page is divided into a 3 by 3 matirix. The top row contains title information and a title banner, the
bottom row contains links to the university department and group pages. The Central Row contains navigation to the left,
content to the center and the right is blank.</p>")
ds.document.write("<b><A HREF='javascript:window.close()'>Close the Document Structure Window</A></b><br>")
ds.document.write("</font></body></html>")
ds.document.close()
ds.focus()
}
function docContent() //Part Ci
{
dc=window.open(' ',"replace","toolbar=no,personalbar=no,directories=no,location=no,resizable=no,width=400,height=180")
dc.document.write("<html><head><title>Document Content Window</title></head><body bgcolor='#00CCCC' text='#000000'
link='#FF8000' vlink='#000000' alink='#FFCCCC' background='bckg.jpg'><font face='Verdana, Arial, Helvetica, sans-serif'>")
dc.document.write("<p>The page contains an intial description of the purpose of the Towel Project.</p>")
dc.document.write("<b><A HREF='javascript:window.close()'>Close the Document Content Window</A></b><br>")
dc.document.write("</font></body></html>")
dc.document.close()
dc.focus()
}
function siteMobility() //Part Di
{
sm=window.open('sitemobility.html',"siteMob","toolbar=no,personalbar=no,directories=no,location=no,resizable=no,width=400
,height=250")
sm.focus()
}
// End Hidding – –>
</SCRIPT>
<STYLE TYPE="text/css"> //Part Ai
<!– –
A.mobility:link     { color: #CCCCFF }
A.mobility:visited { color: #CCCCFF }
//– –>
</STYLE>
</head>
<body bgcolor="#CCCCFF" text="#000000" link="#FF8000" vlink="#000000" alink="#FFCCCC">
<FONT COLOR="#CCCCFF" SIZE="-7"> Mobility Options Begin //Part Aii
//Part Aiii - Part Bii
<A HREF="javascript:docStructure()" CLASS="mobility">Document Structure (42 Words)</A>
<A HREF="javascript:docContent()" CLASS="mobility">Document Content (13 Words)</A>
<A HREF="javascript:siteMobility()" CLASS="mobility">Site Mobility (8 Links)</A>
Mobility Options End. Pause your speech engine now.</FONT>

<table width="100%" HEIGHT="100%" border="0" cellspacing="0" cellpadding="10" bgcolor="#CCCCFF">

Continued…
```

Listing 1: HTML from http://www.man.ac.uk/towel/

otherwise it will take too long to read. This description should also be linked to a preview hyperlink, so that a user is not required to sift through a description of the document if they are already familiar with it (see Listing 1 Part Bii). In our experiments we used JavaScript to open a browser window that is easy to close, displaying the document structure description (see Listing Part Bi). This is also useful because it takes the window focus and then returns it to the opening window when the child generated by JavaScript is closed.

5.1.3. Explicit Content Overview

An explicit content overview is useful so that a user immediately knows what the document in question is and whether they want to investigate it further. This is difficult for visually impaired people, because they are not able to visually glance at the document and so have to wait for a verbose reading of the content, which can be disorientating. A preview should again be used (see Listing 1 Part Bii) and the technique of opening a JavaScript window can also be employed (see Listing 1 Part Ci).

5.1.4 Explicit Site Mobility

Because screen readers use the browsers own mobility mechanisms, site navigation can become complex. This is because many browsers use the 'TAB' key to move from hyperlink to hyperlink in a top left to bottom right motion and so site navigation is mixed up with many other types of hyperlinks. This is especially the case where site links to sections occur on either the right or left of the content, or when the content is split into columns (hyperlinks are discovered out of sequence and not with the content they relate to). Therefore a separate hypertext site navigation and orientation page should be used and again linked to a preview (see Listing 1 Part Bii) and displayed in a JavaScript generated window (see Listing 1 Part Di).

5.1.5 Explicit Page Mobility

Page mobility is by far the hardest issue to address, due to the normally high complexity of visual information found in a hypermedia page or document. This complexity can be simplified by

Figure 3: Towel index page with popup document structure window

fragmenting information on the page into useful mobility objects and then allowing those objects to be quickly derived and interacted with based on previewing information. This cuts down on the information present at once and enables a visually impaired user to orientate themselves to the information more easily. It also enables the preview objects to become cues to aid mobility and allows obstacles to mobility to be excluded from the preview.

Problems exist with the fragmentation of the page using current technology; both in user agents and web language recommendations no suitable fragmenting schema can be derived. Therefore while fragmentation is not currently possible in any meaningful way, new hypermedia resources should be built using XHTML which will eventually allow XFrag to be used to enhance the mobility within these documents.

6. DISCUSSION

Applying knowledge about real world mobility to web based mobility problems can enhance the travel experience for visually impaired users. Because all users share some of the characteristics of visually impaired travellers, when travelling off the viewable area, these enhancements can also be included for sighted users with no additional cost. The use of preview and appropriate knowledge feedback will increase the mobility of users within many virtual journeys, and therefore solve many of the mobility problems encountered frequently in web based travel. Lessons learned with regard to web mobility design can be incorporated into future design methods and best practice. By doing this, solutions established for visually impaired travellers in the real world are applied to all web travel, thereby solving problems faced by all web travellers, sighted or unsighted.

The WAI guidelines propose a long and difficult validation process for inclusive web page and site design. However, they do not examine the concept of mobility to any great extent. The WAI particularly promote the appropriate and extensive use of mark-up and therefore it is envisaged that the use of XML and XSL will become a more appropriate technology for implementing these mobility concepts. XML should encode the travel objects in a DTD, so objects within web content know their travel role. XSL would tell a user agent how to present travel objects, but this is not used in a formal context. It merely enables the user agents to identify and classify objects more easily and less speculatively. Of course as more browsers support both Visual and Aural Cascading Style Sheets visually invisible information should be more easily written into a page.

This work has been successful in leveraging real world mobility solutions to build mobility guidelines for the virtual world. However, solutions for all the problems and guidelines identified are not possible using current technologies. Therefore, further work will be undertaken on developing solutions for XML and XHTML documents and by developing a prototype mobility agent. This will be made available as an opensource implementation so that it can be included (and, or modified) in commercial user agents and web pages if required.

REFERENCES

BENTZEN, B. L. (1979): *Orientation aids. foundations of orientation and mobility,* WELSH, R. L. *et al* (eds). New York: American Foundation for the Blind Press:291-355.
BRAMBRING, M. (1984): Mobility and orientation processes of the blind. *In Electronic Spatial Sensing for the Blind.* WARREN, D. H. and STRELOW, E. R (eds). Dordrecht; Lancaster: Nijhoff, USA:493-508.
CHEN, C. (1997): Structuring and visualising the www by generalised similarity analysis. *Proceedings of the Eighth International ACM Conference on Hypertext and Hypermedia,* New York: ACM Press.
CHIEKO, A. and LEWIS, C. (1998): Home page reader: IBM's talking web browser, *Proceedings of the Closing the Gap Conference.*
DARKEN, R and SIBERT, J. (1996): Wayfinding strategies and behaviours in large virtual worlds. *Proc. ACM CHI'96 Conference on Human Factors in Computing Systems,* New York, USA :132-149 ACM Press.

DILLON, A and VAUGHAN, M. (1997): 'It's the journey and the destination'…. *The New Review of Hypermedia and Multimedia* 3: 91-106.

DIX, A., FINLAY, J., ABOWD, G., and BEALE, R. (1993): Human computer interaction, London: Prentice Hall.

DODDS, A. G. (1982): The mental maps of the blind. *Journal of Visual Impairment and Blindness,* January Pt 76: 5-12. SHIVELY, S. (ed). New York, American Foundation for the Blind Press.

EDWARDS, A. D. N. (1995): Extra-ordinary human-computer interaction, Cambridge, UK: Cambridge University Press.

FARMER, L. W. (1979): Mobility devices. Foundations of orientation and mobility, WELSH, R. L. *et al* (eds). New York: American Foundation for the Blind Press:357-402.

FURUTA, R. (1997): Hypertext paths and the www: Experiences with Walden's paths. *Proceedings of the Eighth International ACM Conference on Hypertext and Hypermedia.* New York: ACM Press.

GARZOTTO, F and MATERA, M. (1997): A systematic method for hypermedia usability inspection. *The New Review of Hypermedia and Multimedia* 3: 39-65.

GOBLE, C., HARPER, S. and STEVENS, R. (1999): Towel: The travails of visually impaired web travellers. *Proceedings of the 11th ACM Conference on Hypertext and Hypermedia* 2000: 1-10. New York: ACM Press.

GREEN, P. and HARPER, S. (2000): An integrating framework for electronic aids to support journeys by visually impaired people. *Proceedings of 7th International Conference on Computers Helping People with Special Needs* (UNESCO).

HARPER, S. and GREEN, P. (2000): A travel flow and mobility framework for visually impaired travellers. *Proceedings of 7th International Conference on Computers Helping People with Special Needs* (UNESCO).

HARPER, S. (1998): *Standardising electronic travel aid interaction for visually impaired people.* MPhil. Thesis. University of Manchester Institute of Science and Technology at Manchester, UK.

HARPER, S., STEVENS, R. and GOBLE, C. (1999): Towel: Real world mobility on the web. In *Computer-Aided Design of User Interfaces II.* VANDERDONKT, J. and PUERTA, A. (eds). Kluwer Academic Publishers:305-314.

HARPER, S., STEVENS, R. and GOBLE, C. (2000): A pilot study to examine the mobility problems of visually impaired users travelling the web. *The ACM Special Interest Group on Computers & the Physically Handicapped Newsletter.* KARSHMER, A.L. (ed). 68, ACM SIGCAPH, September 2000.

NIELSEN, J. (1999): How to write for the web. http://www.useit.com/papers/webwriting/writing.html

SPOOL, J.M. SCANLON, T., SCHROEDER, W., SNYDER, C. and DeAngelo, T. (1997): *Web site usability: A designer's guide, user interface engineering,* Mass. http://world.std.com/~uieweb/index.html

VARIOUS. (1996): *Factsheet, A short guide to blindness,* http://www.rnib.org.uk

WWW CONSORTIUM – http://www.w3c.org

BIOGRAPHICAL NOTE

Simon Harper is a software engineer with Manchester Computing providing integrated hypermedia and database solutions in a teaching and learning environment. He comes from an industrial background as a software consultant with a major energy provider. Simons research interests are a synthesis of real world mobility, pervasive devices, and rehabilitation engineering. He is currently applying these interests to the fields of HCI and hypermedia as part of his PhD thesis with the Department of Computer Science at the University of Manchester.

Carole Goble is a Professor of Computer Science at The University of Manchester where she has been on the academic staff since 1985. Her current interests are in metadata, knowledge representation and ontologies, and their application to multimedia, conceptual hypermedia, information integration, intelligent user interfaces, and intelligent retrieval. Her work is applications driven, chiefly in bioinformatics and document collections. Recently she has been part of the OIL consortium working on a language for representing knowledge for the Semantic Web.

Robert Stevens is a research associate in the Information Management Group of the Department of Computer Science at the University of Manchester. Robert has degrees in the Biological Sciences and Computer Science. Robert's D.Phil and early post-doctorial work was on using prosodic qualities of speech and browsing techniques to promote active reading of complex material by visually disabled people. His current research is in the area of bioinformatics, but he still has a keen interest in HCI issues.