New Review of Hypermedia and Multimedia,
Vol. 11, No. 1, June 2005, 103–128

Open Paper

Augmenting the mobility of profoundly blind Web travellers

SIMON HARPER*, CAROLE GOBLE and ROBERT STEVENS
Information Management Group, School of Computer Science, University of Manchester, Manchester, UK

Use the word 'accessibility' in the presence of any HCI[1] specialist and they will immediately think of creating open interfaces that can be accessed both visually and audibly. Further, mention 'accessability' to any forward thinking group of Web developers and they will start to quote the Web Accessability Initiative Guidelines (WAI) and extol the virtues of accessability checking tools like 'Bobby'. Either way, both groups will focus on the obviously important area of 'sensory translation' but will miss one fundamental truth; profoundly blind people interact with their environment in a markedly different way from that of sighted individuals. We have realized that the ease of movement (mobility) around systems and information space (the hypertext/Web docuverse) is central to good accessibility; and that to achieve this we require additional mobility semantics within systems and information as a way of enhancing the user experience. By adding small amounts of information to existing Web pages (semi-) automatically, we can show significant improvements in the amount of information profoundly blind users are able to access in a given time; in effect 'levelling the playing field' with sighted users. This paper discusses our work and demonstrates how we can make such a claim.

1. Introduction

David doesn't go to his local bookstore on a Sunday afternoon even though he's an avid reader; he doesn't buy newspapers although he's very interested in world affairs; and he doesn't watch TV to find out what is happening locally even though he's very community minded. David is blind and so he uses the largest source of electronic (and therefore accessible) literature, news and information available—the World Wide Web (Web). However, David has a problem, because while many of the sites he visits are becoming more accessible by conforming to the Web Accessibility Initiative (WAI) guidelines he still finds himself disoriented, confused and frustrated. David tells us he is unable to get an overall feel for what's on the page, becomes disoriented in the information, and is unable to decide which option to take when 'tabbing' through links. While his experiences are not common for all profoundly blind (here used to refer to the World Health Organisation (WHO) definition being the inability to distinguish fingers at 3 metres (Authors), 1996)) users they do

*Corresponding author. Email: simon.harper@manchester.ac.uk
[1] Human Computer Interaction or Computer Human Interaction (CHI).

New Review of Hypermedia and Multimedia
ISSN 1361-4568 print/ISSN 1740-7842 online
http://www.tandf.co.uk/journals
DOI: 10.1080/13614560500191220

represent the majority experience of David's peer user group. 'Towel-Transcoding'[2] is our attempt to address David's problems of mobility[3] on the Web, and is based on our previous work in Hypermedia and Real-World Mobility. By using combined Hypertext and mobility paradigms we move toward a system that will assist David in his travels around the Web.

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, Chen and Czerwinski 1997, Furuta 1997). In fact spatial hypermedia systems are designed to enable an overview of the hypermedia environment, but how are these systems useful to profoundly blind users—a group of individuals who most need them? They are not, because most spatial hypermedia systems are based on the structural semantics sighted users require and the requirements of profoundly blind users—whose needs are quite different—are not taken into account. Web users commonly use the slang phrase 'surfing the Web' which implies rapid and free movement, pointing to its importance among designers and users alike. However, it has also been long established (Brambring 1984, Chieko and Lewis 1998) that this potentially complex and difficult movement is further complicated, and becomes neither rapid or free, if the user is profoundly blind.

Work, including ours, has shown that profoundly blind users are hindered in their efforts to access the largest repository of electronic information in the world, namely the World Wide Web (Harper *et al.* 2001). Also that profoundly blind users are at a severe disadvantage, when moving around the Web (Harper *et al.* 2000) compared to their sighted counterparts. We suggest that the 'playing field' is not yet level because the Web usability[4] community has typically concentrated on sensory translation and the implementation of guidelines that focus on this translation. The absence of suitable guidelines (W3c 2005), design and evaluation methodologies (Garzotto and Matera 1997), technical implementations (Spool 1997, Chieko and Lewis 1998) and work on holistic views of the Web docuverse (Darken and Sibert 1996, Dillon and Vaughan 1997, Chen 1997, Chen and Czerwinski 1997, Furuta 1997) all hinder profoundly blind users because they miss the fact that profoundly blind people interact with their environment in a markedly different way from that of sighted individuals (Globe *et al.* 2000, Harper *et al.* 2000). The 'Towel' project was set up to overcome these differences keeping in mind one simple goal:

---

[2] The project name 'Towel' is taken from the popular novel – 'The Hitchhikers Guide to the Galaxy' written by Douglas Adams which states that 'A towel is about the most massively useful thing an interstellar hitchhiker can have. You can wrap it around you for warmth as you bound across the cold moons of Jaglan Beta ... wet it for use in hand-to-hand combat ... wrap it around your head to ward off noxious fumes ... any man who can hitch the length and breadth of the Galaxy, rough it ... win through, and still know where his towel is, is clearly a man to be reckoned with.'

[3] Mobility (as in "quality") *n.* : the quality of moving freely.

[4] ... and therefore the HCI community.

> To make 'surfing the Web' mean 'rapid and free' for profoundly blind users like David.

The Towel project takes its antecedence from hypermedia models which describe hypertext topologies in terms of three- and four-dimensional landscapes (Darken and Sibert 1996, Dillon and Vaughan 1997, Furuta 1997, Jul and Furnas 1997, Lowe and Hall 1999, Maglio and Barrett 1998). These landscape metaphors are part of the overarching view of 'Spatial Hypermedia' which goes beyond the traditional node-link model and uses spatial properties—such as geometric and temporal placement and visual similarity—to express and convey relationships and thus semantics. We reason that if landscapes are present then successful traversal techniques from the real world could be applied to the virtual world—thereby, solving some of the orientation and navigation problems faced by profoundly blind users when moving through hypermedia resources. To this end we have surveyed visually disabled users to ascertain the problems they face (Harper *et al.* 2000) and formulated a framework model to address these problems (Harper *et al.* 2001).

This paper brings us further in that it shows the steps we have now taken towards automatically applying this model through the use of annotation and transcoding techniques to a common use hypertext system—the Web. Here we look at the technical detail behind our solution (the Towel Tool) and analyse the results with a summative evaluation.

1.1 Motivational example/problem

To fully realize the problems users with low vision encounter, we suggest that our sighted readers start their browser and limit the window size to the top left fifth of the screen (see figure 1). Now browse a series of simple and complex Web sites by jumping from hyperlink to hyperlink (Using the tab key) and disregard everything that is not described in that hyperlink. In our example we have used our own university's news site. Now note the problems you have; we believe you will find that:

1. you cannot get an overall feel for what is on the page;
2. you do not know where you are on the page or if you've been there before;
3. consequently, you become disoriented;
4. you cannot tell where you are going to move to once the hyperlink is selected;
5. frames, tables, spacer images and large images become obstacles;
6. there is too much detail for your viewing area and it is too complex;
7. you cannot tell if the information on the target page is the information you are expecting or require;
8. you cannot tell if the link is referential, associative or structural;
9. you often find you tab from link to link all with the same name and in effect become lost in the repeated 'Read More ...', 'Read More ...' anchors;

Figure 1. Where Will I Go? (http://news.man.ac.uk/).

10. the whole movement and travel experience is neither satisfactory nor enjoyable and it certainly is not rapid and free.

The problem is compounded if the reader is profoundly blind and must rely on text-to-speech technology (covered in section 3). In this case the software starts to speak when a webpage is loaded. In the case of a typical site[5] – "support the imdb comma visit our sponsors period [repeatedly until the reader broke out of the frame] graphic5; graphic5; graphic 5; graphic 361; search …" – this is plainly unsatisfactory.

1.2 Synopsis

We have studied the way profoundly blind people move around physical space and interact with the built environment. We have then transposed these models, created for mobility in the physical world, to the Web. These models provide the mobility semantics we need for enhanced interaction and, with some additional coding, can then be inserted into various Web pages. Finally, we create a Netscape plug-in to translate these pages into a format that enhances a profoundly blind users cognitive experience. The paper can be summarized as follows.

What is Browsing We investigate the browsing behaviour of Web surfers to support our assertion that browsing has implicit mobility aspects; and find from related literature that browsing behaviour suggests that mobility is

[5] In this case http://www.imdb.com.

not only involved in the activity, but is also important for the successful completion of the activity.

Mobility and the Web Here we summarize the differences between the way profoundly blind people move and the way sighted users move. We distill this into six mobility principles and assert that these principles should be followed in any system to help profoundly blind users move. Finally, we look at the techniques needed by users and tools if rapid and free movement are to be supported. These principles, captured in our techniques, can then be used within our tool.

On Transcoding Transcoding is a technology used to adapt Web content so that it can be viewed on any of the increasingly diverse small screen devices currently available. Tightly coupled to device independence, we briefly review the state of current research in transcoding. We look at how it has been used for a number of years to make incomplete or badly written hypertext accessible to profoundly blind users and their accessibility technologies. This is important because we can use aspects of this technology to transform Web pages on-the-fly to include more mobility information.

The Towel Tool Next, we present our own experiments which attempt to support enhancements aimed at filling the current gaps in content accessibility and enabling better mobility. Our chief goal was the design of appropriate travel tools within user agents so that page structure can assist the user and not hinder them, as is currently the case. Such tools assist in the transformation of these hindrances and actively support mobility rather than only navigation or content rendering.

Evaluation While extensive evaluations and field trails are to be carried out on the plug-in we have performed some initial testing. These evaluations were performed along the two tracks of: (1) evaluation by the user; and (2) evaluation by reapplication of the mobility framework. Here, we discuss these evaluations and decide on what changes need to be made to our techniques.

Conclusion Finally, we focus on our conclusions from the work undertaken and look at future work including system evaluations.

2. Browsing

Browsing is an activity that is difficult to define (Carmel *et al.* 1992), but there is general agreement that 'we all browse in various context to make sense of the world around us' (Chang and Rice 1993). Some researchers also describe it as a process of 'picking out bits and pieces · · · selecting worthwhile information need or interest' (Cove and Walsh 1998). Different disciplines look at browsing from different perspectives (Chang and Rice 1993). Various reviews suggest that browsing is a kind of searching, in which initial search criteria or goals are only partly defined or known in advance. Browsing involves scanning, which has been described as looking, examining, or sampling, during which the person's body or eyes move smoothly at will

(Morkes and Nielson 1997). Browsing also involves distinct (Chang and Rice 1993) consumer shopping behaviour that is related but not equated with buying behaviour. Methodologically, eye movement can be a useful indicator of browsing and has been used to test the effect of different page layout or catalogues on browsers' attention (Chang and Rice 1993). Browsing is fundamentally scanning and has been related to environmental perception and cognition. For example, sightseeing is environmental browsing as perceptual experience (Chang and Rice 1993).

Current research supports the notion that browsing is movement in the information space where the user is in control of what to read or examine. While chance or synchronicity may have some part to play in browsing behaviour the user is still in control of filtering the information presented. Many studies have addressed different browsing types and strategies (Maglio and Matlock 1998) and support the view that movement is 'the' essential characteristic of browsing (Bates 1989). We all browse in various contexts picking out bits and pieces of information and selecting worthwhile information (Maglio and Barrett 2000); and we accomplish this by using a searching and scanning behaviour over organizations of the material (Obendorf and Weinreich 2003, Morkes and Nielson 1997), interfaces to that material (Barrett *et al.* 1997), and feedback about the material (Marchionini 1995). However, this browsing behaviour can become both complex and confusing due to the nonlinear nature of hypertext structures. To address this confusion a number of solutions have been presented with the main thrust being split into two 'camps': (1) the saving of a series of hypertext links to form a pathway through information (Furuta 1997, Lowe and Hall 1999); and (2) the presentation of a 'Landscape' analogy to aid navigation and orientation in a Web area (Darken and Sibert 1996, Dillon and Vaughan 1997). Navigation by using this second form—the landscape analogy—is advantageous because it plays on an individual's already developed skills of orientation and navigation within a physical environment. These skills can then be carried through to the virtual environment using familiar metaphors to enable orientation in Hyperspace and navigation through complex hyperlinked information. Different mechanisms are used to generate these 'Hyperspace Landscapes', however, the requirements for easy navigation and orientation are still in all cases the goal. While these ideas are useful for sighted users, they do not relate well to the spatial orientation processes of profoundly blind users who demand egocentricity such that descriptions of distance and journey, route and survey knowledge, become associated with the traveller and not the environment. For example, body rotation is used to describe parts of a journey, more temporal and egocentric terminology is used along with less spatial and environmental terminology in defining points, explicit statements on distance are also made more often, and the specification of distance and direction is far more exacting.

We take our antecedence from these hypermedia visualization techniques which describe hypertext topologies in terms of '*N*-dimensional' landscapes. We reason that if landscapes are present then successful traversal techniques from the real world could be applied to the virtual world. In this way some of the orientation and navigation problems faced by profoundly blind users can be solved. We describe these mobility techniques and their relationship to browsing and visualization techniques in the next section.

3. Mobility and the Web

Travel and mobility are both conceptually and practically complex, but are often not considered so during everyday use. To a sighted user, in the real world, travel and mobility seem to be easy and intuitive. It is, however, vision which makes mobility seem easy, as is revealed when the mobility of profoundly blind people is studied. We suggest that it is precisely because travel and mobility are often ignored that journeying terminology needs to be enunciated.

The available literature suggests profoundly blind people interact with their environment in a markedly different way from that of sighted individuals (Lloyd 1972, Downs and Stea 1973, Downs and Stea 1977, Gibson 1979, Howard 1982, Gentner and Stevens 1983, Mc Tear 1988, Farah and Ratcliff 1994, Gazzaniga and Bizzi 1995, Akins 1996). However, similarities in user behaviour do exist between: (1) sighted users interacting with small screen devices or traversing hypertext resources; (2) users with low vision using screen magnification technology[6]; and (3) profoundly blind users interacting via screen readers[7]. In these cases we could say that all users (1 and 2) share commonalities with profoundly blind users (3) when their vision is handicapped by technology. Because the communication bandwidth from web page to blind user is so much reduced, compared with that of a sighted user, the information that is conveyed has to be exactly the information required by the user. Sometimes this will be page description information and sometimes it will be the content of the page. This rationale forces the formulation of a set of concepts describing the way blind people interact with their environment and summarized here into five key mobility requirements.

Information flow Physical travel aids perform a probing task such that a limited amount of preview and feedback is given. Sighted people navigate in physical space by ignoring detail; profoundly blind travellers use simpler information more frequently than complex information.

Granularity Profoundly blind travellers have limited preview of coming objects or obstacles. To compensate for this lack of preview, the route is broken into a greater and more complex number of stages, route descriptions are in finer detail and obstacle information is more explicit.

[6] Screen magnification software enlarges the information on the screen by pre-determined incremental factor [for example, 1x magnification, 2x magnification, 3x magnification, etc,]. Magnification programs run simultaneously and seamlessly with the computer's operating system and applications.

[7] Screen readers are special applications that vocalise the onscreen data. Pages are typically read from the top left to the bottom right, one word at a time.

Egocentricity Many profoundly blind people have a tendency to think of the real world in an 'egocentric' manner, such that descriptions of distance and journey, route and survey knowledge, become associated with the traveller and not the environment.

Mental maps Profoundly blind people have an increased use of cognitive or mental maps of route and survey knowledge.

Regularity and familiarity of environment Profoundly blind people normally only travel independently in environments with regular features and would normally only travel unassisted in areas that were familiar to them. As with preview, this is an issue of predictability.

To cope with these requirements profoundly blind travellers have evolved a set of techniques for enhancing rapid and free movement. These techniques are enacted upon objects within the environment (in the case of the Web—Web page furniture and real-estate) and as such form the basis of any implemented solution. These techniques are used to both decipher the objects within the page and to implement appropriate actions to make them useable by visually impaired users.

Preview and probing: 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 *et al.* 1999). Indeed a user observed traversing the Web 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, which suggests that to avoid unrequired information encountered 'on-the-fly' a Web traveller needs some form of detection and avoidance schema based on accurate and appropriate previews.

Obstacle detection and avoidance: In Web mobility, where to find and avoid obstacles (like Feints[8]—options that are not available can be thought of as obstacles such as disabled menu items) encountered 'on-the-fly' a Web traveller needs some form of detection and avoidance schema. They also need to be supplied with foreknowledge of an area, or be supplied with it in-route, and have some knowledge of one's orientation within an environment. 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.

External memory: A user can be provided with a number of different external memories to facilitate a more informed journey. A user should know the structure of the Web page so that orientation to it becomes easier with both time and familiarity. An explicit content overview is useful so

[8] Feints (as in "fake") *v.* : deceive with a feint.

that a user immediately knows what the document in question is and whether they want to investigate it further. This is difficult for profoundly blind 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. Site navigation can become complex because many browsers mix up site navigation 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).

Orientation by cueing: The similarities between real-world and Web mobility suggests that the provision of some form of explicit and appropriate orientation method (such as explicit cues) might 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 intended destination (Globe *et al.* 2000).

Once we have discerned the status of the encountered object we need to manipulate it into a form consistent with our technique set (above). The manipulation is called transcoding.

4. On transcoding

Simply, transcoding is a technology used to adapt Web content so that it can be viewed on any of the increasingly diverse devices found in today's market. Although the device independence of information has been key in transcoding it has been used for a number of years in the context of making incomplete or badly written hypertext accessible to profoundly blind users and their accessibility technologies. Transcoding in this context normally involves: (1) syntactic changes like shrinking or removing images (Hori *et al.* 2000a), (2) semantic rearrangements and fragmentation of pages based on the meaning of a section (Meyers; Codix), (3) annotation of the page created by a reader (Buyukkokten *et al.* 2000) (4) and generated annotations created by the content management system (Buyukkokten *et al.* 2000).

There are a number of different ways that transcoding can take place. In one example, the original material (an HTML document, for example) is analysed by a program that creates a separate version containing annotations. The annotations include information that will instruct the reformatting process and inaccessible elements will be removed or altered. Systems are often based along similar lines and address set problems, some are annotation based (Hori *et al*. 2000a), others generate text only versions (Myers; Codix 2005), some filter the content (Circa 2004), and others are specifically used for small scale device interaction (Buyukkokten *et al.* 2000). Whatever system is used it invariably does not transform all inaccessible elements but just a subset leaving holes in the accessibility of their transcode.

4.1 Annotation

The goal of annotations for Web content transcoding is to provide better support either for audio rendering, and thus for profoundly blind users, or for visual rendering in small screen devices. The problem of rendering Web pages in audio has some similarities to the problem of displaying Web pages on small-screen devices. For example, in both cases, only the small portion of the page is viewable at any point. However, there are major differences and requirements. Although the amount of information that could be accessed at once in a small-screen device is also limited, the interaction is still visual. The provided visual rendering is still persistent, and the screen acts as an external memory, as opposed to audio rendering which is transient. Additionally, compared to visual rendering, audio is less focused and more serial in nature, the user cannot easily and quickly shift the focus.

Various proxy-based systems to transcode Web pages based on the external annotations for profoundly blind users have also been proposed (Asakawa and Takagi 2000, Asakawa and Takagi 2000). The main focus is on extracting visually fragmented groupings, their roles and importance. Eight different roles such as proper content, header and footer are proposed for annotation. These roles are mainly at an abstract level and are not rich enough to fully annotate the page to enhance the mobility support. They do not support deep understanding and analysis of pages, and in consequence the supported transcoding is constrained by these proposed roles.

Other work centres on small-screen devices and proposes a system to transcode an HTML (Hypertext Mark-Up Language) document by fragmenting it into several documents (Hori *et al.* 2000b). The transcoding is based on an external annotation framework. Since the focus is the small-screen devices, physical and performance constraints of the devices need to be considered, such as screen size, memory size and connection bandwidth. However, these are not the main requirements of the users accessing Web pages in audio and there are differences as explained above.

4.2 Semantic transcoding

In semantic transcoding, the semantics provide the machine understandability and knowledge reasoning and the transcoding provides the transformation technique. However, current systems are at present limited to page analysis (Seeman 2004), where a page built after a set template can be analysed and transformed by semantic or semantic-like technologies. Annotation is another popular way to add semantics to legacy systems. These applications rely on the XHTML page being annotated from a set ontology or taxonomy. These annotations are normally stored in an annotation database and are then used when the page is processed to make the transformations (Yesilada *et al*. 2003). Even the few systems that try and remove themselves from XHTML (Extensible Hypertext Mark-Up Language) presentation (Huang and Sundaresan 2000) to a more service oriented (XML—Extensible Mark-Up Language) model still require the data source to be modified. Therefore, current semantic transcoding techniques do not provide us with replicable solutions to the problems of manipulation of heterogeneous nonlinear information sources. Indeed their over reliance on user annotation can hinder their cognition as these annotations are often stored as part of a separate database which is not updated when the page changes. Therefore disparities between page and annotation can be inadvertently introduced.

5. The towel tool

Even if 'movement' is superficially addressed in some guidelines, design methodologies and technical implementations, their focus is mainly on navigation (and to some extent orientation). However, in many respects even this cursory attention 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 mobility techniques, acted upon objects and conforming to mobility principles. These objects, techniques and principles fit together to form a coherent mobility model that can be transformed to a framework for Web site analysis and modification. The relationships between these mobility items can be used such that a cohesive view of the mobility information, tailored to the needs of a profoundly blind user, can be supplied. W3C[9] WAI, WebABLE[10], RNIB[11], AFB[12] and other guidelines focus on sensory translation by the graceful conversion of visual to auditory information. Hypertext design and evaluation methodologies, like the Hypertext design Methodology (HDM) and Systematic Usability Evaluation (SUE) (Garzotto and Matera 1997), pay little attention to the interaction needs of profoundly blind people. And Web browsers for profoundly blind users focus on examining the Document Object Model (DOM) or the HTML to present this information audibly, and do not consider implicit (or explicit) structural semantics contained within the document (Chen 1997, Chaomei Chen and Mary Czerwinski 1997, Furuta 1997); or how this information can be triaged to provide the user with 'exactly the information (no more and no less) than they require at that particular moment'.

The Towel tool aims to put back these structural semantics by annotating the structural elements (objects) within a page to a level of fine granularity. We then apply our techniques, encoded within our tool, to these pages and transcode/triage them into new accessible forms.

5.1 Towel in practice

It has been shown that the environment has a strong influence on how mobile a traveller can be (see section 3). The hypermedia environment can be defined

9 World Wide Web Consortium.

10 A Web accessibility company supplying De-Facto accessibility guidelines.

11 Royal National Institute for the Blind – UK Organisation.

12 American Foundation for the Blind – USA Organisation.

in terms of the page, site and the user agent (being the combination of the browser and the device used to present the information in an appropriate sensory form, like a screen reader) (Harper *et al.* 1999).

Our preliminary studies suggest that profoundly blind users take far longer to orient themselves within, and navigate around, a Web environment. While sighted users can orient themselves to the page quickly, it often takes profoundly blind users roughly ten times as long to perform the same task. Profoundly blind users also have problems describing the layout and structure of the pages accurately, and this includes correctly distinguishing selectable and non-selectable menu items.

The main problems seem to occur due to the complexity of the screen format and the unavailability of a structural overview of both the page and the site. There is, however, a correlation between the proximity of an element to the top left corner of the screen and the speed of a profoundly blind user's access to that element, which is on average faster than elements to the bottom right corner. These problems faced by profoundly blind users can be directly associated with those encountered in the real world and are concerned with seeing into the distance and the appropriateness of the information either stored (as maps, descriptions, etc.) or returned from the environment as it is traversed.

In order to find solutions to these problems we must convert our techniques, objects and principles to a form that can be used to increase profoundly blind mobility on the Web. We do this by

- creating a set of *'Web Specific Techniques'* for use in formulating virtual world design solutions (see section 3 starting **preview and probing**).
- Defining a minimum set of *'Design Solutions'* (see section 5.1.1 and table 1) that will enable the standard generic mobility items to be specifically applied to the World Wide Web, thus supporting our mobility framework and adapting it to be implementation specific.
- Formulating a set of *'Tests'*—based on our principles—to be met when examining our proposed solution (see table 2). By using this 'Test of Good Mobility' derived from our mobility principles we are able to formulate some answers regarding the results returned from the transcoded pages (see section 6.3).

5.1.1. Design solutions.

Fragmentation Fragmentation of the hypermedia resource is key to facilitating good mobility for profoundly blind people because fragmentation encapsulates and facilitates the ideas of preview and probing by splitting large complex documents into smaller more manageable parts. This fulfils our objectives of making the environment more regular, increasing information flow, supporting granularity and adapting to egocentric behaviour. Therefore, separating content into smaller units makes travelling through it more manageable and meaningful (Harper *et al*. 1999).

Table 1. Design Solutions.

Implementation Method		Addressed Issues
Fragmentation	Principles	Regularity, Info. Flow, Granularity, Egocentricity
	Techniques	Previewing and Probing
	Objects	Out of View, Memory
Overview	Principles	Memory, Spatial, Info. Flow, Egocentricity
	Techniques	Previewing
	Objects	Memory, Cues
Concise Egocentric Descriptions	Principles	Info. Flow, Memory, Spatial, Egocentricity
	Techniques	External Memory, Previewing, Orientation
	Objects	Memory, Cues
Active Cue/Obstacle Detection	Principles	Info. Flow, Granularity, Spatial Awareness
	Techniques	Obstacle Detection and Avoidance, Orientation
	Objects	Cues, Obstacles, Memory
Explicit Cue/Obstacle Elicitation	Principles	Info. Flow, Granularity, Spatial Awareness
	Techniques	Obstacle Detection and Avoidance, Orientation
	Objects	Cues, Obstacles, Memory

Overview Fragmenting a page means that some type of rejoining method must be employed. This means that previews of the fragmented areas will be clustered together to provide an overview and access mechanism for each fragment, aiding memory and spatial awareness and supporting information flow and egocentricity. The overview therefore encourages previewing of the clustered fragment links and consequently lowers the time taken for a visually impaired user to become familiar with the page content and structure. Fragmentation could also be useful in systems that have small viewable areas, like Personal Digital Assistants (PDAs) and mobile phones and communicators, due to their limited screen area.

Concise egocentric descriptions Providing concise descriptions about the page and site, layouts and contents aids memory and spatial awareness, and enables all Web users to immediately better decide on where they want to travel or if they have reached their goal. In the case of profoundly blind users egocentric descriptions should be used to dovetail in to the users' mental and cognitive processes. These can be in the form of pop-up boxes

Table 2. Tests of Good Mobility.

Principle	Test
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 users access memory appropriately and effortlessly at any point in a journey?
Regularity	Are travel objects deployed in a regular manner and can this regularity be recognised and exploited?
Spatial	Can spatial metaphors be reformulated into a non-spatial representation?

to give this preview information, and indeed access to 'meta' tag content provides a useful mechanism of delivering this information, if both standard 'keyword', and custom descriptive contents are used to enhance mobility.

Active cue/obstacle detection Cues and obstacles can be both active (placed specifically to increase mobility, e.g. explicit hypertext markers, explicit markers around obstructions) or passive (placed for some other reason but still used as a mobility object, e.g. hyperlinks to other pages). While passive cues are already generated in the form of headings and hyperlinks as part of the fragmentation and the hypermedia design processes, active cues and obstacle information must be explicitly placed. These should be associated with something very significant on the document that a user must be aware of if they decide not to continue reading the document to its conclusion. This can be done by linking the target item to an item link at the top or within the document. In this way the first thing that a user is aware of is the fact that an active cue or obstacle exists and it points to a piece of important information, from either a content or mobility standpoint.

Explicit cue/obstacle elicitation Obstacle and cue detection relies on knowledge of the page or site, however this knowledge is not normally explicitly present. In the context of profoundly blind travellers, an image that is just used to space other content (commonly used in HTML) is an obstacle to a user's progress regardless of whether it has an associated 'alt' attribute. Therefore a means of skipping over obstacles can be employed. This means that the obstacle is left, as it may be useful for some user groups, but those that wish to ignore it can do so without any bar to their progress. Similarly explicit cueing enhances the orientation of a user already facilitated by fragmentation.

5.1.2. Design and implementation of an experimental travel tool. The Towel mobility plug-in was designed for evaluating our assumptions and design principles and was also intended to be as simple as possible to use. The method used was fourfold, in order to accurately represent the real world mobility analogy.

1. We find out what mobility objects are present (and add those that are not) by application of our mobility framework (described in Goble *et al.* (2000) and Harper *et al*. (2001) and not examined in this paper) (see figure 2).
2. We manually mark-up any mobility objects found and add those that are not found (see figure 2 and section 5.1.3, later).
3. We implement a mobility plug-in (mobility tool containing our techniques conforming to our principles) and apply it to the marked up document (see figure 3 and section 5.1.4, later).
4. We use our resulting mobility rich hypertext within our browser to assist us in travelling around a site or document (see figures 2 and 3).

The plug-in is designed to work with XHTML pages that have been marked up with 'Towel Mobility Extensions' (TME). These are simply

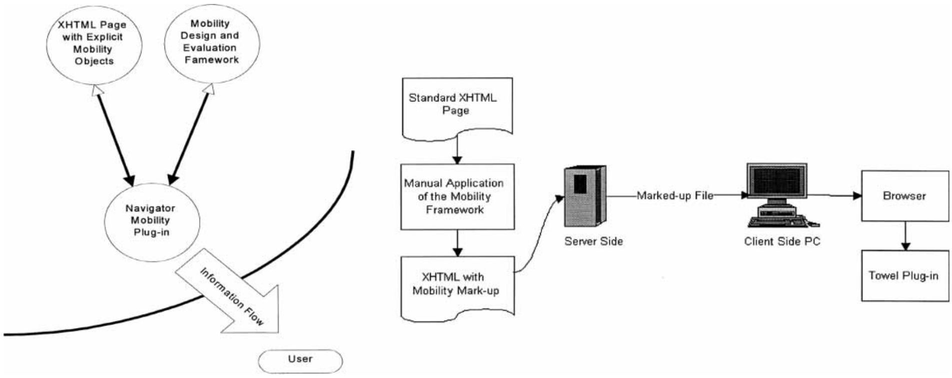

Figure 2. Relationship to the Development.

keywords written into the ID or CLASS attribute of any XHTML element. It is intended that in the future this TME mark-up will be automatically inserted, with only minor designer intervention required[13]. However for our prototype version we manually marked up a page (see a manually decorated example source file of html at http://www.man.ac.uk/towel/TMETest.html). While we used a standard plug-in (as a demonstration tool) this method is not flexible enough to provide a full mobility system; therefore the plug-in should be re-built as a 'helper' program, utility or application. Travel rendering techniques have not been part of this research and therefore all travel information has been rendered textually. However, it is envisaged that more sophisticated techniques will be required.

5.1.3. Creating TME objects. A TME object is created such that information about its placement is encoded into the naming convention for the extension thus: 'TMEi:*Location:Area:Type:Description*'. TMEi stands for **T**owel **Mo**bility **E**xtension **i**nline tag. This tag is used as a prefix so that any user agent that understands TMEs can process them without destroying the HTML, XML or XHTML. Tags can be placed in hypertext class or id attributes or as part of hypertext SPAN or DIV elements.

The Location attribute represents the area, which the TME tag represents. Specifying page as the location signifies that the information should be used when mobility around the page is required, and Site signifies that information about the whole virtual area is encoded within the TME tag.

The Area attribute specifies the part of the virtual environment under investigation and associates it with both the Mobility Framework and the Guidelines in terms of real world analogies.

[13] We envisage an author side annotation tool performing an automated 'best fit' first pass and an optional author directed 'refinement' pass.

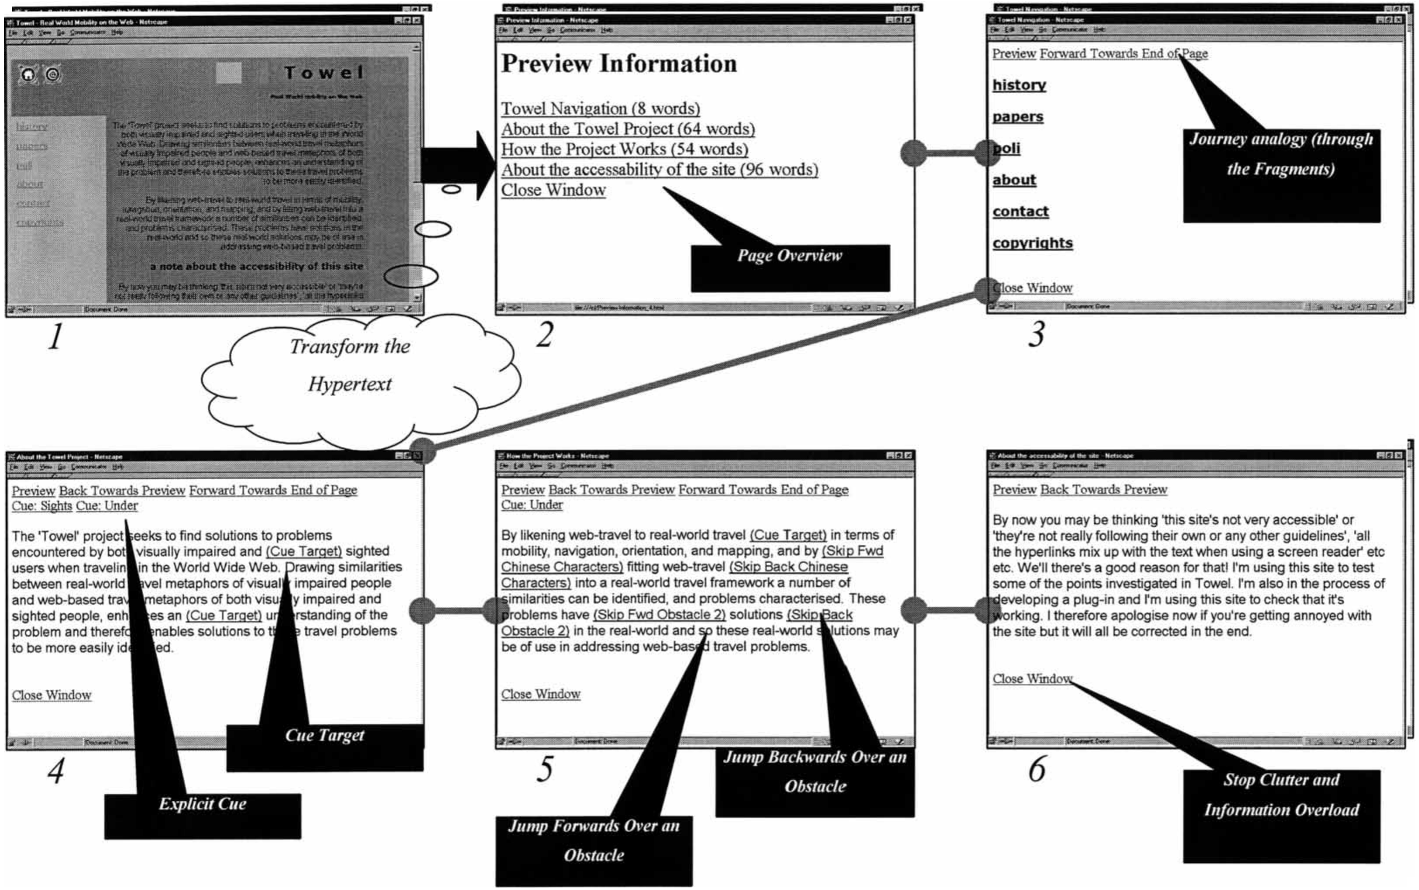

Figure 3. Plug-In Output.

- Memory (mem) represents information that is stored within the page to be referred to when required like route and layout descriptions.
- The environment (env) attribute represents information that is directly part of the actual hypermedia environment like hyperlinks and images etc.

The Type attribute specifies the real world mobility metaphor used within the tag, and represents the context in which the information should be addressed.

- A navigation (nav) context means that the tag should be used for movement from location to location and should answer the question 'Where can I go?'.
- An orientation (ori) context means the tag should be used to 'place' a traveller within a journey and should answer the question 'Where am I?'
- A cue (cue) context is intended to draw the attention of the traveller to a particularly relevant piece of information within the journey. This information may be overlooked if it is not explicitly documented and therefore a user's journey may be more difficult than it needs to be.
- An obstacle (obs) represents a mobility object that could cause problems when a user is journeying. This could be a Java Applet that is not accessible or an HTML frameset that cannot be view properly.
- An out-of-view (vws) context is used when a cluster of information cannot be read within 20 s. It is a way of creating summary information so that a traveller can quickly overview a page or site.

The description represents a simple and explicit textual description of what the tag represents and the mobility information stored within it. An example of a TME object placed within a hypertext FONT element would be:

```
<font id='TMEi:page:env:cue:Important'>in terms</font>
```

or:

```
<font class='TMEi:page:env:cue:Important'>in terms</font>
```

Here the TME object is a cue placed within the mobility environment (env) (the XHTML BODY) and its contents are related to information on that page.

5.1.4. The browser plug-in. The Towel plug-in produces a series of fragmented hypertext files, decorated with generated information derived from the placement of the TME objects. The plug-in is loaded when the browser encounters an XHTML EMBED element within the unprocessed hypertext file signalling that there are TME objects available. Once loaded the plug-in is driven by a series of keystrokes, which both activate and control its operation. These are function keys and are therefore all easily accessed by both profoundly blind and sighted users. When the plug-in is activated the TME objects are processed and the page fragmented and saved as separate local hypertext files with an automatically generated preview page. This page preview is a clustering of hyperlinks to the individual page fragments. Each hyperlink is post-fixed with a word count of the target document so that a user knows what to expect when the fragment is reached. In this way orientation to the page contents is much quicker. The mobility objects representing external memory are also processed so that a memory overview of the page and site content and layout are created. These external memory aids are implemented such that the layout and content are presented as pop up message boxes which can be viewed at any time by using a hotkey. Each generated file has hyperlinks to move through the fragment set and to move back to the preview, so that the travel analogy is maintained along with the regularity of the environment (see figure 3, pane 3). This regularity is important to profoundly blind users because it enables a degree of predictability, and this predictability can often compensate for some of the loss in visual cues.

The preview page, and hence the fragments, can now be accessed from the main browser window (see figure 3, pane 2). Therefore, egocentricity is also supported because the fragments and indeed the preview itself can be viewed in any sequence at any time and movement can occur in any direction, which means that the user chooses the way they move around a document and not the hypertext designer or the user agent developer. Additionally the links on the fragment to fragment journey are described in relation to the previous page and not as absolute references, which naturally aids this sense of egocentricity. Further, as part of the fragmentation process the obstacle and cueing information are inserted into each fragment (see figure 3, panes 4 and 5).

An obstacle component is used to define original mark-up that may pose a problem (from a mobility perspective) to certain user groups. Obstacles are not hierarchical and so a decision as to the atomic nature of obstacles or combinations of obstacles must be made. Obstacles may, however, occur as part of out-of-view and cue components although preview information is not created so as to reduce a user's cognitive overload. Cues are the opposite of obstacles but are placed in much the same way. They are used to pick out important implicit document information and make it explicit. In this way unnecessary information can be excluded by focusing specifically on useful mobility information. Cues may or may not be linked to a preview element and can be sequenced by placing a series within an out-of-view element. Figure 3, pane 4 shows two cue elements both linked back to an element in the page preview line (at the top of the page). Selecting an 'Explicit Cue' link skips the user directly to the start of the cue (the 'Cue Target'); for example selecting Cue: Under moves the user to the sentence beginning 'Understanding . . .'. Information regarding cues and obstacles is required at the time the fragment is accessed and not with the preview, as this would create unwanted information overload for the user.

Obstacles are dealt with by placing a hyperlink to a hypertext anchor 'name' attribute at each end of the obstacle so that by selecting the hyperlink the user can jump either forwards or backwards over the obstacle (see figure

3, pane 5). For example, a sentence in pane 5 has a set of non-displayable Chinese characters which when related to a visually impaired user would cause confusion. These characters can be isolated by bounding them with an object element for easy skipping. Therefore the obstacle is retained in case the user finds it useful, but can be avoided if required. This system requires user intervention in that the user selects the obstacle 'jump' hyperlink to avoid the obstacle itself. However, another solution is that the obstacle could be removed completely from the page and placed in another location, linked by a hyperlink, so that the flow of the page content is not broken and so that a user needs take no action if the obstacle is not required.

Cueing is implemented by placing a hyperlink at the top of the fragment. This link is joined from the top of the fragment to the cue destination within the fragment (see figure 3, pane 4) by an anchor 'name' tag (as for an obstacle). In this way a user immediately knows of any important mobility information contained within the fragment. This cueing information means that users travelling through a journey can rely on explicit cues and not just those that may be provided like titles and section headings. In this case the traveller knows that the hypertext designer has seen fit to explicitly mark cues that they feel will be useful.

6. Evaluation

We have performed some initial testing on original pages and those transformed by the Towel tool. These evaluations were performed along the two tracks of evaluation by the user and evaluation by reapplication of a mobility metric derived from the experience of sighted users compared to that of profoundly blind users (Harper *et al*. 2000, 2001). We have previously performed a formative evaluation (presented here in summary in section 6.1) and include details here as our summative evaluation involves a reapplication of this evaluation framework.

6.1 Formative user evaluation

Formally examining mobility on the Web is by no means a trivial task. However, a first step in looking at Web mobility is the formulation of a formative evaluation. In this way the validity of the premise of the project, and tool, could be supported and the relevance of the responses to the evaluation questionnaire can be examined. These questions are not already addressed by current visualization and modelling techniques but may be evaluated by an examination of the responses of users performing certain directed tasks:

- initial orientation within the environment;
- orientation on the page;
- orientation to a specific and disjoint page component (which may be a cue);

- reorientation to a familiar environment;
- direct searching for both information and structural page components;
- navigation;
- obstacle detection and avoidance.

The formative evaluation was based on the Internet Movie Database (http://www.imdb.com) site (IMDB) because it is an example of a classical, popular and well-organized commercial Web site. The core of the site is the search capabilities over a variety of information; over 200,000 movie & TV titles are catalogued. The site is dynamically generated using Perl as the Server Side Scripting language, linked to a relational database.

The formative evaluation questionnaire was sent to various mailing lists on the Internet and required that an individual complete a series of mobility tasks and answer questions about each task completed (Harper *et al.* 2000). The questionnaire was completed by twelve individuals, six who were completely blind or profoundly blind and six who were sighted. The equipment used varied from a standard computer set up for the sighted subjects to the use of speech and Braille translation devices for those who were blind or profoundly blind. The method of output is not taken as a factor in the analysis of the figures as these systems all perform the task of information output and are not concerned with structural and navigational information. Although the IMDB is dynamic (i.e. one that is generated from a database, when the page is requested by a user) the questions were also formulated so as to be valid for all delivery mechanisms (be it dynamic or static). These questions also cross-reference each other so that any one question is not entirely relied upon to give a single (possibly misleading) result. The questions were all based on real world surveys mainly undertaken with profoundly blind subjects, and they all relate to ascertaining journey information that can be divided into a number of real world inspired mobility groups. A full analysis of the results has been undertaken (Harper *et al.* 2000); an abridged summary of the results can be found in table 3.

6.2 Summative user evaluation

We tested the Towel tool by analysing the responses of a sighted and profoundly blind user when moving around the processed 'Towel' test page[14]. The only requirements were that they had experience in using the World Wide Web, that they each surfed different kinds of websites (not just news, say) and their first language was English or they were fluent in reading English (although their fluency was not formally tested and we relied on their honesty). All testers followed a set script and the whole process was initially tested on four respondents as a way of eliciting comments on the procedure.

[14] Soliciting the initial formative evaluation proved easy as we did not require the user to be present at our testing location. However, finding subjects for the summative evaluation proved more difficult; we therefore had to rely on locally resident users, or perform no summative evaluation at all.

Table 3. Formative Evaluation Summary.

Category	Sighted		Visually Impaired	
	Mean Time (sec.)	User Comments	Mean Time (sec.)	User Comments
Initial Orientation	N/A	Cued to icons/status bars.	N/A	Cued to random hardware noise.
Page Orientation	10	All agreed on layout. Users correctly labelled each area.	120	Agreement on layout was not reached. Users did not label each area correctly.
Component Orientation	N/A	All users easily identified the search box component.	N/A	All users easily identified the search box component.
Re-Orientation	15	The page layout is familiar. All users responded that the page layout was the same.	120	The page layout was not familiar. Only 33% of users recognised that the page layout was the same.
Directed Search	10	The information was easily found.	100	It was difficult to decide if the information was present.
Navigation	10	All users identified which menu items could be traversed and which could not.	95	All users identified that some menu items could be traversed and and that some could not.
Obstacle Detection	5	All respondents correctly identified the obstacles.	95	No respondents correctly identified the obstacles.

In this way the questioning system was tested before the real evaluation (these test scripts were not taken into account in the overall survey). Each user was asked a series of questions as per those in the formative evaluation (Harper *et al.* 2000) and their responses were noted.

It was found that the fragmentation of the page enhanced a profoundly blind users movement around, and cognition of, our test pages. Both users said that they gained a greater understanding of both the page and site structure, and this was particularly apparent for profoundly blind users.

The preview information enabled the user to make a quick decision as to whether to investigate the page further or to move on to a different page. In fact the profoundly blind user found this far more useful than the sighted user, as these preview descriptions made visually implicit information explicit.

When investigating the page the movement cues along the top of the document were found to be particularly helpful in guiding a user through the page fragments. They were useful in enabling orientation to both the page structure and the fragments themselves. This information was mainly helpful to the profoundly blind user as the sighted individual could easily orient themselves by the visual information present within the document.

The notion of obstacles is promising, especially for large obstacles that take up a large amount of screen real estate. However, the users did suggest that the obstacle be removed from the page and a suitably descriptive hyperlink inserted, so that user intervention is not required when travelling through the fragment. The 'skip' metaphor is helpful but was a little laborious when tabbing from link to link; however obstacle removal would solve this problem.

Cues also seem to be useful in drawing a user's attention to specific significant items within the page or journey. This was the case for both users, however they also felt that the inclusion of an annotation mechanism enabling them to define their own cues would be an advantage.

Finally, the user presentation of page and site, content and layout was not found to be as useful as we had initially expected. This seemed to be because either too much or too little information was present and this information did not change with the journey.

6.3 Evaluation by reapplication of the formative evaluation metric

In light of the small respondent group of the summative evaluation, we performed a further evaluation exercise. This consisted of a series of automated tests created by reapplying the formative evaluation metric (Globe *et al.* 2000, Harper *et al.* 2001). We selected our own Towel Web site after the page had been updated with Towel Mobility Extensions (TME). We could then evaluate how effective the mobility tool (our plug-in) implemented the mobility techniques on the structural objects now coded into the page. It also enabled us to see if the mobility metrics were suitably applied to the site by the tool, because if not, the overall scores for that system would decrease.

After placement of the TME there was little change in the usability ratings for sighted users. In fact on examining the usability summary and the usability statistics it could be seen that there was very little room for improvement, as the initial score was 91%. For the profoundly blind user however the story is quite different because with a usability rating of only 51%, before placement of the TME, the scope for change was far higher, and increased to 86%.

Major improvements can be seen within environment mobility objects. These improvements are mainly concerned with making inventory items more visible and explicit if used as mobility objects. There are also increases with regard to memory objects[15] but this is not as pronounced as with objects within the environment. The overall rating for the page has therefore been increased substantially and this increase has brought the mobility information available to a profoundly blind user into line with that useable by a sighted user[16].

7. Making a better mobility tool

Our experimental plug-in cannot provide the full range of functionality to adequately support the complexity found in Web travel. However, it has stimulated ideas about new properties and additions that could be made to the system.

Journeys should be monitored so the positional information is known: in this way cues and obstacles should be decided-on when they are near to the traveller and a user can state that obstacles should no longer be regarded as obstacles for them personally.

The presence of cues and obstacles should be included with the preview information. These should be just either a ‘C’ or ‘O’ merely to indicate their presence and so as not to increase information overload. External Memory should keep journey metrics, dynamic maps, old journeys and allow the exchange of information between travellers.

External Memory should also allow the marking of interesting points for later investigation and should create a ‘road’ by placing junctions of hypertext links that you may want to follow at present or later. Either way this map should be kept for the session and permanently if the user decides.

Further work will also be undertaken on establishing that the solutions contained within the plug-in are both useful and appropriate for profoundly blind travellers, even though our initial studies indicate that they are useful and timely.

8. Conclusions and further work

Applying knowledge about real world mobility to Web based mobility problems can enhance the travel experience for profoundly blind users like

[15] Objects that are used to provide cueing information to a user.

[16] A more complete discussion and statistical analysis as was created in the process of testing this site (and others) (Harper 2001).

David. The use of fragmentation to assist preview and appropriate knowledge feedback may increase the mobility of users within many virtual journeys and therefore solve many of the mobility problems encountered frequently in Web based travel.

The Towel system adds two more levels of complexity to standard Web systems, by manually pre-marking an XHTML page with mobility information, and then by processing that information in an appropriate way using a plug-in when the XHTML page arrives at the client's browser.

Extensions to the system propose that the hypertext page should be automatically marked-up with mobility information and that this page should then be displayed using an integrated browser/mobility application. The extended system should also process the XML and mobility information can also be inserted using XML so that it can be portable between systems. Eventually, it may be possible to process hypertext information on the fly based on a user pre-defining a set of mobility options defining XML elements or combinations of elements that represent mobility objects to them personally.

This work has been useful in leveraging real world mobility solutions to build a prototype mobility plug-in for the World Wide Web. By addressing travel issues to enhance the mobility of all users, especially those who are profoundly blind, users should be able to better surf the Web. Finally, when David 'Surfs the Web' it will be 'Rapid and Free'.

References

K. Akins (Ed.), *Perception*, Oxford: Oxford University Press, 1996.

C. Asakawa and H. Takagi, "Annotation-based transcoding for nonvisual web access", in *Proceedings of the Fourth International ACM Conference on Assistive Technologies*, 2000, pp. 172–179.

R. Barrett, P.P. Maglio and D.C. Kellem, "How to personalize the web", in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1997, pp. 75–82.

M. Bates, "The design of browsing and berrypicking techniques for the online search interface", *Online Rev.*, 13, pp. 407–424, 1989.

M. Brambring, "Mobility and orientation processes of the blind", in *Electronic Spatial Sensing for the Blind*, D.H. Warren and E.R. Strelow, Eds., Dordrecht, Lancaster: Nijhoff, 1984, pp. 493–508.

O. Buyukkokten, H.G. Molina, A. Paepcke and T. Winograd, "Power browser: efficient Web browsing for PDAs", in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2000, pp. 430–437.

E. Carmel, S. Crawford and H. Chen, "Browsing in hypertext: a cognitive study", *IEEE Trans. Systems, Man Cybernetics*, 22, pp. 865–883, 1992.

S.-J. Chang and R.E. Rice, "Browsing: a multidimensional framework", *Annu. Rev. Inf. Sci. Technol.*, 28, pp. 231–276, 1993.

C. Chen and M. Czerwinski, "Spatial ability and visual navigation: an empirical study", *New Rev. Hypermedia Multimedia*, 3, pp. 67–89, 1997.

C. Chen, "Structuring and visualising the www by generalised similarity analysis", in *Proceedings of the 8th ACM Conference on Hypertext and Hypermedia*, 1997, pp. 177–186.

A. Chieko and C. Lewis. "Home page reader: IBM's talking web browser", in *Closing the Gap Conference Proceedings*, 1998.

Codix, *"Textualize"* (Manual—supported on Website circa 2005). Codix.net. Available online at: http://codix.net/solutions/products/textualise/index.html (2005).

J.F. Cove and B.C. Walsh., "Online text retrieval via browsing", *Inf. Process. Manag.*, 24, pp. 31–37, 1998.

R.P. Darken and J.L. Sibert, "Wayfinding strategies and behaviors in large virtual worlds", in *Proceedings of the Conference on Human Factors in Computing Systems: Common Ground*, 1996, pp. 142–149.
A. Dillon and M. Vaughan, "It's the journey and the destination: shape and the emergent property of genre in evaluating digital documents", *New Review, Hypermedia and Multimedia*, 3, pp. 91–106, 1997.
R.M. Downs and D. Stea (Eds), *Image and Environment; Cognitive Mapping and Spatial Behavior*, Chicago: Aldine Pub. Co (Walker De Gruyter Inc.), 1973, pp. 391–423.
R.M. Downs and D. Stea, *Maps in Minds: Reflections on Cognitive Mapping*, New York: Harper Row, 1977, pp. 264–272.
M.J. Farah and G. Ratcliff (Eds), *The Neuropsychology of High-Level Vision: Collected Tutorial Essays*, Hilldale, NJ: Lawrence Erlbaum Associates, 1994.
R. Furuta, "Hypertext paths and the www: experiences with Walden's paths", in *Proceedings of the 8th ACM Conference on Hypertext and Hypermedia*, 1997.
F. Garzotto and M. Matera, "A systematic method for hypermedia usability inspection", *New Rev. Multimedia Hypermedia*, 3, pp. 39–65, 1997.
M.S. Gazzaniga and E. Bizzi (Eds), *The Cognitive Neurosciences*, Cambridge, Mass: MIT Press, 1995.
D. Gentner and A.L. Stevens (Eds), *Mental Models*, Hilldale, NJ: Lawrence Erlbaum Associates, 1983.
J.J. Gibson, *The Ecological Approach to Visual Perception*, New York: Harper Row, 1979, pp. 313–318.
C. Goble, S. Harper and R. Stevens, "The travails of visually impaired Web travellers", in *Proceedings of the Eleventh ACM on Hypertext and Hypermedia*, 2000, pp. 1–10.
S. Harper, "Web mobility for visually impaired surfers (the Towel Project)", PhD thesis, University of Manchester, UK (http://towel.man.ac.uk).
S. Harper, C. Goble and R. Stevens, "A pilot study to examine the mobility problems of visually impaired users travelling the Web", *ACM SIGCAPH Comput. Physically Handicapped*, (68), pp. 10–19, 2000 (http://towel.man.ac.uk).
S. Harper, R. Stevens and C. Goble, "Towel: real world mobility on the Web", in *Computer-Aided Design of User Interfaces II*, J. Vanderdonckt and A. Puerta, Eds, Dordrecht: Kluwer Academic, 1999, pp. 305–314 (http://towel.man.ac.uk).
S. Harper, R. Stevens and C. Goble, "Web mobility guidelines for visually impaired surfers", *J. Res. Practice Inf. Technol.*, 33, pp. 30–41, 2001 (special issue on HCI) (http://towel.man.ac.uk).
M. Hori, G. Kondoh, K. Ono, S. Hirose and S. Singhal, "Annotation-based web content transcoding", in *Proceedings of 9th International World Wide Web Conference*, 2000a.
M. Hori, G. Kondoh, K. Ono, S. Hirose and S. Singhal, "Annotation-based web content transcoding", in *Proceedings of the 9th International World Wide Web Conference on Computer Networks*, 2000b, pp. 197–211.
I.P. Howard, *Human Visual Orientation*, New York: J. Wiley, 1982, pp. 588–669.
A.W. Huang and N. Sundaresan, "A semantic transcoding system to adapt web services for users with disabilities", in *Assets '00: Proceedings of the Fourth International ACM Conference on Assistive Technologies*, 2000, pp. 156–163.
S. Jul and G.W. Furnas, "Navigation in electronic worlds: a CHI 97 workshop", *ACM SIGCHI Bull.*, 29, pp. 44–49, 1997 (http://doi.acm.org/10.1145/270950.270979).
B.B. Lloyd, *Perception and Cognition: a Cross-cultural Perspective*, Harmondsworth, UK: Penguin, 1972, pp. 155–171.
D. Lowe and W. Hall, *Hypermedia and the Web: An Engineering Approach*, New York: Wiley, 1999.
M. McTear, *Understanding Cognitive Science*, New York: Halsted Press and Chichester, UK: Harwood Press, 1988, pp. 242–258.
P. Maglio and R. Barrett, "Intermediaries personalize information streams", *Commun. ACM*, 43, pp. 96–101, 2000.
P. Maglio and T. Matlock, "Metaphors we surf the web by", in *Workshop on Personalized and Social Navigation in Information Space*, 1998 (citeseer.nj.nec.com/maglio98metaphors.html).
G. Marchionini, *Information Seeking in Electronic Environments, Cambridge Series on Human–Computer Interaction*, Cambridge: Cambridge University Press, 1995.
J. Morkes and J. Nielsen. "Concise, scannable, and objective: how to write for the web". Available on line at: http://www.useit.com/papers/webwriting/writing.html (1997).
W. Myers, "BETSIE:BBC Education Text to Speech Internet Enhancer", British Broadcasting Corporation (BBC) Education. Available on line at: http://www.bbc.co.uk/education/betsie/.

H. Obendorf and H. Weinreich, "Comparing link marker visualization techniques: changes in reading behavior", in *Proceedings of the Twelfth International Conference on World Wide Web* , 2003, pp. 736–745.
RNIB (various authors), "A Short Guide to Blindness". Available online at: http://www.rnib.org.uk (February 1996).
L. Seeman, "The semantic web, web accessibility, and device independence", in *W4A: Proceedings of the International Cross-disciplinary Workshop on Web Accessibility*, 2004, pp. 67–73.
J.M. Spool, "Web site usability: a designers guide, user interface engineering". Available online at: http://world.std.com/uieweb/index.html (1997).
H. Takagi and C. Asakawa, "Transcoding proxy for nonvisual web access", in *Proceedings of the Fourth International ACM Conference on Assistive Technologies*, 2000, pp. 164–171.
W3c, "Web content accessibility guidelines". Available online at: http://www.w3.org/TR/WCAG20/ (February 2005).
Various sources, "WebCleaner". Available online at: http://webcleaner.sourceforge.net (circa 2004).
Y. Yesilada, S. Harper, C. Goble and R. Stevens, "Ontology based semantic annotation for enhancing mobility support for visually impaired web users", in *Proceedings of the K-CAP 2003 Workshop on Knowledge Markup and Semantic Annotation (Semannot'2003)* , 2003.