

A User Evaluation of the SADie Transcoder

Darren Lunn, Sean Bechhofer and Simon Harper
Information Management Group
School of Computer Science, University of Manchester
Kilburn Building, Oxford Road, Manchester, M13 9PL, UK

darren.lunn@cs.manchester.ac.uk
sean.bechhofer | simon.harper}@manchester.ac.uk

ABSTRACT

The World Wide Web (Web) is a visually complex, dynamic, multimedia system that can be inaccessible to people with visual impairments. SADie addresses this problem by using Semantic Web technologies to explicate implicit visual structures through a combination of an upper and lower ontology. This is then used to apply transcoding to a range of Websites. This paper describes a user evaluation that was performed using the SADie system. Four users were presented with a series of Web pages, some having been adapted using SADie's transcoding functionality and others retaining in their original state. The results of the evaluation showed that providing answers to a fact based question could be achieved more quickly when the information on the page was exposed via SADie's transcoding. The data obtained during the experiment was analysed and shown to be statistically significant. This suggests that the transcoding techniques offered by SADie can assist visually impaired users accessing content on the Web.

Categories and Subject Descriptors

H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems – Evaluation / methodology; H.5.3 [Information Interfaces and Presentation]: Group and Organization Interfaces – Web-based interaction; H.5.4 [Information Interfaces and Presentation]: Hypertext / Hypermedia – User Issues

General Terms

Measurement, Experimentation, Human Factors

Keywords

Accessibility, Web, Semantic Transcoding, SADie, Visually Impaired Users

1. INTRODUCTION

People with disabilities, in particular visual impairments, are hindered in their access to information on the Web be-

cause it is not designed with their needs in mind. Visually impaired users can make use of tools to access the Web that read aloud the page. These tools, known as screen readers, access the underlying structure of the Hypertext Markup Language (HTML) to create a sequential, audio rendering of the document [11]. Most Web designers, however, are mainly concerned with how content is presented on screen, rather than its structure and meaning. Consequently, implicit information that is available through the visual rendering of the page is lost to the screen reader, and therefore the user. As an example, consider a navigation menu common on many Websites. Typically a navigation menu is contained within its own distinct chunk, separated from the rest of the content, and is located towards the side of the screen. A sighted user is aware of which element on screen is the navigation menu due to the way it is rendered. There is nothing explicitly stating that the element on the Web page is a menu. The knowledge that it is a menu is implicit from the visual presentation. This implicit information is only available to those who can see it, as opposed to those people who use a screen reader to read aloud the page content.

Structural-Semantics for Accessibility and Device Independence (SADie) makes use of Semantic Web technologies, in particular ontologies, so that information available implicitly through its on screen rendering, such as a menu, can be made explicitly available. This is achieved by annotating the rendering information defined within the Cascading Style Sheet (CSS) of the Website. With the structural semantics exposed, transcoding is then applied to the Web page content so that information held within the page is more accessible to visually impaired users.

In previous works, we have demonstrated that annotating the CSS of a Website through an ontology can provide enough information for the page to be transcoded [2]. While these previous demonstrations have validated our approach on a technical level, they have not provided evidence that the adaptations applied to pages are of use to visually impaired users. In this paper, we address previous limitations of our evaluations and present a user-centred qualitative analysis of the SADie approach.

The results of the evaluation show that SADie is effective in decreasing the time required for users to find facts held within the content of the page. This provides evidence that the SADie functionality of removing unnecessary elements and promoting important and useful content is of use to visually impaired users and should appear in future transcodes that our tool provides.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ASSETS'08, October 13–15, 2008, Halifax, Nova Scotia, Canada.
Copyright 2008 ACM 978-1-59593-976-0/08/10 ...\$5.00.

2. RELATED WORK

Efforts exist that attempt to make Web pages accessible to people with visual impairments. A major source of this is Web Content Accessibility Guidelines (WCAG) devised by the Web Accessibility Initiative group at the World Wide Web Consortium (W3C) [4]. These provide Website developers with guidelines that range from advice on the design process, such as separating the structure of the document from its presentation, right through to the testing process, such as testing the Web page on as many system configurations as possible. While the guidelines can go some way to provide accessible content, they are not without problems. Petrie et al. evaluated 100 Websites with 51 users who had a disability [18]. Experts observed the difficulties that the users encountered and of the 585 problems that were observed, 45% were due to problems that were in no violation of any checkpoint found within the WCAG guidelines. This result, in addition to other critiques of accessibility guidelines, such as work by Kelly et al. [13], have led to investigations into alternative techniques, centred around the use of transcoding, to assist users with disabilities navigate the Web.

Transcoding is a way of transforming Web content so that it can be accessed on a diverse range of devices [10]. In adapting Web content, transcoding systems use a variety of architectures, a range of methods and have a diverse target user group. However, they all exude similar characteristics that allow them to be placed into one of three categories; Heuristic, Template and Semantic.

With heuristic transcoding, tools analyse a page and adapt it based on a set of predefined rules. Chen et al. made use of heuristics where it was asserted that a Web page was composed of objects, such as presentation objects, interaction objects and hyperlink objects [3]. By discovering what the functionality of the objects was, suitable transcoding rules could be applied. For example, small images were considered as presentation objects, and therefore removed. Gupta et al. exploited the Document Object Model (DOM) to apply transcoding to the Web document [5]. The Web Document was parsed into a DOM tree and the nodes traversed in order to identify content. For example, table cell nodes that had a large link to paragraph ratio were considered to be “link lists” and removed from the page.

Template based Transcoding, as used by Parmanto et al. , is based on the premise that humans interact with familiar environments better than non-familiar ones [17]. This tendency means that Websites that provide similar functionality all use similar templates. News Websites all have a similar look and feel and e-commerce sites all use similar styles. By discovering the genre of Website that the user is accessing, the transcoding system can match a template to the site in order to exploit the underlying structure of the Web document so that the content can be rearranged. Jang et al. used templates to create pages for different devices [12]. A template that suited the device capabilities was created and used as the basic model when adding and reorganising content from different Websites.

Semantic Transcoding is the adaptation of a Web page by using the semantics of the structure or content to make better adaptation decisions [8]. Whereas heuristic and template transcoding makes educated guesses as to how the page should be adapted, semantic transcoding uses the metadata that describes the structure of the page to make adapta-

tion decisions. To help capture the metadata, many semantic transcoding systems annotate Web pages to capture the structure of the page and the knowledge held within it [16].

Takagi et al. created annotations in three ways: *volunteer-specified*, which identified visual fragments and assigned them a level of importance; *automatic*, which added alternative text to images and warned the user of JavaScript; and *user-specified*, which allowed the user to specify the starting point of the main content [20]. With the annotations in place, transcoding could then be applied to the page. For example, users could simplify the page by reorganising the content into a more suitable format. However, for the transcoding to work, every Web page visited required annotations in order for the page to be adapted by the transcoding proxy.

In addition to identifying content importance, annotations can also identify roles elements play within the page [15]. Towel compared the notion of travelling in the real world to travelling within a virtual environment [7]. To aid travel, the Extensible Hypertext Markup Language (XHTML) of the Web page was annotated so that the roles of the elements could be identified. For example, *cue contexts* were items that drew the user’s attention to relevant pieces of information on the page. *Obs contexts* were obstacles that could cause problems to the traveller. By explicitly stating the purpose of the elements within the page, Towel could help navigate the user through the components and structure of the Web page, including links to other Web pages and resources.

Heuristic, template and semantic transcoding all have benefits and drawbacks. A major advantage of Heuristic and Template based transcoding Systems is that they can be applied to a wide variety of Websites without any additional information with regards to the page content [1]. While manipulations of the structure can aid navigation and display, they have two problems. The first is that the rules or templates have to be general enough so as to apply to every possible Web page. Therefore they do not capture and exploit the semantic information of the page, which is usually only available implicitly through the visual presentation. Secondly, a specific user group is targeted, such as mobile Web users or visually impaired users. This means that for each user group or target device, a different set of heuristics and templates need to be discovered and applied to the page [9].

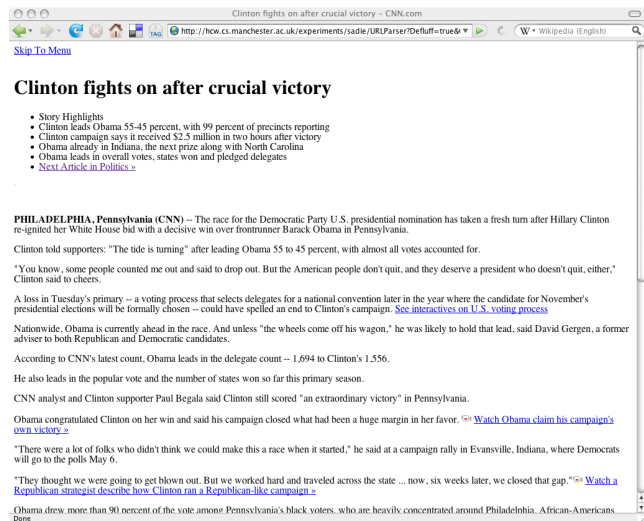
In contrast, semantic transcoding attempts to capture the meaning of the information on the page in order for this meaning to be exploited to suit a variety of different user groups [1]. This involves deeper levels of document understanding and therefore human intervention into machine understanding of documents is required [16]. Annotations can provide the additional information that machines use to perform the semantic transcoding. This tends to produce higher quality transcribed documents because of the additional level of understanding. However, the cost of this is that annotations have to be created by hand, which can be extremely tedious and time consuming [21].

3. SADIE

The SADIE platform provides a solution that combines the benefits of heuristic and semantic transcoding to offer an accurate yet highly scalable transcoding solution. The principle idea behind our research is that the rendering of a Web page element is closely associated with its role. For example,



(a) A News Story From The CNN Website



(b) A News Story From The CNN Website After Being Transcoded By SADIE

Figure 1: Comparison of a Standard CNN News Story and The SADIE Transcoded Version. Web Page Taken From <http://edition.cnn.com/2008/POLITICS/04/23/us.primary.intl/index.html> on 23rd April 2008

sighted users know that a list of links is a menu due to the way it is rendered on screen. This rendering information is defined within the CSS and associated with the HTML via tag attributes such as class or id. A Website typically only has a single CSS file containing all the style definitions, therefore, rather than annotate every page, the CSS classes themselves are annotated. This reduces the annotation overhead as any annotating of the document only occurs in one location. However, as the Website content is associated with the CSS, this too is indirectly annotated. The CSS annotations provide us the accuracy of semantic transcoding as we are explicitly stating the role of each element. Additionally we also gain the high scalability of heuristic transcoding. Defining the roles of CSS elements allows every page within the Website to be transcoded due to the fact that CSS tends to contain site-wide style definitions that all pages use.

The transcoding is driven by an ontology that provides a defined set of terms for classifying the CSS elements within a Website. The ontology consists of two parts. The first is an upper ontology containing high level abstract concepts representing the potential roles of Web page elements. These roles include concepts such as “Removable” which is used to annotate elements that are not important to the Website and “Priority” which identifies elements that contain important information on the page. The second part of the ontology is a Website specific extension to the upper ontology that contains the elements found within the CSS. These elements are annotated with the roles that they play from the terms found in the upper ontology. The benefit of this approach is that the upper ontology acts as an interface between the SADIE transcoding engine and the page we wish to adapt. SADIE requests a list of elements that need to be transcoded and the ontology returns all the CSS elements that satisfy the query. SADIE can query any number of Website ontologies because each site specific ontology uses the same upper ontology, providing a consistent interface to the heterogeneity of the names that can be found within CSS files. The results of the query are then used to apply the transcoding.

The aim of SADIE is to improve access to Web content for visually impaired users. This is achieved by transcoding the page into a format more suited to the sequential audio stream generated by a screen reader. SADIE currently provides the user with three varieties of functionality. These are: **Defluff**: which involves removing elements that provide little or no information to the page; **Reorder**: which involves reordering the page so that the areas that provide the most important information appear near the top of the page; **Menu**: which moves the menu of the Website to the bottom of the page where it can be easily found yet allows the main content to be immediately accessed.

SADIE matches the elements of the Web page to the functionality that the user has selected in order to apply the transcoding. For example, if a user has selected “Defluff” then SADIE will query the ontology for a list of all the CSS classes that have been classified as “Removable”. When the query has been answered, SADIE traverses the Web page’s DOM and removes any element that occurs within the list returned by the ontology query. When the DOM has been traversed, the modified page is returned to the user.

Figure 1 illustrates how SADIE can be used to transcode a Web page. Figure 1a shows the main news story from the CNN Website on the 23rd April 2008. The story, which can be considered to be the main content, is surrounded by banners, search boxes and advertisements. These hinder a visually impaired user as they attempt to access the main content. Figure 1b shows the same page after it has been transcoded by SADIE. This page has been “Defluffed”, “Reordered” and “Menued”. Note how all the clutter has been removed from the original page and the story has been promoted to the top of the page where it can be immediately accessed by a screen reader. The menu has been suppressed to the bottom of the page but is accessible via a link situated near the top. This allows users to easily find the menu when they wish to navigate to other news stories. For further discussion of the SADIE method and architecture, the reader is directed to [2, 6].

Table 1: Participants In The Evaluation

<i>Participant</i>	<i>Gender</i>	<i>Age Range</i>	<i>Disability</i>	<i>Assistive Technology</i>	<i>Internet Usage</i>
P1	M	36 – 45	Profoundly Blind	JAWS	Every Day
P2	F	26 – 35	Partially Blind	JAWS & ZoomText	Every Day
P3	M	36 – 45	Profoundly Blind	JAWS	Every Day
P4	M	36 – 45	Profoundly Blind	JAWS & Keynote Gold	Every Day

4. EVALUATION

An evaluation was performed to compare the time it took to find information on an original Web page compared to the time it took users to find information on pages adapted using the SADie transcoding. The evaluation was conducted with users who regularly accessed the Web using screen reader technology. The evaluation was based upon the hypothesis that the time taken to find information within a Web page is reduced when accessing the content through SADie transcoding. For a full description of the pages used in the study, the tasks that users were asked to perform and the data obtained, the reader is directed to [14].

4.1 Participants

Four individuals participated in the evaluation, all of whom had a visual impairment and required assistive technology to access their computer. The evaluation involved three men and one woman, aged between 26 and 45. All the participants used a computer daily and were experienced at using JAWS. One participant, P2, used JAWS in conjunction with ZoomText, a screen magnification tool, and a second used JAWS in conjunction with Keynote Gold, a braille device. For the purposes of this evaluation, all participants used JAWS. The participant’s details are summarised in Table 1.

All the participants gave their informed consent to take part in the evaluation and were paid £10 for their time and effort. Potential participants were screened before the evaluation took place to ensure that they were suitable for the SADie evaluation. A number of participants applied who, whilst visually impaired, solely used screen magnification software to access their computer. While some of the users were familiar with both screen readers and screen magnifiers, we limited the experiment to screen reader use only. This is the target audience of the SADie transcoding features and, as screen magnification and screen reading are different interaction paradigms, we felt that combining the two groups in the evaluation would not yield fair comparisons.

4.2 Materials

Participants were asked to perform the evaluation on a series of Web pages. The number of tasks used in the experiment was 20 as we felt this struck a balance between gathering enough data to perform statistical testing, yet not providing too many tasks that would frustrate the user and require too much time. While 20 pages may not seem like a representative sample of pages found on the Web, it would be impractical to build a corpus of Web pages which represents a greater number than this. According to Netcraft¹ there are 150,000,000 Websites. Testing 10,000 Websites, rather than 20, means that only 0.007% more of the Web is tested, and this is not significant. Indeed, we would need to

¹http://news.netcraft.com/archives/Web_server_survey.html

test 1,500,000 Websites to achieve just a 1% testing. Given the impossibility of testing any significant proportion of the Web, the key thing is good sampling. A sample must reflect the population: major categories should be represented and this can be narrowed down further by choosing sites which are of a typical representation (sites with large numbers of visitors, and longevity). Our 20 pages were chosen based on a list of the top 100 visited Websites². In this way we ensured that we worked with a set of pages which represented a reasonable percentage of user hits and usage.

All the Websites chosen were predominantly text based and provided content that was updated frequently, for example news pages. This ensured that even if participants were familiar with the sites chosen, as the content often changed they would not be able to pre-empt the answer to the task. The tasks were based upon the entry page to the site where appropriate. For some Websites, such as search engines, this was not appropriate as the front pages contained little content and rarely changed, allowing participants to be able to provide an answer to the question immediately. Instead, the search results page was used. Tasks were designed such that similar pages had similar questions. For example, all blog pages had the task “What is the first blog entry?” and search results pages had the task “What is the first search result?” Tasks were also designed to try and lead users to the main content of the page, so as to try and emulate Web usage. One can assume that users would be interested in reaching search results or reading blog entries when accessing pages.

All the transcoded pages had the identical adaptations of defluff and reorder applied to them. Users were not in control of the transcoding. All the Web pages used in the evaluation were saved as a local copy to reduce the risk of network failure and latency effecting the evaluation.

To avoid an element of bias being introduced into the results, the evaluation was conducted as a double-blind experiment [19]. In addition, the tasks were also performed by a sighted user before the evaluation took place. This was to ensure that the wording of the tasks was clear and unambiguous. The aim of the evaluation was to establish if SADie improved the time it took a user to answer fact based questions about the page. If the wording was not clear, then the evaluation could have become a test of interpretation rather than a test of SADie’s ability to expose information. Performing the experiment on a sighted user ensured that this was not the case.

4.3 Method

To ensure that the participants were comfortable with the evaluation, and to ensure that there was a high level of internal validity, participants performed the tasks using their own equipment in familiar surroundings. The nature of the evaluation was to study the effects of SADie and it was felt

²Taken from <http://www.alexa.com/>




Figure 2: Comparison of the average time to find information on a page using SADie and the average time to find information on a page using the original version.

that accustoming the participant to new surroundings and unfamiliar equipment was an unnecessary burden.

The tasks for each participant were randomised so that no participant completed the task in the same order. In addition, participants were never asked to perform tasks on both versions of a Web page. For example, if a user was asked to find information on the original version of a news Website, then they were not asked to perform the task on the transcoded version. This ensured that users were actively seeking information within the page and not recalling the answer from memory, causing a positive effect on the results on the second version of the task. Due to the small number of users not every combination of Webpages was attempted.

A task page was created in order for participants to perform a task. This was a simple HTML document that had a list of links pointing to the twenty tasks that had been created. The evaluator gave the participant instructions for the next task. When the participant was ready, they clicked on the link that took them to the page required to perform that task. The timer started as soon as the participant had clicked the link. The time stopped when the participant had completed the task. The participant was then allowed to return back to the main page in their own time. When the participant indicated that they were ready, the evaluator gave instructions for the next task and the process was repeated.

To validate the results of the evaluation, the entire session was recorded using an audio tape recorder installed at the back of the room. The tape recorder was also used to record participant feedback and comments after the evaluation had been completed.

5. RESULTS

Not all participants completed the tasks, therefore we only present comparisons of the eight pages where tasks were per-

formed on both the original and transcoded version of the pages by at least one user. Where a task was attempted but not completed, we use the time at which the user gave up and moved onto the next task. For the original version of the page, there were four instances of the tasks not being completed and for the SADie versions, there were no instances where the user failed to complete a task. Figure 2 shows a comparison between the average time taken to complete the task on both the original and transcoded versions of the page. It should be noted that the same user never performed a task on both versions of a page; they were either asked to complete the task on the original version or the transcoded version. With the exception of Page 2 and Page 8, SADie decreased the average time it took users to find information within the Web pages.

The increased time for the transcode version of Page 2 could be attributed to the user moving beyond the information as they did not expect the story to be immediately available. When the task started, the evaluator asked the question “*What is the title of the main news story?*” and the user went to the page as requested. Within three seconds, the news story title was being read aloud, however P2 was following their usual interaction patterns to navigate the Web page. They were not aware that they were accessing a modified version of the page and so quickly navigated from link to link, to reach the content. However, when they came across a link entitled “*Around the World Now – Africa*”, this acted as a cue that they had moved too far down the page. They then had to navigate backwards to the start of the page where the article title was available.

Page 8 was a blog and the participant who accessed the transcoded version was not entirely sure what a blog was. The evaluator asked the question “*What is the title of the first blog entry?*” followed by the user starting the task and moving around the page trying to gain an overview and un-

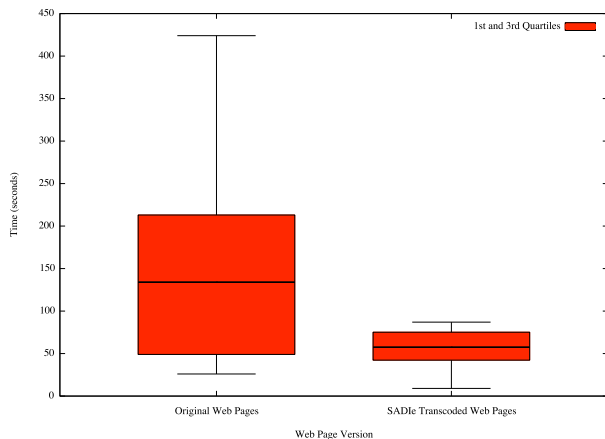


Figure 3: Comparison of the average time to find information on a page using SADiE and the average time to find information on a page using the original version.

derstanding of the page. As the user was unsure of what a blog was, there was confusion between the title of the blog itself and the title for the most recent entry, which contributed to the longer task completion time.

On average, participants took longer to complete tasks with original Web pages ($mean = 163.25$, $standard\ error = 49.87$) than they did to complete tasks using SADiE transcoded pages ($mean = 54.19$, $standard\ error = 10.48$). Analysing the results using the t-test gave a value of $t(7.6) = -2.14$, with $p = 0.033$. As $p < 0.05$, the result is significant which supports our hypothesis that the time taken to find information within a Web page is reduced when accessing the content through SADiE transcoding. A box plot illustrating the average task completion times for both transcoded and original versions of Web pages can be seen in Figure 3.

6. DISCUSSION

There is now evidence to suggest that SADiE is of benefit to visually impaired users and can help them access content on Web pages more easily. The results show a statistically significant difference between the time taken to complete fact-based tasks with SADiE than without. While this is encouraging, there were still instances where SADiE did not seem to benefit users, such as Page 2 and Page 8, however qualitative feedback from participants and observations made during the evaluation sessions provided insight into areas where we can enhance SADiE and further improve the transcoding that we provide.

6.1 The Importance of Structure

During the evaluation, a number of the participants discussed how the structure of pages can help them find content. P1, when discussing search engine results after the session had finished, stated that *“if you told me these were the results, such as a heading stating ‘The results are here’ then that would help me.”* This claim was supported by P3 who commented that *“headings should proclaim what is coming”*. This is consistent with how P3 interacted with and navigated around the Web page. P3 relied heavily on headings, using functionality provided in JAWS to skip between heading sections of the page. P3 found it very frustrating

to interact with pages that had poor structure and which they were not familiar with. Having a well laid out page, using appropriate mark-up, may have aided P3 in completing some of the tasks.

When accessing the transcoded version of Page 2, it took participants longer to find the content than on the original version. P2 actually reached the main headline story within three seconds, but due to rapidly pressing tab to move down the page, failed to identify that it was the main story. This may be because the page was not in a format that P2 expected. If a user is familiar with the layout of a page, then they have a sense of where approximately on the page the information they require is and so, rapidly tab down the page to where they expect the content to be situated. Explicitly stating what content is approaching may be a way of further enhancing the accessibility of content. Inserting a heading with “News Stories” before the news stories started on Page 2 may have helped P2 become aware of what was approaching and ensured they did not jump past the first news item.

Adding explicit headings may also be of particular help for pages that a user is unfamiliar with. Two participants suggested that asking them to perform a task on a blog was unfair as they had never used a blog before. It was observed that during the task on Page 6, which involved a blog, P1 did become very frustrated because they became lost in the page, unable to identify any elements or information. Indeed, P1 described Page 6 as *“rubbish on the screen”*. When the same page was transcoded using SADiE, the participants who performed the task did not have any difficulty with the page. What is interesting to note is that Page 8 was also a blog page yet P1 had no such frustrations even though this too was the original version of the page. Conversely P2, who performed the transcoded task on Page 8 had difficulty. The tasks were randomised so that the participants performed them in different orders. P1 performed the task for Page 8 towards the end of the session, having already completed two additional blog tasks. P2 on the other hand performed the task for Page 8 nearer to the beginning of the session and had only interacted with one blog prior to Page 8. P2 may still have been uncomfortable with the term blog and lacked confidence that they could interact with one, contributing to the overall increase in task completion time. A heading stating “Blog entries” would have explicitly proclaimed the area of the page where the blog entries began and reduced any confusion for the participants, especially those who were unfamiliar with what a blog actually was.

Navigational aids were also important for the participants. During a task on an original page, P3 complained that one of the pages had headings that did not describe what the information was that followed, saying *“when you say ‘heading level 4’ or whatever, it needs to describe what is under the heading.”* This statement was quite noteworthy as headings were often used by P3 to navigate around the page, regardless of the task. Badly used headings hindered their progress and sometimes resulted in P3 guessing what the answer to the task was.

In addition to headings, P3 also discussed the importance of numbered lists. One of the tasks involved finding the first entertainment story on a portal Website (not included in our analyses as only the transcoded version was attempted by users). P3 commented that while he could find the entertainment section because it was marked as heading level 2,

they would have preferred to have “Entertainment” followed by a numbered list of links to the stories. That way a reader can find the stories and know how many stories there are. When they want to read a story, then they can click the link and be taken only to that story and gain immediate access to it. No other participants commented upon the desire for numbered lists so it is difficult to draw solid conclusions from a single participant’s comments. However, Takagi et al. in their study of user behaviour on shopping Websites noted that participants could easily recognise the items that were indicated with numbers, allowing participants to easily navigate between them using navigation commands [22]. Further investigation into the benefit of list structures within Websites is required.

6.2 Using Cues to Locate Information

Structure was not the only aspect of the page that users relied upon to discover information. During the evaluation, participants made use of cues to help them establish their location within the page and then use that to help complete the task. This was particularly useful if participants were familiar with the page and had a sense of the layout and organisation of the content. On Page 4, the task involved finding search results from a search engine. During the task, P1 moved through the content of the page looking for the answer. P1 went beyond the first search result and when they first heard the word “cached” knew that they were in close vicinity of the search result. At this point P1 navigated backwards to find the first search result title and completed the task. During a similar search engine results task, P3 attempted to find the answer to a task by tabbing through the links within the page. When they reached a link entitled “Feedback” they knew that they had reached the end of the search results as feedback occurs at the footer of the page. P3 then started to navigate backwards, reassessing each link until they found the correct answer to the task.

The strategy of looking for cues and keywords was not limited to search engine results. For a task involving a portal Website, P1 was asked “*What is the title of the first headline story?*”. As this was an unfamiliar page, the participant quickly tabbed through the links of the page trying to gain an overview of the content of the page in order to complete the task. This did not provide any clues as to where the information was available, so the participant tried to search for the word “headline” as this word may be found close by the news headlines of the day. In this instance the strategy proved to be unsuccessful. During a later task, the same participant was asked “*What is the title of the first highlight article?*”. Again this was an unfamiliar page but, as in previous examples, P1 searched for a word that may give clues as to where the information might be found, in this case searching for the word “highlight”. These strategies were touched upon during the post evaluation discussion where P1 explained that the “*basic strategy is to use ‘find’ but this only works when you know the page. If you don’t know the page you take a punt, for example, ‘headlines’ but that doesn’t always work.*”.

6.3 User Coping Strategies

Yesilada et al. in their evaluation of DANTE observed phenomena, which they defined as Coping Strategies, whereby users would use a variety of techniques to move around Web pages to reach desired content and information [23]. These

strategies included techniques such as searching for cached to avoid the Google banner and counting the number of tabs to speed up reaching specific areas of the page. In the evaluation of SADIE, similar characteristics were observed from the participants. P1, for example, searched for keywords such as “headline” and “highlight” in order to complete the tasks. When observing P3, it was noted that they made significant use of headings to try and identify areas of the page and reach chunks of information that may contain the answer to the task. P2, on the other hand, quickly tabbed through pages to try and gain an overview of the content. For most pages this strategy worked fine but when the page layout had changed, as with the task for Page 2, this resulted in them missing the main content as it was not in an area that they expected. The disorientation resulted in the task taking longer to complete than what was expected.

An interesting area of research would be to investigate these strategies that visually impaired users exploit when trying to find information within pages. By understanding and modelling the strategies that users employ to find areas of content, it would be possible to create algorithms that adapt pages so that the need for such strategies are reduced and therefore content is easier to reach. While initially this may make familiar pages more difficult to interact with due to the change in layout, overall the transcoding would be of benefit to users. Users who were new to the Web, or were accessing unfamiliar pages, would be able to easily interact with the content from the first instance. This would create a more enjoyable experience than a user repeatedly visiting a page and struggling to interact with it until they had learnt how to deal with the page’s quirks and layouts.

7. CONCLUSION

SADIE is an application designed to help visually impaired users access information on the Web more easily. Previously we have evaluated SADIE from a technical viewpoint and demonstrated that our approach is capable of applying transcodes to a wide variety of pages [2]. In this paper we have progressed the evaluations from a technical perspective to a user-centred perspective. The evaluation was conducted with four participants who were visually impaired and required assistive technology when accessing their computer. The participants were given a series of tasks that involved finding information contained within a series of Web pages. The participants were not aware if they were accessing the original version of the page or the SADIE transcoded version. Analysing the results using the t-test shown that on average, participants took significantly longer to complete tasks with original Web pages than they did to complete tasks using SADIE transcoded pages.

The evaluations have demonstrated that the SADIE approach is capable of supporting visually impaired users access to the Web. By removing unnecessary elements and promoting important information so that it is easier to access, the time taken to complete fact based tasks can be reduced. However, while the evaluation has provided evidence that SADIE is of benefit to users, there still remains further work. It was observed that the participants tried to make use of strategies to find information within the pages that they were unfamiliar with. These included relying on the structure of the page to help comprehend the layout and identifying cues to discover their location within the page. Further investigations are required in order to understand

this phenomenon. We believe that by understanding these strategies, we can develop a model of how visually impaired users navigate through Web content. This model can then be used to provide transcoding more suited to visually impaired users and will further improve the use of SADIE as a transcoding platform that improves access to content within Web pages.

8. ACKNOWLEDGEMENTS

We would like to thank both the sighted and visually impaired users who kindly gave their time to participate in the evaluation. Our thanks also go to Access SUMMIT who helped with the recruitment of participants and the evaluator who created the conditions for a double blind evaluation to reduce the risk of bias entering our results.

9. REFERENCES

- [1] C. Asakawa and H. Takagi. Annotation-based transcoding for nonvisual web access. In *Assets '00: Proceedings of the 4th international ACM conference on Assistive technologies*, pages 172–179, 2000.
- [2] S. Bechhofer, S. Harper, and D. Lunn. Sadie: Semantic annotation for accessibility. In *ISWC '06: Proceedings of The 5th International Semantic Web Conference*, pages 101–115, 2006.
- [3] J. Chen, B. Zhou, J. Shi, H. Zhang, and Q. Fengwu. Function-based object model towards website adaptation. In *WWW '01: Proceedings of the 10th international conference on World Wide Web*, pages 587–596, 2001.
- [4] W. Chisholm, G. Vanderheiden, and I. Jacobs. *Core Techniques for Web Content Accessibility Guidelines 1.0*. W3C, November 2002.
- [5] S. Gupta, G. Kaiser, D. Neistadt, and P. Grimm. Dom-based content extraction of html documents. In *WWW '03: Proceedings of the 12th international conference on World Wide Web*, pages 207–214, 2003.
- [6] S. Harper, S. Bechhofer, and D. Lunn. Taming the inaccessible web. In *SIGDOC '06: Proceedings of the 24th annual conference on Design of communication*, pages 64–69, 2006.
- [7] S. Harper, C. Goble, R. Stevens, and Y. Yesilada. Towel: Experiences of augmenting the web with journey knowledge. In *Proceedings of the The Twelfth International World Wide Web Conference*, 2003.
- [8] M. Hori, G. Kondoh, K. Ono, S. ichi Hirose, and S. Singhal. Annotation-based web content transcoding. In *Proceedings Of The 9th International World Wide Web Conference On Computer Networks : The International Journal Of Computer And Telecommunications Netowrking*, pages 197–211, Amsterdam, The Netherlands, The Netherlands, 2000.
- [9] A. W. Huang and N. Sundaresan. A semantic transcoding system to adapt web services for users with disabilities. In *Assets '00: Proceedings of the 4th international ACM conference on Assistive technologies*, pages 156–163, 2000.
- [10] S. C. Ihde, P. Maglio, J. Meyer, and R. Barrett. Intermediary-based transcoding framework. *IBM Systems Journal*, 40(1):179–192, 2001.
- [11] F. James. Presenting html structure in audio: User satisfaction with audio hypertext. Technical Report 97-201, Stanford University, 1997.
- [12] M. Jang, J. Kim, and J.-C. Sohn. Web content adaptation and transcoding based on cc/pp and semantic templates. In *WWW '03: Proceedings of the 12th international conference on World Wide Web*, page 195, 2003.
- [13] B. Kelly, D. Sloan, L. Phipps, H. Petrie, and F. Hamilton. Forcing standardization or accommodating diversity?: A framework for applying the wcag in the real world. In *W4A '05: Proceedings of the 2005 International Cross-Disciplinary Workshop on Web Accessibility*, pages 46–54, 2005.
- [14] D. Lunn. A quantitative evaluation of the sadie transcoder: Measuring time to task completion. SADIE Technical Report 9, The University of Manchester, School of Computer Science, July 2008. <http://hew-eprints.cs.manchester.ac.uk/54/>
- [15] S. Mukherjee, G. Yang, and I. Ramakrishnan. Automatic annotation of content-rich html documents: Structural and semantic analysis. In *Lecture Notes in Computer Science*, pages 533–549, January 2003.
- [16] K. Nagao. Semantic transcoding: Making the world wide web more understandable and usable with external annotations. In *SSGRR 2000: Proceedings of the International Conference on Advances in Infrastructure for E-Business, Science, and Education on the Internet*, August 2000.
- [17] B. Parmanto, R. Ferrydiansyah, A. Saptono, L. Song, I. W. Sugiantara, and S. Hackett. Access: Accessibility through simplification & summarization. In *W4A '05: Proceedings of the 2005 International Cross-Disciplinary Workshop on Web Accessibility*, pages 18–25, 2005.
- [18] H. Petrie, F. Hamilton, and N. King. Tension, what tension?: Website accessibility and visual design. In *W4A '04: Proceedings Of The 2004 International Cross-Disciplinary Workshop On Web Accessibility*, pages 13–18, 2004.
- [19] F. G. Smith and J. E. Smith. *Key Topics In Clinical Research: A User Guide To Researching, Analyzing, And Publishing Clinical Data*. The key topics series. Taylor & Francis, Mortimer House, 37-41 Mortimer Street, London, W1T 3JH, UK, 1st edition, 2003.
- [20] H. Takagi and C. Asakawa. Transcoding proxy for nonvisual web access. In *Assets '00: Proceedings of the 4th international ACM conference on Assistive technologies*, pages 164–171, 2000.
- [21] H. Takagi, C. Asakawa, K. Fukuda, and J. Maeda. Site-wide annotation: Reconstructing existing pages to be accessible. In *Assets '02: Proceedings of the fifth international ACM conference on Assistive technologies*, pages 81–88, 2002.
- [22] H. Takagi, S. Saito, K. Fukuda, and C. Asakawa. Analysis of navigability of web applications for improving blind usability. *ACM Transactions on Computer-Human Interaction*, 14(3):13, 2007.
- [23] Y. Yesilada, R. Stevens, S. Harper, and C. Goble. Evaluating dante: Semantic transcoding for visually disabled users. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 14(3):14, 2007. ISSN: 1073-0516.