Web Accessibility through Adaptation

Chrisoula Alexandraki[1], Alexandros Paramythis[1], Napoleon Maou[1],
Constantine Stephanidis[1, 2]

[1] Foundation for Research and Technology – Hellas (FORTH)
cs@ics.forth.gr
[2] University of Crete, Department of Computer Science

Abstract. This paper presents the e*AccessibilityEngine* tool, which employs adaptation techniques to automatically render web pages accessible by users with different types of disabilities. Specifically, the e*AccessibilityEngine* is capable of automatically transforming web pages to attain AAA-level conformance to the W3C WAI Web Content Accessibility Guidelines and to "Section 508" of the US Rehabilitation Act. The proposed tool is intended for use as a web-based service and can be applied to any existing web site.

1   Introduction

It is widely acknowledged that people who are faced with difficulties in exploiting the capabilities of the Internet may experience social exclusion. Nevertheless, most web sites are currently partially or entirely inaccessible to people with disabilities [4]. Thus far, the majority of initiatives for promoting web accessibility are oriented towards the design of accessible web sites. These efforts include, for example, guidelines for web design, accessibility evaluation software for checking web page compliance with guidelines, etc.[1] Although this "proactive" approach to ensuring accessibility is, arguably, the most promising long-term strategy, there is a pressing need to re-design existing, non-accessible sites, rendering them accessible in a fully- or semi- automated way.

This paper presents a tool, called e*AccessibilityEngine*, which can automatically introduce web accessibility enhancements in existing web sites. The e*AccessibilityEngine* transforms non-accessible web pages into accessible forms; the actual "output" of the tool can vary in accordance with specific user needs, and the assistive software and hardware available to the user for accessing the web. The primary motivation driving the development of this tool has been to facilitate "reactive" approaches to introducing e-accessibility support in existing web sites. Development to date has fo-

---

[1] A thorough review of existing technologies and tools can be found in "Evaluation, Repair, and Transformation Tools for Web Content Accessibility", available at:
http://www.w3.org/WAI/ER/existingtools.html.

cused on supporting, primarily, the "Web Content Accessibility Guidelines 1.0"[2] and the provisions made in "Section 508" of the US Rehabilitation Act.[3]

The e*AccessibilityEngine* is based on concepts and techniques derived from the field of user-adaptable systems [5]. Specifically, the tool maintains user profiles and uses them to decide upon the type of adaptations to be applied to a document, in order to render it accessible to a particular user. The tool does not engage in active interaction monitoring during run-time, nor does it modify the user model automatically. The later is modified / maintained directly by the users.

The rest of the paper is organised as follows. The next section "Related Work" presents other relevant systems and technologies, while the following one provides an outline of the main operational principles of the tool, its architecture, and the way in which users interact with it. The paper is concluded with a brief overview of ongoing and future work on the tool.

2   Related Work

In recent years, research efforts have focused on design and development tools facilitating the construction of accessible user interfaces for applications and services [6]. Additionally, efforts have been directed towards tools such as Betsie,[4] or the Web Access Gateway[5] that can be installed "on top" of specific web sites with the purpose of providing equivalent but more accessible forms of their pages. However, such tools mainly address the requirements of users with visual disabilities (i.e., colour impairment, low-vision, etc). As a result, their role is typically restricted to removing graphics from web pages, and providing text-only versions, with very limited control and configuration capabilities offered to the user.

A significant body of related work has also been presented in the literature recently, under the general title of "accessibility-oriented transcoding". Although a full review of the area is beyond the scope of this paper, the rest of this section examines a few representative examples for the purpose of pointing out potential limitations or gaps, which the e*AccessibilityEngine* aspires to address.

"Aurora" [2] is an example of a system that uses a semantic-approach to web content transcoding, with the intention of identifying and facilitating multi-step, goal-oriented transactions. The system uses XML-transcoding technologies to extract from web content and present for user interaction only the data relevant to execute these transactions. Aurora, like other tools in this category, provides good results in terms of e-accessibility, but cannot be considered a comprehensive solution that can be applied to a web site as a whole.

At the other end of the spectrum, the framework presented in [1] does represent a comprehensive solution, based on a so-called "ability driven" approach to accessibility-oriented transcoding. Nevertheless, it is not obvious whether users can control the transformation behaviour of the system. Furthermore, several assumptions are made

---

[2] W3C – WAI, Web Content Accessibility Guidelines 1.0: http://www.w3.org/TR/WCAG10/.
[3] Section 508: http://www.section508.gov.
[4] Betsie: http://www.bbc.co.uk/education/betsie/about.html.
[5] Web Access Gateway: http://www.flatline.org.uk/~silas/access.html.

about the way in which the framework is integrated with existing web-based systems, which suggest a rather "intrusive" approach that may not be easily applied to existing web sites.

Another line of work in the field is related to annotation-based transcoding of web documents.[6] External annotations of documents and document transcoding on the basis of these annotations, is a quite powerful approach that can also be used to address accessibility. The downside is that, where existing sites are involved, applying the approach would require not only the integration of a transcoding "engine", but also the creation and maintenance of the external annotations, which may prove impractical as an entry-level solution.

3 An Adaptation-Based Approach to Web Accessibility

3.1 Principles of Operation

The e*AccessibilityEngine* is based on an approach to web accessibility targeted towards addressing physical or perceptual disabilities, or combinations thereof. Disabilities are modeled using stereotypes [3], and each stereotype is associated with a set of accessibility transformations that adapt web pages according to the needs of the respective group of users. An individual user model may comprise of one or more stereotypes (one or more stereotypes may be active for a given user). Users may have full control over the creation of their individual user models, i.e., they can both inspect and manipulate. It should also be noted that the e*AccessibilityEngine* is designed so that new disability stereotypes can easily be introduced and associated with appropriate web document transformations.

The accessibility transformations associated with specific stereotypes are sequences (or "chains") of elementary transformations (Transformation Processing Units, or TPUs for short). TPUs can perform modifications both at the physical level of interaction (e.g., fonts and colours), and at the syntactic level of interaction (i.e., restructuring) of web documents. TPU functionality is based, to a large extent, on the ability to "recognise" and address patterns of inaccessibility in web pages (e.g., HTML tables used for page layout).

For the majority of disability stereotypes, a web page undergoes a first transformation process, which aims at providing Level AAA accessibility according to the W3C - WAI standard. Fig. 1 illustrates the result of an AAA accessibility transformation. Besides complying with the AAA level standard, this initial transformation is designed to have the following supplementary features:

---

[6] W3C, "Annotation of Web Content for Transcoding":
http://www.w3.org/1999/07/NOTE-annot-19990710/.

Fig. 1. Transformation of a web page to an AAA-WAI conformant version

– The resulting page is rendered linearly, in a top-down manner. Any tables used for page layout are removed at this stage.
– "Tab" browsing is supported. A tab index is assigned to every active element of the page (links and form elements), making it possible for the user to navigate through them in a sequential manner using the "Tab" key.
– HTML anchor elements are added to demarcate specific sections of the HTML document (e.g., top of the page, navigation bar, core content of the page, bottom of the page). Additionally, each of these anchors is associated with a specific access key (in effect, a key which acts as a "shortcut" to the anchor). This enables users to jump to these sections with a single key press.
– Icons are added to indicate conformance to standards and specifications (AAA-WAI conformance icon,[7] Valid HTML 4.01 icon,[8] Valid CSS icon,[9] Bobby approved icon [10]).

---

[7] Level Triple-A conformance to Web Content Accessibility Guidelines:
http://www.w3.org/WAI/WCAG1AAA-Conformance.

Usually, the TPU chains defined for certain disability stereotypes perform additional modifications to the AAA-WAI conformant version of the web document. For example, a stereotype profile for blind users could begin from the AAA-WAI conformant version and add, as an extra transformation, the replacement of images by their alternative text. However, certain disability stereotypes do not need to provide AAA accessibility. Consider for example, a user with dexterity impairment. In this case, the full graphic layout of the page may be retained, and the active elements of the page (links and form elements) can simply be enlarged, so as to provide easier interaction to the particular user. The described approach therefore adopts the perspective that interface adaptation is a fundamental mechanism towards achieving optimal accessibility for target user groups with diverse needs and requirements [7].

3.2 Architecture

The e*AccessibilityEngine* is implemented as a web-based service. The communication with the user's web browser and with the web server that delivers the content is done through the HTTP protocol. As a consequence, the tool may be deployed on a separate machine than the web server. The software itself is implemented using the Java™ Platform in conjunction with a variety of XML technologies (XSLT, XPATH), and Application Programming Interfaces (API).

The overall architecture of the e*AccessibilityEngine* (see Fig. 2) comprises of three main components: the Communication Component (CC), the Profile Managing Component (PMC), and the TPU Managing Component (TPUMC). The adaptation process is as follows:

The CC receives a request for a web page, located on a remote web server, from the user's web browser (Fig. 2, (1)), and forwards the complete request (including HTTP headers and parameters) to the remote web server (Fig. 2, (2)). Upon receiving the page (Fig. 2, (3)), the CC forwards the HTTP request headers to the PMC (Fig. 2, (4)). The PMC identifies the user, and transmits the corresponding user profile to the TPUMC (Fig. 2, (5)). Concurrently, the web server transmits to the TPUMC the page to be adapted (Fig. 2, (6)). The TPUMC selects the TPUs to apply on the document, on the basis of the user profile. The result of this selection process is a TPU chain, which is sequentially applied to the document. The result of the step-wise transformation process is a document that has been adapted for accessibility, in accordance with the user profile (Fig. 2, (7)). The adapted document is provided as a response to the client, along with the HTTP response headers previously received from the web server (Fig. 2, (8)).

---

[8] W3C Markup Validation Service: http://validator.w3.org/check/referer.
[9] W3C CSS Validation Service: http://jigsaw.w3.org/css-validator.
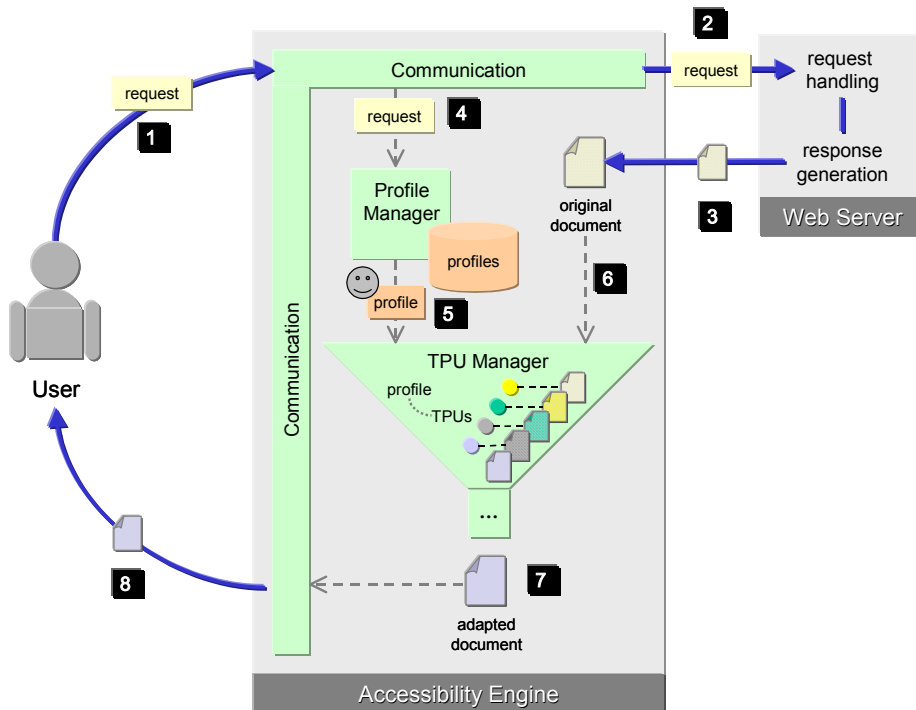[10] Bobby: http://www.cast.org/bobby.

Fig. 2. Overall Architecture of the e*AccessibilityEngine*

Integrating the e*AccessibilityEngine* with an existing web site involves typically the following steps:

– The e*AccessibilityEngine* is installed as a stand-alone web application (installation can be at the same server as the web site is served from, or at a different one). The tool can be deployed as J2EE web application, using any of the available application servers (e.g., Apache Tomcat[11])
– Links are added to the existing web pages, through which the e*AccessibilityEngine* can be invoked. Basically, these links "point" to the corresponding web application and pass the URL of existing page as a parameter. Using these links users can retrieve the "accessible versions" of pages.
– Pages served through the e*AccessibilityEngine* have a one-to-one correspondence to the pages in the original web site. The adapted pages also contain additional links that enable users to interact directly with the e*AccessibilityEngine*, as described in the following section.

It is important to note that the e*AccessibilityEngine* does not pose requirements to the original web site / system, other that the ones outlined above, for its basic operation. It does, however, provide additional facilities that can be integrated on a need-to-have

---

[11] Apache Tomcat: http://jakarta.apache.org/tomcat/index.html.

basis. Such facilities enable, for instance, the manipulation of user models / profiles directly on the associated web system, the configuration of operating parameters of the tool (e.g., caching of adapted documents), the configuration of presentation parameters of the tool (e.g., style sheets, which provide the accessible pages with the same look-and-feel as that of the site), introduce site-specific transformation possibilities (e.g., by defining the standard structural elements of pages), etc.

3.3 User Interaction

Users are not necessarily aware of the presence of the e*AccessibilityEngine*. In fact, if users do not wish to modify their initial selection of an "accessibility profile" (see below), they need to directly interact with the tool only once.

More specifically, the first time the tool is invoked for a particular user, the user is prompted to choose one, or a combination of, the available disability profiles (each corresponding to a stereotype). Invocation of the tool, in this context, is simply achieved by following a link to the "accessible version" of a page.

The user's choice is stored as an individual default profile. Individual profiles may be controlled and refined by the user, by fine-tuning parameters of the transformations that apply for the chosen disability stereotype. For example, a user with vision impairment may choose the preferred font size and colour contrast for viewing the web site (achieved by modifying the options provided by default in the low vision stereotype), or may even define the preferred position for the navigation bar at the top or at the bottom of the page.

4   On-going and Future Work

The e*AccessibilityEngine* is a working prototype system currently under user evaluation. The tool has been successfully employed to provide accessibility support to the SEN-IST-NET[12] web portal (see also acknowledgements).

The main limitation of the current implementation is that accessibility transformations need to be applied to a specific context. In order to register web sites with the e*AccessibilityEngine*, some minimum requirements regarding the content of the pages must be fulfilled. Such requirements include the validity of the markup language (currently HTML, and XHTML) and the use of CSS style sheets. To take advantage of some of the more advanced capabilities of the tool, the "tagging" of the different sections in the web document may be required in some cases (e.g., in order to indicate items of the navigation bar).

Our main objective, which largely defines our ongoing work on the tool, is to build additional facilities into the system, so that comparable levels of functionality can be achieved, without enforcing the aforementioned requirements on the target system. This involves the development of algorithms that can intelligently "tag" or annotate pages at run-time, as well as the integration of functionality that would enable the tool

---

[12] SEN-IST-NET: http://www.senist.net.

to break down pages and serve their individual components to the users (e.g., by utilising a document cache in the system).

5 Acknowledgements

References

1. Balasubramanian, V., & Venkatasubramanian, N.: Server Transcoding of Multimedia Information for Cross Disability Access. ACM/SPIE Conference on Multimedia Computing and Networking (MMCN 2003). Electronically available at: http://www.ics.uci.edu/~dsm/suga/servertrans-mmcn.pdf (2003)
2. Huang, A., Sundaresan, N.: Aurora: A Transaction-Based Approach to Web Accessibility. In Poster Proceedings of the 9th International World Wide Web Conference (WWW9) (2000)
3. Kay, J.: User Modeling for Adaptation. In: Stephanidis, C. (ed.): User Interfaces for All - Concepts, Methods, and Tools. Lawrence Erlbaum Associates, Mahwah NJ (2001) 271-294
4. Lazar, J., Beere, P., Greenidge, K-D, Nagappa, Y.: Web accessibility in the Mid-Atlantic United States: a study of 50 homepages. Universal Access in the Information Society, 2 (4) (2003) 331 – 341
5. Stephanidis, C.: Adaptive techniques for Universal Access. In: User Modelling and User Adapted Interaction International Journal, 11 (1/2) (2001) 159-179
6. Stephanidis, C. (ed.): User Interfaces for All - Concepts, Methods, and Tools. Lawrence Erlbaum Associates, Mahwah NJ (2001)
7. Stephanidis, C.: User Interfaces for All: New perspectives into Human-Computer Interaction. In: Stephanidis, C. (ed.): User Interfaces for All - Concepts, Methods, and Tools. Lawrence Erlbaum Associates, Mahwah NJ (2001) 3-17