

The Meaning of 'Life': Capturing Intent from Web Authors

Rhys Lewis

Volantis Systems Ltd

1 Chancellor Court, Occam Rd,
Surrey Research Park, Guildford UK

+44 1483 739 775

www.volantis.com

rhys@volantis.com

ABSTRACT

Interest in accessing the Web from small, mobile devices, such as cell phones, is increasing rapidly. The challenge of delivering content to such devices is similar in many ways to the challenge of delivering it to users with disabilities. There is a real synergy between these use cases which offers the hope that solutions applicable to one will also be applicable to the other. This presentation will examine the ways in which recent work in standards, being driven by the need to support mobile Web users, may also help to improve accessibility.

Categories and Subject Descriptors

K.4.2 [Computers and society]: Social Issues - *Assistive technologies for persons with disabilities.*

General Terms

Human Factors, Standardization, Languages.

Keywords

Semantics, Adaptation, Web, Authoring, Mobile.

1. INTRODUCTION

For the majority of users, the Web is primarily a visual experience. The availability of network bandwidth and highly capable display technology mean that many web sites are lively, colorful and compelling.

However, in achieving sites that are interesting and usable by the majority, authors and designers often use techniques that rely heavily on visual cues. The placement of material within a page, the color and styling of elements, and juxtaposition are all commonly used to convey meaning. Such techniques work well if the page is used in the context originally envisaged by the designer. However, these techniques tend to fail if the user context differs from that envisaged. For example, users with certain kinds of disability may not be able to view the site at all, or may

need to zoom the page to perceive its text and graphics. In this case, the visual context and cues that rely on juxtaposition may be lost. Similarly, users accessing the site from small, mobile devices will only be able to see a small part of each page at any one time, losing the context and relationship between the sections that it contains.

In a number of cases, the constraints placed on delivery of web pages to small mobile devices mimic those that have been experienced by users with disabilities since the inception of the Web. Improvements in user experience require more explicit representations of meaning than are currently in common use by web authors.

The title of this paper, while more than a little contrived, emphasizes the issue of the lack of semantic information in most web pages. According to one dictionary [11], the word 'life' can be used as a noun or an adjective. There are a dozen or so different meanings for the word when used alone, and several additional meanings when it is paired with other words. The terms 'life preserver' and 'life science' are examples.

The problem of trying to analyze the meaning of the word 'life' within a sentence is akin to the problem of trying to analyze part of a web page. Context is often of vital importance in allowing the true meaning to be understood. The problem with much of today's Web is that the context is not explicit. It relies on a particular interpretation of the rendered page by the user. Anything that alters that interpretation, such as a disability, or a device that has characteristics that the designer did not expect, can inhibit the user's understanding of the page.

It is worth examining for a moment the kinds of semantic information conveyed by web pages. It is possible to think of these in two broad categories. First, there are the semantics associated with the application itself. For example, a web site dedicated to movies is likely to encapsulate concepts such as reviews, tickets, prices, movie theatres and times of shows. These concepts are likely to be represented in the user interface of the site. However, they will almost certainly be expressed as some combination of markup and script. These representations will not directly encapsulate the application level concepts. Rather, they will implement a user interface for viewing or manipulating the data associated with the concepts. For example, the proposed movie site might provide forms from which users can access a movie review or purchase tickets.

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, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

W4A at WWW2006, 23rd-26th May 2006, Edinburgh, UK
Copyright 2006 ACM 1-59593-281-x/06/05...\$5.00.

This specification of the user interface leads to the second general category of semantics. Even at the simplest level, sites embody general user interface semantics. By convention, particular arrangements of user interface elements have become commonly understood. A set of links arranged horizontally across the page near its top edge is understood to be the primary navigation for the site, leading users to major sections. It has become common to term such an arrangement a ‘navigation bar’. However, there is no explicit representation of this abstraction in current markup languages, such as XHTML Version 1. The implication that such a set of links has some specific meaning is conveyed purely by visual cues, such as position, color and background.

Another example of user interface semantics conveyed visually relates to simple XHTML forms. It is very common for the relationship between an input field and its label to be defined simply by juxtaposition. The label is placed so that it is adjacent to the field when the page is viewed in the way envisaged by the designer. As with the navigation bar, unanticipated changes in the way in which the page is rendered can defeat the intent of the author in conveying information.

Logo	My Account Login	
	Book Tickets	
	City <input type="text"/>	Movie <input type="text"/>
New this week	Search for Movies Advanced	
Movie #1	Title <input type="text"/>	<input type="button" value="Go"/>
.....	Character <input type="text"/>	
Movie #2	Featured Movie	
.....		
Movie #3	Most Popular	Vote for your favorite movie
.....		
Movie #4	<input type="radio"/>
.....	<input type="radio"/>
Movie #5	<input type="radio"/>
.....		<input type="radio"/>
.....		<input type="button" value="Vote"/>

Figure 1. Example of a moderately complex web page

These broad categories of application and user interface semantic information have characteristics that differ somewhat and suggest that different approaches to their representation may be appropriate. Application level semantics are essentially unconstrained. After all, a web site might deal with virtually any subject. Representation of

application level semantics effectively requires the full power of semantic web approaches [14]. These representations need to be associated, in some way, with the page markup. We’ll look at some examples in the section on Semantic Enrichment.

In contrast, user interface semantics form a far more restricted set. Annotation may still be appropriate for some representations. However, it seems possible that at least some of these semantics might be represented directly in the markup languages that express the interfaces themselves. Explicit representations tend to be simpler for authors than annotation. However, they also tend to require more comprehensive support in markup languages. We’ll return to this topic in the section on Semantically Rich Languages.

Providing additional semantics in authored materials is one thing. Using them is, of course, quite another. However, once again, the needs of assistive technologies, which help those with disabilities access the Web, mirror those of systems which adapt content for mobile devices. In both cases, the additional semantics help the systems involved interpret the available materials in order to provide a more appropriate user experience. We’ll look further at this topic in the section on Semantics and Adaptation.

The W3C is actively involved in standards development related to web access for the users with disabilities and web access from a wide variety of different types of device. We’ll look at the work of the W3C Mobile Web Initiative (MWI) [13] and of the W3C Device Independence Working Group (DIWG) [12] in the section on Device Independence and the Mobile Web.

2. TRADITIONAL APPROACHES

The traditional approach to solving the problems of web access for users with disabilities and users with mobile devices is essentially the same. It involves processing the markup sent to the device to try and provide a more appropriate representation. In the case of users with disabilities, the function is usually carried out by some assistive technology that runs in the computer being used to access the Web. In the case of mobile devices, it usually runs as a server side process that transforms the material between its origin and the user’s device. This transformation process is normally termed ‘transcoding’ and the processor which performs it is usually termed a ‘transcoder’. The work of an assistive technology can also be viewed as transcoding.

Transcoding relies on the interpretation of markup created for a page by an author who almost certainly assumed that it would be used on a typical desktop or notebook computer by an able bodied person. As we’ve already noted, a significant portion of the semantics of most web pages is conveyed in the visual arrangement of elements. It is not explicitly present in the markup that a transcoder processes. In addition, lively, dynamic, interactive sites that provide compelling user experiences tend to make heavy use of

client-side scripting. This can mean that some of the semantics of the site are embodied within program code that executes within the browser. Not only might a transcoder or assistive technology need to interpret the markup of a page, it may also need to try and interpret the program code within the scripts that are used.

Figure 1 shows the wire frame layout for a moderately complex web page. This hypothetical example is taken from the movie web site mentioned earlier. We'll use it to illustrate some of the challenges that transcoders and assistive technologies can face.

The page is laid out as a series of sections on a two dimensional grid. Within each section, there are further subdivisions. The relationship between sections and subsections is not necessarily explicit. For example, consider the section 'New this Week'. The layout shown in the figure could be achieved by using a table in which each movie has its own row. The title of the section simply occupies the first row of this table. The movies and the title are at the same level of nesting. In this case there is no explicit containment relationship that can be used, by a transcoder or assistive technology, to label this part of the page as 'New this Week'. The grouping intended by the author is achieved implicitly by visual juxtaposition. This same issue might afflict any of the sections shown in Figure 1, of course.

The forms in Figure 1 use a variety of different layouts. Once again, the relationships between labels and fields are achieved visually, making interpretation of the markup difficult. In the case of the ticket booking section, the labels and fields are not even in the same row of the table that the author uses for layout. They are simply text items that happen to appear within the same table as some form fields.

One other relatively common approach to form input can also be illustrated by Figure 1. The author has decided use just a single form processing URL. Consequently, the 'Vote' button, the 'Go' button and the scripts associated with the drop down 'City' and 'Movie' selection lists all submit to the same URL. All of the input controls are associated with the same form. When the form is submitted, the processing code examines the fields that are present to determine the function to be performed.

While this may be convenient for the author, it makes life much more difficult for interpretation. Not only is the relationship between input fields and labels implicit, so is the function of the submission buttons. And, of course, one of the submission methods is hidden within the scripts that process the 'City' and 'Movie' drop down selection lists.

Now of course, there are authoring approaches that are less unfriendly to transcoders and assistive technologies than those described here. However, even really well written pages require some level of intelligent guesswork, often called heuristics in polite circles, to interpret their semantics.

Although this is a contrived example, it should be clear by now that even relatively simple web pages can hold real challenges for automatic interpretation. It should also be clear that at the heart of these challenges is the lack of semantic information in the page.

3. SEMANTIC ENRICHMENT

If the issue of interpretation is caused by a lack of semantic information, then one obvious strategy is to make up the deficit. This approach goes under the general title of semantic enrichment. The approach concentrates on giving the authors the tools needed to make their intentions clearer by adding semantics explicitly. In the last year or so, considerable progress has been made in the definition of the kinds of semantics that can aid interpretation of web pages.

Late in 2004, the W3C DIWG [12] held a workshop on Metadata for Content Adaptation [4]. Mechanisms for describing additional semantics in web pages were discussed and some general principles were defined. The notion of the 'role' of a particular part of a web page was identified as a key item of semantic information. The relationship between parts of a web page was also identified as a key item.

The W3C Web Accessibility Initiative Protocols and Formats Working Group (WAI-PF) [16] is currently developing a taxonomy [9] for the roles associated with particular parts of a web page. The taxonomy identifies roles associated with user interface components, such as those on forms. It also identifies various structural roles, such as menus, toolbars and lists.

Recent versions of the Mozilla Firefox browser incorporate support for additional semantic information, such as the roles defined by WAI-PF [16]. Information about the support in Firefox for accessible dynamic HTML (DHTML) is available [17].

Support for the addition of semantic information is based on a new attribute used in the markup of the page. This attribute, named 'role' was first proposed for inclusion in a major new version XHTML [1]. Subsequently it has been proposed for and included in implementations of other markup languages [17]. The approach has very little impact on languages into which it is introduced. For example, the following XHTML markup fragment identifies a table whose role is to behave as a spreadsheet.

```
<table id="table1" x2:role="wairole:spreadsheet" >  
...  
</table>
```

The prefix values 'x2' and 'wairole' identify XML namespaces that remove problems associated with duplicate names being used for different purposes.

have been omitted. These would apply to the entire page in which such markup appears. The semantic information is conveyed by the value of the role attribute. The meanings

of values like ‘spreadsheet’ are defined in [9]. Using them, authors and those responsible for interpreting markup can gain a shared understanding of the semantics of the user interface.

Although the current work is being driven by the needs of assistive technologies, the roles that are being defined have general utility. They are capable of labeling parts of a page with the author’s intended user interface semantics. As such, they are as applicable and useful for assisting transcoding as they are for supporting assistive technologies.

The ability to annotate sections of web pages with semantic information does not, of course, remove the need to structure the page appropriately. In particular, there must be markup that represents the section to be annotated. As we saw earlier in the discussion of the section ‘New this Week’ in Figure 1, this may not always be the case with existing web pages. Nevertheless, the definition of a taxonomy of user interface roles and the appearance of support for it within a leading browser represents considerable progress.

4. SEMANTICALLY RICH MARKUP LANGUAGES

Semantic enrichment provides a powerful way to extend markup languages by retaining additional information about the author’s intent. As we’ve already noted, a complementary mechanism is the use of semantically rich markup languages. Where a facility can be provided directly by a language, rather than requiring semantic enrichment, it is arguably simpler for authors to use and less liable to error.

4.1 User Interface Semantics

Commonly used markup languages do not, as yet, contain such a rich a set of facilities as are described in [9]. However, newer W3C markup languages are introducing additional abstractions which improve the level of semantic information available within web pages. A good example, which overcomes a number of the issues associated with Figure 1, is XForms [2]. This replaces the traditional HTML forms support with a semantically richer and more capable set of facilities. In addition, XForms provides much more explicit linkage between the various components of a form. Control of the way in which form components are rendered is removed from the markup and made a concern only of styling. For example, the operation of selecting of one option from many is defined by its semantics, ‘single selection’, and not by its representation, ‘radio button’ or ‘drop down list’.

In addition to its interest in semantic enrichment, the W3C DIWG [12] has also been pursuing the notion of semantically rich markup languages to support use of the web on a wide range of devices with different characteristics. Rather than defining entirely new languages, DIWG has based its work on other W3C specifications. In particular, work has focused on XHTML

2.0 [1] and XForms [2] as the basis for a Device Independent Authoring Language (DIAL) profile [10]. To this base, DIAL adds facilities that are particularly useful for authors that must support a variety of different kinds of device.

Some aspects of support for different devices can be automated, but some may require additional work by authors. For example, consider the task of helping a potential customer travel to a shop to collect some goods. This is usually accomplished by displaying a map. Often this is supplied by one of the on-line services readily available in many countries. The map is composed of one or more images showing locations, roads and other features. Such a representation is appropriate for an able bodied person using a typical desktop or notebook computer. For someone with a visual impairment, or who is accessing the site using a mobile phone, the map may be of little or no use. For example, automatic transformation of the map to a smaller size appropriate for the user’s mobile phone may render it illegible.

One solution, in this particular case, is for the author to provide an entirely different form of information in addition to the map. An alternate, textual description of how to find the shop may be more appropriate for the phone user than another visual representation. For a visually impaired user, spoken material may be of more appropriate. In either case, an improved user experience depends on the author providing alternative versions of the materials. The ability to create such alternate representations, and to have them delivered when appropriate, is one of the extensions on which DIWG is currently working [6]. This capability also forms part of the DIAL profile.

DIWG is also working on the provision of explicit mechanisms for authors to define page layout separately from page content. Current practice for page layout often involves the subversion of the XHTML table mechanism or the use of advanced CSS. Using tables means that the layout of a page is embedded in the markup and cannot easily be changed for use on different devices. CSS currently does not include sufficiently subtle means of associating styling with different devices. It also has limited support for coarse-grained definitions of page layout. These are the aspects of particular interest to DIWG.

Often, the kinds of layout change required in order to make a page render well on a different device are rather simple. It may only be necessary to move sections of the page in relation to one another. For example, one way to rearrange the page in Figure 1 might be to convert it to a single column layout with the sections following one another in a particular order. This kind of layout change can be achieved if each section is associated with a specific area in a layout representation held externally to the page markup. Sections can be moved simply by changing the associated layout without affecting the page markup itself. Different layouts can be used to support devices with different

capabilities. Commercial implementations, such as that provided by Volantis, have been available for a number of years. The approach has been found to be versatile and very effective.

4.1.1 Richer User Interfaces

Recently, W3C has initiated new work in the area of richer user interfaces with more capability. The Web Application Formats Working Group (WAF WG) [19] is looking at ways to enhance existing W3C specifications using the results of work by specific browser manufacturers. The aim is to provide more capable platforms for web application development. Part of this work is likely to lead to richer user interface abstractions with higher semantic content.

4.2 Application Semantics

As we noted earlier, application level semantics relate to the concepts embodied in some application, rather than those of the user interface by which it is represented. In the hypothetical movie web site, such concepts include reviews, tickets, prices, movie theatres and times of shows. It is possible to construct markup languages that deal explicitly with such concepts. One common approach is to create a language based on XML by defining an appropriate schema [3]. Such a language would include markup that explicitly represents the key concepts. For a movie, for example, there might be explicit representations of the title, director, leading actors, genre, audience suitability and so on. The important characteristic of such markup languages is that they represent only the semantics and not the associated user experience. Such representations can be adapted to create a user experience, but do not explicitly define it. An adaptation step is used to convert the semantic representation to one that can be used in a web page. This step can create different markup for different classes of device if necessary. However, whereas adaptation of languages that represent user interface semantics can be generalized, adaptation of application level semantics is tightly coupled to the application itself. Applications are too numerous and too varied for there to be much likelihood of general agreement about the form of the semantics employed.

Over the last few years, a number of mobile operators have created languages that include some application level semantics. In their desire to provide data services over their networks, mobile operators have, historically, provided systems that distribute specific web pages to their customers. These pages are accessed via a home page provided by the operator. Applications and pages are provided by companies and organizations that are in partnership with the operator. This sort of arrangement is often known as an operator 'portal'. The operator provides the language used to create pages within this portal environment and is responsible for adaptation of those pages to allow them to work on any device that is supported. Such languages usually provide some form of user interface abstraction, but also provide abstractions that relate to the portal itself, or to applications that run within

it. However, every operator language is different and encapsulates different aspects of the application semantics. This makes it very difficult for application and content partners who wish to provide materials for multiple operators. They may have to rewrite their application for use on different operator's networks.

5. SEMANTICS AND ADAPTATION

In the traditional Web, the concept of a 'page' is fundamental. Users access pages. Authors create pages. A page is the unit returned in response to a request from a user. Actually, that last statement is not entirely true, since images and other media are delivered separately. However, it's broadly true to say that what the user perceives, what the author creates and what the Web delivers in a single request-response cycle are essentially the same.

In its efforts to describe a system that could deliver materials to a wide range of different devices with very different capabilities, DIWG generalized the definitions of the concepts associated with web pages to allow a more precise description of the associated architecture. We've already noted that under certain circumstances, authors may need or indeed want to provide alternative representations of specific materials. This immediately suggests that a page is not an indivisible item. Rather, it is composed of one or more 'authored units' [7], sets of materials from which the actual user experience will be constructed.

Even in the traditional Web, the delivery of materials to a browser usually occurs as several distinct steps.

Logo
Search for Movies
Title
<input type="text"/>
Character
<input type="text"/>
Go
Advanced

Figure 2. The page from Figure 1 after adaptation

Pages may refer to style sheets, script functions, images and other media, that are requested separately. DIWG has defined the term 'delivery unit' [7] to describe a set of material transferred in a single request-response cycle.

Using these definitions, we can describe the process of adaptation in the context of the request for a web page from a particular device. Look again at Figure 1. Let's suppose that the author constructed this page knowing that it would need to be adapted for use on some particular mobile devices. In particular, the author created an authored unit for each of the major sections in the page, such as 'New this Week' and 'Book Tickets'. When the URL is accessed from a traditional web client, all of the authored units are aggregated to form the delivery unit, which is returned to the client. The user sees the entire page as in Figure 1.

Now suppose that the same URL is accessed from a small mobile client. For this client, the author has defined that only the 'Logo' and 'Search for Movies' sections should be returned. The process of adaptation selects just those authored units when composing the delivery unit. In addition, let's suppose that the author has specified an alternate image for the 'Logo' section for use on this particular type of device. The adaptation system selects that version for use in the authored unit that represents the logo. The delivery unit is returned to the client and the user sees a page that contains only the selected sections. Figure 2 shows how the resulting page might appear to the end user. Notice that the layout of the materials in the search section is different from that in Figure 1. The various controls have been moved to positions that are more appropriate for the portrait style display screen of the mobile device.

So far, this example has described an adaptation process in which materials are selected and modified, but has said nothing about how the author makes their intent known. How does the author define which sections are to be used and how do they reorganize the materials within those sections? Even without getting into the details, it is possible to see that the example illustrates transformations based on both application semantics and user interface semantics.

The initial decision about which sections of the full page should be included is based on application semantics. The resulting page on the mobile device must be functional. For example, if the button marked 'Go' were not within the materials selected for the mobile device it would not be possible for a search to be submitted. This may seem rather obvious. However, it's important to remember that the association between controls in conventional web pages often relies on visual juxtaposition. In traditional web pages, there is no guarantee that the markup that generated Figure 1 even contains a single structure that represents the search section. Let's assume that the page is constructed appropriately and that sections are represented explicitly. Each section could be labeled semantically to define its purpose. Since these are application semantics, precise labeling could involve the use of techniques such as the 'role' attribute, as we saw in the section on Semantic Enrichment. The adaptation system might then be sensitive to such roles and might be able to use them in selecting which sections are used for the mobile device in question. Of course, at some point the author would need to create the rules about whether sections with specific roles are delivered to specific types of device.

Alternative mechanisms, based on the notion of the priority or importance of a section were discussed in the workshop described in [4]. In this approach, sections are given numerical priorities. Authors also define some threshold of priority for each type of device to control which sections are delivered. Though less precise than role-based labeling, this mechanism adds semantic information to the sections that can be used by the adaptation process.

In contrast with these selection mechanisms based on labeling, the alteration to the layout within the search

section could be viewed as purely a user interface adaptation. The material in the section has not changed. However, its representation has been altered. In this case, approaches based purely on styling and layout are possible. One approach might be to specify a style sheet for use on the mobile device which differs from the one used for access from desktop and notebook systems.

One additional layout needs to be considered when comparing Figure 1 and Figure 2. Not only does each section have its own internal layout, but in addition, the sections are laid out within the page. This layout differs between the two versions of the page. Some mechanism is needed that allows authors to specify this sort of coarse-grained layout separately from the authored units that define each section. As with the style sheets, different versions of such layouts could be used with different devices. Unlike the case with style sheets, there are no agreed standards for such definitions yet. However, a number of commercial implementations exist and DIWG is working on this topic with other groups in W3C.

In addition to the kinds of technique described here, there are many other ways in which authors can specify materials that are used in adaptation. DIWG is currently working on extensions to W3C markup languages that allow such information to be expressed by authors.

5.1 Delivery Context

As we have seen, adaptation uses the materials provided by authors to create versions of a page, appropriate for use on particular devices. To enable the adaptation to be appropriate, certain information about the target device must be available. DIWG has chosen to call this sort of information the 'delivery context' [7]. This name reflects the fact that information that influences adaptation could be more than just the characteristics of the device itself. In particular, characteristics related to intervening networks might also influence adaptation, as might the personal preferences of the user.

Device-related information in the delivery context might include items such as:

The physical size of the device's display screen in some linear measure

The size of the device's display in pixels

The number of colors that the device can represent

Markup languages supported by the device

Image formats supported by the device

Audio formats supported by the device

etc.

Information in the delivery context is used during adaptation to select or create suitable representations from the materials provided by the author. For example, if an author has created several versions of a company logo, an adaptation processor might use display size information, color capability and supported image formats, to select the most appropriate version from those available. If none of

the available versions were appropriate, the processor might be able automatically to generate a new image from those available.

Likewise, by knowing that a device requires cHTML or WML, for example, an adaptation processor can automatically transform the markup used to create authored units into one appropriate for the delivery units.

The category of personal preferences in the delivery context has been a topic of discussion within the accessibility community. It's relatively common for disabled users to influence the user experience they receive explicitly through settings associated with their browser. This might be as simple as altering text sizes, or as complex as creating alternate style sheets.

The ability for an adaptation process to be influenced by personal preferences offers at least the possibility that, in future, much more sophisticated control of the user experience might be available to users with disabilities. However, there is a challenge. By its nature, adaptation is distributed, and might take place anywhere in the chain of processing between the user's device and the origin server. The basic architecture proposed by DIWG shows delivery context flowing throughout the processing chain. If user preferences are part of the delivery context, this implies that some level of personal information will be transmitted from the device into the network. This raises questions of security and even of personal safety. Already, research has shown that it is possible to reason about a user from the sort of information that would be available in such a context (see for example position paper number 26 in [18]). Clearly, users making this sort of information available need to be assured that it will not be used inappropriately. This leads to the need for trust relationships between users and the systems that they use.

6. DEVICE INDEPENDENCE AND THE MOBILE WEB

In recent years, improvements in the capability of mobile devices and in the networks they use have led to renewed interest in the provision of general web access. In 2005, the W3C formed the Mobile Web Initiative (MWI) [13] to provide a focus for standardization work associated with web access for mobile devices. The initiative has received support from a broad range of organizations.

The overall goal of the initiative is to enable the greater use of the Web on non-traditional devices, such as mobile phones and other handheld systems. It will achieve its goals in a number of ways that are complementary to those of other groups within the W3C. Some of its work builds on W3C specifications by providing help and guidance to authors in the form of best practices (see for example [8]) and other outreach programs. Some of its work will result in additional requirements and clarifications that will be used by other W3C groups. Indeed, DIWG is already being assisted in its work by information provided by groups within the MWI. The method of operation of the initiative is similar in nature to that of the Web Accessibility

Initiative (WAI) [15] at W3C. Groups within both initiatives provide a focal point for expertise in their respective fields. They influence developers of W3C specifications and provide specific help and guidance. Although the ultimate objectives of these two initiatives are different, it is recognized that there are common solutions that can benefit both. There is a close working relationship between their members, as there is with other, relevant working groups at W3C.

7. CONCLUSION

There are many tensions on the Web. At one end of the spectrum are social and political tensions, such as the conflicts between free speech and censorship. At the other, are tensions between different organizations who would like to exercise control at the technical level. Somewhere between these extremes is a tension that has existed as long as the Web itself. It is the tension between authors and end users.

In many cases, authors need to control the precise look and feel of a web site to meet specific requirements of the organizations for which they work. Many organizations have very strong views on how their sites must appear and how they must behave. Style specifications are common, and are often strictly applied. This desire for controlled look and feel is also common in the kinds of operator portal provided for mobile users and discussed earlier.

Users of web sites, on the other hand, may wish to alter the way in which a site renders. Simple changes, such as a modification in overall font size to improve readability or for use during a presentation, are common. As we have seen, users with disabilities may need to alter many aspects of the rendering in order to use a site effectively. Such changes are anathema to an author who has spent a great deal of time and effort in meeting the requirements of a style guide. They may also impair overall usability when compared with the original look and feel. However, they may be crucial for the end user with a disability.

Part of the problem is that those who specify style guides and those who implement them tend to consider only able-bodied users with the same kind of web access as they have themselves. In addition, the technologies that underpin the Web and the tools that have been built on them, have not really given authors the tools they need to cater for a wide range of different kinds of access.

The increasing interest in support for users who are mobile and who use devices very different from those used by authors, is forcing a rethink. There are commercial pressures on organizations to provide a mobile presence, and this is compelling them to devise and use solutions that can support a much wider diversity of access. Adaptation and the provision of alternative content, styling and layout plays a key role in such solutions.

Interestingly, the same kinds of approach to authoring that can support access from diverse devices can also be used to assist in supporting some types of disability. After all, the

display on a small mobile device behaves in a very similar way to the display on a desktop system on which a web page has been zoomed up by a user with low visual acuity.

Work amongst vendors and standards bodies is progressing to the point where the technologies needed to support access from diverse environments will be widely available. If authors adopt such technologies, the very same materials needed to support mobile access may also play an important role for users with disabilities.

At present, support for disabilities tends to place a burden on the end user. Often, the user is required to control web page rendering explicitly to achieve acceptable results. The changes that a user needs to make may affect the page adversely, causing it to be more difficult to understand and use.

An intriguing possibility for the future is that it might be possible to serve appropriate content automatically to users with at least some classes of disability, using the very same techniques used to support mobile devices. Sites would effectively tailor themselves to users via the same techniques by which they tailor themselves to different devices.

Much work remains to be completed. However, the renewed focus on addressing these problems, which has resulted from the need to support small mobile devices, is driving progress in the development of standards that may be applicable in addressing at least some of the needs of disabled users.

8. REFERENCES

- [1] Axelsson, J, Birbeck, M., Dubinko, M., Epperson, B., Ishikawa, M. McCarron, S., Navaro, A., Pemberton, S. XHTML™ 2.0 *W3C Working Draft* (May 2005) publicly available at <http://www.w3.org/TR/2005/WD-xhtml2-20050527/>
- [2] Boyer, J. XForms 1.1 *W3C Working Draft* publicly available at <http://www.w3.org/TR/2005/WD-xforms11-20051209/>
- [3] Fallside, D.C., and Walmsley, P. XML Schema Part 0: Primer Second Edition. *W3C Recommendation* (October 2004) publicly available at <http://www.w3.org/TR/xmlschema-0/>
- [4] Hanrahan, R. *W3C Workshop on Metadata for Content Adaptation: Executive Report* (Oct. 2004) publicly available at <http://www.w3.org/2004/06/DI-MCA-WS/execreport.html>
- [5] Lewis, R. Authoring Challenges for Device Independence. *W3C Note* (Sep. 2003) publicly available at <http://www.w3.org/TR/acdi/>
- [6] Lewis, R., and Merrick, R. Content Selection for Device Independence. *W3C Working Draft* (March 2006) publicly available at <http://www.w3.org/TR/cselection/>
- [7] Lewis, R. Glossary of Terms for Device Independence. *W3C Note* (Jan. 2005) publicly available at <http://www.w3.org/TR/gloss>
- [8] Rabin, J., and McCathieNevile, C. Mobile Web Best Practices 1.0. *W3C Working Draft* (Jan. 2006) publicly available at <http://www.w3.org/TR/2006/WD-mobile-bp-20060113/>
- [9] Seeman, L. Role Taxonomy for Accessible Adaptable Applications. *W3C Working Draft* (Nov. 2005) publicly available at <http://www.w3.org/WAI/PF/GUI/roleTaxonomy-20051106>
- [10] Smith, K., and Lewis, R., Device Independent Authoring Language Profile (DIAL) 1.0, *W3C Working Draft* (March 2006) publicly available at <http://www.w3.org/TR/DIAL/>
- [11] Webster's New American Dictionary. Smithmark, New York, 1995.
- [12] The Device Independence Working Group (DIWG) at W3C <http://www.w3.org/2001/di/>
- [13] The Mobile Web Initiative (MWI) at W3C <http://www.w3.org/Mobile/>
- [14] The Semantic Web Activity at W3C <http://www.w3.org/2001/sw/>
- [15] The Web Accessibility Initiative (WAI) at W3C <http://www.w3.org/WAI/>
- [16] The WAI Protocols and Formats Working Group (WAI-PF) at W3C <http://www.w3.org/WAI/PF/>
- [17] Accessible DHTML at <http://www.mozilla.org/access/dhtml>
- [18] *W3C Mobile Web Initiative Workshop: Position Papers* (October 2004) publicly available at <http://www.w3.org/2004/10/MWIWS-papers/>
- [19] The Web Application Formats Working Group (WAF WG) at W3C <http://www.w3.org/2006/appformats/>