

BEYOND PROXIES: XLINK SUPPORT IN THE BROWSER

Angelo Di Iorio¹, Gabriele Montemari² and Fabio Vitali³

^{1,2,3} Department of Computer Science, University of Bologna, Bologna, IT
 {diiorio,montemar,fabio}@cs.unibo.it

ABSTRACT

In this paper, we describe some issues that every implementation of external linkbases using XLink has to consider, and provide a number of possible solutions. We particularly consider as relevant the issues connected to providing navigation-related features, that is, adding external links to the actual web pages before displaying them, and to providing creation-related features, that is, the user interface elements necessary to users to create new links, or modify and delete existing ones. Examples are carried out based on our experience with two different architectures for XLink application, a proxy-based solution called XLinkProxy, and a browser-based solution for Mozilla browsers called XLinkZilla, which is described in this paper for the first time.

KEYWORDS

XLink, external links, linkbases, Mozilla browsers

1. INTRODUCTION

One of the reasons personalized links are not gaining importance in hypertext systems could be reasonably attributed to the lack of adequate tools to create custom links on hypertext content. It is pretty obvious to say that the expressive power of a hypertext system primarily depends on the expressive power of the underlying linking mechanisms. So the evolution of the hypertext should have been parallel with the evolution of its linking mechanism.

Unfortunately this is not true for the World Wide Web, where the evolution of the linking mechanisms has not gone hand in hand with the impressive diffusion and evolution of the whole system. Drawing from the taxonomy proposed in [7] and [9], in fact, three main models for storing link information on a document set can be identified: the *embedded link model*, the *anchor table model* and the most flexible *external link model*. However the Web still relies on the simple and limited *embedded link model* [1].

A substantial flexibility in link models has been finally made available to the WWW through the XLink standard [11], which enhances links in web pages by introducing many improvements to the plain HTML link model, among which the possibility to express external, multiple and generic links. So far, though, web browsers have been slow to adopt this innovation. Mozilla, for instance, claims support for XLink, but only limited to simple links.

While waiting for support for sophisticated link types to become an integrated feature of popular browsers, many research prototypes have circumvented the problem by converting the set of sophisticated link features available with XLinks into plain HTML anchors (possibly enriched with some JavaScript and CSS to account for the additional services available with XLinks, see [14]) and placing them on-the-fly in the original unmodified document before the result is shown within a standard WWW browser. Another feature to be considered is providing the users with solutions and user interfaces to create and update these links, allowing readers to create personal links during the browsing.

Proxies are a popular architecture for this doctoring of web pages with additional links before sending them to standard browsers. Proxies in fact are easy to write and set up, and they are a regular part of the HTTP chain of connections. Our own XLinkProxy [4], as well as other projects such as Goate [14] or Webvise [15], are proxy-based experiments with XLink. Proxies have also difficulties and drawbacks, though, mainly related to subtle complexities in providing a user interface, and to the restriction policies that some networks might impose on their use.

Many current generation browsers, on the other hand, allow custom widgets such as menu items and sidebars to be installed locally and perform additional functions on documents and web sites. These widgets intrinsically provide a user interface, and would make the whole doctoring process completely client-side, thus completely avoiding the difficulties related to proxies.

This paper introduces and describes XLinkZilla, a linkbase application for XLink-based external links. XLinkZilla is composed of an extremely simple server-side CGI application for storing and delivering XLinks, and two installable add-ons for Mozilla browsers providing all the needed functionalities and interfaces for a complete external link management. Our main idea, in fact, is that browsers widgets can be exploited to create non-invasive, efficient and easy to implement applications for XLink. So we propose such architecture which moves to the client the whole doctoring process, using a server-side support only for storing and delivering XLinks.

In section 2 we examine some related works, especially regarding external linkbases and XLink applications. In section 3 we present a minimal set of features that an XLink application needs to perform when coupled with a standard HTML browser, and the problems that may occur with this. In section 4 we discuss pros and cons of the proxy architecture in providing such features, based on the experience we have had with our own XLinkProxy [4]. Section 5 introduces and discusses XLinkZilla, our new research prototype based on client-side add-ons.

2. RELATED WORKS

In the literature a lot of different works exploited the concept of external linkbases to enhance their linking mechanisms. Microcosm [6] is an open hypermedia system enabling users to browse through and query collections of linked multimedia information. From the same research came out Distributed Link Service [2] a modular link service for the World Wide Web, subsequently, enhanced with alternative transport mechanisms for links [10] and with automatic discovery mechanisms [3]. Hyper Wave [17] enables a more structured organization and visualization of Web content and links, even if there is no direct way to create multi-ended links. The interface is transparent to the user and the presentation of links is in traditional HTML syntax. Aquarelle [20] introduced an application to automatically and centrally verify the correctness of links and MMM [18] addressed the synchronization of multiple views of the same document though external links. The Open Hypermedia Protocol [8] (evolved in the Fundamental Open Hypermedia Protocol [19]) proposed models to express links independently of their storage location.

In the case of the WWW, the proposal that should change things for the better is certainly XLink [11], the linking language for XML that, in addition to simple links, introduces the concept of extended links. An extended link is an XML fragment whose elements are labeled with a specific linking role (through an attribute belonging to a specific namespace): the resources are anchors local to the document where the link resides, the locators are anchors on different documents and the arcs are navigable directions in the link. Thus, external, multi-destination and bidirectional links can be easily expressed by XLink and XPointer[12], the addressing language it is based on.

Many different implementations of XLink exist. Both the Amaya and Mozilla browsers claim to support simple XLinks (and a limited subset of XPointers). X2X [26] is a tool to easily create, manage and store external links that are re-inserted into documents on the fly. X2X does not use XPointer but exploits proprietary functions to allow sub-node links in the anchors. Fujitsu XLink

Processor [13] provides the same specific functionalities as a generic module to plug in more complex systems. XLinkit [27], on the other hand, is a complete application that makes use of XLink to provide a rule-based verification mechanism for link consistency in a document base. Related to XLinkit, XTooX [29] is a free XLink processor that turns out-of-line links into inline links. It takes a linkbase as its input and puts the links into the referenced documents.

Given that current Web browsers only support HTML or only have a limited support for XML technologies such as XLink, several proposals are based on intermediary and transparent application that turn external links into a set of internal links, understandable by the common browsers, and merge them with the original document without altering it. For instance, Goate [14] embodies this solution through a HTTP proxy. Goate does not implement any particular linking language, but instead provides a translation service between various “high-level” languages (the authors have used this term to indicate languages, such as XLink, allowing external, multi-ended and bidirectional links) and HTML.

Finally, Xspect [5] is an implementation of XLink based on standard XSLT and Javascript that provides navigational hypermedia functionalities, a SVG interface to guided tours and an authoring environment for annotations. Xspect uses XLink, but can convert to and from other hypermedia formats, such as OHIF [16] and FOHM. The attention is focused on the browsing and displaying of the linkbases rather than their creation or the surfing monitoring, provided for instance by Goate, XLinkProxy and XLinkZilla.

3. USING XLINKS IN WEB BROWSERS

Current Web browsers only support HTML links or only have a limited support for XLink links (e.g., only for inline simple XLINKs). Any implementation supporting more than the simplest features of XLink on the standard browser is therefore necessarily based on the conversion of external sophisticated links into internal, simple, unidirectional ones, most probably through the intervention of an intermediary and transparent application that convert external data and inserts it into the main document. Obviously the new simple links, understandable by the plain browser, need to be merged within the main document without altering its original content.

Independently of the implementation choices and the whole architecture of the system, it is possible to give a high-level description of the features that any typical XLink application for the plain browser must provide currently. These features can be divided in two classes: navigation-related features and creation-related features.

3.1. Navigation-related features

Assuming that the links already exist and are stored in a well-known external linkbase, navigation-related features include the following services: monitoring the navigation, finding relevant XLINKs, transforming XLINKs into plain HTML links, adding links to the document, providing support for the multi-destination ones.

3.1.1. Monitoring the navigation activities

Pages enriched by external links are composed of two independent parts, the original document on its origin web server and the external links to be added on-the-fly. So the first of the navigation-related features necessarily consists of finding out that a new URL request has been performed by the browser.

3.1.2. Finding all relevant XLINKs from the active linkbases

The appropriate linkbases are selected among all the available ones. This selection can be automatically made by the system based on user’s authentication or explicitly by the user who

identifies the linkbases he/she wants to be consulted. Since all XLink linkbases are actually XML documents, this corresponds to performing a query on an XML database.

3.1.3.Transforming XLink information into HTML

The XLink information drawn out of the linkbases is not immediately usable, given the simplicity of the HTML linking model. Therefore, in order to use current generation browsers, intensional references and multiple anchors have to be resolved into several separate instances of plain HTML anchors. After having detailed each instance of each multiple anchor, other problems need to be solved regarding counting, overlapping anchors and nested links. A detailed discussion about these issues can be found in [4].

3.1.4.Adding the links to the document

Once these situations have been considered and the problems have been resolved, it is time to actually add the elements to the documents. Note that links with multiple anchors or destinations are already split in different items and counting XPointers have been corrected. So the next step consists of inserting a new A element in the correct location within the document for each computed link, and making sure that it contains all the relevant information.

3.1.5.Support for selection of multiple destinations

XLinks leaves to the implementation any decision about the meaning and user interface widgets associated to the activation of a link with multiple destinations. A popular and well-motivated choice [25] is to add a pop-up menu to the link anchor where all destinations are listed and can be individually selected.

3.2. Creation-related features

Besides navigating the web using links that are not specified in the document itself, the real advantage of external links should be to allow each individual reader to add his/her own links and annotations to the pages being visited. The possibility to create new links on documents, and to share them with friends and colleagues, is what makes external linkbases more than a bizarre storage policy for link information. In order to manage user-created links, some basic services need to be provided: selecting the start- and end-points of the link, calculating the appropriate XPointer, setting options for the new link, saving information and managing existing links.

3.2.1.Selecting the start- and the end-points of the link

Whenever a user makes selections on the main text, a JavaScript object is created in modern browsers containing the selected string and the position of the string within the document. This is not true for old browsers (Netscape 4.x has a simpler DOM and does not reveal the position of the selection within the document) and, obviously, for browsers not supporting JavaScript. This data should be made available to the link creation widget.

3.2.2.Calculating the appropriate XPointer of the endpoints

The link creation widget must include a library implementing the conversion of addresses in a HTML DOM to an XPointer address. In creating links, the library need only to deal with the simplest XPointers, explicit address of strings within the document. More complex XPointers (e.g. intensional links and particularly generic links) are difficult to create with direct selections on the documents, and may need to be entered manually through some appropriate forms.

3.2.3.Setting options for the new link

The link creation interface widget needs to include forms with buttons and input fields to set further XLink options for the link. For instance, the user may want to specify the title of the

whole link and of every single anchor, or dictate behavior attributes such as “show” (indicating the desired presentation of the ending resource of the link) or “actuate” (indicating the desired timing of traversal from the starting-point to the ending-point).

3.2.4. Saving link information in the relevant linkbase

Once determined, the link data need to be sent to the external linkbase, expressed according to the XLink specifications. Here it will be saved within the linkbase, which could be as simple as a plain XML file on the file system of the server, or as sophisticated as a full fledged XML database. The linkbase is probably best placed on a third-party server accessed via HTTP.

3.2.5. Modifying and deleting existing links

Although modification and deletion functionalities to existing links could be provided with other and simpler means (e.g., via plain web forms accessible on the web server of the linkbase server), it may be felt elegant to provide some interface to select and access the link to be modified or deleted from within the link-creation widget itself.

4. ISSUES IN THE IMPLEMENTATION OF XLINK APPLICATIONS

Any application providing external linking services is composed of a pair of collaborating modules placed on different machines: a linkbase server, providing some kind of storage of link documents and searching of relevant links to a hypertext document, and a user interface providing access to the creation-related features described in section 3.2.

4.1. Linkbase servers

The monitoring of navigation activities and the nature of the HTTP impact on the possible architectures for linkbases, leaving only two candidates: proxies and plain CGI applications.

4.1.1. Proxy-based solutions

In [4] we presented XLinkProxy, a full-featured proxy-based external linkbase management system completely supporting XLink and a relevant part of the XPointer language. We believe that an analysis of the features of XLinkProxy could be useful in order to determine the advantages and disadvantages of the proxy-based architecture for external linking *per se*.

The basic working of XLinkProxy is as an HTTP proxy, i.e., as an intermediary between a browser and a Web server. XLinkProxy is a transparent proxy (i.e. a proxy that does not change in any way the response of the origin HTTP server) for all resources that are not HTML or XML documents. For documents that are either HTML or XML, XLinkProxy can add external links to it, retrieved from the linkbases locally stored. This architecture has many points in its favor:

- it relies on server-side technologies such as PHP, Perl, ASP or Java, which are sophisticated, easy to code and substantially more stable and reliable than client-side technologies;
- the computation intensive insertion of link data into the requested document is done on the proxy, with minimal workload left to the browser, and with large caching possibilities;
- navigation-related features are completely independent of the client technology, and can work with any browser, since the browser only receives a very plain HTML document (but this is not completely true if we consider multi-destination links, as discussed in section 3.1.5).

Proxy-based XLink applications provide proxy-side support for most of the navigation-related features, as described in section 3.1, but rely on the browser to provide support for the selection of multiple destinations and for all the creation-related features required by all XLink

applications. For these, in fact, browser-specific customization steps need to be considered. Furthermore, there are some drawbacks that need to be considered as well:

Network load and timeouts: the first and most evident drawback regards the time elapsed for each request. Interposing a proxy between the browser and the server doubles all network connections, since the browser will have to ask the proxy, which in turn will ask the server, for every resource; this clearly slows down the response time and doubles the risk of timeouts. Furthermore, the proxy will be consulted not only for HTML and XML documents, which is appropriate, but also for every other web resources, such as images, animations, JavaScript and CSS files, not to mention binaries and downloads of all sorts. It is true, on the other hand, that smart proxy implementations exist that reduce proxy-caused latency in HTTP response-request pairs. In particular, special configuration files (using PAC format [Pac96]) can be specified to have the browser select different proxies (or no proxy at all) depending on properties of the URL to be requested, such as the presence of the strings “.xml” or “.html” at the end of the address. But this solution can give problems, too, for three different reasons: first, PAC is not a standard but a proposal by Netscape silently accepted by some (but not all) other browsers; second, writing PAC files requires a certain technical skill; third, and more importantly, the use of PAC files is in conflict with the general WWW principle of URI opacity [23], according to which it is not correct to infer the nature of a resource by inspection of its URI.

Proxy chaining. A second problem with proxy-based solutions is that a proxy as an elective intermediary for content modification will not work in those situations where another proxy is being used (e.g., when the browser is behind a firewall), unless an explicit chain of proxies is set up and managed. This is not as easy as it seems: the problems associated with proxy chaining need to be solved at the infrastructure level, i.e., convincing the managers of the closest proxy to specify the second one in the connection chain.

Server-side parsing of bad HTML. A third problem with proxy-based solutions is that the proxy needs to parse the original document in order to find the correct positions where the links are to be added. When documents are in XHTML or reasonable HTML, this is hardly a problem. Sometimes, though, the HTML is sufficiently badly formed that no simple approach is possible. A number of tools exist to fix bad HTML documents using a server-side language. For instance, in XLinkProxy we systematically converted the original HTML into XHTML by means of HTML Tidy, the handy utility by Dave Raggett[22]. Unfortunately, not all existing HTML documents can be amended by HTML Tidy, so that in those situations we preferred not to add links, and to leave the document untouched. The same problem exists with many other server-side tools. On the other hand browsers are more sophisticated than server-side correctors and are able to deal with rather bad HTML code, and quite often the document has only been tested to display correctly, and not to be correct according to the HTML specifications. Indeed, exploiting a browser parser means to be able to handle all documents the browser is able to display.

4.1.2.Plain CGI applications

A linkbase server can also be implemented as a CGI application activated by plain HTTP requests. In this case, its functions are limited to creation and deletion of linkbases, and searching and adding of new links in selected linkbases. The most evident disadvantage of such applications for external link management is that they give **no** solution to the issue of monitoring the navigation activities of the browser (and, for that matter, to all navigation-related issues but 3.1.2: finding all relevant XLinks from the active linkbases): CGI applications simply receive HTTP requests for XLinks and give lists of XLinks as the response to those requests.

Moving to the client all navigation-related issues means, more than anything else, that access to XLink linkbases becomes a browser-dependent feature and thus, by definition, needs to be implemented from scratch for every browser (or, at least, for every operating system in which the

browser is run). This opens the possibility of ending up with different, non-interoperable implementations of the same functionalities, and of leaving in the cold the users of some browsers for which the implementation is either impossible or deemed to be too much of a hassle. In the next section we give a few glimpse of the issues connected to client-side support for the navigation-related features that CGI applications provide no support for.

4.2. User interfaces

As mentioned, no XLink application could be acceptable with at least some user interface mechanisms to provide support for the creation of links, and the selection of the destination in multiple destination links. Furthermore, client-side solutions such as XLinkZilla, as described in section 5, also move client side most of the navigation-related features discussed in section 3.1, that also have to be dealt with in the user interface part.

The complex interaction between the main navigation activities and the (most likely sporadic) link-creation activities of the user leave (to our knowledge) three possible architectures for the implementation of the user interface: frames, external applications and browser widgets.

4.2.1. Frames

In a frame-based implementation the main browser window is divided in two frames, respectively containing the XLink interface and the main HTML document being navigated to. Whenever the user navigates and selects links, the response is transformed into a frame request for which the requested page is one of the elements. This is the solution we selected for XLinkProxy, as discussed in [4]. One advantage of this solution is that more or less all browsers now support frames, and with some fiddling and tweaking of the JavaScript scripts it is possible to have the same implementation work on multiple browsers. But there are several drawbacks, too.

The first and most important drawback regards URL management. Suppose in the following that the linkbase server is available at the URL `http://www.linkbase.org/app.cgi` and the main web page the user navigated to is `http://www.site.com/index.html`. A simple implementation is to have the user access a URL like `http://www.linkbase.org/app.cgi?u=http://www.site.com/index.html`. The `app.cgi` application responds with a frameset document with the two individual frames pointing to the appropriate documents. The URL now is significantly different than the one the user expected to see in the address box of the browser, and significantly harder to read. Finally, all relative URLs of the main document had to have their base explicitly specified.

A more sophisticated solution implies fiddling with the actual URL of the main document. XLinkProxy, that implemented this policy, was rather radical: whenever the browser requested the URL `http://www.site.com/index.html`, the frameset document was first returned as the response, and in the meantime the proxy had some time to request the actual document from the origin server and transform it by adding the XLinks. Then the main document was stored internally by the proxy with a local URL that had already been included in the frameset document. The browser would then request it as a URL local to the proxy server and of course relative URLs to images and other documents will also need to be identified and modified. The advantage of this is that the user perceives the page as actually *being* `http://www.site.com/index.html` and that all URL fiddling happens behind the scene.

Unfortunately, this solution does not work well for documents that are framesets themselves, or where the relative URLs are created by computation (e.g. via a JavaScript script) rather than plainly available in the HTML. All in all, the XLinkProxy has taught us that the frame solution is unstable, not general, and rather complex, both client-side and proxy-side. Thus, although we used it for XLinkProxy, we now believe it to be inappropriate for a solid XLink application.

4.2.2.External applications

A different solution would be to create the XLink user interface as an independent application (i.e., a self-standing executable or a plug-in) to be installed on the client machine. Through inter-process communication APIs, this tool would be able to monitor the browser activities and send this information to the XLink server for further processing.

Two points are relevant here: whether the API for inter-process communication exists and how sophisticated it is (not really an issue for more modern browsers) and how would the modalities intrinsic to switching between applications interfere with the way the user worked with the tools: switching between different windows covering each other would most probably be perceived as to cumbersome for the XLink interface to be acceptable, so floating windows or other kind of solutions would be required in this case. All in all, external applications seem to provide no further advantage to browser widgets which will be discussed in section 4.2.4.

4.2.3.JavaScript additions

The effect of the floating window obtainable with the external application can also be obtained by having the proxy add some JavaScript scripts to the requested document. These scripts would then create all the required interface elements, including a floating window where the widget for the creation-related features would be displayed. In part, both XLinkProxy and XLinkZilla use this approach, although limited to the interface elements necessary to select among the destinations of a multiple destination link. In XLinkProxy, in fact, a JavaScript pop-up menu (i.e., the JavaScript Menu Component by Gary Smith [21]) is added by the proxy to every multiple link in the page and it is invoked via the onclick event of the A element. XLinkZilla actually adds some XUL code to the link, as will be discussed in section 5.

Although neither XLinkZilla nor XLinkProxy insert the full XLink interface into the requested document, it would in fact be easy to do so: before displaying the document in the main window it would be possible to attach the JavaScript code to the DOM of the main document, and have the `onLoad()` event of the main window activate the required interface elements. We have decided against this solution for two main reasons: first of all, it acts invasively on the internals of the document, and it may be cause problems when interacting with pre-existing JavaScript code in the main document; second, the solution relying on browser widgets is just as easy to provide, and much less invasive and uncertain.

4.2.4.Browser widgets

Recent versions of both major players (Microsoft IE and Mozilla Navigator) can be extended with new user defined interface objects such as toolbars, sidebars, menu items, etc. These items have full access and control on the document and the URL shown in the main browser window.

Returning to our discussion, this means that it is possible to exploit these extensions to give the users an interface to select anchors, send links to a server, update linkbases and so on. Residing directly in the browser, these objects do not have problems of communication, interaction and security as mentioned in the previous section; furthermore, not being within the document, they do not act invasively and do not interfere with existing scripts in the document itself. The only drawbacks we have found are of course the fact that they are specific to a single browser platform, and that they require an explicit installation for the user to be able to make use of its features. The second problem is not much relevant since the XLinkZilla installation is really simple (the user has only to drop an auto-install package on an open Mozilla window) but the first one needs to be further investigated: XLinkZilla runs on all Mozilla browsers and also on Netscape 7; furthermore few implementation efforts are required to let it run on Firefox and an Internet Explorer could be easily built, based on the same architectural principles.

In conclusion, we have found that browser widgets are easy to create, pluggable into the majority of the browsers, flexible to manage and lacking the limitations that other solutions have to face. For this reason, XLinkZilla has been implemented using this approach.

5. XLINKZILLA: A CLIENT-SIDE XLINK MODULE

XLinkZilla is a client-side application running on Mozilla that provide users an environment for creating and navigating on external links on normal Web pages. XLinkZilla is not based on a proxy, as XLinkProxy, but it relies on a browser addition that realizes the steps described in the section 3, getting through its task with mostly client-side code. XLinkZilla relies on a very simple CGI application (written in Perl) for server-side support (such as linkbases storage and querying).

Exploiting XUL [30], Mozilla allows customization in its interface elements, such as sidebar, toolbars, etc., that can be used for our purposes. Installation and configuration is easy, since one only has to drop an auto-install package on an open Mozilla window. Automatically a sidebar is added to the Mozilla interface, a couple of new commands are added to the right button and the application starts to monitor the navigation being performed on the main browser page.

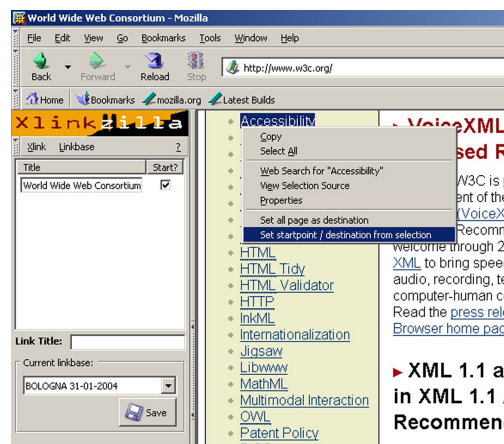


Fig. 1: the sidebar and menus of XLinkZilla

In figure 1 we show the main XLinkZilla sidebar, as well as the right-button command to create XLink locators out of the fragment currently selected in the main window.

5.1. Navigating with XLinkZilla

Navigation features are provided by means of a hidden module, called XMonitor, that watches over the navigation activities of the main page, and calls a registered handler in the sidebar whenever a new URL is loaded in the browser (see section 3.1.1). The XLinkZilla sidebar contacts (with independent HTTP requests) all currently active linkbases, querying for links originating from the URL being loaded in the main browser (see section 3.1.2). URLs are compared in their canonical form, as discussed in [31].

Many parallel and independent HTTP requests are performed: while the browser is loading the actual document from the origin server, the sidebar verifies with all XLinkZilla servers the existence of appropriate links. This architecture saves much of the time we would have spent had we used a proxy architecture: network connections are still doubled, in fact, but they are executed in parallel. Furthermore the sidebar does not interfere with the normal navigation activities (as a frame, for instance, would do) and gives no problem in the management of URLs. This architecture also allows for progressive display of links: the visualization of the main web page, in fact, is performed independently of the addition of external links. As soon as the links are

collected, in fact, they are added on-the-fly to the DOM of the document, even if it is already shown on the browser.

XLinkZilla does not require the document to be well formed according to the XML syntax, since relies on XPointerLib [28], an internal library of Mozilla that can solve XPointers within any web page. This is a great advantage over XLinkProxy, which could only manage a limited set of real web page, while XLinkZilla can deal with all the documents that can be displayed by Mozilla.

Each XLink is transformed into HTML fragments as discussed in section 3.1.3 by means of appropriate JavaScript modules in the XLinkZilla sidebar, and added to the DOM of the document (see section 3.1.4). External links are shown with a different style than plain links (a light yellow background for start points, a light blue-green background for destinations) to let the user understand the links are not originally from the document being shown. In order to manage links with multiple destinations (see section 2.1.5), a XUL menu is added to the contextual menu, invoked by the `onFocus()` event of the A element, as shown in fig. 2.



Fig. 2: the contextual menu for multiple destinations in XLinkZilla

5.2. Creating a new XLink with XLinkZilla

Creating a new link requires selecting the “New Link” command in the sidebar XLink menu. The main section of the sidebar is cleared, and all new locators will be added to the new link. The user selects a text fragment on the currently displayed page, and by opening the contextual menu (see figure 1) selects the command “set startpoint/destination from selection”. The sidebar, called by the appropriate event handler of XMonitor, transforms the information of the current selection into an XPointer, by means of the XPointerLib, and shows it in the main sidebar section. The user then may freely navigate on the web. Once a new interesting fragment is found, the selection is taken and transformed into an XPointer as before. As many locations can be selected and added to the current link (as start- or end-points). Once all necessary link information are inserted, the XLink is saved on the selected linkbase by clicking on the “Save” button.

An additional dialog box can be reached by a contextual menu item associated to the locator line, allowing the user to select advanced options for the XLink, as shown in fig. 3.

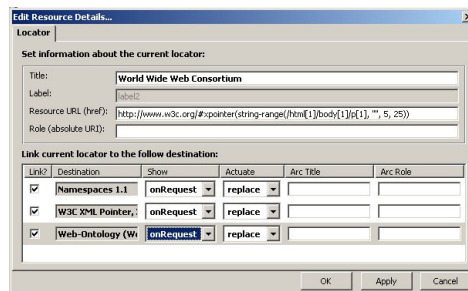


Fig. 3: the XLinkZilla dialog box for advanced link settings

XLinkZilla implements and supports all XLink features, including the “show” and “actuate” parameters. These control when the traversing of the locator should take place, and whether the traversal should substitute the currently displayed document, create a new window or embed the destination fragment within the source document.

6. CONCLUSIONS

XLink-based applications still have to build a market for themselves, and surely have a long way before they become mainstream and widely available. One of the problems they have to face is providing a powerful, yet easy to master interface for the kind of innovative linking functionalities that it aims to provide.

In this paper we have examined a number of services that any XLink linkbase application should aim to provide, and a number of possible implementation choices. The most important of these choices is surely the one between proxy and browser-based implementation. We have been through both types, with XLinkProxy first, and XLinkZilla now.

We have presented these two implementations discussing pros and cons of each architecture in providing the XLink features. In conclusion, we claim that a widget-based solution provides some important benefits that the widely used proxy-based solution lacks and allows implementers to create efficient, simple and easy to use applications. We tend to suggest browser-based architectures just for the ease of implementation, the simplicity of the interface mechanism, and the undeniable fact that at least some browser-based interface has to be supplied, so it may as well be worth implementing it all in the browser.

XLinkZilla our proposal for a browser-based XLink application, is a robust implementation of all XLink functionalities, and has been tested on a number of different operating systems. It can be downloaded and tested at the address <http://xpointer.web.cs.unibo.it/>.

The development of XLinkZilla is not finished yet. In the next months we will be adding an access control mechanism to the linkbases, and support for inline substitutions (as embedding of resource-type elements of the XLink link), as well as providing a linkbase management system for the Perl server-side application. But we are seriously convinced that these are just minor additions to a strong and already quite powerful tool.

REFERENCES

- [1] Bieber M., Vitali F., Ashman H., Balasubramanian V., Oinas-Kukkonen H. (1997) "Fourth Generation Hypertext: Some Missing Links for the World Wide Web", *International Journal of Human-Computer Studies*, 47, 1997, pp. 31-65.
- [2] Carr, L.A., De Roure, D.C., Hall, W., Hill, G.J., "The Distributed Link Service: A Tool for Publishers, Authors and Readers", in: *Proceedings of the Fourth International WWW Conference*, Boston, MA, The Web Journal 1(1), O'Reilly and Associates, 1995, pp. 647-656.
- [3] Carr, L.A., W. Hall, S. Hitchcock, "Link Services or Link Agents?", *Hypertext 98 Proceedings*, ACM Press, pp. 113-122
- [4] Ciancarini, P., Folli, F., Rossi, D. and Vitali F., "XLinkProxy: external linkbases with XLink", *Proceedings of the 2002 ACM symposium on Document Engineering*, ACM Press, 2002, pp. 57-65.
- [5] Christensen B. G., Hansen F. A., Bouvin N. O. "Xspect: bridging open hypermedia and XLink", *Proc. of the twelfth World Wide Web Conference*, May 20-24, 2003, Budapest, Hungary, pp.490-499.
- [6] Davis, H.C., Hall W., Heath I., Hill G., Wilkins R., "Towards an integrated information environment with open hypermedia systems", *Proceedings of the ECHT 92 Conference*, ACM Press, pp. 181-190.
- [7] Davis, H.C., "To Embed or Not to Embed...", *Communications of the ACM*, 38(8), 1995, pp. 108-189.

- [8] Davis H. C., Rizk A., Lewis A. "OHP: A draft proposal for a standard open hypermedia protocol". In *Proc. of the 2nd Workshop on Open Hypermedia Systems*, ACM Hypertext '96, Washington, pp. 16-20.
- [9] Davis, H.C., "Referential Integrity of Links in Open Hypermedia Systems", in: *Proceedings of ACM Hypertext '98*, Pittsburgh, PA, 1998, pp. 207-216.
- [10] De Roure, D.C., N.G. Walker, L.A. Carr, "Investigating Link Service Infrastructures", *Proceedings of the Hypertext 2000 Conference*, ACM Press, pp. 67-76
- [11] De Rose, S.J., Maler, E., Orchard, D., "XML Linking Language (XLink) Version 1.0", World Wide Web Consortium, Recommendation REC-xlink-20010627, 2001.
- [12] De Rose, S.J., Maler, E., Daniel, R., "XML Pointer Language (XPointer) Version 1.0", World Wide Web Consortium, Candidate Recommendation CR-xptr-20010911, 2001.
- [13] Fujitsu laboratory, "Fujitsu XLink Processor", 2003, <http://www.labs.fujitsu.com/en/freesoft/xlip/index.html>, last visited 4 February, 2005.
- [14] Martin, D. and Ashman, H. "Goate: XLink and beyond". *Proceedings of ACM Hypertext '02*. 2002.
- [15] Grønbaek K., Bouvin N.O., Sloth L., "Designing Dexter-based hypermedia services for the World Wide Web", *Hypertext 97 Proceedings*, ACM Press, pp. 146-156
- [16] Grønbaek K., Bouvin N.O., Sloth L., "Open hypermedia as user controlled meta data for the Web". *Computer Networks*, (33), 2000, pp.553-566
- [17] Maurer M. "HyperWave — The Next Generation Web Solution". Addison-Wesley, Reading, Massachusetts, 1996.
- [18] Ladd B., Capps M., Stotts D., "The World Wide Web: what cost simplicity?", *Hypertext 97 Conference proceedings*, ACM Press, pp. 210-211
- [19] Millard D. E., Moreau L., Davis H. C., Reich S. "FOHM: a Fundamental Open Hypertext Model for Investigating Interoperability between Hypertext Domains". In *Proc. of Hypertext 2000*, pp. 93-102.
- [20] Rizk A., Sutcliffe D., "Distributed Link Service in the Aquarelle project", In *Hypertext '97 Proceedings*, ACM Press, pp. 208-209
- [21] Smith G., "The JavaScript Menu Component, Creating Cross-Browser Dynamic HTML Menus", http://developer.netscape.com/viewsource/smith_menu/smith_menu.html, last visited 5 February 2005.
- [22] Raggett D., "HTML Tidy Library Project", <http://tidy.sourceforge.net/>, last visited 5 February 2005.
- [23] Jacobs, I., "Architecture of the World Wide Web", First Edition, W3C Working Draft 9 December 2003, <http://www.w3.org/TR/webarch/>.
- [24] Vitali F., Folli F., Tasso C. "Two implementations of XPointer", in *Hypertext 2002 Proceedings*.
- [25] Weinreich, H., Obendorf, H. and Lamersdorf, W. "The look of the link – Concepts for the user interface of extended hyperlinks". *Proceedings of ACM Hypertext 01*. pp. 200.
- [26] Empolis Hungary, "X2X, the XML XLink Engine", http://support.eqnet.hu/~empolis/products_x2x_en.html, 2002.
- [27] Systemwire, "XLinkit", 2004, <http://www.systemwire.com/xlinkit/>
- [28] Mozilla Organization, "XPointerLib", 2004, <http://xpointerlib.mozdev.org/>
- [29] Nentwich C., "XTooX", <http://www.xlinkit.com/xtoox/index.html>
- [30] Mozilla Organization, "XML User Interface Language (XUL)", 2003, <http://www.mozilla.org/projects/xul/>
- [31] Berners-Lee T., Fielding R., Masinter L., Uniform Resource Identifier (URI): Generic Syntax, Internet Draft, 16 February 2004, <http://www.ietf.org/internet-drafts/draft-fielding-uri-rfc2396bis-04>