Link Lens: An Enhanced Link User Interface for Web Broswers

Dominic Stanyer and Rob Procter
School of Computer Science
Division of Informatics
Edinburgh University
Scotland
Tel: +44 131 650 5165|5177
Fax: +44 131 667 7209
Email: dsjs|rnp@dcs.ed.ac.uk
URL: http://www.dcs.ed.ac.uk/home/dsjs|rnp

Abstract. A number of factors may influence Web users' choice of which links to follow. These include assumptions about document quality and anticipated retrieval times. The present generation of World Wide Web browsers, however, provide only minimal support to assist users in making informed decisions. Web browser 'link user interfaces' typically only display a document's Universal Resource Identifier (URL), whilst a simple binary colour change in the URL's anchor is used to indicate its activation history. The question then is, how do users deal with the problem of having to make such decisions when the information at hand is insufficient?

We have been conducting an investigation of how users make link selections. The results show users often are forced to fall back on heuristics and improvising strategies drawn from past experience. Based upon these results, we present a prototype of the 'link lens', an enhanced link designed to make such decisions easier and more productive for all users and help less experienced ones gain a better understanding of Web behaviour.

Introduction

The present generation of World Wide Web browsers provide users with too little information to help them decide whether to choose a link. Whilst there may be a number of factors that may influence this decision, the only information that the user can rely on being provided by a browser user interface is the document's URL, and an indication of its activation history. Sometimes, the user may be able to determine from contextual information whether to activate it or not (see Figure 1). For example, by choosing an appropriate anchor, the author of the document in which the link is embedded can help the user to determine quality and/or relevance [Nielsen 98]; similarly, a description of the document's size and type may help the user to assess how long it will take to download, and whether he or she will be able to view it locally. Since this kind of information is entirely discretionary, however, authors often do not go to the trouble of including it.

Where such explicit information is lacking, users may still be able to infer useful information about the document and its download overheads. For example, the URL may

contain implicit information such as the document's type and the physical location of its server. Such resourcefulness suggests that users would benefit if this kind of information were made more explicit, perhaps through the agency of an enhanced link. The selection and activation of a link is not the end of the user's dilemma, however. Download times are often long and unpredictable and may leave users uncertain about when — or if — it will be completed. Yet, browsers (see Figure 1) generally provide only limited feedback of download progress; even its placement seems to suggest that it is an afterthought.

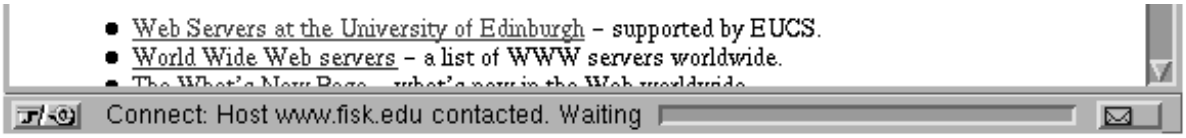

Figure 1: Netscape's 'link user interface'.

We have been conducting an investigation of how users make link selections and of their behaviour subsequent to link activation. The results show how users fall back on heuristics and improvising strategies drawn from past experience. One important issue that emerges is that users' heuristics are often flawed simply because the picture of Web behaviour that they get from the user interface is too abstract. Based upon these results, we discuss specific ways in which links might be enhanced to improve users' understanding of Web behaviour and provide more reliable information about its contents. The concept of enhanced or 'rich' links has been the subject of some discussion in the hypermedia literature (see e.g., [Bieber 97]). Our objective here is to investigate what these might entail from a user perspective.

We begin by discussing some critical determinants of World Wide Web usability. We then present a model of link selection decision-making behaviour and relate this to our empirical observations of users. We then review relevant usability principles and finally, we propose ways in which users' decisions could be better supported through the 'link lens', a prototype of an enhanced link based upon Magic Lens[1] concepts [Bier 94, Stone 94].

Usability and the World Wide Web

The Web raises many usability issues [GVU 98]. In this paper, however, we focus on two specific ones which are fundamental to understanding how to enhance the link. The first of these is content: how users determine that the document referenced by a link meets their requirements. The second is the temporal behaviour of the link: how long the selected document will take to download. These two factors together: quality of the document and quality of service are key determinants of users' selection and use of links. These factors interact, however, so their combined influence cannot be determined a priori. For example, a user may be prepared to tolerate slow download where the document quality is known to be high. Conversely, if time is pressing, then a lower quality — but more speedily downloaded — document may be more acceptable.

Neither document quality, nor its corollary, relevance, can be properly assessed except in the context of use. The only certain way that a user has of determining these properties is to download the document. However, following this strategy blindly may require the

[1]Magic Lens is a trade mark of the Xerox Corporation.

user to invest a lot of time and effort. This is simply impractical when the user has many documents from which to choose, such as when picking from a list of documents generated by a search engine. One solution is to make use of whatever contextual information is available. For example, the choice of anchor may provide a meaningful summary of the content of the document which it references. The wider context in which the anchor appears may also provide further information.

The problems of relying on contextual information partly explain the current wide interest in developing a standardised document markup language such as W3C's Extensible Markup Language (XML) [XML] so that content providers can add meta-data — information about information — to Web documents. One goal of adding meta-data is to facilitate the automatic processing of documents by Web applications such as search engines. We argue, however, that meta-data may also be an important resource for the design of an enhanced link. The first practical objective of our work is therefore to investigate how meta-data might be presented through an enhanced link.

As any Web user soon learns, the simple underlying abstraction of a seamless, distributed information space is rarely sustained in practice. Instead of an instant and deterministic response to link activation, the user typically must cope with unpredictable download times arising from variable latencies and network bandwidths. System response times in general are recognised as a major factor in determining the usability of an interactive system [Conn 95, Dix 87, Dix 94, Johnson 95b, Planas 88] and investigating how Web delays affect user behaviour is an area of growing interest [Guo 97, McManus 97, Johnson 97, Sears 97].

Studies have looked at how users' strategies change to cope with extreme delays or system failures [O'Donnell 95, McManus 97]. The unpredictable nature of Web download times means that users are unable to apply consistent coping strategies which leads to frustration and error. Attempts have been made to formalise this behaviour with an aim to helping developers design for download latencies [Johnson 95a, Johnson 97]. Other studies suggest that users' perceptions of document quality on the one hand, and their tolerance of download times on the other, are not independent. In particular, they indicate that download times may influence users' perception of document quality and relevance. For example, users' ratings of *how interesting* a document have been shown to decrease for longer download times [Ramsay 98, Sears 97]. Conversely, users' tolerance of download delays is influenced by their expectations of document quality [Johnson 97]

Despite the well-known detrimental effects of unpredictable delays on users, browser user interfaces all but ignore the problem. Whilst abstractions are useful tools for protecting users from unnecessary details of system behaviour, rigid adherence to a single abstraction may deny them the information they need to act effectively [Dourish 95, Kieras 90]. The link user interface, we argue, is a good example of this problem. The account it offers of the behaviour of the underlying system is too abstract and so fails to provide users with the resources for informed decisions and effective improvisation when potential problems are detected. Our second practical objective, therefore, is to investigate how to provide all users with a more informative account of download behaviour [Parker 97] which is helpful both for decision-making and for learning.

Terminology

A distributed information space such as the Web consists of documents connected by links. A document may represent different types of content, it may be referenced by links, and also contain links that reference other documents within the information space. A document can therefore be referenced by a number of links each situated in a different document.

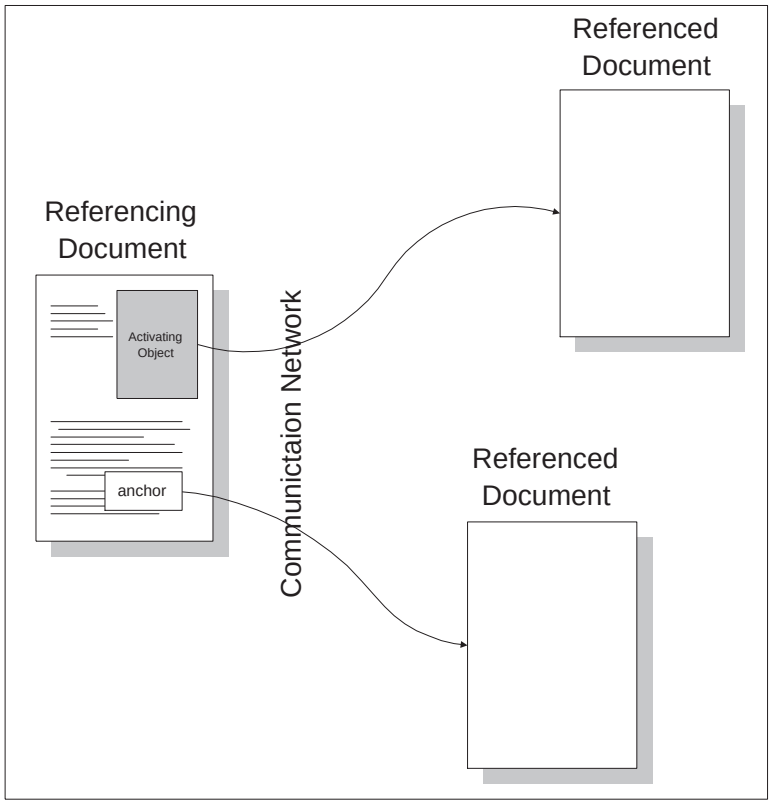

Figure 2: The link.

A link is embedded within a document and attached to part of latter's content by its *anchor* (see Figure 2) [Neilsen 90]. With respect to a particular link, we will define the object in which the link is embedded as the *referencing document*. Similarly, the object that the link points to will be defined as the *referenced document*.

As well as being identified with its anchor, referencing, and referenced documents, a link also has behavioural properties which become evident when it is activated and the referenced document is downloaded from the server to the user's client. This behaviour reflects the communication process which is implemented on the underlying network between the two locations.

Understanding Web User Behaviour

A series of studies of Web user behaviour were undertaken in order to learn more about how users made link activation decisions. A number of subjects were recruited with varying degrees of Web experience and expertise in computer use, ranging from novices to experts. Data about subject browsing behaviour was automatically logged by using

a proxy server connected to the WWW client. This parsed the HTML pages that the subject downloaded and recorded the URL, number of images, number of links, number of applets and the download time from the proxy request to the proxy receiving the data. Each browsing session was also recorded using video and audio.

At the end of the session, all the collected data — video, audio and proxy logs — was reviewed and a detailed transcript of the session was produced. A re-enactment protocol was then employed to elicit from subjects the reasoning behind their actions [Wright 91]. All the utterances of the subject and the investigator were recorded on audio cassette and later transcribed.

Analysis of this data leads us to propose a three phase model of user link activation decision-making behaviour:

Assessment of value The user must decide prior to activating a link whether the referenced document contains information that is useful with respect to the user's overall objective. Reaching a positive assessment will lead to the link being activated and the download of the referenced document will commence; a negative assessment will result in the link being ignored. In addition to the assessment of the referenced document's content, the user also makes an assessment of the likely download quality of service (QoS). The potential of a long wait (e.g., because of server or network congestion, low network bandwidth) for the referenced document may persuade a user to ignore the link, or at least ignore the link for the present until a more favourable time. Our subjects often tried to infer by the URL information about the probable length of the wait. Here, heuristics and rules of thumb prompted by previous experience often come into play. For example:

> "This page is on a University server and that means it should be quick because they have fastish servers."

The assessment of value process that the user undertakes prior to activating a link can therefore, be split into two parts:

1. an assessment of content, and
2. an assessment of download QoS, particularly the anticipated delay in downloading the document.

Download Once the link has been activated, the document download begins. During this phase the user is exposed to the underlying behaviour of the communication network. As we noted above, users of interactive systems are often very sensitive to delays, especially where these are of an unpredictable extent. The problem takes on even greater significance for the Web, however, because browser user interfaces allow users to interrupt and abandon the download at any time. In our studies, observations of users choosing this option were very common.

Johnson has employed micro-economic cost-utility models and the notion of marginal utility to analyse Web user responses to download delays [Johnson 97]. He demonstrated that users' tolerance of delays could be manipulated by increasing their expectations of document quality. Johnson's findings are evidence for the close relationship between the assessment of value phase and the download phase; by increasing the information available to the user before they activate the link, the

user is not only better informed about its nominal value, but is also better able to make an informed decision about whether the document is worth waiting for.

Some of our subjects revealed interesting beliefs about the extent to which their actions could influence the download time. For example, several subjects articulated the view that slow downloads could be speeded up by stopping and restarting them. When asked to explain this, one user drew upon a telephone analogy:

> "...[it] is like a crackerly telephone line, if you ring again the quality is often better."

Others provided more technically sophisticated rationales for their behaviour:

> "...when you stop and reload, it is possible to get a better route for the data, so it may be quicker."
>
> "...there is only a certain number of connections that a server can make, by reloading a page a new connection may have become free due to a time out or other users disconnecting."

Evaluation The final phase allows a user to evaluate the outcome of activating the link. If the activation was successful, the referenced document can be perused and its content can be compared with the expectations generated by the assessment of value phase. Similarly, the download performance can be evaluated. The evaluation process may stimulate the formation of heuristics that will serve to make sense of the experience, and which may then come into play in future cycles of interaction. A flavour of these heuristics is evident in explanations given by our subjects such as:

> "...all sites in Australia are slow."

> "...it is night time in America so all information will be transferred very fast."

We can now consider in more detail how the usability objectives defined earlier may be tackled. Our discussion will focus on two broad areas in turn: documents and links.

Documents

Our subjects sometimes used the URL to draw conclusions about the content of a document. For example, in the case of URL `www.dcs.cam.ac.uk/smith/publications.html`, one user was able to conclude that it referenced a list of publications (`publications.html`) by the author Smith (`smith`). Other parts of the URL may also yield useful pointers to document content, quality and relevance. For example, understanding how to decode `dcs.cam.ac` may lead a user to infer that the author's publications are in the field of computer science (`dcs`) and that they are probably academic papers (`ac.uk`). The user might even infer that, as a member of a prestigious university (`cam`), the quality of the papers is likely to be good.

Despite such user resourcefulness, it is evident that meta-data is the key element in improving users' capacity to make informed assessments of a document's quality and

relevance [Bieber 97]. Of course, the form this meta-data will take must be dependent upon the kind of document it describes. Whilst it is easy to specify that the meta-data for a scientific paper would include its authors, key words and abstract, that of many other document types is more difficult to anticipate. The issue that must be addressed then is how such meta-data may be effectively presented to the user. One possible approach is to use meta-data to provide some insights in to the document structure. A simple illustration of how this might be done is shown in Figure 3. Here we see the document content summarised by:

1. Author, title, keywords, abstract and section names appearing in the document. HTML already provides some support for document meta-data through the use of a number of pre-defined tags. These include author, title, and keywords. The new proposed document markup language XML goes further in that it allows the author to define new tags; in fact, every component of an XML document may be tagged in this way.

2. The distribution of keywords through the document body. This can provide a better indication of relevance than a simple count [Hearst 95].

3. A thumbnail outline of the document. The concept of genre is one possible way in which broader notions of meta-data may be addressed. The use of genre in the print media, for example, enables people to make quick and easy distinctions between the vast variety of publications available. To help the readership distinguish between these news genre, simple visual cues such as the tabloid and broad sheet format have been developed. The thumbnail is an attempt to provide similar visual cues for Web documents.

4. The numbers and sizes of different embedded media types and the number of links. This also may provide the user with genre cues. For example, a document containing many links may appear to be a richer resource for browsing.

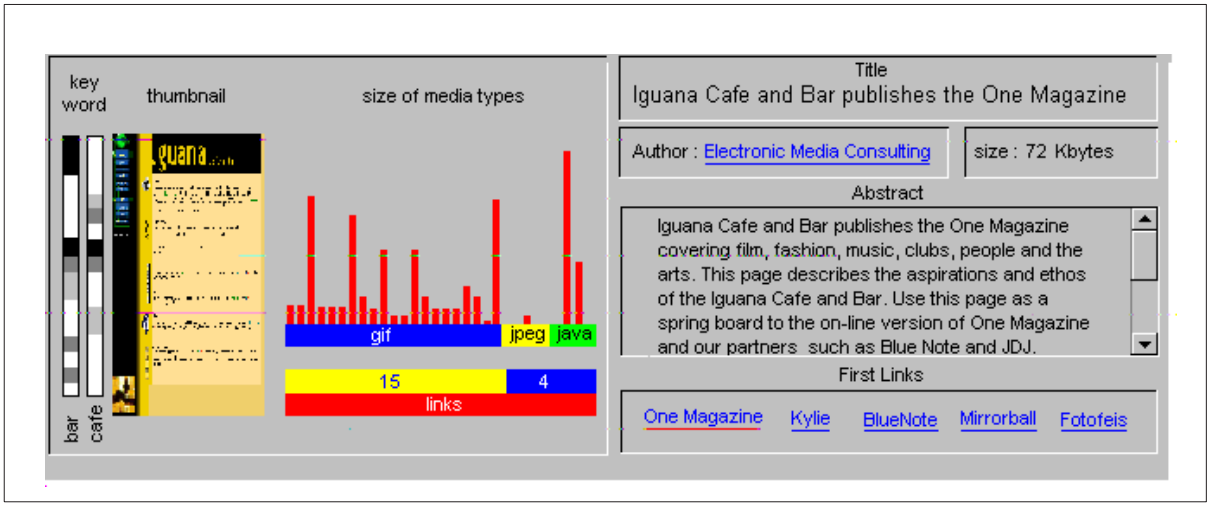

Figure 3: An example of document meta-data: structure and content summaries.

The Link

To provide the user with appropriate information about the link, it is necessary to identify an abstraction — or perhaps a layered set of abstractions — that can explain its behaviour [Dourish 95]. Our chosen link abstraction has two elements. The first is the *channel*, a communication resource which connects two (or more) computers together, and allows information to flow between them in either direction. If we restrict our notion of channel to the overall connection between the browser client and the document server, we have a concept that is synonymous with the communication associated with a download. The other element of our link abstraction is the *site*, a resource for despatching documents provided by the server. A channel and a site can now be associated with a link, and their properties can be presented to the user as affordances for understanding link behaviour.

The key property of a link is its QoS. Definitions of link QoS for multimedia applications cite several contributing factors [Guo 97]. For the Web user, however, the key QoS criteria is time. In simple terms, what the user wants to know before initiating a download is the delay before the download is completed. For the link, this can be defined as a function of site and channel latencies and channel bandwidth. Latencies reflect server and network congestion causing competition for these resources.

A full link time affordance, however, must also provide the user with information about the task in progress. The concept of time affordance is informed by users' subjective desire for closure, the knowledge that a system task has been successfully completed. Until this is evident, users will tend to monitor the task, assessing its progress and preparing contingency plans should they determine that it has failed. Lack of adequate time affordances may mislead users, for example, causing them to terminate the task erroneously.

According to Conn, a complete time affordance should meet eight requirements, representing information about distinct task properties [Conn 95]. We confine ourselves here to those where the link user interface is particularly deficient, and discuss ways in which an enhanced link might satisfy them. Some of the solutions we propose raise functional requirements which are addressed later.

Acceptance: Acknowledgement that the task has been correctly specified and is executable. Arguably, this is essentially a null requirement for browsers if the task 'specification' is defined syntactically as simply the act of pointing at any visible anchor. However, in so-called direct manipulation interfaces, syntactically acceptable action specifications have an implied semantic correctness [Schneidermann 86]. For example, a document's visibility is proof of its existence. In the case of an anchor, however, this relationship between visibility and existence may not hold.

Improved satisfying of acceptance requirements could be achieved by server and document status information being available when the user is in the act of specifying the download task. This is the most timely moment to provide this information and resolves the inconsistency between visibility and existence.

Scope: Measures of task scale, including predicted delay before completion. This requirement is not satisfied by the information provided through the link user interface. It may be partially satisfied by contextual information provided by the author of the referencing document, such as the referenced document's size.

Download scope is a function of the referenced document size and the link QoS. Satisfying scope requirements implies that these factors can be determined in ad-

vance of the download being initiated, In practice, however, only the document size can be accurately determined in this way — given suitable functional extensions.

Link QoS figures should take into account the effects of caching. Caches were introduced to exploit locality in users' document downloading behaviour. For example, there may be a cache on the user's (client) machine and a cache on a local proxy serving a number of users. If a document has been previously retrieved it could be available on these caches. A cached document will almost always download more quickly than the copy on its server.

Progress: Clear feedback as execution of the task proceeds, including its rate of progress. This requirement we consider is only partially satisfied. Although a progress indicator is maintained during the download, it is often difficult to relate the information it provides of the task state to its subsequent behaviour. One reason for this is that a document which may seem to the user to be a single entity may actually be composed of many embedded parts (e.g., images) each of which will trigger its own download process. The order of these embedded downloads is non-deterministic. In most browsers[2] the progress display area is time-multiplexed between these separate downloads, making it impossible for the user to maintain a coherent picture of the overall download progress [McManus 97].

An example of an improved link time affordance is shown in Figure 4. The map display shows the channel as a physical connection between the user's client and the document server. Its colour is used to represent predicted channel QoS (ranging from blue for poor to red for good). The graph shows site QoS information and the predicted download scope. Download progress is indicated by animation of the document media content display in Figure 3.

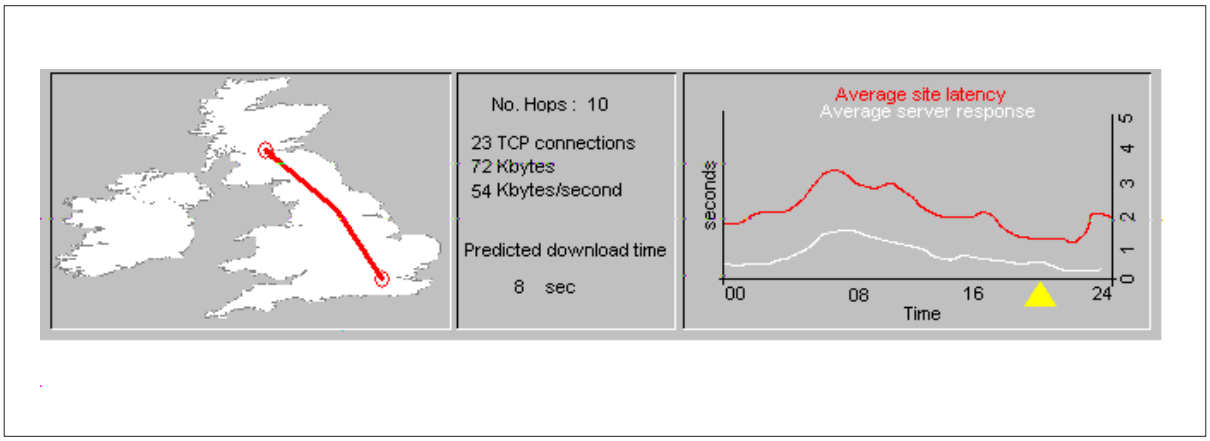

Figure 4: An example of link QoS data: channel and site QoS summaries.

Functional Implications

The possible improvements outlined above are more than cosmetic and raise additional functional requirements of the underlying system.

[2]Sun's Hotjava is an exception in providing progress indicators for embedded downloads.

If the link is to be able to provide the user with meta-data, this information must either be provided by the referencing document's author, or be downloaded from the referenced document's server. The former approach has been applied in some specialised hypermedia publications [Bieber 97]. However, our goal is to make meta-data available for a wider range of documents, and this is only achievable by following the download approach. To avoid this creating its own delays, two requirements must be met. First, the meta-data must be accessible to the client without downloading the whole document. Second, some kind of meta-data pre-loading or transclusion strategy is required so that it is available in a timely way to the user.

The HTTP protocol already supports client access of document parts through the HTTP-EQUIV directive. This allows authors to specify document tags which can be downloaded separately from the rest of the document and so HTTP already has the necessary hooks to meet the first requirement. The second requirement can be met by appropriate client side enhancements. There are a number of possible meta-data pre-load strategies. For example:

- All link meta-data is pre-loaded when the referencing document is opened. This is the least efficient strategy. Its drawbacks include the fact that the number of links may be large. The following three strategies adopt a progressively more conservative approach.
- Pre-loading is limited to those links that are within the currently visible area of the referencing document.
- Pre-loading is triggered when the user moves the link lens tool within a specified distance of the link anchor.
- Pre-loading is triggered when the user passes the link lens tool over the link anchor.

All the above pre-load strategies are assumed to operate automatically. It would be sensible, however, to allow user control over pre-loading, e.g., setting it on or off, and determining thresholds where appropriate. Yet another approach would be to give control of pre-loading to the referenced document's site. For example:

- Pre-loading is determined by the accumulated history of link activation by other users. This strategy would provide user with a form of collaborative filtering mechanism.

The accurate determination of download scope relies on knowing channel and site QoS. Obtaining real-time channel QoS data is problematic given current communication protocol limitations. The Hypertext Transfer Protocol (HTTP) is layered on top of the Transmission Control Protocol (TCP), which itself is layered on top of the Internet Protocol (IP). The latter does not support any QoS functions and so TCP cannot make any assumptions about factors such as latency or bandwidth. The only channel QoS related notion supported by IP is distance as measured by the number of 'hops' between the sender and receiver [Brebner 97]. Support for QoS is planned as part of an enhanced HTTP protocol [Guo 97], but in the meantime we must look to other solutions.

There are now several companies offering Web server QoS data on-line. Such data typically includes the number of hits per day, the average latency, and the numbers of broken links. However, we have decided to use a different approach which, though less

accurate, does not assume the availability of site QoS servers. The link lens utilises the performance of the meta-data pre-load cycle to calculate channel and site QoS for each link. In the limit case where link pre-loading has not been completed before the user wants to have access to it, the link can simply be assigned a low QoS, leaving the user to draw his or her own conclusions.

A Prototype Link Lens

The concept upon which the link lens is based is the so-called Magic Lens or movable filter [Bier 94, Stone 94]. Magic lens filters combine arbitrarily shaped visual field with an operator that changes the presentation of objects within the view field. In the link lens, the view field is bound to the cursor and is sensitive to any anchors that it is passed over. When applied in this way, its two operators, document and link, transform the anchor into views of document meta-data and link QoS. The prototype link lens is shown in Figure 5. The design was evolved from paper-based sketches of possible lens features. These were presented to six subjects who rated them according to perceived usefulness.

Their overall opinion was that the link lens would affect their browsing strategies. With the exception of the document abstract, there was no consensus about which parts of the link lens would be the most useful, but all subjects expressed a desire to be able to choose a subset of meta-data for viewing at one time. Depending on the task being undertaken, the subjects felt that some parts of the link lens could be beneficial to the usability of the system.

We are currently carrying out further testing of several versions of the link lens with groups of Web users. The objectives include evaluating the value to users of the various different kinds of link information. We are also investigating different ways of presenting the information, including enabling users to customise or adapt the link lens to match their needs, and more economical encodings based upon the appearance of the link in the referencing document. A particular emphasis of this work will be to explore how degrees of user experience affect their ability to make use of link information.

Conclusions and Further Work

The prototype link lens presents document meta-data and time affordances in the form of link QoS information. Its goal is to provide Web users with timely and relevant information prior to their activation of links and during subsequent downloads. Underlying our work is the belief that by making a more "truthful" picture of Web behaviour available, users will be able to formulate more appropriate explanations of its behaviour and make more informed decisions about its use.

Future work will focus on evaluation and exploration of improved techniques for link information presentation, including the value of layered abstractions of Web behaviour. The question of how much users need to know about Web behaviour is a priority for further investigation. The evidence from our user studies confirms that users will generate explanations of Web behaviour and will act upon them. Their reasoning and the analogies they draw upon, however, range from the simple to the sophisticated. So, it is probable that novice users will need different kinds of abstractions from those for more experienced users. Our aim in developing the link lens is to provide support for all kinds of users. An important part of our strategy for achieving this is to recognise that in the hands of a

less experienced user, the link lens may serve as a tool for learning about how the Web behaves.

As part of a broader programme of work for improving the value of the Web, we also plan to investigate the use of collaborative filtering — document usage and user ratings [Hill 95, Procter 97] — as additional document quality and relevance information. Finally, as HTTP protocol enhancements become available, we will examine how they may be applied in this context.

Finally, any mechanism for improving Web users' capacity to find relevant information must ultimately rely upon the willingness of content providers to recognise users' needs. In particular, the need for content providers to include meta-data is paramount. The link lens is designed so that the user will find it easy to distinguish between content providers who are conscientious in this regard and those who are not. In this way, we believe that putting tools like the link lens in the hands of users can provide an incentive for content providers to formulate and adhere to "good authoring practice".

References

[Bier 94] E. Bier, M. Stone, K. Fiskin, W. Buxton and T. Baudel, A Taxonomy of See-Through Tools. In *Proceedings of CHI'94*, ACM Press, 1994, pp 358-364.

[Bieber 97] M. Bieber, F. Vitali, H. Ashman, V. Balasubramanian and H. Oinas-Kukkonen, H. Fourth generation hypermedia: some missing links for the World Wide Web, *International Journal of Human-Computer Studies*, 47, 1997, pp 31-65.

[Brebner 97] G. Brebner, Computers in Communication, McGraw-Hill, 1997.

[Conn 95] A. Conn, Time Affordances: The Time Factor in Diagnostic Usability Heuristics. In *Proceedings of CHI'95*, ACM Press, 1995, pp 186-913.

[Dix 87] A. Dix, The Myth of the Infinitely Fast Machine. In D. Diaper and R. Winder (Eds.), *People and Computers III*, Cambridge University Press, 1987, pp 215-228.

[Dix 94] A. Dix, Seven Years on, the Myth Continues. Research Report RR9405, University of Huddersfield, 1994. `http://www.soc.staffs.ac.uk/~cmtajd/papers/myth95/7-years-on.html`

[O'Donnell 95] P. O'Donnell and S. Draper, How machine delays change user strategies. In M. Kirby, A. Dix and J. Finlay (Eds.), *People and Computers X: HCI'95*, Cambridge University Press, 1995.

[Dourish 95] P. Dourish, Accounting for System Behaviour: Representation, Reflection and Resourceful Action. In Proceedings of *The Third Decennial Conference on Computers in Context*, 1995, pp 147-156.

[GVU 98] Graphics, Visualisation and Usability Center, 8th WWW user survey, Georgia Tech, 1998. `http://www.gvu.gatech.edu/user_surveys/survey-1997-10/graphs/use/Problems_Using_the_Web.html`

[Guo 97] X. Guo and C. Pattinson, Quality of Service Requirements for Multimedia Communications. Workshop on Time and the Web, Staffordshire University, June, 1997.

[Hearst 95] M. Hearst, TileBars: Visualisation of Term Distribution Information in Full Text Information Access. In *Proceedings of CHI'95*, ACM Press, 1995, pp 59-66.

[Hill 95] W. Hill, L. Stead, M. Rosenstein and G. Furnas, Recommending and Evaluating Choices in a Virtual Community of Use. In *Proceedings of CHI'95*, ACM Press, 1995, pp 194-201.

[Johnson 95a] C. Johnson, Time and the Web: Representing Temporal Properties of Interaction with Distributed Systems. In M. Kirby, A. Dix and J. Finlay (Eds.), *People and Computers X: HCI'95*, Cambridge University Press, 1995, pp 39-50.

[Johnson 95b] C. Johnson, The Challenge of Time. In P. Palanque and R. Bastide (Eds.) The Design, Specification and Verification of Interactive Systems, Springer Verlag, 1995, pp 345-357.

[Johnson 97] C. Johnson, What's the Web Worth? The Impact of Retrieval Delays on the Value of Distributed Information. Workshop on Time and the Web, Staffordshire University, June, 1997.

[Kieras 90] D. Kieras and S. Bovair, The Role of a Mental Model in Learning to Operate a Device. In J. Preece and L. Keller (Eds.), Human-Computer Interaction, Prentice Hall, 1990, pp 205-221.

[McManus 97] B. McManus, Compensatory Actions for Time Delays. Workshop on Time and the Web, Staffordshire University, June, 1997.

[Neilsen 90] J. Nielsen, Hypertext and Hypermedia, Academic Press, 1990,

[Nielsen 98] J. Nielsen, Using Links to Help Users Predict Where They are Going. Jakob Nielsen's Alertbox, January, 1998. `http://www.useit.com/alertbox/980111.html`

[Parker 97] H. Parker, Temporal Usability and Disturbance Management in Interaction. Workshop on Time and the Web, Staffordshire University, June, 1997.

[Planas 88] M. Planas and W. Treurniet, The effects of feedback during delays in simulated teletext reception, *Behaviour and Information Technology*, 7(2), 1988, pp 183-191.

[Procter 97] R. Procter and A. McKinlay, Social Affordances and Implicit Ratings for Social Filtering on the Web. In L. Kovacs et al. (Eds.) Proceedings of the Fifth DELOS Workshop on Filtering and Collaborative Filtering, Budapest, Hungary, November 10th-12th, 1997. ERCIM Press.

[Ramsay 98] J. Ramsay, A. Barbesi and J. Preece, A psychological investigation of long retrieval times on the World Wide Web, *Interacting with Computers*, 10, 1998, pp 77-86.

[Schneidermann 86] B. Schneidermann, Designing the User Interface, Addison Wesley, 1986.

[Sears 97] A. Sears, J. Jacko and M. Borella, Internet Delay Effects: How Users Perceive Quality, Organization, and Ease of Use of Information. In *Proceedings of CHI'97*, ACM Press, 1997, pp 353-354.

[Stone 94] M. Stone, K. Fiskin and E. Bier, The Movable Filter as a User Interface Tool. In *Proceedings of CHI'94*, ACM Press, 1997, pp 306-312.

[Wright 91] P. Wright and A. Monk, The Use of Think-Aloud Evaluation methods in Design, *SIGCHI Bulletin*, 23(1), ACM Press, 1991, pp 55-57.

[XML] Extensible Markup Language. W3C. `http://www.w3.org/XML/`

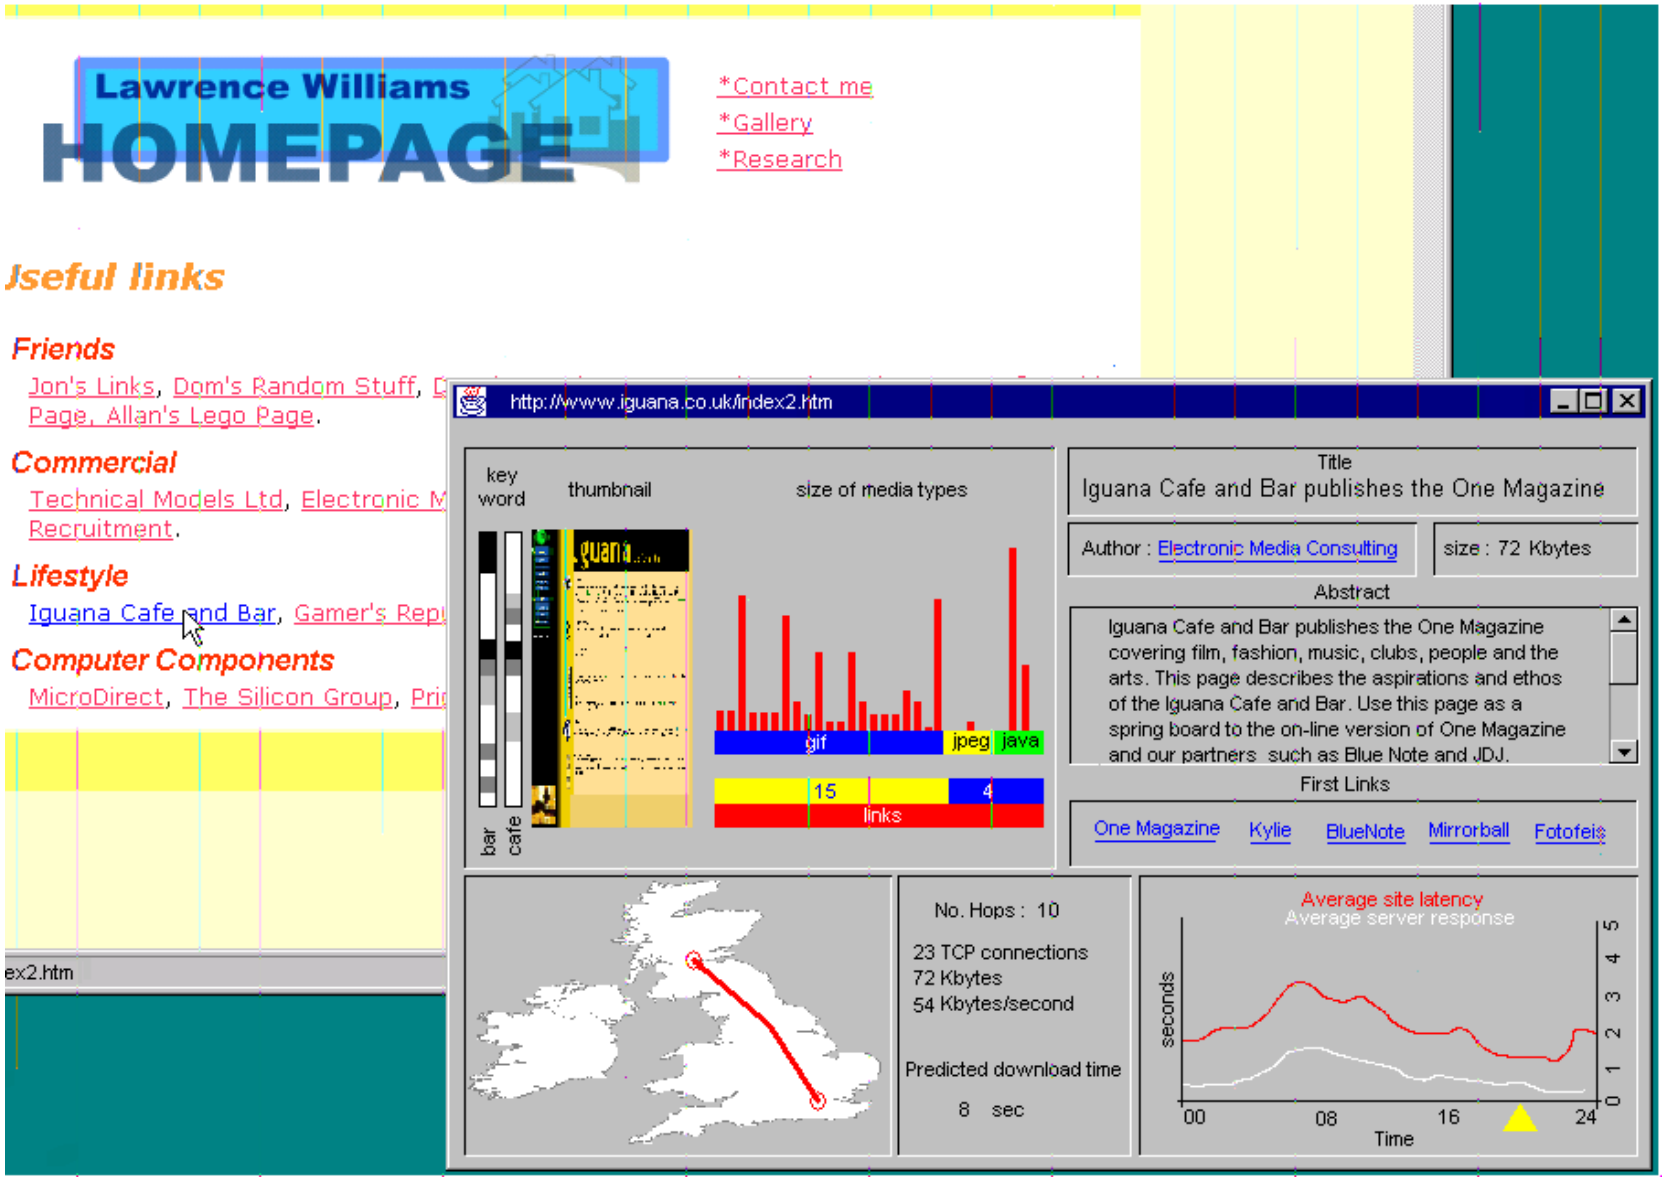

Figure 5: The prototype link lens.