The World Wide Web: What Cost Simplicity?

Brian C. Ladd
Department of Computer Science
University of North Carolina at Chapel Hill
Chapel Hill, NC 27599 USA
E-mail: ladd@cs.unc.edu

Michael V. Capps
MIT
Cambridge, MA 02139 USA
E-mail: capps@lcs.mit.edu

P. David Stotts
Department of Computer Science
University of North Carolina at Chapel Hill
Chapel Hill, NC 27599 USA
E-mail: stotts@cs.unc.edu

ABSTRACT
The ubiquity of the World Wide Web owes much to the simplicity of its graph model. Unfortunately that graph model omits powerful features found in traditional hypertext systems: concurrency and synchronization. These shortcomings are addressed in an extensible manner as part of the Multi-head, Multi-head, Multi-client Browsing Project; our research is focused on extending the Web Web through the use of the more powerful link semantics.

KEYWORDS: Hypertext, World Wide Web, concurrency, synchronization, link semantics, browsing, automata

1 INTRODUCTION

The link model of the World Wide Web is incredibly simple: a hyperlink anchor embedded in one document points, by name at another hypertext document. Unfortunately, this link model is too weak to express two concepts that have been used to great advantage in other, more traditional hypertext systems: concurrency and synchronization. These shortcomings are addressed in an extensible manner as part of the Multi-head, Multi-tail, Mulit-client Browsing Project (MMM) [2, 3, 4]. MMM provides a way to define different and more powerful link semantics as well as a framework for studying the use and usefulness of different link semantics on the Web.

2 THE PROBLEM

The World Wide Web owes much of its meteoric growth to the simplicity of its standards, protocols, and its hypertext model. The links semantics of the Web are simple: a hyperlink anchor is embedded in a hypertext document, enclosing some portion of the document and indicating by name the hypertext document (or other Web resource) at the other end of the link. Any Web browser rendering the document including the link will indicate to the user in some fashion that the portion of the document enclosed by the anchor is the tail of a link; should the user select the anchored text and indicate they wish to follow the link, the browser will fetch the named document [6].

Hypertext 97, Southampton UK

This model is too simple. While anchors can be nested, browsers resolve the activation of a hyperlink to the innermost enclosing anchor. Links in the Web are single-headed. This link model precludes the deliberate use of concurrent streams of browsing (though some browsers permit the reader to launch additional, concurrent browsing streams). The introduction of frames [5], permits a single document, the frameset, to divide up a single browsing window (in compliant browsers) into multiple, independent panes; while these are concurrent browsing streams, frames are not part of the Web standards and they do not permit synchronization between browsing streams.

Multiple concurrent browsing streams are desirable in many situations: the author wishes to present supporting documents in their entirety in parallel with a summation or analysis; an educator wishes to present two different views of the same material such as a graphic and a description of what the graphic represents; any situation where multiple pieces of information should be presented in parallel. The desirability of concurrency can be seen in the original description of hypertext systems by Bush and Nelson.

Splitting the browsing stream into multiple, concurrent browsing streams is desirable; once this is done, however, it is important to be able to synchronize them. Unfortunately, links in the Web are single-tailed. Keeping multiple streams synchronized requires that certain links be "hot" for activation only under certain conditions; the current Web link model presents all links as ready for activation since the only activation requirement it recognizes is the user's input.

Factoring out additional activation conditions has been considered [1] but has not been added to the Web model. Frames add the ability to specify a "target" frame in the current frameset for the presentation of the target object; this permits a link in one frame to update the content of another frame but since links are from single object to single object, there is no way, at the link level, to specify synchronization.

Synchronization is important in many situations: the author wishes to maintain a "seen-before" relation between different sections of a hypertext document; an educator wishes a student to pause and take an exam before proceeding with the course documents; any situation where progress (or lack of same) in one parallel browsing stream should influence progress in another stream. The desirability of synchronization is seen wherever concurrency is advocated; it is also used in traditional hypertext systems which treat the document as an automaton [7].

3 OUR SOLUTION

The overarching goal of the MMM Project is to add concurrency and synchronization to the Web link model. This involves extending the model at the language level and then implementing browsers to handle the extensions. Initially we extended the Hypertext Markup Language (HTML) with multi- head/multi-tail (MHMT) constructs, permitting authors to specify splitting and joining of parallel browsing streams. We also modified the then-popular NCSA Mosaic browser to handle the multi-links. A MHMT tag includes a list of all the target pages in individual anchor tags and a new "in-reference" or IREF tag for each source pages in addition to the current page. The initial semantics of our multi-links was that of a parallel finite automaton; this meant that following a multi-link opened all the target pages, links were only ready for activation when all the source pages were currently open, and that all the source pages were closed [4].

While the modified Mosaic client proved a useful test of concept prototype, the cost of porting our code modifications to each new version of Mosaic at a time when it was consistently losing market share to commercial browsers was too high. In light of this burden, MMM2 was developed as a proxy server consisting of modular components that parse extended dialects of HTML, rewriting them into standard HTML which will display appropriate link semantics [2]. The generated HTML uses a frameset to convince frame-aware browsers to display multiple browsing streams simultaneously. Specially encoded link values are inserted by the proxy server for multi-link constructs which are ready for activation; the proxy server keeps track of the state of the extended graph model and interprets encoded links back into the appropriate changes in the state of pages displayed by the browser. It is this state-consciousness which permits MMM to provide synchronization as well as concurrency.

The MMM2 proxy server has a semantic graph layer which permits developers to plug in and combine semantic engines; we have written and experimented with engine modules which can be composed to deliver the link semantics of various automata such as DFAs, PFAs, and Petri nets. Most interesting is a networking module which allows MMM2 proxy servers to communicate, allowing synchronization and concurrency among multiple users. This has allowed us to instance parallel finite automata and colored Petri nets as multiple-user models.

Work to date has focused on using extended graph models in the classroom; scalability is not an issue since it is assumed that the instructor's machine is sufficiently powerful to run the proxy server for the class. The current implementation of MMM uses a custom-built proxy server library; current research plans include reimplementation using the Jeeves servelet model and possibly browser plug-in technology.

4 CONCLUSION

The simplicity of the Web hypertext model was its greatest boon; those simple standards now stand in the way of certain interesting interaction with Web content. The Multi-head, Multi-tail, Multi-client Browsing Project is developing a way to extend the Web through the use of the more powerful link semantics. This fusion between the Web model and traditional hypertext models provides interesting avenues of research for both communities.

REFERENCES

1. Burchard, "Links and Interactivity" message to HTML-WG mailing list, 16 May 1995.

2. Capps, Ladd, Stotts. "Enhanced Graph Models in the Web: Multi-client, Multi-head, Multi-tail Browsing," *Proceedings of the 5th Annual World Wide Web Conference*, Paris, France, May 1996, pp. 1105-1112.

3. Capps, Ladd, Stotts, Nyland. "Educational Applications of Multi-Client Synchronization through Improved Web Graph Semantics," *Distance Learning Workshop, WETICE '96*, Stanford, California, June 1996.

4. Ladd, Capps, Stotts, Furuta. "Multi-Head Multi-Tail Mosaic: Adding Parallel Automata Semantics to the Web," *Proceedings of the 4th Annual World Wide Web Conference*, Boston, MA, December 1995, pp. 433-440.

5. Netscape Communications Corporation, "FRAMES: AN INTRODUCTION," December 1996, http://cgi.netscape.com/assist/net_sites/frames.html

6. Raggett, "HTML 3.2 Reference Specification," January 1996, http://www.w3.org/pub/WWW/TR/REC-html32.html

7. Stotts and Furuta, "Petri Net Based Hypertext: Document Structure with Browsing Semantics," *ACM Trans. on Information Systems*, vol. 7, no. 1, January 1989, pp. 3-29.