

Annotating the World-Wide Web

Sarah M. Luebke, Hilary A. Mason, and Samuel A. Rebelsky

Department of Mathematics and Computer Science

Grinnell College

Grinnell, Iowa 50112

{luebke,masonh}@ac.grin.edu, rebelsky@grinnell.edu

Abstract: The World-Wide Web has the potential to significantly change the ways in which students interact with texts, making them more active participants in their reading. Unfortunately, while many see the Web as promoting interaction, most pages provide relatively passive modes of interaction, only permitting students to click on one of a series of links to select what to read next. In this paper, we present one step toward making pages more interactive: a system that permits students to *annotate* arbitrary Web pages with notes, glosses, questions and comments; to *share* those annotations with their colleagues; and to develop *discussions* through those annotations.

1. Introduction

The World-Wide Web (Berners-Lee et al. 1994). is revolutionizing the ways in which we share information. In particular, it is affecting the ways in which we teach and learn. Unfortunately, while the Web is often seen as making students more active learners, most students can only interact with pages passively, by reading and clicking links. Some educators suggest that this is no more active than flipping pages or looking through the table of contents or index in a book.

In fact, paper provides many opportunities for active interaction with the material being studied. On paper, students freely mark the pages in their texts, underlining words or phrases, jotting notes in the margins, summarizing key sections, and even folding down corners to remind themselves of particularly important points (see, e.g., (Marshall 1998)). While some of these aspects are included as core Web technologies (e.g., the lists of bookmarks provided by most browsers are similar to the folded corners of printed texts, and perhaps more usable), others are less available (e.g., it is difficult to jot notes electronically and keep them on an existing page).

In this paper, we describe a system we have developed to support *annotations* to pages on the World-Wide Web (WWW). Our system permits learners to add comments, notes, and glosses to existing Web pages, thereby helping students become more active learners. While our system does not provide all of the mechanisms suggested in (Marshall 1998) (e.g., we do not support highlighting), our system provides an appropriate starting point for developing richer systems.

We have developed a system that permits selected *annotation* of arbitrary Web pages. Our system permits students to annotate pages with public and private notes, glosses, questions, answers, and discussions. We expect that as students have increased ability to interact with Web pages, they will become more active learners.

There are other annotation tools (or Web-based systems that include annotations) available, such as ComMentor (Röscheisen 1995), CoNote (Davis 1996), the Group Annotation Transducer (GrAnT) (Schickler et al. 1996), HyperWave (Maurer 1996), Medium (Lapique and Regev 1998), and WebCT (Goldberg 1998). We distinguish our system from by supporting annotations of both local and remote pages; permitting annotations at author-defined and automatically-generated positions on the page; supporting annotations on changing pages; and providing multiple protection levels for annotations.

We developed this annotation system to serve multiple purposes. First, an annotation system would provide a way for a student to take personal notes. Given the increasing number of course pages available, and the increased use of computers in the classroom, it may be possible for students to take online notes on course pages during class. For example, (Rebelsky 1998) reports that in a computer-equipped classroom in which daily course outlines were available online during class, some students opened the daily outline in an HTML editor to add their own notes. Second, an annotation system is also an appropriate way to share notes. For example, one might post questions about an assignment or suggest followup readings. Third, students can use sequences of notes to provide a discussion (similar to a typical threaded bulletin-board, but with focus on a particular point on a Web page).

For an annotation system to be successful, it must not only accommodate the needs, preferences, and habits of its users, but also predict and shape them. A student or instructor new to hypermedia authoring and usage may understand the ability to add notes to a page, but may not think about other issues, such as *protection* (who should be able to see an annotation), selection of *annotation points* (where annotations should go in a document), or *appearance* and *positioning* of annotations. We report on these design issues and the user testing that guided our design choices in the subsequent sections.

2. Architecture

Among our primary goals was to design an annotation system that could be installed and used widely. This led to a number of core requirements, including

- a teacher should be able to install and use the annotation system without making modifications to the Web page server;
- students and teachers should be able to use the annotation system with current Web clients; and
- students should be able to annotated a wide variety of pages, including pages on different servers and pages that neither teachers nor students had direct control over.

These requirements led to an architecture based on a centralized annotator and an extended version of the core of Project Clio (Becker and McLaughlin 1998) (Becker et al. 1999). A single CGI script is used to coordinate all annotation activities. A page request is routed through the CGI script. The CGI script

1. retrieves the page from a server (or from the file system, if the page is stored locally to the CGI script),
2. adds *annotation points* (described in section 2.2),
3. identifies the annotator (including group permissions),
4. fetches annotations from an annotation server,
5. inserts the annotations or links to the annotations into the page,
6. updates the links on the page so that the CGI script is used for further page fetches, and
7. returns the modified page.

In a future version, we expect to be able to use a form of proxy server for the same purpose.

The retrieval of original page, identification, link update, and return are handled by Project Clio. The addition of annotation points, fetching of annotations, and insertion of annotations were added for this project.

Because the annotation system is based on a CGI script, it is relatively easy to install (some servers may need a flag set to permit use of CGI scripts). Since it does not require modification to the underlying pages, it permits annotation of pages that neither teachers nor students control. Since the determination and fetching of annotations is done at the server side, and not the client side, it requires no extensions to clients. However, the current mechanism for displaying annotations does require JavaScript, which is supported by both of the most popular browsers (Netscape Communicator 4.0 and Internet Explorer 4.0). The display mechanism is described in section 3.

A settings file can be used to customize some aspects of the CGI script. In particular, it can be set to accept requests from only certain domains or certain accounts. This permits instructors to limit access to annotations. (The protection system described in section 2.1 provides additional forms of access control.) The settings file can also influence which pages can and cannot be annotated. For example, some instructors might wish to limit annotations to their own pages, or to a particular set of pages they have already identified.

2.1. Protection

There are three trends for protection of annotations: annotations are either private (accessible only to the user), group/semi-private (restricted to a group of users), or public (accessible to anyone reading the page). Many systems permit only one type of annotation. For example, CoNote provides only group annotations. Similarly, systems treat the different types of annotations quite differently. For example, WebCT (Goldberg 1998) separates annotations, which are private, from discussions, which are semi-private, using different interfaces for the two.

We chose to provide all three “types” of annotations with a consistent interface. When a user creates a new annotation, (s)he may choose to make it private, group, or public. Because we could foresee different groups of people using the same page (e.g., faculty, teaching assistants, and students might all have discussions on a course page), group protection also includes a note as to which groups may view the annotation. When browsing annotations, readers can see the headers of all the annotations, but may only view the bodies of annotations for which they have appropriate permissions.

2.2. Annotation Points

An annotation system for the Web must accommodate the differences between the Web and paper. One particularly important difference is that while paper remains fixed (and any errata are separate documents), Web pages often change. This means that annotations cannot be conveniently tied to a particular piece of text. Where should annotations appear?

Some annotation systems, such as the original NCSA Mosaic annotations (NCSA 1997), place annotations at the page level, with one annotation per page. Such placement seems overly coarse given the more typically fine-grained annotations students typically make. Discussions with faculty and students suggested that there is need for more specific placement of annotations. For example, a question about a particular problem in a homework assignment best belongs near that problem. More importantly, an instructor interested in having students gloss a work (i.e., add commentary) will not be satisfied with annotations that are so coarse-grained.

A finer-grained ideal would be to permit users to select the text to be annotated. Such a solution is used in the Multivalent Document system’s markup system (Phelps and Wilensky 1997). Unfortunately, standard HTML browsers will not easily support this choice. One alternative, used in the Group Annotation Transducer (Schickler et al. 1996), is to allow users to specify the text that will be tagged. Unfortunately, entering the appropriate text requires more effort on the part of the user and also suggests difficulties if the page changes or if the same phrase appears multiple times. Both of these fine-grained solutions may lead to annotations at awkward, confusing, or otherwise inappropriate points on the page.

Another alternative is to allow authors to specify the annotation points in a document, as in CoNote (Davis 1996). The disadvantages of such an alternative is that it places all the burden on the author, and provides little freedom for the annotator. An advantage is that it permits annotations to remain through changes of the document (something particularly difficult in phrase-based annotations).

We chose a hybrid solution, in which annotation points can be added by the author, but are also generated by the system for “natural” points in the document. At present, these points are currently set to be section headings, which often designate a logical break from one thought to another. For typical documents, section headings are fine-grained enough to permit annotation at a reasonable level but coarse-grained enough that the document is not over-crowded with annotation points. We use heuristics based on approximate text matching to determine the

(re-)placement of annotation points in changed documents.

3. The User Interface

A similarly interesting problem was where to put the annotations themselves. A number of alternatives suggested themselves (many of which have been used in other systems). One might make links to the annotations and bring them up in the same window when clicked. Such a solution is easy to develop, but makes it more difficult for the novice Web user to view both text and annotations. One might make links to the annotations and bring them up in the same window. One might insert the annotation text into the same page. However, this can make the original text more difficult to read, as the annotation interrupts the flow of the text. One might use popup windows, in which the annotation appears over the annotated text when one moves the mouse over that text or clicks on it.

3.1. Surveying Potential Users

We decided that this was not an issue we could approach in the abstract, after the system was constructed. Hence, we surveyed potential users about the strengths and weaknesses of a variety of designs. We developed eight different prototype interfaces for an annotation system and surveyed students with a variety of computer background. The examples varied in appearance and technical design issues. One example displayed annotations in a separate pop-up window. Another brought the user to a new page which displayed the annotation and the corresponding paragraph (the annotated text). In a third design, the annotations were visible on the page, within the document. For another the annotations were in the margin. The colors and icons changed from example to example, while the text and annotations were constant. Since we worried about interrupting the flow of text, a few examples permitted the users to turn the annotation feature on or off. We did not test the interface for entering the system (in which students enter their account and password) which is taken from Project Clio (Becker and McLaughlin 1998) (Becker et al. 1999).

Most participants preferred a pop-up window to view and add annotations, since they did not want to disrupt the page they were on. Annotations directly on the Web page seemed to add clutter and make the page less readable. Also, the less obtrusive the annotation buttons were, the better. At the same time, many found the ability to hide and show the annotation buttons (or the annotations themselves) confusing.

Based on this feedback, we developed a hybrid of the eight used in the survey. We chose to use two small buttons to indicate the annotation points. One button is used to add annotations and appears at every annotation point in the document. The second button is used to indicate the presence of annotations at that particular point in the document, and is only used if there are annotations for that annotation point.

3.2. Interacting with the system

If one clicks on the **View Annotations** button, a separate window appears showing a collapsible tree of annotations corresponding to that annotation point. There is a tree of annotations because we permit readers to write followups to other annotations, and found a tree the most natural way to represent the many discussions that may come from one annotation point.

If there is only one annotation available for a particular annotation point, that annotation is used as the “alt text” for the annotation button. In modern browsers, such as Netscape Navigator 4.0, this provides a quick pop-up view of the annotation. If multiple annotations are available for a particular annotation point, the number of annotations is used as the “alt text”.

When the user selects the **Add Annotation** button or the **Followup** button, a new window appears. The annotator can enter a name, title, and the annotation. Annotators can also select permissions: self, selected groups, everyone on system, world. Finally, an annotator can choose to make an annotation anonymous (so that, for example, students can post comments without fear of appearing less confident or successful they wish to appear).

While the annotations do not preserve text formatting (e.g., whitespace), they can be written in HTML, the HyperText Markup Language (Raggett et al. 1998). This gives authors freedom to format their comment or question by adding emphasis, tables, lists, and other structure.

3.3. Followup Surveys

After the new interface was complete and the underlying system implemented (so that subjects could add their own annotations and not just view existing annotations), a second group of subjects was brought in to test the new interface. This group of subjects included some of the first group, but also some new subjects.

While participants were generally satisfied with the structure of the implemented annotation system (e.g., the pop-up windows), many identified some technical problems that would not be obvious in a prototype. For example, many pointed out that the windows should auto-reload, so they can see the changes they made immediately. It was also noted that a close button was needed in the annotation window. Finally, many subjects needed more ready access to a help system.

4. Future Work

We have developed and begun testing of a system that we expect will enhance the effectiveness of course pages by permitting students to add notes and questions to Web pages in and out of the classroom. Our initial user-testing suggests that this will be a usable system, but we do not yet know how much or how it will be used. Because the system depends on parts of Project Clio, we must wait until Fall 1999 for our initial course trials. Based on feedback from those trials, we expect to make extensions and modifications to the system.

We also expect to refine the interface for the “annotation tree” (the set of annotations for a particular point in the document). While we conducted regular testing of the placing and appearance of annotations relative to the page, we were able to do less testing of the appearance of this tree.

Finally, we intend to provide alternate “views” of the annotations so that, for example, a student who wishes to print a document with annotations may conveniently do so. We expect to test at least two alternates for this “all in one” view: one in which the annotations are placed near the text they annotate and one in which they are used as footnotes (with footnote numbers placed in the text).

References

Becker, R. and McLaughlin, K. (1998). Tracking student use of course webs. Presentation at the 1998 Consortium for Computing in Small Colleges Midwest Regional Conference (September 25-26, 1998, Spring Arbor, MI).

Becker, R., McLaughlin, K., and Rebelsky, S. (1999). Project Clio: Tools for tracking student use of course webs. In *Proceedings of the 1999 EdMedia World Conference on Educational Multimedia and Hypermedia*. Charlottesville, VA: Association for the Advancement of Computing in Education.

Berners-Lee, T., Calliau, R., Luotonen, A., Nielsen, H. F., and Secret, A. (1994). The World-Wide Web, *Communications of the ACM* 37, 8 (1994), pp. 76-82.

Davis, J. (1996). CoNote - small group annotation experiment.
<http://dri.cornell.edu/pub/davis/Annotation/annotation.html> (accessed 20 July 1998; last modified 10 June 1996).

Golberg, H. (1998). WebCT - World Wide Web course tools. Online document at <http://homebrew.cs.ubc.ca/webct/> accessed 19 March 1998).

Lapique, F. and Regev, G (1998). An experiment using document annotation in education. In H. Maurer and R. G. Olson (Eds.) *Proceedings of WebNet 98 - World Conference of the WWW, Internet & Intranet* (pp. 539-544). Charlottesville, VA: Association for the Advancement of Computing in Education.

Luebke, S. M. and Mason, H. A. (1998). An annotation system for the World-Wide Web. Presentation at the 1998 Consortium for Computing in Small Colleges Midwest Regional Conference (September 25-26, 1998, Spring Arbor, MI).

Marshall, C. C. (1998). Toward an ecology of hypertext annotation. In K. Grenbaek, E. Mylonas, and F. M. Shipman, III (Eds.), *Proceedings of the Ninth ACM Conference on Hypertext and Hypermedia*, June 20-24, 1998, Pittsburgh, PA, (pp. 40-49). New York, NY: The Association for Computing Machinery.

Maurer, H. (1996). *Hyper-G, Now HyperWave: The Next Generation Web Solution*. Reading, MA: Addison-Wesley.

NCSA - National Center for Supercomputing Applications (1997). Mosaic User's Guide: Annotations. <http://www.ncsa.uiuc.edu/SDG/Software/Mosaic/Docs/help-on-annotate-win.html> (accessed 22 October 1998; version of 18 June 1997; last modified 2 February 1998).

Phelps, T. A. and Wilensky, R. (1997). Multivalent annotations. *Proceedings of the First European Conference on Research and Advanced Technology for Digital Libraries* (1-3 September 1997, Pisa, Italy).

Raggett, D., Le Hors, A., and Jacobs, I., editors (1998). HTML 4.0 Specification (version REC-html40-19980424), World Wide Web Consortium. (accessed 22 October 1998, last modified 24 April 1998).

Rebelsky, S. (1996). "Evaluating and improving WWW-aided instruction", *Journal for Universal Computer Science*, 2 (12).
http://www.iicm.edu/jucs_2_12/evaluating_and_improving_www/html/paper.html
(accessed 28 July 1998).

Rebelsky, S. (1998). In-class use of course webs: A case study. In T. Ottmann and I. Tomek (Eds.), *Proceedings of the 10th EdMedia World Conference on Educational Multimedia and Hypermedia* (pp. 1115-1120). Charlottesville, VA: Association for the Advancement of Computing in Education.

Röscheisen M., Mogensen, C., and Winograd, T. (1995) Beyond browsing: Shared comments, soaps, trails, and on-line communities, *Computer Networks and ISDN Systems*, 27, 739-749.

Schickler, M. A., Mazer, M. S., and Brooks, C. (1996). Pan-browser support for annotations and other meta-information on the World Wide Web. In *Proceedings of the Fifth International World Wide Web Conference* (May 6-10, 1996, Paris, France),
http://www5conf.inria.fr/fich_html/papers/P15/Overview.html (accessed 22 October 1998; last modified 17 June 1996).

Acknowledgments

We particularly thank Raphen Becker and Kevin McLaughlin, whose work made our project possible and whose comments also proved quite valuable. We also thank our survey participants for giving us very helpful comments. This work was supported by Grinnell College, the Grinnell College Noyce Science Summer Research Fund, and the Robert N. Noyce Faculty Study Grant.