

Hypertext Paths and the World-Wide Web: Experiences with Walden's Paths

*Richard Furuta, Frank M. Shipman III, Catherine C. Marshall, Donald Brenner, and
Hao-wei Hsieh*

Center for the Study of Digital Libraries and Department of Computer Science
Texas A&M University
College Station, TX 77843-3112, USA

E-mail: {furuta, shipman, marshall, dbrenner,
haowei}@csdl.tamu.edu

Citation: Furuta, Richard, Shipman III, Frank M., Marshall, Catherine C., Brenner, Donald, and Hsieh, Hao-wei, "Hypertext Paths and the World-Wide Web: Experiences with Walden's Paths," *Hypertext '97: the Eighth ACM Conference on Hypertext*, Southampton, U.K., April 6-11 1997, pp. 167-176.

Copyright ©1997 by the Association for Computing Machinery, Inc. Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Publications Dept, ACM Inc., fax +1 (212) 869-0481, or permissions@acm.org.

This work is supported by the DARPA CAETI program. Approved for Public Release, Distribution Unlimited. This material is based in part on work supported by the Texas Advanced Research Program under Grant Number 99903-230. The project's World-Wide Web pages are located at <http://www.csdl.tamu.edu/walden/>.

Hypertext Paths and the World-Wide Web: Experiences with Walden's Paths*

Richard Furuta, Frank M. Shipman III, Catherine C. Marshall, Donald Brenner, and Hao-wei Hsieh
Center for the Study of Digital Libraries and Department of Computer Science

Texas A&M University

College Station, TX 77843-3112, USA

E-mail: {furuta, shipman, marshall, dbrenner, haowei}@csdl.tamu.edu

ABSTRACT

Walden's Paths applies the concept of hypertextual paths to the World-Wide Web. Walden's Paths is being developed for use in the K-12 school environment. The heterogeneity of the Web coupled with the desirability of supporting the teacher-student relationship make this an interesting and challenging project. We describe the Walden's Paths implementation, discuss the elements that affected its design and architecture, and report on our experiences with the system in use.

KEYWORDS: Walden's Paths, hypertext tours and paths, meta-structure, educational applications

1 INTRODUCTION

Hypertext paths, also called tours, provide a meta-structuring mechanism that allows the affiliation of elements from many different hypertexts. The association of paths with hypertext is a concept as old as hypertext itself. Bush's Memex [2] incorporated paths for two purposes: first, to provide a personal means for remembering and organizing found information, and second, to provide a means to communicate that information to friends and associates.

Paths appear in implemented systems beginning around 1988 in Trigg's Guided Tours [14], integrated into NoteCards [4], and in Zellweger's Scripted Paths [15, 16, 17], a stand-alone environment. Guided Tours, centered around a graphical representation, provided the means to associate sets of cards with each stop along the tour. Scripted Paths allowed stops in documents or applications supported by "compliant" programs specified with a programming-language-like specification language.

Both Guided Tours and Scripted Paths were specialized mechanisms built for their own customized environments. The

published accounts suggest that they were intended for use by a sophisticated audience of hypertext authors and readers. The documents or nodes included in a Guided Tour or a Scripted Path were considered to be under local control—created in the same environment and acknowledging the same sensibilities as the tour or path meta-structure document. Experiences with Guided Tours noted, however, that a tour author needed to contribute significant annotative material to pitch an existing hypertext to a new audience of peers, and that the Tours themselves tended to have a linear spine, a minimal subset of tour stops that were assumed to be necessary for explication and intelligibility [6]. These use characteristics foreshadow in a limited way what we see today in a more heterogeneous environment of documents, authors, and readers.

The World-Wide Web provides an entirely new context for hypertext application. Although the technology of the Web is largely derivative from earlier applications, the environment that it has created is novel. Unlike its predecessors, the Web is highly heterogeneous, both in readers but also in information provided. Readers range from highly-educated academicians to elementary school students. Ages range from retirees to pre-schoolers. People make material available on the Web for all conceivable reasons—to communicate, educate, persuade, promote, defraud, and delude, to name just a few of the possibilities.

Path-like mechanisms that support personal-space organization are commonplace in Web browsers. The popular Web browsers all retain bookmark lists, sometimes hierarchical in structure. Brown and Shillner have described a new browser called Deckscape [1] that represents path-like groups of Web pages as decks, with particular implementation support for collecting and browsing decks.

Less common, however, are Web applications that use paths to support interpersonal communication. Sharing of bookmark files is provided in some systems, but paths have additional communicative potential that is not being tapped.

In this paper we discuss Walden's Paths, a Web-based path implementation.¹ Walden's Paths are intended to be commu-

*This work is supported by the DARPA CAETI program. Approved for Public Release, Distribution Unlimited. This material is based in part on work supported by the Texas Advanced Research Program under Grant Number 99903-230. The project's World-Wide Web pages are located at <http://www.csdl.tamu.edu/walden/>.

¹The Walden's Paths name was inspired by Henry David Thoreau's classic book *Walden*, first published in 1854. In *Walden* Thoreau describes his life for the two years and two months in 1845 and 1846 he spent living in the woods a mile from the nearest neighbor in a self-built cabin at Walden Pond, located outside of Concord, Massachusetts. Although Thoreau sought the

nicated from path author to path readers. The basic metaphor is that of a network-based meta-document: a document whose elements are themselves documents. In general the author of the path is *not* the author of the supporting documents, a characteristic that raises interesting rhetorical and societal issues. Moreover because our current application environment is in the primary and secondary schools (i.e., in support of curriculum for younger students, most of whom are between 5 and 18 years old), we found additional needs for contextualization perhaps not as important in the peer-to-peer environments present in earlier systems.

The next section of this paper describes the characteristics of the target student population. Sections 3 and 4 describe the Walden's Paths implementation and the architectural considerations behind it. We then discuss our experiences with Walden's Paths in use in sections 5 and 6. These experiences have lead to further extensions and developments, summarized in section 7. Finally we conclude with more global observations in section 8.

2 STUDENTS AND THE WEB

How do students experience the World-Wide Web as it is today? We had the opportunity to observe four classes of sixth graders from a local middle school who were using Netscape's browser to explore the Web [8]. The classes took turns in a campus engineering computer lab at Texas A&M University for these sessions; each student who participated in a session sat at his or her own PC, so each was free to explore at will. And explore they did.

At the start of each session, all the Netscape browsers in the room displayed the same starting page, a page that had been explicitly designed by the instructors for this instructional setting. The page included "starting out" links to different kinds of on-line places: Museums and public institutions (e.g., the Franklin Institute Science Museum, the Smithsonian Institute, and the White House), collections of pictures (e.g., the entomological collection at University of Illinois), peer-related Web activities (e.g., Chico Jr. High goes on-line), and entertainment-oriented sites (e.g., the Cartoon Factory Animation Art Gallery).

Instruction at the beginning of the sessions covered the "how-tos"—how to navigate and how to search; the instructors encouraged active exploration beyond the initial set of links. The computers in this lab imposed some restrictions on what the students could access and run (for example, they could not download games or play audio files although some tried); but in general the browsers provided much unfettered access to popular Internet resources.

We observed a variety of phenomena that suggest some design elements and considerations of a path facility that will meet the needs of elementary and secondary students. In general, the students' experiences were positive. Many were pleased with what they found, and the apparent ease with which they

solitude of life at Walden Pond, he quickly discovered that his daily activities created paths in the woods, and he suspected that others adopted his paths, keeping them from fading back into the woods. It is this relationship between early trailblazers and later explorers that motivated our choice of the Walden's Paths name.

could navigation among pages and sites. Of course, a few were not engaged, and were more or less passive, clicking a few times, and deciding they were getting nowhere; a few "got stuck" on pages, and could not return to a familiar orienting place. Five phenomena stood out: the difficulties associated with getting started and reaching navigational dead ends; the sociability of Web use and the collaborative nature of exploration; the success of simple navigational modes over more sophisticated ones; the confusion in interpreting the signs and signals of network performance; and finally, the compelling quality of participation, rather than just interaction.

Starting out and getting stuck. Some students hesitated when they were confronted with the unfamiliar software and range of navigational choices; they found it difficult to set out on an unstructured exploration (undertaken for no apparent reason other than to explore). Students occasionally "got stuck" as well. That is, they reached an apparently nonsensical page which offered them no easy way back to the special page set up for the class. In most cases, this happened when students pressed the Home button, which brought them to the home page for the engineering lab that normally took place in the room, an unfamiliar landmark which did nothing to reorient them. A stuck student would usually remain stuck unless there was some sort of intervention.

Active collaboration. We were struck with the sociability of Web use; students showed each other what they'd found and how they got there. Many of the students chose to work together, which included showing each other paths to interesting Web sites (one particularly popular site—trophy pages for many of our young cybernauts—was MTV's Beavis and Butthead collection). When we listened to the students talking to each other, we heard things like, "How did you do that?" and "How did you get there?" Students helped each other navigate in a very literal sense: they would help another student retrace their own steps to get to a particular page on a Web site. Most did not use the shortcut of telling the other student the URL of the page of interest, although during the sessions that took place late in the day, it became clear that such a strategy was growing in popularity.

Navigational strategies. Most students—perhaps just because they were neophytes—appeared to prefer the standard hypertext navigation that the Web offers rather than using search engines, or the shortcut of typing in a URL. Students were shown how to search at the outset of the class, but only a few tried the search engines; still fewer tried them with search terms of their own, preferring to use instead the reliable keywords demonstrated by the instructors. This said, the classes near the end of the day were affected by the discoveries of the students in the earlier sessions. During a morning session, one student discovered MTV's Beavis and Butthead site. Soon other students were following his path there, and apparently by the end of the day, students had figured out how to pass the URL to one another to navigate there directly.

Interpreting net delays. Network delays (and occasional unpredictability in response) caused the students some confusion. Because most of them did not have much (if any) experience with networked computers, they were unable to make sense of this sort of behavior. Characteristics of net-

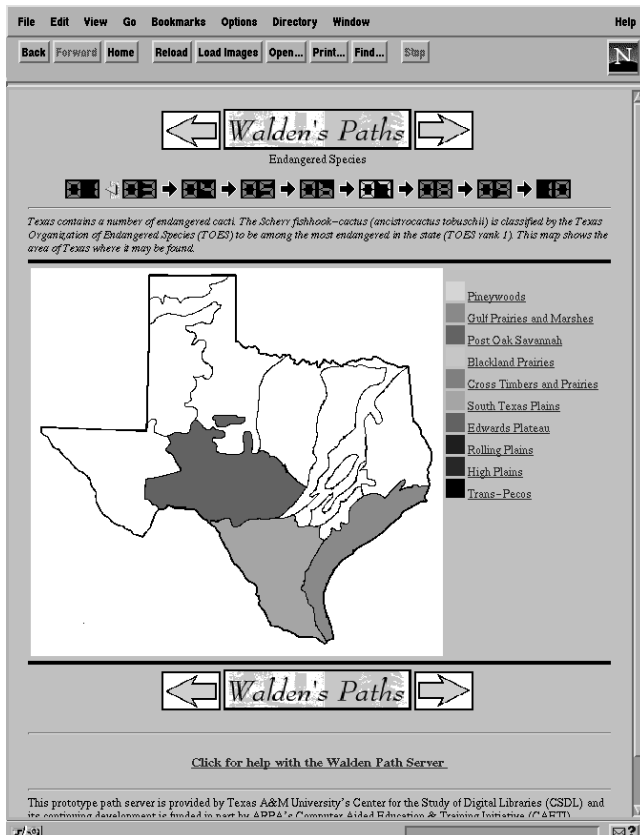


Figure 1: Page from path on endangered species in Texas

work and Web use—slow response, missing pages, and access denied messages—that are merely annoying to experienced computer-users becomes baffling to students (and others) with a less complete understanding of distributed systems.

Participation, not just interaction. The students seemed delighted when they were given the opportunity to participate, rather than just navigate. Of course, this was most evident in the active collaborative partnerships that formed, but it also manifested itself in other ways. One student was particularly pleased that she had located an on-line pen-pal service. With some guidance, she submitted her name and U.S. Mail address, and requested a pen-pal in China. She was careful to explain that she needed to send and receive paper mail, since she had no permanent email address. Her enthusiasm was considerable, and she was anxious to share her experience with her classmates.

Anticipating these phenomena, and addressing them through the design of the path mechanism, and the design of individual paths, will help primary and secondary students use the Web effectively. Certain principles emerge: first, make it easy to get started and to stay oriented; acknowledge the social and participatory nature of exploration; keep navigational options accessible; and finally, develop caching and path validation schemes that hide some of the complexity (and frustration) of navigating the Web. Next, we describe the Walden's Paths mechanism, which takes into account these principles.



Figure 2: Off-path pages include button to return to path

3 MECHANISM

While the Internet provides access to a wealth of new multimedia materials, especially through the extensive, encyclopedic materials available on the World-Wide Web, the students need to be provided a focus based on the curricular objectives.

Walden's Paths is designed to enable teachers to make use of materials available via the Internet by creating directed paths over the World-Wide Web [13]. Besides providing an ordering of pages, Walden's Paths allows teachers to provide additional context for the page through annotation. By providing text or other annotations in addition to the content of the page, the teacher may provide a rhetorical structure to the path as a whole, create transitions to fill in informational gaps between pages, and create emphasis to particular aspects of the materials.

Figure 1 shows a page from a path on endangered species in Texas. At the top of the page is a header including control buttons for moving back and forth in the path, numbers indicating the current position and length of the path, and the annotation/introductory text provided by the teacher. The introductory text on the pages of this path bridge together information on specific plant and animal species with information about wildlife preservation and ecological policies.

Students following a path may follow encountered links to freely examine the information space. A student's need for detailed explanation, alternative discussion, or details on re-

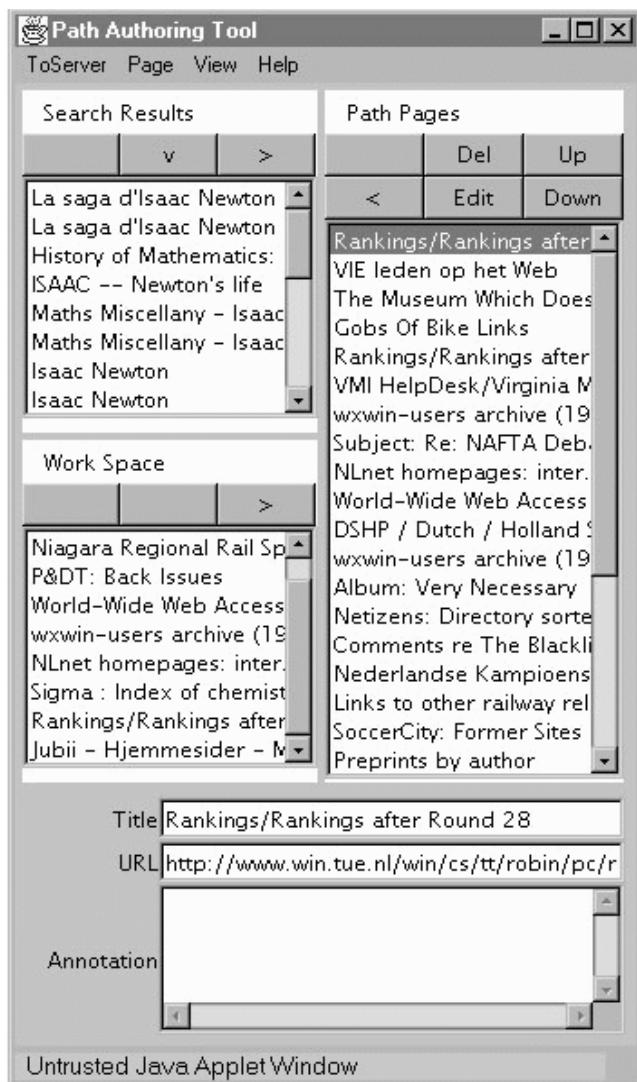


Figure 3: Prototype of the path authoring tool

lated topics are satisfied by his or her own desires to explore, to construct knowledge, to find information. Such explorations don't cause the student to lose the organizing context provided by the path as Walden's Paths augments off-path pages with controls that allow immediate return to the path. Figure 2 shows a page available to the student by browsing off from the endangered species page in Figure 1.

Keeping the focus of the project on providing teachers an easy way to collect materials for use in the classroom has motivated a KISS ("keep it simple, stupid") strategy. While earlier work on guided paths allowed the creation of generic directed graphs, Walden's Paths currently uses a strictly linear representation—a path consists of pages 1 through n in a strict sequence.

To support teachers authoring paths we have created a Path Authoring Tool. It supports the process of locating sites with relevant information, selecting individual pages from those sites, and ordering and annotating those pages to form the path. Figure 3 shows a prototype of this tool.

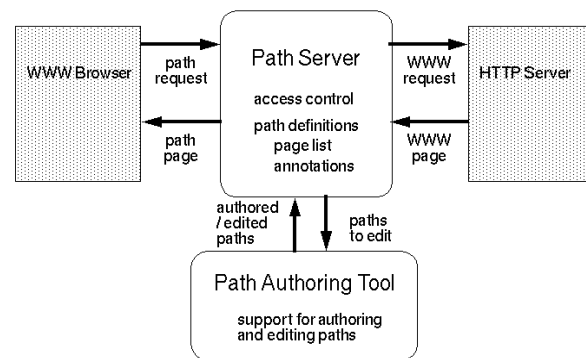


Figure 4: Diagram of the Walden's Paths architecture

Prior to the existence of the Path Authoring Tool, several of the paths were authored by collecting and organizing Web-based materials in VIKI [7], a spatial hypertext system designed to support information analysis. From these experiences functionality was added to VIKI so information could be arranged into spatial structures (e.g., lists or outlines) and exported as Walden's Paths path files [12].

4 ARCHITECTURE

Figure 4 shows the architecture for Walden's Paths. Walden's Paths consists of two main components: the Path Authoring Tool and the Path Server. As described above, the Path Authoring Tool enables the teacher to create, modify, validate, and reuse paths. When a path is saved by the teacher it is stored with the Path Server. The Path Server is the implementation means through which students navigate the paths.

In order to avoid requiring students or teachers to use specific software, the Path Server works as an intermediary between the students' Web browser (e.g., Netscape) and the http servers providing the original material being used in the paths. The Path Server is a Common Gateway Interface (CGI) that recognizes requests for path information. When it receives a request, the Path Server determines the path and page number of the request, or if the request is to go off of the path, and retrieves the needed information to dynamically generate the page.

Because schools frequently have limited connectivity and bandwidth, the Path Server can also cache pages of the path prior to their use in the classroom or cache pages as they are browsed by the class. Since a single server is used by all the students, the cache is shared rather than being per individual like the browser caches. Of course this assumes that the school has a local area network with one machine running the Path Server (currently this machine must be a Sun workstation.)

The Walden's Paths architectural design was influenced by both practical and pedagogical considerations. Our project team is relatively small in number—particularly when compared with the giant programming teams occupied in producing the currently-popular Web browsers (e.g., Netscape). We had no desire to attempt to track the rapid changes taking place in browser technology, no resources to attempt to support the

multiple hardware platforms that are commonplace in school environments, and no interest in fighting the frequently political battles necessary to get schools to adopt standard software configurations in their client machines. Consequently we decided to implement our system as a protocol filter—tapping into and modifying the standardized protocol stream that is exchanged between browser and server in a way that is transparent to both. As such we are buffered from the effects of updates to the browser, the part of the Web environment that is currently most subject to rapid change.

In related work, the Footsteps project [10] also provides a Web-based path mechanism and independently has developed some of the same implementation and architectural components as Walden's Paths, although they have not yet addressed issues of authoring, annotation, and genre. We expect that our experiences in use of Walden's Paths will be of interest in that project.

5 EXPERIENCE WITH USE OF WALDEN'S PATHS

The Path Server has been available to the public since March 1996, accessible via the World-Wide Web and supporting an increasingly-large collection of paths written by authors both within the project but also authors recruited from the teaching community. The first path contributed from outside of the project was written in April 1996 by a teacher who saw one of our early presentations. Additional paths have been contributed by teachers recruited from elementary and secondary school technology coordinators from schools around Texas; in particular six teachers wrote nine paths during a three-day period in June 1996 as part of an extended workshop. Since the Authoring Tool was not available at the time of this workshop, we defined an ASCII-based input representation and asked the teachers to use generally-available text editors to create paths in the representation. The teachers were uniformly successful in doing this, and primarily relied on cut-and-paste operations that mimic functions now provided automatically by the Path Authoring tool. Walden's Paths currently is being considered for deployment in several K–12 schools.

The teachers who have written paths in the system have been enthusiastic in their evaluations. Comments indicate that they view paths as an effective way to draw together materials from disparate sources. Indeed one teacher from the extended workshop was enthusiastic enough that on his return to his home school, he recruited the assistance of a graduate student and implemented his own subset version of Walden Paths. The primary requests by the teachers for enhancement have been for mechanisms to increase the level of interaction students have with paths to help maintain and measure student engagement with the material. Two approaches that we are now considering include providing opportunities for associating student-generated annotations with paths and extensions to our path specification to permit conditional branches based on student responses to questions.

One surprise for us has been the rapidity with which our technology “disappears”—becomes an element of the landscape rather than a focus of attention. We had our first hint that this was occurring shortly after we placed our public version on the Web. This initial release carried notifications that the

software and paths were prototypes and should be viewed as examples of the capabilities of the system rather than as finished products. One of the strongest comments we received at the feedback email address concerned not the technology but instead the *content* of one of the pages included on one of the paths. The commentator clearly had become engaged enough with the path material that distinction between the contents (i.e., the original material) and the container (i.e., the path system) had disappeared.

We saw similar effects with the teacher-authored paths. The extended workshop concluded with project presentations by the participants. After a brief, almost cursory, acknowledgment of the technology, most of the teachers focused on the content contained within the path. Once again, the material of importance was the contents rather than the container.

We find the message of these effects to be both reassuring but also cautionary. The effects are reassuring because they are precisely what we wish to attain—the provision of new layers of meta-structure that enhance rather than diminish access to the content that they place into new contexts. But these reactions also remind us of the changing nature of ownership and attribution for networked information; effects that have implications for the models we implement. We return to this point subsequently.

6 META-DOCUMENTS, DOCUMENTS, AND WALDEN'S PATHS

Paths are an interesting form of document (i.e., a meta-document): one in which the constituent parts are themselves documents, not the more primitive elements that make up most documents. The rhetorical purpose of a path is likely to be distinct from that of its components. Components frequently are part of a tightly interlinked network devoted to elucidation of a particular point of view about a selected set of topics. On the other hand the path is a unifying structure, directed towards contextualizing multiple points of view about a topic into an organizing framework.

Because of Walden's Paths' use in conveying curriculum from teacher to student, there is a strong need to be able to make explicit the path's rhetorical structure; since the path is a meta-document, the student often needs help in understanding why an element has been included. This characteristic differs from a bookmark facility, in which the more homogeneous nature of the environment allows implicit expression of structure. Furthermore because the rhetorical goals of the constituent documents are independent and distinct from that of the path, there is often a desire to adapt the underlying document to bring it into closer agreement—this is a question of interest not only from a technical perspective but also from a societal one. Finally, from a practical point of view the dynamic nature of the Web guarantees that the constituent documents will be evolving, mutating, and disappearing. Successful implementation of paths must maintain the integrity of the meta-document in the face of this change.

Paths that have been authored in Walden's Paths to date appear to have been created with one of three overlapping rhetorical structures: extended bookmark lists, extended tours, and curriculum-based expositions.

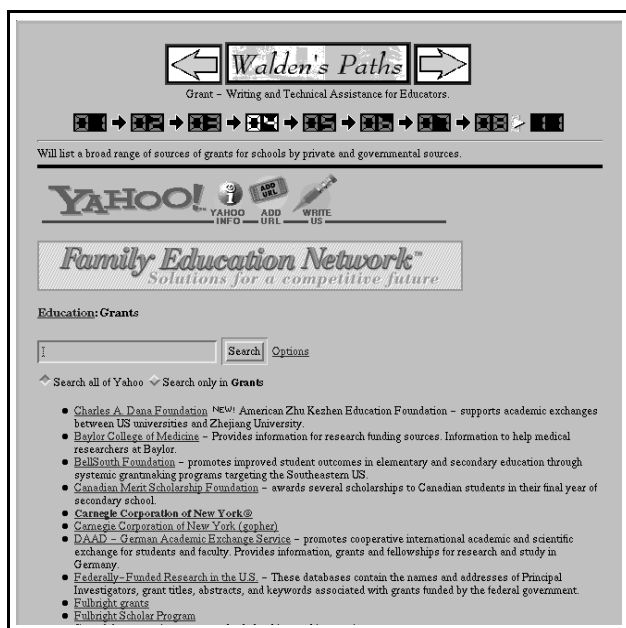


Figure 5: Path as extended bookmark list

Extended bookmark lists appear to be generated as an organizing mechanism—a translation of a “hot list” into a sequentially-traversable form. Figure 5 shows an entry taken from one such path (indeed in this path, the content is a “virtual” hot list—the dynamically-constructed results of a Yahoo search). In extended bookmark lists, annotations tend to be brief or absent. If present they often provide a short abstract of the page’s contents—almost as a memory-jogger to aid the author in remembering what is important about the page. Frequently the entries on the extended bookmark list are intended to serve as starting points for further exploration—multiple stops within a particular neighborhood are infrequent.

Extended tours resemble the narrative that one might get from an audio guide while walking through a museum. Here the role of the annotation is to contextualize the individual stops within the larger framework of the collection as a whole. Figure 6 shows an entry taken from this kind of path. In these paths multiple stops within a neighborhood are common as the extended tour seeks to point out the key locations explicitly in order to guarantee that they have all been seen.

The curriculum-based exposition both contextualizes the material it organizes but also provides an interpretation. The structures seen here have been more complex than those of the other forms. Annotations are used for curricular goals—for example questions may be posed to attempt to retain student attention (Figure 7). Introductory and interpretive material is often added as preface and delimiter for path segments. Figure 8 shows one such example that is *entirely* annotation—this stop along the path includes no reference to WWW page. Path authors often wish to modify distracting backgrounds or to “rework” inappropriately-focused source material. Figure 9 shows an example in which the size of the original page’s textual material has been minimized and figure 10 an example where the material included is illustration but not the original surrounding text.

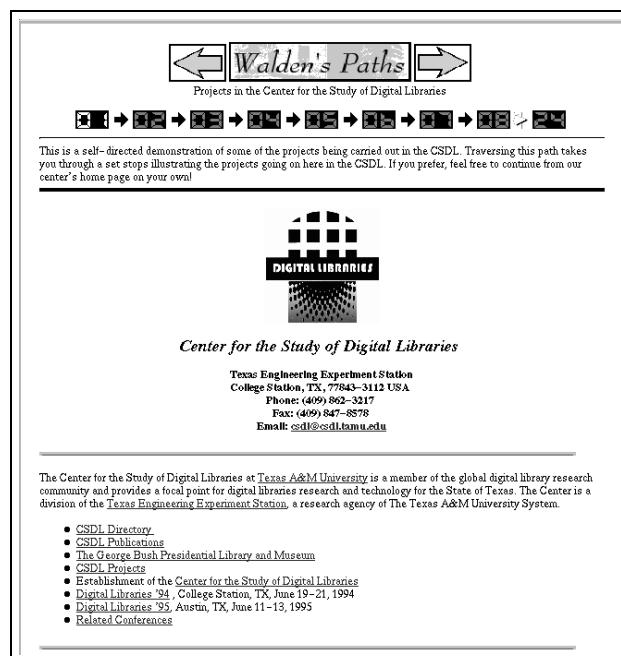


Figure 6: Path as extended tour

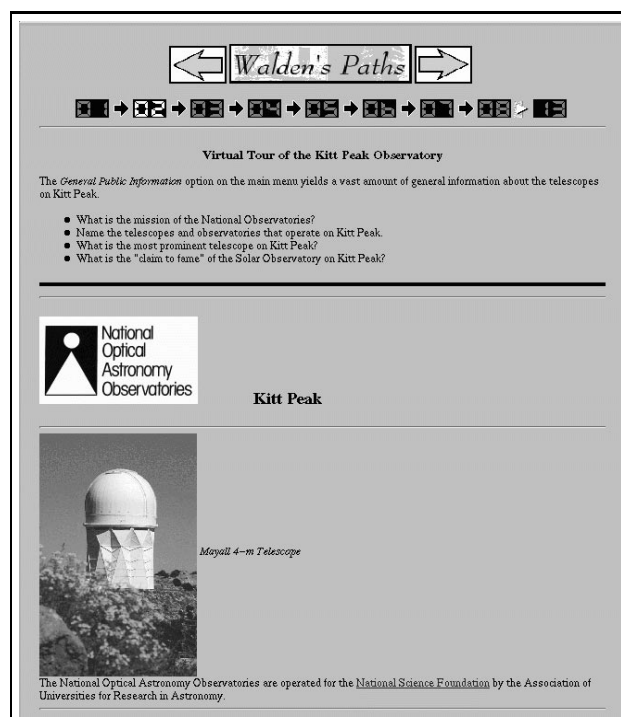


Figure 7: Path annotation in support of student engagement

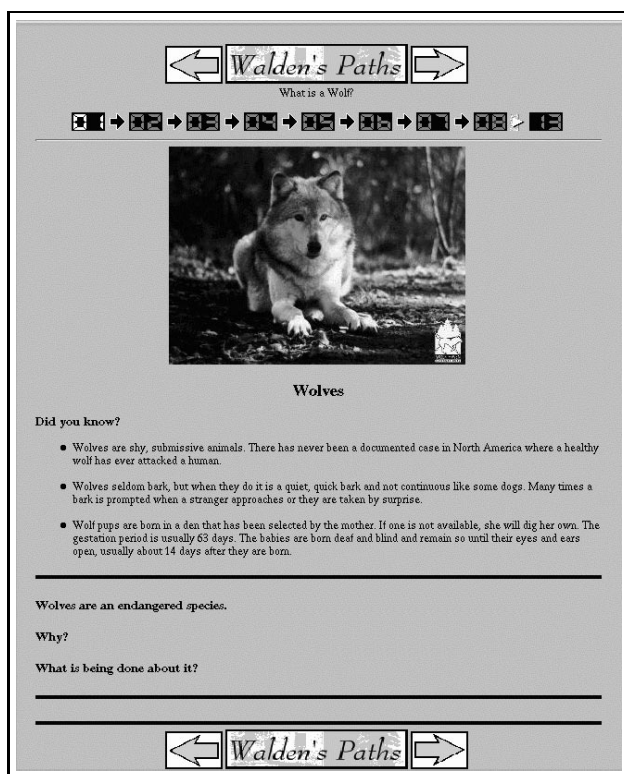


Figure 8: "Preface" material

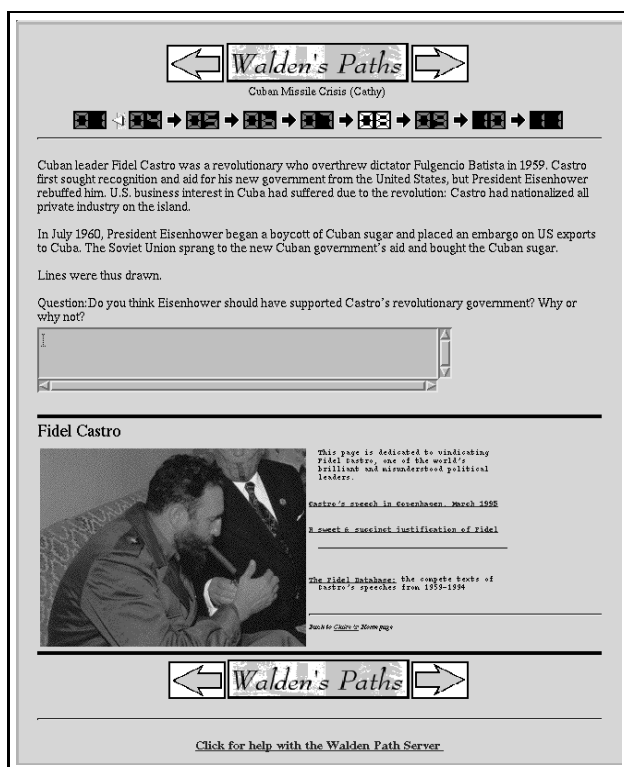


Figure 9: Minimized content

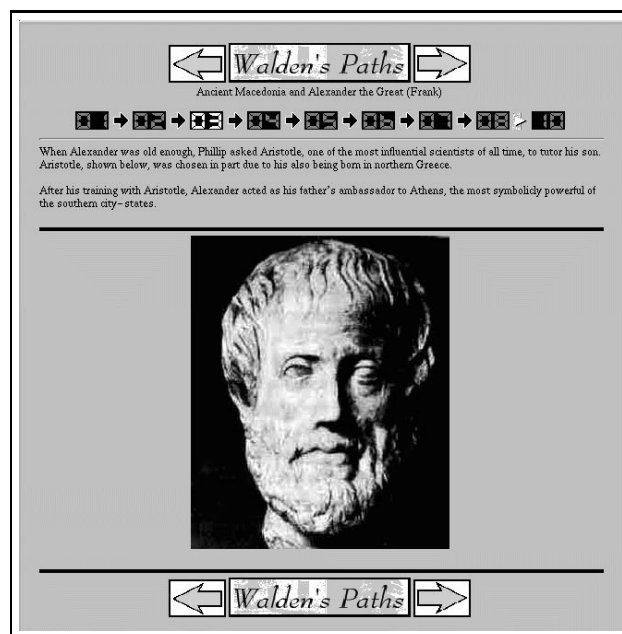


Figure 10: Partial content

The intellectual contribution of the Walden's Path author is the path and the annotation. However, curricular goals sometimes require modification or interpretation of others' intellectual contributions. Such adaptations are commonplace and accepted in print form (e.g., in textbooks). However, pages contained on a path are images of source material, not reformulations cast into uniform format, and hence disagreements are possible about appropriate and fair use of the material.

Discussions centered around transclusion [9, 11] suggest that acknowledgment and renumeration of original authors is an appropriate management model, but our experiences in the Digital Libraries domain [3] suggest that such acceptance will not be sufficient while transition is taking place. Recently botanists with whom we are working in the digital library domain encountered a sticky problem of reuse of some of their botanical images. The original source of the reused images was appropriately attributed. However the surrounding text incorrectly identified the species represented in the picture. The botanists were unhappy with this situation, even with accessibility of the original identifications and "renumeration" in the form of academic visibility, as they felt that the graphic juxtaposition of image and text was sufficient to cause confusion in the reader's mind as to *who* made the incorrect identification. In a sense, the *form* is as much a part of the intellectual contribution as is the *content*.

We encounter a combination, consequently, of competing and conflicting user requirements in developing specifications for Walden's Paths prototypes. Curriculum-based exposition benefits from extensive and flexible mechanisms to alter the appearance and content of path pages. Content authors, on the other hand, are often most comfortable with display mechanisms that fix both the content but also its context and appearance (consider, for example, the use of Adobe's Acrobat language for document dissemination).

Change is a further factor introduced by the distributed nature of Web documents that, unless countered, limits the useful lifetime of paths. Taken as a whole, material on the Web provides a dynamically changing collection of material—both content but also structure are undergoing continuous change. The environment on which paths are built represents a snapshot of the Web, frozen at the time of a path's creation. Without attention, the path-defined meta-structure rapidly loses synchrony with the external world's dynamically-defined Web structure.

We take two approaches towards limiting these entropy-like effects: detecting change and controlling the time of change. Detection of change depends on discovery of modifications to the underlying substrate (e.g., changes to addresses and to referenced pages) and initiation of corrective action. At present these activities are carried out manually, but we are in the process of implementing path validators that verify the correspondence between the current environment and the environment that existed when the path was written. Control of change centers around “freezing” the state of the path and its contents at well-defined points in time. Towards this end, we have implemented and use a “path compiler.” This allows us to give demonstrations without concern about changes to the substrate or to the network environment. Once again, however, a potential conflict exists between path compiler and content author since some Web pages are inherently dynamic (e.g., real-time or customized information) and compilation violates that author-desired characteristic.

7 EXTENSION AND CURRENT DEVELOPMENTS

The Walden's Paths system continues to evolve, driven by our user's requirements and needs. This section describes some of the mechanisms that are currently being prototyped, relating them to the issues discussed in the previous section.

7.1 Attribution and separation of content from annotation

In the previous section we discussed attribution of content from the viewpoint of the content author. As accurate attribution is a keystone of academic life, it is not surprising that it is also a goal of the teachers who are prospective path authors. Similarly, clear separation of annotation and original content also is desired by teachers—while the content author may desire this because it avoids confusion with other's words, teachers also benefit because it more clearly identifies *their* comments to *their* students.

Our ability to unambiguously separate annotation from content inherently is compromised by our system's architecture. The architecture's advantage of using the client-supplied user interface also requires that the annotation share display space with original content. While we can use graphical conventions to reduce confusion (and we continue to refine our designs to accomplish this), the possibility remains that *some* referenced page could be designed in a way to cause confusion.

7.2 Encouraging student engagement

Teachers have told us that they want active ways of generating student engagement with Web-based materials—they wish to avoid passive browsing behaviors, but instead want to make

sure that the students are reading and comprehending the material presented. Based on teacher suggestions, we have implemented several mechanisms and are in the process of evaluating them.

Figure 7 showed one example in which a teacher associated a set of questions with each path stop. We are presently evaluating a test version of Walden's Paths in which such questions are provided as an explicitly-supported annotation construct. Students are requested to “log in” when first encountering these paths. The teacher can associate a multiple-choice question with each page. Students see the question/answers as part of the annotation and can select one of the alternatives presented. The answers can be collected or summarized for the teacher's use.

A second mechanism, also teacher-suggested and also currently implemented and under evaluation, is the provision of student-authored annotations. Here, a student can enter an annotation at any stop along the path. These annotations are made available for reading to later visitors to the path.

7.3 Richer path structures

Adopting the KISS approach in designing our path model has been a good choice in most cases. As discussed earlier, teachers with a wide variety of computer experience were effective in authoring paths that they seemed to find of use. The simple form seemed particularly effective for paths that were based on the extended bookmark metaphor, as each stop in these paths tended to point to a large number of off-the-path resources. However, teachers with more programming-related experience have tended to ask for more complex structures, particularly if they were creating curriculum-based expositions with an extended tour flavor.

Initially we have decided to focus on appropriate introduction of sub-path mechanisms (i.e., incorporating hierarchy into the path specification). From the point of view of implementation, we view sub-path traversal as an analogue of off-the-path traversal. That is to say that selection of a sub-path allows independent traversal with single-button return to the point of departure.

We currently are comparing two different implementations of sub-paths. The first implementation provides sub-paths as a distinctive element of the annotation. Here the display of sub-paths is separated graphically from the rest of the annotation. The representation permits immediate visual location of the sub-path markers, a consistent graphical representation of the markers, and easy specification of sub-paths by the path author. The second implementation permits the embedding of sub-path references into the content of the annotation. Here the responsibility for distinguishing sub-paths (if at all) is that of the path author, a potential disadvantage. The primary advantage here is the high degree of flexibility in use of sub-paths afforded the path author. This approach also corresponds more closely to the philosophy of providing options to the path author rather than embedding them in the system's implementation.

8 DISCUSSION AND FUTURE DIRECTIONS

Our initial focus in the Walden's Paths project has been supporting teacher to student discourse. As the project matures, we are also interested in investigating the mechanism's potential for supporting communication in the other direction—from student to teacher. Perseus' use of paths [5] provides one example of paths in this context. Perseus paths can be a student-prepared artifact that communicates the student's understanding of a topic in a form that can be submitted to the instructor for evaluation.

We believe that an appropriate use of Walden's Paths would be in the context of a broader environment for information discovery, analysis, organization, and communication. In the same way that Perseus uses paths, the Walden's Path could substitute for the traditional term paper. In addition it could serve as the organizing framework for in-class presentation of information. However it is also clear to us that support of the broader goal requires significant additional thought, design, and implementation.

On a more general note, in this project we have found the World-Wide Web provides a provocative environment in which to conduct research. A strong opportunity exists in the Web for transfer of the results of earlier hypertext research. However the distributed, unregulated, heterogeneous nature of the Web raises issues that often did not appear in the original environments, which were more homogeneous both in the sense of users but also in the sense of software architecture. It is interesting to note that even with its heterogeneity, the Web remains surprisingly dependent on the good will of its users in support of its continuing technological survival. Much of the infrastructure of the Web appears to depend on adherence to "gentlemen's agreements," primarily relying on peer pressure for enforcement. Discovering, respecting, and negotiating such agreements adds an extra element of novelty to research in the Web environment.

REFERENCES

1. BROWN, M. H., AND SHILLNER, R. A. Deckscape: An experimental Web browser. SRC Research Report 135a, Digital Systems Research Center, Mar. 1995. Also available at <http://www.research.digital.com/SRC/publications/src-rr.html>.
2. BUSH, V. As we may think. *The Atlantic Monthly* 176, 1 (July 1945), 101–108.
3. FURUTA, R., MARSHALL, C. C., SHIPMAN III, F. M., AND LEGGETT, J. J. Physical objects in the digital library. In *Proceedings of the First ACM International Conference on Digital Libraries* (1996), pp. 109–115.
4. HALASZ, F. G., MORAN, T. R., AND TRIGG, R. H. NoteCards in a nutshell. In *Proceedings of the ACM CHI + GI Conference* (1987), pp. 45–52. Toronto, Ontario, (5–9 April, 1987).
5. MARCHIONINI, G., AND CRANE, G. Evaluating hypermedia and learning: Methods and results from the Perseus project. *ACM Transactions on Information Systems* 12, 1 (Jan. 1994), 5–34.
6. MARSHALL, C., AND IRISH, P. Guided tours and on-line presentations: How authors make existing hypertext intelligible for readers. In *Proceedings of Hypertext '89* (1989), pp. 15–26. November 5–8, 1989, Pittsburgh, Pennsylvania.
7. MARSHALL, C., AND SHIPMAN, F. Spatial hypertext: Designing for change. *Commun. ACM* 38, 8 (Aug. 1995), 88–97.
8. MARSHALL, C. C. Observations of sixth graders exploring the World-Wide Web, Jan. 1996. <http://www.csdl.tamu.edu/~marshall/caeti/observations.html>.
9. NELSON, T. H. *Literary Machines*, 87.1 ed. 1987.
10. NICHOL, D. 'Footsteps' linearising hypertext, 1996. <http://www-interact.eng.cam.ac.uk/footsteps/footsteps.html>.
11. SAMUELSON, P., AND GLUSHKO, R. J. Intellectual property rights for digital library and hypertext publishing systems: An analysis of Xanadu. In *Third ACM Conference on Hypertext Proceedings* (Dec. 1991), ACM, New York, pp. 39–50.
12. SHIPMAN, F. M., FURUTA, R., AND MARSHALL, C. C. Generating web-based presentations in spatial hypertext. In *Proceedings of the 1997 International Conference on International User Interfaces* (1997). To appear.
13. SHIPMAN III, F. M., MARSHALL, C. C., FURUTA, R., BRENNER, D. A., WEI HSIEH, H., AND KUMAR, V. Creating educational guided paths over the world-wide web. In *Educational Telecommunications, 1996: Proceedings of ED-TELECOM 96* (June 1996), P. Carlson and F. Makedon, Eds., Association for the Advancement of Computing in Education, pp. 326–331.
14. TRIGG, R. H. Guided tours and tabletops: Tools for communicating in a hypertext environment. *ACM Trans. Office Inf. Syst.* 6, 4 (Oct. 1988), 398–414.
15. ZELLWEGER, P. T. Directed paths through collections of multi-media documents. In *Hypertext '87* (Nov. 1987). Position paper.
16. ZELLWEGER, P. T. Active paths through multimedia documents. In *Document Manipulation and Typography*, J. C. van Vliet, Ed. Cambridge University Press, Apr. 1988, pp. 19–34. Proceedings of the International Conference on Electronic Publishing, Document Manipulation, and Typography, Nice (France), April 20–22, 1988.
17. ZELLWEGER, P. T. Scripted documents: A hypertext path mechanism. In *Hypertext '89 Proceedings* (Nov. 1989), ACM, New York, pp. 1–26.