Instructional Information in Adaptive Spatial Hypertext

Luis Francisco-Revilla and Frank Shipman
Center for the Study of Digital Libraries and Department of Computer Science
Texas A&M University
College Station, TX, 77843-3112, USA
1+(979)845-4924
{luis, shipman}@csdl.tamu.edu

ABSTRACT
Spatial hypertext is an effective medium for the delivery of help and instructional information on the Web. Spatial hypertext's intrinsic features allow documents to visually reflect the inherent structure of the information space and represent implicit relationships between information objects. This work presents a study of the effectiveness of spatial hypertext as medium for delivery of instructional information. Results were gathered based on direct observation of the people reading a spatial hypertext document which was used as informational support for a complex task. Two versions of the spatial hypertext document were used: a non-adaptive and an adaptive. The document was adapted based upon the inferred relevance of information to the user's knowledge and task requirements. The study produced insights on emergent reading strategies such as informed link traversals and the use of collections as bookmarks. Observations and evaluation of how people interacted with both document versions showed that the spatial layout and the use of collections as a way to encapsulate information allowed people to read, browse and navigate very large information spaces while maintaining a clear understanding the structure of the information. Finally, several differences between the adaptive and non-adaptive versions were identified, showing that adaptation alters not only the display of information but the way that people read spatial hypertext documents.

Categories and Subject Descriptors
H.5.4 [**Information Interfaces and Presentation**]: Hypertext/Hypermedia – *architectures, navigation, user issues.*

I.7.2 [**Document Preparation**]: document Preparation – *Hypertext/hypermedia, Multi/mixed media.*

General Terms
Documentation, Design, Experimentation, Human Factors.

Keywords
Information delivery, spatial hypertext, adaptation.

1. INTRODUCTION
Hypertext has proven invaluable for the delivery of information. The most popular hypertext implementation is the Web. It allows authors to create, organize and associate documents regardless of their ownership and location. It results in a link structure that readers navigate in order to access the desired information.

Navigational hypertext, such as the Web, bolstered the delivery of information to increasingly larger audiences. However, in navigational hypertext, readers are subject to the structure provided by the authors, who are responsible for creating structures that meet their audience's needs. As the size and diversity of the audience increases, it became harder to predict the needs of its members.

The pre-authored nature of the linking structure in conjunction with the vast differences across readers results in a rigid approach for information delivery. This has driven the development of new approaches that seek to provide increased flexibility in delivering hypertextual information such as Adaptive Hypertext and Spatial Hypertext.

Adaptive hypertext modifies the content and structure of documents in response to user characteristics such as preferences of knowledge. These characteristics are abstracted in a user model that guides the adaptation.

Spatial hypertext provides an expressiveness that augments the representation capabilities of hypertext. In contrast to navigational hypertext, where links are either present or not, spatial hypertext can represent *possible* relationships and implicit links. This provides a medium where authors can delegate to the readers the proper interpretation of the association between objects. Spatial hypertext supports direct manipulation of the objects and relationships present in the document. Thus readers interact with the document, actively cooperating in the process of attaining a proper presentation of the information.

Both of these approaches improve the basic hypertext concept by augmenting it in different ways that can complement each other. Allowing readers to participate in information structuring results in presentations that better fit individual needs. The visual nature of spatial hypertext allows readers to perceive the structure of the information, facilitating in turn the assimilation of the information. Similarly, adapting the presentation in response to aspects such as user knowledge is particularly useful when delivering instructional information.

Given these strengths, we are investigating Adaptive Spatial Hypertext to deliver instructional information. In particular, we created an adaptive spatial hypertext document containing approximately half of a book on HTML/XHTML. The approach was tested by presenting the document as instructional support to people conducting a complex task. This study had two goals: first, to evaluate the effectiveness of the spatial hypertext in general as a medium for the delivery of instructional information, and second, to explore the effect of adaptation in spatial hypertext.

The next section presents relevant research previously conducted in the areas of adaptive hypertext and spatial hypertext. The paper continues with the design of the spatial hypertext document. The evaluation and results follow. The paper concludes with a discussion of the results and observations of emerging interaction strategies.

Figure 1. Hierarchy and Containment

2. RELATED WORK

Prior work on adaptive hypertext and spatial hypertext heavily influenced the development of the present work. A brief overview of each is presented here to provide grounding for the description of the adaptive spatial hypertext document.

2.1 Adaptive Hypertext

In information-rich environments, such as the Web, the simple approach of gathering and presenting all available information can easily overwhelm readers. One way to alleviate this is to contextualize the delivery by filtering and presenting only the *right* pieces of information. This requires the ability to differentiate the relevant from the irrelevant. However, this is not a straightforward process, as the assessment of what is relevant varies from person to person. Consequently, research in adaptive hypertext has focused on customizing the presentation according to a user model that represents significant user characteristics such as goals, knowledge and preferences [1, 4].

Human actions, however, are situated and depend heavily on the particular context [21]. Thus, in addition to the user, it is necessary to take into consideration other relevant factors that also demand the adaptation of the presentation of the information. As a result, adaptive hypertext research now often includes additional adaptation models, e.g. task and situation models [3, 7,8].

Adaptive hypertext has evolved into a field in its own right. It produces a large corpus of documented work and holds periodical international conferences (AH 2000, AH 2002, and AH2004). While, as noticed by Calvi [2], research in adaptive hypertext has been focused mainly in the domain of education it has also explored areas such as intelligent tutoring [5], context-sensitive help [3] and information retrieval [10].

2.2 Spatial Hypertext

Spatial hypertext originated evolved from observations of how people use systems such as Aquanet [15]. Researchers discovered that explicit links were cumbersome for many activities. People often employed only the relative spatial position between objects in order to express relationships among them [13, 14, 19]. These observations prompted the development of Spatial Hypertext systems like VIKI [16] and VKB [20]. These systems provide a two dimensional space for information objects to be placed. The relative positioning of objects implicitly creates relationships such as lists, piles etc. Spatial hypertext systems use parsers in order to recognize these implicit composite objects.

In addition to implicit composites, spatial hypertext supports the ability to group objects into collections. Collections are explicit composite objects that instantiate a hierarchy of navigable information spaces [16] as shown in Figure 1. They provide a "window" into a sub space that can contain more objects and collections. Similar to window systems, users can navigate into a collection by "maximizing" the collection such that it expands to fill the space available on screen. Similarly "minimizing" a collection causes it to collapse into its minimum size. Figure 1 shows the collection "10. Tables" maximized. The collection "10.3 Advanced Table Tags" remains in its normal size, and the collections with titles beginning with "10.1", "10.2" and "10.4" all appear minimized.

In order to provide a smooth navigation of the spatial hierarchy, spatial hypertext systems often *animate* the maximizing, minimizing and normalizing of collections.

Spatial hypertext's ability to define a hierarchy of spaces and to visually represent relationships between objects makes it potentially well suited for the delivery of instructional information. This is discussed in the following section.

3. DOCUMENT DESIGN

Creation of the spatial hypertext required reflection on the inherent structure of the information in the document and the visualization of the relationships between the different components. Careful consideration was required in deciding the dynamic and adaptive behaviors of the document and which

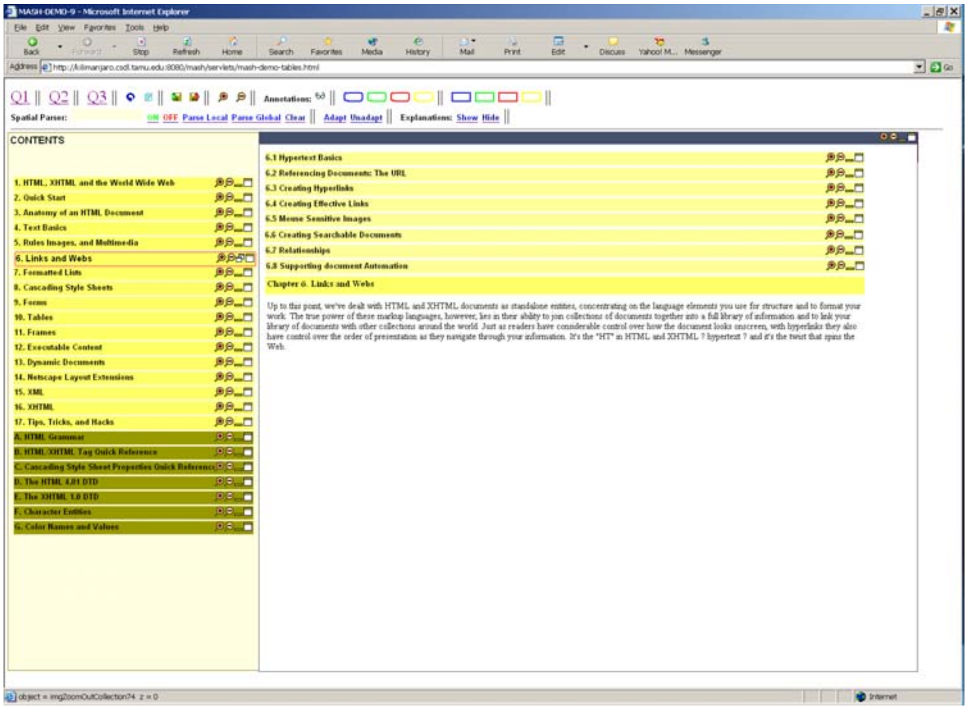

Figure 2. Emulating navigational hypertext in spatial hypertext

aspects should be considered for adaptation. Detailed discussion of these reflections is provided in the following sections.

3.1 Topic and source of information

The first decision was to select a domain and collect the information to be delivered. An extensive domain was desirable, as it would help to test the limits of the approach. Availability of previously-authored and reputable sources of information that ensured the quality of the contents was a must. Finally, the domain should facilitate easy evaluation conditions.

After considering and discarding several domains, content authoring in HTML/XHTML was selected as it facilitated the desired experimental conditions, and required the presentation of relatively large amounts of information. Also, quality content was readily available on the Web.

Several reputable sources were evaluated and finally O'Reilley's "HTML & XHTML: The Definitive Guide" [11] was selected. Its coverage of the domain and clear writing fit readers of all levels. Additionally, by virtue of being available on-line through O'Reilley's Safari digital library [18], it provided a navigational interface that aided in the design of the spatial hypertext document.

There was a challenge that needed to be tackled. For the spatial hypertext document to be useful to novice and expert readers, it needed to contain the whole book, or at least a large portion of the on-line book. This content was distributed over a 100 Web pages with text and images in the Safari digital library [18]. This was an excellent test of spatial hypertext's ability to properly encode and support navigation of large information spaces within a single document. Due to system and time constraints, we included 10 of the 17 chapters and one of the seven appendices.

3.2 Reflecting structure of the information

Instructional information often has a strong hierarchical nature that divides the contents into sections and subsections. Many Web sites serving instructional information provide a list of links that reflect this inherent structure. For instance, the Safari digital library [18] provides a "Table of Contents" Web page for each of its on-line books. Clicking on an item of the table of contents returns the associated Web page containing the selected chapter or subsection. This approach forces readers to discover the underlying structure of the information throughout navigation and inference. In contrast, spatial hypertext visually reflects the structural organization of the information. For instance, hierarchy can easily be mapped into a hierarchical arrangement of collections containing collections, as shown in Figure 1.

Sections and subsections are encapsulated into collections and sub-collections. As a result, readers directly perceive the structure. Visualization of the relationship between objects and their location within the information hierarchy can be reinforced by color coding such that major sections have darker colors than sub sections.

From a user perspective, maximizing a collection is similar to traversing a link. Both cause the currently visible information to be replaced by alternate information either within the collection or at the other end of the link. However, collections and transclusion links [17] in spatial hypertext are powerful alternatives to links in navigational hypertext. While the latter requires the readers to traverse links in order to discover what lays ahead, the former offers readers a view of the destination, allowing them to make informed decisions about whether they should traverse the link or not.

3.3 Document layout

Design of the document layout commenced by considering the transfer of layouts frequently encountered in navigational hypertext. This resulted in experimenting with layouts such as the one showed in Figure 2, which shows the document structure on the left with the particular content displayed to the right.

In order to imitate the functionality of Web pages, behaviors attached to objects automatically move and open the collections inside the work area (shown on the right side of Figure 2).

However, this approach did not take advantage of many features of spatial hypertext and required the implementation of additional dynamic behaviors. Thus, alternative layouts were designed to take advantage of the spatial features of the medium being used and to avoid requiring the definition of navigational controls

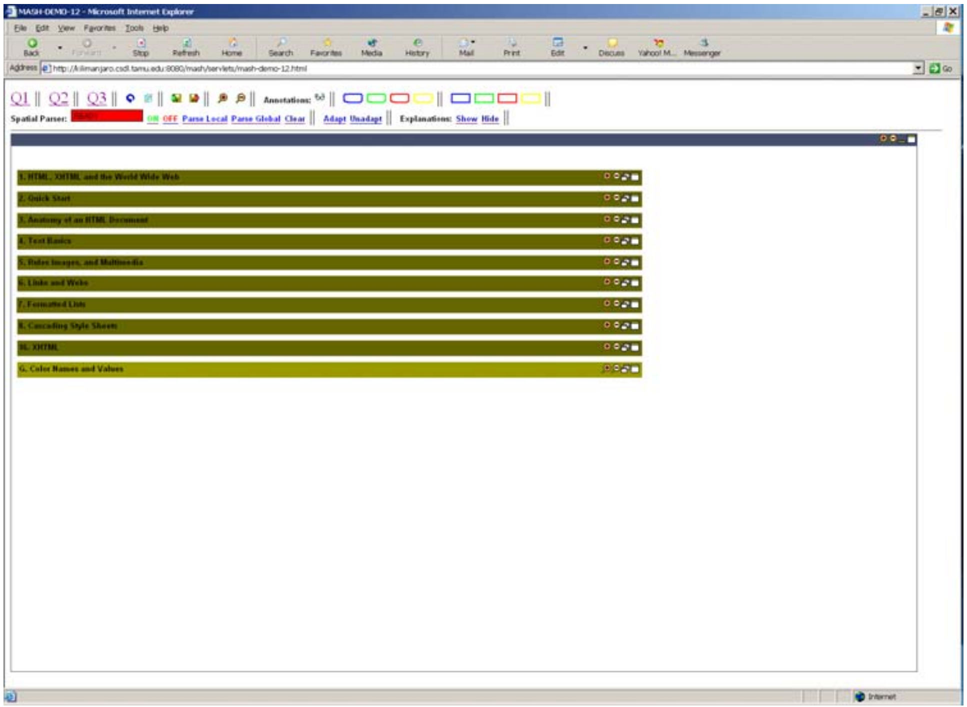

Figure 3. Basic spatial arrangement of a hierarchical information space

extraneous to spatial hypertext. The simplest and most obvious arrangement, shown in Figure 3, was to represent one section as a collection and arrange all sub sections as a vertical list of collections.

Using minimized collections that only take a single line allow all sections in the same hierarchical level to be shown as a list, which implicitly expresses the sequential ordering of its elements. This basic construct of hierarchically nested lists of collections fits very well with the ordering of sections and subsections of a book.

3.4 Dynamic behaviors

While minimized collections provide a compact way to represent the information contained in a given level of the hierarchy, people reading the spatial hypertext need to "open" the collections in order to access their contents. Collections can be opened by maximizing or resizing them. However, opening a collection will obscure collections that follow it in the list. Rather than occluding other collections, it is often preferred to *push* the list down, as shown in Figure 4.

This emulates Nelson's notion of stretch text for the document structure [17]. In spatial hypertext, stretching can be performed simultaneously on the vertical and horizontal dimensions. Stretch space is one example of dynamic behaviors available in spatial hypertext. There is a vast range of useful behaviors that enhance the presentation of information by animating typically static media such as text [12]. However, dynamic behaviors risk becoming excessive as they can distract readers from the reading process itself. Thus, only stretch space was initially selected.

3.5 Adaptive behaviors

Differentiating relevant from irrelevant information often plays a critical role in the reading of instructional information. Systems that adapt documents in such a way that relevant information is more noticeable than irrelevant information facilitate the assimilation of the required knowledge for the task. Consequently, we decided to focus our research on documents that adapt in response to the user's knowledge and task requirements.

Part of the adaptive document creation was deciding how to "emphasize" and "de-emphasize" objects. The philosophy was to employ multiple visual cues that could represent varying degrees of relevance. After some experimentation, relevant objects were marked with a red glow, increased size and font size, and a higher

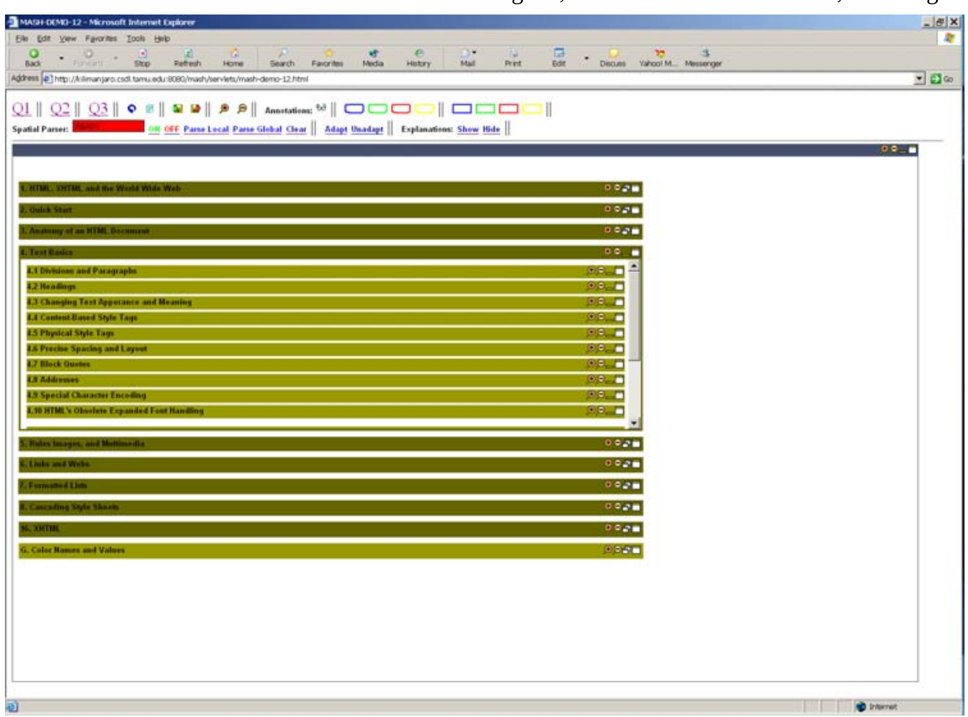

Figure 4. Stretch lists

zoom factor. Objects deemed irrelevant were visually altered by fading them out, reducing their size and font size, and zooming them out. Figures 6, 7 and 9 show several examples of relevant and irrelevant objects.

3.6 Final interface

During the process of authoring, the spatial hypertext was naturally rearranged in order to keep track of the progress and to compare sections. During this task it became clear that the previous layout shown in Figure 4 was too linear, not taking advantage of the second spatial dimension. Therefore, it was decided to create a horizontal list of vertical lists. The resulting document is shown in Figure 5.

Interaction with the space revealed some trade-offs associated with stretch space. As all collections expanded vertically, the stretch list improved the management of the space. However, it also restricted the reader's flexibility to re-arrange objects, and hindered the simultaneous view of all sections. The reduction of collection width allowed the simultaneous presentation of all sections and subsections, but also resulted in the instantiation of horizontal scrollbars. Finally, this interface was used for the evaluation of spatial hypertext and adaptive spatial hypertext as a medium for the delivery of instructional information.

4. EVALUATION

The evaluation procedure presented the spatial hypertext document shown in Figure 5 to a set of 15 subjects, randomly selected from the Texas A&M University and neighboring areas of Bryan and College Station. All participants were proficient in English and their ages ranged from 20 to 40 years old.

The experimental goal was to observe and gather insights of how subjects read and interacted with spatial hypertexts documents. The experiment included two types of spatial hypertext: adaptive and non-adaptive. Participants were randomly divided into two groups such that 8 subjects used the adaptive and 7 the non-adaptive spatial hypertext document.

The spatial hypertexts were shown to the participants using WARP, a multi-model spatial hypertext system [9], which was configured to support adaptive and non-adaptive spatial hypertext.

The adaptive version used two models in order to guide the adaptation, a task model and a user's knowledge model. The task model suggested adaptations to the document (shown in Figure 5) based on what topics and subtopics are useful for authoring different types of Web pages. The user model suggested adaptations in response to user knowledge, emphasizing or de-emphasizing topics and subtopics that best fit the knowledge level

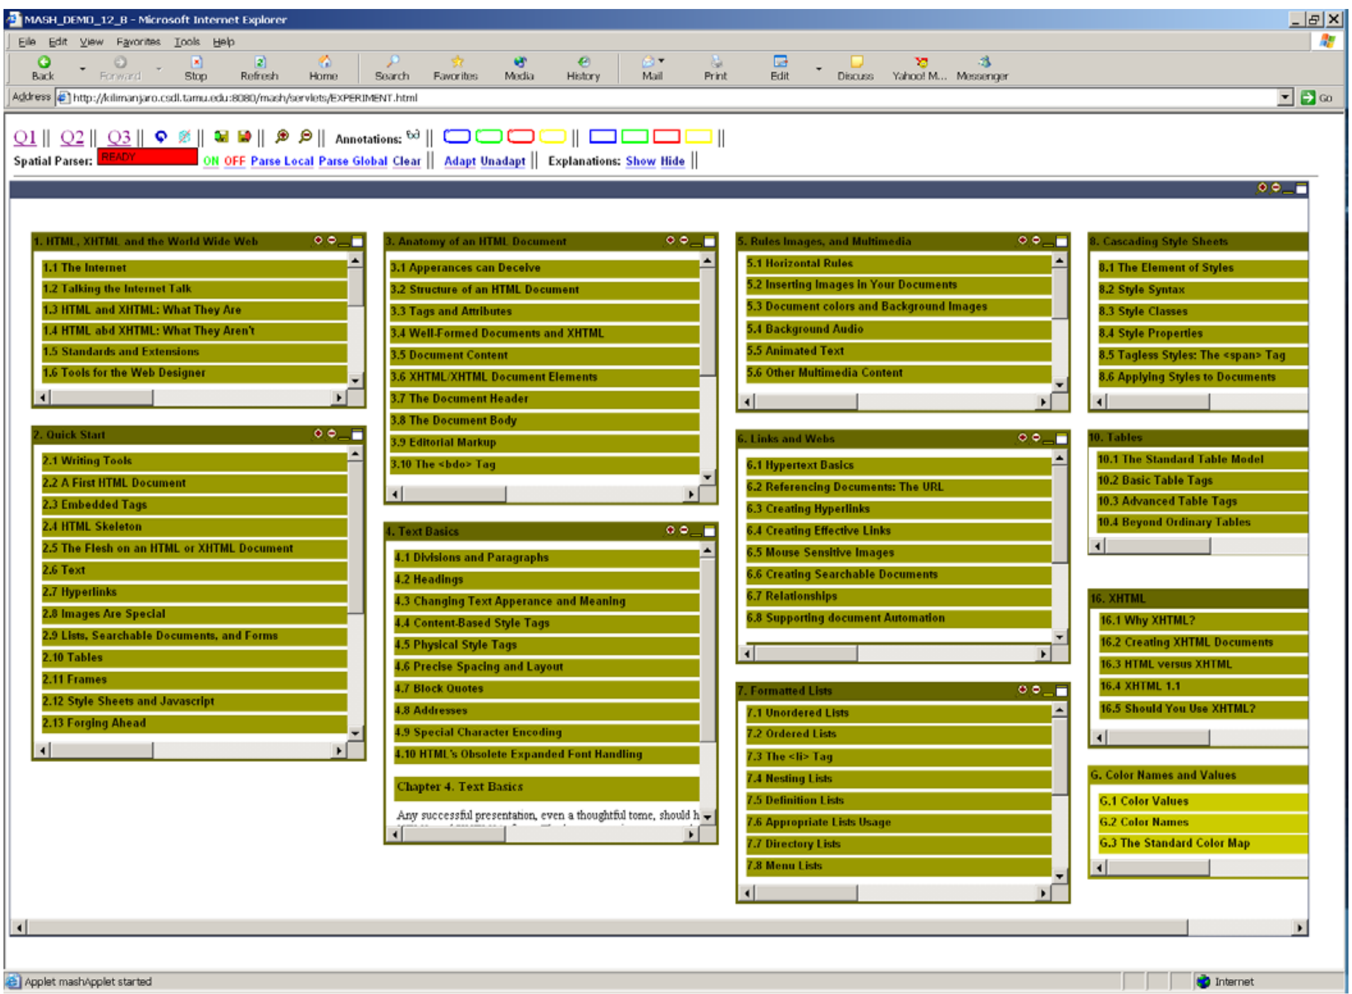

Figure 5. Final Interface

of the user. In order to infer the knowledge of the user, the system presented the user with a questionnaire about HTML and analyzes his/her answers. Figure 6 shows the interface after the completing the adaptation process for one of the participants.

In addition to the HTML questionnaire, subjects answered another on-line questionnaire that gathered demographic information. After completing a 20 minute training session in the use of the software applications employed in the experiment, participants were asked to author a Web page within 90 minutes. The task was carefully crafted such that novices and experts were able to author the Web page. At the same time the task was demanding enough such that both novices and experts were asked to include features previously unknown by them that required the use the spatial hypertext as informational support.

The interactions of the study participants with the spatial hypertext document were closely monitored by the investigators. This provided valuable observations of how the people read spatial hypertext documents.

After completing the Web page or the allotted time expired, the subjects were asked to answer an on-line questionnaire about their experience using the spatial hypertext document.

Evaluation sessions concluded with the investigators interviewing the participant in order to collect additional comments and suggestions, and to answer subject questions. The commentaries gathered this way provided valuable insights about reading spatial hypertext documents

5. RESULTS

The study produced several results in regard to the activity of reading spatial hypertext documents. These are insights that emerged from interviews, analysis of the on-line questionnaires, and observations of the subjects' interactions with the document. The results show that while some reading strategies apply to spatial hypertext in general, others are used depending upon whether the spatial hypertext is adaptive or non adaptive. The following sections first present the general results and then discuss the differences in reading behaviors between adaptive and non-adaptive spatial hypertext.

5.1 Layout

The two-dimensional arrangement of the information was very well received by participants using the adaptive and non-adaptive interfaces. When asked if they had any comments, participants often began with expressions like: *"it was very well organized"* and *"this is so much better than Tables of Contents"*. Even if the reduced width of the collections resulted in a limited view of the

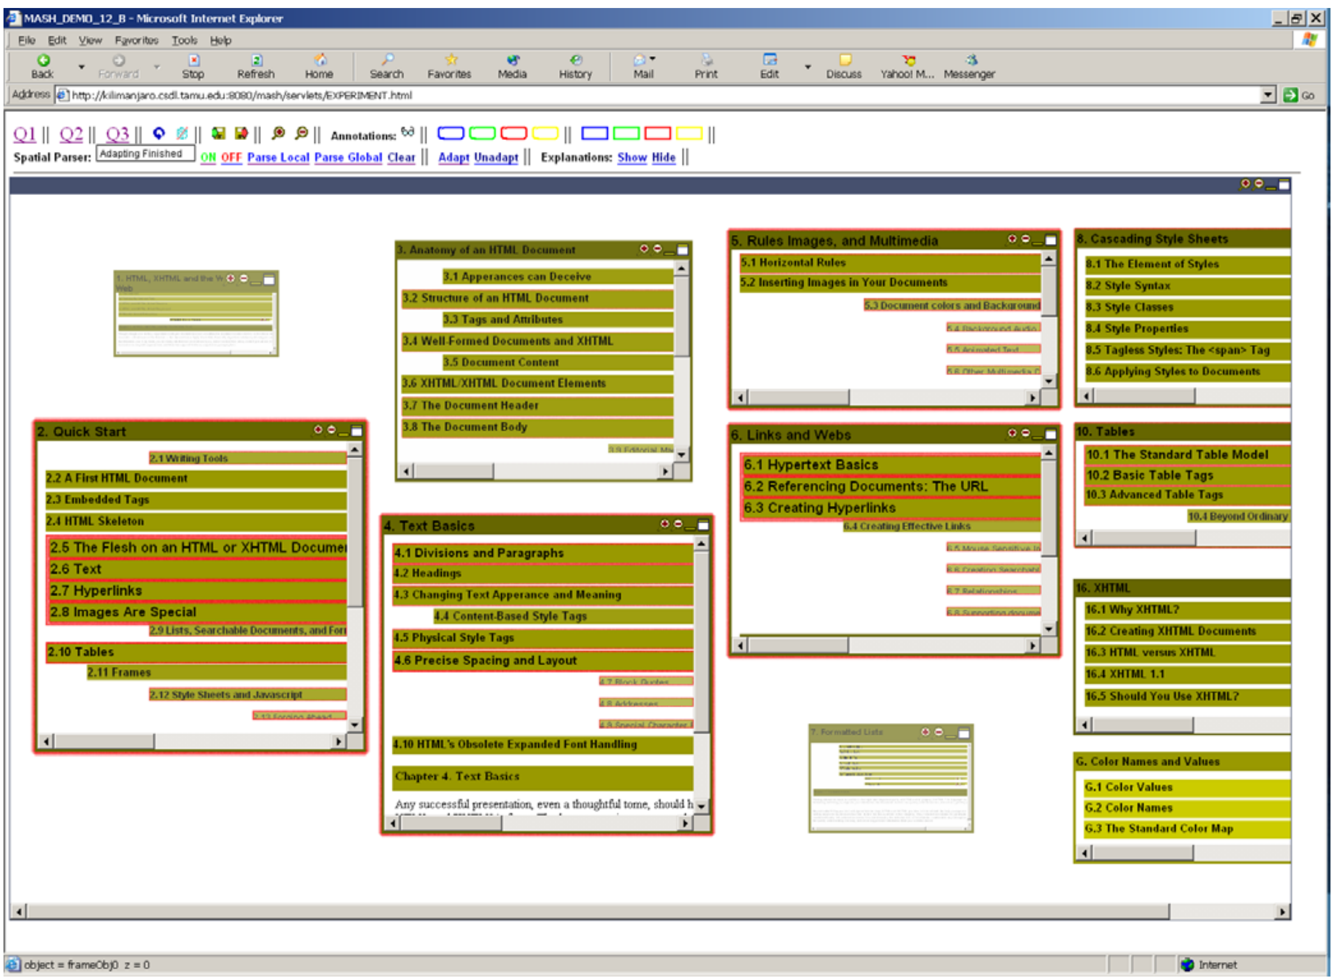

Figure 6. Adapted interface

contents, requiring them to scroll horizontally, participants appreciated the fact that they could quickly see the complete structure of information, saying: *"I really like that I can see all of the chapters"*.

5.2 Moving and re-arranging

The participants' appreciation of the layout did not prevent them from moving objects around the space. Overall, 73% of the participants moved the collections during their reading of the space. Examples of the interfaces after completion of the task are shown in Figures 7 and 8.

Often, moving collections was related to the act of reading rather than to the restructuring of the space. For instance, participants often dragged objects to the center of the screen *"for convenience"*, or moved objects around as they resized collections. In some cases, after completing their exploration of the collection's contents, participants went to considerable lengths in order to maintain the pristine state of the document. After completing the task, their interfaces looked exactly the same as in the beginning.

Other participants, however, consciously moved collections in order to group *"what is more important*", to *"see both and compare"*, and to signify what is being read or has been read *"for reference"*. For instance, in Figure 7 the participant has moved (and piled) several collections to the top left corner. In the process of reading, the subject has resized these collections and scrolled to display the relevant information. In contrast, other collections remain in their location and initial state.

More participants using the non-adaptive interface rearranged the layout than those using the adaptive interface (86% vs. 63%). The difference was not statistically significant (p = 0.35) but it can be observationally explained. In the case of the non-adaptive interface, all collections are potentially relevant. Thus, when relevant information is found, participants wanted a visual cue that facilitated revisiting that information. As moving is the easiest action to visually change a spatial hypertext and is often used in the course of the normal reading activity, users chose to move collections as a way to identify the important information. In contrast, the adaptive interface already provided a visual encoding of objects' relevance, diminishing the incentive for additional user actions.

5.3 Navigation and Collections

The experiment centered heavily on the use of containment as a metaphor for navigation. This proved quite powerful as all participants were comfortable with the concept and understood the structure of the information. On a scale from 0 to 6 – 0 being the worse and 6 the best – the participants' evaluation averaged 5.60 regarding the ease of understanding how information was contained in collections.

Participants reported that they could navigate with ease through the entire document even though the document included content originally divided across 90 Web pages. On average they judged the ease of navigation to be 4.65 on the 0-6 scale. However, some participants expressed that they lowered their score in response to a glitch in the software, namely that the animation to open and close collections was too slow.

5.3.1 Informed link traversals

The study revealed that collections can be used in many ways that had not been anticipated. The original assumption was that subjects would tend to proceed by first maximizing a section, then exploring its contents, and finally maximizing the subsections as appropriate. However, rather than maximizing a collection and then reviewing its contents with the extra space provided by a full screen display, most participants explored the contents using the normal size of the collection. Even the necessity to scroll horizontally did not appear to obstruct their tasks. Furthermore, relevant subsections were often maximized inside the normal-sized parent collection, as shown in Figure 7.

When asked about this, participants commented that they wanted to review the contents before committing to opening collections that would fill the entire screen. This strategy allowed them to quickly explore and compare multiple collections, selecting to open only useful collections. One participant clearly expressed this as:

"You are not clicking on a bunch of links that may or may not have what you are looking for"

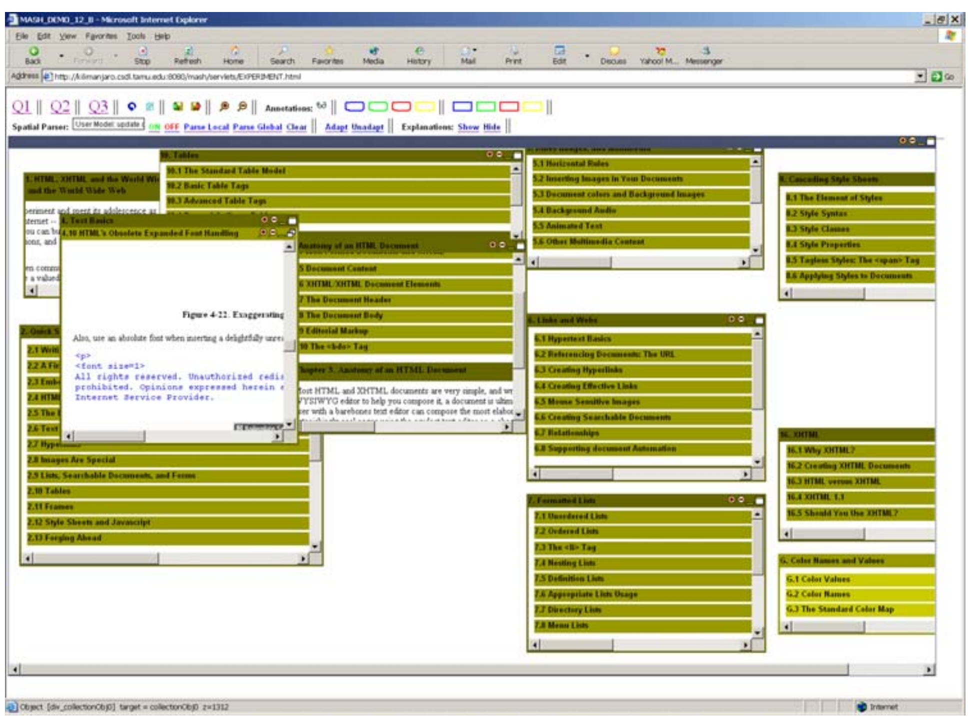

Figure 7. Non-adaptive interface after completing the task

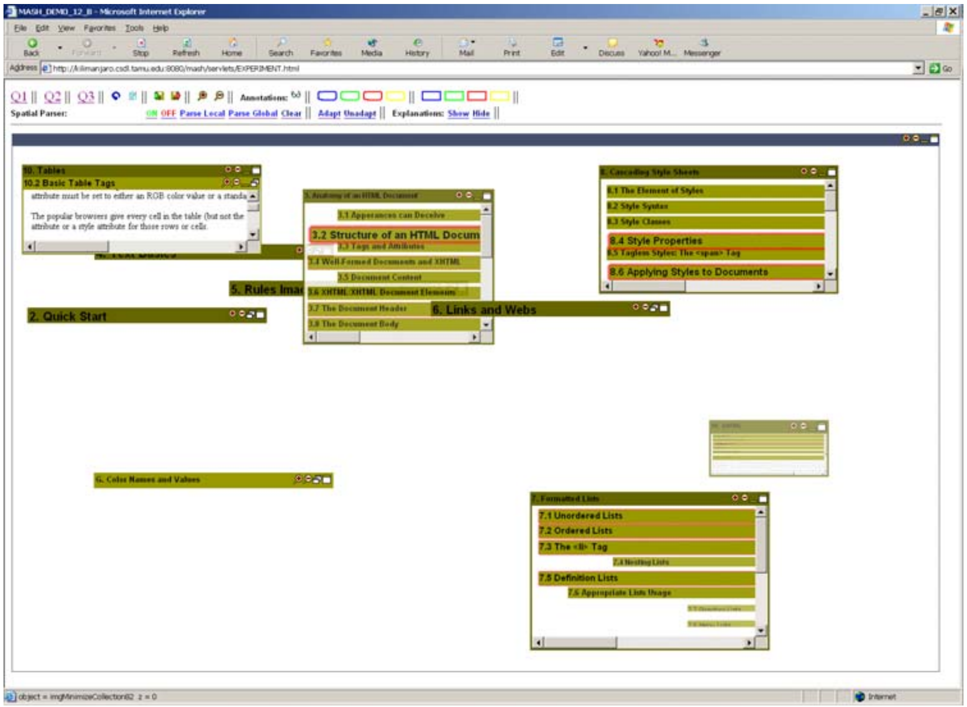

Figure 8. Non-adaptive interface after task completion

This ability to partially take a link by maximizing within a parent collection instead of the entire window was important to the work practices participants developed.

5.3.2 Using collections as bookmarks

A second discovery that emerged from the handling of collections and sub-collections is the use of collections as bookmarks. Having located relevant information in certain subsections, users often left these open to that location, while they browsed on other collections for additional information.

Figure 8 shows an example of this practice. In this case, the participant identified relevant content inside section 10.2. The collection containing 10.2 is maximized. However, not only is section 10 not maximized but its size has been reduced and it has been located out of the way in the top left corner. The participant employed this strategy in order to conduct other reading activities such as browsing and exploring other collections while keeping track of the content and the fact that section 10 was valuable.

5.3.3 Minimizing

This bookmark strategy was not limited to normal sized collections. Investigators also noticed that many minimized collections, such as the ones visible in Figure 8, also contained maximized sub-collections. After interviewing and observing several participants it was discovered that this was an intentional extension of the bookmark strategy. After finishing working in a section, participants minimized the collections in order to make space for the reading of new sections. However, predicting that they might need to access those same contents later, they decided to keep the sub-collections open. This behavior is in stark contrast with cases when the participants assumed that they would not need to revisit the section any time soon. In such cases, participants tended to either restore the original structure of the section or document or chose to minimize the section or sub-section.

Interviews also revealed that participants consciously used the minimization of collections to signify completion of a section. Figure 9 illustrates this point.

Sections 2, 4, 8 and 10, visible in Figure 9, have been minimized. However the overall structure of the document has been left constant and no actions have been taken to utilize the unused space. As the participant in this case stated, minimizing meant *"I am done with that"*.

5.4 Adaptation

In addition to the moving of collections previously mentioned, the experiment revealed other behavioral differences between adaptive and non-adaptive spatial hypertext. These included the changing of the object's visual appearance, and zooming.

When using the adaptive document, 63% of the participants changed the visual look of objects. As for participants using the non-adaptive version, only 43% changed the original appearance of objects. In regards to the use of the zooming, 88% of the participants in the adaptive case used the zoom feature while only 57% of the subjects in the non-adaptive case used it.

This is not unexpected, as the additional effort from the subjects was motivated by visually excessive adaptations, such as extreme font sizes (large or small) and zoom factors. The glow effect was effective as an initial way to get the attention, but it became obtrusive subsequently. Hence participants chose to cancel it out. Changes of objects' visual features in non-adaptive cases focused on resizing collections, although one subject experimented with emphasizing important objects, before starting to re-locate important objects on the top left-corner of the screen.

The design decision to use multiple cues was validated by the participant interviews. Participants distinguished relevant from irrelevant objects very easily (5.1 on the 0-6 scale). However, not all participants were conscious about the different visual cues used, asking questions such as *"What red glow?"* Having multiple cues facilitated the use by different people. However, different subjects complained about liking of disliking the modification of different features. Discussion with the participants led to the conclusion that adaptation should take into consideration the user's preferences for the selection of appropriate and meaningful visual cues.

Irrelevant objects played an important role. The adaptation mechanisms were designed to never hide objects. Instead objects were de-emphasized in such a way that they were visually less prominent than the rest of the objects in the document. This approach was very successful as most readers always chose to explore and navigate into emphasized collections before de-

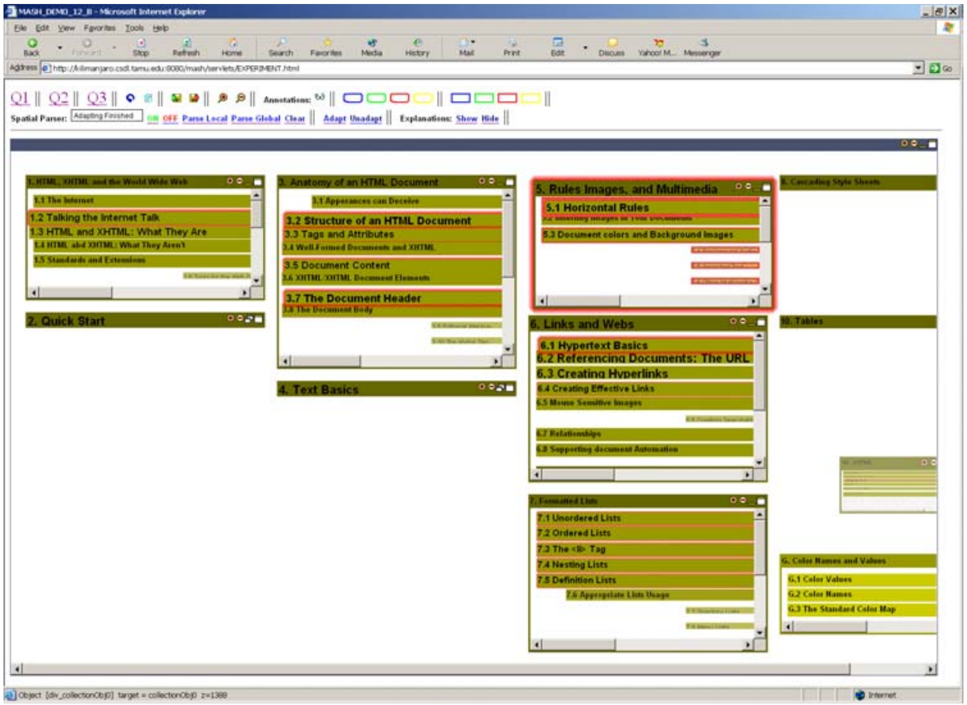

Figure 9. Simple use of minimization to signify completion of reading

emphasized ones. This was the case even when the participants thought that they were navigating the list of sub-collections sequentially. An explanation for this is that perceptually, participants were filtering out the de-emphasized collections before making cognitive decisions about exploration and navigation. This hypothesis seems validated by participants' comments like: *"There were not many de-emphasized objects"*

Finally, adaptation is not valued or desired by all users. For instance, one participant actually *fought* the adaptation, commenting that *"I didn't like that the system was trying to make me look into certain sections, so I decided to check all of them sequentially"*. This was lamentable failure, as the design of the spatial hypertext and the adaptation was trying to present the adaptations as suggestions rather than commands. The participant's comment serves to stress the important observation that document adaptation should be optional and, when used, clearly presented as such.

6. Conclusions

Spatial hypertext is an effective medium for the delivery of help and instructional information on the Web. It facilitated the creation of efficient layouts that allowed readers to navigate through large amounts of information. These layouts are not mere replicas of navigational hypermedia, as they rely on intrinsic features of spatial hypermedia such as the readers' ability to re arrange and manipulate objects.

Spatial hypertext has often been used in order to organize information, with spatial objects providing links to external documents that contain the desired information. In contrast, the document used in this work contained all the information in itself. To the best of our knowledge this is the largest spatial hypertext document to date in terms of sheer textual content.

The study showed the emergence of navigation and orientation strategies performed spontaneously by people reading spatial hypertext documents. Users often moved and modified the information objects as part of their reading process. Readers used changes in object appearance or location to group important objects, to compare objects and to keep track of what had been read.

The encapsulation of information into hierarchically ordered collections was shown to be an effective way to visually convey the structure of the information. Observations and evaluation of how people interacted with both spatial hypertext versions showed that people readily understood the structure of the information space, and were able to read, browse and explore book-sized documents with ease.

In addition to reflecting the information structure, collections and transclusion links augment navigation links as they support readers making of informed decisions about the usefulness of traversing a link. While both collections and navigational links provide a transition into a separate information space, collections and transclusion links provide a partial view of their contents, allowing readers to see what there is at the other side of the links before having to traverse them. Interviews and observations of the interactions with the spatial hypertext document confirmed the readers' appreciation of this feature as they often based their decision about maximizing a collection based on the previous exploration of its contents.

Adaptation of the document by directly relating the visual prominence of objects and collections to their relevance provided an effective way to facilitate the navigation of spatial hypertext documents. The study revealed the value of representing these adaptations using multiple visual cues. Comments from participants indicated that document adaptation should be optional and selection of the particular adaptation techniques should be made cooperatively with the reader.

As part of this first study, dynamic behaviors were kept to a minimum facilitating the observation and analysis of the effects of spatial hypertext basic features upon the reading activity. Further experimentation will be required to assess the value and effectiveness of other more dynamic behaviors.

7. ACKNOWLEDGMENTS

This work is partially funded by NSF grants. IIS-9734167, DUE-0121527 and DUE-0226321. We also thank Pratik Dave for providing valuable feedback on an earlier draft of this paper.

8. REFERENCES

[1] Brusilovsky, P. *Methods and techniques of adaptive hypermedia. User Modeling and User-Adapted Interaction*, 6, 2-3 Kluwer academic publishers, 1996, pp. 87-129

[2] Calvi, L. The Future of AH. In Third Workshop on Adaptive Hypertext and Hypermedia (Århus, Denmark, August 14-18, 2001). http://wwwis.win.tue.nl/ah2001/

[3] Conlan, O., Wade, V., Bruen, C., and Gargan, M. Multi-model, Metadata Driven Approach to Adaptive Hypermedia Services for Personalized eLearning. In *Proceedings of the Second International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems* (May 29 - 31, 2002). Springer-Verlag London, UK. pp. 100-111.

[4] De Bra, P., Houben, G-J., and Wu, H. AHAM: a Dexter-based reference model for adaptive hypermedia. In *Proceedings of HT'99* (Darmstadt, Germany, February 21-25, 1999) ACM Press, pp. 147-156.

[5] Encarnação, L. M. Multi-Level User Support through Adaptive Hypermedia: A Highly Application-Independent Help Component. In Proceedings of Intelligent User Interfaces '97 (Orlando, Florida - January 6-9 1997) ACM Press, New York, USA, pp. 187-194.

[6] Fisher, S., and Steinmetz, R. 2000. Automatic Creation of Exercises in Adaptive Hypermedia Learning Systems. In Proceedings of the Eleventh ACM Conference on Hypertext and Hypermedia (San Antonio, Texas, USA, May 30 – June 4) ACM Press, New York, USA, pp. 49-55.

[7] Francisco-Revilla L. *User and Situation Models for Medical Information Delivery*. Masters Thesis, Dept. of Computer Science. Texas A&M University, College Station, TX, USA, 1998.

[8] Francisco-Revilla L., and Shipman, F.M. Adaptive Medical Information Delivery: Combining User, Task and Situation Models. In *Proceedingsof IUI 2000* (New Orleans, USA, Jan 9-12, 2000) ACM Press, pp. 94-97.

[9] Francisco-Revilla L., and Shipman, F.M. MASH: A Framework for Adaptive Spatial Hypermedia. In *the Third Workshop on Spatial Hypertext* (Nottingham, UK, August 26, 2003). http://www.csdl.tamu.edu/~shipman/SpatialHypertext/SH3/luis.pdf

[10] Höök, K. Evaluating the Utility and Usability of an Adaptive Hypermedia System. In Proceedings of Intelligent User Interfaces '97 (Orlando, Florida - January 6-9 1997) ACM Press New York, USA, pp. 179-186.

[11] Kennedy, B., and Musciano C., *HTML & XHTML: The Definitive Guide*, 5th Edition, O'Reilly

[12] Lee, J.C., Forlizzi, J., and Hudson, S.E. The Kinetic Typography Engine: An Extensible System for Animating Exspressive Text. In *Proceedings of Uist'02*, (Paris, France, October 27-30, 2002) ACM Press, pp. 81-90.

[13] Marshall C.C., and Rogers R.A. Two Years before the Mist: Experiences with Aquanet, In *Proceedings of HT'92* (Milano, Italy, November 30-December 4, 1992) ACM Press, pp. 53-62.

[14] Marshall C.C., and Shipman, F. M. III. Searching for the Missing Link: Discovering Implicit Structure in Spatial Hypertext. In *Proceedings HT'93* (Seattle, Washington, USA, November 14-18, 1993) ACM Press, pp. 217-230.

[15] Marshall C.C., Halasz, F.G., Rogers, R.A., and Janssen W. C. Jr. Aquanet: A Hypertext Tool to Hold Your Knowledge in Place. In *Proceedings of HT'91* (San Antonio, Texas, USA, December 15-18, 1991) ACM Press, pp. 261-275.

[16] Marshall C.C., Shipman, F. M., and Coombs, J. H. VIKI: Spatial Hypertext Supporting Emergent Structure. In *Proceedings of the 6th ACM European Conference on Hypermedia Technology* (Edinburgh, Scotland, September 18-23, 1994) ACM Press, pp. 13-23.

[17] Nelson, T. H. The Heart of Connection: Hypermedia Unified by Transclusion. In *Communications of the ACM* 38, 8 (August 1995) ACM Press. Pp 31-33.

[18] ProQuest, *Safari Tech Books Online.* http://proquest.safaribooksonline.com/

[19] Shipman, F. M., Marshall, C. C., and Moran, T. P. Finding and Using Implicit Structure in Human Organized Spatial Layouts of Information. In *Proceedingsof CHI'95* (Denver, Colorado, USA, May 7-11, 1995) ACM Press, pp. 346-353.

[20] Shipman, F., Hsieh, H., and Airhart, R., Maloor, P., and Moore, J.M. The Visual Knowledge Builder: A Second Generation Spatial Hypertext. In *Proceedings of HT 2001* (Århus, Denmark, August 14-18, 2001) ACM Press, pp. 113-122.

[21] Suchman, L. A. Plans and Situated Actions. Cambridge University Press, New York, USA, 1987.