

Graphical Multiscale Web Histories: A Study of PadPrints

Ron R. Hightower, Laura T. Ring, Jonathan I. Helfman,
Benjamin B. Bederson¹, James D. Hollan²

Computer Science Department

University of New Mexico

Albuquerque, NM 87131

(505) 277-3112

www.cs.unm.edu/pad++

{high, lring, bederson, jon, hollan}@cs.unm.edu

ABSTRACT

We have implemented a browser companion called PadPrints that dynamically builds a graphical history-map of visited web pages. PadPrints relies on Pad++, a zooming user interface (ZUI) development substrate, to display the history-map using minimal screen space. PadPrints functions in conjunction with a traditional web browser but without requiring any browser modifications.

We performed two usability studies of PadPrints. The first addressed general navigation effectiveness. The second focused on history-related aspects of navigation. In tasks requiring returns to prior pages, users of PadPrints completed tasks in 61.2% of the time required by users of the same browser without PadPrints. We also observed significant decreases in the number of pages accessed when using PadPrints. Users found browsing with PadPrints more satisfying than using Netscape alone.

Keywords

World Wide Web, Web Navigation, Web Browser, Usability, Pad++, Zooming User Interface (ZUI), Hypertext, Multiscale Interfaces, Information Visualization.

INTRODUCTION

Navigating hypertext structures like the World Wide Web (WWW) is difficult for users. After following a number of links, people can have difficulty remembering where they have been and returning to previously visited pages. According to a WWW usability study, 13.4% of subjects report not being able to find pages recently visited [9].

The same usability study found that while only 0.1% of page accesses were through the history list, 42% of page accesses used the Back-Button. So, while pages are revisited with a high frequency, the history list is largely

To appear in the proceedings of the ACM Hypertext 98 conference.

HCIL Technical Report No. 98-06 (November 1998);

<http://www.cs.umd.edu/hcil>

Proceedings of ACM Conference on Hypertext (Hypertext 98), ACM Press, 58-65




Figure 1: The PadPrints Browser Companion



Figure 2: The PadPrints display, showing close-up of nodes in the graphical history hierarchy.

1. Current Address: Computer Science Department., Human-Computer Interaction Lab, University of Maryland, bederson@cs.umd.edu

2. Current Address: Cognitive Science Department, University of California, San Diego, hollan@cogsci.ucsd.edu

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE NOV 1998		2. REPORT TYPE		3. DATES COVERED 00-11-1998 to 00-11-1998	
4. TITLE AND SUBTITLE Graphical Multiscale Web Histories: A Study of PadPrints				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Defense Advanced Research Projects Agency,3701 North Fairfax Drive,Arlington,VA,22203-1714				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT We have implemented a browser companion called PadPrints that dynamically builds a graphical history-map of visited web pages. PadPrints relies on Pad++, a zooming user interface (ZUI) development substrate, to display the history-map using minimal screen space. PadPrints functions in conjunction with a traditional web browser but without requiring any browser modifications. We performed two usability studies of PadPrints. The first addressed general navigation effectiveness. The second focused on history-related aspects of navigation. In tasks requiring returns to prior pages, users of PadPrints completed tasks in 61.2% of the time required by users of the same browser without PadPrints. We also observed significant decreases in the number of pages accessed when using PadPrints. Users found browsing with PadPrints more satisfying than using Netscape alone.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 8	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

unused. This suggests a problem with the history mechanisms found in current browsers. Three shortcomings of the common history mechanism immediately suggest themselves. First, the history list is incomplete—large parts of the history disappear when users follow a branch point. Second, the history list is textual and page titles may lack the cues needed to find a particular page. Third, the history list is cumbersome to use. A user must pull down a menu before finding and following the desired entry.

Better history mechanisms could benefit millions of WWW users. There are a variety of human-factors issues to be considered [11]. Some issues are addressed by proposed alternatives and extensions to browsers. WebMap is a browser extension that shows a graphical relationship between web pages [6]. Each page is represented by a small circle that can be selected to display the actual page. Links between pages are colored to indicate information such as whether it is a link to a different server or whether the destination page has already been read. Webmap graphs may be saved and used by others. WebBook is a three-dimensional Web browser that allows multiple pages to be viewed simultaneously and supports collecting pages in books to replace bookmark files [4]. MosaicG is a modified version of the Mosaic browser that incorporates a graphical history mechanism [1]. In contrast to other systems that build maps for individual websites, MosaicG displays pages as users access them, resulting in personal maps of web usage.

WWW pages often have links to pages linking back to a starting point. It is difficult to visualize these cycles, so many web visualizations are based on hierarchies extracted from graphs of web structure. Some interesting work focuses on alternative visualizations of such structures. Furnas [8], for example, demonstrates how multi-trees can be used to represent a collection of hierarchies sharing parts of underlying structure. One application of multi-trees is visualization of bookmarks from multiple individuals [17].

In addition to problems associated with extracting hierarchies from general graph structures, layout and graphical display of large numbers of nodes is inherently difficult due to limited screen space. As we discuss below, zooming multiscale user interfaces [2][3][15][18] provide one viable approach to this problem.

A ZOOMING GRAPHICAL HISTORY

We have built a browser companion, called PadPrints, to aid web navigation. Figure 1 shows the PadPrints history-map to the left of a browser window. Figure 2 shows a detail of the history-map and some of the PadPrints controls. PadPrints is similar to the MosaicG graphical history (MGGH) [1] in depicting a left-to-right hierarchical history of web pages visited during a browsing session. Also in both systems a node in the hierarchy displays the title of the web page and a small picture associated with the

page. Finally, both systems construct the hierarchy as users traverse links from one page to another, as opposed to prebuilding a hierarchy for a single website as WebMap and other systems do [6].

As with MGGH, the PadPrints browser companion monitors and controls the web browser. When users access pages from the web browser those pages are added to the PadPrints display. Pages are added as children of the current node in the hierarchy, unless that page is already present in the hierarchy. A single click on a page in the PadPrints display sends the browser to the corresponding URL.

Three characteristics distinguish PadPrints from MGGH. First, PadPrints was implemented in Pad++, a substrate for building zooming user interfaces [2][3]. Pad++ allows users to zoom in and out to explicitly control how much context is viewed at any time. The multiscale nature of Pad++ allows geometric objects, text, and images to be displayed at arbitrary size on the screen. This permits the user to view the graphical history in its entirety or zoom in to focus on a particular part of the hierarchy. We conjecture that multiscale contextual display of the graphical history can provide important support for navigation. In addition, multiscale facilities of Pad++ allow natural ways to temporarily remove sub-hierarchies from view by shrinking them until they occupy minimal screen area. PadPrints was written as a Pad++ application using the Tcl scripting language [14].

The second difference between PadPrints and MGGH is the use of a proxy server to monitor browser behavior. Our proxy server notifies PadPrints when the browser (Netscape Navigator) retrieves a page. PadPrints then adds a new page to the graphical hierarchy or, if the URL is already present in the hierarchy, marks the existing page with a yellow outline to denote it as the current page. To achieve similar functionality in MGGH, Ayers and Stasko had to modify the Mosaic browser code to monitor page accesses. This has two drawbacks. Their graphical history is limited to use with one browser and their code changes must be reintegrated as the browser changes to support new HTML features. PadPrints was designed to work with any browser that supports a proxy server and should not require modification as browsers change to support new HTML additions.

Both systems display a small thumbnail sketch of a web page on each node in the hierarchy. The source of this small image represents the third difference between the two systems. Because they had direct access to the implementation of the Mosaic browser, Ayer and Stasko were able to re-render the web page into a hidden bitmap, subsample the resulting image, and display it in their graphical history. PadPrints, without direct access to the renderer, instead uses the X server to grab that portion of

the screen associated with the browser. The resulting image data is subsampled and displayed with the node. We note that the multiscale nature of Pad++ supports displaying the full resolution view of a web page with each node, limited only by available memory.

USABILITY TESTING

We performed two usability studies to test the effectiveness of PadPrints. We hypothesized that a graphical hierarchical history mechanism would improve users' web browsing efficiency in terms of speed and number of pages accessed as compared to having only traditional history mechanisms as available in Netscape Navigator 3.0. We also hypothesized that users of PadPrints would find browsing more satisfying than using Netscape Navigator 3.0 alone.

Equipment

Experiments were performed on personal computers, running Linux. All systems had 21-inch monitors. Questions and tasks were presented to the subjects by a program that recorded statistics such as the time subjects spent on each question, pages accessed, and questions answered correctly.

Stimuli

Two different sets of web pages were used in the testing. The first set was based on the database from the "browse off" at CHI'97 [13]. This is a hierarchical database of people, places, things, events, and facts. Each page contains a word or a list of words, possibly with links to other pages lower in the hierarchy. There are no images or graphics. This database was chosen because it provides a means to test simple navigation. Throughout the rest of the paper this database is referred to as the *CHI Database*.

Another web page collection was taken from the National Park Service (NPS) website. This website contains photographs and information about national parks in the United States. In contrast to the CHI database, the NPS database is a large real-life website containing images and has a complex structure that is not entirely hierarchical. A copy of the website was downloaded to a local server to avoid network problems during the experiments and to help minimize variability in access times.

EXPERIMENT 1

Method

Subjects were asked to find pages, answer questions about individual pages, and make comparisons between pages. Generally, all questions except the first required some backtracking. Example questions used with the CHI database include: *Go to the page that contains the Kangaroo. How many moons does the planet Uranus have? Are there more whales or more apes listed in this database?* Some example questions used with the National Parks website are: *How many National Parks are in California? Find a picture of the General Sherman Tree in*

Sequoia National Park in California. Which has more islands, Channel Islands National Park or Apostle Islands National Lakeshore? Subjects completed 14 questions using the CHI database and 12 questions using the National Parks website.

A 2x2x2 incomplete-block design was used. With this design we were able to collect more information from each subject than if we had used a between-subjects design and reduce some of the between-subject variability. The independent variables were, whether PadPrints was used, which website (CHI database or National Parks Site), and task order. Each subject was randomly assigned to one of four experimental groups. All subjects browsed both websites and did one set of questions with PadPrints and one set of questions without PadPrints. As a result of the incomplete-block design employed in the experiment, the PadPrints-by-Website interaction was confounded with the Subject-by-Block effect.

Three dependent variables of interest were average time to completion of each question, average number of pages accessed for each question, and user satisfaction ratings. User satisfaction was measured using selected questions from the Questionnaire for User Interaction Satisfaction (QUIS) developed at the University of Maryland [10], as well as some browsing specific questions written in the same format. The QUIS uses a 9 point rating scale.

Only questions a subject completed correctly were included in the time averages. Furthermore, the first question in each task was treated as a practice trial and not included. Questions that over 25% of subjects could not answer correctly were also not included. Two questions from the National Park's task were thrown out for this reason and one question from the CHI database task was thrown out.

Training

Prior to beginning the experiment subjects were trained in use of the Netscape browser and PadPrints. In the first experiment this training was done using the University of New Mexico website. Training for Netscape Navigator, consisted of having the subjects use the forward and back keys, the bookmark list, and the go list. Subjects were told they could make their own bookmarks if they felt it would help them with the task.

PadPrints training included telling subjects a little about PadPrints and how it could be used for navigation. It was explained that PadPrints was a visual history mechanism. Subjects were then instructed to visit a series of pages on the UNM website and return to those pages using PadPrints. Subjects also practiced zooming and panning within the Pad++ environment.

Subjects

Thirty-seven subjects participated in Experiment 1. One subject's data was not used because of failure to complete over 70% of the questions correctly. Subjects ranged in age from 18 to 50 years old. Mean age was 26. All but two subjects were students at the University of New Mexico. 83% of subjects had previously used Netscape Navigator to browse the Web. 14% of subjects reported never browsing the World Wide Web. 50% of subjects reported browsing the web an average of two hours a week or less. All but two subjects used computers more than two hours a week. 56% of subjects used computers over 10 hours a week. Of the student subjects, 24% were computer science majors, 21% majored in other sciences or mathematics, 27% majored in health-related areas, 12% majored in education or educational technology, 9% majored in the arts, and 3% majored in business. 44% of subjects were female. All subjects were paid for their participation.

Results

In Experiment 1 we observed a statistically significant improvement in average number of pages accessed (i.e., fewer pages were accessed) when using PadPrints, increased user satisfaction ratings with PadPrints, but no difference in average time to complete each question

The subject effects were confounded with the Website-by-PadPrints-used interaction. Thus main effects only ANOVA models were fit, with effects estimated for Subjects, Website, Task-Order and PadPrints, for each dependent variable. The nine dependent variables were: average time on task, average number of pages accessed, total QUIS score, and 6 scores for the individual parts of the QUIS. Since 9 ANOVAs were run Bonferroni adjusted critical values were used. For an effect to be considered significant p had to be less than 0.006, but $p > 0.09$ for all tests so there were no significant effects due to task-order for any of the dependent variables.




Figure 3. Mean Number of Pages Accessed for Each Question in Experiment 1.

Although we found using PadPrints resulted in significantly fewer page accesses ($F_{1,30} = 18.78$, $p = .0002$), this did not lead to faster task completion (with $F_{1,33} = .93$, $p = .343$). Mean pages accessed are displayed in Figure 3.

User satisfaction was measured using questions from the Questionnaire for User Interface Satisfaction (QUIS). Subjects indicated increased satisfaction when using PadPrints as is illustrated in Table 1.

QUIS Section	Average Satisfaction		$F_{1,33}$	$p <$
	Without PadPrints	With PadPrints		
Overall	5.748	7.506	29.49	0.0001
Screen	6.778	7.181	2.80	0.1038
Feedback	6.148	7.037	14.76	0.0005
Learnability	6.081	7.608	32.14	0.0001
Speed	7.602	8.120	3.50	0.0701
Browsing	6.053	7.681	29.97	0.0001
Total	6.402	7.522	29.58	0.0001

Table 1. User Satisfaction Ratings for Experiment 1

Using Bonferroni adjusted critical values, a significant effect was found for Overall User Satisfaction, Satisfaction with Feedback, Satisfaction with Learnability, and Satisfaction with Browsing. The total User Satisfaction Score was also significantly higher when PadPrints was used.

EXPERIMENT 2

In Experiment 1 subjects using PadPrints expressed greater satisfaction and accessed fewer pages. However, we did not observe any savings in task completion times. We conjectured that although PadPrints did not save users time in finding information on pages not already visited that PadPrints would save time in returning to pages previously visited. In Experiment 2 we specifically tested this conjecture.

Method

In Experiment 2 subjects were instructed to visit a set of pages and to answer questions about certain pages. The CHI database and the National Parks Site were again used. For the National Parks website, subjects performed a pre-task with 16 questions. For the CHI database subjects were asked 23 questions. After the pre-tasks 10 questions were asked that required subjects to return to already visited pages. Half of these questions could be answered with information from one page, the other half required information from multiple pages. Subjects were instructed to try again if they answered a question incorrectly or failed to visit the correct page. After three tries if a subject could not complete the question they were allowed to move on.

For each subject, time to answer each question and number of pages accessed were recorded. Only questions subjects answered correctly within three attempts were included in computing averages. Only 1.4% of questions were not correctly answered within three attempts.

Training

The training in Experiment 2 was very similar to the training for Experiment 1. Prior to beginning the National Parks task, subjects were trained on the Olympic National Park website. Prior to beginning the CHI database task, subjects were trained using a portion of the CHI database. In each case, training involved portions of the databases not used in the experimental tasks.

Subjects

Thirty-six subjects participated in Experiment 2. Subjects ranged in age from 19 to 60 years old. The average age was 33. All but four subjects were students at the University of New Mexico. Five subjects had participated in Experiment 1. 92% of subjects had previous Netscape Navigator experience. 33% of subjects reported browsing the Web less than two hours a week. All subjects had some web browsing experience and reported using computers more than two hours a week. 72% of subjects used the computer more than 10 hours a week. Of the student subjects, 31% were computer science majors, 22% were education majors, 19% were social science majors, 16% were majoring in mathematics or engineering, 3% were art majors, 3% were business majors, and 6% were non-degree students. 53% of subjects were female. All subjects were paid for their participation.

Results

Data analyses for Experiment 2 were identical to Experiment 1, except we looked at ten dependent variables instead of nine. Bonferroni adjusted critical values were employed (p had to be less than .005 for an effect to be considered significant). The 10 dependent variables were: average time on the pretask questions, average time on the experimental questions, average number of pages accessed on the experimental questions, the total QUIS score and 6 scores for each individual part of the QUIS. There was one significant effect due to task order for the user satisfaction ratings, $F_{1,33} = 9.92$, $p = .0035$. A breakdown of means suggests this was due to users giving Netscape Navigator alone a lower rating after they had used PadPrints than if they had never used PadPrints. It is not possible to test this explicitly in the current model because of the confounding inherent in the design.

Consistent with Experiment 1 there was no significant difference in time to answer questions for the pre-tasks between subjects using PadPrints and those not using PadPrints, $F_{1,33} = 0.08$, $p = 0.78$. However, there was a significant decrease in time to answer questions using PadPrints for the experimental tasks requiring revisiting

pages, $F_{1,33} = 66.23$, $p < 0.0001$. Mean times to complete questions are displayed in Figure 4.




Figure 4. Mean Response Times in Experiment 2.

Fewer pages were accessed when PadPrints was used, $F_{1,33} = 133.61$, $p < .0001$. Means are shown in Figure 5.




Figure 5. Mean Number Pages Accessed in Experiment 2.

Consistent with Experiment 1 subjects expressed greater satisfaction when using PadPrints. The User Satisfaction Ratings are given in Table 2.

QUIS Section	Average Satisfaction		$F_{1,33}$	$p <$
	Without PadPrints	With PadPrints		
Overall	5.106	7.977	101.82	0.0001
Screen	6.363	7.704	24.98	0.0001
Feedback	5.458	7.972	55.59	0.0001
Learnability	5.550	8.189	72.32	0.0001
Speed	7.592	8.611	21.25	0.0001
Browsing	5.229	8.069	83.44	0.0001
TOTAL	5.883	8.087	87.78	0.0001

Table 2. User Satisfaction Ratings for Experiment 2

The high satisfaction ratings were further supported by comments made by subjects. Examples of user comments are:

Navigating Without PadPrints

The [NPS] information was easy to read once I found it. It was frustrating that Go saved only my main link and not the pages I had been to.

[With Netscape Navigator] I have to rely on my own memory too much to find a way back to already visited pages.

Navigating With PadPrints

Enjoyed the ease of finding old pages.

It was so easy. There was no confusion whatsoever. It was simple to go where I needed to go and get back to my starting point. I never got lost and I loved the fact that everything was right in front of you without having to go and scroll down "book mark" or some other option. I would really enjoy using this system more often.

Very helpful I used it mostly for the big leaps - when I just needed to go back or forward just one page I used the netscape buttons - I like this global view and it is fun too.

This is a totally cool system, I'm going to miss not being able to have it on my desktop.

FUTURE DIRECTIONS

The development and testing of PadPrints raised a number of issues that should be resolved in future research. These issues fall into the three categories of system structure, interaction, and visualization issues.

The primary issue with system structure is the use of a proxy server as a way of monitoring user-browser interactions. The advantage of this technique is that the browser need not be modified to work with PadPrints. The disadvantage is that user actions must be deduced from the request stream passing through the proxy server. This makes it difficult to distinguish certain actions. For example, we could not distinguish between a user following a link and typing in a URL. Ideally browsers would supply this type of detailed information to external programs. However, we expect that we could deduce this with limited accuracy, based on parsing the HTML source.

Interaction issues are one of the most important parts of future work. One can imagine many different ways to interact with PadPrints and choosing among the variations will require additional usability studies. What information is most useful to the user? What history structures are most

appropriate? Is there a single history structure that is sufficient for all web tasks, or is there a handful of different tasks that each require a unique history structure? Should the history-map include only webpages from the current browsing session, all sessions of the last week, or all pages every visited by the user? The history-map could display all webpages visited by a group of users. How should the history-map be organized if a large number of pages are involved? Test subjects expressed a desire to delete nodes and modify the content of a history-map. What types of editing operations would be most useful?

Finally, there are many visualization issues. Are the thumbnail pictures essential to the success of the history-map, or are webpage titles sufficient? How should pages be arranged and connected in the history-map? Sometimes it is appropriate to display pages just once in the history-map, but other times multiple instances of the same page are appropriate. Also, the multiscale nature of Pad++ allows nodes in the history-map to be displayed at different sizes, but which nodes should be small and which should be enlarged? Is page size controlled automatically or should the user be responsible for controlling this feature of the display? Which of the many display options will be most beneficial to the user?




Figure 6. Using the multiscale capabilities of Pad to enhance the layout of the hierarchy.

The version of PadPrints used in the two experiments reported here was frozen for testing prior to adding a number of extensions. Among these are alternative tree layout algorithms and mechanisms for tree editing.

Figure 6 shows two views of a PadPrints graphical history hierarchy. The nodes in the left hierarchy are displayed at a constant scale factor while the nodes in the hierarchy on the right are scaled according to tree depth. This simple layout algorithm provides a fisheye view of the tree. Pad++ permits users to zoom into the tree and view the nodes at a legible size.




Figure 7. Manual editing of the tree layout. Two subtrees in the middle have been scaled very small to reduce their consumption of screen space.

Figure 7 shows a PadPrints graphical history after a user has edited the layout manually. Here two subtrees have been scaled very small in order to compress their space on the screen. Extensions like these need to be tested in order to discover which will make PadPrints more effective for the user's needs.

CONCLUSION

We have shown that the PadPrints graphical history-map is an effective way to support web navigation. In two usability studies, PadPrints reduced the number of page accesses during general navigation and significantly reduced access time for tasks requiring page revisitation. The results from questionnaire ratings showed subjects experienced higher degree of satisfaction navigating the web using PadPrints than without PadPrints.

In addition to these quantitative results, many subjects commented that they enjoyed using PadPrints. A number of our subjects expressed an interest in using the software for their own work. One subject was a teacher who suggested that PadPrints might be particularly useful in an educational setting. He indicated that while the web is becoming very useful in the classroom, teachers and students frequently experience problems navigating the web while trying to find specific information. He thought PadPrints might help reduce navigational problems. Furthermore, teachers could provide a condensed PadPrints history-map to students to summarize useful websites. In

the future it would be informative to bring PadPrints into the classroom and empirically evaluate these conjectures.

PadPrints is a promising graphical multicale history technique to aid web navigation. Providing users with a visual history-map allowed them to more effectively return to information found in previously visited webpages. By reducing the burden of navigation more cognitive resources can be spent on finding, integrating, and processing new information, potentially making the World Wide Web a more useful information source.

ACKNOWLEDGEMENTS

We acknowledge generous support from DARPA's Human Computer Interaction Initiative, Contract #N66001-94C-6039, and from the United States Department of Energy, Contract DE-AC04-96AL85000. This work was also supported by a gift from the Intel Corporation. In addition, we acknowledge our Pad++ collaborators at New York University Media Research Lab, especially Jon Meyer, and at the University of Michigan.

REFERENCES

1. Eric Ayers and John Stasko. "Using Graphic History in Browsing the World Wide Web", The 4th International World Wide Web Conference, www.w3.org/Conferences/WWW4/Program_Full.html, December 11-14, 1996.
2. Benjamin B. Bederson and James D. Hollan. "Pad++: A Zooming Graphical Interface for Exploring Alternate Interface Physics", *Proceedings of UIST'94* ACM Press, 17-26.
3. Benjamin B. Bederson, James D. Hollan, Ken Perlin, Jon Meyer, David Bacon, and George Furnas. "Pad++: A Zoomable Graphical Sketchpad for Exploring Alternate Interface Physics", *Journal of Visual Languages and Computing* (7), 1996, 3-31.
4. Stuart K. Card, George G. Robertson, and William York, "The WebBook and the Web Forager: An Information Workspace for the World Wide Web", *Proceedings of CHI'96*, ACM Press, 111-117.
5. Andy Cockburn and Steve Jones. "Which way now? Analysing and easing inadequacies in WWW navigation", *International Journal of Human-Computer Studies* (45), 1996, 105-129.
6. Peter Doemel, "WebMap - A Graphical Hypertext Navigation Tool", 2nd International Conference on the World Wide Web, Chicago, IL, 1994, 785-789.
7. William C. Donelson. "Spatial Management of Information", *Proceedings of SIGGRAPH'78*, ACM Press, 203-209.
8. George W. Furnas and Jeff Zacks. "Multitrees: Enriching and Reusing Hierarchical Structure", *Proceedings of SIGCHI'94*, ACM Press, 330-336.

9. GVU's WWW Surveying Team. "GVU's 6th WWW User Survey", www.cs.gatech.edu/gvu/user_surveys/survey-10-1996, 1996.
10. B. Harper and K. Norman. "QUIS 5.56: Questionare for User Interaction Satisfaction", University of Maryland at College Park, 1994.
11. Wendy A. Kellogg and John T. Richards. "The Human Factors of Information on the Internet", in *Advances in Human Computer Interaction*, Ed. J. Nielsen, Ablex Press, (5), 1-36.
12. Sougata Mukherjea, James D. Foley, and Scott Hudson. "Visualizing Complex Hypermedia Networks through Multiple Hierarchical Views", *Proceedings of CHI'95*, ACM press, 331-337. .
13. Kevin Mullet, C. Fry, and Diane Schiano. "On Your Marks, Get Set, Browse!", *Proceedings of CHI'97*, Human Factors in Computing Systems, Extended Abstracts, 113-114.
14. John K. Ousterhout. *Tcl and the Tk Toolkit*, Addison-Wesley, 1994.
15. Ken Perlin and David Fox. "Pad: An Alternative Approach to the Computer Interface", *Proceedings of SIGGRAPH'93*, ACM Press, 57-64.
16. David Ungar and Randy B. Smith. "Self: The Power of Simplicity", *Proceedings of OOPSLA'87*, 227-241.
17. Kent Wittenburg, Duco Das, Will Hill, and Larry Stead. "Group Asynchronous Browsing on the World Wide Web", *Proceedings of the 4th International World Wide Web Conference*, Boston, MA, International WWW Conference
18. Benjamin B. Bederson and James D. Hollan. "Pad++: A Zoomable Graphical Interface System", *Demonstration, SIGCHI'95 Companion*, 1995, 23-24.