Hypermedia vs. Paper:
User Strategies in Browsing SNA Materials
TR92-036

December, 1992

Dana Kay Smith
John B. Smith

The University of North Carolina at Chapel Hill
Department of Computer Science
CB#3175, Sitterson Hall
Chapel Hill, NC 27599-3175
919-962-1792
jbs@cs.unc.edu

A TextLab/Collaboratory Report

Portions of this work were supported by the National Science Foundation (Grant #IRI-9015443 and by IBM Corporation (SUR Agreement #866).

Introduction

This study is part of an on-going program of research to develop better tools for collaborative work as well as a better understanding of that complex social and intellectual process. To address the first concern, we are developing a distributed graph storage system and a collection of browsers and application with which to develop and access electronic hypermedia "documents." To address the second concern, we are conducting a series of studies of groups collaborating under a variety of conditions, including working without computers, working with conventional computer and communications resources, and groups working with the systems we are building. A premise of our work is that data stored as hypermedia documents can be used under real-world conditions as effectively as paper documents. Indeed, as the size of the data increase, we anticipate that hypermedia documents may be even more effective for some purposes than this comfortable, time-tested medium.

The specific goal of this study was to answer the question,

> *Can individuals working with a hypermedia system and a hypermedia document of sophisticated technical content locate the information needed to synthesize answers for typical, real-world questions?*

There were two primary motivations behind the study. First, we wanted to add to what is known about the pros and cons of hypermedia vs. paper documents, in general. Second, we wanted to address the more specific concerns of one of our sponsors. For the past several years, we have worked with a group of network architects at IBM to develop a prototype hypermedia system to test the feasibility of using such a system to access and maintain the technical documentation concerning IBM's Systems Network Architecture (SNA), IBM's principal standard for network communications implemented across many of its product lines.

In this report, we first provide background for the study, including a review of related work, a description of the SNA documentation, and a description of the ProtoABC hypermedia system used in the study. Second, we describe the experiment we conducted to address the question listed above and discuss some of its more important results. The report concludes with a brief discussion of possible future work.

Previous research on hypertext vs. paper

Earlier studies comparing hypertext materials to paper materials have yielded mixed results. Marchionini and Shneiderman (1988) report a set of studies that compare paper materials to hypertext materials implemented within their Hyperties system. One study showed that performance with paper materials was faster than performance with the hypertext system. However, another study, also using the Hyperties system, showed that as complexity of a task increased, performance was nearly equal between the two types of materials. Marchionini and Shneiderman also found that regardless of performance, users preferred the Hyperties materials to the paper materials.

Egan, Remde, Gomez, Landauer, Eberhardt, and Lochbaum (1989) compared performance using the "Superbook" system (Remde, Gomez, & Landauer, 1987) to performance using paper materials. They were interested in how performance changes as a function of the type of task required. The results of this study showed that when answering questions that required detailed searching through a text, there was an advantage for the Superbook materials. This advantage was especially pronounced for questions that were not anticipated by the author of a text (i.e. questions whose answers were not referred to in the headings for the text). As in the Hyperties studies, subjective ratings of the Superbook were overwhelmingly positive.

More recent work with Superbook compared the hypertext browser to paper materials and to an electronic document retrieval text indexing system (Egan, Lesk, Ketchum, Lochbaum, Remde, Littman, Landauer, 1991). In this study, each set of materials contained information relevant to basic chemistry. Participants were asked to use the different types of materials to answer typical chemistry questions. The results of this research replicated the results from the Egan et. al. (1989) study. For search and essay questions, the electronic systems led to better performance. When comparing the two electronic versions, Superbook users were more accurate than the users of the document retrieval system when searching for specific information.

Given that these past studies showed that hypertext users perform better for only certain types of tasks, Mynatt, Leventhal, Instone, Farhat, and Rohlman (1992) decided to address this issue in a more systematic way. They examined users performing information-seeking tasks that varied in the amount of information needed to accomplish the task. These tasks were performed using either a hypertext encyclopedia or a paper encyclopedia, both of which contained information about Sherlock Holmes. The results of this study replicated the Superbook research and showed that the hypertext group performed better when searching for information embedded in text. However, the groups did not differ significantly in the time it took them to perform the tasks. Mynatt et. al. argue that these results suggest that with the right support and system functionality, hypertext systems can "equal or surpass conventional books as an information seeking medium (p. 19)."

Overall, the results of these studies suggest that improved hypertext systems can lead to better performance than traditional paper materials. More specifically, prior results show that hypertext performance is better when the task being performed requires (a) searching for information in embedded text or (b) going beyond the text organization imposed by the author of the materials. However, only one of these studies examined the use of hypertext systems in "real world" scholarly activities. That is, the Egan et. al. (1991) study investigated how chemists used a hypertext system to access chemistry literature relevant to their research. We propose that an important part of arguing for the use of hypertext systems is proof that these systems are useful for "real world" problems. The size and complexity of the SNA documentation and the way it is used and maintained by SNA architects makes it an ideal application through which to address this issue.

System Network Architecture

For nearly 20 years, IBM has relied on the same, albeit evolving, architecture as the basis for its communications product line. This architecture, called the System Network Architecture (SNA), is implemented across many different product lines within the company, ranging from mainframe computers to PCs. Currently, there is a small group of highly technical specialists whose primary responsibility it is to maintain the integrity of this architecture under evolving technical conditions and requirements. These people work with other IBM employees across the company who are either representing the various product lines or who are actually designing and programming these product lines. Since SNA is a published standard, they also work with representatives of outside vendors who include SNA-compatible resources in their products. In their interactions with these various individuals, SNA architects are often asked questions about the architecture. For example, designers might ask for interpretations of technical specifications or system builders might ask for changes to be made to accommodate a new product line. To answer these questions, SNA architects consult the SNA documentation, currently maintained in multiple paper manuals that occupy several feet of shelf space. Within these materials, they must locate relevant information, often in several different volumes, and then synthesize this information into a coherent answer for their fellow employees.

Given the current arguments made in support of replacing large sets of paper materials with on-line databases, we, with the support of IBM, were interested in investigating whether or not a hypertext version of the SNA materials could help SNA architects with their day-to-day activities. We built a sample hypermedia document based on an approximately 100-page SNA paper manual, which we then used in an experiment to compare performance using a hypermedia system to performance using the traditional SNA materials. The results of this study not only address the specifics of using a hypertext system to support SNA tasks, but also the more general issue of comparing hypertext and paper materials in "real world" tasks.

The ProtoABC System

The hypermedia system we used in this study was built as a prototype for the Artifact-Based Collaboration (ABC) system. ABC is intended to support collaborative work among a group of distributed users. A primary goal is to allow members of a collaborative project to share an extensive database of material that can be accessed from remote locations via a communications network. Overall, this system has five major components: a distributed graph storage system, a middle level of infrastructure called the Matrix that supports computer conferences, the hypermedia browsers and applications, a video/voice communications component, and tools for recording users' interactions with the system. Only two of these components were included in the ProtoABC system -- the graph server and the hypertext browsers.

ProtoABC allows a user to build, browse and modify a hypermedia data structure that may contain several different types of information. Information is stored in nodes that are connected to one another by links. The structures produced by the node/link relationships can be networks, trees, or lists and special browsers are provided for each type of graph structure. Two additional browsers were built to allow users to work with decomposition diagrams and message flow diagrams included in the SNA material. To browse the data, users either open a node using the menu associated with the node or follow a hyperlink from one node to another. Hyperlinks are different from structural links in that they cut across the formal node/link structures and allow the user to move through the hypermedia graph structure in a non-linear fashion.

Because this is a prototype, we focused our attention on the functionality of the system and developing the specialized tools needed for SNA, rather than performance. As a result, ProtoABC has a slow response time when retrieving nodes that contain complex information. However, performance did not inhibit use of the system by subjects in this experiment.

Information in the System

For purposes of this experiment, we used the general hypermedia data model supported by ProtoABC to construct a specific hypermedia document. That structure was based on materials from an approximately seventy-page chapter on SNA Path Control. It included five types of information: a table of contents, text, figures, index and hyperlinks. Figures 1 and 2 show overviews of the screen layout and the different types of information included.

The Table of Contents shown in the upper right corner of Figure 1 contains the same information as the Table of Contents for the original chapter. It is organized into a tree structure with each node corresponding to a section of the original text All Table of Contents nodes are text nodes. There are three types of text nodes. *Basic* text nodes shown in the upper left corner of Figure 1 contain small sections of prose that range from a single paragraph to four or five paragraphs. Text nodes for topic headings contain all the information written on that topic. *FAPL* text nodes shown in the bottom right corner of Figure 1 contain procedures written in FAPL code, an IBM internal programming language used to describe System Network Architecture procedures. Finally, there are *type definition* text nodes shown in the bottom left corner of Figure 1. They contain definitions for the data structures used in the FAPL code. Within the text nodes, users could use a search function to locate specific words or phrases. However, this search function did not work across nodes and, therefore, did not provide a way to retrieve all nodes that contained a specific reference.

There are two type of figures in the system -- decomposition diagrams and flow diagrams. These figures, shown in Figure 2, are graph structure representations of the figures included in the original chapter. Decomposition diagrams shown on the left side of Figure 2 are tree structures that use a nested box notation to represent parent/child relationships. Flow diagrams shown along the bottom of Figure 2 show the flow of messages among the procedures and processes in Path Control.

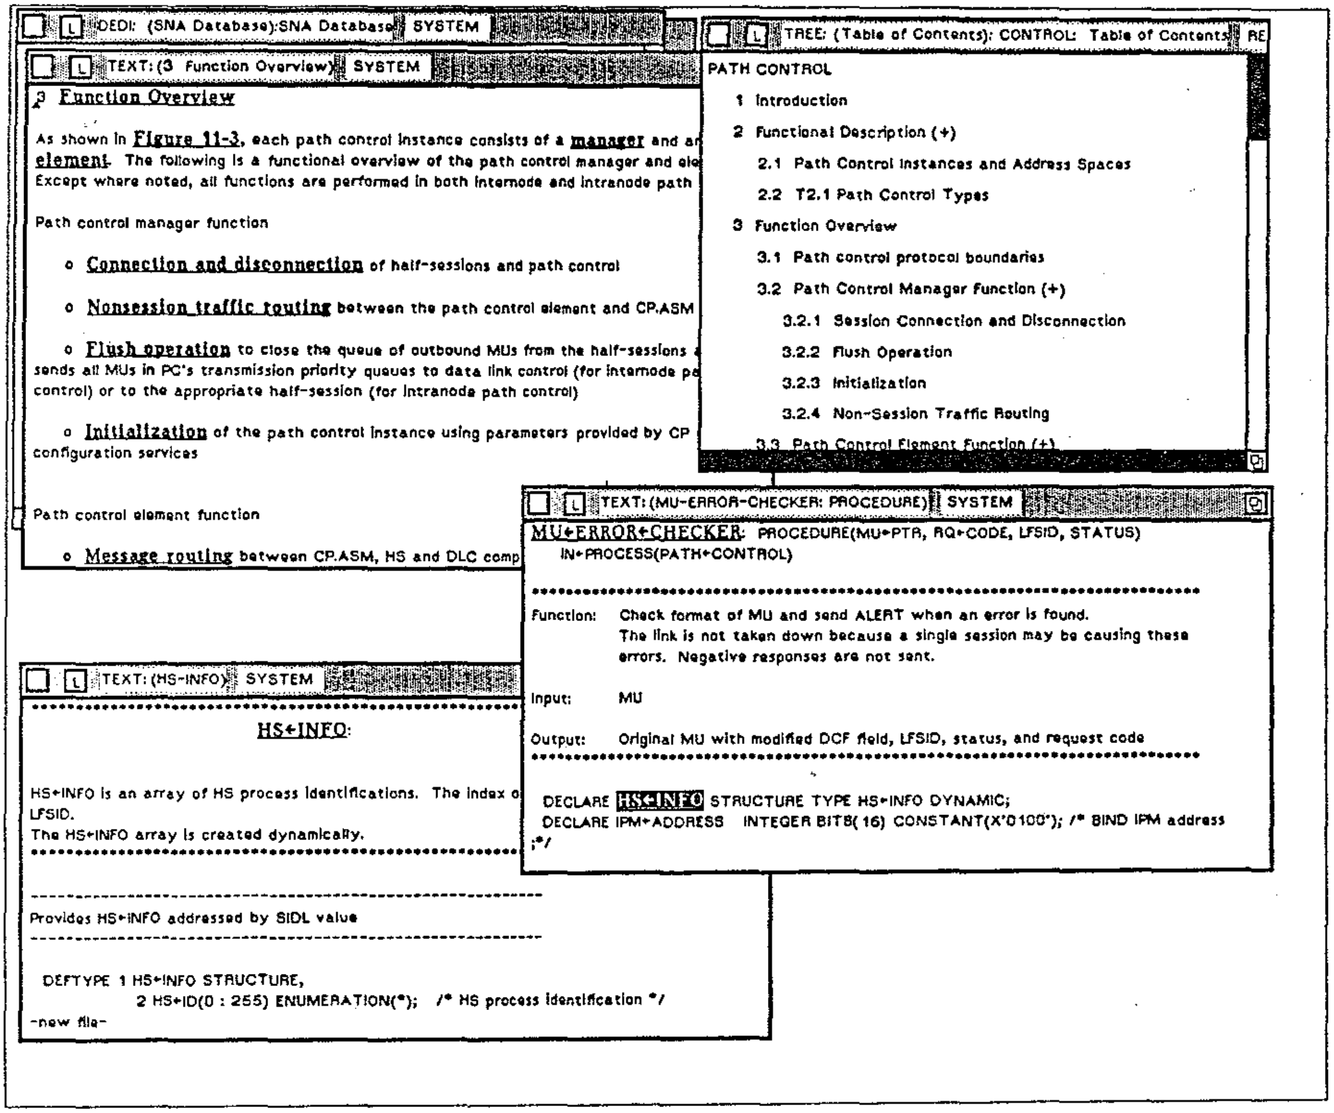

Figure 1: Example screen from ProtoABC system

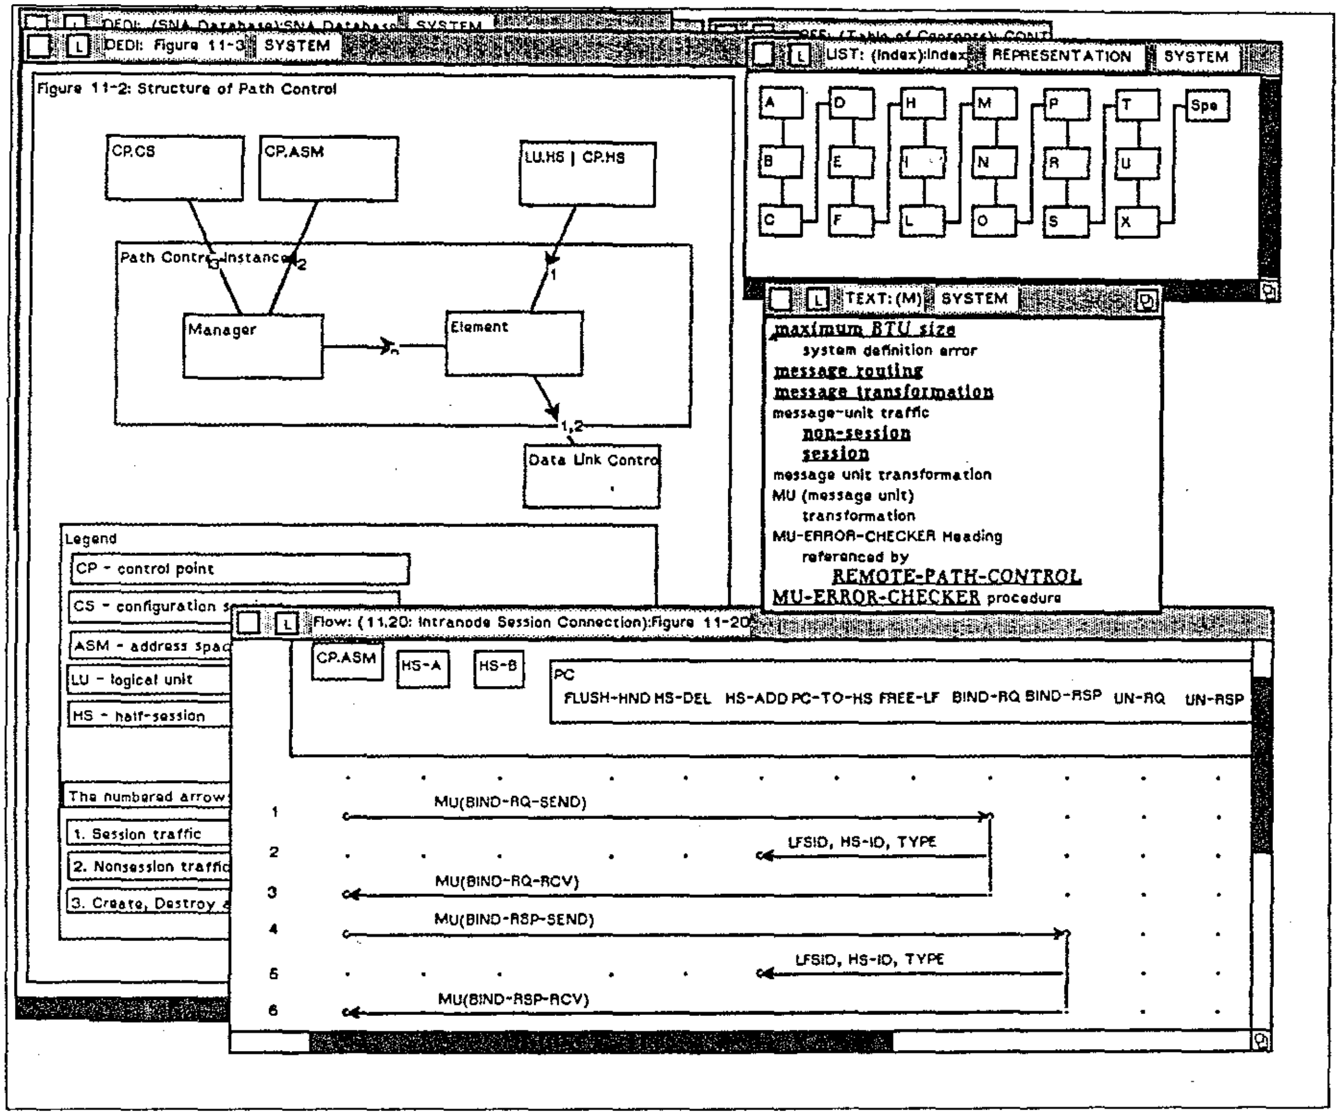

Figure 2: Example screen from SNA hypermedia system.

Another type of information is the Index shown in the upper right corner of Figure 2. The index was stored as a list of entries grouped within nodes corresponding to individual letters of the alphabet. Thus, for example, all the index entries for words beginning with an "m" are contained in the node in the list labeled "M." In Figure 2, the M node has been opened.

The final type of information included in the system is hyperlinks. Hyperlinks are relations that cut across the structure imposed by the structural links (e.g., tree) or connect nodes in two different subgraphs. Using hyperlinks, users can navigate through the database using any strategy that they choose. That is, users can build their own mental structures for moving through the database. There are five different types of hyperlinks that take the user to different types of information. The most basic hyperlink is the *index* hyperlink. Each entry in the index has at least one hyperlink that takes the user to the text where the entry is referenced. Once in the text node, the user can use *reference* hyperlinks to move to additional discussions of topics. Reference hyperlinks are noted by text that is in boldface and underlined. From the text the user can also follow *figure* hyperlinks. These hyperlinks go to the decomposition and flow diagrams in the system. From both the text and the figures, there are *FAPL* hyperlinks that go the description of a specific FAPL procedure that is referenced in the text or figure. Finally, there are *intro* hyperlinks. These hyperlinks provide the user with easy access to introduction for the current section of the document.

Method

Subjects and Design

Six IBM employees participated in this study. All participants were familiar with System Network Architecture, but no participant was formally trained in Path Control, the topic covered in the experimental materials. Two subjects participated in a paper condition in which they used paper manuals to solve problems. Four subjects participated in a hypertext condition in which they used the SNA hypertext system containing the same information as the paper materials, but in an on-line hypermedia format.

Materials

For the paper condition, we wanted to simulate the way that SNA workers currently perform their job. That is, in their everyday activities, these workers use multiple manuals to locate the information required to answer questions asked by colleagues who are using the SNA information to build computer systems. Therefore, we took the IBM chapter on Path Control and broke it into four manuals that contained the same information as presented in the Path Control Chapter. The breakdown was linear and topics within the chapter were not split across manual boundaries.

For the hypertext system, the same IBM chapter on Path Control was used. However, it was divided into individual text pieces and figures and placed into the nodes described above. The SNA system allows users to both browse through the database and to build their own node/link relationships. For this study, we limited participants to the browsing functions and tasks required only browsing activities. In the future, we hope to do other studies in which users may modify this database or build their own database.

Six SNA test problems were used. These problems and their answers were written by SNA experts in Path Control. There were three general types of problems. One type required looking in the text to find the answers to the problem (e.g., Who creates PC?). More complex problems required searching through parts of the text to collect and synthesize the information needed to form the answer (e.g., How does PC determine if something is a BIND?). The most difficult type of problem required searching through the entire text to locate the necessary places in the text and figures where specific changes are needed so that the text would reflect a more general change in the architecture (e.g., What changes would have to be made to the chapter to remove segmentation?). Two of each type of question were used and counterbalanced across subjects.

Procedure

Paper materials: Participants, in this condition, were seated at a table and given the paper manuals. After discussing a few example problems with the experimenter, subjects were then given a workbook that contained the six experimental problems. The problems were randomized within difficulty level for each participant. Participants were given as much time as they needed to solve the problems and experimenters were present to answer any questions. All participants completed the test problems within two hours.

Hypertext system: For this condition, participants were first trained on the hypertext system. This training consisted of a half hour walk-through of the system functions and information types. After training, participants were given the same example problems as in the paper condition. Then participants were given the experimental problems. These problems were the same as in the paper condition and were randomized within difficulty level for each participant. Again participants were given as much time as needed to complete the problems and experimenters were present to answer any questions. All participants completed the test problems within two hours.

Results

For all participants, we recorded their time to complete each exercise, their score on each question, the pieces of information accessed to answer the questions and for the computer participants, we recorded the functions used to access the computer information and had them fill out a questionnaire on how they liked the system. We used these measures to address two general issues -- (1) How did computer performance compare to performance using the paper materials? and (2) How was the hypertext system used and viewed?

Hypertext vs. Paper

Overall problem performance

All participants completed six problems. However, one of these problems proved to be very confusing and, therefore, was eliminated from the analysis.

Table 1 presents the amount of time each participant (hypertext and paper condition) took to answer each question. Due to the small number of participants in this study, we will not present parametric statistical analyses of this data, but rather discuss the mean performance for the two groups with the understanding that these means represent trends and not definitive statistical results. Table 1 also presents individual performance times to illustrate the general lack of variation across subjects and conditions. Mean total time spent on the problems for the hypertext condition was 83 minutes. For the paper condition, the mean time was 85.5 minutes. Thus the average total time for the two conditions was roughly equal. Looking at the performance on each individual question, we see further evidence in support of the equal performance. In four of the five problems, hypertext and paper users took approximately the same amount of time to complete the problems. Although the participants were more familiar with the paper materials and one might predict that this familiarity would lead to better performance, our results show that all participants, regardless of condition, performed the tasks with equal efficiency.

Table 2 shows total score for each participant on each question and the individual scores for each problem and each subject. Problems were scored based upon a comparison of the participant answers with answers that had been previously generated by experts in Path Control in SNA. For each problem, a participant could receive a maximum score of one point for a complete answer. Partial credit was awarded based on the percentage of information that was given when the answers were compared to the expert answers. Table 2 shows that for the hypertext condition, the mean total score was 4.06 out of a possible five points; for the paper condition, the mean total score was 4.08. Thus, as was the case with respect to time, participants using the hypertext were able to perform the tasks as well as the participants using the paper materials. Scores for individual problems also show that performance was roughly equal across the two groups.

In summary, the results show that hypertext users were able to solve typical, real-world SNA problems as well as and as efficiently as those participants who used the original paper materials. These results suggest that although the SNA materials were presented in a new format, users were able to quickly and easily adapt to the new materials. The next analysis delves deeper into this issue and examines the problem solving methods used by the two groups of participants.

SS	Prob#1	Prob#2	Prob#3	Prob#4	Prob#5	Total
system						
SystSS#1	9	14	52	5	23	103
SystSS#2	15	14	31	3	8	71
SystSS#3	12	9	26	3	24	74
SystSS#4	12	18	21	9	24	84
Total	48	55	130	20	79	332
mean	12	13.75	32.5	5	19.75	83
paper						
PPSS#1	8	14	34	6	23	85
PPSS#2	11	31	27	3	14	86
total	19	45	61	9	37	171
mean	9.5	22.5	30.5	4.5	18.5	85.5

Table 1: Time (in minutes) spent on experimental problems.

SS	Prob#1	Prob#2	Prob#3	Prob#4	Prob#5	Total
system						
SystSS#1	1	1	1	1	1	5
SystSS#2	1	0.5	0.67	0.33	1	3.5
SystSS#3	1	0.5	0.5	1	1	4
SystSS#4	1	0.75	0.5	1	0.5	3.75
Total	4	2.75	2.67	3.33	3.5	16.25
mean	1	0.6875	0.6675	0.8325	0.875	4.0625
paper						
PPSS#1	1	0.5	0.5	1	1	4
PPSS#2	1	0.5	1	0.67	1	4.17
total	2	1	1.5	1.67	2	8.17
mean	1	0.5	0.75	0.835	1	4.085

Table 2: Score received on experimental problems.

Problem Solving Methods

The solutions to the different problems required different amounts of effort in navigating through the materials and finding the required pieces of information to answer the questions. We were interested in whether the navigational methods would be different between the two conditions. To address this issue, the solution process for each question was analyzed.

As discussed previously, participants were asked to solve three types of problems. Regardless of the form of the SNA materials, there are at least three strategies that one could use in solving these problems. First, participants can use the *Table of Contents* headings to locate sections of the text that are relevant to the topics in the question. An alternative to using the Table of Contents is to use the *Index* and look up words that are mentioned in the question. Finally, participants can *randomly* move through the materials in hopes of finding a relevant piece of information. In addition to using the Table of Contents and Index to solve these problems, participants using the SNA system could also use the hyperlinking facility. For example, users could choose a FAPL code procedure name from within a figure and follow a hyperlink to the code associated with that feature. The paper participants could get to this same information by looking up the procedure name in the figure and then going to the code in the FAPL procedure manual. For each problem type discussed in the materials section of this paper, we looked at the Paper and Hypertext user strategies and compared the two groups.

The first type of problem required looking through the text and locating a specific piece of information. Our analysis for this type of question showed that the different types of participants used very similar strategies. Most people used the Table of Contents to get to a relevant section where the specific information was located. Two hypertext and one paper participant tried to use the index, but they were not satisfied with the information that it provided and switched to the Table of Contents. (Note: The index used was that found in the original materials.) There was one paper participant who on one problem flipped through the pages of the manual and randomly or opportunistically found the solution to the problem. It is important to note that this behavior is not easy to simulate in a hypertext system and we are currently discussing how we might create a computer-based form of page flipping.

The second type of problem given to the participants required that they put several different pieces of information (text, figures, code) together to solve the problem. Our analysis of user strategies again showed that participants in both conditions had similar behavior. When text was required, participants continued to use the Table of Contents. However, when looking for code, hypertext and paper users went to a figure and then to the code for the necessary procedure. What is most interesting about these results is that the paper participants were in some ways hand-simulating the hyperlink feature implemented in the hypertext system.

The last type of question, required searching through the entire text and noting the places where changes in the materials had to be made. Our analysis of this type of problem showed that both hypertext and paper participants used a combination of the Table of Content and Index strategies. One half of the participants used the Table of Contents for a first pass through the materials and then used the Index to double check. The other half of the participants used the Index to locate initial information and then used the Table of Contents to complete their solution.

Given that the participants in this study were experienced with SNA materials, we assumed that they would bring their typical strategies to the experimental setting. The results of our analysis of their navigational strategies suggest that the participants did use their personal strategies and these strategies were the same, regardless of the form of presentation for the materials. These results suggest that users could easily adapt to the hypertext system and did not have to alter their normal problem solving behavior.

Hypertext System Performance

Hypertext Feature Usage

In addition to comparing performance between the two types of materials, we were also interested in which features of the system were useful and which were not. To address this issue, we focused on the specific operations used by the hypertext users. In this analysis, we looked at whether people were using simple text nodes or more complex figure nodes. We also examined how this information was accessed -- open node action or hyperlink action. The open node actions came from the Table of Contents and the hyperlink actions came from the Index and the within text and figure relations.

Table 3 shows the average percentages for the different types of actions. Not surprisingly, these results show the same trend as in the strategic analysis. The Table of Contents was used almost twice as much as the hyperlinking. This result is easy to explain, since users are much more familiar with using a Table of Contents than with hyperlinking. However, the fact that they used the hyperlinks with any regularity is encouraging and suggests that with time and increased familiarity, this feature might become even more useful.

Within the Table of Contents, text information was used four times as much as figure information. This is not a surprising result given that the ratio of text to figures present in the original materials is almost three to one; therefore, users had more text than figures to choose from. Within the hyperlinking, most of the actions were of three types. The most frequent type of actions were those used to go from the Index to a specific text reference. Users also hyperlinked from one text reference to another text reference and they hyperlinked from a figure to the FAPL code associated with a specific procedure. The Index result is easily accounted for by the use of the Index strategy previously noted in the discussion of navigational strategies. The latter two types of hyperlinks were put in the system to allow users to peruse the materials in a non-linear fashion that provided a direct route to related materials. This result suggests that our intuitions about the benefits of this type of facility were correct and that users had no trouble making use of this new method of information access.

In summary, we found that all aspects of the hypertext system were used by the participants in the hypertext condition. The differences in frequency with which these features were instantiated depends upon (a) previous familiarity with functionality and (b) the amount of information available to the users. Participants had no trouble adapting to the system and even the most novel feature in the system, namely hyperlinking, was used with some degree of frequency.

	Table of Contents		Hyperlinking			
	Text	Figure	Text to Text	Index	Figure to Text	Figure to Figure
	80	20	30	38	23	8
Total	66		34			

Table 3: Percentage of use for computer actions.

Subjective View of the System

Regardless of performance, it is also important that users like a system and find it useful. Therefore, we administered an exit questionnaire to all hypertext users. This questionnaire asked the users for their subjective views on their experience. We then pooled all the answers and gathered the participants for a focus group discussion to get some consensus of opinion. In general, our results show that all hypertext users enjoyed using the system and would like a similar system in their offices.

More specifically, users liked a number of the hypertext features. All users commented on the intuitive feel of the system and how they could easily adapt their personal style to this new form of information. Our empirical results on strategy supports this notion. Users also enjoyed the ease with which they could search through the materials and didn't need to "use their fingers to mark their place." All information could be displayed on the screen at once. Finally, but perhaps most importantly from a hypertext perspective, users found the hyperlinking capability to be very useful. They liked being able to move through the text and figures in a non-linear fashion and felt that this feature might replace the functionality of an index.

When asked what they didn't like about the system, most user comments reflected problems with the computer platform and system performance rather than the actual format of the system. That is, users did not like the three-button, round mouse. (Note: most of these users were more familiar with the two-button, rectangular mouse and use it in their daily activities.) Users were also dissatisfied with the slowness of the system response time. This was due to the implementation of the prototype. (The current ABC system is written in C++ using the Interviews toolkit and has much faster response time.)

Conclusions and Future Directions

The results of this study are extremely promising. We found that even with a slow system that was unfamiliar, users could perform as well as participants using familiar paper materials. We believe that with a better system and more information in the database (i.e. all the SNA materials), performance using the system could become even better than performance using the paper manuals. Thus, it appears that the functionality of a system of this type is useful to SNA architects and might make them more productive.

In addition to raw performance, we also examined how the users adapted to the system and we found that users could easily adapt the system to their existing problem solving strategies that they use daily with traditional SNA materials. We believe using their existing strategies allowed the users to apply a familiar structure to an unfamiliar format and, thus, yielded successful performance. We hypothesize that with increased familiarity with the system, users might develop new strategies and these new strategies might in turn lead to the development of new system functions that better support the user.

In the future, we would like to address more creative aspects of the system. For the existing system, we created all the system hyperlinks. In the next phase, we would like to let users create their own hyperlinks. Previously, we mentioned that some of the system users remarked that the hyperlinks might replace the index. By examining how users create and modify hyperlinks, we could address this issue as well as issues such as using an index to check for completeness of a set of hyperlinks. In addition to creating hyperlinks, users can also create new documents from an existing database. We are interested in how people might do this (e.g., hyperlinks, new graph structures, etc.) and how a system like the ABC system might facilitate this job.

We would like to do a comprehensive study of the evolution of a hypermedia system and the adaptation that grows between system and user. This study would address questions such as (1) How does an electronic document (like the SNA materials) get built from a paper document? (2) How does this electronic document change over time and finally, (3) If we extract a paper document from the evolved electronic document, how does it compare to the original paper document.

As mentioned previously, we believe that the results of the current study give us confidence in the design decisions we made in building a hypermedia system to facilitate the use of complex materials such as those used by SNA architects at IBM. These results and the choices that underlie them are taking us in new and exciting directions.

References

Bush, V. (1945) As we may think. *Atlantic Monthly, 76,* 101-108.

Egan, D.E., Lesk, M.E., Ketchum, R.D., Lochbaum, C.C., Remde, J.R., Littman, M. & Landauer, T.K. (1991). Hypertext for the electronic library? CORE sample results. In *Proceedings of Hypertext '91.* San Antonio, TX, 299-312.

Egan, D.E., Remde, J.R., Gomez, L.M., Landauer, T.K., Eberhardt, J., & Lochbaum, C.C. (1989). Formative design-evaluation of Superbook. *ACM Transactions on Information Systems, Vol. 7, No. 1,* 30-57.

Marchionini, G. & Shneiderman, B. (1989). Finding facts vs. browsing knowledge in hypertext systems. *IEEE Computer, 21,* 70-80.

Mynatt, B.T., Leventhal, L.M., Instone, K., Farhat, J., & Rohlman, D.S., (1992). Hypertext or book: Which is better for answering questions? In *Proceedings of CHI '92* Monterey, CA, 19-25.

Remde, J.R., Gomez, L.M., & Landauer, T.K. (1987). Superbook: an automatic tool for information exploration -- hypertext? In *Proceedings of Hypertext '87.* Chapel Hill, NC, 307-323.