

A Hypertext Writing Environment and its Cognitive Basis*

John B. Smith, Stephen F. Weiss, & Gordon J. Ferguson

Department of Computer Science
University of North Carolina
Chapel Hill, NC 27599-3175
919-962-1792

Abstract

WE is a hypertext writing environment that can be used to create both electronic and printed documents. It is intended for professionals who work within a computer network of professional workstations. Since writing is a complex mental activity that uses many different kinds of thinking, WE was designed in accord with an explicit cognitive model for writing. That model raises several important questions for both electronic and printed documents. The paper includes a discussion of the underlying cognitive model, a description of WE as it currently exists and as it will be extended in the near future, as well as a brief outline of experiments being conducted to evaluate both the model and the system. It concludes by re-examining some of the issues raised by the cognitive model in light of WE, especially the role of constraints in hypertext systems.

1. Introduction

Hypertext is a form of electronic document in which data is stored as a network of nodes connected by links. Nodes can contain text, source code, graphics, audio, video, or other forms of data. Hypertext documents are normally meant to be written, stored, retrieved, and read within a computing environment. Thus, they spend their entire life on-line rather than on paper.

We are building a system that differs from most hypertext systems. It regards the network or directed graph form of information as one (early) stage in the development of a document rather than as its final form. Our system, which we call the Writing Environment or WE for short, helps writers transform loose associative networks of ideas into a hierarchical structure and then write a document in accord with that structure. The product that results can remain in electronic form but it can also be printed out to produce a paper document. Thus, the system can

* Parts of this research were sponsored by The National Science Foundation (Grant IRI-8519517), The Army Research Institute (Contract MDA 903-86-C-0345), and The International Business Machines Corporation (SUR Project 423).

be used both as a conventional hypertext system and as an authoring system with advanced graphical, direct manipulation structure editing capabilities.

Supporting both electronic and conventional paper forms of documents is a key aspect of WE. While electronic documents may eventually replace paper ones, that day is not at hand. Even in organizations in which professionals work within a network of workstations, paper documents continue to be important. Many users prefer to edit on paper rather than on screen. Most internal documents – memoranda, proposals, reports, etc. – must be printed for upper management to read them. And most documents that go outside the organization still go out through the mails or, more likely, Federal Express, than through a network. Thus, in building a system that supports both electronic and printed forms of documents, we have attempted to provide the best of both worlds.

A second major concern of our research group is the relation between WE and the cognitive processes of its users. We are particularly interested in the cognitive strategies writers use to transform information in one form into another. Consequently, a second line of research we are carrying out is a series of experimental studies to first map and then differentiate between the strategies used by expert vs. novice writers and those that lead to effective vs. ineffective documents.

This interest in the relation between cognitive process and system functions is shared with a number of hypertext developers. From the beginning, those working on hypertext systems and concepts have been keenly interested in the relation between thinking and computing. Vannevar Bush called his theoretical system *memex* and saw it as an environment that would enhance the thinking of knowledge workers [Bush, 1945]. Doug Engelbart called the first actual hypertext system *The Augmented Human Intellect System* [Engelbart, 1968]. Another, more recent system is called *Knowledge Garden* [Thompson & Thompson, 1987].

While the hypertext systems that are emerging offer many new opportunities for structuring and using information, they also raise a number of new questions concerning how best to create and use those resources. Much of the work we have done to understand and model the cognitive processes of writers applies equally well to the authors and users of hypertext “documents”. Hypertext authors must still transform inchoate ideas into coherent structures that can be comprehended as well as traversed. Users of hypertext documents must still understand what they read (or see, or hear, . . .) and must construct relations between new information and old, one idea and another. Thus, a second part of our discussion will be a consideration of the cognitive processes that underlie WE and that apply to other hypertext systems, as well.

In describing our work, we will look first at the cognitive basis for WE that includes a Cognitive Framework for written communication. In doing so, we will point out issues that have long-term implications for hypertext systems. Next, we describe WE. Following a brief discussion of some of the experimental studies we

are conducting to evaluate both the cognitive model and the system, our discussion concludes by reconsidering several of the questions raised by the Cognitive Framework in light of the description of WE.

2. Cognitive Basis for WE

In this section, we discuss the cognitive processes involved in writing. We describe those processes, first, as they are used for conventional paper documents and, then, for electronic or hypertext documents.

2.1. Cognitive Modes

Writing is a complex process that draws on many different cognitive skills. Not just translating ideas into words but retrieving information from the writer's memory or from external sources, identifying associative relations among ideas, drawing inferences and making deductions, building larger hierarchical structures, as well as reading, analyzing, and rewording during the editing process. Some writers even report using visual and kinesthetic thinking.

We view these processes as constituents of a set of *cognitive modes*. A mode consists of three components: one or more cognitive processes, a product produced and/or operated on by those processes, and a set of rules that govern the kinds of products that can be produced within the mode and the relations that can exist among the parts of the product(s). Writers use different cognitive modes to produce different forms of information or to transform one intermediate product into another.

For an intuitive sense of modes, consider the following examples. During early exploratory thinking, many writers adopt a mode of thinking in which the primary purpose is to identify ideas and data that may be included in the document and to consider various relations among them. The tenor of exploratory thinking is often relaxed and creative as the individual generates and considers alternative possibilities for the document. However, the mode of thinking used for organizing the content of the document is different. As the writer shifts to building a single integrated structure, he or she is likely to work with more focused attention and a stronger sense of purpose. Writing, *per se*, and editing involve still other modes of thinking.

2.2. Cognitive Framework for written communication

Figure 1 shows the flow of information through the different cognitive modes used for conventional written communication. The model includes both reading and writing. The areas of the "hourglass" denote different cognitive modes. The modes are shown across the top of the figure, the products along the bottom; the tapered areas of the hourglass form, itself, indicate relative differences in the constraints

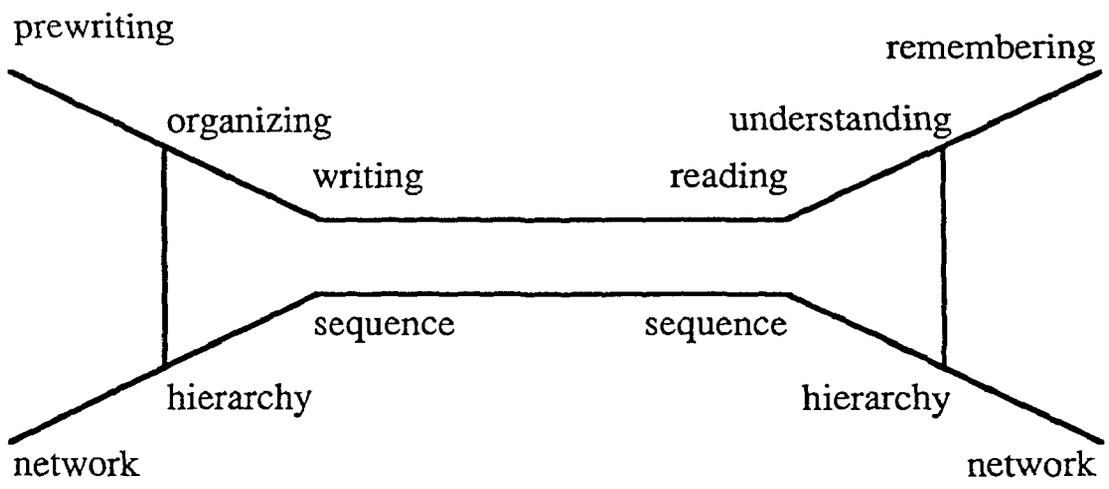


Figure 1
Cognitive Framework for Written Communication

imposed by the different rules for the various modes. (A smaller area implies more constrained options; a larger area, more relational possibilities.) We can now define more precisely the processes, products, and rules for the major modes used for writing.

During the early exploration phase of writing, represented on the left of Figure 1, the writer retrieves potential content from long-term memory or from external sources, considers possible relations among ideas, and, perhaps, groups related ideas and constructs small hierarchical structures. In that mode, the underlying rules are those associated with a network: any idea can be related to any other idea through simple association. Thus, the intermediate product is a network or directed graph of ideas.

Organization is the task of constructing an integrated structure for the document. For many documents, particularly those written by professionals, that structure will be a hierarchy. The product is, thus, a hierarchical structure, and the rules are those that govern hierarchies. That is, each concept or node in the hierarchy can be subordinate to at most one other concept/node, but it may be superordinate to many concepts/nodes. Building such a structure requires a different set of cognitive processes from those used during exploration. The critical one is the process of abstract construction that includes perceiving subordinate/superordinate relations, comparative levels of abstraction, sequencing, proportion, and balance. This mode is shown as smaller than that for exploration since a hierarchy is a restricted form of network and is, thus, more constrained.

Writing, itself, involves still a different set of cognitive processes. Here, the primary task is encoding the abstractions of content and the relations of the hierarchical structure into a sequence of words, drawings, or other explicit forms. The structure of the encoded text is linear and represents a path through the hierarchy. Consequently, it is even more constrained than organization mode.

Editing is not shown in the Figure, but would represent still further constraining of the linear sequence and would include additional reading and analysis, as well as encoding, processes.

Reading, shown on the right half of the Figure, employs an analogous set of processes and forms. Whether the reader reads the document from beginning to end or jumps around from place to place, when that reader settles to read, the text that is read or decoded is a linear sequence of words. The text that is comprehended, however, is a hierarchy. That is, the reader sees that several points do, indeed, add up to the conclusion drawn by the author or that a generalization is supported by the facts or argument cited [Kintsch, 1974; Meyer, 1975; Kintsch & van Dijk, 1978]. What is remembered, though, is that portion of the text hierarchy that is integrated into the network of long-term memory.

Thus, both writing and reading involve a series of transformations in which different cognitive processes transform information in one structural form into a

different structural form. For writing, that dynamic is network to hierarchy to linear sequence. For reading, the dynamic is reversed. (A more through discussion of cognitive modes can be found in [Smith and Lansman, 1987].)

2.3. Implication for printed documents

Central to both writing and reading is the hierarchical form of information. Perhaps that is not surprising since research in reading comprehension has shown scientifically what many writers have known intuitively: that hierarchy is an optimal form for most expository documents. Consequently, features that highlight a document's (hierarchical) structure increase its comprehensibility.

More specifically, thematic titles presented prior to a well-structured text significantly increase free recall of the content of that text [Schwartz & Flammer, 1981]. Within a text, advance organizers – passages containing the main concepts of a text or section of text but at a higher level of abstraction – positively affect comprehension [Ausubel, 1963]. Hierarchical texts in which the structure is signaled or cued are comprehended more effectively than texts in which the structure is not signaled [Meyer, Brandt, & Bluth, 1980]. At the paragraph level, inclusion of a topic- or theme-sentence in the initial position, rather than in an internal position or not at all, results in more accurate comprehension [Kieras, 1980; Williams, Taylor, & Ganger, 1981].

Consequently, for efficient, effective communication, a writing environment should support and encourage development of documents with these characteristics.

2.4. Implication for hypertext documents

This research has significant implications for developers of hypertext systems. Readers of hypertext documents are likely to have problems comprehending what they read similar to those of readers of conventional documents. The same features that facilitate comprehension there are also likely to apply to electronic documents – a well-defined structure that is clearly signaled, advance organizers such as overviews and descriptive titles, and topic statements within individual paragraphs or content units. All of these help the reader develop a high-level understanding of the document's content and purpose that serves as a framework in which to understand and attach its details.

The underlying model for most hypertext systems is a directed graph in which content units are associated with the nodes and the sequences in which the reader may access them determined by the links. However, a *network* of information has properties very different from those of a *hierarchy*. By definition, a hierarchy addresses a single, high-level concept or purpose. Thus, it is well-suited for writers who wish to argue a single point or produce a specific action by their document. A network has no such central thrust. Rather, it is an *environment* in which different

readers may immerse themselves for different purposes and with different expected results. Thus, the emphasis is on the experience of the reader rather than any specific motivation or action. We can easily imagine new forms of entertainment, new literary genres, or even bodies of research materials with directed graph structures. But we cannot foresee purposeful, action-oriented communications in this form. Hierarchical documents, on the other hand, provide the reader with a sense of the whole by including high-level overviews that describe what will follow. Structural information of this sort does not exist in a directed graph. Most network-based hypertext systems have ignored the issue of global structure. Instead, they simply provide for each unit of information the links in and the links out. And most readers quickly get lost in the tangle.

While visual tools for helping readers grasp large graph structures are promising, (see Figures 5 and 6, below), the issues of purpose and focus are inherent. In the section that follows, we will describe the WE system, pointing out as go along how it has addressed these issues. After that, we will return to these same concerns, suggesting a somewhat different perspective of hypertext that may help resolve some of these problems.

3. Description of WE

WE was designed to be congruent with the cognitive theory of writing outlined above. In describing the system, we will first discuss its multimodal design, then the function it provides for moving intermediate products from one system mode to another, and, finally, several special features including a zoom and roam function for searching a large graph structure and controlling the display, WE's interface with an underlying database, and print options. We will also describe WE's hypertext features and several extensions we plan to make to the system in the near future.

3.1. Modes

WE supports each of the major phases of writing in a separate window or system mode. The rules that underline each cognitive mode are reflected in the operations WE supports in the corresponding system mode. WE's structural modes, shown on the left of Figure 2, support representing information units as nodes, moving these nodes from one place to another, and defining relationships among them in the form of directed links. WE's encoding and editing modes, shown on the right side of the screen, only permit manipulation of the content (currently, text) associated with these nodes; structural operations are not allowed in these. A more detailed explanation of each of WE's modes follows.

The user interface provides direct manipulation of visual objects. Objects are selected by pointing with a mouse. Pressing a mouse button provides a pop-up menu specific to the type of object selected. Thus, user operations are organized around a taxonomy of visible object types.

Writing Environment		emptyWS		Work Space		Holding Areas		System	
								10 September 1987	
NETWORK MODE: Net a			View Control		Display/Print		TEXT MODE		
							⌘		
TREE MODE: Tree a			View Control		Display/Print				
							EDIT MODE		
							View Control		
							Display/Print		

Figure 2
WE: Default Screen Layout

3.1.1. Network Mode

Network mode, shown in the upper left quadrant of Figure 2, is intended to support the early exploratory phases of document development. It is also the mode normally used for hypertext. The cognitive processes, described above, for the corresponding cognitive mode are retrieving potential concepts from long-term memory and/or from external sources, representing these concepts in tangible form, clustering them into related groups, defining specific relations or associations between pairs of concepts, and constructing small hierarchical structures.

The system functions for network mode were designed to support the corresponding cognitive processes as directly and as unobtrusively as possible. To represent a concept, the user may point anywhere in the visual space of WE's network mode and select *create node* from the menu. He or she is then prompted for a brief title to label the node. As the set of nodes/ideas grows, the user's cognitive orientation is likely to shift to building small clusters of ideas. The *move* option, selected from the node menu, can be used to gather concepts into spatial groups. To make relations between nodes explicit, the user may link them and give the link a title. The writer who thinks of links as indicating a super/subordinate relation may use this options to build small hierarchies. Figure 3 shows a network constructed by exploring the cognitive concepts related to WE.

As patterns of nodes emerge on the screen, they produce a similar change in the pattern of concepts in the writer's mind. At some point, the author is likely to shift from exploring the ideas and relations inherent in the data and his or her mind to constructing a single, integrated structure for the document. Thus, the writer's intention shifts from possibility to commitment. At this point, he or she may pause to tidy up the exploratory clusters in preparation for moving into a different mode of thinking and working.

3.1.2. Tree Mode

Tree mode helps the user build a single, integrated hierarchical structure for the document. Noting super - and subordinate relations as was done during exploration is frequently almost a reflexive cognitive process, but constructing a large integrated hierarchical structures is not. It requires additional processes. The writer must think on a broader scale, noting relations among not just small groups of concepts, but whole substructures of ideas. He or she must also note parallel relations among similar configurations as well as balance the overall structure. In short, organization is a process of conscious, deliberate construction.

WE represents the hierarchy as a tree, as seen in the lower left quadrant of Figure 2. Figure 4 shows an expanded tree mode. The constraints for tree mode are different from those of network mode. It is no longer possible to create isolated nodes; new nodes can be created only in relation to the tree. To add a node, the writer first selects a node in the tree. He or she can then add a 'child' (subordinate

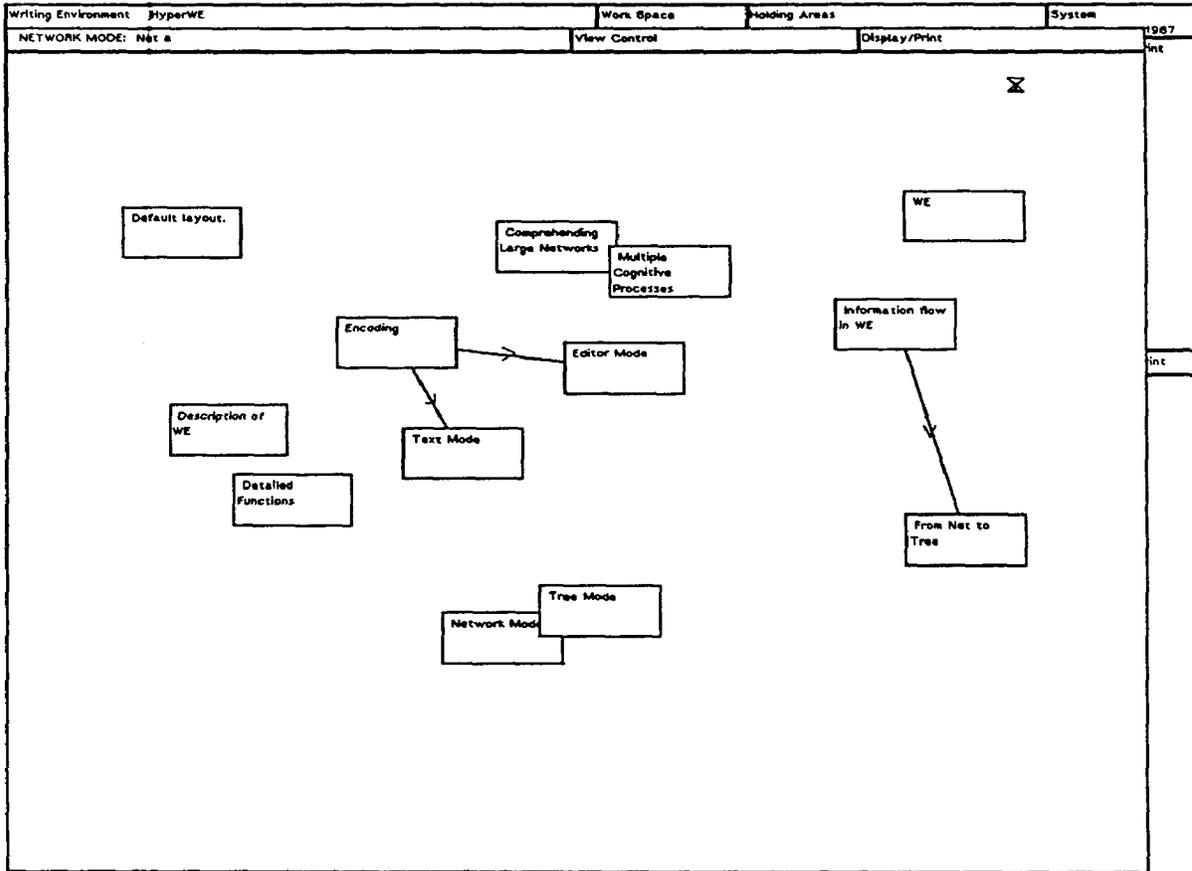


Figure 3
WE: Network Mode

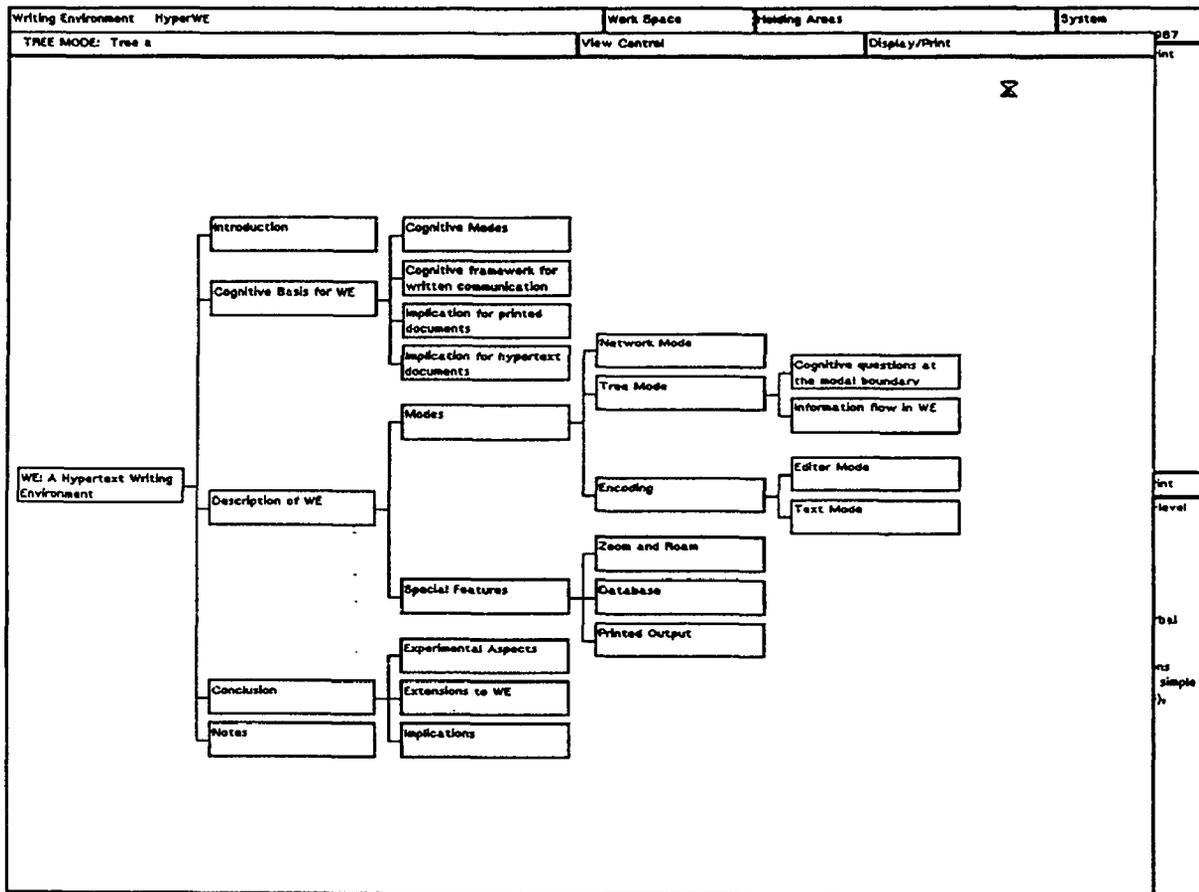


Figure 4
WE: Tree Mode

concept), a 'parent' (superordinate concept), or a 'sibling' (parallel concept). Nodes may be moved from our place to another in the tree. In fact, entire subtrees (a node and all its descendants) may be manipulated: moved, deleted, or made the focus of display. But neither nodes nor branches may be moved out of the tree and remain in tree mode, since the rules of this mode constrain the product to a single hierarchical structure, not a forest.

One of WE's strength is its support of information flow across the modal boundary between network and tree modes. Moving concepts is simply a matter of copying and pasting nodes. Both operations may be done to and from either tree or network mode simply by selecting the node to be moved in one mode and then pointing to the position where it is to be pasted in the other mode. Moving structures of nodes is done the same way. If the structure in Network Mode is hierarchical, the operation is straight forward. If it is not a hierarchy - e.g., a graph containing a cycle - WE transforms the graph into a hierarchy by applying a depth first algorithm that breaks links that cross the hierarchy.

A *subtract tree* operation in network mode provides a form of negative information flow. When a branch of the hierarchy is selected in tree mode, the subtract tree operation removes from the display in network mode all the nodes contained in the branch. Thus, only those ideas/nodes will remain displayed in network mode that have not yet been integrated into the document's hierarchical structure.

3.1.3.1. Editor Mode

Editor mode, shown in the lower right quadrant of Figure 2, provides access to a standard text editor. It is used to encode the concept represented by a node into text. In the current system, that editor is the Smalltalk text editor. In future extensions of WE, the system will support additional text editors as well as editors for other kinds of data, such as graphics, sound, and video. At that time, the editor invoked will be keyed to the data type of the particular node.

To begin writing, the user points to the node in either tree or network mode and selects the edit option on the menu. This transfers control to editor mode. Text may then be keyed in, deleted, and so on. The user leaves editor mode simply by moving the cursor from that area of the screen into any of the other mode windows.

3.1.4. Text Mode

In text mode, shown in the upper right quadrant of Figure 2, the document is presented in linear form much as it would appear in printed form. Text mode constructs a representation of the continuous document by stepping through the tree - from top to bottom, left to right - inserting node labels as section headings, followed by the blocks of text associated with the nodes. A scroll bar permits the writer to move forward and backwards through the document as a whole (the path through the tree). A long-term goal is to make the representation identical to final

formatted output. Currently, text mode provides three editing regions within its window. As the tree is traversed using the scroll bar, the blocks of text associated with the various nodes are moved into the three areas of the window.

Within each area, a second scroll bar permits the user to move through the text for the individual node displayed there. Thus, by scrolling to the bottom of one section and the top of the following section, the writer can see how the text for the two nodes fits together. The writer can edit the text for each node using the Smalltalk editor, just as in editor mode. Text can also be moved from one area/node to another, and the section headings (node label) can be edited, as well. However, the node itself can't be deleted or edited structurally from text mode. This can be done only from tree mode.

3.2. Special Features

WE provides several additional featuring that are not, strictly speaking, part of the writing process. These include a zoom and roam option for managing a group of nodes too large to fit on a single screen, an interface for a supporting database, and options for printed output.

3.2.1. Zoom and Roam

Navigation through the two dimensional space of computer displays has typically involved some form of scroll bar. Unfortunately, these do not present any overview of the space being explored. WE uses a different technique, called roaming, that was originally developed by other members of our research group [Beard & Walker, 1987]. The user can invoke the roam and zoom display from either network or tree modes: the system will then display in a pop-up window a very small representation of the entire graph or hierarchical space with the area of the current display indicated by a box, (see Figure 5). This box can be directly manipulated to change the scale or position of what is then displayed in the mode window. Figure 5 shows a stretch box and Figure 6 shows the resized network mode that was produced as a result.

3.2.2. Database

WE is intended to be used in conjunction with an object oriented database system in which all structural information is stored. To support this interface, WE uses low-level data objects that correspond with database operations. These objects are currently of three types: structures, nodes, and links. Structures are typed, named sets of links (and, by implication, associated nodes). The type indicates whether the structure is a graph, hierarchy, or path. This information is used by the system to determine the operations that can be performed on the particular structure. Each node is also viewed as a typed object. Associated with it are various attributes that identify the type of content "within" the node and, thus, bind it to

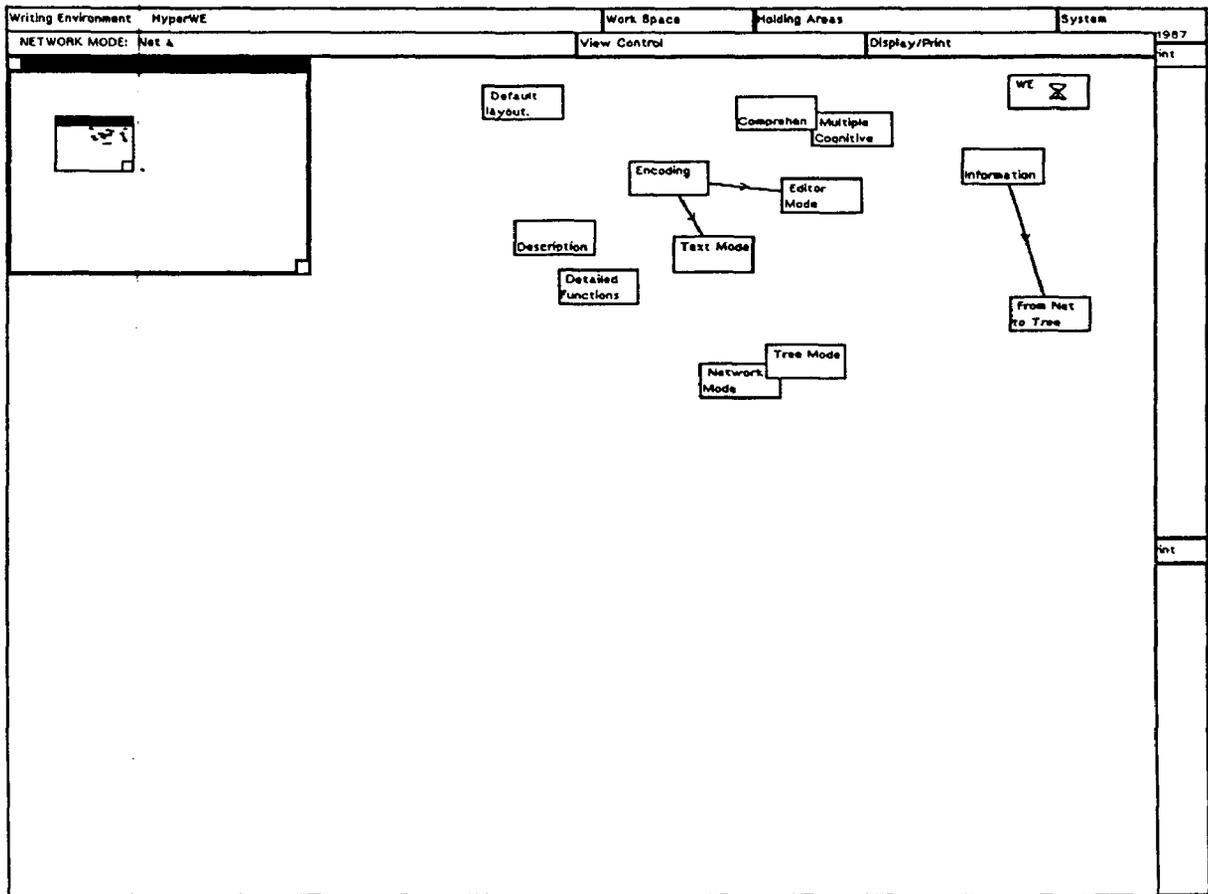


Figure 5

WE: Roam and Zoom View of Network Mode

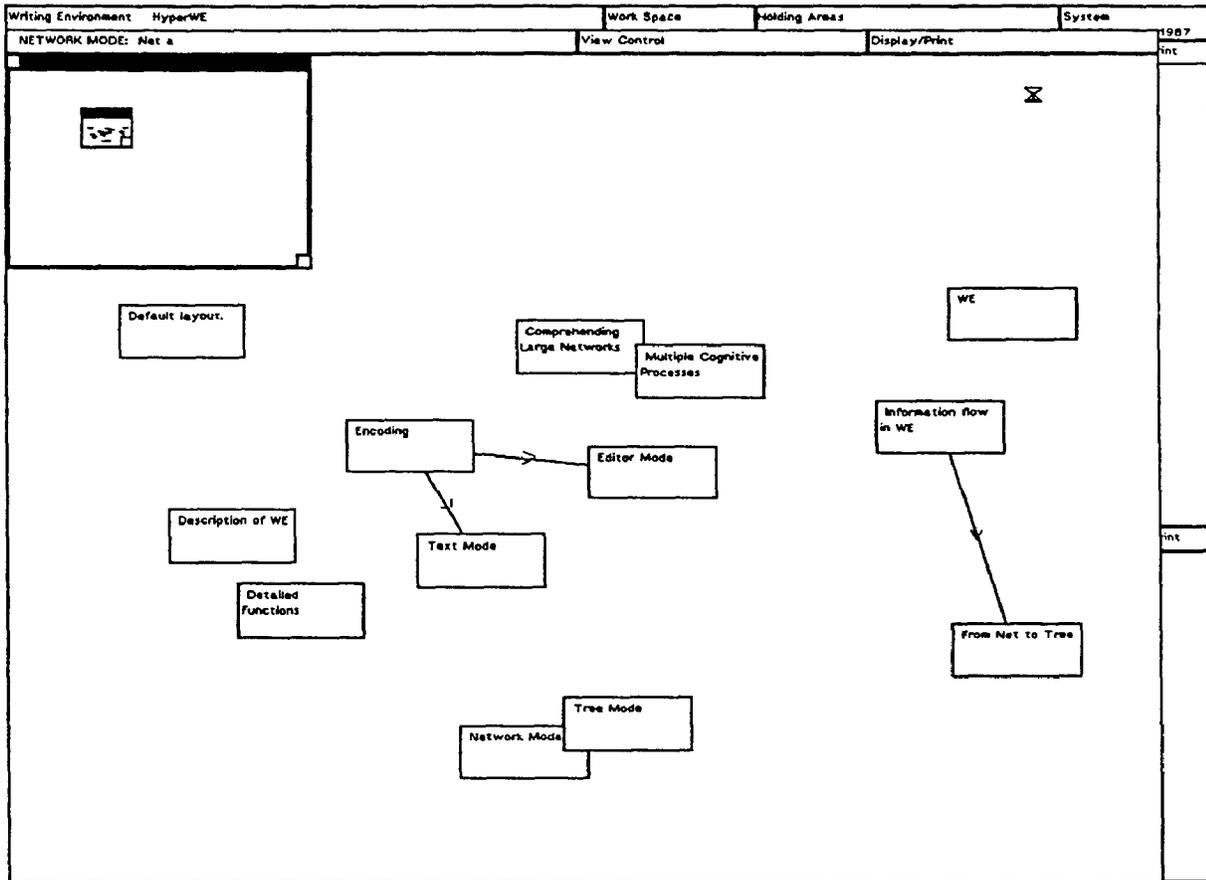


Figure 6

WE: Resized View of Network

a particular editor/display program; its spatial dimensions in network mode space; and both its associative and hierarchical links. Links are attributed pairs of node identifiers that define the directed arc. Attributes indicate the structure of which the link is a part and additional system information.

Currently, the database is confined to a single document, but we will extend its definition to permit teams and departments to store collections of documents and other kinds of data. Thus, future users will be able to search the database for information relevant to a current project. Once a usable node or structure is found, it can be imported into the environment and included in the structure being developed.

3.2.3. Printed Output

WE produces output for a laser printer, although actual formatting is done by TeX using commands inserted into the text by WE. The detailed mechanics of the printing process are, of course, installation dependent. However, two levels of intervention are available to users. First, the TeX files, themselves, can be saved and modified as necessary. Some sections of this document were prepared using WE, others with a conventional editor; the two groups were integrated in this way. Second, the TeX macros that format the headings, select the fonts and spacing between sections, and so on, are stored in a dictionary and can be changed by the users.

3.3. Extensions to WE

WE is an evolutionary system. We will continue to enhance it, providing additional functions and additional capability with respect to the size and number of documents. High on the list of priorities are functions to help users manage multiple drafts of single documents. Most of these extensions will be implemented in terms of the database design when WE is merged with an underlying object-oriented database system.

A more fundamental enhancement will allow WE to handle distributed writing. There are many (perhaps most) projects that are too big for a single person. The next major step in the development of the Writing Environment will be to support collaborative writing. A group of writers, possibly widely separated geographically, will be able to work together to produce a single product. Each user will be able to see a single shared workspace and will be able to manipulate the workspace under a managed collaborative paradigm.

A longer term goal is to merge another system we are developing – MICROAR-RAS [Smith, Weiss, & Ferguson, 1986], an advanced full-text retrieval and analysis system – with WE. This will provide fast, flexible content-based searches as well as other analytic functions.

While we have chosen to characterize the system in terms of writing, it is actually a more general tool. It is useful in many other information management applications where the major steps are:

- conceptualization – the generating of ideas;
- organization – imposing structure on those ideas;
- modification – refining the ideas and structures;
- and linearization – defining linear paths through the structure.

Such applications include designing a building, planning a logistical operation, or writing a large computer program. In each case, the process begins with the creation of a graph structure or hypertext of content units. The nodes represent the individual components of the operation, and the links represent the dependencies. Implementing the operations requires that the hypertext be linearized, for example, along a single time line. WE will provide tools to develop additional modes tailored to particular applications.

4. Experimental Studies

4.1. Protocol Tracker and Cognitive Grammar

In addition to serving as a tool for writing, WE can also be used to observe how people write. We have implemented an on-line tracker that captures a user's interactions with the system. That information is represented as a sequence of symbols, with attributes, that constitutes a history of the session. We have also built a replay function that permits us to replay the session in time proportional to the original session, in uniform time, and in manually controlled steps.

We are developing a more powerful tool for analyzing these data. To be comprehensible, the low level data must be transformed into symbols that are more general and more indicative of the user's strategy. That is, users think in terms of high level conceptual phrases but they enact those phrases as a sequence of lower level operations. For example, the conceptual phrase might be to create a cluster of ideas. This is accomplished in WE by creating a set of nodes, labeling them appropriately, moving them near one another, and, perhaps, linking them together. We are developing a cognitive grammar by which low level operations can be mapped to a small set of conceptual phrases and higher level constructs. The parse trees produced by the grammar will provide insight into the user's overall writing strategy. By comparing the strategies of different classes of users (for example, expert technical writers vs. novice writers), we hope to develop more effective and efficient writing methodologies and tools.

5. Conclusion

We conclude by returning to some of the questions and issues raised at the beginning of this paper. While the processes of reading and writing conventional documents have been studied in considerable detail, we still have a very limited understanding of the cognitive processes and strategies that produce effective information transfer. We know even less about such communications for electronic documents. A major line of research that should go hand in hand with the development of hypertext and other electronic document systems is formal, controlled experimental studies of users' interactions with these systems followed by actual-use studies to confirm results. We are committed to this approach as an integral part of the development method for WE; we know of at least one other research group (Xerox PARC) that, we believe, shares this concern. But this is a large and complex area of inquiry that will require additional researchers as well.

In many respects, hypertext is a state of mind. It has been described frequently as a tool to enhance the user's mental abilities, as an environment in which to think, etc. It is essential, however, to remember that human beings don't exist in only one state of mind. We use multiple cognitive modes for different intellectual tasks and purposes. (Figure 1 showed the organization of those modes for written communication.) But, as suggested above, hypertext in its fundamental form – a directed graph of information components – is most consistent with one particular mode of thinking – exploration. Exploratory thinking usually occurs early in the development of a set of ideas. Such thinking is an integral part of the overall cognitive process not just for writing but for many forms of productive, professional work. But it is an end in itself for only certain situations. Great for an aesthetic experience – James Joyce, or more likely, John Fowles, would have loved it as a literary medium. Great for an undirected, free-flowing learning experience, analogous to spending an evening browsing through an encyclopedia. But as a tool for professionals, hypertext, we believe, will become a supporting utility over which more constrained applications will be developed rather than the primary application system, itself. To be truly effective, hypertext applications must match additional power with additional control and structure. In the long term, constraints may turn out to be more important than raw power.

Looking further into the future to a time when large distributed databases of hypertext documents will exist, we don't see (or don't want to see) a flat, hyperplane of spaghetti. Rather, we believe that out of that hyperplane will emerge peaks of understanding and purpose created by professionals using powerful new tools. These peaks will be criss-crossed, to be sure, by multiple paths and relations, sometimes visible, sometimes not. But each peak will be supported by a single, integral hierarchical structure.

We share the enthusiasm for hypertext that is growing daily. But we hope that trail-blazers will think about *where* they are going in addition to *how* to get there. And that those that follow them will do so for purpose as well as for possibility.

6. Acknowledgments

A number of individuals and organizations have contributed to the work described here. We wish to thank our sponsors for both their financial support and the advice and encouragement provided by their program officers. These include The National Science Foundation, The Army Research Institute, and The IBM Corporation. We also wish to thank our faculty colleagues, Profs. Marcy Lansman (Psychology) and Jay Bolter (Classics), for their contributions to the ideas described here. We also wish to thank the following graduate students who have helped to develop WE: Paulette Bush, Yen-Ping Shan, Irene Jenkins (Psychology), Valerie Kierulf, and Greg Berg (Psychology).

7. Notes

Ausubel, D. P. (1963). *The Psychology of Meaningful Verbal Learning*. New York: Grune & Stratton.

Beard, D. V. & Walker, J. Q. (1987). *Navigational techniques to improve the display of large two-dimensional spaces*. Chapel Hill, NC: UNC Department of Computer Science Technical Report 87-031.

Bush, V. (1945). As we may think. *Atlantic Monthly*, 176(1), 101-108.

Engelbart, D. & English, W. (1968). A research center for augmenting human intellect. *Proceedings of 1968 FJCC*. Montvale, NJ: AFIPS Press, pp. 395-410.

Kieras, D. E. (1980). Initial mention as a signal to thematic content in technical passages. *Memory and Cognition*, 8(4), 345-353.

Kintsch, W. (1974). *The representation of meaning in memory*. Hillsdale, NJ: Erlbaum Associates.

Kintsch, W. & van Dijk, T. A. (1978). Toward a model of text comprehension and production. *Psychological Review*, 85, 363-394.

Meyer, G. J. F. (1975). *The organization of prose and its effects on memory*. Amsterdam: North Holland Publishing Company.

Meyer, G. J. F., Brandt, D. M., & Bluth, G. J. (1980). Use of top-level structure in text: key for reading comprehension of ninth grade students. *Reading Research Quarterly*, 1, 72-103.

Schwartz, M. N. K. & Flammer, A. (1981). Text structure and title-effects on comprehension and recall. *Journal of Verbal Learning and Verbal Behavior*, 20, 61-66.

Smith, J. B. & Lansman, M. (1987). *A theoretical basis for a computer writing environment*. Chapel Hill, NC: UNC Department of Computer Science Technical Report 87-032.

Smith, J. B., Weiss, S. F., & Ferguson, G. J. (1986). *MICROARRAS: An overview*. Chapel Hill, NC: UNC Department of Computer Science Technical Report 86-017.

Thompson, B. & Thompson, B. (1987). KnowledgePro. Software distributed by Knowledge Garden, Nassau, NY.

Williams, J. P., Taylor, M. B., & Ganger, S. (1981). Text variations at the level of the individual sentence and the comprehension of simple expository paragraphs. *Journal of Educational Psychology*, 73(6), 851-865.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© 1989 ACM 089791-340-X/89/0011/0214 \$1.50