Using the WRITING ENVIRONMENT to Study Writers' Strategies

Marcy Lansman, John B. Smith, Irene Weber

Writers rave about their word processors. Some are so dependent on the computer that writing with paper and pencil has become an intolerably slow and laborious process. But word processors are limited. They're good for composing and editing text, but they're not so good for brainstorming and exploring the relationships between ideas (Bridwell-Bowles, Johnson, & Brehe, 1987; Lutz, 1987). And it's hard to get a sense of the shape of the whole text when you see it only one screenful at a time (Haas & Hayes, 1986). To expand the capabilities of the computer as a tool for writers, we developed an experimental computer program called the WRITING ENVIRONMENT (Smith, J. B., Weiss, Ferguson, Bolter, Lansman, & Beard, 1987). The program was designed to help nonfiction writers create well-structured texts.

In designing the WRITING ENVIRONMENT, we assumed that writing involves several different kinds of thinking, which we have called *cognitive modes* (Smith, J. B. & Lansman, 1989, 1991). We also assumed that these different kinds of thinking should be supported by different *system modes* on the computer. Consequently, the program has four system modes, which appear in four windows on the computer screen. Each mode allows the user to view the text in a different way and to perform different operations on it.

In "Network Mode," writers represent their ideas as separate *nodes*, which are shown as labeled boxes on the screen. They move these boxes around and link them to indicate the relationships between ideas. In "Tree Mode," they use their idea nodes to develop an organizational structure for their texts. In "Edit Mode," they open individual nodes and write the text that is to be associated with those nodes. In "Text Mode," they view and revise text from adjacent nodes in a single window.

Perhaps the most important feature of the WRITING ENVIRONMENT is that it allows researchers like ourselves to study how people use the program. The WRITING ENVIRONMENT records users' activities as they work, creating computer-generated protocols that can be analyzed to determine writers' strategies. These protocols provide an alternative to think-aloud protocols (Hayes & Flower, 1980; Swarts, Flower, & Hayes, 1984) or videotapes (Gould, 1980, 1981; Matsuhashi, 1982) as a source of information about the writing process.

In this paper, we first describe in more detail the WRITING ENVIRONMENT and the protocols it creates. Then, to illustrate the capabilities of the WRITING ENVIRONMENT as a research tool, we discuss a study in which we asked technical writers and graduate students to use the program to write technical reports and then analyzed the protocols to discover their strategies.

The WRITING ENVIRONMENT

The WRITING ENVIRONMENT is implemented on UNIX workstations with 19-inch, high-resolution screens. The four system modes are shown in four windows on the computer screen. The user chooses a particular mode by moving a mouse-controlled cursor into one of the windows. Within each mode, the user interacts with the program by selecting options from a series of pop-up menus (Smith et al., 1987).

Figure 1 shows the computer screen as it might look in the middle of a writing session. In the figure, the four modes are laid out in four quadrants on the screen, but users can resize the windows as they like. For example, they can expand one mode to fill the entire screen when they are working in that mode, and then shrink it back to its original size when they are through.

"Network Mode," shown in the upper left quadrant of Figure 1, was designed to support writers as they generate the ideas to be included in their texts and as they explore the relationships between those ideas. The ideas appear as labeled rectangular nodes on the screen. The nodes can be dragged around the screen with the mouse and can be linked using directional arrows. In "Network Mode," users

may place nodes anywhere in the window and may link any two nodes, regardless of the form of the resulting structure. The program places no constraints on the arrangement of nodes or the links between them.

"Tree Mode," shown in the lower left quadrant, was designed to help writers develop an organizational structure for their texts and edit that structure. As in "Network Mode," users can generate and link idea nodes. But in "Tree Mode" there are constraints as to the form of the resulting structure. In "Tree Mode," the structure must be hierarchical—as in an outline. Furthermore, the form of the structure is dictated by the program. "Tree Mode" represents the organizational structure of the writer's text as a right-branching tree like the one shown in Figure 1. Nodes, branches, and entire trees can be moved from "Network Mode" to "Tree Mode."

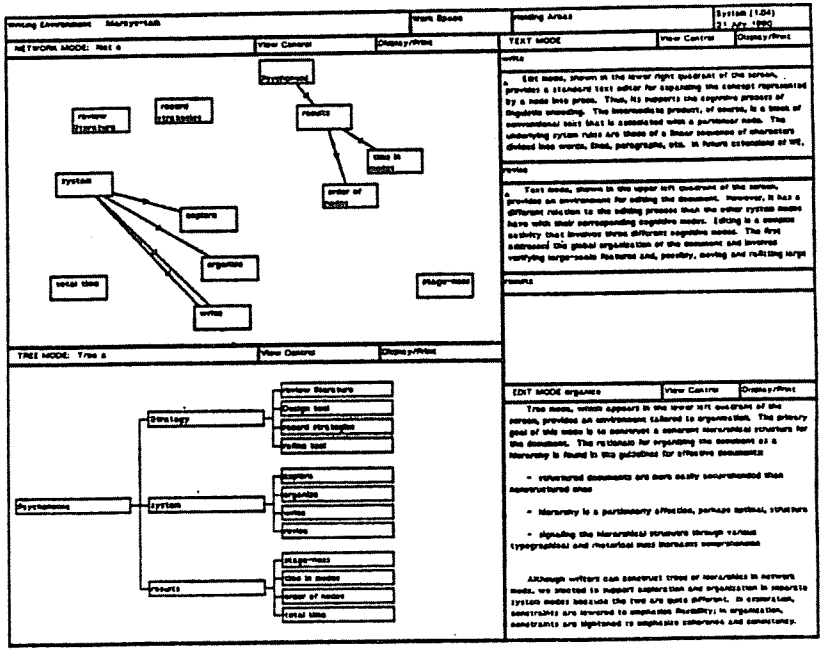

Figure 1. The computer screen might look something like this while a writer is using the WRITING ENVIRONMENT to compose a document. "Network Mode" is shown in the upper left, "Tree Mode" in the lower left, "Edit Mode" in the lower right, and "Text Mode" in the upper right. The user may enlarge any of the modes to fill the screen.

"Edit Mode," shown in the lower right quadrant, enables writers to write the text that will be attached to the nodes in the organizational structure. If a node is moved from one position in the structure to another, the text attached to that node moves with it. In order to write text in "Edit Mode," users first open up a node in either "Network Mode" or "Tree Mode" and then compose the text for that node in the "Edit Mode" window. Text is written using a rudimentary text editor in which the cursor is controlled by the mouse. Using this text editor, writers may insert, delete, and change text, as well as cut and paste it. At present, however, the program does not have the formatting capabilities of commercially available word-processing programs.

"Text Mode," shown in the upper right quadrant, allows writers to examine and revise the transitions between paragraphs and sections. In "Text Mode," the user can scroll through and edit the text attached to a sequence of nodes. The order of the sequence is dictated by the structure indicated in "Tree Mode."

An important difference between "Edit Mode" and "Text Mode" is that in "Edit Mode," the writer can create and edit the text attached to only one node at a time. In "Text Mode," the writer can view and edit the text attached to several nodes and can move text back and forth between nodes.

Thus, the four system modes support several different types of thinking: "Network Mode" supports brainstorming or exploratory thinking, "Tree Mode" supports planning the structure of the document and structural editing or reorganizing, "Edit Mode" supports text generating and text editing, and "Text Mode" supports editing section and paragraph transitions for coherence. In this first version of the system, we consciously decided to omit modes for rhetorical, situational, or audience analysis. However, heuristics that support these types of thinking have been described in J. B. Smith and C. F. Smith (1987) and may be included in future versions or subsequent systems.

There are several important differences between the WRITING ENVIRONMENT and more conventional outline processors, such as the one provided by Microsoft WORD. Perhaps the most important difference is that the WRITING ENVIRONMENT provides graphical representations of the relationships between ideas and allows writers to directly manipulate those representations in two dimensional space. Dragging idea nodes around the screen with a mouse and linking them gives writers more freedom to explore relationships than moving headings up and down in an outline. Another important difference between the WRITING ENVIRONMENT and outline processors is that in the WRITING ENVIRONMENT, planning the structure of a text and writing that text are carried out in

spatially separate locations, allowing writers to view the structure as they write the text.

The WRITING ENVIRONMENT has some of the characteristics of a hypertext system such as HYPERCARD because it represents a document as a set of independently manipulable blocks of text. Again, however, there are important differences. In the WRITING ENVIRONMENT, the structure created in "Tree Mode" dictates a single, linear ordering for the blocks of text. Given this linear ordering, the WRITING ENVIRONMENT enables writers to view transitions between blocks of text, to compile the blocks into a single document, and to print the document as a whole. HYPERCARD does not provide these capabilities.

Computer-Generated Protocols

As writers use the WRITING ENVIRONMENT, the computer records their activities. The initial result is a computer-generated protocol, called an *Action Transcript*. The Action Transcript shows writers' activities at a detailed level; for example, it records each time writers open a menu, select an option from the menu (e.g., "Create Node"), label a node, and so forth.

The Action Transcript makes it possible to replay writing sessions (Smith, J. B., Smith, D. K., & Kupstas, 1991). During replay, the researcher can see on the screen the same displays that the writers saw as they wrote. Replaying a session shows when and where on the screen writers created, deleted, and linked nodes in "Network Mode" and "Tree Mode" and how they labeled those nodes. It also shows when writers opened nodes to write text. So far, however, it does not show the text that was written or revised. Thus, as it stands, the replay function provides detailed information on how writers developed the structure for their texts but less complete information on how they composed and revised text. In future versions, we hope to expand the replay function so that it will provide more detailed information on creation and revision of text.

A computerized grammar (Smith, Rooks, & Ferguson, 1989) condenses the Action Transcript into a less detailed list of activities called the *Operation Transcript*. For example, in the Operation Transcript, the three actions listed above—raising a menu, selecting the "Create Node" option, and labeling the node—would be condensed into a single operation. In the study reported below, the Operation Transcript provided the raw data for the analyses. (The grammar further refines the Operation Transcript into more condensed lists of activities. However, because only the Operation Transcript was used in the analysis of this study, these additional levels will not be further described here.)

A Study of Writers' Strategies

The purpose of our initial study was to discover how writers used the WRITING ENVIRONMENT to develop the structure for a technical document. In this study, we asked graduate students and technical writers to use the system to write reports from source materials we provided.

The study focused on several dimensions of writers' activities. First of all, we wanted to know whether writers would indeed be able to use this new tool to develop structures for their reports. We also wanted to know what proportion of their time they would spend on these structures, how that time would be distributed across the writing sessions, and what kind of strategies they would use. Finally, we wanted to know whether individual differences in strategies would be associated with variations in the quality of the final reports.

Method

Participants

There were two groups of participants. The first group consisted of 9 experienced technical writers, 8 men and 1 woman from a large computer company. Their ages ranged from 26 to 69 with a mean age of 46. They had worked as technical writers for 3 to 30 years, with a mean of 10 years. The second group consisted of 9 graduate students, 6 men and 3 women, in the social sciences and humanities at the University of North Carolina, Chapel Hill. Their ages ranged from 24 to 31 with a mean age of 28. All participants in the study had previous experience using word-processing packages.

Source Materials

All the participants used the same five articles as source material for the reports that they wrote during the experiment. The topic of the articles was the manufacture and application of shape-memory alloys, a group of metal alloys that resume their original shape when heated. The topic was chosen to be unfamiliar yet understandable to all the participants.

Tutorial

A written tutorial (Jenkins, Lansman, & Smith, 1989) instructed users in how to use the four modes of the system and guided them through a series of exercises. The tutorial did not suggest strategies for writing a document with the WRITING ENVIRONMENT.

Procedure

Participants came to the lab for two half-day sessions. During the first session, they went through the tutorial individually and familiarized themselves with the WRITING ENVIRONMENT. It took them about 1.5 hours to learn to use the system. During the first session, participants also read the source articles on shape-memory alloys.

During the second session, participants used the WRITING ENVIRONMENT to write a summary report based on the source articles. Source materials and highlighters were available while participants were writing, but they were not allowed to use paper and pencil. They were told that they could take as much of the half-day session as they needed to write their reports.

Evaluation

Two judges evaluated each report. Each of the judges was a writer with a master's degree in English, and each had taught courses in technical writing. They judged the overall quality of the reports on a 5-point scale.

Results and Discussion

Participants' Comments

At the end of the second session, writers were asked to comment in writing on what they liked and disliked about the WRITING ENVIRONMENT. Responses indicated that participants adapted easily to the multimodal nature of the system and to the fact that different system modes appeared in different windows on the screen. They particularly liked the fact that they could see the organizational structures of their papers in a separate window as they wrote text. As one user wrote, "The multiple window display is the most useful feature. The ability to see the organization of the document while editing a node is unique."

Users were enthusiastic about the structure-planning modes of the system. They felt that the spatial representation helped them develop the structure of their reports: "The ability to group & regroup ideas is tremendous, especially when these groups can be visually seen. I think the longer the document the more valuable this feature can be." Some thought that "Network Mode" and "Tree Mode" were redundant, but they did not agree on which one should be eliminated.

On the negative side, some participants were frustrated by the "Edit Mode" and "Text Mode." They wanted to see the whole text contiguously, not always divided into separate nodes:

> Also, not being able to view the whole text w/out the node headings was a little disabling in terms of judging the continuity of what was being written. When I was nearly finished I wanted to be able to see the document as it will look on paper but was unable to do so.

In general, participants appreciated the novel characteristics of the structure-planning modes—"Network" and "Tree"—but missed some of the editing capabilities of their accustomed word-processing packages.

Reports

Each of the subjects but one wrote a report using the WRITING ENVIRONMENT. One technical writer left the study at the beginning of the second session explaining that she could not write under lab conditions. Reports ranged in length from 450 words to 1735 words with a mean of 1051 words ($SD = 332$). They varied in quality from well organized and polished to quite rough.

Many reports contained an unusually large number of headings. When writers use the WRITING ENVIRONMENT, they create text by opening up a node in either "Network Mode" or "Tree Mode" and composing sentences that will be attached to that node. When the document is printed, the node label appears as a heading that precedes the associated text. Each node in the organizational structure becomes a heading in the final document whether or not any text has been written for it. In these reports, the number of headings ranged from 4 to 46, with a mean of 16.2 headings ($SD = 9.8$). Thus there were, on the average, only about 65 words (six lines) of text associated with each heading. For many headings, there was no text at all, indicating that the writer had created a node but failed to write text for it.

The fact that the number of headings was so large relative to the length of the final reports suggests that writers devoted a large amount of effort to planning the organizational structures of their documents relative to the amount of effort they spent composing and revising text.

Evaluations of the Reports

The two judges were moderately consistent in their ratings of overall quality, $r = .67$. According to this criterion, the graduate students wrote significantly better reports than the professional writers. On a 1- to 5-point scale, the quality ratings of the graduate students ranged from 1.5 to 4.5 with a mean of 3.39 ($SD = 1.05$). The quality ratings for the technical writers ranged from 1.0 to 3.0 with a mean of 2.06 ($SD = .62$). This difference was statistically significant, $t = 3.1$, $df = 15$, $p < .01$.

We had expected that the professional writers, with their years of experience writing technical documents, would do better on this task. We cannot explain why they did not. They were older and may have had a harder time adapting to a new computer tool. They may have been less motivated than the students to work on an assignment that was unrelated to their work. Or they may have differed from the students in cognitive ability.

Strategies for Developing Structures

Order of Creation Versus Order of Appearance

If writers used the WRITING ENVIRONMENT for exploration (i.e., to generate ideas freely and then to try out various organizational schemes), then the order in which ideas were generated might be quite dissimilar from the order in which they appeared in the final text. On the other hand, if writers organized their papers in their heads before they began to use the computers, then order of generation might be quite similar to order of appearance in the text.

Our impression, based on replaying the sessions, was that both the order and the superordinate-subordinate relationships among nodes were often roughly worked out in writers' heads before they were recorded on the computer screen. In order to look at this issue more quantitatively, we computed for each participant the correlation between the order in which nodes were created and the order in which they appeared in the final report. The correlations ranged from –.16 to .99 with a mean of .48 ($SD = .34$). The fact that all but two of the correlations were positive and that the mean was fairly high indicates that the order in which participants generated topic headings was, in most cases, far from random.

There were no differences between graduate students and professional writers on this measure, nor was there a significant correlation between it and quality of the report.

Top-Down Versus Bottom-Up Generation of Ideas

One method of creating a hierarchical structure for a document is to write down superordinate topics first, followed by the subordinate topics that go under them. We will call this a *top-down* strategy. It is probably the most common strategy among writers creating an outline with paper and pencil or with standard outline-processing packages because writers generally work from top to bottom and superordinate topics always appear above their subordinates in a standard outline. Another method is to generate subordinate topics first and then to

group these subordinate topics and create appropriate superordinate topic headings for each group. We will call this a *bottom-up* strategy. Although the bottom-up strategy gives the writer more flexibility in grouping subordinate topics under superordinate topic headings, the strategy is quite clumsy to carry out using word- or outline-processing packages. (Some of us have resorted to moving yellow Post-It notes around on our desks.) Using the WRITING ENVIRONMENT, a writer can easily use either a top-down or bottom-up strategy.

As a group, our writers were much more likely to follow a top-down strategy. As a quantitative measure of *top-downness*, we used the percentage of nodes (omitting the root node) that were generated *after* their superordinate node. This measure ranged from .5 to 1.0 for individual participants, with a mean of .80 (SD = .15).

Technical writers adhered more closely than graduate students to a top-down strategy. The mean top-down score for technical writers was .86 (SD = .14) and the mean for graduate students was .74 (SD = .15). The difference was marginally significant, $t = 1.83$, $df = 15$, $p < .10$.

Furthermore, those participants who deviated more from a strictly top-down strategy tended to write higher quality reports, as evaluated by our judges. There was a significant negative correlation between top-down score and quality of the report, $r = -.54$, $p < .05$.

Distribution of Time

Total time spent writing the reports ranged from 2 hours 10 minutes to 4 hours with a mean of 3 hours 11 minutes (SD = 35 minutes). This total did not include time writers spent reviewing the source articles or mentally planning their reports before they began to plan or write using the computer.

Time Per Mode

Table 1 shows how participants distributed their time among the four system modes. Included are the mean and range for time spent in each mode, percentage of time spent in each mode, and number of episodes in each mode. (An episode was defined as a series of operations on one system mode uninterrupted by visits to other modes.)

The mean times spent in "Network Mode" and "Tree Mode," 30 minutes and 22 minutes respectively, indicate that participants spent considerable time building the organizational structures of their reports. The mean number of episodes in the various modes indicates that participants moved often among the modes rather than using them in a strictly sequential manner.

Table 1
Mean and range of time spent in each mode, percentage of total time spent in each mode, and number of episodes in each mode

Mean time (min)	30	22	79	59
Range	6-68	10-45	0-166	0-161
Standard deviation	16	9	58	56
Mean percent of total time	16%	12%	41%	31%
Range	3-37	6-34	0-82	0-82
Standard deviation	8	6	28	28
Mean number of episodes	12	18	14	9
Range	1-34	8-35	0-32	0-17
Standard deviation	9	8	9	5

"Edit Mode" was intended to be used for working within individual nodes, and "Text Mode" was intended for coherence editing across nodes. The range of values in Table 1 indicates that some writers spent no time in "Edit Mode" and others spent no time in "Text Mode." These writers preferred to do all their writing and revising in either "Edit Mode" or "Text Mode."

Division of Time Between Planning Structures and Writing/Revising

We can estimate the amount of time writers spent planning the structures of their reports by adding their times in "Network Mode" and times in "Tree Mode." Similarly, we can estimate the amount of time they spent writing and revising by adding their times in "Edit Mode" and times in "Text Mode." A structure-planning episode was defined as a series of episodes in "Network Mode" or "Tree Mode" that was uninterrupted by episodes in either "Edit Mode" or "Text Mode." Similarly, a writing/revising episode was defined as a series of episodes in "Edit Mode" and/or "Text Mode" uninterrupted by episodes in "Network Mode" or "Tree Mode." Table 2 summarizes the distribution of time between structure-planning and writing/revising for all 17 participants.

Planning the structure for a text is a very specific kind of planning. Other kinds of planning include working out the wording of individual sentences, considering the intended audience, deciding how much time

Table 2
Mean and range of time spent structure-planning and writing/revising, percentage of total time spent structure-planning and writing/revising, and number of episodes spent structure-planning and writing/revising

	Structure-planning	*Writing/Revising*
Mean time (min)	53	138
Range	22-88	63-194
Standard deviation	19	35
Mean percent of total time	28%	72%
Range	11-54	46-89
Standard deviation	11	11
Mean number of episodes	19	19
Range	6-46	6-45
Standard deviation	10	10

to allot for various sections, and so forth. Given the fairly narrow range of activities we considered under our definition of *structure-planning*, it is noteworthy that writers spent quite a large proportion of their time on it: an average of 53 minutes out of 191 minutes, or 28% of their time.

Other researchers have used different methods to study planning. Kellogg and Mueller (1989), for example, estimated the percentage of time their college student writers spent planning by interrupting them at random intervals and asking them to classify their current activity as planning, translating, reviewing, or other. By this method, they estimated that their participants spent about 25% of their time planning. Gould (1980, 1981) studied videotapes of professional researchers writing short business letters and classified all pauses when participants were not actually composing text as planning. By this method, he estimated that his participants spent two thirds of their time planning. Selzer (1983), on the basis of interviews, estimated that his single participant, a professional engineer, spent 80% of his writing time "inventing and arranging" the ideas he would include in his written texts. Because the definitions of planning used by these researchers and by our project are all so different, it is difficult to draw meaningful comparisons.

Distribution of Structure-Planning Time

According to a traditional *stages* model of composition, writing should take place in three sequential stages: planning, writing, and revising. In this model, then, planning occurs (or at least *should* occur) before writing and revising begin. On the basis of think-aloud protocols, Flower and Hayes (1981a) have concluded that planning, writing, and revising processes are called up recursively rather than serially. By studying how our participants moved back and forth between the structure-planning modes ("Network" and "Tree") and the writing/revising modes ("Edit" and "Text"), we can find out whether their strategies were consistent with a strict stage model of composition.

The data from four individual writers are shown graphically in Figure 2, panels A through D. Each of the four panels in the figure represents an individual subject. In each panel, time is represented on the horizontal axis, and time elapsed in the session is shown by the small numbers along the bottom. The broken horizontal line labeled *plan* represents the time the writer spent planning the structure of the text. The line labeled *write* represents the time the writer spent writing and revising the text. Each vertical tick shows the beginning of a planning or writing episode. The horizontal line attached to the tick represents the duration of that episode. For example, the writer shown in Figure 2, panel a spent the first third of the session planning the structure of his text in three long episodes, and the last two thirds of the session writing and revising the text, again in three long episodes.

Panels A and B of Figure 2 show participants who conformed quite closely to a stage model. They did almost all of their planning before they began to write. Panels C and D show participants who used a quite different strategy: They alternated between planning and writing throughout the session. If we consider these two pairs of participants as defining the two ends of a continuum, the participants in our study were spread out across the entire range of the continuum with no noticeable clusters. Only a few could be said to have conformed closely to a stage model. The others all deviated to various degrees from this pattern. Thus there was tremendous variation in the extent to which participants planned the structure of their texts before they wrote.

In order to measure more rigorously subjects' conformity to a stage model of writing (i.e., planning structure first, followed by writing/revising), we developed a *stage index*, which was computed for each subject. This index was designed to assess the extent to which structure-planning time preceded writing/revising time. In order to understand the *stage index*, imagine computing for every minute of writing time the proportion of total planning time that preceded that minute of writing.

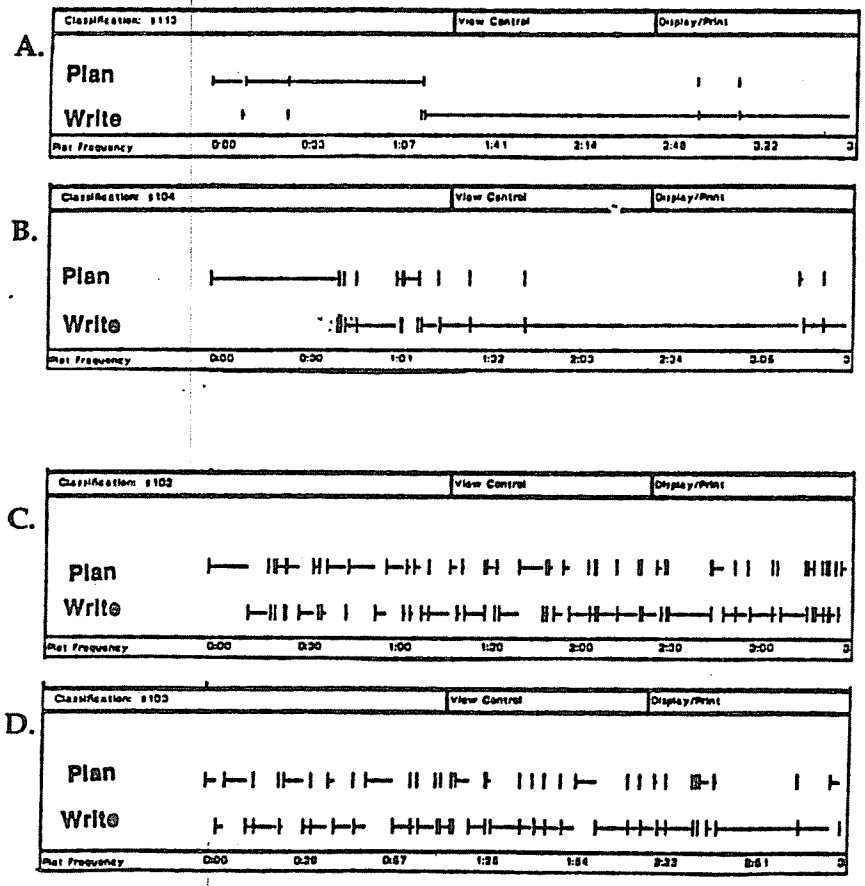

Figure 2. Each panel of this figure shows how an individual subject distributed his or her time between planning ("Network Mode" and "Tree Mode") and writing ("Edit Mode" and "Text Mode"). Time since the beginning of the session is represented on the horizontal axis. Each vertical tick represents the beginning of a planning or writing episode. The length of the horizontal line attached to the tick represents the duration of the episode. Panels A and B represent writers who did almost all of their planning before they began to write. Panels C and D represent writers who alternated often between planning and writing.

In order to compute the *stage index*, these proportions are averaged across all the minutes of a writing session. To take the simplest example, if a subject completed all planning before beginning to write, then for each minute of writing, the proportion of planning that preceded that minute would be 1.0 and the average, the *stage index*, would be 1.0. The index can vary from almost 0 to 1.0. (It can't be zero because subjects using the WRITING ENVIRONMENT must create at least one node in either "Network Mode" or "Tree Mode" before beginning to write.) The *stage index* for each of the subjects in Figure 2, panels A, B, C, and D was .95, .98, .72, and .67, respectively. For the group as a whole, it ranged from a minimum of .58 to a maximum of .98, with a mean of .78 (SD = .11). Although these writers did not, as a group, conform very well to a stage model of writing, they did concentrate the majority of their planning time before their writing time, as one would expect.

Group Differences

There were no statistically significant differences between the graduate students and the professional writers on any of the measures related to distribution of time. There was, however, a tendency for the technical writers to spend a greater percent of their time planning structure (33%) than did the graduate students (24%), $t = 1.74$, $df = 15$, $p = .103$. There were no differences between the two groups in the *stage index*.

Correlations with Quality

Although there was wide variation in the total time participants spent on their reports, there was no relationship between the total time spent and overall quality, $r = -.02$. Participants who spent more time did not necessarily write better reports.

Nor did writers who planned their structures before they wrote (as many composition teachers advise) write better reports. Writers who distributed their structure-planning time throughout the session were just as likely to write good reports as those who concentrated their planning at the beginning of the session. In fact, the correlation between the overall quality and the *stage index* (which indicates the extent to which participants planned before they wrote) was slightly negative, $r = -.17$.

There was, however, a significant *negative* relationship between the quality of the reports and the time spent planning structure. Writers who spent *less* time planning tended to write better reports. Both total time spent planning and percentage of time spent planning were

negatively correlated with overall quality, $r = -.44$ and $-.48$ respectively, $p < .05$. The same general pattern held when planning time was broken down into time spent in "Network Mode" and "Tree Mode."

In light of the many claims that planning is good and that good writers plan more (Bereiter & Scardamalia, 1987; Flower & Hayes, 1980, 1981b; Hayes, 1989; Kellogg, 1986, 1987, in press), the negative relationship between structure-planning time and quality was very surprising. We had expected to find that those writers who spent more time planning would write better reports.

As always, it is important to note that correlation does not imply causality. It is not necessarily true that spending more time in "Network Mode" and "Tree Mode" *caused* writers to write lower quality reports. It could be, for example, that ability was a third, causal variable. For instance, low ability may have caused some writers to have a hard time learning to use "Network Mode" and "Tree Mode" and thus to spend more time in these modes, and low ability may have caused these same writers to write lower quality reports.

Nevertheless, replay of the writing sessions caused us to speculate about other possible explanations for the negative relationship between time spent in the structure-planning modes and overall quality of the reports. Some of our participants appear to have spent more time experimenting with the novel capabilities of the WRITING ENVIRONMENT, especially "Network Mode" and "Tree Mode," than was practical, given that they were required to complete their reports in a single writing session. Figure 3, which shows the organizational structure of one report at two stages in its development, illustrates an extreme case. This writer experimented with the idea of using nodes to represent individual readers of the report. (The experimental instructions specified that the report was intended for a management team including two MBAs, an accountant, an electrical engineer, and the owner.) In the end, he gave up the idea and reverted to a more standard organizational structure. But the experimentation had taken precious time away from other activities.

A number of writers seemed to have created structures that were overly elaborate for the reports they wrote. As one writer commented,

> I should have made a simple hierarchy in the tree mode and not made so many nodes. It got to a point where I would have less than a sentence in each node. Better to have much larger chunks in each node.

There seems to be a tendency for people to spend more time than they otherwise would doing what a writing tool makes easy, especially

A.

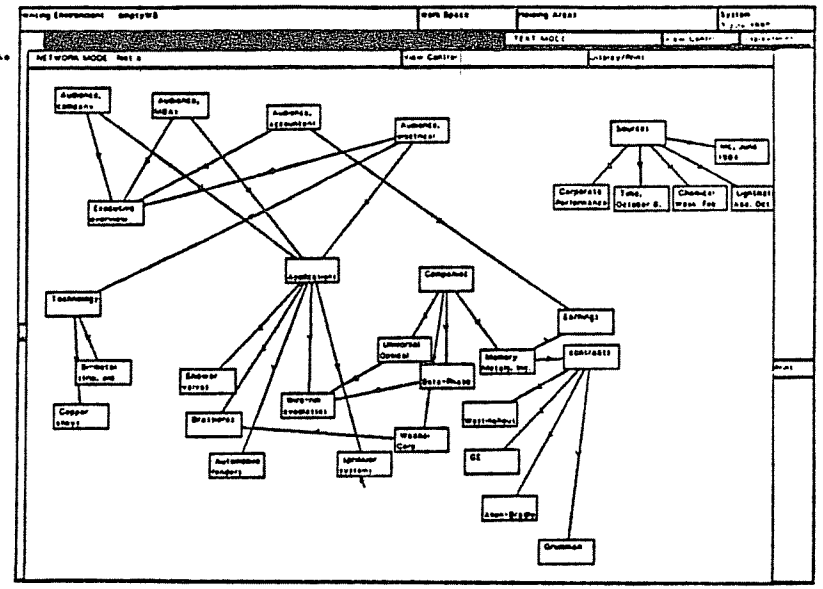

B.

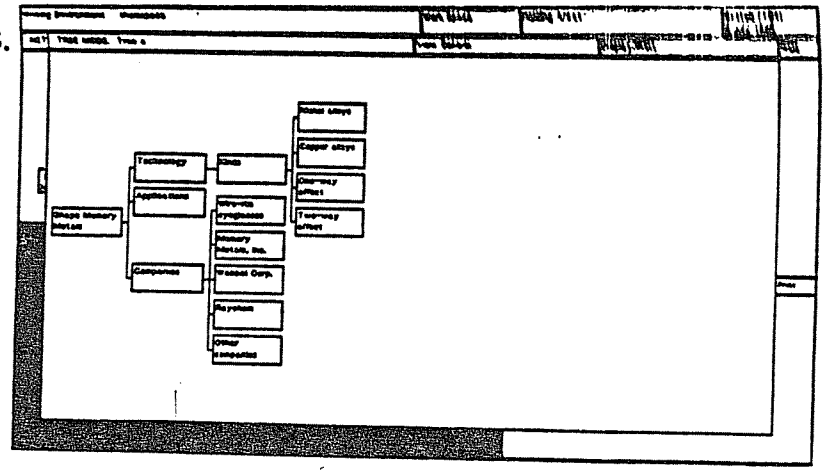

Figure 3.
The two panels of the figure represent two phases in the development of a single subject's organizational structure. In panel A, the subject was working in "Network Mode" experimenting with the idea of using nodes to represent the readers as well as the content of the report. Panel B shows the final structure of the report in "Tree Mode." The structure in B is much simpler, and the *reader nodes* of panel A were omitted.

when they are novice users. For example, some researchers have found that writers spend more time on sentence-level editing when they are using a word-processing package than when they are using paper and pencil (Collier, 1983; Gould, 1981; Haas, 1989; Lutz, 1987). This finding, in conjunction with our own, suggests that as writers learn to use the new capabilities of a writing tool, they must learn to let their writing strategies be dictated by their writing goals rather than by the capabilities of the computer program they are using.

In interpreting the findings of the study reported here, remember that our participants had only 90 minutes of practice using the WRITING ENVIRONMENT. Part of the time they spent using the program, especially time spent in "Network Mode" and "Tree Mode," was undoubtedly spent learning the program. As we do studies with more experienced users, we may find that the negative relationship between planning time and quality disappears or is reversed. In a more recent study, in which each participant wrote two reports using the WRITING ENVIRONMENT, both the amount of time spent planning and the number of nodes in the final structure decreased from the first to the second report.

Keep in mind also that the WRITING ENVIRONMENT constrains users to write in a way that may be inconsistent with their accustomed strategies. For example, some writers freewrite before they begin to build an organizational structure for their texts (Elbow, 1981; Lansman, 1991). A freewriting strategy is certainly possible within the WRITING ENVIRONMENT, but the system was not designed to support such a strategy. The constraints imposed by the WRITING ENVIRONMENT may have caused some writers to produce lower quality documents than they otherwise would have.

Some Comments on Computer-Generated Protocols

Many empirical studies of writing research fall into two categories: studies that examine the thought processes of individual writers and studies that use standard statistical techniques to generalize over groups of writers (Strong, 1985). Representative of the first category is Emig's (1971) study of eight gifted high school writers, in which the strategies of only one student were given extensive, detailed analysis. Representative of the second category are Kellogg's (1987, in press) randomized, controlled studies of college students who were instructed either to outline or to begin writing immediately.

While case studies may give the researcher a good feel for the conscious cognitive processes of individual writers, it is difficult to generalize from these studies. Analysis of think-aloud protocols is so labor intensive that researchers rarely report detailed data from more

than a few participants (Hagge, 1987). Although the informal insights derived from these protocols have changed the shape of the field, it is impossible to judge the validity of their more objective claims, such as the claim that expert writers elaborate the rhetorical problem more extensively than novices (Flower & Hayes, 1980).

On the other hand, quantitative studies, such as those of Kellogg, often leave the reader wondering whether, in averaging over groups of participants, the researcher has lost track of the often idiosyncratic nature of individual writers' strategies.

Computer-generated protocols, such as those produced by the WRITING ENVIRONMENT, offer a way to bridge the gap between the case study and quantitative approaches. These protocols support both detailed study of individual writers' strategies *and* quantitative analyses of groups of writers. Watching the replay of individual writing sessions has given us a sense of the struggles individual writers go through as they try to shape their ideas into coherent organizational structures. Replay of individual sessions also reveals episodes during which the writer was struggling not with ideas, but with the computer system (e.g., trying and failing to resize the screen so that both the organizational tree and the text are visible). Through replays we gain an intimate view of the participant, not only as a writer, but as a computer user.

But computer-generated protocols also tell the other side of the story: They allow us to summarize the data from groups of participants and thus to judge the generalizability of our observations. Because the protocol is readable by the computer, we have been able to write computer programs that condense a detailed list of user actions into a higher level summary of user operations, classify and count operations, calculate the time spent on various types of operations, and summarize all these types of data over groups of participants.

In this particular study, replay of individual writing sessions and quantitative analysis of summary data led to a coherent picture of user strategies. We discovered that writers had indeed learned to use the WRITING ENVIRONMENT during the brief training session, and that, given the task of summarizing a set of source materials, they used the planning modes of the system extensively to organize their reports. In fact, both replays of individual writing sessions and quantitative summaries indicated that at least some writers devoted too much of their effort to planning the organizational structures of their reports and too little to writing and revising the text.

Only further studies will reveal whether the findings reported here will generalize to experienced users writing longer texts.

Marcy Lansman taught in the Department of Computer Science at the University of North Carolina (UNC) at Chapel Hill until she left the project in the fall of 1992. She is now working on a book concerning a family's experience with a serious illness.

John B. Smith teaches in the Department of Computer Science at UNC Chapel Hill, where he is the director of Text Lab.

Irene Weber is a graduate student in the Department of Psychology at UNC Chapel Hill and a Research Associate in the Department of Computer Science.

Note

This research was supported by Army Research Institute Contract #MDA 903-86-C-0345 and the National Science Foundation Grant #IRI-8519517. We would like to thank Gordon J. Ferguson, who was the principal technical designer of the WRITING ENVIRONMENT and provided assistance with the system throughout the experiment and data analysis. We are also grateful to Matthew Barkley, Paulette Bush, Karen Harrower, Mark Rooks, and Yen-Ping Shan for technical support, as well as Gregg Berg for assembling materials and instructions.

The first author was a visitor in the Department of Technical Communication at the University of Washington in Seattle while writing this paper.

Requests for reprints should be sent to John B. Smith, Department of Computer Science, CB #3175, University of North Carolina, Chapel Hill, NC 27599-3175; or by e-mail at <jbs@cs.unc.edu>

References

Bereiter, C., & Scardamalia, M. (1987). *The psychology of written composition.* Hillsdale, NJ: Erlbaum.

Bridwell-Bowles, L., Johnson, P., & Brehe, S. (1987). Composing and computers: Case studies of experienced writers. In A. Matsuhashi (Ed.), *Writing in real time* (pp. 81-107). Norwood, NJ: Ablex.

Collier, R. M. (1983). The word processor and revision strategies. *College Composition and Communication, 34,* 149-155.

Elbow, P. (1981). *Writing with power.* New York: Oxford University Press.

Emig, J. (1971). *The composing processes of twelfth graders.* (Research Report No. 13). Urbana, IL: National Council of Teachers of English.

Flower, L., & Hayes, J. R. (1980). The cognition of discovery: Defining a rhetorical problem. *College Composition and Communication, 31,* 21-32.

Flower, L., & Hayes, J. R. (1981a). A cognitive process theory of writing. *College Composition and Communication, 32,* 365-387.

Flower, L., & Hayes, J. R. (1981b). Plans that guide the composing process. In C. H. Frederiksen & J. F. Dominic (Eds.), *Writing: The nature, development and teaching of written communication.* (pp. 39-58). Hillsdale, NJ: Erlbaum.

Gould, J. D. (1980). Experiments on composing letters: Some facts, some myths, and some observations. In L. W. Gregg & E. R. Steinberg (Eds.), *Cognitive processes in writing* (pp. 97-127). Hillsdale, NJ: Erlbaum.

Gould, J. D. (1981). Composing letters with computer-based text editors. *Human Factors, 23,* 593-606.

Haas, C. (1989). Does the medium make a difference? Two studies of writing with pen and paper and with computers. *Human-Computer Interaction, 4,* 149-169.

Haas, C., & Hayes, J. R. (1986). What did I just say? Reading problems in writing with the machine. *Research in the Teaching of English, 20,* 22-35.

Hagge, J. (1987). The process religion and business communication. *Journal of Business Communication, 24,* 89-120.

Hayes, J. R. (1989). Writing research: The analysis of a very complex task. In B. Simon, D. Klahr, & K. Kotovsky (Eds.), *Complex information processing: The impact of Herbert Simon* (pp. 209-234). Hillsdale, NJ: Erlbaum.

Hayes, J. R., & Flower, L. S. (1980). Identifying the organization of writing processes. In L. W. Gregg & E. R. Steinberg (Eds.), *Cognitive processes in writing* (pp. 3-30). Hillsdale, NJ: Erlbaum.

Jenkins, I., Lansman, M., & Smith, J. B. (1989). *A tutorial for WE 1.0 for use with experimental studies.* (Tech. Rep. No. TR89-016). Chapel Hill, NC: University of North Carolina Department of Computer Science.

Kellogg, R. T. (1986). Writing method and productivity of science and engineering faculty. *Research in Higher Education, 25,* 147-163.

Kellogg, R. T. (1987). Writing performance: Effects of cognitive strategies. *Written Communication, 4,* 269-298.

Kellogg, R. T. (in press). The effectiveness of prewriting strategies as a function of task demands. *American Journal of Psychology.*

Kellogg, R. T., & Mueller, S. (1989, November). *Cognitive tools and thinking performance: The case of word processors and writing.* Paper presented at the annual meeting of the Psychonomic Society, Atlanta, GA.

Lansman, M. (1991). Organize first or write first? A comparison of alternative writing strategies. (Tech. Rep. No. TR91-014). Chapel Hill, NC: University of North Carolina Department of Computer Science.

Lutz, J. A. (1987). A study of professional and experienced writers revising and editing at the computer and with pen and paper. *Research in the Teaching of English, 21,* 398-421.

Matsuhashi, A. (1982). Explorations in the real-time production of written discourse. In M. Nystrand (Ed.), *What writers know* (pp. 269-290). New York: Academic Press.

Selzer, J. (1983). The composing processes of an engineer. *College Composition and Communication, 34,* 178-187.

Smith, J. B., & Lansman, M. (1989). A cognitive basis for a computer writing environment. In B. K. Britton & S. M. Glynn (Eds.), *Computer writing environments* (pp. 17-56). Hillsdale, NJ: Erlbaum.

Smith, J. B. & Lansman, M. (1991). *Cognitive modes and strategies for writing.* (Tech. Rep. No. TR91-047). Chapel Hill, NC: University of North Carolina Department of Computer Science.

Smith, J., Rooks, M. C., & Ferguson, G. J. (1989). *A cognitive grammar for writing: Version 1.0.* (Tech. Rep. No. TR89-011). Chapel Hill, NC: University of North Carolina Department of Computer Science.

Smith, J. B., & Smith, C. F. (1987). *A strategic method for writing.* (Tech. Rep. No. TR87-024). Chapel Hill, NC: University of North Carolina Department of Computer Science.

Smith, J. B., Smith, D. K., & Kupstas, E. (1991). *Automated protocol analysis: Tools and methodology.* (Tech. Rep. No. TR91-034). Chapel Hill, NC: University of North Carolina Department of Computer Science.

Smith, J. B., Weiss, S. F., Ferguson, G. J., Bolter, J. D., Lansman, M., & Beard, D. V. (1987). WE: A writing environment for professionals. *Proceedings of the National Computer Conference '87* (pp. 725-736). Reston, VA: AFIPS Press.

Strong, W. (1985). Linguistics and writing. In B. W. McClelland & T. R. Donovan (Eds.), *Perspectives on research and scholarship in composition* (pp. 68-86). New York: Modern Language Association.

Swarts, H., Flower, L. S., & Hayes, J. R. (1984). Designing protocol studies of the writing process: An introduction. In R. Beach and L. S. Bridwell (Eds.), *New directions in composition research* (pp. 53-71). New York: Guilford.