CHAPTER 13

Psychological Issues in Support of Multiple Activities

YOSHIRO MIYATA and DONALD A. NORMAN

That people engage in multiple activities at the same time should not be surprising. Nonetheless, little is known in contemporary psychology about either the phenomena that result or of the underlying mental structures. In this chapter we present a quick tour of the psychological theory relevant to the understanding of people's actions and of multiple activities. Consider this an "approximate theory" that is sufficiently accurate and relevant for the analysis of the basic phenomena. We conclude with a discussion of how the theoretical ideas can be applied for system support of multiple activities.

THE PSYCHOLOGY OF MULTIPLE ACTIVITIES

There are three major areas of psychological studies relevant to the study of multiple activities:

- Studies of memory, especially studies of short-term and working memory and of the organization of knowledge in long-term memory.

- Studies of attention, including studies of "controlled" (conscious) and "automatic" (subconscious) behavior, studies of "simultaneous attention" (how much can be done at the same time), and studies of "selective attention" (what activities interfere with one another).

- Studies of action, how people actually accomplish tasks.

Although these topics constitute three distinct areas of study within psychology, we believe the three to be intimately related. In particular, the theoretical tools of memory schemas and activation values underlie all three. These topics are all under active study and there is no general resolution of the appropriate theoretical mechanisms. But for the analyses that follow, we need only an approximate model, and for this purpose, much is known.

Memory

A reasonable approximate model of the Human Information Processing system divides processing structure into conscious and subconscious operations and memory into two classes of structure, short- and long-term memory (STM and LTM). In this approximate model, we treat working memory and STM as the same structures, and in the chapter we primarily refer to working memory. For current purposes, all that matters about STM is to recognize its limited capacity—for example, as described by the "5-slot model of memory" in Norman's Chapter 3, working memory can be thought of as having a capacity for only 5 items at any one time. Long-term memory (LTM) consists of organized knowledge units, called schemas, that structure knowledge and also contain the procedural information necessary to control actions. Several aspects of LTM are important, including the difficulty of acquiring new information and the problems and issues in retrieval of information, once acquired. An important aspect of memory is "reminding," the manner by which one event may cause retrieval of the memory for another. (See Card, Moran, & Newell, 1983, for further discussion of approximate models of memory; see Norman & Bobrow, 1979 and Schank, 1982, for a discussion of reminding and memory retrieval.)

Two Control Systems: Conscious and Subconscious

The large body of psychological research on attention allows us to develop a two system approximation to the control of behavior. One system, conscious control, has limited resources, especially that of STM or working memory. In general, the resources required to do any particular task subtract from those required for the simultaneous conduct of other tasks. In practice, these resource limits are severe enough that we can assume that only a single task can be under conscious control at any one time. The other system, subconscious control, seems to develop specialized procedures for tasks that are relatively independent of one another. As a result, we can treat subconsciously performed tasks as resource unlimited, so that several can be done simultaneously (as long as they do not require joint use of the limbs or sensory organs). Only well-learned, routine tasks can be done subconsciously. Note that the subconscious control system does not appear to use STM.

The conscious limitations do not preclude a person from simultaneous conduct of subconscious or automatized activities. Indeed, skilled people make use of automated actions to allow themselves to do several things at the same time. A good typist thinks of other things while typing: A skilled computer user plans ahead for several activities while doing the first. Again, although the physical limits of our limbs and sense organs are real, skilled practitioners (such as circus performers) can do things we never would have thought possible. As Buxton shows in Chapter 15, proper development of input devices allows us to extend the number of operations we can perform simultaneously.

Note too that the definition of "task" changes as skill develops. Thus, a beginning piano player treats control of the two hands as separate tasks, but the accomplished pianist does not. We return to all these issues in the discussion of "backgrounded" activity.

Conscious control is used primarily in four situations:

- When the task to be performed is novel or ill-learned.

- When the task is perceived to be especially critical, difficult, or dangerous.

- When there is a need to override the automatic control, either to cause actions that would not otherwise take place or to inhibit those which would take place (but are not desired).

- When there is a need to resolve conflict among schemas (or activities), especially when an ongoing activity is interrupted by demands from a different activity.

These properties of conscious control are discussed in more detail by Norman and Shallice (1986). Norman and Shallice assume that conscious control occurs by increasing or inhibiting the current activation values of schemas.

Planning

A primary need for planning comes about because of what Cypher (Chapter 12) calls "linearization": the act of taking the several activities that might be competing for attention and arranging to do them in sequence. That is, if we can do only one task at a time, then when many need to be done, they must be arranged in a linear sequence. Planning needs to occur for other reasons, however. Some tasks must wait until prerequisite conditions are satisfied. Planning helps discover these situations and arrange the order of task execution so that prerequisites are satisfied. Some tasks may have critical timing requirements: Again, the planning stage helps arrange matters appropriately. Finally, some tasks should be postponed (even if one wishes to do them) when more important work should be done instead, and some unpleasant tasks should be done, even if one would rather not.

One of the major difficulties in planning comes about because of the limitations of human processing resources. The limits on conscious resources and working memory capacity mean that in-depth planning is often not possible without external aid. Moreover, at the start of an activity, many different ideas are often jumbled together. The problem is that concentration upon one tends to cause the others to be forgotten. The primary planning aids are those that support memory. In addition, various retrieval aids are required to allow further work on the ideas and to act as reminders, to make sure that once generated, plans are not forgotten. We treat these aspects under the discussion of reminders.

Interruptions

Interruptions can be both external and internal. External interruptions result from events in the environment. Internal interruptions come from our own thought processes—new ideas that draw attention from the current activity. Interruptions introduce new tasks on top of the ongoing activity, often unexpectedly. As a result, conflicts arise. Because of a person's limited processing and memory capacity, one suspends work on current activity at the risk of losing track of the current activity by failing to resume the work where it was interrupted.

But taking the time to make the current activity recoverable runs the risk of losing the new idea.

There is a conflict in processing between the need to concentrate upon one thing in order to give it full processing capability and the need to be alert for unexpected, but relevant and important, thoughts and external events. As a result, the human information processing system seems to exhibit conflicting properties: continual concentration and continual distraction. The first property moves the system toward operation as a single purpose, dedicated processor, primarily attending to one task and ignoring other events: We call this *task-driven processing.* The other property moves the system toward operation as a responsive processor, continually changing its activities to reflect new thoughts and ideas or to respond to events in the environment: We call this *interrupt-driven processing.* When people are deeply engrossed in a book or movie, they are task-driven. When people are in a job that requires constant interaction with others, whether by telephone or in person, they are apt to be interrupt-driven.

Task-driven processing. In a task-driven state, people are so occupied by the processing of the ongoing task that there is apt to be an effective decrease in sensitivity to events external to the activity. This lack of sensitivity to external events has been widely studied by psychologists interested in "attention." A typical experimental paradigm investigates a person's ability to process information in a secondary task while concentrating upon a primary task. These studies show that there is a severe limitation in the degree to which information is processed from the secondary activity. Although people remain aware of the existence of activity around them while concentrating upon a task (while under task-driven processing), they can do only minimal processing of this external activity, not enough to draw meanings and implications. Thus, while deeply engrossed in a book, people can be quite unaware of questions addressed to them. People do note the presence of activity, but not the content: Signals intrude, but meanings are only minimally extracted, if at all. (A reasonable review of these phenomena, although a little old now, is Norman, 1976. See also Shiffrin, 1986).

The implications of these findings for our analyses are that if people become too engrossed in the task to which they are paying conscious attention, they will not process other events that occur on the computer screen. Note that they will detect gross signals, but not their meaning. Thus, even when people are heavily engrossed in task-driven situations, they will notice abrupt sensory signals such as flashes of light or auditory tones. Unfortunately, the type of signals needed to attract the attention of people heavily engaged in task-driven activities are also the kinds of signals that are most intrusive and annoying to those who are not in this state, for they draw away processing resources from the main task.

Interrupt-driven processing. The phrase "interrupt-driven" processing refers to the situation in which people are especially sensitive to extraneous events, easily distracted by extraneous thoughts and external signals. The normal working environment is full of potential interruptions: phone calls, talking, requests for help from colleagues, and so on—unpredictable and often irresistible interruptions. Internal interruptions occur as a natural result of thought, as new ideas and new topics get suggested by the processing for the current topic. The result can be captured by some other thought, perhaps unrelated to the task that is supposedly being performed. It seems more difficult to maintain attention on a topic that is purely internal than on one which has external support.

Whether a person is in a state of task- or interrupt-driven processing is a function of both the person and the activity. Differences among individuals play a role: Some people are more distractable than others, some are more easily controlled by task-driven structures, others are more distractable by extraneous events or thoughts. The amount of external activity clearly is relevant: It is obviously difficult to maintain a task-driven state in the presence of external events irrelevant to the main task. Task-driven processing continues when then processing is dominated by the schemas relevant to the activity. The result is that there are few resources available for other activities. Interrupt-driven processing occurs when the activity does not have much structure, or if the external support for the activity is such that schemas are not always kept activated. The result is that there will be idle resources that will tend to get used for extraneous activities.

Multiple Activities: Current and Suspended

We now turn to an analysis of the different types of multiple activities. First, activities are either currently controlling actions or are suspended: Call the first case "current activities," the second "suspended activities." Second, there are two forms of current activities, those that are in the foreground of conscious attention ("foregrounded activities") and those that are not ("backgrounded activities"). And finally, there are two forms of backgrounded activities: external and internal. The result is three major classes of activities:

I. Foregrounded Activities:
A current activity under conscious control.

II. Backgrounded Activities:
There are two classes of backgrounded activities: external and internal. Externally backgrounded activities are those done by some other agency. Internally backgrounded activities represent ongoing activities under "automatic" or subconscious control.

III. Suspended Activities.

Current activities. We distinguish among several different types of ongoing, current activities, divided into two major categories. One category is reserved for the primary activity, the activity that is the focus of conscious attention: We call this the "foregrounded" activity. The other category is used for ongoing, active tasks that receive little or no conscious attention: We call this the "backgrounded" activity.

> *Thus, in writing a paper (or a program), there are at least two tasks going on: One is the development of the ideas, the other the act of typing them onto the computer keyboard. The development of the ideas should be the foregrounded activity and the act of typing the backgrounded activity. This holds only for skilled typists: The act of typing is backgrounded; typing takes place simultaneously with the development of the ideas and without interfering with them. For nonskilled typists, this is not possible. They must focus so much attention upon the typing that it becomes the foregrounded activity, interrupting and suspending the development of the ideas. This leads to severe disruption of the task of idea development, and, as a result, many nonskilled typists cannot compose at the keyboard, but prefer other means of composition, one where the translation of thoughts to symbols is more automatic (for example, dictation or handwriting).*

Backgrounded activities. Backgrounded activities result whenever a task is performed "automatically," without conscious supervision, thus allowing other activities to be done at the same time. There are two classes of backgrounded activities: *external* and *internal.*

Externally backgrounded activities.

The phrase "externally backgrounded" comes from the field of computer science and the development of computer systems that allow a program to be run in "background": performed by the computer as an independent job, without interaction or supervision by the user (and often at a lower priority than "foregrounded" jobs). We generalize this concept to refer to any task that is being performed by an external system, such as another person or a computer, without requiring supervision. The important point is that when a task is backgrounded, it no longer requires conscious attention and other activities can be started. Examples in the domain of computers are frequent. Thus, in writing a paper, the getting the "hardcopy" printout of the final product is often "externally backgrounded." In similar way, compiling and loading a long program is "externally backgrounded." This allows the user to start another job before these backgrounded tasks are completed. Often it is desired to resume the backgrounded task as a foregrounded activity when the external system has completed the tasks. But if no appropriate signal or reminder is presented to the user signaling the completion, there may be problems in recalling that a task was backgrounded and, hence, a failure to resume.

Internally backgrounded activities.

We generalize the notion of "backgrounding" still further for the category of "internally backgrounded" tasks. Here, we refer to the situation where the person continues the task, but in situations where it is so well practiced and learned that performance can take place with minimum conscious control, thus allowing other things to take place at the same time. Thus, skilled practitioners can do one task (e.g., talk or type) while their main attention is devoted to something else. Although this kind of "automatic" or "backgrounded" performance only takes place with highly skilled, well-practiced behavior, it is still a common, frequent occurrence.

Errors, especially "slips," are likely to occur when attention is not focused upon the current task. Thus, "capture errors" are most likely to occur in "internally backgrounded" activities. An expert typist intending to copy a limited part of a written

> *page is apt to end up having copied the whole page: The more distractions or other activities present during the action, the more likely the task will be internally backgrounded, and the more likely the error.*

Suspended activities. Activities that are not current are called "suspended." The schemas relevant to the activity remain activated, but do not control performance. Activities can be suspended for a variety of reasons, including both external and internal interruptions, boredom or fatigue with the task, the need to delay until some prerequisite condition has been satisfied, or even from a judgment that things are not going well and that delay would help. Some of the issues confronting the suspension or postponement of activities are discussed by Cypher in his chapter. Whatever the reason for the suspension, the result is a task awaiting the appropriate time to resume execution. The critical issue with suspended activities, of course, is how, when, and whether they will be resumed.

Reminding

Reminding is required if suspended activities are to be resumed at the appropriate time or place. The reminder, therefore, is a signal that indicates that a suspended or backgrounded task still exists or that it is ready for further processing. From the practical point of view, reminders help the person remember what is to be done. From the theoretical point of view, the problem is to reactivate relevant schemas and to re-establish whatever information is needed in STM or the environment so that the appropriate trigger conditions will be satisfied. It is good to keep in mind that although reminders are valuable in re-establishing a planned or suspended activity, they often interrupt ongoing activity.

Reminders as signals and as descriptions. There are two aspects to a reminder:

1. *Reminder as signal*: to indicate that something is to be remembered;

2. *Reminder as description*: to aid in retrieving what was to be remembered.

These two aspects of reminders can be quite independent: A reminder can succeed at one aspect while failing at the other. For example, a cooking timer is a common aid in preparation of a meal. But the ringing of the timer acts only as a signal that some externally backgrounded activity has been completed. The timer is nondescriptive, and it is up to the cook to retrieve what activity is being referred to. A list of activities to be performed during the day provides a good description of each individual activity, but if the list is hidden from view, it acts as a poor signal: The list may not be examined because it was never noticed.

For a reminder to be effective in specifying the activity, it must act as a memory retrieval cue. The reminder can be thought of as a partial *description* of the to-be-remembered materials (Norman & Bobrow, 1979). The relevant factor is specificity of the cue. Sometimes the cue itself specifies the task. Dirty dishes in the kitchen sink serve as excellent descriptions of the task to be done—wash the dishes—but they are apt to have poor signaling qualities if they are hidden from casual view. A boiling tea-kettle is intermediate: The sound of a tea-kettle boiling is an effective signal, but it may not always be a good description of the activity. A telephone call from a friend reminding of a dinner appointment serves both as a good signal and a good description. Messages on display screens can be effective reminder cues. A light, a nondescriptive message or icon, or the common mnemonic of tying a string around the finger all provide only very partial descriptions of the items to be remembered, oftentimes requiring deliberate and difficult retrieval processes.

> *An example of a reminder that acts as a good signal but a poor description comes from the job control features of the Berkeley distribution of UNIX. On our computer system, attempts to execute the "logout" command are sometimes rewarded with the message:* "There are stopped jobs." *This message is a reminder that some tasks have been suspended but it gives no indication of what those tasks are. (It also gives no indication of the possible courses of action, and beginning users have reported feeling frustrated by the presence of a reminder and the failure of the "logout" command to work, combined with a lack of information about the cause of the problem and no hint as to what they are supposed to do about it.)*

For a reminder to be a good signal, it must be conspicuous. This can be accomplished in several ways. The literature on selective attention indicates that simple sensory signals can be especially effective, being noted even when a person is deeply engaged in activity (flashes of light for a person engaged in auditory task: a tone or bell for a person engaged in a nonauditory activity). Discrepancy is especially relevant: An event that is unusual or not expected acts as a good signal. A frequent or common event does not. An alarm is effective only in an environment in which the sound of the alarm is infrequent. People seem especially sensitive to change and to violation of expectancy. When a discrepancy is noted, the resulting processing can interrupt ongoing activity, drawing attention and processing resources to the discrepant part of the environment. All this suggests that a cue is most effective when it is discrepant from one's ongoing expectation. Tying a string on a finger is an effective cue only if there is not usually a string around the finger. Display of a message or signal light is effective only if there is not usually a message or light in that area.

SYSTEM SUPPORT FOR MULTIPLE ACTIVITIES

Several different aspects of multiple activities require support. Cypher has discussed a number of the issues in his chapter: Here we expand upon the notions, with special emphasis on support for reminding (resumption of tasks following interruptions and suspensions). First, we discuss support for transitions between one activity to another focusing on two aspects: suspension of activity and reminding of activity. In addition, we discuss some aspects of support during execution of an activity, especially with regards to the execution of simultaneous activities.

Support for Suspensions of Activities

What happens when the decision is made to suspend current tasks? This usually takes place when some internal or external event occurs that makes it desirable to change the task currently foregrounded: See the description of these interrupting events in Cypher's chapter. This means that a new task is foregrounded and the previously current task becomes backgrounded or suspended. If the change occurs at the conclusion of the current task or at a natural breaking point, then there is probably no difficulty.

Unpredictable interruptions are likely to occur at potentially disruptive times. As a result, support for suspensions of activities has three aspects. First, the system should be designed so that it is easy to suspend an activity when this is desired, without interfering either with memory for the current task or with the thoughts relevant to the interrupting task. The suspension should not require much activity, or else thoughts in working memory will tend to be lost, thus interfering with the activities to be performed on the interrupting activity. Second, sufficient information should be saved with the suspended task so that when the activity is resumed, it can be continued where it left off (recovering the active thoughts is the hard part). Third, a reminding structure should be established so that the user does not forget that the task is still unfinished. Making the suspension easy to accomplish (the first point) while simultaneously saving enough context and unfinished ideas to allow smooth resumption (the second point) are somewhat opposing requirements: yet another tradeoff to worry about.

> *Cypher's* Notepad *program (Chapter 12) addresses all three aspects of the suspension of activities. It is easy to interrupt a task—a single keystroke will do it.* Notepad *saves all context and continually displays (and thereby reminds one of) the titles of the tasks that have been interrupted. The user can review the list of unfinished and interrupted tasks, and the program provides a simple means to step back through them. Completed tasks do not appear on the list. This procedure also provides the user with a way to deal quickly with potentially disrupting events. The user can decide to interrupt the main task only briefly enough to "jot down" a simple reminder of the interrupting event, then to resume the main task. This avoids major disruption by the new event while still noting its existence.*

Support for Reminding

When an activity is suspended, it needs to be resumed at a later moment when time is appropriate, when some prerequisite conditions are met, or when the user is free. The task-driven aspect of processing suggests the need for reminders to overcome the limits on memory and processing capability. The problem is that focusing upon one activity makes a person unresponsive to other events, including reminders of the activities that have been suspended or backgrounded. As a result, once an activity has been started, it is possible (likely) to forget about other things that should be done. If an action is to be started at some specified time or when a specified condition arises, the intended activity is likely not to be done if the person is engaged in some other activity at the time.

What would an ideal reminder look like? It should:

1. Inform the user when conditions are ready for resumption of a suspended or backgrounded activity.

2. Remind the user when something has to be done immediately.

3. Not distract from the current activity.

4. Continuous or periodically list activities that have been suspended or backgrounded.

5. Help resumption of an activity by retrieving the exact previous state of the activity and making it available to the user.

When should reminders occur? The problem with reminders is that they are also interrupts: One task's reminder is another task's interruption. Suppose that there is some control over when a reminder is presented to the user. When would be the ideal time to do so? What are the natural breaking points in activities? Deciding when to remind is a very sensitive problem. The importance and relevance of the conditions are important factors, but how can these be determined by the system? If the user is producing new ideas constantly, or if the user is working on a very complicated process like programming and has to keep a lot of information in working memory, it is not a good idea to interrupt the train of thought. An experienced assistant knows when to interrupt the boss or not. How can this expertise be captured?

In general, the major factor that determines whether a reminder should interrupt the current activity is the relative importance of the two activities. But even if the current activity is less important than the interrupting one, it is still necessary to examine the state of the current activity before interrupting: There are some states at which interruption would be very disruptive, and some where it would not disrupt as much. (We discuss this factor in the next section—the stage analysis of user activities.)

Sometimes, it is possible to have the user specify which activities can (or should) interrupt the current activity. In extreme cases, people lock their doors, turn off their telephones, and put up "do not disturb" signs. Some systems allow the user to control whether or not system messages or announcements of arriving mail will be permitted to intrude. It is often desirable to take into account the importance of the interrupting activities. Thus, although irrelevant messages might not be wanted, if a task has been externally backgrounded, the message that reports its completion is a desired interruption. Similarly, if in the writing of a paper, information is sought from some other person, then that person should be allowed to interrupt, but only for the relevant topic.

Reminders and interrupts provide a large set of tradeoffs among options. One problem is to determine how priorities can be established without excess effort by the user. Some means have to be devised for handling interrupts and reminders that do occur, but that are withheld from the user: Interruptions that are not important should be put off until the user finishes the current activity, but they should never be discarded. The two separate aspects of reminders can be handled separately: reminders as signals and reminders as description (the message part of the reminder). Different kinds of reminders might be allowed to use different levels of signals, from covering the whole screen, overwriting whatever is there (the most obtrusive, attention-getting form) to a subtle indicator in one corner of the screen indicating that reminders are awaiting user-initiated action.

Interruptions and the stage analysis of user activities. The analysis of seven stages of user activity (in Norman's Chapter 3) suggests that an interruption would be least disruptive if it occurred between the completion of the last stage—evaluation—and the formation of a new goal or intention. However, it is also clear that some points of interruption between stages should be less disruptive than others. In particular, interruptions where memory load is high should be disruptive, whereas interruptions where load is low (perhaps because of the reliance on external cues) should be less disruptive. This suggests that interruptions should be most disruptive while in the planning or evaluation stages: i.e., during the formation of the intention and the development of the action, and during the interpretation and evaluation of the outcome. Interruptions should be least disruptive at the juncture between execution and evaluation, where there is maximum use of external information. It is clear that an interrupting message should not be presented in the stages of execution or perception, stages where the user is directly interacting with the system.

> *Of course, these observations are not easy to follow because the stages where disruption is apt to be most serious consist of mental activities (planning, interpreting, evaluation), and these activities are usually not visible to the system. However, a good procedure might be to present relatively unobtrusive reminders to the user just following the completion of an action. If timed right, the reminder will not disrupt the*

ongoing task, but will be present and visible when the user completes the cycle and is starting to think of the next sequence of activities. The visibility of the reminder has to be carefully selected so that it does not disrupt if it comes on during evaluation, but yet is noticeable when the user completes the tasks. Think of the user buried in the task, unwilling to be interrupted, but every so often finishing a cycle and "coming up for air," quickly breaking from the task and taking a quick look around. The reminder should only be noticeable during that "breathing spell."

Memory Aids. The old folk saying, *Out of sight, out of mind* makes a good slogan for designers. The slogan speaks directly to the data-driven aspect of physical reminders: Something that is physically present keeps its memory schema in an activated state through data-driven activation. This means that one way of reminding is to keep visible the activity that is to be remembered. In computer systems, this could be implemented in a variety of ways, but the most common today is through the use of windows, lists, menus, and icons. (Windows are important enough that we give them their own special section.) Keeping things to be remembered constantly present has several advantages and disadvantages.

The advantage of continual presence is continual reminding. The deficits are distraction and clutter, as well as potential loss of working space. Keeping piles of work to be accomplished on a desk diminishes the usefulness of the desktop as a working space: The same is true of the computer screen. The problems of distraction are real. Being reminded of the important tasks that remain undone can have serious implications: It can demoralize, it can cause rapid switching among ideas. (This leads to the human state analogous to the computer system state of "thrashing": So many resources are used in contemplating all the tasks yet to be done, that there are no resources left to do the task.)

Making reminders nonvisible avoids all the deficits of visibility, but also avoids the virtue as well: Out of sight, out of mind. A compromise solution is possible. Reminders can be simplified, text can be abbreviated, windows "closed" or "shrunken." A group of descriptive reminders could be replaced with a simple flag or note that informs (reminds) the user that more specific reminders are awaiting attention. Then, when the user is ready to deal with the interrupts, the flag or note could be expanded into its set of more complete, more detailed reminders.

> *There is an important set of issues here that should be explored: The interaction of stress and performance. When under high stress, human performance deteriorates. The conscious resources available decrease. There tends to be a focusing upon one task to the exclusion of others, even if the task is not directly relevant to the problem (and sometimes even when execution of the task makes the difficulties more severe). This is a special problem with the ability to deal with interruptions. The ability to perform backgrounded jobs and to plan decreases. Working memory capacity seems to decrease. The more stressful the situation, the more disruptive interruptions become, which increases the stress.*

Support for Concurrent Activities

One point that should be emphasized is that many complex tasks require simultaneous, concurrent activities. Thus, as we write this paper, it is often useful to have an outline of the entire paper as we write any particular section—an outline that changes as the writing continues reflecting the ongoing writing activity. Moreover, because this paper is so closely related to Cypher's chapter, we need a copy of that chapter in front of us as we write. Similarly, we need to refer to other sources of information, such as reference lists, or working notes. Finally, because this is a joint activity, at times we need to interrupt the writing activities in order to see if the other person is logged on to the computer network or if mail relevant to the writing of the paper has been received (while somehow avoiding the distraction of the other messages that would be discovered, irrelevant to the writing activity, but probably a more attractive pursuit of time).

So far, the best way we know to provide this kind of support is to use multiple displays. This can be done either through multiple terminals or by using windows on a terminal that has a sufficiently large screen to allow simultaneous display of the output of several ongoing programs. The problem with this solution is to avoid the surplus of riches: Some of the displays may be useful in the conduct of the activity, but they also act as dangerous lures, leading one to stray from the unpleasant duties of the current task.

> *The problem of distraction by other, more attractive messages or thoughts that are not relevant to the current activity is a pervasive difficulty. It leads to the types of interruptions Cypher called " While-I'm-At-It." The problem is exacerbated by the manner in which information is usually displayed: No attempt is made to suppress irrelevancies. This characteristic is common to many information systems: It is especially problematic in the use of dictionaries, atlases, and encyclopedias. It happens mentally as well: The attempt to think of one topic often leads to thoughts of others, at first related to the item of interest, but possibly leading to thoughts far astray. The extraneous retrievals can be both beneficial and disruptive. But the important point is that it would be useful to have a better control over the tradeoff: To be able to choose to be diverted or not, depending on the context.*

This paper was written with two different kinds of terminals, each of which supported a different style of support for concurrent activities. One system, the *SUN Workstation*, provides for a large, bit-mapped display with multiple windows, allowing for simultaneous display of a number of components (see Figure 13.1). Thus, a typical session on the *SUN* has one window devoted to the text, one to an outline and another to an outline of Cypher's paper. In addition, other windows are used for other, suspended or backgrounded activities.

Most of the work, however, was done on 24-line by 80-character terminals. Here the support for concurrent activities came primarily through the "jobs" facility of Berkeley UNIX, so that although only the foregrounded activity was visible on the screen, it could be suspended immediately by the typing of a single character (control-Z) and a previously suspended program made visible by the typing of a relatively short command (*fg %N*, where *N* is its "job control" number). Still, this is disruptive and it prevents simultaneous viewing of different components of the task. As a result, work on these normal terminals had to rely more on external support, such as printed copies of the papers and outlines. In many ways this was more convenient, for external reminders were easier to read and scan than even the multiple windows on the SUN, but in other ways it was less convenient, for the paper was static and did not reflect changes that took place over the course of a work session.

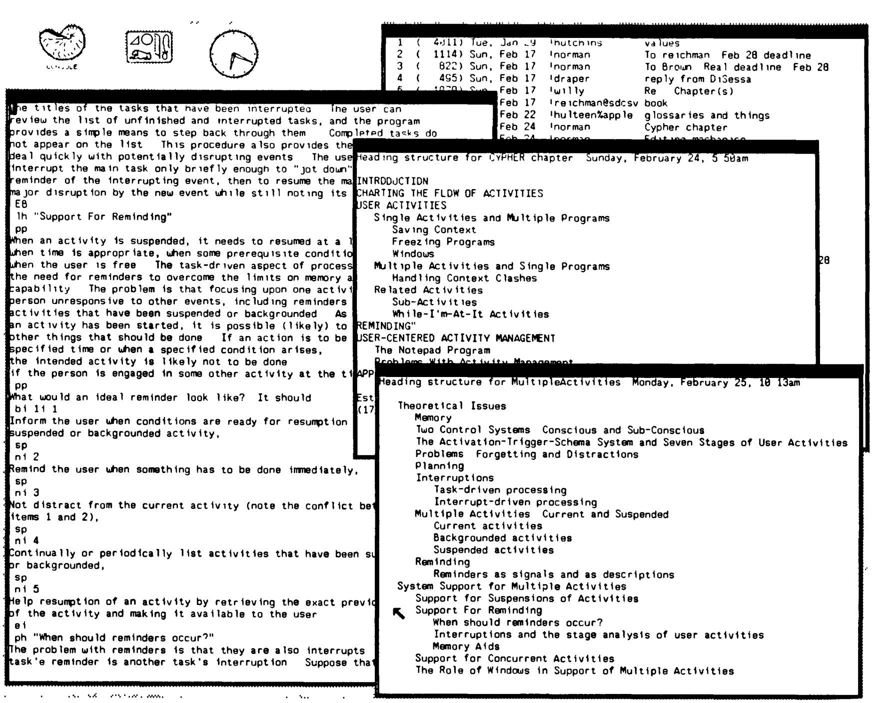

FIGURE 13.1. A typical screen on the *SUN Workstation* being used in the writing of this paper. The large window on the left supports the foregrounded activity: the writing of the text. Other windows provide other information related to main task, such as the outline of this and Cypher's chapter (the two medium sized, overlapping windows on the right). Other windows provide support for and reminders of suspended and backgrounded tasks. Thus, one window shows the time of day (as a clockface), another shows mail received by the message system. Several icons represent closed windows, acting as "signals" of other activities that have been suspended. Finally, and not visible, an alarm clock program runs in background that will signal at the requested time, and (one hopes) cause the user to stop work and get on to some other activity.

How well a system supports multiple activities can depend on the details of the hardware as much as the software support. Thus, we found dramatically different support was provided by the very same computer system, terminal, and software as communication rate (Baud rate) was changed. Thus, the job-control facility of Berkeley UNIX works reasonably well to support multiple activities when the terminal is in the office, operating at 9600 Baud—it takes about 2 seconds to display a full screen. However, when working at home, at 1200 Baud, it takes as much as 16 seconds to display the screen. As a *result, the job control facility was not used much to bounce back and forth between jobs: To leave the paper in order to check one item on an outline (saved in a suspended job) and then return could take as long as 32 seconds—a long time to spend simply waiting. When one of us switched to a 2400 Baud modem at home, the difference was dramatic: It cut display times in half, and the result was a substantial increase in the amount of switching back and forth among activities.*

The Role of Windows in Support of Multiple Activities

Windows were originally designed as explicit supports for the conduct of multiple activities. Some of these aspects have been forgotten in the current craze for window systems. Moreover, there has not developed any systematic body of knowledge about the properties of windows, either as reminders or just as useful interface structures.

Window systems make possible the display of considerable information for each of the multiple activities that are currently active, subject to limitations on the size of the screen and the memory space allotted to handing the screen map. All sorts of reminders can be presented on the screens because a major portion of them are continually visible. Windows themselves can serve as reminders of the existence of the activities contained within them. Life, however, is a series of tradeoffs, and window systems are no exception.

Windows distract. Too much extraneous information on the screen distracts from the main task. It distracts both by drawing attention to subsidiary, irrelevant activities and also by cluttering the visual appearance, thus increasing the difficulty of finding the point of interest. People are easily distracted, especially as work on a long, tedious (but necessary) task drags on and attention momentarily wanders from the work. At that point, it is *not* desirable to have a subsidiary window that contains an ongoing game display, or a list of new messages received by computer mail, or partially completed other tasks. In this case the window system, is much *too* effective as a reminder. A slogan summarizes the tradeoffs inherent in all reminding schemes (not just those involving windows): *One task's reminder is another task's distractor.*

Organization of windows. Major discussions in the window-design community take place on the desirability of the various ways of organizing windows: Should they be neatly placed on the screen, or should they be allowed to be put wherever they are handy? Should the user have control, or should the system provide neat partitioning of the screen space in some predetermined manner? Should cleanup be done automatically, or only when requested in some predetermined manner?

Much of this controversy is outside the scope of this chapter. Some of it is unsolvable, both because different tradeoffs occur in different circumstances and also because different people have different preferences and work habits. However, we do wish to remind of one important principle: *exploit spatial location*. People can use space as effective memory aids. This allows reminding to take place simply by the use of markers in space, with minimum or litle need for verbal labels or special icons. Even where icons or full text is available, spatial location acts as powerful cues to the contents without the need for the disruptive processing required to interpret the icons or text. But space can serve as an effective reminding tool if it is consistent. This means that windows, icons, or other reminders should have fixed positions, and that each time the computer system is used, the same positions are always used for the same information. Automatic systems that restructure the screen each time the system is called can be as disruptive as the well-meaning cleaning person who "tidies up" a person's private possessions, desktop, or bureau drawer.

ACKNOWLEDGMENTS

We thank Allen Cypher and Steve Draper for their comments on the manuscript.