9

Executive Control: Attention Switching, Interruptions, and Task Management

Introduction

Aviation accidents are often the result of poor task management (Dismukes and Nowinski 2007); the operator switches attention from critical tasks of airplane guidance and stability control to deal with an interruption (e.g., a communication from air-traffic control; a possible failure of landing gear) and then fails to bring attention back to the high-priority safety-critical task. In 1991 in Los Angeles, an air-traffic controller positioned a plane on an active runway, switched attention to a number of unrelated items, and then failed to return attention to the vulnerable airplane and move it to a different runway. Another plane was then cleared to land on the runway where the first plane had been left. Several fatalities resulted from the ensuing crash. In 1987 in Detroit, pilots configuring the airplane for takeoff switched attention to address a request from Air Traffic Control and then returned attention to the checklist-guided preparation activities after missing the critical step of setting the flaps, which were necessary to gain adequate lift on takeoff (Degani and Wiener 1993). In the resulting crash, more than 100 lives were lost.

These are examples of breakdowns in selective attention—the attention element described in chapter 4. However, the present chapter refers to attention directed to tasks rather than to perceptual channels, and the topic thereby can be relabeled as *task management* (Adams, Tenny, and Pew 1991; Damos 1997; Dismukes and Nowinski 2007; Dornheim 2000; Funk 1991; Wickens 2003). Also, in contrast to chapters 7 and 8, where the concerns with task management were those of the allocation of resources during parallel processing activities, here the focus is very much on sequential activities, in which parallel processing either does not or cannot take place.

In many activities, the environment may seem to dictate behaviors and tasks to the operator. This can occur because the appropriate response to a signal of some form is largely reflexive or because it has been inculcated by experience—as, for instance, with the tendency to orient toward the source of a loud noise or the tendency to brake in response to a red light, respectively. In other cases, however, the environment may present a range of potential behaviors, none more urgent or important than another in any perceptually obvious way. In such cases, the operator may easily fail to select or prioritize tasks optimally. After sitting down at the computer, for example, we may choose any of a large number of tasks to perform. Ideally, we might open a word-processing document to carry on writing an unfinished paper. Too often, we may instead double-click a Web browser and proceed to surf the Internet or might set to work on the paper and find ourselves distracted by a seemingly constant arrival of new e-mail and phone calls. The potential result in either case may be that a deadline for finishing the paper is missed. As shown already, the importance of appropriate task prioritization is higher still in dynamic and complex environments such as aviation or the hospital operating room, where the status of the system changes rapidly, the number of tasks to be juggled is large, and the consequences of poor management can be fatal (Chou, Madhavan, and Funk 1996). The typical nurse may have as many as ten tasks in a queue waiting to be performed, and the delay of some of these could have serious consequences for patient safety (Wolf et al. 2006).

This chapter first describes some of basic research on executive control and task-switching, processes that underlie the metatask of task management; then it turns to more applied research that deals with this issue in complex real-world domains and places particular interest in recent work on the psychology of interruptions.

Executive Control

When faced with variety of potential behaviors, how are we able—at least sometimes—to willfully choose to perform those that are most urgent or important and to suppress those that are not? How do we manage, likewise, to suspend or abandon an ongoing behavior when an alternative task assumes a new, higher priority? Models of cognition typically assign the intentional management of thought processes and behaviors to an executive or supervisory attentional component (Baddeley 1986; Norman and Shallice 1986). Here, the terms *executive attention* and *supervisory attention* are used interchangeably. The Norman-Shallice model illustrates the role of supervisory attention well. In this account, a repertoire of cognitive and motor behaviors—routine thought processes and actions that the operator is capable of performing—exists as a set of programs or schemas in long-term memory. A given activity is performed when its schema is triggered. Because many behaviors are mutually incompatible, only one or at best a small number of congruous schemas should be allowed to operate at a given time. Schemas are triggered through a competitive process known as contention scheduling. Here, individual schemas receive activation from the operator's perceptual system. Schemas that represent congruous behaviors then facilitate one another whereas schemas that are incompatible inhibit one another. A given schema finally assumes control of behavior when it achieves a dominant level of activation relative to the competing schema.

By itself, however, this process of stimulus-driven contention scheduling explains only the moment-by-moment, bottom-up control of routine behaviors. To allow an influence of top-down, willful processes and the capability for action planning, the Norman-Shallice model incorporates the supervisory attentional system. One role of the supervisory system is to bias the interactions between schema in a goal-driven manner. This entails holding the current goals in working memory and then providing activation to schema that match the operator's goals and suppressing those that do not. Failure to suppress unwanted schema can lead to capture errors or slips (Reason 1990), in which a familiar stimulus triggers a habitual response that is inappropriate under the circumstances. Driving to the store on a Saturday morning, for example, we may unthinkingly take a turn toward the office, carrying out a routine behavior that is triggered by the context. Not surprisingly, capture errors become more common when the executive attentional system is burdened (Roberts, Hager, and Heron 1994). Additional functions of the executive attentional system are to generate novel patterns of behavior (Baddeley 1986, 1996) or to plan extended behavioral sequences (Shallice and Burgess 1993). Under high levels of cognitive load, therefore, when the executive system is heavily burdened, behaviors tend to become less flexible and more stereotyped (Baddeley 1986). Damage to the executive system, moreover, impairs the ability to preplan sequences of behaviors needed to carry out many complex tasks (Shallice and Burgess 1993).

The operation of executive attentional processes is well illustrated within the real-world situation where an operator engaged in an ongoing task is interrupted by a second task, as happened in the Detroit crash. Either the need to switch attention to the second task may be announced by a signal to the operator, or the operator may decide perform the task wholly of his or her own volition, without an explicit cue to do so. In either case, the operator will be required to suspend the first task and switch attention to the second, a process that in and of itself can consume several tenths of a second (Monsell 2003). While performing the new task, however, the operator must maintain the goals of original task so that it can eventually be resumed (Altmann and Trafton 2002). Finally, when the interruption has been dealt with, the operator should be able to switch attention back to the original task, picking up as fluidly as possible where it was left off.

Task Switching

To study task management at the smallest time scale, we can examine the cognitive mechanics of switching attention from one task to another. Remarkably, even this simple process can entail a substantial time cost, a fact first demonstrated by Jersild (1927). Subjects in Jersild's study were presented lists of items and were asked to work their way through each list performing either or both of two different tasks. In some cases, for example, the stimuli were lists of numbers, and the subjects' task was either to add or subtract

18	60	32 + 3
26	COLD	80 – 3
41	44	38 + 3
45	NEAR	50 – 3
73	77	68 + 3
69	TOP	79 – 3

Figure 9.1 Stimuli like those used by Jersild (1927) and Spector and Biederman (1976). Subjects proceed down each list as fast as possible, performing a mental operation on each item, or trial, in sequence.

three from each item on the list (Figure 9.1, left column). In pure blocks, the subject performed the same task on each item in the list. In mixed blocks, the subject alternated back and forth between two different tasks while working through the list. Jersild found that the time necessary to complete the mixed blocks was substantially longer than the average time needed for the corresponding pure blocks. That is, the need to alternate back and forth between different tasks imposed processing demands beyond those associated with the mathematical operations themselves. Jersild's experimental procedure has become known as the task-switching paradigm, and the response-time (RT) increase produced by the alternation between tasks has become known as a switch cost. Switch costs can be measured by examining trial-by-trial RTs for blocks in which tasks alternate in pairs (e.g., A, A, B, B); the switch cost is the difference between RT following a switch and RT following a task repetition (Figure 9.2). In addition to the trial-by-trial switch costs, RTs in mixed tasks blocks may also show a more general mixing cost.* This can be measured by comparing RTs for task repetitions in mixed blocks with the mean RTs for pure blocks. Frequently, RTs for repetitions in the mixed blocks are longer than pure block RTs, indicating an additional slowing of mixed block performance even after trial-by-trial switch costs are accounted for (Kray and Lindenberger 2000; Monsell 2003).

What is the cause of these switch costs? Data indicate that multiple effects contribute. One is uncertainty about which task to perform on a given stimulus. Jersild (1927), after discovering the task-mixing effect, demonstrated that the costs of task mixing were eliminated when the stimuli for the alternating tasks were mutually incompatible so that the nature of the stimulus implicitly defines the nature of the task—for example, when one task was to

* In the basic attention literature, switch costs as described here are sometimes referred to as local or specific switch costs, whereas mixing costs are referred to as global or general switch costs.

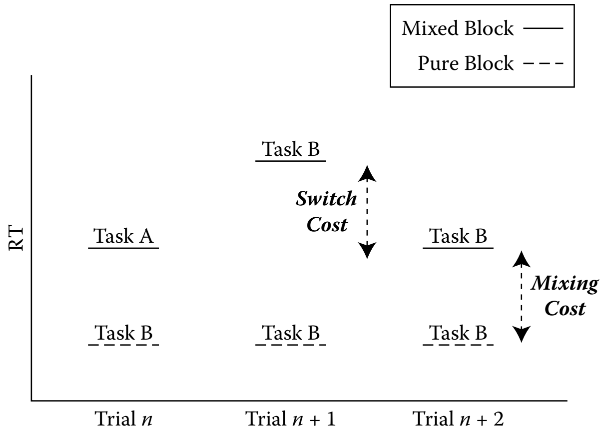

Figure 9.2 Hypothetical data illustrating switch costs and mixing costs. A switch cost is the difference between RT for a given task following a task alternation and RT for the same task following a repetition. Here, the switch cost is the difference in RT for task B on trial $n + 1$ and trial $n + 2$ of the mixed trial block. A mixing cost is the difference between nonswitch RTs for a given task in the mixed block and RTs for that task in the pure block.

add three to a two-digit number and the second task was report the antonym of a common word (Figure 9.1, middle column). An experiment by Spector and Biederman (1976) replicated this effect and also found that the effects of mixing addition and subtraction within blocks were reduced when a +3 or –3 were placed alongside each stimulus as a cue to indicate which task should be performed (Figure 9.1, right column). Spector and Biederman concluded that uncertainty about which task to perform on a given stimulus was one source of the switch cost, suggesting that the cost will be attenuated to the extent that an external cue is available to indicate which task should be performed on a particular stimulus. Thus, switch costs are greatest when stimuli are compatible with either task and when no cue is provided to signal which task is appropriate (Figure 9.1, left column), are reduced when a disambiguating cue is provide (Figure 9.1, right column), and are minimized when the stimulus unambiguously specifies which task to perform (Figure 9.1, middle column). In the absence of external cues, operators tend to rely on rehearsal in verbal working memory to remind themselves which task to perform on each new stimulus. Thus, when operators are prevented from talking to themselves, either out loud or subvocally, the costs of uncued task switches increase (Baddely, Chincotta, and Adlam 2001). The need to maintain multiple-task sets in memory may also contribute to general mixing switch costs shown in Figure 9.2, since additional memory load would be present even on task repetition trials. In dual-task, divided-attention paradigms discussed in the previous two chapters, this cost is sometimes referred to as a *cost of concurrence*, referred to in chapter 7, Figure 7.5.

The need to remember or determine which task to perform on a give stimulus, however, does not entirely account for task switch costs. A transition from one task to another also appears to necessitate a task set reconfiguration, "a sort of mental 'gear-shifting'" (Monsell 2003, p. 136) that may including changing goals, activating new stimulus-response mappings (Rubinstein, Meyer, and Evans 2001), and adjusting the parameters of subordinate perceptual and attentional processes (Gopher, Armony, and Greenshpan 2000; Logan and Gordon 2001). This reconfiguration accounts for at least part of the specific switch cost. The time needed for reconfiguration can be estimated using a procedure in which the order of tasks varies randomly and a cue is presented before each target stimulus to indicate which task should be performed. Data from such experiments indicate that switch costs increase with task complexity. It takes longer, for example, to establish the proper mental configuration for a task with difficult stimulus-response mappings than for a task with simpler or more natural mappings (Rubinstein et al. 2001). Conversely, switch costs tend to decrease as the interval between the cue and target grows longer. However, even when the operator is given a long preparation period, the switch cost is not entirely eliminated. These results suggest that the operator can begin task-set reconfiguration when cued but that the process cannot be completed until the target stimulus arrives to provide an exogenous event-driven trigger (Meiran 1996; see also Rogers and Monsell 1995).

Because task-set reconfiguration is a responsibility of the executive attentional system, it is hindered by a secondary task that also burdens executive attention processing—for example, a task that requires the juggling of information in working memory or retrieval of information from long-term memory (Baddeley et al. 2001). Switch costs can likewise be inflated by interference from carryover activation of a task set that is no longer appropriate, or by carryover inhibition of a set that was previously inactive but is now required—effects that have been labeled task-set inertia (Allport, Styles, and Hsieh 1994). The phenomenon of task-set inertia implies that even after an interrupting task has been completed, the act of having performed it may continue to hinder the ongoing task, contributing to general switch costs of the form just described. Such an effect parallels the well-known effect of proactive interference in memory, in which earlier but no longer relevant material continues to interfere with the ability to learn and to remember later arriving material.

Importantly, the phenomenon of switching costs scales up very nicely from the basic laboratory research to more applied environments. For example, Wickens, Dixon, and Ambinder (2006) observed relatively large (> 1 second) costs as pilots switched between subtasks of controlling and supervising two simulated unmanned air vehicles. These costs were inferred from the difference in task times when completed singly and in combination, much like the psychological refractory period research described in chapter 2.

Task Interleaving

Switch costs are but one, albeit important, component of task management, when divided attention between two or more tasks, involves switching back and forth between them—that is, interleaving between two tasks. A related phenomenon concerns the circumstances of return back to one task (A) after having dealt with another (B). Returning to the interrupted task, the operator may resume performance at the point where it was interrupted or may pick up the task at an earlier or later point. For example, a skilled musician who is interrupted in midphrase might return to the beginning of a phrase or perhaps even to the beginning of the piece when resuming his or her performance. Alternatively, as in the tragic aircraft accident in Detroit described at the outset of the chapter, an operator might resume performance of the interrupted task beyond the point at which it was suspended, omitting an intended step. Importantly, the double switch (A‡B‡A) may be a natural part of time sharing, or time swapping, two intended tasks, as switching gaze in a repeated cycle between the roadway and a head-down display; alternatively, B may be a specific event-driven interruption (Trafton 2007; McFarlane and Latorella 2002). Finally, earlier we spoke of returning eventually to the original task. However, sometimes this return may never take place at all, which is considered a failure of prospective memory (Dismukes and Nowinski 2007; Harris and Wilkins 1982; McDaniel and Einstein 2007), a form of memory that describes remembering to do something in the future.

A context for discussing the research findings on task interleaving is presented in the ongoing-interrupting task (OT-IT) diagram shown in Figure 9.3. Here the operator is performing some ongoing task, shown by the two OT boxes at the top of the figure in panel A. This task is interrupted by an interrupting task. The arrival of the IT may or may not be announced, and in either case it may take some time before the operator disengages from the OT to initiate the IT—or, of course, the operator may ignore the IT entirely, a strategy not shown in Figure 9.3. Once the operator leaves the OT, there may be some attention-switch time before the IT can be performed; this is called switch 1. Correspondingly, after completion of the IT, there may be an attention-switch return time before the OT can be resumed; this is called switch 2. Trafton (2007) referred to these two times as the interruption lag and the resumption lag, respectively. The amount of time that the IT is performed is analogous to the dwell duration discussed in chapter 4—the delay in return being analogous to the first passage time. Once attention is returned to the OT, there may be an initially degraded quality of performance, as it may take time to get engaged again if we forgot where we left off the OT when we departed, or we may make mistakes on return that are a carryover from the IT. Finally, at the bottom of the figure in panel B, we show the total time to do the OT within the dual-task, or task-switching, context and compare this with the time it would have taken to complete the OT in single-task conditions. The difference is the cost of interruptions, a cost that has been

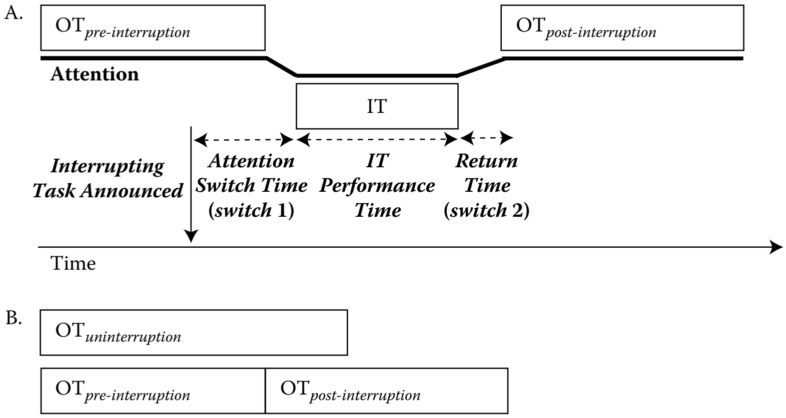

Figure 9.3 Effects of an IT on performance of an ongoing task OT. Panel A presents the time course of task performance, with the switch from and back to the OT. The total time needed to perform the interrupted OT includes the time spent on the task preinterruption; the time necessary to switch attention from the OT to the IT, to complete it, and to switch attention back to the OT; and then the time needed to resume and complete the OT. As depicted in panel B, the total time needed to perform the OT itself may increase under interrupted relative to uninterrupted conditions even when excluding attention switch times and IT task performance. This increase is a result of the time needed to resume the OT after returning from the IT.

well documented in the literature (see Trafton 2007). So, too, has been documented the frequency of interruptions and task switching in the workplace such as the aircraft cockpit (Dornheim 2000). Gonzales and Mark (2004), for example, found that task switching of information workers took place about every three minutes, and Wolf et al. (2006) observed that nurses were interrupted in their activities at least every twenty minutes. McFarlane and Latorella (2002) describe a large number of factors that influence the fluency of interruption management.

As noted already, the representation in Figure 9.3 can be applied in one of two contexts. On one hand, it may characterize the interleaving of two relatively ongoing tasks so that the distinction between IT and OT is somewhat arbitrary (e.g., the cycle shown in Figure 9.3 continues); this may characterize map checking while driving. On the other hand, it can characterize the response to a particular single interruption, such as a phone call while working on a word processor. The following section describes the effects on this process from the standpoint of the second of these contexts, recognizing that this analysis can generalize to the first as well.

Interruption Management

One can analyze the processes in panel A of Figure 9.3 in terms of four sequential task properties: those that influence attention switch 1 (interruption), that are characteristic of the OT and the IT, and those that effect switch 2 (return) that are properties of the IT and the OT. Each of these is described in turn.

Switch 1: OT Properties

Three factors directly influence the likelihood and speed of leaving the OT to deal with an IT—or, in the interleaving case, the tendency to stay longer on an OT before switching.

Engagement and Cognitive Tunneling on the OT

Though it is intuitively obvious that a task will be less susceptible to interruption when the operator is highly engaged in it, this effect has been somewhat difficult to capture experimentally or parametrically for the purposes of modeling. Interest plays a role. Interesting tasks are engaging, and we are reluctant to leave them, just as boring tasks can be easily interrupted. The role of interest is revealed by the results of a recent meta-analysis of research on cell-phone disruption of driving performance. Horrey and Wickens (2006) found that studies using actual conversations tended to disrupt driving more than those simulating the information-processing demands of such conversations with cognitive tasks but that were less engaging. That is, the latter tasks did not involve interesting semantic content, whereas the former were generally explicitly designed to attract the interest of the participant (e.g., discussing current topics, personal stories).

In addition to interest, at least two other sources of cognitive tunneling can be identified. First, the phenomenon has been observed in highly immersive realistic three-dimensional displays (Wickens 2005b). For example, a form of cockpit navigational display known as the three-dimensional highway in the sky (Alexander, Wickens, and Hardy 2005) was found to engage head-down attention to a degree that pilots sometimes failed to notice critical events that took place in the outside view of the world but that were not rendered on the engaging three-dimensional virtual reality display in the cockpit. Such failures were not observed when pilots flew with more conventional two-dimensional flight instruments (Wickens 2005b). Second, cognitive tunneling has also been observed during critical problem-solving or trouble-shooting operations, with people failing to notice other important events (Dismukes and Nowinksi 2007; Moray and Rotenberg 1989). This phenomenon was manifest in the Eastern Airlines Everglades crash in 1972, when pilots, trying to diagnose a landing-gear problem, were entirely unaware of a salient auditory alert that signaled their impending crash into the ground.

Strategic Factors and the Stopping Point

Operators may choose to remain with an OT for a while before switching until they get to a stopping point, perhaps as a way of resolving some ongoing subtasks and thereby of avoiding the need to hold information in working memory during the IT period or avoiding the need to reacquire forgotten information when the OT is resumed. Altmann and Trafton (2002) and McFarlane (2002) modeled the OT in terms of hierarchical goal structures, noting that the OT is more likely to be left after a goal is completed than in the middle of a completion sequence (e.g., carrying out a sequence of programming subtasks that lead to a task goal, like finishing writing a sentence or paragraph). The OT is more rapidly resumed on return when it was left at a good stopping point (e.g., between tasks or subtasks) than in the middle of a subtask (Bailey and Konstan 2006; Monk, Boehm-Davis, and Trafton 2004). It is also better resumed when a delay is taken after the interruption and before the OT is left—for example, a large switch 1 or interruption lag (Dismukes and Nowinski 2007).

Presumably part of the benefit of this delay—the interruption lag—is that it allows people to better encode the state of OT into memory before leaving it so that the point of departure will be well remembered and, therefore, the point of return easily found (Trafton, Altman, and Brock 2005). In this light, it is noteworthy that when people intentionally take a longer time with switch 1, this strategy of prolonging switch 1 will have beneficial effects on return, giving people time to think about where they left off or possibly placing a visible reminder in the OT workplace (Dismukes and Nowinski 2007). Also, designers of computing systems are considering ways in which human cognitive operations can be monitored to adaptively impose interruptions only at optimal stopping points during ongoing tasks (Bailey and Konstan 2006).

A closely related task characteristic found to influence switching strategy is the working memory demand of the OT. For an auditory working memory task, like dialing a long phone number just heard on an answering machine, there should be reluctance to leave it until the dialing task is completed because some of the dialed digits may have been forgotten after an interruption-induced switch. This would not be the case when dialing from a visual phone number on a device with visual feedback of the dialed digits. Here it is noteworthy that visual—rather than auditory—presentation of complex information allows better task management and more optimal task switching (Wickens and Colcombe 2007b) because with a permanent visual text display of an OT, switch 1 can occur without fear that the information will be gone on return. This is, of course, not the case with working-memory-challenged auditory information.

In dynamic-process or vehicle-control tasks, operators should be more reluctant to abandon the OT control task if the systems involved are high bandwidth or unstable; they also may choose to switch at moments of time when the system is most stable (e.g., the car is in the center of the lane and on a forward trajectory, and the road ahead is straight). In this case, how much neglect will lead to a significant state change depends on the inertia and dynamics of the system (Wickens 1986, 2003). A large aircraft flying in still air will allow longer neglect than a light aircraft flying in turbulence. Furthermore, some tasks, like flying a helicopter or riding a bicycle, are inherently unstable and lead to time-dependent divergence if neglected. Finally, it should be noted that the consequences of not following a deferred switch strategy and therefore of abandoning the OT (switch 1) in the midst of a goal pursuit (subtask), of a working memory-loading operation, or in an unstable dynamic situation will be realized on return to the OT, as discussed below (Monk, Boehm-Davis, and Trafton 2004).

Importance and Priority

Like engagement, it is also intuitive that an OT with higher priority than an IT should sustain its performance longer in the face of an interruption and generally be less interruptable than a task of lower priority. This effect was found by Iani and Wickens (2007) as a primary flight task was interrupted by the delivery of discrete weather information. Turning to ground transportation, the fact that driving is as safe as it is—despite all of the multitasking that goes on, much of it head down—suggests as well that people's switching strategies tend to prioritize out the window viewing over head-down activity (Wickens and Horrey in press), as if the out-the-window tasks are less interruptible. However, evidence that such prioritization does not always hold comes from the numerous examples in which this sampling strategy fails and in which accidents occur as a result of in-vehicle distractions (Dingus et al. 2006). One important issue addressed in chapter 3 is the ability of operators to know the priority of the IT before the full interruption has taken place and the OT has been abandoned. This points to the importance of preattentive referencing in alarms (Woods 1995) as discussed in chapter 3, and systems that contain preattentive referencing have been found to be effective in task management (Ho et al. 2004).

Switch 1: IT Properties

Importance

As with the OT properties, the IT importance is again a relevant factor for switch 1, as it is indeed the relative importance between the two tasks that matters most. The less important IT will prolong switch 1. The role of relative importance in governing scanning between head-up and head-down driving tasks was well documented by Horrey, Wickens, and Consalus (2006), and chapter 4 showed for visual tasks the prominent role of task value, or importance, in the SEEV model.

Salience

As chapter 3 demonstrated, the salience of stimuli and events that accompany the IT will prominently influence the switching speed. In particular, auditory events are typically more salient than visual ones and will lead to faster switching speeds, a phenomenon referred to as the *auditory preemption* effect (Ho et al. 2004; Wickens and Colcombe 2007b; also see chapter 8 in this volume)—although this effect is not always observed, particularly when visual events have direct access to foveal vision. Another distinction of IT salience is between what we call announced and unannounced interrupting tasks. The announced tasks have a stimulus event associated with it—typically visual or auditory—and these are clearly more salient than unannounced ones in which IT initiation must depend on some form of prospective memory. The unannounced tasks, like remembering to turn the heat down before the pot boils over rather than afterward, are more likely to get missed or delayed.

Switch 2: IT Properties

The case can be easily made that properties of the IT that will lead people to delay switching back to the OT are very much the same properties that lead people to stay on the OT prior to switch 1, as discussed previously. Indeed, there is a modest, but imperfect, reciprocity between IT and OT properties in their influence on switching performance (Iani and Wickens 2007).

Switch 2: OT Properties.

Finally, we address characteristics of the OT that influence the circumstances of return to it after switch 2.

Strategies that Were Carried Out at Switch 1

Here research indicates that the most important variable is that associated with the strategies adopted on leaving the OT for switch 1, as discussed already, and there are several of these. For example, to the extent that OT was left in the middle of a subtask—rather than between subtasks—the OT is more likely to be degraded on return (Miller 2002; Monk, Boehm-Davis, and Trafton 2004). Indeed, Miller (2002) found that the interruptions were often so disrupting that some OTs needed to be restarted from the beginning such that the time taken to complete OT postinterruption (Figure 9.3) was fully as long as OT uninterrupted; that is, nothing was accomplished by OT preinterruption—operators needed to start from scratch. As noted already, an important strategy at switch 1 that affects performance in OT resumption is the placement of an intentional delay at switch 1. This delay can be used to accomplish either of the following beneficial acts: (1) It can be used to rehearse the state of the OT at the interruption (Trafton et al. 2003); and (2) it can be used to provide an explicit, usually visual, placeholder such as the mark on the page where text editing was left off when answering a phone call. Both of these actions are helpful (Dismukes and Nowinzki 2007; McDaniel and Einstein 2007).

Delay in Return

Delaying the resumption of the OT can degrade the quality of return in two ways. First, there will be a simple decay of working memory, of where the task was left off; or if the OT had loaded working memory at switch 1, then there will be a decay of the material in the task itself (remember the phone dialing example). Second, if the OT is a dynamic one, like vehicle control, the delay or neglect may increase the likelihood that the system itself will have evolved toward an unstable or undesirable state while attention, and therefore control, was absent. Thus, a car will be more likely to have diverted toward the ditch the longer the head stays down. As described in the context of scanning in chapter 4, a long first-passage time away from the OT leads to increasing vulnerability (Horrey and Wickens 2007; Sheridan 1970). Third, the longer one stays on an IT, the greater is the possibility that the OT will have been forgotten entirely: Its goal memory will have decayed below threshold (Altmann and Trafton 2002), and a failure of prospective memory will have occurred (Dismukes and Nowinski 2007; McDaniel and Einstein 2007).

IT–OT Similarity

Finally, as a property not associated exclusively with either the IT or the OT, Gillie and Broadbent (1989), Dismukes and Nowinksi (2007), and Cellier and Eyrolle (1992) all documented the degrading role of similarity between OT and IT on the resumption of the OT. This effect appears to reflect the same information-processing mechanisms described in chapter 8 related to confusion and cross-talk. Greater similarity between the OT and the still-active IT will cause greater confusion and interference from the still-active traces of the IT when the OT is reinitiated (i.e., proactive interference) and greater disruption of relevant OT information that needed to be retained during the IT period (i.e., retroactive interference). This will be particularly true if the OT required retention of material (i.e., rehearsal) during the OT1–OT2 interval given the resource demands of working memory.

Task and Workload Management

The area of strategic task or workload management integrates the switching and interruption findings discussed already and places them in a broader context. This context might be represented as zooming out away from the single OT‡ IT‡OT element to consider lots of tasks in sequence so that the distinction between OT and IT is blurred. Historically, this approach received a major boost when Hart (1989) and Hart and Wickens (1990) noted that all of the extensive work being done on measuring mental workload (see chapter 7) and evaluating multiple-task interference through resource models failed to account for how people managed multiple tasks when workload was so excessive that concurrent processing was impossible. Thus, research focus was required on how people deal with these overload situations.

Freed (2000) proposed a reactive prioritization model to account for task management with similarities to Sheridan's (1970) earlier modeling on supervisory sampling of input channels (for a good review, see Moray 1986; see also chapter 4 in this volume). Freed (2000) considered four factors that should optimally influence decisions to switch between tasks:

(1) Urgency: How long is it until the deadline by which a task must be completed, and how long does it take to complete the task? For example, five minutes remaining until the deadline of a four-minute task has an urgency of one minute. If one minute passes and the task is not initiated, it will be too late to finish it on time.
(2) Importance: What is the cost of not doing the task? It is acceptable not to switch to an urgent (factor 1) task if there is no cost in missing its deadline. This factor will lead to a greater proportion of time spent doing more important tasks (Raby and Wickens 1994).
(3) Duration: Longer-duration tasks will, of course, increase the urgency if not yet performed, but they will also be more likely to disrupt performance of other tasks once they are initiated, assuming that there is a task-switching cost to leaving them temporarily uncompleted (factor 4). This penalty will make operators reluctant to leave the longer-duration task.
(4) Switching or Interruption Cost: This concept has been discussed repeatedly already. A high switching cost will lead to task inertia and a likelihood of continuing without a switch (Tulga and Sheridan 1980; Ballard, Mayhoe, and Pelz 1995). Note, by the way, how this cost of switching is tied to the concept of information-access effort or cost, as discussed in chapters 4 and 7 in this volume. Greater distance between visual sources of task information will lead to greater costs of switching.

Added to this mix is, of course, the role of uncertainty. Sometimes we do not know before switching how long it will take to do the task once initiated, nor do we always know about the impending arrival of additional tasks. In the latter regard, Tulga and Sheridan (1980) demonstrated the value in optimal task management of knowing in advance (i.e., preview) what arriving tasks will be and how long they are likely to take.

Freed's (2000) model does not appear to have been fully validated regarding the extent to which people follow its optimal prescriptions or the circumstances that make one factor dominate over others. Tulga and Sheridan (1980) did provide some data, but only using very generic pseudotasks: computer-displayed bars that can be activated to simulate attention directed to the bars. Raby and Wickens (1994) performed an empirical task-management study with trained airplane pilots to examine some aspects of optimal scheduling. Their pilots flew a simulated approach in a realistic airplane simulator under three conditions of increasing workload, varied by the amount of time pressure on the pilots to complete all the tasks necessary to plan for and to accomplish a landing at an unfamiliar airport. Prior to the study, the nineteen tasks to be performed were categorized into three categories of priority, or importance, labeled *must, should,* and *can.* These corresponded somewhat to the more generic aviate-navigate-communicate-system management task importance hierarchy traditionally used in aviation (Schutte and Trujillo 1996). The investigators then evaluated the overall flight quality (i.e., precision of flying) as well as the specific timing and performance of the three categories of tasks. The results revealed that pilots generally behaved appropriately, as characterized by the following:

- Performing tasks at more optimal times, prioritizing must tasks over should tasks, and prioritizing should tasks over can tasks
- Abandoning, or shedding, can tasks more frequently than should tasks as workload increased and abandoning should tasks more frequently than must-do tasks

However, Raby and Wickens (1994) found that when workload increased, pilots did not not optimally reschedule the higher-priority tasks in response to the dynamic workload change. From this they concluded that pilots do not maintain perfectly optimal strategies for the plausible reason that the task scheduling itself demands resources that should otherwise be devoted to performing the tasks themselves. Such a conclusion is consistent with Kahneman's (1973) effort-conserving view of heuristics (chapter 7 in this volume) and with the results of another scheduling study carried out by Moray et al. (1991).

Another aspect of Raby and Wickens's (1984) study examined differences in task-switching behavior between the better-performing and worse-performing pilots. Three conclusions emerged here regarding better-performing pilots, both also supported by other research:

(1) They tended to be more proactive, initiating high-priority tasks earlier (see also Laudeman and Palmer 1995; Orasanu and Fischer 1997). Procrastination hurts. However, too much early preparation in an uncertain world can be counterproductive, especially if formulated plans are rigidly maintained despite changing circumstances. This issue of the plan continuation error is beyond the scope of the current chapter (McCoy and Mickunas 2000).
(2) They tended to switch attention more rapidly between tasks—hence being less inclined toward cognitive tunneling or task inertia.
(3) They tended to be more flexible in attending to tasks as their priority may momentarily change (Schutte and Trujillo 1996).

Conclusion

In conclusion, executive control and task management play the dominant role in multiple-task performance, when parallel processing is no longer possible and when serial processing predominates. This parallels the role of the resource-allocation policy, discussed in parallel processing in the previous two chapters. There is some convergence in results describing these task-management processes when we look at the three domains of research covered here: (1) basic attention switching; (2) more complex interruption management; and (3) the complex level of real-world task and interruption management. However, more research in the third domain is certainly needed to better understand task-management strategies here. All three domains point to the costs of attention switching and interleaving compared with staying on a single, uninterrupted task. All point to the role of strategies and working memory in managing these multiple tasks well.

Also emerging from the literature are findings that certain task- and display-related factors *drive* task management in nonoptimal ways, analogous to the nuisance properties of salience and effort in the SEEV model discussed in chapter 4, whereas other, more optimal, top-down properties can be achieved by the well-calibrated task manager. More optimal task management is possible to the extent that the operator has the information and knowledge available to know things like calibrated task importance, expected value (risk of failing to do the task), duration, and arrival time. This knowledge is analogous to the E and V properties of the SEEV model. In particular, this latter class of top-down knowledge-driven factors in task management points to the possible role of attention and task-management training in improvement of multitasking, resource allocation, and interruption management (Dismukes 2001; Gopher 1992; Hess and Detweiller 1994), an issue addressed in the next chapter.