

30

Modern Computational Perspectives on Executive Mental Processes and Cognitive Control: Where to from Here?

David E. Kieras, David E. Meyer, James A. Ballas,
and Erick J. Lauber

ABSTRACT Formal concepts and algorithms from contemporary computer operating systems can facilitate efforts to precisely characterize the supervisory functions of executive mental processes. In particular, by helping to advance work with the "executive-process interactive control" (EPIC) architecture, a theoretical framework for computational modeling of human multitask performance, operating system fundamentals provide insights about how people schedule tasks, allocate perceptual-motor resources, and coordinate task processes under both laboratory and real-world conditions. Such insights may lead to discoveries about the acquisition of procedural task knowledge and efficient multitasking skills.

Following the cognitive revolution in scientific psychology (circa 1950), many experimental psychologists and cognitive scientists have assumed that human cognition shares fundamental similarities with symbolic information processing by electronic digital computers (Lachman, Lachman, and Butterfield 1979; Newell 1990). Although the operations of such computers are serial in some respects, they can emulate parallel processing of multiple information streams and implement algorithms for modeling the performance of perceptual-motor and cognitive tasks. As a result, the computer metaphor has inspired significant discoveries about perception, attention, learning, memory, language, and problem solving. Furthermore, as computational hardware and software continue to evolve, the computer metaphor may become increasingly apt.

Encouraged by this prospect, our work has focused on characterizing executive mental processes with a particular theoretical framework, the "executive-process interactive control" (EPIC) architecture. Using EPIC, we have formulated precise computational models of human multitask performance under both laboratory and real-world conditions (e.g., Kieras and Meyer 1997, forthcoming; Meyer and Kieras 1997a,b, 1999). EPIC models account well for quantitative data, predict new phenomena, and point toward promising directions for future research on cognitive control.

The functions of executive processes in EPIC correspond closely to ones provided by a computer operating system (OS) that supports parallel information processing for concurrent execution of multiple task pro-

grams (Stallings 1998). This correspondence suggests that studying the fundamentals of contemporary OSs may facilitate the development of EPIC. Such study may also advance the conceptualization of executive mental processes in other theoretical frameworks (e.g., Baddeley 1986; Braver and Cohen, chap. 31, this volume; Kimberg and Farah, chap. 32, this volume; Norman and Shallice 1986), thereby helping to banish the "homunculus" of cognitive control about which previous pundits have complained vociferously (e.g., Newell 1980; Neisser 1967).

In our opinion, the modern computer metaphor is relevant to answering several related questions: Do people have general executive processes that are used across many contexts? Exactly what functions do these processes serve? How might they influence the representation and acquisition of procedural task knowledge? Are there task-specific aspects of cognitive control for which general executive processes must be supplemented through special training? Which experimental procedures are especially suited for eliciting and analyzing particular control operations? Does the human brain really implement the types of function that an OS provides?

Toward answering these questions, section 30.1 introduces EPIC; section 30.2 describes results from applications of EPIC to modeling multitask performance and characterizing particular executive mental processes; section 30.3 presents additional relevant concepts from contemporary computer technology and OSs; section 30.4 discusses how these concepts may promote research with EPIC and guide theorizing about cognitive control; and section 30.5, in summarizing our conclusions, offers final thoughts on the directions of future research.

30.1 THE EXECUTIVE-PROCESS INTERACTIVE CONTROL ARCHITECTURE

Extending proposals by previous theorists (e.g., Anderson 1983; Card, Moran, and Newell 1983; Newell 1990), we have designed EPIC to integrate cognitive and perceptual-motor operations with procedural task analyses of skilled performance.

Components

EPIC has a central cognitive processor with a production-rule interpreter and a multipartition working memory (WM) surrounded by peripheral sensors, perceptual processors, motor processors, and effectors that all operate in parallel. These permanent interconnected components constitute EPIC's "hardware." Each perceptual and motor processor functions as a distinct limited-capacity channel of input or output. Task performance is modeled by programming the cognitive processor with production rules that make decisions and generate responses based on the

contents of WM. The production rules, stimulus codes, and response codes may vary depending on specific task requirements.

Consistent with basic periodicities of human information processing (Kristofferson 1967), EPIC's cognitive processor operates in cycles that have stochastic durations whose mean is 50 msec. While doing so, the cognitive processor enables a high degree of parallelism in multitask performance. On each cycle, its production-rule interpreter tests the conditions of all rules in procedural memory, and executes the actions of all rules whose conditions match the current contents of WM. There is no set limit on how many rules can be applied simultaneously. Thus cognitive processes involving distinct sequences of rules may progress simultaneously, sharing system resources as time passes.

Basics of Control

The flow of information processing in EPIC is controlled with production rules like the following one, which selects and initiates a manual "poke" response to a red target stimulus during a tactical decision task (Kieras and Meyer 1997, forthcoming; Meyer and Kieras 1999):

```
IF
  ((GOAL DESIGNATE-TARGET-FOR-TACTICAL-TASK)
   (STRATEGY MAKE-POKE-RESPONSE-IMMEDIATELY)
   (STEP MAKE-POKE-RESPONSE)
   (TAG ?OBJECT IS STIMULUS)
   (VISUAL ?OBJECT COLOR RED)
   (NOT (VISUAL ??? SIZE LARGE))
   (STATUS TACTICAL-TASK-PROCESS-HAS-EYE)
   (MOTOR MANUAL PROCESSOR FREE))
THEN
  ((SEND-TO-MOTOR-MANUAL-PROCESSOR PERFORM-POKE-(LEFT INDEX)
   ?OBJECT)
   (ADDWM (GOAL WATCH-FOR-DESIGNATION-EFFECT))
   (DELWM (STEP MAKE-POKE-RESPONSE))
   (ADDWM (STEP WAIT-FOR-WATCHING-DONE)))
```

Sequential Rule Execution As illustrated here, EPIC production rules have conditions and actions that contain goal and step items. Adding and deleting step items in working memory enables the rules to be executed in particular sequences. For example, the preceding rule would be enabled by putting "STEP MAKE-POKE-RESPONSE" in WM with an add-to-WM (ADDWM) action. Taking this item out of WM with a delete-from-WM (DELWM) action would disable the rule, and then putting "STEP WAIT-FOR-WATCHING-DONE" in WM would enable another subsequent rule. Because information in WM is subject to loss or corrup-

tion, errors of sequencing may occur under EPIC, as they do under real-world circumstances.

Subroutine Calls Using the same goal item in a set of EPIC production rules lets them function like a computer program subroutine. The subroutine would be "called" by adding its shared goal item to working memory. After the call, a start-up rule in the subroutine would "fire" and add its first step item to WM. When the subroutine finishes, its termination rule would delete the routine's goal and last step items from WM, and signal that the subroutine has finished. For example, the preceding rule calls a subroutine for watching the visual effects of the manual poke response. This entails adding two items to WM: "GOAL WATCH-FOR-DESIGNATION-EFFECT," which is the goal item for the subroutine; and "STEP WAIT-FOR-WATCHING-DONE," which is used by another rule that waits for the subroutine to be completed.

Interrupts Thus EPIC implements capabilities analogous to computer interrupts. A production rule can have conditions such that it waits for a certain future event to occur regardless of other intervening activities. When these conditions are satisfied, the rule may start the execution of other rule sequences to deal with the interrupting event.

Task Processes

Procedural knowledge for performing tasks is represented by EPIC production rules that fire in particular sequences. Our models embody programming-style principles like those applied in computer software design. Each task and subtask has a set of rules with standard formatting of control items and input/output (I/O) information. Standard protocols are used for task start-up, completion, error detection, abort, and restart procedures.

Executive Processes

In modeling multitask performance, we formulate distinct sets of supervisory production rules that implement supraordinate executive processes, whose function is to add and delete working-memory items for controlling the execution of various task and subtask procedures. Under EPIC, an executive process may suspend a task process by deleting its goal item from working memory, and then resume the task process by adding its goal item to WM again. Similarly, an executive process may use strategy items to instruct a task process about which of several alternative paths to take. These control operations can be accomplished through rules whose conditions match status items that the task process adds to WM along the way.

30.2 APPLICATIONS OF EXECUTIVE-PROCESS INTERACTIVE CONTROL TO MULTITASK PERFORMANCE

To illustrate more fully how we characterize executive mental processes, this section describes four cases of several for which EPIC models of multitask performance have been developed: (1) discrete successive tasks; (2) discrete concurrent tasks; (3) elementary continuous tasks; and (4) compound continuous tasks. From them, it will become clearer how EPIC enables task coordination and scheduling to be described under a variety of conditions. Also, the stage will be set for examining cognitive control from the perspective of computer operating systems.

Discrete Successive Tasks

In the discrete successive-tasks procedure, also known as "task switching," participants either alternate between two different choice-reaction tasks or perform one task repeatedly during a series of discrete trials, with a *response-stimulus interval* (RSI) separating each response from onset of the next stimulus. Reaction time (RT) and accuracy are measured as a function of trial type, RSI, and other factors. Switching time costs (STCs) are calculated from differences between mean RTs on alternating-task and repeating-task trials (for a review, see Pashler, chap. 12, this volume).

According to some theorists, executive mental processes contribute substantially to STCs (e.g., Meiran 1996; Rogers and Monsell 1995; Rubinstein, Meyer, and Evans forthcoming; see also Goschke, chap. 14, De Jong, chap. 15, Meiran, chap. 16, and Keele and Rafal, chap. 28, this volume). Following their lead, we have formulated an EPIC model to account for some results from the successive-tasks procedure. The details of this formulation concern both the representation of procedural task knowledge and the cognitive control of task switching.

Lauber 1995 For now, our model deals with data from Lauber 1995 (exps. 4 and 5), which varied response-stimulus intervals, stimulus-response compatibility, and practice orthogonally. Additive and interactive effects of these factors strongly constrain the type of model that may account for them.

Twenty undergraduate students participated in Lauber's study. They were divided into two groups that performed basic choice-reaction tasks with different S-R mappings. Members of each group were tested individually during three 1-hour sessions. The stimuli for each task were printed digits. The responses were keypresses made with fingers of the right hand. Stimuli and responses were paired to form four alternative S-R mappings, each of which was used in one of four different tasks: compatible task A, compatible task B, incompatible task C, and incompatible task D. For task A, the digits 1, 2, 3, and 4 were mapped respectively to

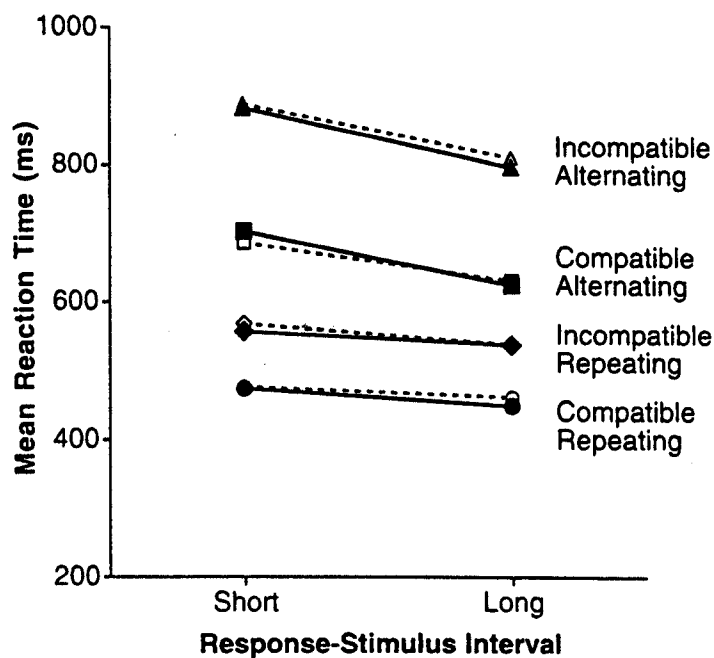


Figure 30.1 Second-task reaction times from the first session in Lauber 1995. The dark points connected by solid lines represent observed mean reaction times as a function of response-stimulus interval, task difficulty (compatible versus incompatible S-R mappings), and trial type (alternating-task versus repeating-task trials). The light points connected by dashed lines represent simulated mean RTs produced by the EPIC model in figure 30.2.

the index, middle, ring, and little fingers; for task B, this mapping was reversed. For task C, the digits 1, 2, 3, and 4 were mapped respectively to the middle, little, index, and ring fingers; for task D, this mapping was reversed.

During each test session, there were two types of trial block. One type contained a series of alternating-task trials, and the other contained a series of repeating-task trials. On each alternating-task trial, participants in group 1 performed task A followed by task B, or vice versa; on each repeating-task trial, they performed one of these tasks twice. A similar arrangement of tasks C and D was used for group 2. Before each trial block, subjects were told what their tasks would be. Each block included two RSIs, 50 and 750 msec, which varied randomly across trials. The intertrial intervals equaled 1 sec.

Empirical Results Figure 30.1 shows some results from the first session. Mean RTs of second-task (post-RSI) responses were reliably longer for alternating-task trials, incompatible S-R mappings, and short RSIs. Although some reliable two-way interactions occurred between these factor effects, S-R compatibility and RSI affected mean switching time costs almost additively. Furthermore, despite these effects, large switching time costs persisted after the longer RSI, as other investigators have found (e.g., Allport and Wylie, chap. 2, this volume; Allport, Styles, and Hsieh

1994; De Jong, chap. 15, this volume; Rogers and Monsell 1995).¹ It is this overall pattern for which EPIC accounts.

EPIC Models Of course, there are various ways that we could model task switching with EPIC. For example, one conceivable model would have two sets of task-specific, goal-sensitive production rules available simultaneously in procedural memory. In this case, the rules used to select responses for Lauber's incompatible tasks C and D might have the following forms:

```
IF
((GOAL PERFORM TASK C)
 (STEP MAKE PRESS-RESPONSE TO DIGIT 1)
 (VISUAL ?OBJECT DIGIT 1))
THEN
((SEND-TO-MOTOR MANUAL PERFORM PRESS (RIGHT MIDDLE))
 (DELWM (STEP MAKE PRESS-RESPONSE TO DIGIT 1))
 (ADDWM (STEP WAIT-FOR PRESS-DONE)))

IF
((GOAL PERFORM TASK D)
 (STEP MAKE PRESS-RESPONSE TO DIGIT 1)
 (VISUAL ?OBJECT DIGIT 1))
THEN
((SEND-TO-MOTOR MANUAL PERFORM PRESS (RIGHT RING))
 (DELWM (STEP MAKE PRESS-RESPONSE TO DIGIT 1))
 (ADDWM (STEP WAIT-FOR PRESS-DONE)))
```

Given the simultaneous availability of such rules, an executive process could switch tasks simply by changing the task goal items in working memory, disabling one task's rules and enabling the other's.

Yet this type of model would fail to account for persistent large switching time costs such as Lauber observed. Under EPIC, changing goal items takes only one cognitive-processor cycle, which should be completed within about 50 msec regardless of other prevailing factors. However, Lauber's STCs ranged from 200 to 300 msec, they endured after a relatively long (750 msec) RSI, and S-R incompatibility affected them reliably. Thus additional delays associated with other control operations besides changing goal items presumably contributed to task switching here. Perhaps these contributions occurred because the tasks had different S-R mappings but involved the same stimuli and responses. Such mapping conflicts might substantially increase the amount of practice needed to learn adequate task-specific, goal-sensitive production rules (Anderson 1983), requiring participants to rely initially on other types of procedural and declarative knowledge instead.

Thus our modeling of Lauber's results has taken an alternative direction. Consistent with some other theorists (e.g., Rubinstein, Meyer, and

Evans forthcoming), we assume that to reduce conflicts in switching between similar tasks, five constraints are imposed: (1) at each moment, symbolic S-R mapping information for performing just one task is kept in WM; (2) switching tasks involves removing currently irrelevant information from WM; (3) the irrelevant information is replaced with relevant information for the next task; (4) these "cleanup" and "setup" operations entail relatively slow interactions with long-term memory; and (5) setting up for the next task is triggered by its stimulus onset.

On the basis of these assumptions, we have formulated a model with a single set of generic production rules that perform both of Lauber's incompatible tasks. For each incompatible task, these rules select responses by using a particular list of S-R pairs in WM. This involves checking the stored S-R pairs serially to find one whose stimulus term matches the presented stimulus (cf. Theios 1973). When the match is found, its associated response term is sent to the manual motor processor. Given this protocol, task switching requires not only changing task goal items but also retrieving the next relevant S-R pairs from long-term memory.

For performing both of Lauber's compatible tasks, our model has another set of generic production rules. They assume that EPIC's visual perceptual processor directly recodes each presented stimulus into two response symbols appropriate for the alternative compatible S-R mappings (e.g., "1" → "index finger" and "1" → "little finger"). A task rule then chooses and sends one or the other of these response symbols to the manual motor processor. This choice is made by referring to a WM strategy item that indicates which S-R mapping is currently relevant. Given this protocol, task switching requires not only changing task goal items but also retrieving the relevant strategy item from long-term memory.

These operations are controlled by an executive process that takes different paths for alternating-task and repeating-task trials (figure 30.2). At the start of repeating-task trial blocks, the executive process calls a subroutine that sets up WM to perform a particular task, and then lets this task be performed twice during each trial. In contrast, at the start of each alternating-task trial, the executive process waits until the first-task stimulus has been recognized, next calls the subroutine that sets up WM for the first task, and then lets the first task be performed. After the first-task response has been made, the executive process calls another subroutine that cleans up WM, waits until the second-task stimulus has been recognized, calls the setup subroutine for the second task, lets the second task be performed, and finally cleans up WM again. Fitting our model to Lauber's data required adjusting the times taken by the WM setup and cleanup subroutines.

Simulated Results Figure 30.1 shows the mean second-task RTs produced by our model, which accounts well for the main effects of trial

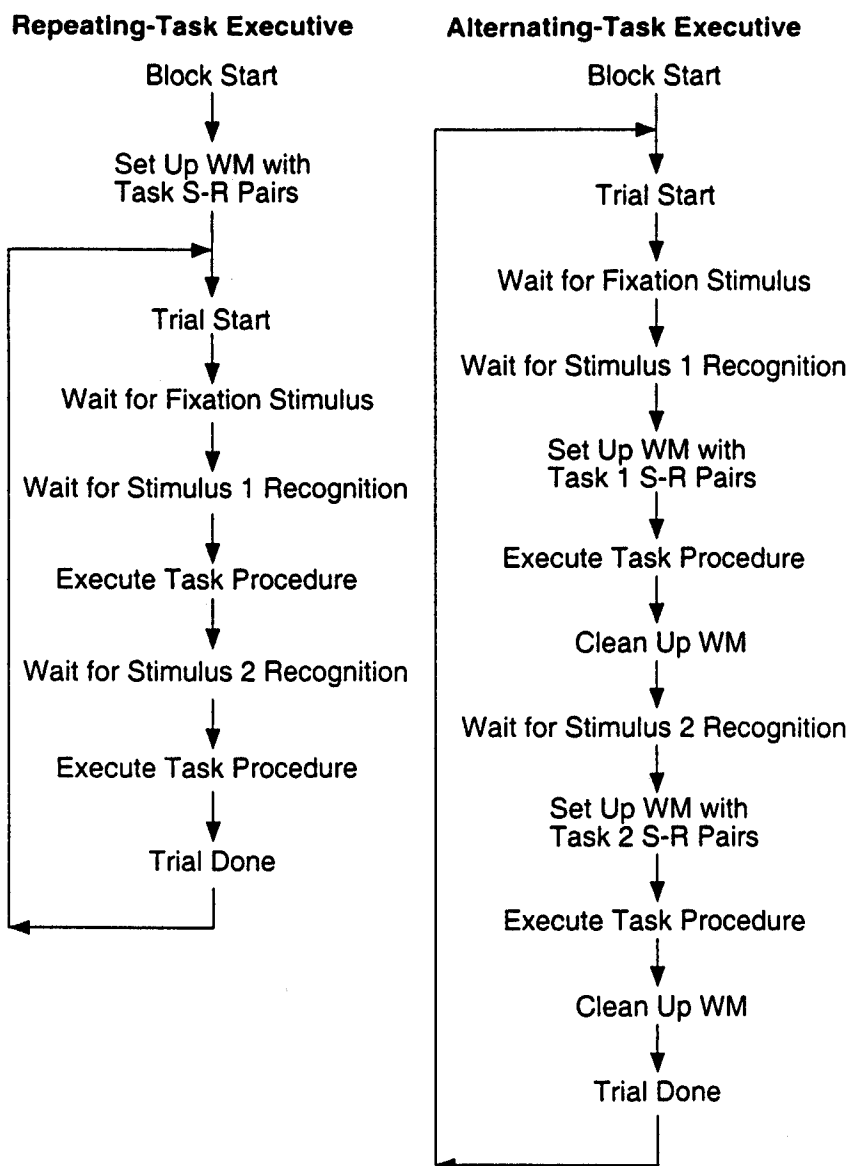


Figure 30.2 Flowchart of executive processes on repeating-task trials (left) and alternating-task trials (right) in the EPIC model for Lauber 1995. Mean reaction times produced by this model appear in figure 30.1.

type, RSI, and S-R mapping, as well as their additivities and interactions.² Our model succeeds much better than one that switches tasks simply by changing goal items in working memory.

Theoretical Implications The working-memory setup and cleanup operations that we needed to fit Lauber's data each took about 150 msec. Why so long? One possible answer is that in reality, these operations entail gradually activating relevant and inhibiting irrelevant symbolic long-term memory representations (cf. Allport, Styles, and Hsieh 1994; Anderson 1983; Goschke, chap. 14, this volume). This would explain why STCs persist at long RSIs and why WM setup is not started until the next

task's stimulus has been recognized. Perhaps the executive process waits to start setting up WM because stimulus recognition helps amplify requisite memory activation. At present, EPIC does not implement such activation explicitly. Thus supplementing EPIC with appropriate activation mechanisms could prove worthwhile.

From our present perspective (figure 30.2), however, the executive processes for task switching seem relatively simple. Other than calling WM setup and cleanup subroutines, they contribute very little to STCs. This is consistent with claims of Allport, Styles, and Hsieh (1994), who questioned whether task-switching studies reveal much about executive mental processes per se. Nevertheless, such studies could have further benefits in other respects. For example, they may yield new insights about the representation of procedural task knowledge, extending what we have discovered already through EPIC modeling.

Discrete Concurrent Tasks

A second context in which EPIC has enabled us to learn more about executive mental processes is the "psychological refractory period" (PRP) procedure (Pashler 1994, chap. 12, this volume). In this procedure, subjects perform two concurrent choice-reaction tasks during series of discrete trials. Typically the tasks involve different stimuli and responses. On each trial, a first-task stimulus is followed by a second-task stimulus. Because the stimulus onset asynchrony (SOA) is relatively short, the second-task stimulus may precede the first-task response. However, subjects are instructed to give task 1 higher priority, and they may be encouraged to make the first-task response before the second-task response. RTs and response accuracy are measured as a function of the SOA and other task factors. The PRP procedure interests us because, despite its task prioritizing and stimulus sequencing, there is potentially ample opportunity for tasks 1 and 2 to be performed at least somewhat in parallel. By formulating EPIC models under these conditions, we can better understand how such cognitive control is achieved.

EPIC Model For example, figure 30.3 outlines the executive process of a model that has been tested extensively in our research concerning the PRP procedure (Meyer and Kieras 1997a,b). Here the executive process puts tasks 1 and 2 respectively in "immediate" and "deferred" modes at the start of each trial. This is done by adding strategy items (e.g., "STRATEGY TASK 1 IS IMMEDIATE") to WM. Putting task 1 in immediate mode lets its responses be selected and sent to their motor processor as quickly as possible for movement production. While task 2 is in deferred mode, its production rules can select symbolic identities of second-task responses and store them in working memory, but the selected second-

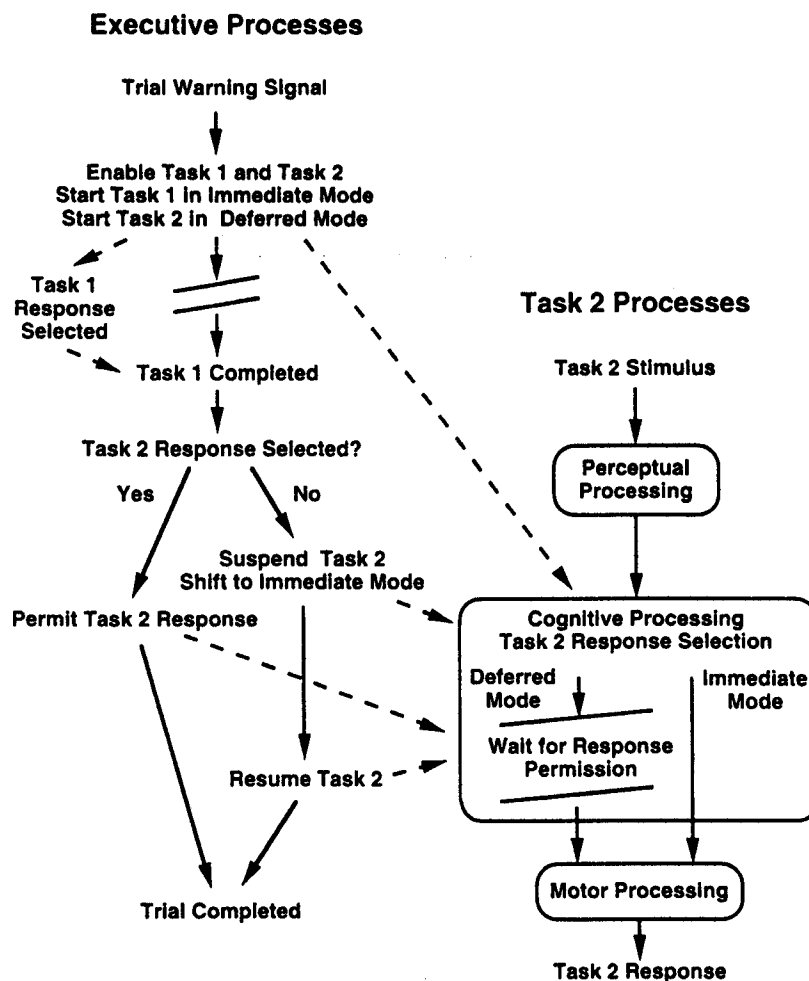


Figure 30.3 Flowchart of executive and secondary-task processes in the EPIC strategic response-deferment model for the psychological refractory period (PRP) procedure.

task response identities are not sent to a motor processor, and they are not produced as overt movements. When, however, a prespecified “unlocking event” occurs subsequently (e.g., the overt first-task response is initiated), the executive process shifts task 2 to immediate mode. Following this shift, previously selected second-task responses may be sent from WM to their motor processor for movement production. If response selection has not yet finished for task 2 before it is shifted to immediate mode, then subsequently the second-task production rules will both select and send the second-task responses directly to their motor processor.

Simulated versus Empirical Results Comparisons between simulated and empirical results from various studies with the PRP procedure have been encouraging. Our EPIC strategic response deferment model accounts accurately for differences between observed mean first- and second-task RTs as well as additive and interactive factor effects on them.

The model's goodness of fit is typically high ($R^2 > 0.95$) and involves only modest numbers of "free" parameters.

Theoretical Implications Our research has revealed that people schedule the tasks of the PRP procedure through a combination of various mechanisms. Symbolic response codes for tasks 1 and 2 may be selected concurrently under flexible strategic control, whereby physical movements are produced in proper serial order. Contrary to traditional response-selection bottleneck hypotheses (cf. Pashler, chap. 12, Jolicoeur, Dell'Acqua, and Crebolder, chap. 13, and Ivry and Hazeltine, chap. 17, this volume), we have found no evidence that skilled dual-task performance is constrained by immutable "hardware" decision or response-selection bottlenecks.

Elementary Continuous Tasks

The preceding conclusions based on the EPIC architecture have been strengthened by formulating computational models of executive mental processes for elementary continuous tasks (Kieras and Meyer 1997). Here the focus is on visuomanual tracking and choice-reaction tasks that must be performed without predictable pauses along the way. By fitting quantitative results obtained under such conditions, we further demonstrate the existence and generality of strategic cognitive control that judiciously overlaps stages of processing in human multitask performance.

Martin-Emerson and Wickens 1992 For this demonstration, our research has dealt especially with Martin-Emerson and Wickens 1992, in which subjects viewed upper and lower windows on a display screen. In the upper window were a circular target and crosshairs cursor. During 1-minute test intervals, the cursor's location was perturbed haphazardly by an accelerative forcing function. The subjects performed a compensatory tracking task, moving a right-hand joystick to keep the cursor on target. The tracking task was either hard or easy, requiring more or less frequent joystick movements. Meanwhile, in the lower window, horizontal arrows appeared intermittently. Depending on whether an arrow pointed right or left, subjects pressed a left-hand index or middle finger key. The centers of the task windows were separated by a visual angle that varied systematically across test intervals. As this angle increased, eye movements that traveled greater distances were required for the stimuli to be identified correctly. Both the tracking and arrow-discrimination tasks were supposed to receive high priority.

Empirical Results As shown in figure 30.4, mean RTs for the arrow discriminations increased reliably with the visual angle between display windows but were relatively unaffected by tracking difficulty. In contrast,




Figure 30.4 Results from Martin-Emerson and Wickens (1992). *Top*. Observed mean reaction times (dark points on solid lines) and simulated mean reaction times (light points on dashed lines) produced by the EPIC model in figure 30.5 for the arrow-discrimination task when it was performed concurrently with either an easy or hard visuomanual tracking task. *Bottom*. Observed and simulated root mean square (RMS) errors for the visuomanual tracking task when it was easy or hard.

root mean square (RMS) tracking errors were reliably greater for hard tracking, but the visual angle affected them relatively little. This occurred even though the tracking errors were measured during 2 sec intervals that started at the onsets of the stimuli for the arrow-discrimination task.

EPIC Models To account for these results, we first formulated an EPIC model that uses inefficient “lockout” scheduling, which let us test predictions based on the traditional response-selection bottleneck hypothesis (cf. Pashler, chap. 12, this volume). According to this model, whenever an arrow occurs, tracking is suspended as soon as possible, performance of the arrow-discrimination task proceeds until completion, and then tracking is resumed. Given realistic delays in EPIC’s motor processors, such lockout scheduling yielded excessively large RMS tracking errors.

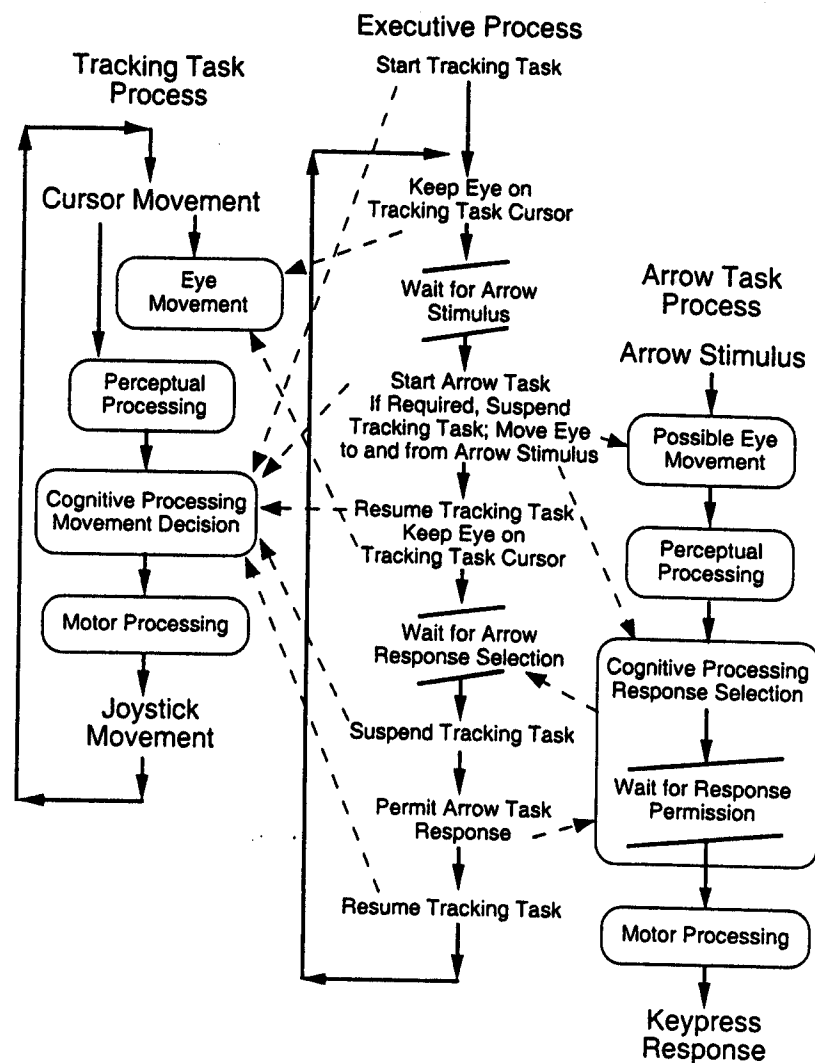


Figure 30.5 Flowchart of an EPIC model with a customized executive process that implements overlapped task scheduling for Martin-Emerson and Wickens 1992. Dashed diagonal arrows from the executive process to the concurrent tracking and arrow-discrimination task processes represent context-dependent supervisory control imposed under these conditions. Mean reaction times and root mean square (RMS) tracking errors produced by this model appear in figure 30.4.

These discrepancies led us to reject this first model and to formulate a second model, with more efficient overlapped task scheduling.

Figure 30.5 shows the task and executive processes of our second model. Here the executive process initially starts the tracking task and enables decisions about joystick movements to be made on the basis of perceived cursor movements. Next, the executive process enters an iterative loop in which it sends commands to the ocular motor processor for keeping the eyes on the tracking task cursor while waiting for an arrow to occur. During this wait, cursor movements may trigger the production rules of the tracking task, which send commands to the manual motor processor for producing joystick movements that keep the cursor on target. When the onset of an arrow is detected, the executive process starts

the arrow-discrimination task and enables its production rules to select a keypress response in deferred mode. For an arrow in foveal or parafoveal vision, perceptual identification proceeds without further ado, and a keypress response is selected while tracking continues until the response's identity becomes available in working memory. Otherwise, for an arrow in peripheral vision, the executive process takes several additional steps: it suspends tracking, moves the eyes to look at the arrow so that its identification can proceed, returns the eyes to look at the cursor, and resumes tracking until a deferred-mode keypress response to the arrow has been selected. As far as possible, this lets tracking continue simultaneously with perceptual identification and response selection for the arrow. Furthermore, as soon as possible after a keypress response has been selected, the executive process also suspends tracking and permits the keypress's identity to be sent to the manual motor processor. Then the keypress response is produced, the arrow-discrimination task is terminated, and tracking is resumed again. Thus, this overlapped task-scheduling model is similar to our previous model for the PRP procedure (cf. figure 30.3).

Simulated Results Figure 30.4 shows simulated results from the present model, whose mean RTs and RMS tracking errors closely approximate those produced by actual participants. Unlike lockout scheduling, overlapped scheduling does not yield excessively large tracking errors.

Theoretical Implications The present model's success supports our claims about how executive mental processes may temporally overlap visual, response selection, ocular motor, and manual motor operations in multitask performance. Apparently, the types of control mechanisms and scheduling strategies we have proposed for discrete concurrent (e.g., PRP) tasks also contribute to efficient performance of elementary continuous tasks. These mechanisms seem to be used regardless of whether the tasks involve the same (e.g., visuomanual) or different (e.g., auditory-vocal and visuomanual) perceptual-motor modalities.

Compound Continuous Tasks

Our characterization of executive mental processes applies not only to elementary but also to compound continuous tasks that entail several distinct subtasks. For example, Ballas, Heitmeyer, and Perez 1992 studied concurrent visuomanual tracking and tactical decision making during simulated military aircraft operations. In tracking, subjects plied a joystick to superimpose a cursor over an evasive target plane. In tactical decision making, subjects pressed finger keys to designate the hostility of numbered icons that depicted jet fighters, bombers, and missile sites. Because there were various types of icon and designation criteria, this decision making constituted a compound task.

To account for performance under these quasi-realistic conditions, we have found that an EPIC model with a three-level hierarchy of executive and task processes fits empirical data well (Kieras and Meyer 1997, forthcoming; Meyer and Kieras, 1997b, 1999). As part of this model, a supra-ordinate dual-task executive process provides overall supervision for a tracking process, a display monitoring process, and a tactical executive process that coordinates three subprocesses—stimulus icon selection, hostility response selection, and track-number response selection—in tactical decision making. Through this hierarchical control, the relative priority of tactical decision making and the temporal overlap of its subprocesses are varied dynamically, contingent on the numerosity of potentially hostile icons in the display. The model, with its adaptive scheduling mechanisms, accounts well for observed sequences of tactical-decision RTs and RMS tracking errors.

Interim Status Quo

From the preceding illustrations, it should be clear that EPIC yields significant theoretical insights about executive mental processes. However, our progress thus far has been limited in some major respects.

Limitations of EPIC Models One limitation is that the executive processes of our models have been customized for particular task combinations. Although these processes may be somewhat similar across contexts, their formulation has incorporated considerable task-specific knowledge. For example, in modeling Martin-Emerson and Wickens 1992, we had the executive directly control eye movements from the stimulus arrows to the tracking cursor (figure 30.5). This enhances tracking performance, consistent with available data, but makes the executive context dependent and nonmodular. To be strengthened further, EPIC needs general executive processes that are context independent.

Previous theorists have also stressed the importance of general executive processes, as in proposals about the “central executive” (Baddeley 1986) and “supervisory attentional system” (Norman and Shallice 1986). Yet they have not provided explicit computational algorithms that achieve the required generality. Thus we must look elsewhere for ways to fulfill this need.

Accompanying EPIC’s lack of general executive processes is a second, related deficiency. Competition among processes for access to limited “hardware” resources may cause miscommunication or deadlock, in which wrong information is transmitted or processes become perpetually stalled (Stallings 1998). EPIC does not yet solve these concurrency problems adequately. Without adequate solutions, veridical modeling of complex adaptive multitask performance will be impossible.

A third limitation is that EPIC does not yet deal with procedural learning in multitask performance. How do people learn to schedule and coordinate concurrent tasks efficiently? How are their multitasking skills transferred across situations? Deeper answers are needed for modeling skill acquisition and developing effective instructional techniques in practical applications (Gopher 1993).

Potential Contributions of Operating System Fundamentals Fortunately, contemporary computer operating systems may stimulate further theorizing. Fundamental principles that underlie their operation provide basic ways for implementing context-independent control and for solving problems of task concurrency (Stallings 1998). By considering these fundamentals, we may augment EPIC with needed general executive processes, concurrency solutions, and multitasking skill acquisition.

30.3 CONTEMPORARY OPERATING SYSTEMS AND COMPUTER TECHNOLOGY

Contemporary operating systems supervise information processing for task programs that are executed virtually or actually in parallel. However, limited capacities of computer hardware impose constraints on an OS trying to maximize process throughput. Consequently, we next consider aspects of both hardware design and OS functions that bear on these matters.

Hardware Design

Starting with early computers like ENIAC, hardware design has become increasingly sophisticated (Tucker 1997). As a result, modern computers typically have at least one central processing unit (CPU), at least one memory unit, and various input/output (I/O) peripherals. The CPU executes sequences of instructions for system and task programs. The memory unit stores programs and data, letting them be manipulated in similar ways. Thus generic information-processing capabilities are implemented by the hardware, whereas overall system control and task procedures are provided by the software.

Uniprocessor Architecture Many operating systems and task programs are used on computers with one CPU. Although this uniprocessor architecture executes instructions sequentially in some respects, its components enable extensive parallelism. For example, separate streams of data may be transmitted simultaneously to or from different I/O peripherals, and the CPU may perform multiple suboperations in parallel. Exploiting such capabilities, an OS can sustain concurrent threads of processing at least somewhat as if each program had its own CPU.

Multiprocessor Architectures Moreover, some operating systems and task programs have been implemented with multiple CPUs. These multiprocessor architectures enable true parallel processing and provide enormous, relatively inexpensive, computational power. Particularly relevant for us is the *shared-memory symmetric multiprocessor* (SMP), in which multiple CPUs function as equivalent "peers" that share one memory unit and I/O peripherals. This corresponds at least approximately to EPIC's organization. Although EPIC has one cognitive processor, it tests conditions and executes actions of multiple production rules in parallel. When programmed with two or more rule sets, the cognitive processor emulates a collection of peer CPUs; as in a SMP, these rule sets share WM and I/O peripherals.

Thus contemporary OS fundamentals should be applicable to EPIC. Indeed, computer scientists have discovered that OS fundamentals are extremely general, applying across many uniprocessor and multiprocessor architectures. This suggests that what OSs and EPIC teach us will likely hold as well for the human mind and brain, which also implement forms of multiprocessor parallelism. To appreciate OS fundamentals, more background about them is in order (see Stallings 1998; Tucker 1997).

Operating System History

Like computer hardware, operating systems have become increasingly sophisticated. For early computers (circa 1950), people loaded and started programs manually. Subsequently (circa 1960), primitive OS resident monitors were developed to automate these processes. Following this development, OS capabilities were gradually extended to enable overlapping CPU and I/O operations so that the CPU would not have to wait idly on slow mechanical devices. These advances led to multitasking, an overarching OS function (circa 1970).

In multitasking, an OS interleaves or overlaps execution of task programs requiring certain limited hardware resources. When an execution process has taken a set time or must wait for pending I/O, it is suspended, and the CPU is allocated to another process. After completion of I/O or other prerequisites, the suspended process is resumed. Consequently, multiple processes may advance efficiently without individual users' intervention. Software for multitasking on uniprocessors has been gracefully adapted for multitasking on multiprocessors.

Operating System Objectives

Systems programmers developed operating systems to keep CPU and memory hardware as busy as possible, increasing process throughput. OSs have also made it simpler and faster to formulate *noncooperating task*

programs, which are executed asynchronously and compete for hardware resources. Given OS services, such a program can be formulated as if it were the only one executed and no intricate control of I/O were required. Furthermore, OSs have facilitated the formulation of *cooperating task programs*, which are executed synchronously and share their products interactively.

However, OSs are neither logically necessary nor maximally efficient in every respect. Nonhierarchical "flat" programs can be formulated to perform multiple tasks concurrently on "bare" computer hardware without OS support. Through this formulation, the computational overhead of hierarchical software can be eliminated, and even faster performance achieved. Nevertheless, such improvement has serious costs. Because it requires dealing directly with many levels of control, the time and effort needed to formulate flat programs can be exorbitant. Also, flat programs do not readily generalize beyond their original applications. In contrast, OSs provide a better compromise between speed of execution, on the one hand, and ease and generality of software development, on the other.

Operating System Functions

This compromise is enabled by operating system functions that solve a basic problem: detailed sequences of execution for independent task programs cannot be predicted accurately. An OS must ensure that execution proceeds correctly and rapidly despite unpredictable interruptions and resumptions. The solution entails judicious task scheduling, resource allocation, process coordination, and conflict resolution.

Task Scheduling In task scheduling, an operating system must make and implement decisions about when programs will be executed. Doing so requires prioritizing, preparing, initiating, suspending, preserving, resuming, and terminating each execution process at apt moments. OSs use various scheduling algorithms for this. Among them are "first come, first serve," "round robin," "shortest remaining time," "shortest process next," "highest response ratio," and "least-time-consumed scheduling," each of which may produce relatively high or low performance, depending on nuances of the prevailing context. Task scheduling by an OS must therefore be "tuned" adaptively to maximize overall throughput.

Resource Allocation An operating system must also allocate hardware resources judiciously to individual processes, depending on resource availability and process needs. For example, during execution, a process may request resources. If these are available, the OS may comply by allocating them immediately. Alternatively, if they have been committed to other processes already, then the OS may deny the current request temporarily, and perhaps suspend the requesting process until its needs can

be satisfied. Exactly when processes request and release their resources, and how the OS handles them, contribute significantly to attained performance.

Process Coordination Among the processes being executed, some may need to share intermediate products of their computations. For this sharing to succeed, these cooperating processes must be coordinated, because interprocess communication involves writing to and reading from the same memory locations in proper serial order.

To facilitate interprocess communication, an operating system performs several coordinative functions, including mutual exclusion, process synchronization, and message passing. Relying on these functions, a receiving process may request that the OS suspend it until an expected message arrives from another sending process. When the sending process is ready to transmit this message, it may request that the message be passed to the receiving process. The OS may then pass the message and resume the receiving process.

Conflict Resolution Because concurrent processes impose high loads on hardware resources and may be noncooperative, serious conflicts can arise. An operating system has to avoid these conflicts as best it can, and resolve them gracefully when need be. This function is crucial for dealing with *deadlocks*, which entail closed chains of processes such that each process currently has exclusive ownership of some resource needed by the next process in the chain. Adaptive conflict resolution also helps deal with other undesirable situations such as *starvation*, in which some low-priority process is perpetually preempted by higher-priority processes.

30.4 COGNITIVE CONTROL AND OPERATING SYSTEM FUNDAMENTALS

Contemporary operating systems embody precise and comprehensive instantiations of executive processes. Such instantiations are scarce in current psychological theories. Thus, to promote further progress, we next discuss some stimulating theoretical concepts, multitasking models, and explanatory hypotheses inspired by these considerations.

Theoretical Concepts

The concepts that interest us here involve distinctions between various types of executive and task processes.³

Customized Executives One major distinction concerns customized versus general executives. By "customized executive" (CE), we mean a modular set of supraordinate mental processes that manage multitask

performance based on unique context-dependent knowledge about the particular tasks and their temporal interrelations. A CE works for only one task combination and cannot be transferred readily across different situations.

Thus far, EPIC models have all used CEs. An instructive case is our model of performance in Martin-Emerson and Wickens 1992. The executive process of this model preallocates resources (i.e., ocular and manual motor processors) to tracking and arrow discrimination without these task processes requesting them explicitly (figure 30.5). The preallocation is possible here because the executive already "knows" the task processes' needs and satisfies them in proper sequence. Such use of context-dependent knowledge may be common after extensive practice under conditions in which high performance is desired.

Our theorizing need not be confined, however, to models with CEs. New EPIC models may be formulated on the basis of general executives that function at least partly like contemporary OSs. From testing them empirically, we learn more about the extent to which OS fundamentals characterize how human multitask performance is controlled.

General Executives A general executive (GE) is a modular set of supra-ordinate mental processes that manage multitask performance without using unique context-dependent knowledge about the tasks and their temporal interrelations. Given such generality, cognitive control can be achieved for different task combinations through standard functions like those of contemporary OSs. Implementing these functions in EPIC is straightforward because it resembles a shared-memory symmetric multiprocessor.

Nevertheless, determining whether a GE should be added to EPIC requires answering a fundamental question about cognitive control: Do people have GEs and use them for multitask performance? We might expect an affirmative answer, given the potential ease of preparing and efficiency of executing task programs based on GE functions. Yet the only way to be sure about this is to formulate and test EPIC models that rely on a GE. We take this course after introducing more distinctions that will facilitate our pursuits.

Managerial Styles Another relevant distinction concerns managerial styles of general executives. At one extreme, a *conservative* GE can have a strict regimented style of scheduling task processes and allocating limited resources to them. Under such regimentation, task processes may have to request resources before using them; processes may be suspended when their requested resources are unavailable; and processes not prone to make deferent resource requests may be kept from starting (i.e., locked out) until others have finished. Alternatively, a *liberal* GE can have a tolerant, laissez-faire managerial style, under which task pro-

cesses may be allowed to proceed at least partially unabated while their requested resources are unavailable, and processes prone to use resources without requesting them may also be accommodated insofar as possible. In principle, a GE's managerial style is adaptable to particular situations. Such adaptability, contingent on the "manners" of task processes, will determine the attained level of multitask performance.

Process Manners and Etiquette Task processes can have various manners of interaction with a general executive. Proper etiquette for a task process entails requesting resources (e.g., motor mechanisms) immediately before they will be used, waiting for the GE's permission to use them, and then releasing the resources immediately after their use is complete. A "*polite process*" conforms to all of these rules. This establishes favorable circumstances for a *laissez-faire* managerial style through which relatively high multitask performance is attainable.

Theoretically, however, some task processes may be impolite. For example, a *presumptuous* process might use crucial resources without requesting them. An *impatient* process might request resources but not wait for permission to use them. A *greedy* process might request resources too early and release them too late. Such inconsiderate conduct will force a GE to be more conservative, curtailing the processes' temporal overlap and impeding their progress.

Cost-Benefit Assessment To assess the costs and benefits of alternative general executive managerial styles, various factors are relevant. One is *interaction overhead*, which includes scheduling, allocation, and abdication costs for supervising task processes. *Scheduling costs* are amounts of time consumed by adding and deleting goals in working memory to start, suspend, resume, and terminate processes selectively. *Allocation costs* are amounts of time consumed by making and fulfilling resource requests. *Abdication costs* are amounts of time consumed by releasing resources. Ideally, these costs should be paid in ways that decrease *resource possession times*, the amounts of time during which a task process possesses crucial resources. Also, as best they can, the payments should increase *process overlap intervals*, the intervals during which multiple processes are advancing simultaneously.

Taking these factors into account, impolite task processes may escape some interaction overhead, but they increase resource-possession times and force the GE to eliminate process-overlap intervals. In contrast, a liberal GE and polite task processes make an attractive compromise. Their process-overlap intervals and resource-possession times may be relatively long and short, respectively, thereby more than compensating for the GE's moderate interaction overhead.

Nevertheless, there are other ways to perform better on all scores. CEs (customized executives) tuned for particular task combinations can achieve even lower interaction overhead, shorter resource-possession

times, and longer process-overlap intervals. As we shall see, this leads to interesting hypotheses about multitasking skill acquisition.

New Multitasking Models

To illustrate how these theoretical concepts help clarify the nature of cognitive control, we have implemented them in two new EPIC models for Martin-Emerson and Wickens 1992. Model 1 has a conservative general executive that supervises two impolite task processes. Model 2 has a more liberal GE that supervises two polite task processes. By comparing these models to our previous one that has a customized executive (figure 30.5), we examine the effects of managerial style and process manners on multitask performance.

Model 1: Conservative General Executive with Impolite Processes

In model 1, tracking and arrow discrimination are assumed to be impolite processes. They do not request or release resources for producing eye and hand movements. Instead, each process tries to move the eyes and hands without regard for what is happening elsewhere in the system, creating prospects for "jams" in EPIC's motor processors.

To cope with this impoliteness, model 1 has a general executive that uses a first come, first serve (FCFS) algorithm for scheduling the tracking and arrow-discrimination task processes in strict lockout mode. Under it, these processes may be started optionally when their stimuli (arrows and suprathreshold tracking errors) are detected. However, the GE lets only one process proceed at a time. If stimuli for both processes occur simultaneously, then the lower-priority one (tracking) is postponed until the higher priority one (arrow discrimination) has responded to its current stimulus.

This protocol resembles the one of Norman and Shallice's "supervisory attentional system" (SAS; 1986), in which action schemata are activated by "trigger" stimuli and contend for limited response mechanisms. Precluding conflicts from this "contention scheduling," the SAS transmits top-down activation to the highest-priority schema, favoring it over lower-priority schemata. In our model 1, the lockout scheduling is like the selective prioritization imposed by the SAS. Thus we may test both model 1 and the SAS by comparing the performance of model 1 to real data.

Table 30.1 shows results of this comparison. When a small visual angle (<5 degrees) separates the displays of the tracking and arrow-discrimination tasks, simulated RTs from model 1 are considerably less than observed ones (mean difference = 103 msec), but at larger angles (>10 degrees), simulated RTs are considerably greater than observed ones (mean difference = 97 msec). Furthermore, the simulated RMS tracking errors of model 1 are much larger than the observed ones; when tracking is difficult, they differ by more than a factor of 2 at large visual

Table 30.1 Reaction Times and Tracking Errors Produced by Four Alternative EPIC Models to Account for Observed Data from Martin-Emerson and Wickens 1992

Tracking Task	Visual angle	Mean arrow-discrimination reaction times (msec)				Root mean square tracking error (pixels)					
		Model 1	Model 2	Model 3	Model 4	Data	Model 1	Model 2	Model 3	Model 4	Data
Easy	small	608	690	686	612	697	19.4	17.2	13.4	13.5	13.3
	medium	781	836	782	696	787	22.5	18.0	13.6	13.6	13.3
	large	1,070	1,106	978	886	997	28.6	20.1	14.3	14.5	15.0
Hard	small	632	688	686	608	749	28.7	25.2	19.9	19.7	18.6
	medium	802	842	787	700	786	34.8	27.6	20.6	20.9	19.6
	large	1,097	1,112	996	887	976	43.2	30.0	22.7	22.9	21.4

Note: Reaction times and tracking errors were measured as a function of the tracking-task difficulty (easy versus hard) and the magnitude of the visual angle ("small" < 5 degrees; "medium" 6 to 10 degrees; "large" > 10 degrees) between the display windows for the tracking and arrow discrimination tasks (cf. figure 30.4).

angles. Model 1 performed very poorly even though under it, tracking and arrow discrimination progress as fast as reasonably possible while they are under way, and there are no resource allocation or abdication costs of supervising them. Instead, the poor performance of model 1 stems from an absence of process overlap caused by its GE having to cope conservatively with the impoliteness of the task processes in their use of motor resources.

These results disconfirm both model 1 and the SAS with respect to Martin-Emerson and Wickens 1992. Contrary to these models, under at least some conditions, cognitive control for multitask performance is more efficient than a conservative GE and impolite task processes allow. We investigate the sources of this efficiency more fully by considering a second new model.

Model 2: Liberal General Executive with Polite Task Processes In model 2, tracking and arrow discrimination are assumed to be polite processes. Each task process requests motor resources immediately before it would use them, does not use them until the general executive grants permission, and releases them immediately after they have been used. Given this politeness, the GE lets these processes advance simultaneously insofar as possible, even after one of them has requested resources that the other is currently using. Such liberalism is feasible because the task processes make eye and hand movements in a considerate manner that avoids motor-processor "jams," thereby enabling more process overlap than model 1 allows.

Another virtue of model 2 is its straightforward flow of control. Compared to our original model for Martin-Emerson and Wickens 1992 (figure 30.5), model 2 has a relatively simple flowchart (figure 30.6). Consequently, during multitasking practice, the skill embodied in model 2 should be fairly easy to acquire.

Consistent with these points, table 30.1 shows that model 2 produces somewhat better performance than model 1 does. Especially when tracking is difficult, simulated RMS errors from model 2 are markedly smaller than those from model 1. Nevertheless, there remain significant discrepancies between the performance of model 2 and the observed data. Both the simulated tracking errors and simulated RTs are still excessively large, suggesting that actual participants achieved even more process overlap than model 2 allows.

Why and how might this be? An answer may come from reconsidering our original model for Martin-Emerson and Wickens 1992 which we now call "model 3."

Model 3: Customized Executive with Resource Preallocation and Enhanced Task Processes As depicted before (figure 30.5), model 3 uses a customized executive that exploits context-dependent knowledge

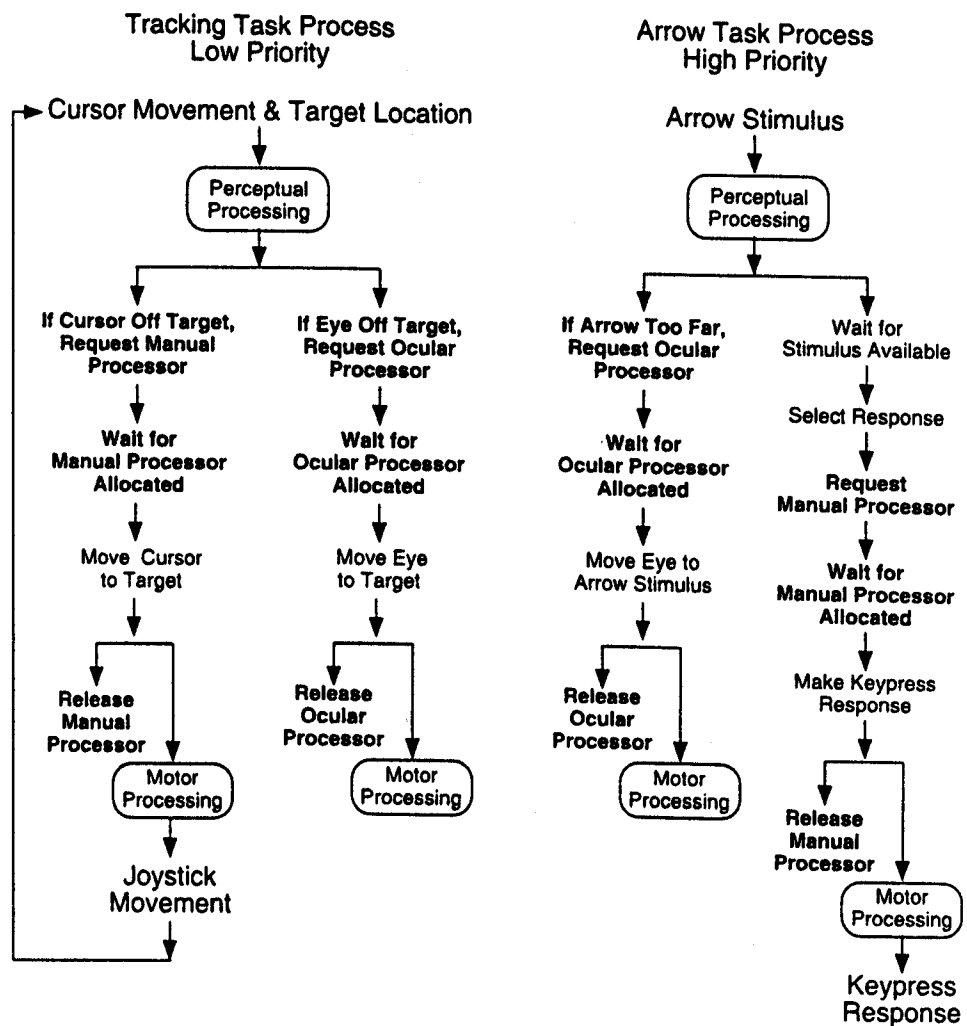


Figure 30.6 Flowchart of an EPIC model that performs the tracking and arrow-discrimination tasks of Martin-Emerson and Wickens 1992 with polite task processes and a general executive whose managerial style is liberal in task scheduling and resource allocation (cf. figure 30.5).

about the tasks and their temporal relationships. Based on this knowledge, the customized executive preallocates resources (i.e., ocular and manual motor processors) to tracking and arrow discrimination without being requested to do so. This enables the task processes to advance even more quickly than under model 2. Under model 3, the task processes also prepare eye movements beforehand. Together, these enhancements further facilitate performance so that the simulated RTs and tracking errors of model 3 are considerably less than those of model 2, closely approximating observed data (table 30.1).

The good fit of model 3 suggests that participants in Martin-Emerson and Wickens 1992 achieved excellent multitask performance through especially efficient cognitive control. Without this efficiency, limitations of perceptual-motor mechanisms would have precluded such performance. The customized executive of model 3 overcomes these limitations more so than a general executive can. Nevertheless, during the course of prac-

tice, participants may have relied on a GE to acquire their high level of multitasking skill. How this could happen is considered next.

Hypotheses about Skill Acquisition

Taken together, our results from models 1, 2, and 3 lead to hypotheses that explain various major aspects of multitask performance and skill acquisition.

Multitasking Skill-Acquisition Stages We hypothesize that multitasking skill acquisition progresses through five stages: preprocedural interpretative multitasking (stage 0); general hierarchical competitive multitasking (stage 1); general hierarchical cooperative multitasking (stage 2); customized hierarchical multitasking (stage 3); and customized heterarchical multitasking (stage 4). Each of these stages can be characterized with respect to its degree of efficiency, types of interaction between executive and task processes, and exploitation of context-dependent procedural knowledge.

Preprocedural interpretive multitasking is necessitated by a fundamental dependence between procedural and declarative task knowledge. We call this "stage 0" because it occurs at the start of practice before sets of production rules for the particular tasks have been created. During stage 0, people must use a generic interpretive process to execute propositional instructions about how the tasks should be performed. Here performance is presumably slow and error prone, placing heavy loads on working memory as people "think" their way verbally through each task. Nevertheless, it is from this explicit directed intentional activity that more efficient procedural knowledge for subsequent task performance emerges (Anderson 1983; Bovair and Kieras 1991; Kieras and Bovair 1986).

Once such knowledge becomes available, general hierarchical competitive multitasking may ensue. We call this "stage 1" because it is the first stage during which a general executive supervises task processes that are executed through individualized sets of production rules. Also during stage 1, task scheduling and coordination are managed as in our model 1 for Martin-Emerson and Wickens 1992. Here performance presumably entails a conservative GE with strict lockout scheduling of impolite task processes whose manners in using perceptual-motor resources are impulsive, presumptuous, and greedy. This impoliteness may be attributed to a need for more practice in order to acquire rules that conform with proper task etiquette.

As practice continues, general hierarchical cooperative multitasking may come next. During what we call "stage 2," task scheduling and coordination would be managed as in our model 2. Here performance presumably entails a liberal GE with temporal overlapping of task processes that request, use, and release system resources politely. This politeness

enables the GE to be more permissive in letting these processes advance rapidly toward completion.

Customized hierarchical multitasking would involve an even higher skill level. During what we call "stage 3," task scheduling and coordination may be managed as in our model 3. Here unique context-dependent knowledge about the particular tasks and their temporal interrelations presumably is exploited to preallocate system resources without time-consuming requests for them, thereby further increasing temporal overlap among task processes. Also, as in model 3, these processes may be enhanced to prepare their motor responses anticipatorily.

Culminating this evolution is customized heterarchical multitasking. During what we call "stage 4," performance presumably is controlled without supraordinate executive processes. Instead, the task processes interact directly with each other, self-governing their resource usage as efficiently as possible. This interaction optimizes overall system throughput, completely eliminating scheduling, allocation, and abdication time costs that contribute to the transaction overhead of hierarchical cognitive control.

Table 30.1 shows some benefits of such optimization. Here we have included results from a fourth model ("model 4") that uses the customized heterarchical multitasking of stage 4 to simulate performance in Martin-Emerson and Wickens 1992. The RMS tracking errors of model 4 closely approximate the data, and its mean RTs are even shorter than observed ones. Although the subjects in Martin-Emerson and Wickens 1992 were highly skilled, they apparently had not yet reached this ultimate asymptotic performance level.

Executive Learning Mechanisms Operations within and transitions between the preceding five stages of multitasking skill acquisition may be mediated by various executive learning mechanisms (cf. Anderson 1983; Bovair and Kieras 1991; Chong and Laird 1997; Kieras and Bovair 1986). These mechanisms may entail several components: a *task interpreter*, which executes propositional instructions for performing single and multiple tasks during stage 0; a *task compiler*, which creates rudimentary sets of production rules for the initially impolite task processes of stage 1; a *task socializer*, which makes these processes more polite in stage 2; an *executive modulator*, which tailors the general executive's managerial style to be either conservative or liberal, depending how polite the task processes are; an *executive customizer*, which creates customized executives to enable even more efficient control in stage 3; and an *executive integrator*, which "flattens" the CEs, converting their flow of control from a hierarchical to heterarchical organization in stage 4.

We hypothesize that such mechanisms are sensitive to the evolving characteristics of performance. For example, during stage 1, simultaneous attempts by multiple impolite task processes to produce movements

in the same response modality could generate motor-processor "jams." These jams might be detected by the executive modulator, leading it to have the GE be conservative during the period of time when the task socializer works toward making the task processes more polite. The task socializer and executive modulator also could operate partly on the basis of noticing that the task processes do not request and release resources properly. Later, after the task socializer achieves its objectives, the executive modulator perhaps would adjust the GE to be more liberal because motor-processor jamming has ceased. Accompanying the latter adjustment, the executive customizer might start creating a CE that later triggers hierarchical-to-heterarchical flattening by the executive integrator. Of course, future research will be needed to understand and model the details of such hypothetical learning mechanisms.

Multitasking Skill-Acquisition Phenomena By doing so, we may eventually explain and predict many empirical phenomena of multitasking skill acquisition. For example, Gopher (1993) has found that multitasking performance is better after variable-priority rather than fixed-priority training. In his fixed-priority training condition, one group of participants gave equal priorities to visuomanual tracking and choice-reaction tasks throughout a series of practice sessions. In his variable-priority training condition, a second group of participants also gave the two tasks equal priorities on some occasions, but devoted higher priority to either tracking or choice reactions on other occasions. After variable-priority training, the second group performed better than the first group even when the two tasks received equal priorities. Similar results have been reported by Meyer et al. (1995). The benefits of variable-priority training could stem from the task socializer and executive modulator receiving a wider range of feedback, which guides them more quickly through successive stages of skill acquisition.

Our hypotheses likewise account for results obtained with some other laboratory paradigms. For example, RTs from the PRP procedure sometimes manifest a response-selection bottleneck (Pashler 1994, chap. 12, this volume). This seems to occur especially when participants receive relatively little practice at coordinating their primary and secondary tasks (Schumacher et al., 1999). A possible reason is that participants lack sufficient opportunity to socialize initially impolite task processes, so their GE has to deal with this impoliteness through strict lockout scheduling (cf. Meyer and Kieras 1997a,b).

30.5 CONCLUSIONS

Assimilating the fundamentals of contemporary computer operating systems into theories of cognitive control will make it possible to character-

ize a wider range of control functions more precisely, and to test more definitively for the existence of general as well as customized executive processes. These advances also will lead to more detailed and veridical analyses of multitasking skill acquisition. Computational modeling based on the EPIC architecture provides one vehicle whereby this progress can occur.

For the present prospects to be fully realized, future research must use a wide variety of empirical procedures to investigate multitask performance. This investigation should extend beyond basic laboratory paradigms like the task-switching and PRP procedures, which are helpful for isolating particular elementary control functions, but come nowhere near to engaging the whole host of executive mental processes that people presumably have. Rather, to explore these processes more completely, overlapping-task procedures with complex realistic tasks and unpredictable stimulus-response event sequences will be needed (e.g., Ballas, Heitmeyer, and Perez 1992).

Another major path for future research will involve identifying systematic relationships between underlying brain mechanisms and the executive mental processes revealed by taking operating system fundamentals into account. Because OS fundamentals apply quite generally to shared-memory symmetric multiprocessors, of which the brain is perhaps one type, it seems reasonable that the brain implements these fundamentals as well. If so, then insights from EPIC computational modeling, applied to results from studies of brain imaging and focal lesion analysis, could eventually yield fundamental solutions to the mind-body problem of cognitive control.

NOTES

Funding for this research was provided by U.S. Office of Naval Research grant N00014-92-J-1173 to the University of Michigan. We thank David Fencsik, Darren Gergle, Jennifer Glass, Leon Gmeindl, Cerita Jones, Shane Mueller, Eric Schumacher, Mollie Schweppe, and Travis Seymour of the Brain, Cognition, and Action Laboratory at the University of Michigan for their helpful assistance. Comments by Leon Gmeindl, Stephen Monsell, Travis Seymour, and two anonymous reviewers on drafts of this chapter are greatly appreciated.

1. First-task responses yielded a similar pattern of results. Although mean reaction times decreased across sessions, their pattern did not change qualitatively with practice. No significant asymmetries occurred in switching time costs. Error rates were moderately low (<10%) on average and correlated positively with mean RTs, suggesting no systematic speed-accuracy trade-offs.
2. Our model also accounts well for mean first-task RTs and the factor effects on them.
3. Insofar as we know, the distinctions described here have not been made explicitly in operating system textbooks. They are introduced here to address issues about human cognitive control, which extend well beyond those associated with computer applications where experienced task programmers adhere consistently to common *a priori* conventions.

REFERENCES

- Allport, A., Styles, E., and Hsieh, S. (1994). Shifting intentional set: Exploring the dynamic control of tasks. In C. Umiltà and M. Moscovitch (Eds.), *Attention and performance XV*, pp. 421–452. Cambridge, MA: MIT Press.
- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Baddeley, A. D. (1986). *Working memory*. Oxford: Oxford University Press.
- Ballas, J. A., Heitmeyer, C. L., and Perez, M. A. (1992). *Direct manipulation and intermittent automation in advanced cockpits*. Technical Report NRL/FR/5534—92-9375. Washington, DC: Naval Research Laboratory.
- Bovair, S., and Kieras, D. E. (1991). Toward a model of acquiring procedures from text. In R. Barr, M. L. Kamil, P. B. Mosenthal, and P. D. Pearson (Eds.), *Handbook of reading research*, vol. 2, pp. 206–229. New York: Longman.
- Card, S. K., Moran, T. P., and Newell, A. (1983). *The psychology of human-computer interaction*. Hillsdale, NJ: Erlbaum.
- Chong, R. S., and Laird, J. E. (1997). Towards learning dual-task executive process knowledge using EPIC-Soar. In *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society*, pp. 107–112. Hillsdale, NJ: Erlbaum.
- Gopher, D. (1993). Attentional control: Acquisition and execution of attentional strategies. In D. E. Meyer and S. Kornblum (Eds.), *Attention and performance XIV: Synergies in experimental psychology, artificial intelligence, and cognitive neuroscience*, pp. 299–322. Cambridge, MA: MIT Press.
- Kieras, D. E., and Bovair, S. (1986). The acquisition of procedures from text: A production-system analysis of transfer of training. *Journal of Memory and Language*, 25, 507–524.
- Kieras, D. E., and Meyer, D. E. (1997). An overview of the EPIC architecture for cognition and performance with application to human-computer interaction. *Human-Computer Interaction*, 12, 391–438.
- Kieras, D. E., and Meyer, D. E. (Forthcoming). The role of cognitive task analysis in the application of predictive models of human performance. In J. M. C. Schraagen, S. E. Chipman, and V. L. Shalin (Eds.), *Cognitive task analysis*. Mahwah, NJ: Erlbaum.
- Kristofferson, A. B. (1967). Attention and psychophysical time. In A. F. Sanders (Ed.), *Attention and Performance*, pp. 93–100. Amsterdam: Elsevier.
- Lachman, R., Lachman, J. L., and Butterfield, E. C. (1979). *Cognitive psychology and information processing: An introduction*. Hillsdale, NJ: Erlbaum.
- Lauber, E. J. (1995). Executive control of task-switching operations. Ph.D. diss., University of Michigan, Ann Arbor.
- Martin-Emerson, R., and Wickens, C. D. (1992). The vertical visual field and implications for the head-up display. In *Proceedings of the Thirty-sixth Annual Symposium of the Human Factors Society*. Santa Monica, CA: Human Factors Society.
- Meiran, N. (1996). Reconfiguration of processing mode prior to task performance. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 1423–1442.
- Meyer, D. E., and Kieras, D. E. (1997a). EPIC—A computational theory of executive cognitive processes and multiple-task performance: Part 1. Basic mechanisms. *Psychological Review*, 104, 3–65.

- Meyer, D. E., and Kieras, D. E. (1997b). EPIC—A computational theory of executive cognitive processes and multiple-task performance: Part 2. Accounts of psychological refractory-period phenomena. *Psychological Review*, 104, 749–791.
- Meyer, D. E., and Kieras, D. E. (1999). Précis to a practical unified theory of cognition and action: Some lessons from computational modeling of human multiple-task performance. In D. Gopher and A. Koriath (Eds.), *Attention and Performance XVII*, pp. 15–88. Cambridge, MA: MIT Press.
- Meyer, D. E., Kieras, D. E., Lauber, E., Schumacher, E., Glass, J., Zurbriggen, E., Gmeindl, L., and Apfelblat, D. (1995). Adaptive executive control: Flexible multiple-task performance without pervasive immutable response-selection bottlenecks. *Acta Psychologica*, 90, 163–190.
- Neisser, U. (1967). *Cognitive psychology*. New York: Appleton Century Crofts.
- Newell, A. (1980). Harpy, production systems, and human cognition. In R. A. Cole (Ed.), *Perception and production of fluent speech*, pp. 289–380. Hillsdale, NJ: Erlbaum.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Norman, D. A., and Shallice, T. (1986). Attention to action: Willed and automatic control of behavior. In R. J. Davidson, G. E. Schwartz, and D. Shapiro (Eds.), *Consciousness and self-regulation*, vol. 4, pp. 1–18. New York: Plenum Press.
- Pashler, H. (1994). Dual-task interference in simple tasks: Data and theory. *Psychological Bulletin*, 116, 220–244.
- Rogers, R., and Monsell, S. (1995). Costs of a predictable switch between simple cognitive tasks. *Journal of Experimental Psychology: General*, 124, 207–231.
- Rubinstein, J. S., Meyer, D. E., and Evans, J. (Forthcoming). Executive control of cognitive processes in task switching. *Journal of Experimental Psychology: Human Perception and Performance*, accepted pending revision.
- Schumacher, E. H., Lauber, E. J., Glass, J. M. B., Zurbriggen, E. L., Gmeindl, L., Kieras, D. E., and Meyer, D. E. (1999). Concurrent response-selection processes in dual-task performance: Evidence for adaptive executive control of task scheduling. *Journal of Experimental Psychology: Human Perception and Performance*, 25, 791–814.
- Stallings, W. (1998). *Operating systems: Internals and design principles*. 3d ed. Upper Saddle River, NJ: Prentice Hall.
- Theios, J. (1973). Reaction time measurements in the study of memory processes: Theory and data. In G. H. Bower (Ed.), *The psychology of learning and motivation*, vol. 7, pp. 43–85. New York: Academic Press.
- Tucker, A., Ed. (1997). *The computer science and engineering handbook*. Boca Raton, FL: CRC.