CHAPTER 4

Task Interruptions

By J. Gregory Trafton & Christopher A. Monk

We review both the theoretical and applied research on task interruptions. We provide a brief task analysis of interruptions and resumptions, discuss how and why they are disruptive, and show how multiple theories attempt to explain the interruption and resumption process. We also review a great deal of the empirical work and show how it fits into previous theoretical accounts. Finally, we discuss what factors make interruptions more or less disruptive as well as theory-based recommendations for reducing the disruptiveness of interruptions.

Without question, interruptions are a hot topic in both the popular press and scientific communities. With the rapid rise of communication technologies that keep people accessible at all times, issues of interruptions and multitasking have caught the attention of the scientific community and the popular press alike. Recently, *Time* magazine (Wallis & Steptoe, 2006) and the *New York Times* (Thompson, 2005) both reported stories about interruptions and multitasking and how they affect performance. In 2005, the information technology research firm Basex issued a report on the economic impact of interruptions, which they estimated to be around $588 billion a year (Spira & Feintuch, 2005). Interruptions have received increased attention by researchers from a variety of backgrounds, including cognitive psychologists, human factors researchers, computer scientists, and marketing research specialists, to name a few.

Broadly speaking, there are two major threads of investigation related to interruptions. These two areas correspond roughly to applied research and theoretical investigations. By no means are these two approaches mutually exclusive. Rather, each is necessary and highly valuable to understanding what makes interruptions disruptive; how performance is affected by various interruption characteristics, such as complexity, duration, timing, and frequency; and ultimately what can be done to mitigate the disruptive outcomes.

Our goal with this chapter is to review the contributions of both literatures with greater emphasis on the latter. Because most of the research on interruptions has centered on computer-related interruptions, we will keep our focus there as well. We begin with the characterization of interruptions and their effects and then transition to candidate explanations rooted in cognitive theory and the evidence that supports each. We stress the theoretical portion because we believe the greatest promise for diminishing the disruptive impact of interruptions on performance lies in understanding why interruptions affect people in the ways that are seen in the various applied studies.

ARE INTERRUPTIONS DISRUPTIVE?

A great deal of evidence suggests that in most instances, interruptions are disruptive. Several characteristics of interruptions in particular have been shown to be disruptive. These characteristics include interruption complexity (Cellier & Eyrolle, 1992; Gillie & Broadbent, 1989; Speier, Valacich, & Vessey, 1999; Zijlstra, Roe, Leonora, & Krediet, 1999), similarity of the interrupting task to the primary task (Cellier & Eyrolle, 1992; Czerwinski, Chrisman, & Rudisill, 1991; Edwards & Gronlund, 1998; Oulasvirta & Saariluoma, 2004), how closely the interrupting and primary tasks are related (Cutrell, Czerwinski, & Horvitz, 2001, control over interruption engagement (McFarlane, 2002), and the availability of retrieval cues in the primary task (Cutrell, Czerwinski & Horvitz, 2001; Czerwinski, Cutrell, & Horvitz, 2000b).

Supplementing these studies are those that investigated the nature of disruption beyond interruption characteristics. Most of this work has focused on how interruptions disrupt the amount of time it takes to complete a primary task (Bailey, Konstan, & Carlis, 2000; Cutrell, Czerwinski, & Horvitz, 2001; Czerwinski et al., 2000b; Gillie & Broadbent, 1989; McFarlane, 2002) and on performance immediately after an interruption (Hodgetts & Jones, 2006a, 2006b; Monk, Boehm-Davis, & Trafton, 2004; Trafton, Altmann, Brock, & Mintz, 2003). Note that in some tasks, it may be quite difficult or irrelevant for people to resume where they left off, leading to different types of resumption strategies (e.g., restarting the entire subtask). Additionally, people may experience source confusion (Johnson, Hashtroudi, & Lindsay, 1993), think that they completed the step, and upon resumption actually skip that step. Finally, some researchers have shown that in some workplace situations, the primary task is never actually resumed (Mark, Gonzalez, & Harris, 2005). Many of these types of disruptions are discussed by Dismukes, Berman, and Loukopoulos (2006).

It is important to note that the studies mentioned in this chapter not only focus on the characteristics, impacts, and explanations of interruptions, but they also represent the fact that the issue is being addressed from a broad range of domain perspectives. A great deal of the literature is from the human-computer interaction domain, but aviation, medicine, marketing, military, and driving researchers are also studying the impacts of interruptions. Indeed, it is critical to remember that interruptions are not merely nuisances that people deal with each day in their professional and personal lives; interruptions can result in devastating consequences in contexts such as aviation, emergency rooms and hospitals, driving, and nuclear power control rooms. Pilots experiencing interruptions during preflight checklists have been blamed for multiple aviation crashes (National Transportation Safety Board, 1969, 1988). In addition, recent studies have shown that interruptions may be an important factor in driving (Monk et al., 2004), emergency room care (Chisholm, Collison, Nelson, & Cordell, 2000; Chisholm, Dornfeld, Nelson, & Cordell, 2001), and nursing errors (Tucker & Spear, 2006).

How People Deal With Interruptions

Research has examined how people deal with interruptions in a number of real-world domains, including meteorology, office work, driving, and scientific data analysis. Other researchers have also performed in-depth real-world analyses of the prevalence of interruptions and how they affect performance (Czerwinski, Horvitz, & Wilhite, 2004; Gonzalez & Mark, 2004; Mark et al., 2005). By shadowing, interviewing, and naturalistic observation, Gonzalez and Mark (2004) found that information workers switched tasks extremely frequently (approximately once every 3 min). They suggested that in many of today's settings, work is extremely fragmented.

To buttress this point, recent research on the frequency of instant messaging behavior suggests that conversations last nearly 4½ min on average, with exchanges every 15 s or so (Isaacs, Walendowski, Whittaker, Schiano, & Kamm, 2002). The study also showed that heavy instant message users in the workplace covered multiple topics within each exchange and frequently attended to other tasks during instant message conversations. Avrahami and Hudson (2006) found that 92% of instant messages were responded to within 5 min, with 50% of responses occurring within 15 s. This fragmentation and multitasking is not necessarily a bad thing, but knowing about how people work should allow other systems to be built to facilitate this type of work practice. In similar work, Mark et al. (2004) found that people resumed an interrupted task on the same day approximately 77% of the time.

Czerwinski et al. (2004) made similar observations in a diary study. They found not only that information workers were interrupted most often in high-priority, complex tasks, but also that these tasks were negatively impacted the most by interruptions. Additionally, they found that it is quite difficult to return to these complex tasks.

Studies of video data on naturalistic interruptions and task switches have also been conducted (Cades, Boehm-Davis, & Trafton, in preparation). Several things became obvious while these naturalistic situations were examined. First, people switched tasks quite frequently (e.g., meteorologists observed in a naturalistic setting were being interrupted every 40 s over a 2-hr period). Second, when an information worker was interrupted by an external source (i.e., another person coming in to ask a question), upon resumption the worker looked back at the computer screen and was able to figure out approximately what he or she was working on because of the state of the screen. Third, people sometimes tried to finish a subtask before turning their attention to the interrupting task. Finally, very short interruptions seemed much less disruptive than very long ones.

Before we transition to a discussion of the candidate theories to best explain why interruptions are disruptive, it would be helpful to discuss the time line of an interruption in general and the corresponding terminology that will be used throughout the remainder of this chapter. Much interest has been focused on the resumption process after a person is interrupted; thus, attention will focus on that aspect of interruptions and resumptions.

ANATOMY OF AN INTERRUPTION

In the interruptions domain, there are relatively few reported task analyses. However, a number of simple task analyses were conducted across several different domains to capture the critical aspects of the tasks. These analyses are described in more detail in Trafton et al. (2003). Figure 4.1 shows several key features of the interruption process, based on these naturalistic observations.

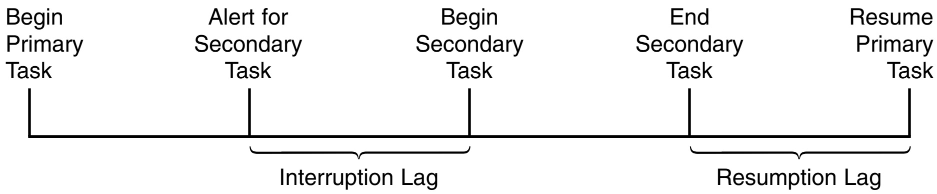

Figure 4.1. *Time line: Anatomy of an interruption.*

A person is working on a primary task, which can be thought of as similar to the complex, long-lasting task that Czerwinski et al. (2004) described. Next, an alert for a secondary task occurs. Alerts come in different forms—for example, a phone ringing, a person coming into the room to ask the person a question, or a fire alarm. During the interruption lag, the person has a moment (or longer) before turning his or her attention to the interrupting task. Then the person starts the secondary task. After the secondary task is completed, the person must resume the primary task. During the resumption lag, the person must figure out what he or she was doing during the primary task and what to do next. Finally, the person resumes the primary task.

From this task analysis and the real-world examples, it is clear that different aspects of the cognitive system are relevant to the study of interruptions and resumptions. First, executive control is very important for all interruption/resumption tasks. Second, upon completing the secondary task, the person's main goal is to remember what task he or she was working on and what to do next (though in some real-world situations, new task demands occur or the environment may have changed so that significant replanning may need to occur). Third, people may or may not use some sort of cue in the environment to actively help them remember what they had been doing. Fourth, there may or may not be a link between the primary and secondary tasks. Fifth, in some situations (e.g., an emergency), cues may not have been thought about—there may be relatively different preparatory processes that occur.

THEORETICAL EXPLANATIONS

Because task switching is such a large component of interruptions, one approach to understanding interruptions comes from the area of executive control and task switching. Within this domain, available theories vary from "loading" a task set after every task switch (Monsell, 2003) to increasing levels of task activation (Altmann, 2004). However, only Altmann has attempted to apply his task-switching theory to the study of interruptions.

A second approach comes from the perspective of prospective memory (Einstein, McDaniel, & Brandimonte, 1996). Within this domain, Dodhia and Dismukes (2003a, 2003b) and Edwards and Gronlund (1998) suggested that prospective memory is no different from other forms of memory (e.g., retrospective) but that any time a task is interrupted, a prospective memory task is created.

Typically, researchers who use this approach examine both interruption tasks and prospective memory tasks. This approach has allowed the researchers to make two predictions based on their theory that have applied implications. First, reminders should cause a person to explicitly encode an intention to resume, which should facilitate performance. Consistent with this prediction, Dodhia and Dismukes (2003a, 2003b) found that a reminder did facilitate resumptions. Coming from a more applied perspective but a similar theoretical paradigm, McDaniel, Einstein, Graham, and Rall (2004) focused on providing a simple external clue to facilitate prospective memory failings after an interruption. They constructed a task in which a single blue dot was presented as a cue to participants to perform a single specific task. Thus, the cue provided a 1:1 mapping to a particular task and proved to be effective in eliminating interruption-related prospective memory decrements.

The second prediction coming from the Dismukes and Dismukes (2003a, 2003b) approach is that the more similar an interrupting task is to the primary task, the more disruptive the interrupting task will be. However, contrary to their predictions, Dodhia and Dismukes (2003b) found that similarity of the interruption did not impact resumption of the primary task.

A third approach comes from the perspective of long-term working memory (Ericsson & Kintsch, 1995). Researchers have worked on applying these theories to the study of interruptions as well. Oulasvirta and Saariluoma (2004, 2006) have proposed using long-term working memory (LTWM) as a theoretical memory construct to reduce the problems associated with interruptions. The LTWM model suggests that experts are able to store and retrieve extremely large amounts of data quickly and accurately (Ericsson & Kintsch, 1995; Ericsson, Krampe, & Tesch-Roemer, 1993). Critical for the study of interruptions, experts are able to remember information over very long periods of time and across interruptions. There are several key components to LTWM, but the most important for interruption research are

- for experts, once information has entered LTWM, it can be retrieved at a later time regardless of interruptions; and
- for all users, information can enter LTWM if there is enough encoding time; the amount of encoding time needed for information to enter LTWM varies depending on the task, but it is typically no more than a few seconds per object or visual chunk (Oulasvirta & Saariluoma, 2006).

Following this paradigm, researchers have investigated whether interruptions are disruptive because short-term working memory is overloaded or information was poorly encoded (Oulasvirta & Saariluoma, 2004). They showed participants a short video and interrupted them with a 30-s video of spoken text or 30 s of silence. They found not only that the interrupting task disrupted recall of the information from the primary task, but also that the type of errors that occurred were from prior knowledge rather than from the information contained in the interruption. From these results, Oulasvirta and Saariluoma (2004) suggested that interruptions disrupt the encoding of information rather than overloading short-term working memory. They went on to suggest that these results are consistent with the LTWM theory.

Using the same theoretical background, Oulasvirta and Saariluoma (2006) examined why people are able to resume with relatively little disruption after an interruption. They asked participants to read expository texts and periodically interrupted them with multiplication verification problems. When they presented the sentences of the text with enough time for participants to encode that information, participants' memory for the text was as good as that in a control condition, which had no interruption. They found that when they reduced the amount of time to read the sentences so that encoding was hampered, the interruption did disrupt memory of the text, as compared with a control condition.

Oulasvirta and Saariluoma (2006) have suggested that people can use the few seconds during interruption lag in order to move information to LTWM. They have also shown that LTWM may be applicable in situations in which people are not experts at the domain itself (e.g., the participants they used were experts at reading but not experts in the subject domain). It seems it may be possible to move information to LTWM when that information is coherent and well understood.

A final potential approach derives from activation-based memory accounts (Anderson et al., 2004; Anderson & Lebiere, 1998). Specifically, a theory called *memory for goals* (Altmann & Trafton, 2002) has been developed and applied to the study of interruptions. A brief description of the theory itself is presented here, but a more detailed description can be found in Altmann and Trafton (2002) and Trafton et al. (2003).

The goal activation model is based on the hypothetical construct of activation of memory items—in particular, activation as construed in the ACT-R, Adaptive Control of Thought-Rational) cognitive theory (Anderson & Lebiere, 1998). A basic processing assumption in this theory is that when central cognition queries memory, memory returns the item that is most active at that instant. Activation thus represents relevance to the current situation. To capture the relevance of any particular item, the memory system computes that item's activation from both the item's history of use and from its associations to cues in the current mental or environmental context. In Bayesian terms, the logic is that history of use and current context together serve to predict the current relevance of that item (Anderson, 1991; Anderson, Conrad, & Corbett, 1989). In functional terms, the implication is that the cognitive system should be able to exploit the predictive computations of the memory system to overcome decay and keep certain information active for use in the future.

In the model, the history-of-use factor is captured by the following equation, which is adapted from the base-level learning equation in ACT-R.

$$Activation = \ln\left(\frac{n}{\sqrt{T}}\right) \quad (1)$$

Equation 1 computes activation as a function of frequency of use. The quantity n is the total number of times the memory item has been retrieved in its lifetime, and T is the length of this lifetime, from the item's initial representation in the system to the present. Thus, as time passes without use of an item, T for that item grows, whereas n does not, producing decay (a decrease in activation). Conversely, if concentrated use of an item causes n to grow rapidly, activation will increase. These dynamics are similar to the usual algorithm for deciding which item to replace in cache memory in a computer: Recency is a good predictor of future need, so the least recently used item in cache is the one replaced when a new item is brought in from slower memory.

Equation 1, a simplified version of ACT-R's base-level learning equation, appears to model behavior reasonably well at time scales of up to a few tens of seconds. For longer time scales, the regular base-level learning equation may provide a better account. The regular base-level learning equation in ACT-R is

$$m = \ln\left(\sum_1^n t_j^{-d}\right) + \beta,$$

in which t_j is the time lag between retrieval j and the present, d is a decay parameter that is typically set to 0.5, and β is an initial activation parameter that was set to 0. With these settings the equation simplifies to

$$m = \ln\left(2n/\sqrt{T}\right)$$

on the assumption that the t_j are evenly spaced; T represents the interval from the very first retrieval to the present (i.e., the lifetime of the trace). For convenience, Altmann and Trafton (2002) then omitted a factor of ln(2) to scale activation to be zero when $n = T = 1$, at the start of the trace's lifetime.

The memory for goals theory suggests that there are two primary ways of reducing or slowing down decay: rehearsal and using environmental cues (the two other components to the theory). Rehearsal can be either retrospective (e.g., "What was I doing?") or prospective (e.g., "What was I about to do?"). Both retrospective and prospective rehearsal are important in the model, though both the model (Altmann & Trafton, 2002) and empirical data (Trafton et al., 2003) suggest that people perform more prospective rehearsal when given the chance.

The second factor that contributes to the activation of a goal in our model is priming from contextual cues to which a goal is linked. When such cues are attended to in the environment, they spread activation to any goals with which they are associated. This spreading activation is simply added to the activation produced by the history-based mechanism represented by Equation 1. In the model, priming is necessary to retrieve a suspended goal because that goal will have decayed during the interval of the interruption.

Because the most active goal governs behavior, the most recent goal (associated with the secondary task, upon returning from an interruption) would continue to govern the system's behavior in the absence of priming of some other goal, because it has decayed less. Thus, upon resuming from a suspended goal, the system must have access to appropriate cues to allow it to remember anything about the to-be-retrieved goal.

As the memory theory is an activation-based theory, it is a straightforward prediction that the longer the interruption, the more disruptive the interruption will be. The theory makes some specific predictions about the shape of the curve as well—it is an exponential decay curve, not linear. However, the critical prediction here is that longer interruptions are more disruptive than shorter interruptions.

This prediction, however, is tempered by one of the processes within the model: rehearsal (called the *strengthening constraint*). If a goal is rehearsed, its activation increases. Thus, if there is opportunity to rehearse and rehearsal occurs, it is theoretically possible to decrease the rate of decay of the goal. When, then, does the theory claim that rehearsal can occur?

Rehearsal occurs several times throughout the model, but critically when a goal is about to be suspended. Because the original model dealt with the suspension and resumption of goals for making subgoals, not for interruptions, it was not clear initially where rehearsal would occur. Additionally, multiple types of rehearsal are available to any cognitive system: retrospective or prospective. When the model was going to suspend a goal, it performed primarily prospective rehearsal (e.g., rehearsing what the next action would be) rather than retrospective rehearsal (what the previous action was). For the interruptions task, there were multiple times to rehearse (e.g., during the interruption lag and during the secondary task). The theory would predict that given an opportunity, prospective rehearsal would occur during the interruption lag. However, it is possible within the model that rehearsal could also occur during the secondary task. So according to the theory, people will likely rehearse during the interruption lag if given the opportunity, though there may also be some rehearsal during the secondary task. This rehearsal should facilitate resumption by mitigating the decay that occurs to goals during an interruption.

The memory for goals theory thus suggests that the three most important determinants for the disruptiveness of an interruption are (a) the length of the interruption (longer interruptions are more disruptive); (b) the amount of rehearsal people do during the interruption lag; and (c) the amount of rehearsal allowed by the secondary task (i.e., the less a secondary task allows rehearsal, the more disruptive it will be).

The third major prediction (called the *priming constraint*) of the model is that environmental cues will prime the goal, facilitating resumption. This prediction suggests that when the environment changes substantially or there is little environmental context, the goal will be more difficult to resume. However, if there is a way to highlight the context via some sort of environmental cue, resumption should be facilitated. The theory provides a very specific process description: The environmental cue adds activation to the goal that was suspended by the interruption.

In this brief description, the theory makes three high-level predictions for reducing the disruptiveness of interruptions: (a) interruptions are disruptive because goals decay—the length of the interruption will impact how disruptive an interruption is; (2) the disruptiveness of interruptions can be reduced by rehearsal; and (3) the disruptiveness of interruptions can be reduced by environmental cues.

These predictions focus on "disruptiveness," which is typically operationalized as the resumption lag and the number or type of errors. However, because these predictions are based on a process theory, relatively sophisticated explanations naturally follow from the theory and provide insight into the interruptions literature (e.g., rehearsal can occur either at the interruption lag or during the secondary task, and rehearsal may be differentially useful to resuming the primary task).

One thing to note is that the memory for goals theory was originally developed using the Tower of Hanoi task, a classic psychological laboratory task within cognitive science. Whether the principles and theoretical constructs developed in that domain would transfer to different domains is primarily an empirical question. Additionally, these factors and predictions may have differential impact in different real-world tasks. For example, if someone is interrupted while talking to a friend about an event that occurred last week, there may be relatively few environmental cues because of the ephemeral nature of language; this lack of environmental cues may make resuming the conversation even more difficult than resuming a computer-related task, in which many environmental cues are available.

FACTORS INFLUENCING DISRUPTIVENESS

Length of Interruptions

As previously described, the memory for goals theory predicts that the longer the interruption, the more disruptive an interruption will be, assuming that additional rehearsal does not occur during the longer interruption. Surprisingly, other researchers have found that the duration of an interruption did not affect the disruptiveness of an interruption (Gillie & Broadbent, 1989; McDaniel et al., 2004). There were, however, several methodological issues in those studies that may not have properly captured the disruptiveness of the interruptions. For example, some researchers did not examine resumption lag; instead they measured total time on task or overall error rates, which would be highly dependent on the number of interruptions as well as the disruptiveness of those interruptions (e.g., a relatively small number of interruptions could be swamped by task or participant variability). Another feature that these experiments shared was the rather long interruption duration itself. Because the theory predicts a steep drop-off in activation early in a goal's history and a leveling off over time, it could be that previous experiments simply tested the "wrong" part of the curve (i.e., too late).

Monk, Trafton, and Boehm-Davis (under review) addressed these issues in a series of experiments. They examined the resumption lag at several durations of interruptions, both short and long, and found, consistent with the memory for goals theory, that the duration of an interruption does play a role: The longer the interruption, the more disruptive it is, as measured by resumption lag. Critically, the resumption lag slope was much steeper for short interruption durations (e.g., 3, 8, and 13 s) than it was for longer durations (e.g., 23, 38, and 58 s). Other researchers have also found a similar effect for interruption length (Hodgetts & Jones, 2006b).

Rehearsal

A second factor predicted to have an impact on performance is rehearsal. If people are given an opportunity to rehearse during the interruption lag, they will do so. Furthermore, the type of rehearsal will be prospective in nature rather than retrospective. People may also rehearse during the secondary task, but that is not a strong prediction of the model.

To examine if and when people rehearse, and whether rehearsal would facilitate an interruption, Trafton et al. (2003) conducted a study in which some participants were given an opportunity to rehearse during the interruption lag and other participants had no such opportunity. They found that people did, in fact, rehearse during the interruption lag and that this rehearsal facilitated the resumption process. People also rehearsed during the secondary task, but it was much less common. Finally, they found, consistent with the memory for goals theory, that participants performed more prospective rehearsals than retrospective rehearsals. Specifically, the most common form of preparation involved prospective encoding of a goal to be achieved at resumption, with the goal formulated in terms of specific values or items visible on the display for the primary task. For example, a participant might rehearse the next specific button to click on an interface in order to make that goal more "active" upon resumption. So the second major prediction of the memory for goals theory, that rehearsal facilitates task resumption, was supported by this study.

Role of Environmental Cues in Resuming a Task

According to the memory for goals theory, environmental cues can facilitate task resumption by priming the suspended goals. However, the type of environmental cue that will facilitate resumption is an open question. Would a very subtle cue be enough? Would a very "obnoxious" cue (e.g., a blinking cue) be too much? Would paying attention to the cue itself take so much time that there is no overall net gain?

Fortunately, there has been some evidence on this point. There is some evidence that a subtle environmental cue (e.g., a highlighted cursor) does not facilitate task resumption (Czerwinski, Cutrell, & Horvitz, 2000a). Using a noninterruption paradigm, Chung and Byrne (in press) found that subtle cues were far less effective than a more obvious or blatant environmental cue. Trafton, Altmann, and Brock (2005) found that a subtle cue (a cursor marking the spot of the last action) did not facilitate resumption any better than no cue at all. However, a more blatant cue (a large red arrow pointing to the last action taken) did facilitate performance, as compared with no cue. In the case of interruptions, a blatant cue that helps people get back on track does seem to facilitate resumption, consistent with the memory for goals theory. The three main predictions of the memory for goals theory do seem to have some empirical support. There are, of course, other predictions, but these three predictions form the core of the theory.

Speedup After Interruptions

One finding that most researchers in the interruption literature agree on is that interruptions are disruptive in diverse ways (Altmann & Trafton, 2004; Trafton et al., 2003, 2005), although see Gillie and Broadbent (1989) for a slightly different view. In some cases, however, interruptions may actually increase performance, at least for simple tasks (Speier et al., 1999; Speier, Vessey, & Valacich, 2003). In exploring the disruptiveness (or lack thereof) of interruptions during very simple tasks, Ratwani and Trafton (2006a) interrupted participants during a very simple task: transcribing odd numbers from a list of random numbers. They found that the first step performed after the interruption (measured as the resumption lag) was disrupted. However, on subsequent steps, these participants actually performed faster than their noninterrupted counterparts! Ratwani and Trafton (2006a) showed that the speedup after the interruption is attributable to perceptual speedup (quicker fixations), rather than to a general speedup in the entire cognitive system. This work and that done by Speier et al. (1999, 2003) suggest that during simple or boring or repetitive tasks, a simple interruption can actually facilitate performance for a short time.

WHAT CAN BE DONE? THEORY-BASED RECOMMENDATIONS

Training

One of the most striking findings in many studies is that the effectiveness of strategies in mitigating the immediate disruptiveness of interruptions (e.g., rehearsal, environmental cues) is strong early in the experiment but much less later on. For example, when compared with the absence of environmental cues (in a control condition), the environmental cues discussed previously (Trafton et al., 2005) were significantly more effective early in the experiment (i.e., during the first of three 20-min sessions) but much less effective later in the experiment (i.e., during the last of three 20-minute sessions). An almost identical pattern was found for the "preparing to resume" study discussed previously (Trafton et al., 2003). This decline in effectiveness could be attributable either to increased expertise in the task or to increased expertise in how to resume.

The implications for this difference are quite profound from a training viewpoint. If people are simply learning the task, training people on the task itself would reduce the disruptiveness of the interruptions. However, if people are learning how to resume, then training on interruptions and resumptions should be built into current training regimens. To explore this training issue, Cades, Trafton, and Boehm-Davis (2006) ran a study in which participants were given the same amount of training and experience on the primary task but different amounts of experience on the interruption-resumption process. The only difference between conditions was the number of times that people were interrupted. They found that interruptions became less disruptive over time with experience and practice on the resumption process itself; experience on the primary task alone (without interruptions) did not reduce the disruptiveness of interruptions. This study strongly suggests that for training people in complex domains, training scenarios should include occasional naturalistic interruptions in order to reduce the disruptiveness of interruptions during actual performance.

Guidelines

Each theoretical perspective can be used to develop guidelines for reducing the detrimental effects of interruptions. For example, within the prospective memory framework, McDaniel et al. (2004) found that the use of a blue dot cue could improve performance upon resumption of the task. This suggests that providing an external mnemonic may greatly benefit performance for people who deal with real-world interruptions and prospective memory tasks.

Using the LTWM perspective, Oulasvirta and Saariluoma (2006) also made several applied suggestions. Based on the results of their experiments, they suggested that system designers should keep "interaction chains" (the number of interface actions that lead to a goal or subgoal) quite short. The amount of time does not seem to be theoretically determined, but 20 s seems to be a heuristic used by some designers. They also suggested preventing interruptions on tasks that require large amounts of encoding time (e.g., certain checklists that airline pilots go through). Finally, they suggested that user control of interruptions is beneficial (consistent with McFarlane, 2002) because it allows the person to have control over the encoding time, which is critical under their framework.

Finally, Table 4.1 presents guidelines arising out of the memory for goals theory and the theoretical justifications for why they should be followed.

Tools

Czerwinski's group at Microsoft Research probably has the best track record of building good tools that are based on theoretically grounded applied research principles. For example, the diary study described previously (Czerwinski et al., 2004) had a major impact on the creation of a prototype tool called the GroupBar (Smith et al., 2003). GroupBar allows people to save and retrieve application and window management setups, which can be extremely useful when switching tasks. (The Czerwinski et al., 2004, study found that window management was a tedious aspect of task switching and task resumption.)

Bailey and his colleagues also have built several tools based on empirical work (Adamczyk & Bailey, 2004; Bailey & Konstan, 2006; Bailey et al., 2000; Iqbal & Bailey, 2005; Iqbal, Zheng, & Bailey, 2004). They suggested that the best place to interrupt people is between "coarse" breakpoints between tasks (Adamczyk & Bailey, 2004; Monk et al., 2004). They have used an empirical approach to explore the linkages between traditional task-analytic approaches (e.g., goals, operators, methods, and selection rules, or GOMS) and pupil size as a measure of mental workload (Iqbal et al., 2004). Bailey, Busbey, and Iqbal (2007) have created a tool that is able to automatically detect times of high and low workload. They have suggested that interrupting people at times of low workload is best. They currently have several demonstration systems that perform components of this task.

The memory for goals theory has spawned several applied pieces of work. One of the more interesting ones is a tool that keeps track of a user's actions. When the user switches tasks, the tool notices it and displays a series of blatant environmental cues showing the last three actions performed on the new task before the switch. For example, imagine a

Table 4.1. Interface Guidelines Inspired by the Memory for Goals Framework

Suggestion	***Theoretical Reason***
Minimize interruptions	Interruptions are disruptive (whole theory)
Do not set your e-mail to automatically alert you when you get e-mail	Minimize external interruptions
Turn off all "intelligent" agents that interrupt you	Minimize external interruptions
When you get interrupted, take 2 s to figure out what you will do next	Rehearsal during interruption lag facilitates resumption
Make the next thing you do an action on a visible object	Environmental cues can prime previously suspended goals
Highlight the next thing you want to do upon resumption	Explicit, blatant environmental cues help the resumption process

user switching between two tasks (e-mail and making a graph in Microsoft Excel). When switching to the graphing task after an interruption, it would be helpful to present some environmental or contextual cues. Figure 4.2 shows a gray-scale version of a color screen snapshot of the graphing task interface immediately after an interruption. The three arrows show the user's previous three steps in color order from light red (top arrow) to bright red (bottom arrow). This type of color order has been shown to be a natural way of presenting ordered and quantitative data (Spence, Kutlesa, & Rose, 1999; Trafton, Ratwani, &, Breslow, under review).

Upon resumption, the user can determine what he or she had been working on and, from that, determine what should be done next. This type of interface facilitation was shown to strongly support resumption in earlier empirical work (Trafton et al., 2005), although the task was quite different from traditional office work.

CONCLUSIONS

This paper has reviewed the interruption literature from several different theoretical and applied perspectives. A number of strong theories in the literature describe some of the cognitive processes that occur when one is interrupted and the resumption process itself. One of the most interesting similarities between the theories is that they are all basically memory-based theories. None of the theories has an especially strong link to perception, action, or other cognitive functions. At one level, this suggests that most researchers believe that interruptions and resumptions are memory-based phenomena. All these theories capture interesting, unique components of the interruption-resumption process, although all of them could clearly be improved in many ways.

From a human factors point of view, an enormous amount of data have been collected using an interruption paradigm. The data come from both relatively simple tasks and

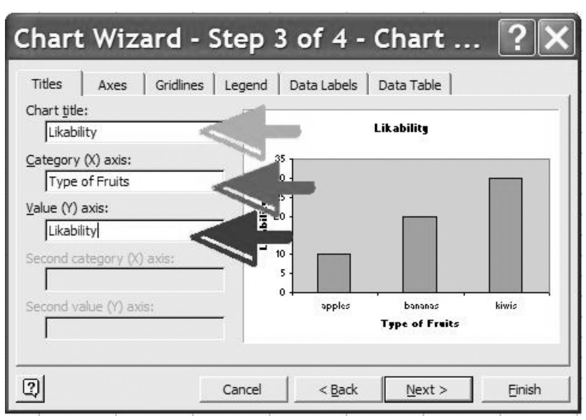

Figure 4.2. *Example of resumption cues in the interface (gray-scale reproduction of color screen shot.)*

very complex tasks and situations. Fortunately, from an applied point of view, most of the experiments have been performed using complex tasks, allowing strong generalizability to real-world tasks. There have also been some interruption studies in the workplace (real life). Again, this further strengthens the generalizability of both the theories and the empirical results.

Although the research on interruptions has recently become a popular topic to study, the topic area is still relatively new, and much work still needs to be done at both theoretical and applied levels.

ACKNOWLEDGMENTS

This work was supported by grants from the U.S. Office of Naval Research to the first author (N0001405WX20011, N0001405WX30020). The views and conclusions contained in this document should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Navy. Thanks to Yi-Fang Tsai, Susan Trickett, Peter Squire, Raj Ratwani, Sarah Kriz, David Kidd, David Cades, and Debbie Boehm-Davis for comments on a previous draft.

REFERENCES

Adamczyk, P. D., & Bailey, B. P. (2004). If not now, when? The effects of interruption at different moments within task execution. In *Human factors in computing systems: Proceedings of CHI'04* (pp. 271–278). New York: Association for Computing Machinery Press.

Altmann, E. M. (2004). The preparation effect in task switching: Carryover of SOA. Memory & Cognition, 32, 153–163.

Altmann, E. M., & Trafton, J. G. (2002). Memory for goals: An activation-based model. *Cognitive Science, 26,* 39–83.

Altmann, E. M., & Trafton, J. G. (2004). Task interruption: Disruptive effects and the role of cues. In *Proceedings of the 26th Annual Conference of the Cognitive Science Society* (pp. 43–48). Chicago: Erlbaum.

Anderson, J. R. (1991). Is human cognition adaptive? *Behavioral and Brain Sciences, 14,* 471–485.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglas, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of mind. *Psychological Review, 111,* 1036–1060.

Anderson, J. R., Conrad, F. G., & Corbett, A. T. (1989). Skill acquisition and the LISP tutor. *Cognitive Science, 13,* 467–505.

Anderson, J. R., & Lebiere, C. (1998). *Atomic components of thought.* Mahwah, NJ: Erlbaum.

Avrahami, D., & Hudson, S. E. (2006). Responsiveness in instant messaging: Predictive models supporting interpersonal communication. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'06)* (pp. 731–740). New York: Association for Computing Machinery Press.

Bailey, B. P., Busbey, C. W., & Iqbal, S. T. (2007). TAPRAV: An interactive analysis tool for exploring workload aligned to models of task execution. *Interacting with Computers, 19*(3), 314–329.

Bailey, B. P., & Konstan, J. A. (2006). On the need for attention-aware systems: Measuring effects of interruption on task performance, error rate, and affective state. *Computers in Human Behavior, 22,* 685–708.

Bailey, B. P., Konstan, J. A., & Carlis, J. V. (2000). Measuring the effects of interruptions on task performance in the user interface. *IEEE Conference on Systems, Man, and Cybernetics* (pp. 757–762). Los Alamitos, CA: IEEE Press.

Cades, D. M., Trafton, J. G., & Boehm-Davis, D. (2006). Mitigating disruptions: Can resuming an interrupted task be trained? In *Proceedings of the Human Factors and Ergonomics Society 50th Annual Meeting* (pp. 368–371). Santa Monica, CA: Human Factors and Ergonomics Society.

Cades, D. M., Boehm-Davis, D. & Trafton, J. G., (in preparation). Exploring interruptions on the real-world.

Cellier, J., & Eyrolle, H. (1992). Interference between switched tasks. *Ergonomics, 35,* 25–36.

Chisholm, C. D., Collison, E. K., Nelson, D. R., & Cordell, W. H. (2000). Emergency department workplace interruptions: Are emergency physicians "interrupt-driven" and "multitasking"? *Academic Emergency Medicine, 7,* 1239–1243.

Chisholm, C. D., Dornfeld, A. M., Nelson, D. R., & Cordell, W. H. (2001). Work interrupted: A comparison of workplace interruptions in emergency departments and primary care offices. *Annals of Emergency Medicine, 38,* 146–151.

Chung, P. H., & Byrne, M. D. (in press). Cue effectiveness in mitigating postcompletion errors in a routine procedural task. *International Journal of Human-Computer Studies*.

Cutrell, E., Czerwinski, M., & Horvitz, E. (2001). Notification, disruption, and memory: Effects of messaging interruptions on memory and performance. In *Interact 2001: IFIP Conference on Human-Computer Interaction* (pp. 263–269). Tokyo: IFIP.

Czerwinski, M., Chrisman, S., & Rudisill, M. (1991). *Interruptions in multitasking situations: The effects of similarity and warning* (Tech. Rep. JSC-24757). Houston, TX: National Aeronautics and Space Administration, Johnson Space Center.

Czerwinski, M., Cutrell, E., & Horvitz, E. (2000a). Instant messaging and interruption: Influence of task type on performance. In *Proceedings of OZCHI 2000* (pp. 356–361). Sydney, Australia: Academic Press.

Czerwinski, M., Cutrell, E., & Horvitz, E. (2000b). Instant messaging: Effects of relevance and time. In S. Turner & P. Turner (Eds.), *People and computers XIV: Proceedings of HCI 2000* (Vol. 2, pp. 71–76). [**author: city of publication, not meeting?**]: British Computer Society.

Czerwinski, M. P., Horvitz, E., & Wilhite, S. (2004). A diary study of task switching and interruptions. In *Human factors in computing systems: Proceedings of CHI'04* (pp. 175–182). New York: Association for Computing Machinery Press.

Dismukes, R. K., Berman, B., & Loukopoulos, L. D. (2006). *The limits of expertise: Rethinking pilot error and the causes of airline accidents.* Presented at the 2006 Crew Resource Management/Human Factors Conference, Denver, Colorado.

Dodhia, R., & Dismukes, K. (2003a,). *A task interrupted becomes a prospective memory task.* Poster presented at the 15th Annual Convention of the American Psychological Society, Atlanta, GA.

Dodhia, R., & Dismukes, K. (2003b). *A task interrupted becomes a prospective memory task: Encoding and retrieval manipulations.* Poster presented at the 44th Annual Meeting of the Psychonomic Society, Vancouver, British Columbia.

Edwards, M. B., & Gronlund, S. D. (1998). Task interruption and its effects on memory. *Memory, 6,* 665–687.

Einstein, G. O., McDaniel, M. A., & Brandimonte, M. (1996). *Prospective memory: Theory and applications.* Mahwah, NJ: Erlbaum.

Ericsson, K. A., & Kintsch, W. (1995). Long-term working memory. *Psychological Review, 102,* 211–245.

Ericsson, K. A., Krampe, R., & Tesch-Roemer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review, 100,* 363–406.

Gillie, T., & Broadbent, D. E. (1989). What makes interruptions disruptive? A study of length, similarity, and complexity. *Psychological Research, 50,* 243–250.

Gonzalez, V. M., & Mark, G. (2004). Constant, constant, multi-tasking craziness: Managing multiple working spheres. In *Human factors in computing systems: Proceedings of CHI'04* (pp. 113–120). New York: Association for Computing Machinery Press.

Hodgetts, H. M., & Jones, D. M. (2006a). Contextual cues aid recovery from interruption: The role of associative activation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32,* 1120–1132.

Hodgetts, H. M., & Jones, D. M. (2006b). Interruption of the Tower of London task: Support for a goal-activation approach. *Journal of Experimental Psychology: General, 135,* 103–115.

Iqbal, S. T., & Bailey, B. P. (2005). Investigating the effectiveness of mental workload as a predictor of opportune moments for interruption. In *Conference on Human Factors in Computing Systems* (1489–1492). New York: ACM Press.

Iqbal, S. T., Zheng, X. S., & Bailey, B. P. (2004). Task-evoked pupillary response to mental workload in human-computer interaction. In *Conference on Human Factors in Computing Systems* (pp. 1477–1480). New York: ACM Press.

Isaacs, E., Walendowski, A., Whittaker, S., Schiano, D. J., & Kamm, C. (2002). The character, functions, and styles of instant messaging in the workplace. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW'02)* (pp. 11–20). New York: Association for Computing Machinery Press.

Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin, 114,* 3–28.

Mark, G., Gonzalez, V. M., & Harris, J. (2005). No task left behind? Examining the nature of fragmented work. In *Human factors in computing systems: Proceedings of CHI'05* (pp. 321–330). New York: Association for Computing Machinery Press.

McDaniel, M. A., Einstein, G. O., Graham, T., & Rall, E. (2004). Delaying execution of intentions: Overcoming the costs of interruptions *Applied Cognitive Psychology, 18,* 533–547.

McFarlane, D. C. (2002). Comparison of four primary methods for coordinating the interruption of people in human-computer interaction. *Human-Computer Interaction, 17,* 63–139.

Monk, C., Boehm-Davis, D., & Trafton, J. G. (2004). Recovering from interruptions: Implications for driver distraction research. *Human Factors, 46,* 650–663.

Monk, C., Trafton, J. G., & Boehm-Davis, D. (under review). *The effect of interruption duration on resuming suspended goals.* Manuscript submitted for publication.

Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences, 7,* 134–140.

National Transportation Safety Board. (1969). *Aircraft accident report: Pan American World Airways, Inc., Boeing 707-321C, N799PA, Elmendorf Air Force Base, Anchorage, Alaska, December 26, 1968* (NTSB/AAR-69/08). Washington, DC: Author.

National Transportation Safety Board. (1988). *Aircraft accident report: Northwest Airlines, Inc., McDonnell Douglas DC-9-82, N312RC, Detroit Metropolitan Wayne County Airport, Romulus, Michigan, August 16, 1987* (NTSB/AAR-88/05). Washington, DC: Author.

Oulasvirta, A., & Saariluoma, P. (2004). Long-term working memory and interrupting messages in human-computer interaction. *Behaviour and Information Technology, 23*, 53–64.

Oulasvirta, A., & Saariluoma, P. (2006). Surviving task interruptions: Investigating the implications of long-term working memory theory. *International Journal of Human-Computer Studies, 64,* 941–961.

Ratwani, R. M., & Trafton, J. G. (2006a). Helpful or harmful? Examining the effects of interruptions on task performance. In *Proceedings of the Human Factors and Ergonomics Society 50th Annual Meeting* (pp. 372–375). Santa Monica, CA: Human Factors and Ergonomics Society.

Smith, G., Baudisch, P., Robertson, G., Czerwinski, M., Meyers, B., Robbins, D., et al. (2003). GroupBar: The TaskBar evolved. In *Proceedings of OZCHI, 2003*

Speier, C., Valacich, J. S., & Vessey, I. (1999). The influence of task interruption on individual decision making: An information overload perspective. *Decision Sciences, 30,* 337–360.

Speier, C., Vessey, I., & Valacich, J. S. (2003). The effects of interruptions, task complexity, and information presentation on computer-supported decision-making performance. *Decision Sciences, 34,* 771–797.

Spence, I., Kutlesa, N., & Rose, D. L. (1999). Using color to code quantity in spatial displays. *Journal of Experimental Psychology: Applied, 5,* 393–412.

Spira, J. B., & Feintuch, J. B. (2005). *The cost of not paying attention: How interruptions impact knowledge worker productivity* (Technical Report). New York: Basex.

Thompson, C. (2005, October 16). Meet the life hackers. *New York Times Magazine,* pp. 40–45.

Trafton, J. G., Altmann, E. M., & Brock, D. P. (2005). Huh, what was I doing? How people use environmental cues after an interruption In *Proceedings of the Human Factors and Ergonomics Society 49th Annual Meeting* (pp. 468–472 .). Santa Monica, CA: Human Factors and Ergonomics Society.

Trafton, J. G., Altmann, E. M., Brock, D. P., & Mintz, F. E. (2003). Preparing to resume an interrupted task: Effects of prospective goal encoding and retrospective rehearsal. *International Journal of Human-Computer Studies, 58,* 583–603.

Trafton, J. G., Ratwani, R. M., & Breslow, L.(under review). *A color perceptual-process theory.* Manuscript submitted for publication.

Tucker, A. L., & Spear, S. J. (2006). Operational failures and interruptions in hospital nursing. *Health Services Research, 41,* 643–662.

Wallis, C., & Steptoe, S. (2006, January 16). Help! I've lost my focus." *Time,* 167, pp. 72–79.

Zijlstra, F. R. H., Roe, R. A., Leonora, A. B., & Krediet, I. (1999). Temporal factors in mental work: Effects of interrupted activities. *Journal of Occupational and Organizational Psychology, 72,* 164–185.