

Disruption and Recovery of Computing Tasks: Field Study, Analysis, and Directions

Shamsi T. Iqbal

Department of Computer Science
University of Illinois
Urbana, IL 61801 U.S.A
siqbal@cs.uiuc.edu

Eric Horvitz

Microsoft Research
One Microsoft Way
Redmond, WA 98052 U.S.A
horvitz@microsoft.com

ABSTRACT

We report on a field study of the multitasking behavior of computer users focused on the suspension and resumption of tasks. Data was collected with a tool that logged users' interactions with software applications and their associated windows, as well as incoming instant messaging and email alerts. We describe methods, summarize results, and discuss design guidelines suggested by the findings.

Author Keywords

Interruption, Attention, Task Switching, Notifications.

ACM Classification Keywords

H.5.2 [Information Interfaces and Presentation]: User Interfaces.

INTRODUCTION

Task switching is common in computing. Several decades ago, Bannon *et al.*[3] noted that computer users often switched among multiple active tasks. The diversity and numbers of applications supported by personal computers has grown since the Bannon study and multitasking has now become a salient feature of modern computing. Today, computer users often run programs simultaneously to support multiple tasks, including word processing, financial analysis, searching, browsing, and communications.

Card and Henderson [4] attempted to characterize useful attributes of designs for computer-based task management, stressing the need to allow for efficient task switching and resumption, and to provide methods for assisting with refreshing a task context. Today's major operating systems include tools in line with these recommendations, such as providing multiple means for switching among tasks. However, efficient shifting ability does not mean that a suspended task will be resumed efficiently. Multiple active

computing tasks and the opportunity to spawn new tasks compete with returns to specific tasks, interfering with the resumption of tasks following their disruption. The timing of shifts among related and disjoint computing tasks is often self-directed, occurring in the absence of explicit external influences. However, task switching may be affected by external signals and events [9]. Such influences include alerts delivered to computer users from applications that are not at the focus of a user's attention. For example, a computer user may be drawn to switch from a spreadsheet program to their email application after hearing or seeing an alert about incoming email or receiving an instant message.

We sought to characterize task suspension and recovery among information workers in the course of their normal daily computing tasks. We developed and deployed a disruption and resumption tracking tool to monitor the use of software applications and associated windows at the focus of computer users' activities, as well as to log incoming instant messaging and email alerts. Rather than seek only to measure the specific effect of an alert on a task at focus, we also pursued patterns and understanding of user behavior before and after interruptions. We have particularly worked to understand the chain of diversions whether likely caused by an alert or by a self interruption, and the path and timing back to the resumption of tasks. The work includes an analysis of behaviors of users before they suspend tasks, and to examine behaviors that would suggest a preparation for more efficient resumption of a task upon return. We also sought to better understand the relationships between actions prior to the suspension and time taken to resume suspended tasks, and factors that promote returns to suspended applications.

We first review related work. Then, we review the methods that we used to study task interruption, diversion, and resumption in real-world computing situations. We summarize results of analyses of the logged activity and of interviews of subjects. Finally, we provide a set of design guidelines based on the lessons gleaned from the data and from interviews of participants.

RELATED WORK

We first motivate our work by providing some background on several studies of interruption and recovery.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2007, April 28–May 3, 2007, San Jose, California, USA.

Copyright 2007 ACM 978-1-59593-593-9/07/0004...\$5.00.

Interruptions and their disruptive effects

Maintaining information awareness and near instant communication in the workplace is becoming increasingly important to knowledge workers to support collaborative practices and increase productivity [6, 10, 11]. However, the pursuit of awareness and rapid communication often injects increased numbers of notifications and potential interruptions into ongoing tasks [9, 15, 24]. Several research teams have explored interruptions of computing tasks. Recent efforts come in the context of a rich history of research in cognitive psychology on the influence of interruptions on human memory and planning, going back to the protean efforts of Zeigarnik and Ovsiankina [25, 32].

Czerwinski, Cutrell, and Horvitz in a series of studies have explored the effects of external interruptions on task switching behavior and performance, and have also investigated the impact of varying timing and type of interruption [5, 6, 9]. Iqbal and Bailey have shown that interruptions during periods of higher mental workload cause users to take longer to resume their suspended tasks and have larger negative affect [19]. Mark et al. have sought to understand the influence of interruptions on task switching and found that users frequently switch between tasks and 57% of their activities are interrupted [22]. Others have also investigated effects of interruption on error rates [21], decision making [28] and affective state such as frustration, annoyance and anxiety [2, 33]. Researchers have also investigated methods that could decrease the cost associated with communication alerts. Approaches explored to date include reducing the frequency and costs of interruptions through identifying the attentional state of users [13, 16, 18] and deferring or scheduling notifications in an intelligent, strategic manner [13, 20, 23], and providing support for recovering from interruption [27, 29].

Recovering from interruptions

Research has shown that inopportune interruptions can increase task performance time, primarily due to increases in the time to resume suspended tasks [1, 19, 20]. Cognitive models suggest that when the workload of the ongoing task is high, interruptions cause users to divert cognitive resources to the interrupting task [31]. On return from the interrupting task, users have to reallocate resources to the suspended task, which becomes increasingly difficult if the resource demands were high to begin with. The result is higher resumption lag, which affects recovery. With users typically suspending sets of applications [9, 22], recovery is often confounded with a cycling through and visiting of multiple suspended applications on the way to resuming a task.

Designs for recovering from interruptions

Software tools have been developed across a variety of application domains with the goal of supporting ease of resuming suspended applications after responding to an interruption [12, 14, 26, 27]. These tools provide visualizations of suspended application states and group

applications based on time proximity, leveraging the use of visual cues, or interfaces for interactive dialogue to help users quickly regain the suspended task context.

We believe that the challenge is not only one of resumption of the suspended application but also one of turning one's focus of attention to the suspended tasks, given other tasks competing for the user's attention. Beyond characterization of the suspension and resumption behavior, we seek to determine (i) how we might best help people to break away from potentially costly 'chains of diversion' following suspension, so as to return to suspended tasks within a time they would desire and (ii) how to help them quickly resume where they left off, once they return to continue on a task.

OVERVIEW OF STUDY

We conducted a field study to better understand task suspension and resumption in practice. We were interested in the influences of computer-based alerts on users' task execution behaviors. Specifically, we sought to explore effects of interruption on task switching and the path and timing back to the resumption of suspended primary tasks. By *primary tasks* we refer to normal daily tasks that users perform as their primary responsibility while in the computing environment. For our study population, this typically entailed programming or content generation tasks, e.g., document editing or creation of presentation material. By *alerts*, we refer to notification cues generated by email clients and instant messaging applications.

In the simplest case, an alert influences the probability that a user will switch to the alerting application with a concomitant suspension of the ongoing primary task and, some time later, will resume the primary task after responding to the alert. However, when users suspend a task because of an alert or for other reasons, they may take advantage of the break in the execution of the primary task offered by the switch to interact with other peripheral applications, and perhaps turn to other tasks. We sought to gain a deeper understanding of how users prepare for the context switch from the primary task to the alert response, how a succession of diversions after a task switch may interfere with a return to their primary tasks, and how they eventually pass through a chain of diversions on the way to resuming a suspended primary task. More specifically, we explored the following hypotheses:

H1: Users prepare to address alerts in their regular task execution by stabilizing their current task state before switching to the alerting application.

H2: Users are less focused on applications visited during the 'chain of diversion' and during resumption.

H3: The chain of diversion mostly consists of rapid interactions with communication and awareness applications.

H4: Availability of cues about suspended tasks assists with resumption of tasks.

H5: Users have difficulty with resuming interrupted desktop computing tasks.

H6: The time to resume a primary task is influenced by the recency and focus of attention on a task before suspension.

Our study was designed to gather evidence from users *in situ* to investigate these hypotheses as well as gain a basic understanding of the prevalence of alerts in practice and the length of time users typically spend on chains of diversions initiated by these alerts. We began with defining a task disruption and resumption lifecycle. Each phase in the cycle signifies a distinct user goal along the path of suspending and returning to an interrupted task. We then defined a set of task state attributes to characterize behaviors across these different phases. We developed a disruption and recovery logging tool by extending an existing user-activity monitoring system. The tool was deployed to log data from users over a period of two weeks. The collected data was analyzed and findings were corroborated through interviews of the study participants. Finally, the findings from the study were distilled into key results and a set of design guidelines for enhancing the recovery of suspended tasks.

PHASES OF AN INTERRUPTION LIFECYCLE

We divide the time following an alert into distinct temporal segments or phases. Our intent is to measure the impact of the interruption by comparing behavioral changes across these phases as users sequentially move through a cycle including focused attention on a primary task, alert arrival, response and diversion, return from diversion and the resumption of original task. A related categorization of aspects of interruption is provided in [14]. One of our key goals was to better understand natural user behavior during each phase so as to inform the design of tools that might assist computer users with multitasking.

Figure 1 displays phases of the interruption lifecycle. We define the initial phase of the interruption lifecycle, which we call the *pre-interruption* phase, as the time between an alert and the concomitant suspension of ongoing tasks. Based on prior research showing response time to be a function of task state [13], we hypothesize that during this phase, the user may consciously or subconsciously perform activities that leave the primary task in a more stable state, before switching to the alerting application.

Phase 2 is the *diversion* phase, defined as the time between the switch from the primary task to respond to the alert and the return to the primary task after the response. During this period the goal is to access the interrupting application but users may also explore other peripheral applications.

Phase 3 is the *resumption* phase, where the user finishes interactions with interrupting and peripheral applications and seeks to a return of conceptual context and focus to become active once again in the primary task. Since it is difficult to identify exactly when the resumption phase may begin, we used cues indicating user intent to terminate the diversion and resume suspended work, *e.g.*, minimizing or




Figure 1. Phases of the interruption lifecycle. a) User begins interaction with two applications on a primary task, continuing through a pre-interruption phase; b) alert arrives and user enters a response preparation phase; c) user suspends primary task and switches to interrupting application, and may become diverted to other peripheral applications; d) user returns to resume primary task.

closing applications accessed in the diversion phase and starting to resume applications from the suspended group. As users can be active in a task in the absence of computing activity (*e.g.*, reading text), we used a simple heuristic to determine resumption: we considered users to resume a suspended task if they had spent more than 15 seconds on the suspended application, which is more time than required for rapid application switches, *e.g.*, tabbed browsing.

To compare users' actions in the aforementioned phases to behaviors seen during task execution behavior, we defined an additional phase, *pre-interruption*, which refers to a predefined time segment of activity before the arrival of an alert. Behaviors during this period provide a baseline for comparison to the same activities in the other phases, thereby providing an opportunity to demonstrate the influences of interruptions on users' task execution.

FORMULATING TASK ATTRIBUTES

With an eye to investigating our hypotheses, we defined a set of events that promised to allow us to probe sets of relevant activities during the different phases of the interruption lifecycle. The definition of events was guided by intuitions about behavioral patterns we expected to see at different phases. For example, we were interested to see if users tended to perform activities during the preparation phase that provided evidence that they were attempting to leave the ongoing task in state that could be resumed with more efficiency (H1). We defined a set of events that promised to provide insights about such potential efforts. These include the number of save operations, the completion of edit operations (sentence or paragraph

completion), and the completion of pending tasks lacking physical representation (pasting content copied into a memory clipboard).

Also, since alerts originate largely from communication applications, we wished to explore whether suspension of primary tasks was associated with users interacting with communication-centric applications, beyond switching to the alerting application, during the diversion phase (H3). To gather evidence on the diversion phase, common email interactions, including mail opens, writes, and sends, were monitored

To study the potential influence of visual cues on the timing of transitions from the diversion to the resumption phase, we developed attributes for capturing the visibility of open windows. Such monitoring promised to allow us to examine the effects of cues associated with suspended task windows on the time to return to those tasks (H4).

We sought to better understand the difficulty that people had with resuming applications that has been suspended as a result of responding to an alert (H5). We decided to quantify difficulty in terms of the time taken by users to not only return to the suspended application after responding to an alert, but also to restore context and state, and to resume the activity they were engaged in before switching tasks in response to the alert. We were also interested in exploring how the intensity of the focus of attention on a task and task recency influenced the time until task resumption after an interruption. We created attributes capturing the duration of time and the last time that users were focused on a particular task, with a goal of exploring the relationship between these coarse measures of focus and salience to the time until a user would completely return to a task (H6). As a related task state attribute, we defined the rate of task switches at each phase - a potential indicator of user focus. We hypothesized that users would switch tasks at a higher rate as they sought to return to suspended applications, especially if they were searching for a window associated with a primary task application among multiple open windows, as part of an attempt to regain task context (H2).

A DISRUPTION AND RESUMPTION TRACKING TOOL

We developed a monitoring tool named DART (for Disruption and Recovery Tracker). The tool was developed on top of the Eve event-monitoring infrastructure, a constellation of user and system monitoring components that have been evolving for over a decade at Microsoft Research [17]. Eve components have been employed in prior research on interruptions in computing settings, including efforts on inferring the cost of interruption and on guiding alerts and information awareness based on cost-benefit analyses [16, 18]. DART runs as a background process, and continues to logs the name, size, and location of all windows on a computing system, noting the opening and closing of windows. The system also logs user activities, including when users are actively engaged with the software, keyboard and mouse activity, and switches

among windows as well as actions of saving, cutting, and pasting. To protect privacy, window titles were truncated, and only a subset of keyboard events were recorded, including the input of periods and carriage returns (which can indicate sentence or paragraph completion and data entry), and shortcuts for saving, cutting, and pasting. The tool also logs alerts from email and IM systems.

DART logged user actions relevant to the aforementioned task state attributes and corresponding timestamps. Logged data files were periodically flushed to a server, where the data was preprocessed to generate the attributes and stored in a SQL database for further analysis.

DEPLOYMENT OF TOOL AND COLLECTION OF DATA

We deployed DART on the primary machines of 27 people at our organization, whose job descriptions ranged from program manager, administrator, and researcher to software developer. On recruiting subjects, we sought a balance of people who focused on different kinds of tasks as primary, including software development, working on productivity applications, and/or managing large numbers of external communications. We did not screen users for whether they used alerts in communications, but excluded from analysis those subjects who did not have alerts enabled.

We collected 2,267 hours of activity data over a period of 2 weeks, resulting in 974 sessions ($M(\text{session length})=2\text{h}, 17\text{m}, S.D.=410.37\text{ m}$). A session was defined as delimited by either the logging on and off or by the unlocking and locking of a machine. Collected data included logs of application access, window sizes and configurations, file retrieval and archival, percentage visibility of open application windows, key events corresponding to content manipulation, *e.g.* cut, copy and paste, file open and save, and completion of text generation.

Users were informed a priori about the overall nature of the data that was to be collected and informed that they would be able to quit the study at any point if they were not comfortable or if the software was perceived to influence the performance of their computers. Users were not informed that the study was investigating disruption and resumption of tasks. As in any field study, the knowledge of being studied potentially may have had influences on the behavior of the subjects. However, we believe the study had little influence on participants. A number of the subjects mentioned during interviews that they had forgotten about the tool running in the background.

ANALYSIS AND RESULTS

We focused on characterizing the behaviors of users in response to alerts generated by Outlook, a widely used email client, and IM clients, including Windows Messenger, MSN Messenger, and Office Communicator. It is difficult in the general case to determine with certainty whether an interruption of a current task is a direct consequence of an alert or if a switch results from users making a decision to switch away from a task largely

independent of recent notifications. We employed a simple heuristic to identify suspensions likely to have been driven by alerts: switches to a notifying application (Outlook or IM) occurring within 15 seconds of the alert were considered as being caused by that alert. Our later interviews with users further raised our confidence about the robustness of this heuristic. Our analysis showed that for such switches, users take on average 2.35 seconds (S.D.=1.39s) to switch to Outlook and 1.72 seconds (S.D.=1.1s) to switch to the IM client.

We distinguish between immediate and delayed responses to alerts in our presentation of the results so as to explore differences in activities during the preparation phase. As user actions for email and IM alerts could vary based on such influences as social conventions and expectations, we analyze email and IM results separately. We also examine results across developers, researchers and managers and note if a significant influence of job role is found.

Rather than measure effects over the entire pre-interruption phase, as the baseline condition we consider 5 minutes of activity prior to the interruption. Preliminary analysis showed that, on average, the maximum time spent on an application before switching to another is just above 4 minutes, with an average of below a minute.

Distribution of alerts

Overall, we found that, on an hourly basis, a user's primary tasks were interrupted by an average of 4.28 email (S.D.=5.56) alerts and 3.21 IM (S.D.=4.31) alerts, with an overall average rate of 3.74/hour (S.D.=4.94). For IM alerts, our system did not discriminate between conversational pings, and sign-in and presence status alerts. We are more interested in attempts to initiate or continue conversation as these pings pose a social obligation to respond; nonetheless, sign-in alerts also affect awareness of the user and may serve as a subtle trigger to self-interruption, *e.g.*, if the user wishes to communicate with the person who just signed in. Job roles did not significantly affect the number of alerts.

Time to respond to alerts

After an alert was delivered, users took, on average, 4 minutes, 59 seconds (S.D.=8m, 43s) to suspend their primary task and switch to Outlook, 7 minutes, 54 seconds (S.D.=16m, 50s) to switch to MSN messenger, 7 minutes, 13 seconds (S.D.=13m, 49s) to switch to Communicator, and 34 seconds (S.D.=1m, 5s) to switch to Windows Messenger. There were no significant differences in the response times of the different alerts, nor were there any significant effect of job roles on the response times.

Immediate and Delayed Responses to Alerts

For 40.8% (2344/5747) of the email alerts, users responded immediately (<15s), leaving on average 3 (S.D.=1.92) task windows suspended. The mean response time for such immediate responses was 2 seconds (S.D.=3.77s). Users did not interact with any other application between the alert and the response for these switches, but did appear to quickly

tab through on average 7.5 applications (S.D.=2.88s), in their pursuit of the application that generated the alert. For the remaining 59.2% of email alerts, the average response time was 7 minutes, 32 seconds (S.D.=11m), suggesting the switch was self initiated and that time was spent reaching a stable state. The delay in response also alludes to an internal deliberation about when users desired to switch to the alerting application, as opposed to switching immediately as a direct effect of being interrupted by the alert. For example, users may have a background rate of checking email independent of alerts or may delay their switch if the alert provides information that the incoming email is of low priority. On average, 3 (S.D.=2.1) task windows were left suspended.

71.01% (3186/4487) of the IM alerts resulted in an immediate response, the average response time being 1.72 seconds (S.D.=1.1s). No other application was accessed in between. The remaining 28.99% (1301/4487) of IM alerts had an average response time of 8 minutes, 48 seconds. As with email, the number of suspended windows also averaged 3 (S.D.= 2.3).

Time spent on responses and diversions

Regardless of the delay in responding to alerts, the time spent on responding to alerts and subsequent diversions to peripheral applications were found to be similar for both email and IM. For email, the average time to return to any suspended application (time spent on the diversion) was 9 minutes and 33 seconds (S.D.=13m, 15s). For IM, the return time was 8 minutes (S.D.=11m, 32s) on average. We note that return times refer to the time until accessing the application associated with the primary task, not resumption of a suspended task. However, we took such returns as proxies for the intent to resume tasks.

Stabilizing task state before responding to alerts (H1)

We found that users selectively perform certain operations at a higher rate during the preparation phase than in the pre-interruption phase. For example, users completed paragraphs at a rate of 0.78/min (S.D.=0.49) during the pre-interruption phase. On receiving an IM alert, the paragraph completion rate was 10.9/min (S.D.=5.2) during the preparation phase if the response was immediate. On receiving an email alert, the paragraph completion rate was 12.8/min (S.D.=13.75) if the response was immediate and 1.17/min (S.D.=0.49) if the response was delayed. All rates during the preparation phase were significantly higher than rates for the pre-interruption phase ($p < 0.038$ for all cases). These results suggest that users typically prefer to complete conceptual and/or motor subtasks before switching and do so quickly before responding to an alert.

The number of paste operations during the preparation phase for delayed email responses ($M = 0.69/\text{min}$, $S.D. = 0.62$) was marginally higher ($p < 0.067$) than those in the pre-interruption period ($M = 0.37/\text{min}$, $S.D. = 0.12$). We take these results as evidence that users may selectively work to externalize pending copy-paste goals, perhaps since

Phases	Email Alerts Mean(S.D)	IM Alerts Mean(S.D)
Pre-interruption	0.84(0.6)	0.84(0.6)
Preparation: Immediate Response	0	0
Preparation: Delayed Response	0.69(0.62)	0.49(0.45)
Diversion: Immediate response	1.33(1.95)	1.42(2.36)
Diversion: Delayed Response	1.24(2.01)	1.36(2.33)
Resumption: Immediate Response	2.34(2.71)	2.56(2.82)
Resumption: Delayed Response	2.42(3.26)	2.05(2.14)

Table 1. Task switches per minute across different phases for email and IM alerts. The diversion and resumption phases have higher switch rates, suggesting less engagement with any one task during these phases.

such tasks do not provide recognizable visual clues, and since clipboard buffers may be disrupted by downstream copy-paste operations. As many users habitually perform save operations, we were not surprised to find no significant difference between saving operations in the two phases.

These findings support H1 in that users appear to perform state-stabilizing actions before switching to the application that generated the alert, presumably to leave the primary task in a state that allows for more efficient resumption.

Task Focus (H2)

We explored task focus in terms of task switch rates for each phase. Our intuition was that higher switch rates might indicate lower levels of focus on any one task—behavior that may be associated with higher likelihoods of browsing among several applications or rapidly switching across open applications in search of a previously suspended task.

Significant differences were found for task switches per minute during different phases of the interruption lifecycle. Table 1 shows the switch rates for email and IM alerts. Pairwise *t*-tests showed that the switch rate (on a per minute basis) during the diversion phase was significantly higher than the switch rate in the pre-interruption phase ($t(18)=6.787$, $p<0.001$), and that the switch rate during the resumption phase was higher than the switch rates in the pre-interruption ($p<0.0001$) and diversion ($p<0.0004$) phases. Users did not switch tasks between alerts and response during the preparation phase for the immediate response situation, reported as 0 for these cases in Table 1.

Job roles had a marginal effect ($F(2,15)=3.442$, $p<0.059$) on the mean rate of application switches during the prepare-to-respond period for email alerts. Researchers switched at a higher rate than developers ($p<0.06$). No other effects of job roles were found.

These findings support H2 in that users appear to be less focused on tasks during the diversion and resumption

phases with rapid interactions with peripheral tasks seen in both. For the resumption phase, this may indicate user intent to regain suspended task context as soon as possible. Lower rates of switching of windows are seen when users are working on primary tasks.

Interactions with communication applications (H3)

During the diversion phase, users largely interacted with communication applications (Outlook and IM clients). User traversed these applications with a mean rate of 0.77/minute (S.D.=1.56), a significantly higher rate ($t(19)=4.283$, $p<0.001$) than that observed during the baseline pre-interruption phase ($M=0.475$, S.D.=0.21). Users spent on average 66 seconds (S.D. 165.9), also significantly higher ($t(10)=3.674$, $p<0.002$) than dwells in the pre-interruption phase ($M=33s$, S.D.=7.5s).

Actions performed within Outlook during the diversion phase are summarized in Table 2. Note that rates of performing monitored operations in the diversion phase were significantly higher than in the pre-interruption period ($p<0.008$ for all operations). We believe that this indicates that alerts influence users to interact in a more rapid, less focused manner with Outlook than they might via default patterns of inspecting the state of their inboxes in the absence of alerts. The rate of web-mail visits ($M=1.76$, S.D.=1.91) during the diversion phase were also significantly higher ($t(7)=3.357$, $p<0.012$) than in the pre-interruption phase ($M=0.41$, S.D.=0.27), indicating that, once users switch to Outlook, they took the opportunity to check for other email too. These findings reinforce the hypothesis that breaks in activity associated with alerts from an application provide opportunities for performing communication operations with others, at significantly higher rates than usual.

Influences of visible cues on return (H4)

Overall, for email alerts, suspended application windows that were less than 25% visible because of obscuration by other windows took significantly longer to return to as compared to application windows that were more than 75% visible ($t(20)=3.131$, $p<0.005$). Similarly, for IM alerts, suspended application windows that were less than 25% visible took significantly longer to return to as compared to application windows that were more than 75% visible ($t(23)=2.503$, $p<0.02$). This finding suggests that the visibility of windows may serve as a reminder to users to break out of the diversion chain and return to suspended applications, thus, lending support to H4.

Actions	Pre-interruption Mean(S.D.)	Diversion Mean(S.D.)	p
Mail open	0.53 (0.5)	3.66(6.7)	0.001
Mail write	0.44(0.31)	3.15(5.71)	0.002
Mail send	0.35(0.27)	1.69(2.89)	0.008

Table 2. Actions per minute of email operations during pre-interruption and diversion phases.




Figure 2. Cumulative probabilities of suspended active windows being resumed given time until resumption across different response types.

Difficulty in resuming suspended applications (H5)

An immediate response to an email alert was associated with users taking, on average, 16 minutes and 33 seconds (S.D.=27m, 20s) in the resumption phase before returning to the state of the application in advance of the suspension. For delayed responses, the resumption phase spanned 15 minutes and 50 seconds (S.D.=25m, 5s) on average, not significantly different from the resumption time for immediate responses.

For suspensions following IM alerts responded to in an immediate fashion, users spent on average 10 minutes and 58 seconds (S.D.=14m, 16s) within the resumption phase. For delayed responses, the resumption phase spanned on average 12 minutes and 2 seconds (S.D.=14m, 58s), again, not significantly different from the resumption time for immediate responses.

A portion of the time spent in the resumption phase can be attributed to cycling through a set of suspended applications. However, even when the user finally returned to the suspended application, substantial time appeared to be devoted to resuming the task state and, presumably, the mental state that they had been in prior to suspension. We speculate that time and effort with resumption may involve reacquiring memories about the task and, more generally, refocusing cognitive resources that may have been usurped during the diversion phase. The results indicate that the diversion, starting off with a seemingly innocuous alert, can result in substantial lag in the resumption of primary tasks. This finding corroborates evidence found in research on interruption effects, as shown in [2, 33].

Influence of time on primary task on resumption (H6)

We pursued whether recency of activity and signs of focus of attention on tasks had influence on the efficiency of




Figure 3. Cumulative probabilities of resuming work within a suspended active window as a function of minutes into the resumption phase.

resuming suspended applications after the diversion phase had ended. To explore such relationships, we examined windows that users had been interacting with during the 15 second period prior to a suspension. On average, 27% of the alerts resulted in users being diverted from these prior active windows for more than 2 hours into the resumption phase. Figure 2 illustrates the breakdown across the different types of alerts and types of responses. Note that delayed email responses had the lowest probability of resulting in a 2 hour or lower resumption time for the suspended applications. This suggests that delayed email responses may be associated with users leaving a task in a more stable state and becoming more engaged in email responses and in subsequent peripheral tasks, hence deviating further away from the suspended primary tasks.

We were also interested in the influence of the duration of the focus of attention on tasks just prior to suspension on the time until resumption. Figure 3 illustrates the results of this exploration. Active windows where users spent between 5-30 minutes before the suspension were typically resumed within 5-15 minutes into the resumption phase. In contrast, if users spent less than 5 minutes on a task before suspension, they had a 10% probability of not resuming the task within 2 hours into the resumption phase.

USER INTERVIEWS

On conclusion of the study, we sought interviews with all participants. We conducted face-to-face interviews with 14 of the participants, based on availability and interest shown in the outcome of the study. The goal of the interview sessions was to both convey to participants their own work activity patterns as compared to other subjects and to query them about behavioral patterns we had observed in the analysis. This gave us a deeper understanding of the participants' behaviors.

In our interviews, we found that participants were generally aware of the frequency of alerts they received during workdays, though they did express surprise at the rate. Two subjects stated that they were seldom influenced by alerts to switch to email or IM and that they would usually decide based on their task context when they wanted to switch. However, these participants were also found to switch to email or IM immediately on a number of occasions.

In the interviews, users who responded to IM more quickly than to email cited two reasons for such quicker responses: one, they could quickly respond and switch back to what they were doing and two, they felt the social obligation of responding quickly as someone was waiting on the other end. They expressed their sense that responding to email would be more time consuming. Hence they would delay with switching to email. We assume that the interim time was spent making progress on the task and nurturing it into a state that they could resume later with ease.

Participants who responded to email more quickly than to IM typically used email extensively as part of their work routines and felt compelled to respond as soon as possible. However, during the responses, they often had to open other applications, resulting in spending more time on email responses than on IM. To them, IM was more of a social tool. When they did spend longer times in IM, it was mostly because they had time to spare, or, on some occasions, because they had forgotten about their suspended tasks.

Users mentioned that they believed that they were subconsciously aware of tasks left suspended. Deadlines and importance of ongoing tasks would often enhance their sensitivity of these tasks, serving as subtle reminders to break away from the chain of disruption. Whether or not participants indeed had the ability to maintain awareness of suspended tasks, our study suggested that visibility of the suspended application windows may have often served as a reminder to return to tasks. We found that people who used multiple displays would often leave their email client open on a secondary display and use an IM client on their primary display. Several participants mentioned that, while engaging in an IM session, the visibility of the suspended application windows on the primary display prompted them to return to the suspended application more quickly than when they were using an email application on the secondary display. Users also mentioned difficulties in restoring task context on return from the alert response, especially if there were multiple suspended applications. All users mentioned that they would habitually self interrupt themselves to access Outlook regardless of alerts, often to read new or previous mails or to send mails.

Participants mentioned that they typically did not feel the need for minimizing windows when they responded immediately to email or IM, because they intended to return to their primary task as soon as possible. When they did minimize windows, it was because they had too many applications open and wanted to increase screen space.

Most interviewees stated that they did not save documents categorically before switching since they were in the habit of saving periodically. They did feel, however, that an auto-save upon switch feature would be helpful.

LESSONS LEARNED

From analysis of the collected data and user interviews, we learned several lessons about how users switch tasks after receiving an alert, how they continue to interact with other applications during the switch, and how they return to resume activity on the suspended applications:

Users view alerts as an awareness mechanism rather than a trigger to switch tasks, but the alerts often cause them to do otherwise. Users prefer having control over when they switch to the application that is generating an alert so that their primary tasks are less disrupted. They do not disable alerts entirely as they would like to be aware of incoming information. Often, the identity of the sender of the alert and the content can influence how the user responds and may stimulate immediate task switches.

Immediate responses indicate alert-driven interruptions and delayed responses are indicative of self-initiated interruptions. Switching within 15 seconds of the arrival of the alert is considered to be a direct consequence of the alert whereas delayed responses (up to 2 hours) tend to indicate self interruptions. For the former, users appear to seldom perform actions on the application to be suspended that might facilitate easy resumption. For delayed responses, the lag in switching appears to often include time spent shaping the state of the suspended task into one that is more stable and more efficiently resumed, e.g., saving the document, completing a sentence, etc.

Users spend more time than they realize responding to alerts. Even though users feel that they are in control of when they switch tasks due to an alert, they appear to be largely unaware of the amount of time they end up spending on the alerting application, on other tasks they invoke as a result of responding to the alert, and on browsing through other peripheral applications before resuming the suspended task. Even when users respond immediately with the intention of resuming the suspended current task as soon as possible, they often end up taking significantly more time to return than the time to respond.

Importance of suspended task is associated with early recovery. Participants reported that deadlines and importance of ongoing tasks often enhance memories about their tasks and promote breaking away from the current chain of disruption to resume the suspended task as soon as possible. The observations about the influence of duration of recent focus of attention on a task on the time to return support this reflection.

Visibility of suspended application windows is associated with faster recovery. The visibility of cues associated with a suspended application affects how quickly a user returns to that application. Our analysis showed that windows largely occluded by application windows that users access during

the disruption chain took longer to recover. Similarly, as reported by the users, windows on a display where the user was not currently focused also had longer recovery times. We found that participants often keep windows of ongoing applications open and leave visual indicators (*i.e.*, set cursor locations, highlights) within the application window to help them resume quickly upon return.

DESIGN IMPLICATIONS

Based on our findings and feedback from users, we envision that designs for recovery tools might benefit by considering the following ideas:

Provide visual indicators of occluded application windows to assist with the recovery of suspended tasks.

As users were found to take significantly longer to resume applications whose windows were largely occluded, a recovery tool might provide value by generating some visual cues when application windows become occluded. A plethora of designs are feasible. For example, cues might be provided by progressively fading in occluded windows so as to have them peek through occlusions based on the time since last access and predetermined importance of the suspended application. In another design, a recovery system might maintain reminder icons of suspended applications or tasks in a visible area that can be scanned by the user.

Automatically save task context on suspension.

A major problem faced by users attempting to recover suspended applications was restoring the context of the suspended tasks, especially when there were multiple applications in the suspended task context. A useful design feature, would be to automatically save not only individual documents but also the broader context in which they were being used. For example, a user writing a research paper may have a Word document, an Excel spreadsheet, and a statistical software package open, and interactions on any of the three counts as her primary task. On suspension, a recovery tool might save this broader task context, as well as the state of all applications within the context, so that users could easily restore all applications belonging to this context and quickly resume where they left off.

Provide easy access to suspended task context.

While on a chain of diversion, users should be provided with easy access to the suspended application contexts. Multiple designs are feasible. As an example, access might be provided in the form of thumbnails with views of the suspended states of each application.

Provide playback of actions within task contexts.

Some users were found to use the undo key sequence on recovery and the interviews and survey questions revealed that they employed sequences of undo to recall the last actions they had taken before suspension. It may be valuable to provide a playback of the last n actions or of the actions that had occurred within the last t seconds, ensuring replay of both content and context [8].

The challenges, of course, are not only implementing the above mentioned features but also ensuring that the reminders and recovery cues do not cause further disruptions, especially if the current task context is more important than the suspended tasks. Determining appropriate timing strategies [5, 7, 20], display techniques [11, 30], and strategies for prioritizing tasks are paramount for the success of future disruption and recovery tools.

SUMMARY AND FUTURE WORK

We performed a field study of the computing activities of 27 users over a two-week period, exploring the suspension, recovery, and resumption of tasks in participants' natural work settings. We found that participants spent on average nearly 10 minutes on switches caused by alerts, and spent on average another 10 to 15 minutes (depending on the type of interruption) before returning to focused activity on the disrupted task. We discovered that, following an alert-based suspension, subjects would often visit several applications in addition to the notifying application. We found that 27% of task suspensions resulted in more than two hours of time until resumption. In interviews, users attributed long delays to the loss of context associated with the task switch. Findings about the association between greater visibility of windows of suspended applications and faster resumption of tasks suggest that visual cues may serve as reminders to return to suspended applications. After reviewing sets of results gleaned from monitoring users,, we presented design implications for reminder and recovery tools and discussed research directions.

In future work, we are interested in pursuing an understanding of the influence of face-to-face and phone-based interruptions on task disruption and recovery and to investigate differences in the disruption, recovery, and resumption of tasks for these interruptions versus computer-based alerts. Studies have shown that such social interruptions are commonplace in work environments [9, 13]. We are also interested in pursuing more deeply the influence of visual cues on task resumption, given the significance of the results on the effects of window visibility that we found in our field study. Also, we would like to better understand the influence of recency and focus of attention on task resumption, as well as other characteristics of tasks such as deadlines and priority.

We believe that continuing careful study of disruption and recovery of people in the course of daily life will reveal valuable insights about the challenges of multitasking over time, and that such insights will invariably shape the design of more productive and rewarding computing experiences.

ACKNOWLEDGMENTS

We thank Muru Subramani for providing assistance with Eve components and their modification. We are indebted to the subjects of our study for their participation and valuable feedback. This research was performed during the first author's summer internship at Microsoft Research.

REFERENCES

1. Altmann, E.M. and J.G. Trafton Memory for Goals: An Activation-based Model. *Cognitive Science*, 26 (1). 39-83.
2. Bailey, B.P. and J.A. Konstan On the Need for Attention Aware Systems: Measuring Effects of Interruption on Task Performance, Error Rate, and Affective State. *Journal of Computers in Human Behavior*, 22 (4). 709-732.
3. Bannon, L., A. Cypher, S. Greenspan and M.L. Monty, Evaluation and analysis of users' activity organization. *CHI 1983*, 54-57.
4. Card, S.K. and A.H.H. Jr., A multiple, virtual-workspace interface to support user task switching. *CHI 1987*, 53-59.
5. Cutrell, E., M. Czerwinski and E. Horvitz, Notification, Disruption and Memory: Effects of Messaging Interruptions on Memory and Performance. *INTERACT 2001*, 263-269.
6. Czerwinski, M., E. Cutrell and E. Horvitz, Instant Messaging and Interruption: Influence of Task Type on Performance. *OZCHI 2000*, 356-361.
7. Czerwinski, M., E. Cutrell and E. Horvitz, Instant Messaging: Effects of Relevance and Timing. *People and Computers XIV: HCI 2000*, 71-76.
8. Czerwinski, M. and E. Horvitz, An Investigation of Memory for Daily Computing Events. *HCI 2002*.
9. Czerwinski, M., E. Horvitz and S. Wilhite, A diary study of task switching and interruptions. *CHI 2004*, 175-182.
10. Dabbish, L. and R.E. Kraut, Controlling interruptions: awareness displays and social motivation for coordination. *CSCW 2004*, 182-191.
11. Dey, A.K. and G.D. Abowd, CybreMinder: A Context-Aware System for Supporting Reminders. in *Proceedings of 2nd International Symposium on Handheld and Ubiquitous Computing*, (2000), 172-186.
12. Dragunov, A.N., T.G. Dietterich, K. Johnsrude, M. McLaughlin, L. Li and J.L. Herlocker, TaskTracer: a desktop environment to support multi-tasking knowledge workers. *IUI 2005*, 75-82.
13. Fogarty, J., A.J. Ko, H.H. Aung, E. Golden, K.P. Tang and S.E. Hudson, Examining task engagement in sensor-based statistical models of human interruptibility. *CHI 2005*, 331-340.
14. Franke, J.L., J.J. Daniels and D.C. McFarlane, Recovering context after interruption. *CogSci 2002*, 310-315.
15. Gonzalez, V.M. and G. Mark, "Constant, Constant, Multi-tasking Crazy": Managing Multiple Working Spheres. *CHI 2004*, 113-120.
16. Horvitz, E. and J. Apacible, Learning and Reasoning about Interruption. *ICMI 2003*, 20-27.
17. Horvitz, E., J. Breese, D. Heckerman, D. Hovel and K. Rommelse, The Lumiere Project: Bayesian User Modeling for Inferring the Goals and Needs of Software Users. *UAI 1998*, 256-265.
18. Horvitz, E., A. Jacobs and D. Hovel, Attention-Sensitive Alerting. *UAI 1999*, 305-313.
19. Iqbal, S.T. and B.P. Bailey, Investigating the Effectiveness of Mental Workload as a Predictor of Opportune Moments for Interruption. *CHI 2005*, 1489-1492.
20. Iqbal, S.T. and B.P. Bailey, Leveraging Characteristics of Task Structure to Predict Costs of Interruption. *CHI 2006*, 741-750.
21. Latorella, K.A., Effects of modality on interrupted flight deck performance: Implications for data link. *HFES 1998*, 87-91.
22. Mark, G., V.M. Gonzalez and J. Harris. No task left behind?: examining the nature of fragmented work. *CHI 2005*, 321-330.
23. McFarlane, D.C., Coordinating the Interruption of People in Human-computer Interaction. *INTERACT 1999*, 295-303.
24. O'Conaill, B. and D. Frohlich, Timespace in the workplace: dealing with interruptions. *CHI 1995*, 262-263.
25. Ovsiankina, M. Die wiederaufnahme unterbrochener handlungen. *Psychologische Forschung*, 11. 302-379.
26. Scott, S.D., S. Mercer, M.L. Cummings and E. Wang, Assisting Interruption Recovery in Supervisory Control of Multiple UAVs. *HFES 2006*.
27. Smith, G., P. Baudisch, G. Robertson, M. Czerwinski, B. Meyers, D. Robbins and D. Andrews, GroupBar: The TaskBar Evolved. *OZCHI 2003*, 34-43.
28. Speier, C., J.S. Valacich and I. Vessey The influence of task interruption on individual decision making: An information overload perspective. *Decision Sciences*, 30 (2). 337-360.
29. Trafton, J.G., E.M. Altmann, D.P. Brock and F.E. Mintz Preparing to resume an interrupted task: effects of prospective goal encoding and retrospective rehearsal. *International Journal of Human-Computer Studies*, 58. 583-603.
30. Van Dantzich, M., D. Robbins, E. Horvitz and M. Czerwinski, Scope: Providing Awareness of Multiple Notifications at a Glance. *AVI 2002*.
31. Wickens, C.D. Multiple Resources and Performance Prediction. *Theoretical Issues in Ergonomic Science*, 3 (2). 159-177.
32. Zeigarnik, B. Das behalten erledigter und unerledigter handlungen. *Psychologische Forschung*, 9. 1-85.
33. Zijlstra, F.R.H., R.A. Roe, A.B. Leonora and I. Krediet Temporal Factors in Mental Work: Effects of Interrupted Activities. *Journal of Occupational and Organizational Psychology*, 72. 163-185.