Self-Interruption on the Computer: A Typology of Discretionary Task Interleaving

Jing Jin & Laura A. Dabbish
Human-Computer Interaction Institute, School of Computer Science
Carnegie Mellon University
5000 Forbes Ave, Pittsburgh, PA 15213, USA
jingidy@alumni.cmu.edu, dabbish@cmu.edu

ABSTRACT
The typical information worker is interrupted every 12 minutes, and half of the time they are interrupting themselves. However, most of the research on interruption in the area of human-computer interaction has focused on understanding and managing interruptions from external sources. Internal interruptions – user-initiated switches away from a task prior to its completion – are not well understood. In this paper we describe a qualitative study of self-interruption on the computer. Using a grounded theory approach, we identify seven categories of self-interruptions in computer-related activities. These categories are derived from direct observations of users, and describe the motivation, potential consequences, and benefits associated with each type of self-interruption observed. Our research extends the understanding of the self-interruption phenomenon, and informs the design of systems to support discretionary task interleaving on the computer.

Author Keywords
Interruption, attention, multi-tasking, work fragmentation, task switching, work spheres, self-interruption

ACM Classification Keywords
H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION
Popular press descriptions of today's information-centered workplace portray increasingly fragmented attention and multi-tasking "craziness." Recent studies on workplace multi-tasking behavior have substantiated these notions, showing that people are interrupted, and switch tasks, about every 12 minutes [16].

A substantial body of research in human-computer interaction has focused on interruption and attention management on the computer [1, 3, 4, 6, 8, 9, 13, 14, 16, 18]. This research has consistently documented the negative consequences of interruptions to ongoing work tasks including errors, work delay, stress, and context-switching costs [5, 6, 7, 9, 11, 18, 20]. System development efforts have tried to alleviate the negative side-effects of interruptions by filtering unwanted notifications and messages or tailoring interfaces to the users based on sensed context. However, the majority of this research has focused on understanding and managing interruptions from external sources, such as notifications [18] or communication interruptions [6].

However, a large portion of the task-switches that users experience are in fact internally generated. In their studies of multi-tasking in the workplace both [8] and [13, 16] found that almost 50% of observed task switches were initiated by the user. They also found that self-interrupted tasks were less likely to be resumed compared to externally interrupted tasks [16].

Self-interruption, also referred to as discretionary task interleaving [23], is a self-initiated switch away from a task prior to its completion. For example, a user may become bored while writing a report and self-interrupt this task by turning on a music player.

Self-interruption is an important type of task-switching behavior and part of the multi-tasking phenomenon. Despite its noted prevalence [8, 13, 16] and potential negative side-effects [16], almost no research has examined in detail what this type of internal task-switching looks like [23]. Our work addresses this gap, extending the research on self-interruptions. The specific research questions we address in this paper are as follows: Why do people interrupt themselves on the computer and switch to doing something else? What do internally generated interruptions look like in practice? What are their potential negative and positive side-effects?

In this paper, we contribute to the literature on multi-tasking and interruption by introducing a typology of self-interruption on the computer. This typology was developed by analyzing observations and retrospective interviews of people using a computer to complete their normal work tasks. Using a grounded theory approach [10, 12, 17, 30], seven types of self-interruptions were identified from the data, and these types were verified using a set of independent coders. The typology we present offers a first

CHI 2009, April 4–9, 2009, Boston, Massachusetts, USA.
Copyright 2009ACM 978-1-60558-246-7/09/04…$5.00

look into the defining characteristics of self-interruption and provides a foundation for future research and system development to support multiple activities on the computer.

RELATED WORK

Multi-tasking & Work Fragmentation

Patterns of work fragmentation due to frequent task-switching and high communication frequency have been observed in both managerial and non-managerial work settings [19, 24, 26, 29]. Human-computer interaction (HCI) researchers have long recognized the importance of designing systems to support the way users naturally interleave activities while working on the computer. Bannon et al [2] and Miyata and Norman [20] were among the first to decree that computer systems needed to support switching between tasks and resuming interrupted tasks.

More recent work has highlighted the importance of self-interruption and its prevalence in information work. Czerwinski, et al. conducted a diary-based study on task-switching [8] and found that 40% of the task switches reported in their study were self-initiated and did not involve proceeding to the next logical task.

In their study on multi-tasking, Gonzalez and Mark [13] refined the notion of a task-switch, introducing the concept of work spheres defined as: “a set of interrelated events, which share a common motive” (p. 117). The task-switches observed in their study indicated switches between work-spheres or tasks with a common motive. In their study of information workers inside a large corporation, they found people switched tasks or were interrupted every 12 minutes on average [13, 16]. Gonzalez and Mark’s [13, 16] observations confirm that the patterns of intense work fragmentation observed 20 years ago persist to this day.

Mark et al’s study also differentiated between self-interruptions and external interruptions [16], defining self-interruptions as “those in which one stops a task of their own volition.” (p. 322) They found that and that 50% of the interruptions people experienced in their study were internal, or people interrupting themselves. They also found that tasks that were internally interrupted in their study were less likely to be later resumed than externally interrupted tasks [16], suggesting that self-interruptions may be detrimental to ongoing in work tasks.

Self-interruptions represent a significant proportion of the task-switches a user experiences, with potentially negative consequences for productivity. The concept warrants more in-depth research. In particular we need a better understanding of why people self-interrupt, and the positive and negative effects of self-interruptions.

Consequences of Self-Interruption

We know from a history of research in psychology and human-computer interaction that external interruption to an ongoing work task can result in degraded performance on the interrupted task [1, 5, 6, 7, 11, 18, 20, 22] difficulty resuming the interrupted task [6, 14], and increased user-frustration [1, 15]. There are several reasons interruptions can have such a negative impact on the tasks they disrupt. The first is the cognitive overhead involved in a context switch between tasks [5, 6, 20]. When a user switches to a different task because of an interruption, they must load the context of the new task into working memory which can be cognitively demanding and time-consuming [5, 20]. The time away from the interrupted task will delay the completion of that task, and decay their working memory representation of information in the interrupted task [20].

When a user finally returns to the task they left behind, they must reload the context of the original task into their working memory, which means recalling their position in the task and relevant task related information [5, 6]. Delay returning to an interrupted work task, combined with the time for a user to recall their position in the task and relevant task context can result in a resumption lag [6, 20].

It is important to note that the majority of the previous work on the task consequences of interruptions has focused predominantly on external interruptions. The potentially unique task and well-being side-effects of self-interruptions are not well understood. It may be that the impact of a self-interruption may depend on the motivation for the task-switch or the form that the interruption takes. For example, Jett and George [15] speculate that if people interrupt themselves to take a break, because of boredom or fatigue, it can potentially lead to increases in productivity.

System development efforts have aimed to facilitate task switching and resumption of interrupted tasks by organizing files and information into task groupings [3, 4, 9, 25, 27]. These systems employ a fairly course-grained notion of task as a long-term project, often combined with a “workspace” metaphor devoting a unique virtual location to each project [4,9,27]. However, these systems have not been empirically validated [8] and for most users the difficulty in resuming externally interrupted tasks persists [14].

System designers may benefit from a more detailed understanding of the different reasons users switch between tasks and the unique consequences of those self-interruptions. Thus we wanted to begin to document self-interruptions and their task and well-being impacts. Our goal was to integrate the existing cognitive psychology and HCI research on interruption with a grounded understanding of self-interruption in context.

METHODOLOGY

Because so little is known about the phenomenon of self-interruption, we took a grounded theory approach [10, 12,, 30] with the goal of generating a descriptive theory of the causes and effects of self-interruption. Grounded theory is "an inductive, theory discovery methodology that allows the researcher to develop a theoretical account of the general features of a topic while simultaneously grounding the account in empirical observations or data" [17] (p. 141).This approach was chosen specifically because there is so little existing research on how self-interruption occurs in context. Following the grounded theory approach, we looked to describe the self-interruption process in terms of contextual conditions, actions, and consequences surrounding the phenomenon [30]. We recorded what self-interruption events looked like in context across a set of computer users.

Observations& Retrospective Interviews

We obtained data on both the observable behavior surrounding a user task-switch and the corresponding motivations users had for switching tasks. We chose to observe task-switching behavior as it naturally occurred by shadowing users as they completed their usual work tasks. This technique has been successfully employed in previous work on attention and multi-tasking, starting with Mintzberg's original study on managerial attention [13, 16, 19].We followed closely the shadowing methodology employed by Mark et al [16] to capture the task switch events of interest.

Participants were shadowed by an observer for approximately one hour while doing their normal work to gather real-life occurrences of self-interruptions. During the observation period, the participant was asked to conduct his or her work on the computer as usual, while the researcher observed and took notes on the tasks they performed. Consistent with the methodology in [13, 16, 19], the researcher noted any instances when the participant switched between different activities, computer applications, windows, or files within the same application. When a switch occurred, the researcher noted the time, preceding activity, and activity following the switch. At this point, only user behavior was recorded and no distinction was made between external or internal interruptions. An example of the data recorded during the shadowing observation is shown in Figure 1.

To capture the motivation behind task-switches and isolate internal task-switches in the data, we also conducted retrospective interviews with participants. These interviews immediately followed the shadowing portion of the study so the experiences would be fresh in participants' minds. Each participant was interviewed for 30 minutes to one hour. During the interview, participants were asked about the task switches the researcher had observed during the shadowing period. For each switch observed, the participant was asked about the pre-switch activity, post-switch activity (i.e. the interruption), and the reason for switching between the two at the moment they did. The interview data helped us later determine whether an observed task-switch was a self-

5:07pm	Back to e-mail go to orkut click on recent visitor's name
5:08pm	Read notepad* back to orkut
5:09pm	Read benefit summary

Figure 1. Example of the task switches recorded by the observer. The task marked with an asterisk is a possible interruption.

Pre-switch activity	Didn't know how much salary was, so was looking for friend who started working already
Post-switch activity	To-do list
Reason for switching	While waiting for page to load, was deleting to-do items

Figure 2. Participant responses in the interview about the possible interruption in Figure 1.

interruption or not. It also provided information about the task context surrounding participants' behavior, and data on the reason for switching between tasks. At the end of the interview, participants were debriefed and paid 20 dollars for their time. A sample of the data obtained from the interviews is shown in Figure 2.

Participants

Participants were recruited from the area surrounding a local university using online postings, and were selected to participate based on the nature of their work and their availability. Our sample included thirteen participants in total: six university students, three office workers, three software developers, and one writer. Six were male and seven were female. Participants ranged in age from 20 to 55, with the majority in their twenties. Observations and interviews were conducted in participants' offices or places they normally did their work (such as the library).

Identifying Self-Interruptions in Our Dataset

After collecting observational and interview data on task-switches, we followed the grounded theory approach of theoretical sampling [30] focusing on identifying and analyzing only the internal interruptions in our data [16].

Our purely observational dataset of task switches included every case where the user switched from using one application to using a different application. However, this type of application switch does not necessarily correspond to a task-switch or an interruption. The applications may be

Work Context	Task Flow	Preceding Task	Interrupting task	Reason for Switch	Duration of Interrupting Task(s)	Task Returned To
Find out about reimbursement	Go to Orkut ⇓ read notepad ⇓ back to Orkut	Go to Orkut	Read notepad	was deleting to-do items while waiting for page to load	1 min	Go to Orkut

Table 1. Information given to independent raters for the interruption in Figure 1.

needed for the same purpose or the switch may indicate completion of the first task.

By modifying the notion of task as defined in the work of Mark et al [16], we developed a more representative way to define user task boundaries. A task in their study was defined as a continuous use of the same device or application or a continuous conversation with the same person [13]. We noticed, however, that users can continue using the same application but switch to a different task. So a more precise definition of a task would also indicate the activities revolved around the same information. Thus, we defined a task-switch as a change in the nature of the user's activity (e.g. use of a different device or conversation with a different person) or a change in the information the activity revolves around. We defined self-interruptions in our data set as the task-switch instances where participants switched away from an ongoing task prior to completion but later returned to that task during the observation session.

We integrated the interview responses with the application switches observed during shadowing to determine the similarity in informational content between different activities, and ultimately whether a task switch was in fact an interruption. In addition, because we were interested only in self-interruptions, the observations combined with the interview data allowed us to focus on task-switches that were internally as opposed to externally initiated. We excluded from our final data set instances of external interruption observed, such as someone walking into a participant's office to talk, for example.

We created a data set containing each self-interruption observed, along with the participant's broader work context, the activities preceding and following the switch, the duration of the interrupting task, and the participant's stated motivation for switching between tasks. This data set was next analyzed using open and axial coding techniques to identify categories of self-interruption. An entry in this data set is shown in Table1 for illustrative purposes.

Data Analysis

We observed 13 participants for a total of 11 hours and isolated 39 self-interruptions activities on the computer. On average, participants interrupted themselves three times per hour, about once every 17 minutes. Despite our limited observation period, the self-interruption frequency in our dataset is consistent with previous work [13, 16].

We used *open coding* [30] to identify possible categories in our self-interruption data and their properties and dimensions. We sequentially examined each self-interruption instance collected and compared it with all of the preceding instances in the data, looking for similarities and differences. Self-interruptions with similar traits were grouped in the same cluster in an affinity diagram. During this process, when we encountered an example that did not fit well into existing clusters, we reconsidered the clustering and in some cases created new groupings of items that seemed more appropriate. We used this comparison and grouping process to organize the 39 self-interruptions, by similarity, into an affinity diagram.

The *open coding* process described above revealed seven distinct categories of self-interruptions which, in accordance with grounded theory development [30], were given the following descriptive labels: *adjustment, break, routine, wait, inquiry, trigger,* and *recollection* (Table 2). The defining differences between each of these categories were carefully considered to identify the properties and dimensions of each category. This process, called *axial coding* [30], identified the following key properties of the self-interruption categories: primary task status, motivation for the interruption, frame of mind, duration, environmental influence, and effect on the primary task (Table 3).

Primary task status denotes what was happening in the primary task when the user interrupted themselves. In our data, for example, we noticed that in some cases there was a stall or slow-down in the primary task, for instance when waiting for a program to respond or a web page to load.

Motivation for the switch refers to the reason users stated for switching from the primary task to the interrupting task during the retrospective interview. For example, one reason users gave for switching away from a particular work task was that they were bored with the task and wanted to do something different to refresh themselves.

Frame of mind refers to whether the switch between the primary and interrupting tasks required the user to be in a different work context. For example, if a user switched from typing a report to checking e-mail, the frame of mind would be different or unrelated. If a user switched from working on a molecular analysis program to reading about an algorithm used in the program, although the information they were dealing with had changed, the frame of mind would be closely related.

Duration of the interruption refers to the time that elapsed between when the user switched away from the primary task to when the user returned to the primary task.

We also observed that the interruptions could significantly delay the primary task. The primary task was only considered delayed in our analysis if the interruption took longer than expected. For example, one of our participants was waiting for a slow program to respond, and performed a series of different tasks for 13 minutes before checking back on the program. In this case, the interrupting tasks significantly delayed the primary task because the user returned to the program much later than he originally intended.

Type of Self-Interruption	Description	Positive Consequences	Negative Consequences
Adjustment	Improving or changing an aspect of the environment with the intention of increasing productivity on the primary task.	Directly or indirectly improves productivity (by improving the user's comfort) without a change in working context.	Attempt to improve the environment delays the primary task and can lead to abandonment of the primary task. Failure to improve the environment increases stress & frustration.
Break	Temporarily switching to a more desirable task because of frustration or fatigue with the primary task.	Alleviates stress, provides mental stimulation, and can increase enjoyment of a routine primary task. Can improve user's mood.	Delays the primary task, or can lead to procrastination if user fails to return to the primary task. Changes the user's context, so requires context switch to return to the primary task
Inquiry	Seeking information that will aid in completing the primary task.	Facilitates completion of the primary task when useful information is located.	Delays the primary task if the information seeking process is time-consuming, can result in abandonment of the primary task, or can make it difficult to refocus on primary task's context if information seeking is mentally intensive or results in another context switch.
Recollection	Remembering the need to perform a different, unrelated task while performing the primary task.	Prevents future neglect of the recalled task and creates a sense of accomplishment.	Likely to cause other self-interruptions. Hard to refocus on the context of the primary task
Routine	Performing a task as a habit of time and sequence.	Streamlines work flow based on prior experience	Decreases efficiency if primary task is mentally intensive or if the routine is time-consuming
Trigger	Performing a new related task as a result of a stimulus encountered in the primary task.	Can generate novel ideas	Likely to cause other self-interruptions. Hard to refocus on the primary task's context if Trigger is time-consuming
Wait	Performing a different task to fill time because of a temporary roadblock preventing the continuation of the primary task.	Maximizes productivity by multi-tasking, and can serve as an implicit break.	Delays resumption of primary task if the user overestimates waiting time. Likely to cause other self-interruptions.

Table 2. Description, positive, and negative consequences of each of the seven types of self-interruptions we identified.

RESULTS: SELF-INTERRUPTION TYPES

Based on our data, the seven types of self-interruption we defined were: adjustments, breaks, routines, waits, inquiry, triggers, and recollections (Table 2). We discuss the unique properties and observed process associated with each self-interruption type (Table 3) and present examples of each type from our data set. Consistent with the grounded theory approach, to offset our small sample size we relate the typology we have developed to existing theory on attention, memory, and task activity [12, 30]. We have undertaken a literature comparison process, because "tying the emergent theory to the extant literature enhances the internal validity and generalisability…because the findings often rest on a very limited number of cases" ([10], p.545). We also suggest some very preliminary design ideas based on our observations.

Adjustment

Adjustment self-interruptions occurred when participants acted to change the environment, either physical or virtual, to improve their productivity on the primary task. In our observations, one user directly improved his productivity by closing a program that was slowing down the computer. Another participant stopped his work on the computer to turn off a different computer that was making the room too hot. This adjustment improved his comfort, thus indirectly improved his productivity by allowing him to concentrate more fully on the primary task.

Impact

Similar to the two examples above, all of the observed adjustments seemed to improve productivity, either by allowing more efficient performance on the primary task or by improving the participant's comfort and mood. Adjustments also tended to be short in duration (Table 3) and required little cognitive processing. Because of this, users easily returned to their primary tasks. It is also possible for adjustments to have a negative impact on the primary task. If the user does not immediately succeed in improving his or her environment, he or she could become frustrated and distracted by the undesirable aspect of the environment, or sidetracked from their primary task for an extended period of time. In this case, adjustments could potentially decrease performance on the primary task because of decreased concentration and negative affect.

Design Implications

Adjustments can be caused by both the physical and the digital environment. In general, providing adaptable physical and digital workspaces that can be easily customized should minimize disruption associated with adjustments and result in more comfortable and productive environments.

Interface designs can also intelligently adapt the digital environment to avoid specific types of adjustments. For example, we observed one user self-interrupt to remove window clutter by closing programs they weren't using. To avoid this type of adjustment, windows can be dimmed with respect to how recently they were viewed. Recently accessed windows, which are more likely to be part of the active work context, would then be the brightest and most noticeable. This kind of change might decrease the chance that programs unrelated to the primary task would cause dissatisfaction to the user.

Break

Breaks occurred when the user decided to perform a more desirable task instead of continuing to work on the primary task. For example, in one break we observed the participant decided to surf the web because he found his primary task, making a tutorial for Excel, to be extremely boring. Breaks occurred because the primary task became undesirable for different reasons: it was too physically or mentally demanding causing the user fatigue, too mundane or routine causing boredom, or too frustrating causing negative affect. The breaks we observed varied greatly in duration and were not influenced by the environment.

Impact

Break interruptions have the potential to refresh the user's mind and improve concentration when the user returns to the primary task. However, since breaks were typically unrelated to the primary task, users had to spend time recalling the primary task context when they returned. Because of the motivation associated with breaks, they led to procrastination, where users intentionally or unintentionally delay the return to the primary task.

Design Implications

Designs should account for the context switching costs and possibility of procrastination associated with breaks, while allowing for and potentially facilitating benign breaks that can improve the user's concentration. Past research has identified potential ways to preserve context during task switches and developed systems that will retain the user's task state [3, 4, 9, 25, 27]. Adding something as simple as a replay of the last few actions the user performed prior to a break may help users recall their primary task context.

Procrastination occurred in our data because the user lost track of time or the user had a strong desire not to return to the primary task. Procrastination can perhaps be mitigated by providing tools that increase awareness of the time spent on a break, facilitate goal setting, and attach incentives to primary tasks. For example, a tool could allow the user to set a goal time of completion for a task, and the closer their actual finish time is to the goal, the more rewards the user could receive upon completion. Rewards can be as simple as a plant widget blooming on the desktop or as complex as to influence the planning of the user's future tasks.

Inquiry

Inquiries occurred when users actively sought out additional information that would aid them in completing the primary task. In our observations one participant, for example, asked a friend for some HR benefits information, because he could not continue the primary task without that information. Users may chose to look for information because it is vital to the continuation of the primary task or because he or she was curious. In another inquiry we observed, a participant decided to look up a product while instant messaging his friend simply because his friend mentioned the product and he was curious about it.

Impact

Throughout our observations, inquiries positively impacted the participants' primary tasks because the participants located the information they were seeking. However, we recognize this may not always be the case. If the information is not easily located, users may become frustrated or get sidetracked from their primary task. Inquiries that involve a complex or long process will also require a more substantial switch in work context, making it harder to resume the primary task [5, 6, 20].

Design Implications

Inquiries are helpful to the primary task if they are successful, thus systems could focus on reducing the incidence of unsuccessful inquiries or ameliorating their effects. For example, users may be less frustrated at unsuccessful inquiries if they are able to easily keep track of the information they need to locate so that it is easy to remember to perform the inquiry again at a later time.

Recollection

Recollections occurred when users suddenly remembered they needed to perform a different task while engaged in the primary task and interrupted the primary task to perform the remembered task. Recollections seem to be prospective memory events, recall about things that need to be done in the future. In our observations we found that recollections sometimes did not appear to have an explicit cue, but sometimes were cued by stimuli in the primary task or the environment. For example, one participant was checking his e-mail and self-interrupted to book a ticket for a trip he was planning in a few months. He was reminded about the booking task because the e-mail he was reading contained a flight itinerary for a different trip.

Another participant experienced a recollection because she heard a song playing in another room and remembered that she wanted to e-mail her friend to ask for some mp3s (including that song). In this case, the stimulus that caused her to remember the other task came from the environment and was unrelated to the primary task.

The research on prospective memory has predominantly focused on prospective memory failure, when people forget to do things in the future. But in the cases we observed prospective memory events were recalled successfully. Perhaps further investigation into recollections in context,

Type of Self-Interruption	Number Observed	Defining Characteristics			
		Motivation	**Frame of Mind**	**Duration**	**Environmental Influence**
Adjustment	5	To make environment more comfortable to work in	Unrelated to primary task	Mean=36s SD=13s	Yes, gradually noticed
Break	8	To rest and refresh mind	Unrelated	Mean=2m30s SD=2m33s	No
Inquiry	4	To find information helpful to primary task	Related	Mean=1m45s SD=1m39s	No
Recollection	4	To prevent forgetting the remembered task	Unrelated	Mean=53s SD=45s	Varies, depending on source of the interruption
Routine	10	None	Unrelated	Mean=57s SD=36s	No
Trigger	3	Varies, depending on what the task switched to is	Related	Mean=3m50s SD=3m49s	Varies, depending on source of the interruption
Wait	5	To utilize the waiting time productively	Varies	Mean=1m SD=37s	No

Table 3. Observed trends for the seven types of Self-Interruption with respect to the defining characteristics used in the

when and how prospective memory events successfully occur, can extend our understanding of human memory.

Impact

Recollections did not have a significant impact on the primary task in the observation sessions. Participants always performed the remembered task and returned to the primary task soon afterward. It is possible, however, for recollections to cause other self-interruptions related to the remembered task. And since recollections usually involved tasks that were different in context than the primary task, users needed additional time to recall the primary task context upon completion of the remembered task.

Design Implications

A potential negative impact of a recollection is difficulty resuming the primary task. Therefore, systems such as [3, 4, 9, 25, 27] intended to help users retain and recall the primary task context can reduce the negative impacts of recollections. Another way to design for recollections is to allow the user to easily keep track of remembered tasks, without needing to switch contexts and perform the task immediately. For example, a lightweight and persistent to-do list application may be useful for this purpose.

Routine

Routines occurred when the user interrupted themselves out of habit. For example, in one self-interruption instance in our data the participant switched tasks to open up a selection of websites. They indicated always opening these websites at the beginning of a work session. For this participant, the routine we observed was relative to the event of starting work. We found that routines could occur relative to an event, such as in the example above, or occur on a regular time interval. Routines were the most frequently observed type of self-interruption in our dataset, since they tended to occur periodically. Routines also seemed to occur without explicit motivation; participants performed routines reflexively rather than because of a desire or need to do so.

Impact

In another routine we observed, a participant refreshed her favorite social networking site every few minutes. She explained that she is "addicted" to checking the site every fifteen minutes. She indicated that even when she has a tight deadline, she still refreshes the site, just less frequently. In this case the routine had a negative effect on the user's work, since the user was interrupted from the primary task on a fixed interval, making concentration difficult [7]. In general, though, routines did not appear to have a major impact on the primary task because they were integrated into the user's daily work and in many cases were performed seemingly without conscious thought or cognitive overhead.

Design Implications

Since routines in general did not appear to impact the primary task, there may not be a significant need to design technology that prevents routine interruptions. However, in the previous example, the routine of checking the social network every fifteen minutes did not yield any value to the user's primary task or general work context. In this case, a goal-based system could be designed to reward users for not performing a routine or performing it less frequently. It may be valuable to investigate in future work the utility of such a system for removing undesirable habits in general.

Trigger

Triggers occurred when the user initiated a new task that was cued by a stimulus the primary task. For example, one participant in our sample was typing an article and suddenly had an idea about a new article they could write, so they interrupted their primary task to write down that idea.

Impact

As illustrated in this example, triggers can be helpful in general because they can facilitate idea generation. Compared to the other types of self-interruption, triggers occurred least often in our sample but were the longest in duration. So even though triggers didn't occur as frequently as other types of interruptions, they consume more time in total because of each trigger's duration (Table 3). If the trigger interruption was not a beneficial task for the user, it negatively impacted the primary task much more than other types of interruptions because of the extended delay and subsequent increased primary task resumption costs.

Design Implications

A trigger's effect on the primary task is solely dependent on the nature of the trigger task, so it may be improper to design a single solution to problems that may be caused by triggers. Triggers that take the user away from the primary working context may be ameliorated by methods that preserve and aid the recall of work contexts. Designers can also encourage performing the trigger after the completion of the primary task by facilitating, for example, easy recording of the triggered task and the context in which the trigger formed.

Wait

Waits occurred when workload decreased in the primary task and the user needed to wait for a specified or unspecified length of time in order to continue engagement in the primary task. Users subsequently decided to utilize that time on a different task. For example, one of our participants was working with complex computational models, and the software program they were working in was extremely slow and stopped responding during the observation. The participant decided to do research online while waiting for the program to finish processing.

Impact

By multi-tasking, the participant in the example above was maximizing his utilization of time and increasing overall productivity. This is similar to the notion of time-sharing in perceptual psychology, where skilled users can effectively optimize the attention resource by switching between tasks at opportune moments [20]. However, we noticed in our observation that inaccurate predictions of the wait time can delay returning to the primary task which can be harmful depending on the primary task urgency. For the participant in the example above, his primary task was delayed by inaccurate predictions of the wait time, but it did not need to be completed until several days later; therefore the delay was not a significant cost for the participant. However, if the primary task had a more proximate deadline, the participant may have preferred to return to the primary task more quickly.

Design Implications

Most negative impacts of wait interruptions stemmed from misestimating of the amount of time the user needed to wait to continue the primary task, or forgetting to return to the primary task. If the wait is caused by a program's slow processing, a system or tool that estimates and displays the amount of time needed to process certain information may help users make more accurate estimations of the wait time. In addition, if the system keeps the user updated on the status of the program's processing, without being intrusive, the user might be able to constantly update their wait estimation and return to the primary task more promptly.

Inter-rater Reliability of Self-Interruption Types

To verify the reliability of the self-interruption types we identified, four independent raters were asked to label the 39 observed self-interruption instances using the seven categories from our affinity diagrams described by their defining properties (Table 2). For each self-interruption instance in our data, the raters were provided with the information shown in Table 1: work context of the primary task, the task flow surrounding the switch, the task directly preceding the interruption, the interrupting task, the reason the participant provided for the task switch, the approximate duration of the interruption, and the task engaged in directly following the interruption. Evaluators were asked to decide which category each self-interruption instance belonged to and provide a brief justification for their categorization.

Inter-rater reliability between the four independent evaluators was calculated using Intraclass Correlation (ICC), which indicates the proportion of variance in the ratings due to between-rater variability. Higher ICC indicates little variance between the scores given to each item by the raters [28]. A two-way mixed effects model on the categorization the evaluators assigned to each self-interruption showed that a substantial level of agreement (80%) was achieved by the raters ($p<0.001$).

DISCUSSION

We considered similarities and differences across the self-interruption types to identify higher order concepts [10, 30]. Several overarching themes emerged from this comparison.

Motivation

As illustrated in Figure 3, breaks, recollections, and routines were primarily initiated by the user's internal cognitive processes, whereas adjustments, triggers, and waits, and inquiries were motivated by the situation. For example, several users in our sample initiated a break because of their internal feelings of boredom, whereas waits were initiated because of the primary task status.

Another theme that differentiated interruptions was whether the initiating motive is bottom-up or top-down in nature. In perceptual psychology, researchers have examined how internal and external factors influence focus of attention. This research has delineated two broad motivations for shifts in focus of attention: bottom-up - or stimulus based direction of attention- versus top-down- or goal-oriented direction of attention [22]. Breaks, adjustments, and triggers are all bottom-up, or stimulus based changes in focus of attention, whereas recollections, routines, and inquiries are top-down, or goal-directed switches. Wait falls in between bottom-up and top-down because it is both stimulus-based, such as waiting for a program to respond, or goal-directed, such as utilizing the time to perform a specific task.

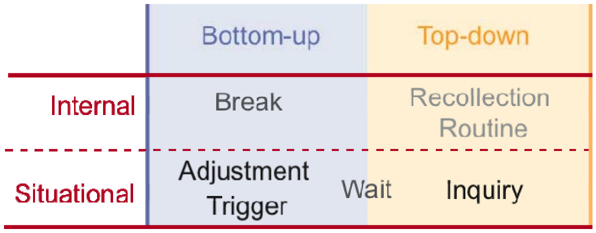

Figure 3. Types of self-interruption across dimensions. New tasks are in dark text, returned-to tasks are in light gray.

A third dimension differentiating the self-interruption types is whether or not the interruption is creating a new task. For adjustments, triggers, and inquiries, a new task is created by the interruption. For recollections and routines, the user is executing a preexisting or outstanding task. For example, one participant initiated a recollection when he remembered that he needed to check his earnings account. The need and method of checking his earnings account was already determined before the interruption occurred. Types such as waits and breaks can both generate a new task or involve a preexisting task. We can support interruptions that create new tasks by allowing users to easily record information about the need to perform the new task, giving them option to defer those tasks instead of interrupting the primary task.

Impact

Self-interruptions in our sample had several over-arching effects on the primary tasks. Inquiries, breaks, and adjustments facilitated the primary task by providing valuable information or creating an environment that encourages increased productivity. But, self-interruptions can also harm the primary task because of context-switching costs and delay. Past work has already started to explore how to facilitate switches between work contexts [3, 4, 9, 25, 27]. Finally, the user may forget to return to the primary task following a self-interruption. In our study, it was not possible to isolate interruptions that caused the user to forget to return to the primary task because of our limited observation period duration. One way systems might prevent forgetting is to provide an alert that reminds the user about an abandoned task in an unobtrusive way.

Since self-interruptions, and interruptions in general, can both positively and negatively affect the primary task, we cannot design systems to simply eliminate interruptions. Instead, we must address the unique characteristics of different types of interruptions and support the process of returning to the primary task following an interruption.

Limitations & Future Work

By observing participants directly we were able to gather rich data about all of their activities during the session. However, observer presence may have influenced participant behavior. In our case, participants could have been more focused on their work, and biased towards performing more "productive" tasks than they normally would when working alone. In addition, since observations are time-consuming, we were only able to study a limited number of participants, and there is a possibility our participant's were not representative. Our limited observation window and limited number of participants meant our sample of observed interruptions was also small, meaning there may be rare or less frequent types of self-interruption we did not observe. The small sample size also raises the question of whether our theory is saturated in the grounded theory sense; therefore additional work should test if our categories remain stable in a larger sample and in different contexts.

In the future, we would like to increase our sample size and conduct user studies with a larger pool of participants who vary both in terms of age and profession. Also, in order to collect a larger sample of self-interruption data uninfluenced by the presence of a researcher, we would like to collect behavioral data in a naturalistic way, such as using activity logging on the participant's computers. One challenge is that the participant's explanation of the task and the work in general during the interview was very important in identifying and classifying self-interruptions. Therefore, even though activity logging can occur anytime, vital information about the task must be obtained from the participants when it is still fresh in their memory. With a larger amount of unbiased data, it will be possible to examine in greater detail the consequences and benefits of each type of self-interruption. Finally, we would like to apply the typology to automatically detect different types of interruptions based on observed user behavior.

CONCLUSION

We have introduced a typology of self-interruption on the computer grounded in qualitative data. Prior to this study, to our knowledge, no one has examined this topic in-depth. Our work greatly extends the understanding of self-interruption on the computer. This increased understanding of the phenomenon and its consequences contributes to the literature on attention and interruption, and should inform the design of interactive systems intended to support multiple activities on the computer. We have only begun to scrape the surface on this important and complex phenomenon, and we hope that this study will inspire future research on the topic of self-interruption.

ACKNOWLEDGEMENTS

Our thanks to Mark Stehlik for his support, and to our study participants for their time and cooperation. This research was funded by NSF grant IIS-0325351, the CMU School of Computer Science, and a CMU SURG grant. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

REFERENCES

1. Adamczyk, P.D., & Bailey, B.P. If not now, when?: the effects of interruption at different moments within task execution, In *Proc. CHI 2004*, ACM Press (2004), 271-278.
2. Bannon L., Cypher A., Greenspan S. & Monty M. L. Evaluation and analysis of users' activity organization. In *Proc.CHI 1983*, ACM Press (1983), 54-57.
3. Bardram, J., Bunde-Pedersen, J., & Soegaard, M. Support for activity-based computing in a personal computing operating system. In *Proc. CHI 2006*, ACM Press (2006), 211-220.
4. Card S. K., & Henderson A. A multiple, virtual-workspace interface to support user task switching, In *Proc. CHI 1987*, ACM Press (1987), 53-59.
5. Cellier J.M., & Eyrolle H. Interference between switched tasks, *Ergonomics, 35*, 1 (1992), 25-36.
6. Cutrell, E.B., Czerwinski, M., & Horvitz, E. Effects of instant messaging interruptions on computing tasks, In *Proc. CHI 2000,* ACM Press (2000), 99-100.
7. Czikszentmihalyi, M. *Flow: The Psychology of Optimal Experience*. Harper Collins, NY, NY, 1990.
8. Czerwinski, M., Horvitz, E., & Wilhite, S. A diary study of task switching and interruptions. In *Proc. CHI 2004*, ACM Press (2004), 175-182.
9. Dragunov A.N., Dietterich T.G., Johnsrude K., McLaughlin M., Li L. & Herlocker J.L. TaskTracer: A desktop environment to support multi-tasking knowledge workers, In *Proc.IUI 2005*, ACM Press (2005), 75-82.
10. Eisenhardt, K.M. Building theories from case study research, *Academy of Management Review, 14,* 4 (1989), 532-550.
11. Finstad K., Bink M., McDaniel M., & Einstein G. O. Breaks and task switches in prospective memory, *Applied Cognitive Psychology, 20*, 5 (2006), 705-712.
12. Glaser, B.G., & Strauss, A.L. *The Discovery of Grounded Theory: Strategies for Qualitative Research,* Aldine Publishing Co, Chicago, IL, US, 1967.
13. Gonzalez, V., & Marks, G. “Constant, constant, multitasking craziness”: Managing multiple working spheres. In *Proc. CHI 2004*, ACM Press (2004), 113-120.
14. Iqbal, S., & Horvitz, E. Disruption and recovery of computing tasks: field study, analysis, and directions. In *Proc. CHI 2007*, ACM Press (2007), 677-686.
15. Jett Q. R., & George J. M. Work interrupted: A closer look at the role of interruptions in organizational life, *Academy of Management Review, 28,* 3 (2003), 494-507.
16. Mark, G., Gonzalez, V., & Harris, J. No task left behind? Examining the nature of fragmented work. *In Proc. CHI 2005*, ACM Press (2005), 321-330.
17. Martin, P.Y., & B.A. Turner. Grounded theory and organizational research," *The Journal of Applied Behavioral Science,22*, 2 (1986) 141-157.
18. McFarlane, D. Comparison of four primary methods for coordinating the interruptions of people in human-computer interaction. *Human-Computer Interaction, 17,* 1(2002), 1-61.
19. Mintzberg, H. *The Nature of Managerial Work.* Prentice Hall, Englewood Cliffs, N.J, USA, 1973.
20. Miyata Y., & Norman D. Psychological issues in support of multiple activities, In: *User Centered Systems Design: New Perspectives on Human-Computer Interaction*, Hillsdale: Lawrence Erlbaum Associates, (1986) 265-284.
21. O'Conaill B., & Frohlich D. Timespace in the workplace: Dealing with interruptions, In *Ext. Abstracts CHI 1995*, ACM Press (1995), 262-263.
22. Paschler, H., Johnston, J.C., & Ruthruff, E. Attention and Performance, *Annual Review of Psychology, 52* (2001)*,* 629-651.
23. Payne S. J., Duggan G. B., & Neth H. Discretionary task interleaving: Heuristics for time allocation in cognitive foraging, *Journal of Experimental Psychology: General, 136*, 3 (2007), 370-388.
24. Perlow L. The time famine: Toward a sociology of work time, *Administrative Science Quarterly, 44,* 1 (1999), 57-81.
25. Rattenbury, T., & Canny, J.CAAD: An automatic task support system. *In Proc. CHI 2007,* ACM Press (2007), 687-696.
26. Reder, S., & Schwab, R.G. The temporal structure of cooperative activity. In *Proc. CSCW 1990*, ACM Press (1990), 303-316.
27. Robertson, G., van Dantzich, M., Robbins, D., Czerwinski, M., Hinckley, K., Risden, K., Thiel, D., & Gorokhovsky, V. The Task Gallery: A 3D window manager. *In Proc. CHI 2000*, ACM Press (2000), 494-501.
28. Shrout P. E., & Fleiss J. L. Intraclass Correlations: Uses in Assessing Rater Reliability, *Psychological Bulletin, 86*, 2 (1979), 420-428.
29. Sproull L. S. The nature of managerial attention, *Advances in Information Processing in Organizations*, 1 (1984), 9-27.
30. Strauss, A., & Corbin, J. *Basics of Qualitative Research: Techniques and Processes for Developing Grounded Theory*. Sage Publications, Thousand Oaks, CA, USA, 1998.