The Effects of Time Constraints on User Behavior for Deferrable Interruptions

Peter Bogunovich

Department of Computer Science
Drexel University
pjb38@cs.drexel.edu

Dario Salvucci

Department of Computer Science
Drexel University
salvucci@cs.drexel.edu

ABSTRACT

Previous studies of multitasking have highlighted the importance of cognitive load in interruptibility by showing that forced interruptions are least disruptive when cognitive load is low, and also that users prefer to address interruptions at low-load points when given a choice. We present an empirical study that uses a ringing-phone scenario to examine how users manage deferrable interruptions in the presence of varying time constraints. We found that while cognitive load did influence multitasking as expected, the time constraints placed on the user also had a significant impact. In particular, we observed three distinct strategies for addressing interruption: the expected strategy of switching at low-load points, but also two other strategies of continuing on after a low-load point or giving up at a high-load point. The presence of the latter two strategies strongly suggests that users can adapt their multitasking behavior with respect to the time constraints of the interrupting task.

Author Keywords

Multitasking, interruption, attention, workload

ACM Classification Keywords

H.5.2 User interfaces: theory and methods; H.1.2 User / machine systems: human factors, information processing

General Terms

experimentation, human factors

INTRODUCTION

The interleaving of several tasks, commonly known as multitasking, is an important human ability: it is difficult to imagine the execution of some tasks, like driving a car to a destination, without some ability to manage and interleave lower level subtasks. People are generally inclined to perform some degree of multitasking, and it has been observed that computer users tend to switch tasks every few minutes [6]. Multitasking and interruptions have been researched extensively, especially with respect to the impact of *forced*

interruptions on task performance. In a forced interruption scenario, a user performing a primary task is interrupted with a secondary task that must be addressed immediately. The impact of interruption is typically evaluated according to a metric such as time on task [3], user error rate [9], frustration with the system [2], or resumption lag in returning to the primary task after the interruption [10]. A consistent finding of these studies is that interruptions occurring at times of higher cognitive workload or memory requirements in the primary task are more disruptive than those occurring when workload is reduced [5, 4, 7, 1].

Recent research has also looked at *deferrable multitasking* scenarios, in which an alert of a pending secondary task is issued, but the user is free (to some degree) to choose when to deal with the interruption. This work has focused mainly identifying patterns in user preferences. One observed pattern is the tendency for users to achieve a “stable state” in the primary task before switching [8]. A similar observation was made in [11] where the authors formalized memory requirements through *problem state*, or information that is temporarily relevant to subtasks [12], and showed that in the deferrable interruption scenario users prefer to “mono-task” with the primary task until a drop in problem-state requirements is reached. Deferrable multitasking scenarios are commonly found (e.g., email, chat messages, phone calls, etc.) and these results provide important insights into multitasking behavior in these contexts.

In this paper we expand on these results by examining the influence of factors beyond cognitive load, namely the influence of time constraints. Many real-life scenarios involve time constraints for dealing with interruption, such as a ringing phone (which is missed if unanswered) or a chat message (whose sender may go offline if a response is not promptly received). We present an experiment in which people experienced time-constrained deferrable interruptions using a simulated ringing-phone task. The experiment manipulated the time constraints (by using either two or six rings), and the “problem-state” memory requirements at different points in the task sequence. We examined the strategies that were employed in this scenario and found that while cognitive load remains a major factor in task switching, users exhibited strategies that are also clearly influenced by time constraints.

EXPERIMENT

We modified the code used in [11] for the present experiment. In the primary task, users answered emails about con-

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, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2011, May 7–12, 2011, Vancouver, BC, Canada.

Copyright 2011 ACM 978-1-4503-0267-8/11/05...\$10.00.

Figure 1. The multitasking simulation. On the left the user is alerted of an incoming phone call during the mail task. On the right the call has been answered and a new chat window grabs focus over the mail task until the call is completed.

sumer products and prices. The user first selected and read an email from an email browser window. Each email requested the price of a particular product. All product names comprised a real manufacturer name and a fictitious model consisting of one letter and two digits (e.g. “Canon H-44” or “Sony M-76”). The user then proceeded in a separate product browser window by first selecting the manufacturer from a list of manufacturers and then selecting the model from a list of models, at which point the price in question was revealed. The user then composed a response email where the price was relayed to the customer. Finally, the email was archived and a new email was received in the browser. Throughout the experiment, the mail-task windows were arranged in a fashion so that only information displayed on the focused window could be viewed.

The secondary task took the form of a simulated phone call and brief conversation. An incoming call alert was issued to the user in the form of audible ringing and the appearance of an “Answer Phone” button. While the phone rang, the user could continue working on the primary task as desired. To answer a call, the user clicked on the answer button, at which point a chat window appeared and grabbed focus from the primary mail task. The user was then required to read and respond to two questions from a simulated interlocutor. The first question asked the user whether or not she had seen a movie (with movies randomly selected from the Academy Award Best Picture nominees of the last 10 years). The second question was a follow-up about the user’s opinion of the same movie. The users were told to respond to each question with a “yes” or “no” response. Once the second question was answered, the chat window disappeared, and only then could the user continue with the primary mail task. A view of the tasks can be seen in Figure 1: the left panel shows the interface after the user is alerted to an incoming call, and the right panel shows the interface after the call has been answered and the conversation has begun.

Procedure

A single mail trial comprised the user reading an email, looking up a price, responding to the email, and archiving the email. One incoming phone call alert was issued during each trial, triggered by a randomized, uniformly chosen event. A delay between 50-200 ms was used to avoid tying an event to the call, and to avoid any subsequent events before the alert.

The task simulation was coded in Java, and throughout the experiment, user mouse action and keystroke events were logged. Copy and paste and other keyboard shortcuts were disabled. Along with the forced window arrangement, this scheme allowed the system to log switches between tasks and forced users to remember the problem-state information within certain subsequences of events (discussed shortly).

Each user performed two sessions: (1) a *short-ring* session (2 rings ≈ 4.5 s), and (2) a *long-ring* session (6 rings ≈ 14 s). Users were told the number of rings and approximate time constraints before each session, and to answer the calls whenever they wished. There were no penalties for missed calls. Each session ended after 12 calls were successfully answered. The order in which the sessions were administered was randomized. A total of 20 users (6 female and 14 male) participated in the study.

Primary-Task Structure

As in [11], the mail task was designed to include two subsequences of events in which users needed to maintain problem-state information in memory, specifically the name of the product, and its price. The most common mail-trial sequence and these subsequences are illustrated by Figure 2. The user first reads the email, including the name of the product. This information must be remembered and used while the user looks up the price in the product browser. Once the model link is selected, the name of the product is no longer needed, but the price is revealed and must be remembered as the response is composed. After the email is sent no problem state is maintained for the remainder of the trial.

Given the steps of a mail trial, it is possible to label each event according to whether or not problem state going into the event must be maintained and carried-over to the next event. This labeling is illustrated in Figure 2 where gray shaded events require information carry-over while unshaded events do not. Based on the results of [11] that show a tendency of users to favor task switching at points of lower cognitive load with no information carry-over, we refer to the dark shaded events as “*bad breakpoints*” and the unshaded events as “*good breakpoints*”.

RESULTS

As an initial validation of our results, we performed the same analysis on the task-switching data that was performed in

Figure 2. A typical mail-task strategy. Events shaded in gray can be considered bad breakpoints because they require information carry-over from the previous to the next state, while events shaded in white can be considered good breakpoints because they involve no information carry-over. The distance to the next good breakpoint is also shown.

Figure 3. Proportion of task switches at good and bad breakpoints, as a function of multitasking session and distance score.

[11], specifically examining whether task switches tended to occur at points of lower workload with no information carry-over. The previous study [11] reported that 94% of all task switches occurred at these “good” breakpoints of lower workload. In the current study, we found that under the looser time constraints of the long-ring condition, users followed an almost identical pattern, switching tasks 93% of the time at points of lower workload. Even under the stricter time constraints of the short-ring condition, users switched 80% of the time at points of lower workload.

To more closely examine the interaction between time constraints and task switching, we assigned a *distance score* to each mail-task event that represented the “distance,” in terms of number of events, to the next good breakpoint. This scoring is illustrated in Figure 2: events considered good breakpoints have a distance of 0, bad breakpoints just before good ones have a distance of 1, and so on.

Figure 3 shows the proportion of switches at good (and bad) breakpoints as a function of distance score. We performed a repeated-measures ANOVA on these data with session (long- or short-ring) and distance (0, 1, or 2) as within-subject factors (omitting users with missing data cells), and the effects of session, $F(1, 17) = 20.30$, distance, $F(2, 34) = 18.59$, and their interaction, $F(2, 34) = 8.74$, were all significant, $p < .01$. When an alert occurred after a distance-0 event, users overwhelmingly chose good breakpoints, but this ten-

dency decreased as the distance increased. For a distance of 2, the users chose bad breakpoints about 20% of the time in the long-ring session. The results in 3(b) show a similar pattern, but the stricter time constraints of the short ring session act as a further influence. When the distance was 0, users in this session again exhibited an overwhelming preference for good breakpoints. When distance was 2, however, users chose bad breakpoints a majority (66%) of the time.

More insight into user behavior and strategies can be derived from a different view of the data. Figure 4 shows the distribution of the number of events from ring to answer broken down by session and distance. We performed a repeated-measures ANOVA on these data, specifically on the average number of events between ring and answer, and again the effects of session, $F(1, 13) = 37.06$, $p < .01$, and distance, $F(2, 26) = 5.23$, $p < .05$, were significant, although their interaction was only mildly significant, $F(2, 26) = 2.65$, $p < .10$. Looking more closely at the distributions, three distinct strategies can be identified and are labeled in the figure. The columns labeled *expected* illustrate a simple, expected strategy: users worked until the next good breakpoint and switched there. The *continue-on* label suggest a different strategy in the distance-0 case: users did not immediately switch, but worked for a short amount of time (1.5-8 s) and switched at a *subsequent* good breakpoint. Finally, the *give-up* columns suggest a third strategy: when close to a good breakpoint (distance 1), users often proceed to that good breakpoint before switching; however, if the good breakpoint is more distant (distance 2), they are likely to give up and switch immediately at the bad breakpoint. The strategies are evident (in differing frequencies) in both sessions.

A fascinating element of the *continue-on* and *give-up* strategies is the evidence that users are keenly aware of good breakpoints and use this knowledge to behave *adaptively* in anticipating them. In the *continue-on* strategy, users know that another good breakpoint is not far ahead, and thus they have time to perform another step in the primary task before switching. In the *give-up* strategy, users sense that a good breakpoint is far away, and decide simply to switch immediately rather than risk missing the time window for responding to the interruption.

DISCUSSION

Whereas prior work has generally tended toward an all-for-one explanation of task switching—that is, people wait until a point of lower workload to switch—our work demonstrates

Figure 4. Number of mail-task events between the onset of phone ringing and the actual task switch to the phone task, as a function of multitasking session and distance score. The labels indicate data suggestive of user strategies (see text).

that there is at least one other contributing factor, namely time constraints of the interrupting task, and that there is variability in user behavior as evident in three distinct user strategies for dealing with the interruption. The *expected* strategy correlates most with prior research [4, 8, 11] showing that users tend to process interruptions at points of lower mental workload—in our case, workload that arises from maintaining information to carry over from one task state to the next. While workload clearly was an influence, the *continue-on* and *give-up* strategies suggest that other factors can contribute to behavior, specifically that a combination of time constraints and distance to the next good breakpoint can also guide behavior for deferrable interruptions.

Of course, it is likely that still more factors influence task switching. A non-exhaustive list of potential factors includes task priorities, hierarchical task structure, stress factors, alert type, task length and complexity, and similarity between tasks. In our experiment, task priority was completely determined by users, and the tasks were relatively short with a simple structure. If these variables were changed, it is possible that different strategies would be exhibited. Furthermore, the phone-call alerts may have invoked a more immediate response based on users’ prior experience with telephones. Nevertheless, we expect the same strategies (*continue-on* and *give-up*) to manifest themselves in many situations, albeit potentially in different frequencies or forms. Further investigation would provide a more complete understanding of how strategies and time constraints more generally interact with other influences on multitasking behavior.

REFERENCES

1. B. P. Bailey and S. T. Iqbal. Understanding changes in mental workload during execution of goal-directed tasks and its application for interruption management. *ACM Trans. Comput.-Hum. Interact.*, 14(4):1–28, 2008.
2. B. P. Bailey and J. A. Konstan. On the need for attention-aware systems: Measuring effects of interruption on task performance, error rate, and affective state. *Computers in Human Behavior*, 22(4):685–708, 2006.
3. E. Cutrell, M. Czerwinski, and E. Horvitz. Notification, disruption, and memory: Effects of messaging interruptions on memory and performance. pages 263–269. IOS Press, 2001.
4. M. Czerwinski, E. Cutrell, and E. Horvitz. Instant messaging and interruption: Effects of relevance and timing. In *Proc. HCI 2000*, pages 71–76, 2000.
5. M. Czerwinski, E. Cutrell, and E. Horvitz. Instant messaging and interruption: Influence of task type on performance. In *Proc. OZCHI 2000*, pages 356–361, 2000.
6. V. M. González and G. Mark. Constant, constant, multi-tasking craziness: managing multiple working spheres. In *Proc CHI 2004*, pages 113–120. ACM, 2004.
7. S. T. Iqbal and B. P. Bailey. Investigating the effectiveness of mental workload as a predictor of opportune moments for interruption. In *Proc CHI 2005*, pages 1489–1492. ACM, 2005.
8. S. T. Iqbal and E. Horvitz. Disruption and recovery of computing tasks: field study, analysis, and directions. In *Proc CHI 2007*, pages 677–686. ACM, 2007.
9. K. A. Latorella. Investigating interruptions: An example from the flightdeck. *Proc. Human Factors and Ergonomics Society*, 40:249–253(5), 1996.
10. C. A. Monk, J. G. Trafton, and D. A. Boehm-Davis. The effect of interruption duration and demand on resuming suspended goals. *Journal of Experimental Psychology: Applied*, 14:299–313, 2008.
11. D. D. Salvucci and P. Bogunovich. Monotasking and multitasking: the effects of mental workload on deferred task interruptions. In *Proc. CHI 2010*. ACM, 2010.
12. D. D. Salvucci and N. A. Taatgen. *The Multitasking Mind*. Oxford University Press, New York, 2011.