Recovery from Interruptions: Knowledge Workers' Strategies, Failures and Envisioned Solutions

Simone Stumpf, Margaret Burnett, Thomas G. Dietterich,
Kevin Johnsrude, Jonathan Herlocker, Vidya Rajaram
Oregon State University
Corvallis, OR
{stumpf, burnett, tgd, johnsrud, herlock, rajaramv}@eecs.oregonstate.edu

ABSTRACT

This document presents qualitative results from interviews with knowledge workers about their recovery strategies after interruptions. Special focus is given to when these strategies fail due to the nature of the interruption and existing computer support. Potential solutions offered by participants to overcome some of these problems are presented. These findings have implications for researchers and designers of task-centric applications, especially in the area of support for recovery from interruptions.

1. INTRODUCTION

Tasks play an important role to knowledge workers [3, 14], a term used to describe computer users who spend the majority of their working hours processing information. It is commonly perceived that knowledge workers switch tasks because they are interrupted. Switching tasks involves unavoidable costs by potentially leaving the interrupted task not completed, by forcing users to reorient themselves and by necessitating the retrieval of appropriate task-related information. Since many interruptions are due to external sources [14], one approach to reduce this cost is to prioritize and schedule interruptions [20]. However, interruptions cannot be avoided in working environments, even though some mechanisms may exist for managing them. Therefore, recovery from interruptions is an important issue and solutions need to be developed to minimize the costs involved.

Knowledge workers constantly have to deal with interruptions, and it would be surprising if they had not developed some way of coping with them. After all, they recover from them day after day, so who better to ask about what works for them currently or what they think may work?

This document reports on an interview study with knowledge workers and presents findings on what strategies they currently employ to recover from interruptions. In particular, we highlight instances when existing strategies fail. We also describe knowledge workers' suggested solutions for recovery from interruptions and address potential implications for improved software tool support in the future.

In Section 2, we briefly summarize others' work relating to interruption recovery. Section 3 describes our interview study's setup, and Section 4 presents the results. Section 5 then draws from the results implications for researchers and designers of task-centric environments, and Section 6 summarizes the main findings.

Our results contribute to the understanding of recovery from interruptions and solution development. These results will help researchers and designers in the area of task-centric applications, especially in the area of support for recovery from interruptions.

2. RELATED WORK

2.1 Tasks and Multi-tasking

Previous research showed that projects, activities, and to-do items, which we will term tasks, are important to knowledge workers as structuring devices [3, 5], and this forms the backdrop of our research. As a result of these and other efforts, there are now a number of solutions that aim to support task-centric work [18, 11, 4, 15, 22, 6, 21, 4, 13, 10].

It has been noted that interruptions frequently cause knowledge workers to switch tasks [19, 8, 14]. Some of these interruptions are internal and initiated by knowledge workers themselves. Other interruptions are forced on them from external sources, such as phone calls, emails, or face-to-face meetings. Each entails some cost; this has driven research on the management of interruptions to reduce the frequency of having to switch between tasks.

2.2 Management of Interruptions

Previous work has identified immediate, negotiated, mediated and scheduled interruption styles [20]. The timing and amount of attention that the user must direct to either the task or the interruption vary with each of these styles and affects their performance. Therefore, none of these interruption styles is suitable for every task.

The point during a task at which interruptions occur can be crucial. Tasks are hierarchically composed of subtasks, and it has been found that an interruption between high-level task components is less disruptive than between low-level components, and that interruptions during subtasks are more challenging than interruptions occurring in natural divisions between subtasks [16, 7, 1].

Research has also been directed at making interruptions coincide with opportune times for users. For example, it has been investigated how to help people become aware when it is a good time to initiate an interruption in communications tools, such as Instant Messaging (IM), email, and phone calls [24]. However, it has been pointed out that an indication of presence does not equate with availability, and that people tend to ignore signs of unavailability [12].

While certain mechanisms may reduce external interruptions, they will still happen. Hence, research directed at helping users recover from these interruptions could significantly reduce overhead costs, errors and frustration.

2.3 Interruption Recovery

Recovery on tasks that span a longer time frame is perceived as more difficult since they usually are more complex and require more, and more diverse, information [8]. To overcome these difficulties, people use a variety of artifacts as reminders to pick up where they left off [5]. These reminders are sometimes placed before an interruption actually occurs, to maintain attention for returned-to tasks [19]. Consequently, computer-based solutions have involved the use of visual groupings to act as reminders and containers for project-specific information [22].

Two variables of interest in understanding the effects of interruptions are *interruption lag* and *cues*. Interruption lag is the time between an interruption occurring and the user responding to that interruption. If this time is long enough, the user has a chance to form or associate cues, which in turn shortens the *resumption lag*, i.e., the time it takes between ending one task and resuming another task [2].

3. STUDY SETUP

To add to this body of knowledge, we conducted semi-structured interviews with nine participants who were invited via an email recruitment notice. The participants had an expressed interest in our research area in that they were writers on the topic of multi-tasking, people interested in task-centric software, or recognized by peers as highly multi-tasking knowledge workers. All participants had been knowledge workers for a considerable number of years (mean=14.57, SD=7.32) in a variety of professions (e.g., software engineers, technical writers, conference managers, researchers). One participant was excluded from analysis because he had retired from professional work some time ago.

The interviews lasted from 45 minutes to 1 hour 30 minutes. Each concluded with a brief background questionnaire to gather data on professional background, computer experience and details of their computer environment. All phone interviews were audio-taped and transcribed for data analysis.

Our goal was to develop a rich understanding of the topology of recovery from interruptions, and, therefore, we adopted a qualitative approach [23]. In line with this approach, the questions were carefully worded to ensure that answers were grounded in specific examples and that the questions did not tamper with users' concepts in their own language (e.g., "how many things did you work on today?"). Our results are exemplified with quotes from participant transcripts, edited lightly for reasons of space.

4. RESULTS

4.1 Problems with Interruptions

Interruptions and multi-tasking have become so common in knowledge workers' lives that we had to modify our language that we initially used in pilot trials. On being questioned about their problems with multi-tasking, pilot participants reported that they don't have any real problems. It suggested to us that task switching has become so habitual to them that they no longer perceived this as being outside the norm. Furthermore, it suggests that they have devised strategies to minimize difficulties. (As a result of the pilot, we changed our language from asking about "problems" to asking them to report what happened the last time that they had to switch between units of work, and participants were more forthcoming after that modification was made.)

Our study confirmed that participants indeed faced a variety of interruptions. A very large number reported *internal* interruptions, such as "dread" or "guilt," which made them switch tasks. All of the participants were also interrupted through *external* sources, such as IM, phone calls, emails, and impromptu face-to-face meetings. A few of our participants managed external interruptions explicitly by delaying or refusing interruptions, such as by checking emails only at certain times, demonstrating that negotiated and scheduled interruption styles are sometimes adopted in a real environment. Another strategy that was reported by a participant for decreasing the severity of an interruption was by switching to a low information-carrying mode, both in the mathematical and social sense. In the particular example, he substituted IM for face-to-face meetings, thereby reducing the "social bandwidth" [9].

4.2 Level of Difficulty of Recovery from Interruptions

Some interruptions were perceived as being complex to recover from, but recovering from other interruptions was described as tedious and time-consuming but not mentally complex, as the following emblematic view describes:

"Well, to do the software it takes you awhile to get in the mind set with all. There's like 10 things you're trying to remember while you are writing it. And you just get to the sweet spot and then someone is calling you away to switch to a totally different task. And it's a totally different set of information. Then you have to track down a lot of information. And for me I will have other utilities that I run that are my personal utilities. There's a checklist and I will have to go find my checklist or keep it in my brain of things I'll do. And when you get done with that then you have to remember 'shoot, I was working on what part of that module...I had a strategy in mind two hours ago now I can't remember what I am doing.' I would also say that, I'm actually spending time in three different environments. So I have to go back, reconnect to the database server, go back open up my sources code and then pull up all my notes whether they be on a remote or back on Outlook or—it just takes time and also it's just a matter of getting your screens configured right."

The difficulty of recovering from an interruption is also strongly tied to the complexity of the task, as previous researchers have indicated [8]:

"Certainly going from coding to answering questions for sales or helping them with some document that they need to prepare is a lot easier, because it tends to just be at a higher level. I don't have to keep as much in my head at the same time. Usually when coding, there's just so many little details across many things. Whereas the other work, that is just kind of crafting individual answers to questions. So that way is a lot easier. Switching back to coding is definitely a lot harder."

A particular difficulty for participants appeared to be recovering after a face-to-face interruption. It could be that this particular difficulty is due to not having enough interruption lag to establish cues that could allow one to recover easily.

"Um, shoot. What was I working on when the IT guys showed up unexpectedly? That's what I'm trying to remember."

4.3 Current Strategies for Recovery from Interruptions

The participants in our study were able to provide strategies that they have used for recovering from interruptions. Some of them employed *internal* cues that they had memorized during the interruption lag. Others used explicit *external* cues that they prepared either in anticipation of an interruption or because the interruption lag was long enough. For example, people sent emails to themselves, or created post-it notes or markers within computer files. These function very much like reminders or to-dos in task management [5].

A very popular method was to leave visual cues, such as open windows or icons on desktops, to allow recovery of details on a task, such as the last state before the interruption or where to find information that allows one to resume:

"It's one of the things that I'll do a lot. I leave Exceed, that's our way of connecting to UNIX, up usually the entire week because that way I've got like six different desktops going. And I can open it and say, 'Okay, on [project X], I was doing this.'"

Participants reported that visual cues from screen configurations helped, since state information is encapsulated in visual cues. Costs increased when they did not have these visual cues available and instead had to rely on searching or browsing through folders, files and emails to retrieve their last current state.

Previous work has also pointed out that folder structure carries important information to users [17] and this was exploited by our participants.

4.4 Failure of Strategies

All of these recovery strategies rely on some kind of cue, and only rarely did participants report that there is no time to leave a cue at all. However, they were aware that sometimes their strategies would fail because the cue was impossible or difficult to retrieve. For example, most participants voiced that they had problems finding their reminders, such as notes, again. When these reminders were available, a new challenge presented itself to our participants. The shorthand reminders sometimes did not carry enough information to recover their last state, or the names or locations of the files that they needed to resume their work.

Although visual cues seemed very important to the participants, many of the participants pointed out that visual cues have drawbacks as well. First, there may not be enough screen real estate to display all the cues needed. Second, and worse, visual cues on the computer screen do not currently persist long enough for them to be used reliably:

"The problem is, whenever you have to stop what you were doing or restart your machine, all that gets lost, so you have to bring it back up and reposition it."

Even when there was a visual cue, often the most recent cue displayed was not valuable to our participants, as it did not carry the right state information:

"In what I'm doing, I'm making changes to several files at the same time, and just because something was the last thing doesn't mean I was really working there. I'm coding in one file and I needed to look at something quick in another file. And then I've got a phone call, and when I come back here I see this file in front of me but it wasn't really what I was working on. It just happened to be the last thing I had on my screen."

4.5 Knowledge Workers' Envisioned Solutions

Not surprisingly, some of our participants were concerned with finding their cues again. Managing their explicit reminders, such as notes, on a computer in a persistent way appeared to them as a possibility. Furthermore, they wanted the ability to put quick and easy markers on a variety of information. Some expressed an interest in different kinds of visualizations to provide time-based cues.

In short, in keeping with the importance of visual cues, many of our participants would welcome some kind of persistent, visual grouping of information:

"I could hit a button and they would open up my data base connection, my two software projects and, uh, some structured stuff. It would say 'here is what you were working on,' maybe give me structured text that would maybe be tagged in a way that would be useful."

This suggests that it could be valuable to support grouping of task information across applications and to recreate the environment that existed before the last interruption on that task.

5. IMPLICATIONS FOR RESEARCHERS AND DESIGNERS OF TASK-CENTRIC ENVIRONMENTS

Task-centric environments and their user interfaces need to support recovery from both internal interruptions (i.e., when the user controls when to switch tasks) and external interruptions (i.e., when the user does not control when to switch tasks). For external interruptions in particular, users appear to prefer negotiated or scheduled interruption styles, returning some form of control back to them. This is important to the user because it enables them to create and manage interruption lag.

Once an interruption occurs, it is of importance that adequate interruption lag is provided to enable the user to form cues that enable them to recover from the interruption. These cues could be internal, memorized cues or external cues (such as explicit notes, or visual cues). Also, external cues that are held by the system itself could be made easier to customize. For example, users might like to mark and group a number of documents during the interruption lag, perhaps including what action was being performed in that group of resources.

Existing strategies by knowledge workers underline the importance of visual cues. These visual cues encapsulate the last computer state before the interruption and hence function as a breadcrumb trail to information on the interrupted task. However, to have reliable value, these visual cues need to be persistent on the system and compact enough to not take up screen real estate. Ideally, visual cues could facilitate recovering not only computer state, but also the user's last *mental* state in an interrupted task.

The general theme of our participants was that the recovery from interruption needs to be made easier. To enable the reduction of costs, the user should be able to quickly retrieve information and recover state (both computer and mental) in relation

to the interrupted task. In considering strategies to support interruption recovery, designers need to also keep in mind that the more detailed and/or complex the interrupted task is, the harder it will be for a user to recover from an interruption.

6. CONCLUSION

In this document, we have described how knowledge workers recover from interruptions. Interruptions occur however much knowledge workers (or researchers or developers) attempt to manage them out of their lives. Certain aspects of interruptions pose particular challenges:

- Face-to-face interruptions appeared to be more difficult for our participants. This could be due to a lack of interruption lag.
- While visual cues were a popular way for our participants to encode information about where in a task a participant had been interrupted, any cue can fail if the cue becomes unavailable or difficult to retrieve. A dramatic failure reported by participants is that the persistence of visual cues over long periods is not supported.
- Although knowledge workers demand the support of visual cues, their last, most recent screen display may not point to the activity where they want to resume.

Taken together, these findings suggest the need for design features in task-centric environments that support reasonable and productive interruption lag combined with persistent, task-centric visual cues that enable recovery of both computer and mental state.

REFERENCES

1. Adamczyk, P.D. and Bailey,B.,P. If not now, then when? The effects of interruption at different moments within task execution. Proc. CHI 2004, ACM Press(2004), pp.271-278.
2. Altmann, E. M. & Trafton, J. G. Task interruption: Resumption lag and the role of cues. Proceedings of the 26th annual conference of the Cognitive Science Society 2004, 42-47.
3. Bannon, L., Cypher, A., Greenspan, S., Monty, M. Evaluation and analysis of users' activity organization. *Proc. CHI 1983*, ACM Press (1983), 54–57.
4. Bellotti, V., Ducheneaut, N., Howard, M., Smith, I. Taking email to task: The design and evaluation of a task management centered email tool. *Proc. CHI 2003*, ACM Press (2003), 345–352.
5. Bellotti, V., Dalal, B., Good, N., Flynn, P., Bobrow, D. and Ducheneaut, N. What a to-do: Studies of task management towards the design of a personal task list manager. Proc. CHI 2004, ACM Press(2004), 735-742.
6. Card, S. and Henderson, A. (1987). A multiple, virtual-workspace interface to support user task switching. Proceedings of the SIGCHI/GI conference on Human factors in computing systems and graphics interface, ACM Press(1987), 53-59.
7. Czerwinski, M., Cutrell, E., and Horvitz, E. 2000a. Instant messaging and interruptions: Influence of task type on performance. In Proceedings of the Australian Conference on Computer-Human Interaction (OZCHI 2000). 356--361.
8. Czerwinski, M., Horvitz, E., Wilhite, S. A diary study of task switching and interruptions. *Proc. CHI 2004*, ACM Press (2004), 175–182.
9. Donath, J. Sociable Media. Encyclopedia of Human-Computer Interaction., Berkshire Publishing Group 2004.
10. Dourish, P., Edwards, K.W., LaMarca A., Salisbury,M. Presto: An experimental architecture for fluid interactive document spaces. *ACM Transactions on Computer-Human Interaction 1999*, 6(2), ACM Press (1999), 133–161.
11. 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. *Proc. IUI 2005*, ACM Press (2005), 75–82.
12. Fogarty, J., Lai, J., and Christensen, J. 2004. Presence versus availability: The design and evaluation of a context-aware communication client. Int. J. Hum.-Computer Studies 2004, 61(3), 299-317.
13. Freeman, E., Gelernter, D. Lifestreams: A storage model for personal data. In *SIGMOD Record*, 25(1), ACM Press (1996), 80–86.
14. Gonzalez, V., Mark, G., "Constant, constant, multi-tasking craziness": managing multiple working spheres. *Proc. CHI 2004*, ACM Press (2004), 113–120.
15. Gwizdka, J., TaskView: Design and evaluation of a task-based email interface. *Proceedings of the 2002 conference of the Centre for Advanced Studies on Collaborative research.* IBM Press, 2002.
16. Iqbal, S.T., Adamczyk,P.,D., Zheng,X.,S. and Bailey,B.,P. Towards an index of opportunity: Understanding changes in mental workload during task execution. Proc. CHI 2005, ACM Press(2005), 311-320.
17. Jones, W., Phuwanartnurak. A.J., Gill, R., Bruce, H. Don't take my folders away! Organizing personal information to get things done. In *Ext. Abstracts CHI 2005*, ACM Press (2005), 1505–1508.
18. Kaptelinin, V. UMEA: Translating interaction histories into project contexts. *Proc. CHI 2003*, ACM Press (2003), 353–360.

19. Mark, G., Gonzalez, V. and Harris, J. No Task Left Behind? Examining the Nature of Fragmented Work. Proc. CHI 2005, ACM Press(2005), 321-330.

20. McFarlane, D. and Latorella, K.The Scope and Importance of Human Interruption in Human-Computer Interaction Design. Human-Computer Interaction 2002, 17(1), 1-61.

21. Robertson, G., van Dantzich, M., Robbins, D., Czerwinski, M., Hinckley, K., Risden, K., Thiel, D. and Gorokhovsky, V. The Task Gallery: A 3D Window Manager. Proceedings of the SIGCHI conference on Human factors in computing systems 2000, ACM Press (2000), 494-501.

22. Smith, G., Baudisch, P., Robertson, G., Czerwinski, M., Meyers, B., Robbins, D., Andrews, D. GroupBar: The taskbar evolved. *Proc. OZCHI 2003*.

23. Strauss, A. and Corbin, J. (1998). Basics of Qualitative Research. London, Sage.

24. Tang, J. C., Yankelovich, N., Begole, J., Van Kleek, M., Li, F., and Bhalodia, J. 2001. ConNexus to awarenex: extending awareness to mobile users. In Proc. CHI 2001. ACM Press, 221-228.