Current Research in Workplace Interruption Management

COMP 171

James Kebinger

April 19, 2005

Introduction

This paper will survey some research into the effects of interruptions, specifically focusing on interruptions triggered by instant messages in the workplace. The paper begins with a brief introduction to instant messaging in the workplace, then look at the effects of interruptions on workflow. Next some methods for coordinating interruptions are described and compared, followed by two areas in which researchers are working to better support interruption-limitation for remote collaboration.

Collaborative Nature of the Modern Workplace

The modern workplace is ever more distributed, with "virtual" teams distributed across buildings, countries and time zones. At the same time there is a growing wealth of technology to deal with the problems of interacting over a distance. Chief among them in many companies is Instant Messaging [IM]. IM allows one to place a textual note on the screen of another user wherever he or she may be. Rennecker and Godwin [15] provide an overview of the features of contemporary instant messaging systems.

Nardi and Whittaker [14] observed side effects of instant messaging's awareness features (features that show a list of users as online or offline) like "waylaying" where users would watch for a user to come online, and immediately contact the user vi IM, a phonecall or even stopping by his or her office. Another aspect they termed "outeration" is the use of IM to facilitate communications via other media. An example of this is instant messaging someone to determine if it is a good time to call or stop by.

Effects of Interruptions

The negative effect of an interruption on productivity is much more than the time spent handling the interruption. Nielsen [6] notes: "… even a one minute interruption can easily cost a knowledge worker 10 to 15 minutes of lost productivity due to the time needed to reestablish mental context and reenter the flow state." Picard [4] concurs, "We take a major productivity hits with each interruption"

Quantitative studies of interruption effects note the negative effects on interruption. McFarlane's [2] study involving a simple game in which figures jumping from a window had to be caught showed an average 36% decrease in effectiveness between the no-interruption and interrupted trials. The correct handling of the interrupting task decreased 70 % in trials between trials without and with the game running.

Horvitz [8] found that subjects searching a list of book titles for a particular title or subject were reliably slower in trials with interruptive notifications. The trial also showed that test participants were significantly more likely to request a reminder of the task to be performed in trials with interruptions.

The subject matter of an interruption can alter its effect on performance. [7] Increasing relevancy between the content of the interruption and the primary task tends to moderate the impact on the primary task. Horvitz [7] also found that by breaking down the performance of a search task into planning, execution and evaluation stages, it could be shown that the worst time for interruptions is during the evaluation of search results.

Categories of Interruptions

In order to design effective solutions to the problems related to interruptions, it is useful do have a vocabulary for understanding the types of interruptions and a model for how humans mentally process interruptions. McFarlane [1] creates taxonomy for interruptions along many different axes. The factors used in the taxonomy are: source of interruption, individual characteristic of person receiving interruption, method of coordination, meaning of interruption, channel of conveyance, human activity changed by interruption, and effect of interruption.

Latorella [1] describes a model for the human mental processing of interruptions called the interruption management stage model (IMSM). The model describes the four primary effects of interruption as diversion, distraction, disturbance and disruption. IMSM further describes the stages of the interruption handling process during which these effects are primarily manifested.

Methods for Coordinating Interruptions

As it is useful to describe taxonomy for types of interruptions, it is also useful to group the approaches used to handle interruptions into a number of categories. McFarlane [2] groups the approaches into four primary methods for coordinating user interruptions: immediate, negotiated, mediated and scheduled.

The immediate (or no coordination) method describes the usual approach for many systems, which is to immediately interrupt the user's current task and force the user to switch his or her attention to process the interrupting task. (This is primarily valid for systems in which only one task can be displayed and performed simultaneously; most task notifications in multitasking operating systems wouldn't apply)

In the negotiated method the user is notified that a task is in need of attention, but instead of being forced to handle it immediately, the user can choose to either handle or ignore it until sometime in the future. There may be some attention cost for the user to deal with the notification, but the user can then determine when the best time to divert his or her attention to the waiting task and handle it at that time.

The mediated method is one in which a heuristic or a set of rules determines when best to interrupt the user. This decision can be based on a wide set of criteria including how busy the user currently is in performance of the primary task, how important the

primary task is, how important the interruption is, and if the user is alone or in a meeting, to name just a few potential measurements.

The scheduled method involves queuing interruptive tasks until a time period has elapsed, at which time the user can be made to handle the entire queue of tasks before returning to performing his or her primary work.

In practice many interruption management systems [3,4,5] are based on the mediated approach but also blend aspects of each of the other methods.

Comparison of Methods for Coordinating Interruptions

McFarlane [2] performed an experiment to compare the efficacy of each of the four identified methods for interruption coordination. The experiment consisted of primary task, which was an absorbing interactive (but non-pauseable) game, and an interruptive task, which was a simple shape-matching task. The interruptive task obscures the entire screen while being performed, leaving the game running in the background. As in the real world, accurate performance on both the primary task and the interruptive task are important.

The experiment showed the expected effect that any interruptions hinder performance on the primary task. McFarlane generated a set of tentative design guidelines indicating that the best coordination method for accuracy on the primary (continuous) task is negotiated, while the worst is scheduled. The best method for accuracy on the interruptive (intermittent) task is anything but immediate, that is one of scheduled, mediated or negotiated. The guidelines are interesting in that they show it is difficult to build a system that enhances accuracy for both the continuous and immediate tasks.

The results of this study presenting negotiated coordination as a best practice should be tempered by the naïve measure used in by the mediator to determine if the user was in a high or low cost-of-interruption state. It may be that negotiated shows better results because the user is a better judge of how able he or she is to handle an interruption at that point in time than the mediator. With a much-improved busyness evaluation model, mediated may turn out to be better than negotiated.

Horvitz [9] presents an approach combining scheduling and mediation named *bounded deferral* for coordinating interruptions. This approach is similar to scheduled, only it waits at most a given time interval to alert the user. If in the interim, the user transitions to a low cost-of-interruption stage, the notification is delivered at that time. User studies [4,9] showed that users spend only between twenty and thirty percent of their day in a high cost-of-interruption state and that there is a high probability of a user in the high cost state transitioning to a lower cost state in at little as one to two minutes.

Approaches to Handling Collaborative Interruptions

People in the workplace already have a variety of ways to signal low-interruption receptivity, including closing office doors, using headphones and wearing frowned masks of deep concentration. People can usually sense how busy a person is from prior interactions or from social cues emitted by each person. There are not in wide use suitable mechanisms to be able to signal how receptive one is to communication and block or filter communication attempts through remote mechanisms like the telephone or instant messaging.

There are two primary facets to the research in this area. The first is to make as many social cues available online as possible, and rely on people to evaluate those cues before negotiating for interaction. The second is to use as much data as possible to evaluate the availability of the user and use that as input to "attention firewall" systems that will redirect or refuse incoming interruptions if the user is busy.

Publish Availability Information

In person, people are generally excellent at evaluating the busy state of people with whom they wish to interact. In a remote situation, these social cues are lost, requiring more work on the part of both parties to better negotiate availability.

Erickson and Kellogg [10] present the idea of *social translucency*, with its three facets of visibility, awareness, and accountability. (The term translucency is used rather than transparency to account for the tension between privacy and information availability.) The hope is that if one makes known his or her level of availability in an easy to understand way, and one knows that others know this information, people will self-regulate their willingness to interrupt the user depending upon the available information.

Current Instant Message (IM) systems indicate availability on a course-grained scale that often only has two or three levels such as "Active", "Away", and "Do Not Disturb". It is tedious for a user to ensure that these systems are up to date, and therefore very likely that the information will be inaccurate [3]. It would require a much more rich set of data that is automatically updated to provide a level of function sufficient and reliable enough to deter others from interacting with a given user.

A study by Dabbish and Kraut [11] bears out the idea that in two person teams, a person given information about the busyness of another person and an enough incentive to perform well will effectively throttle the rate of interruptions at high cost-of-interruption times. The results showed a 43 % decrease in interruptions generated at high-cost moments, and an overall decrease in interruption rate as the level of abstraction about the user's bust state decreased from none to a full display of the user's activity. There was no significant difference in performance between an abstract and full representation of the busy state, which is reassuring for alleviating privacy concerns in real world implementations. Only just over 60% of the users actually reported using the

busyness displays during the experiment when available, showing that even under experimental conditions, people do not always use all the cues available to judge the best time to interact with another person.

A number of experiments have expanded from what is on the user's screen as a proxy for busy-ness to what is going on in the user's office in its entirety. Portholes [12] involved still video images to provide the social cues normally absent in remote situations. NYNEX Portholes [13] further improved on the model by providing capabilities to enhance privacy and augment the information content of the images with graphs of sensed activity levels.

Personal agent approaches

Neilsen [6] notes "IM… lets your agenda be controlled by anybody who has your screen name". Executives have the benefit of an administrative assistant or receptionist that can server as a buffer against those seeking to interact; the assistant can weigh how busy the executive against the potential benefit of dealing with the information in scheduling the executive's time. This process of evaluating the cost of the interruption versus the benefit of the interruption is at the heart of any effective interruption management solution. There are many efforts underway to automate this attention mediation process and make the benefits available to a wider user population in order to empower people to better manage their own attention-agenda. [4]

The benefit of a personal agent approach over the information publishing approach is not only that it enforces the otherwise voluntary system that relies on the good judgments of others, but also that interruptions can be handled far more flexibly. Notifications can be suppressed entirely, queued for processing later, shown with some ambient indicator, or brought to the users immediate attention, to name a few variations on the theme.

Dekel and Ross [3] present a rule-based interruption-mediator, *Gatekeeper*, that was built as part of effort, *Jazz*, to integrate collaborative features into an integrated development environment (IDE). Jazz generates a rich stream of information useful for team software development, such as who is looking at or has modified the resources a given user may be currently working on, as well as providing standard text based chat, voice over IP (VOIP) telephony and screen sharing. No person can efficiently process so much information, and during the process of software development, there are particularly bad times for interruptions, such as debugging sessions.

Gatekeeper works by evaluating an incoming message against a list of policies that specify what action to take with a given interruption based on the kind of interruption, whom the interruption is from, and the current state of the user. The current state is based on the task being performed in the IDE or another window based on a set of known applications. Interruptions can be allowed to proceed immediately, allowed to prompt to negotiate availability, shown in an ambient manner, or suppressed completely. Gatekeeper simplifies the iterative creation of more effective rules over time with a

capability that shows for a given interruption, which policies influenced (or failed to influence) the processing of that interruption.

Horvitz has developed a system *Bestcom/Enhanced Telephony* that is being used by 3,800 people at Microsoft to manage interruptions from phone calls, instant messages and emails. [4] The system has a wide range of sensors to evaluate the availability of a user such as microphones to detect if there is a conversation underway and the use of inexpensive web cameras that can determine if there's anyone in the office, or more than one person having a meeting. *Bestcom* the importance of incoming calls based on if the user has recently called the number from which the call is coming, if the number is in the user's personal address book or belonging to a member of the user's team or in the user's family. The system can also use the caller and user's calendar to schedule a call for a mutually convenient time.

Gatekeeper and *Bestcom* differ in an important way: *Bestcom* uses a heuristic modeling approach whereas *Gatekeeper* uses a rules-based system. The problem with statistical modeling approaches is that they are sometimes wrong: this is frustrating to users who may not be able to understand or determine why an interruption was handled in a particular (incorrect) way. [4]. As systems become more sophisticated to the point where users are increasingly reliant on its correct operation this is an important problem to solve. Ben Schneiderman highlights this problem, [4] "The more 'attentive' things become, the more unpredictable they are … We have a history in this community of creating 'smart' devices that people don't use because they can't understand how they operate"

Conclusions

In the future as all humans become increasingly connected and reachable, it will become increasingly important to be able to moderate the level of interruptions experienced while doing important things. A future where cell phones know they're at the symphony certainly would be a start. There's a wealth of interruption-related research happening already; with more powerful devices emerging constantly and increasing availability of semantic content the accuracy and sophistication of interruption mediation can only increase. This looks to be an exciting research area for some time to come.

References

[1] McFarlane D. C. and Latorella K. A. (2002) The scope and importance of human interruption in human-computer interaction design, Human-Computer Interaction, 17 (1), 1-61

[2] McFarlane D. C. (2002) Comparison of four primary methods for coordinating the interruption of people in human-computer interaction, Human-Computer Interaction, 17 (1), 63-139.

[3] Dekel U and Ross S(2004) A Plug-In Based Framework for Research on Interruption Management in Distributed Software Development. To appear in *Proceedings of the 2004 OOPSLA workshop on eclipse technology exchange*. ACM Press, not yet published.

[4] Gibbs W. (2005) Considerate Computing. Scientific American. January 2005, p 55-61

[5] Avrahami D. and Hudson S. E. (2004) Balancing performance and responsiveness using an augmented instant messaging client. In *Proceedings of ACM Conference on Computer Supported Cooperative Work (CSCW'04)*, New York: ACM Press.

[6] Nielsen J (2003) IM, Not IP (Information Pollution). *Queue*, 76-75

[7] Cutrell E. B., Czerwinski M. and Horvitz E. (2000) Effects of instant messaging interruptions on computing tasks. In: *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2000 Extended Abstracts)*. New York: ACM Press, 99-100.

[8] Cutrell E., Czerwinski M. and Horvitz E. (2001) Notification, disruption, and memory: Effects of messaging interruptions on memory and performance. In: M. Hirose (Ed.), *Human-Computer Interaction – INTERACT 2001 Conference Proceedings*. Amsterdam: IOS Press, 263-269.

[9] Horvitz E., Apacible J. and Subramani M. (2005) Balancing awareness and interruption: Investigation of notification deferral policies. *In: User Modeling 2003: Proceedings of 10th International Conference (UM 2003)*, Berlin: Springer-Verlag

[10] Erickson T and Kellogg W. (2000) An Approach to Designing Systems that Support Social Processes. In *ACM Transactions on Computer-Human Interaction*, Vol 7, No 1 March 2000 p 59-83

[11] Dabbish L. and Kraut R. (2003) Coordinating communication: Awareness displays and interruption. In*: Proceedings of ACM Conference on Human Factors in Computing Systems (CHI'03): Extended abstracts*, New York: ACM Press, 786-787

[12] Dourish,P. and Bly, S. (1992). Supporting awareness in a distributed work group. *Human Factors in Computing Systems, CHI '92 Conference Proceedings*, pp. 541–547. New York: ACM.

[13] Lee, A and Girgensohn, A (2002) Design, experiences and user preferences for a web-based awareness tool. *Int. J. Human-Computer Studies* 56, 75–107

[14] Nardi B. A., Whittaker S. and Bradner E. (2000) Interaction and Outeraction: Instant messaging in action. In: *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW'00)*, New York: ACM Press, 79-88.

[15] Rennecker J. and Godwin L. (2003) Theorizing the Unintended Consequences of Instant Messaging for Worker Productivity, *Sprouts: Working Papers on Information Environments, Systems and Organizations*, 3 (Summer).