The Effects of Interruptions on Task Performance, Annoyance, and Anxiety in the User Interface

Brian P. Bailey, Joseph A. Konstan, and John V. Carlis

University of Minnesota
Department of Computer Science and Engineering
Minneapolis, MN 55455

{bailey, konstan, carlis}@cs.umn.edu

Abstract: When an automating application needs a user's input or has feedback or other information for that user, it typically engages the user immediately, interrupting the user's current task. To empirically validate why unnecessarily interrupting a user's task should be avoided, we designed an experiment measuring the effects of an interruption on a user's task performance, annoyance, and anxiety. Fifty subjects participated in the experiment. The results demonstrate that an interruption has a disruptive effect on both a user's task performance and emotional state, and that the degree of disruption depends on the user's mental load at the point of interruption. We discuss the implications of these results in terms of building a system to better coordinate interactions between the user and applications competing for that user's attention.

Keywords: Agents, Annoyance, Anxiety, Automated Tasks, Awareness, Interruption, Intrusion

1 Introduction

As users continue offloading more control and responsibility to automating applications such as interface agents, softbots, and peripheral information displays, these applications must increasingly compete for user attention. User attention must be periodically gained in order for an automating application to receive additional guidance from the user (Horvitz, 1999; Maes, 1994), provide feedback regarding decisions made on the user's behalf, or keep the user aware of peripheral information (Bailey et al, 2000b; Maglio & Campbell, 2000). When an automating application wants user attention, it can either engage the user immediately, interrupting the user's current task, or wait for a more opportune moment.

Waiting for an opportune moment before interrupting someone's task is a social behavior commonly found in human-human interaction. Interrupting a person who is visibly concentrating on a task, except in the most extreme circumstances, is considered rude and socially unacceptable behavior, as it disrupts that person's concentration. Analogously, we argue that it is equally rude and distracting for an automating application to unnecessarily interrupt a user's current task. An application must adhere to the same interruption protocols already established in human-human social interaction, following the theme that computers are social actors (Nass et al, 1994).

Although interrupting a user's task is rude behavior, the goal of this work is to provide quantitative evidence of the disruptive effects of an application-initiated interruption on a user's task performance, annoyance, and anxiety. By providing this evidence, we lay an empirical foundation from which to justify building systems that observe or predict opportune moments for gaining user attention.

Although other researchers have measured the effects of different peripheral information displays on a user's task performance, awareness, and distraction (Maglio & Campbell, 2000), our work is the first to measure the quantitative effects of manipulating the *time* of a peripheral information display on a user's task performance, annoyance, and anxiety. Peripheral information is nonessential information that is helpful or of interest to the user but not necessarily related to the user's current task.

In our experiment, we used two categories of peripheral information, breaking news headlines and stock market updates. To ensure the peripheral information was read and comprehended by a user, each was structured in the form of a *peripheral task*. An *interruption* was a peripheral task presented to a user while performing a primary task.

A user from either of two groups, a control and experimental group, performed eighteen primary tasks, three from each of six task categories. A user from the control group was presented with a peripheral task *just after* completing two of three primary tasks in each category, while a user from the experimental group was interrupted *during* two of three primary tasks in each category. In both conditions, the user attended to the peripheral task immediately.

The key findings of this work are that (i) a user performs slower on an interrupted task than a non-interrupted task, (ii) the level of annoyance experienced by a user depends on both the category of primary task being performed and the time at which a peripheral task is displayed, (iii) a user experiences a greater increase in anxiety when a peripheral task interrupts her primary task than when it does not, and (iv) a user perceives an interrupted task to be more difficult to complete than a non-interrupted task.

The rest of this paper is organized as follows. In section 2, we explain the rationale for our experiment and define our experimental hypotheses. In section 3, we define our experimental method and then report on the analysis of the data in section 4. In section 5, we discuss the implications of our results in terms of building an attention manager that coordinates interactions between the user and applications competing for her attention. In section 6, we summarize our key findings.

2 Rationale for the Experiment

Various effects of interruptions have previously been studied in both psychology and human-machine interaction. (Zijlstra et al, 1999) measured the effects of interruption frequency and complexity on a user's emotional state and task performance in the context of document-editing tasks. Although a significant difference in anxiety was detected, the difference was attributed to the difference in interruption complexity and *not* to the event of being interrupted per se. The effect on a user's anxiety due to the time at which a peripheral task is presented has not previously been studied, which is a goal of our experiment.

The authors also found that interrupting a user during the document-editing tasks caused that user to complete the tasks *faster* than when performing the same tasks without interruption. The more often a user was interrupted during the editing tasks, the faster that user completed those tasks.

In contrast to Zijlstra et al, (Kreifeldt & McCarthy, 1981) found that interrupting a user while performing a series of calculator-based tasks caused that user to complete those tasks *slower* than when performing the same tasks without interruption. Clearly, the conclusions derived from these two independent studies are inconsistent and further investigation into the effects of an interruption on a user's task performance is warranted.

To the best of our knowledge, our experiment is the first to quantitatively measure the subjective level of annoyance experienced by a user due to interruptions in the user interface.

In sum, our experiment should (i) help resolve the contrasting conclusions regarding the effect of an interruption on task performance and (ii) provide a first attempt at measuring the effect of manipulating the time of an interruption on a user's anxiety and subjective level of annoyance.

2.1 Experimental Hypotheses

Based on related work, we formulated five hypotheses for our experiment, structured in terms of task performance, annoyance, and anxiety.

Task Performance. Although previous research has provided contrasting conclusions, it does seem reasonable that task re-orientation after an interruption would cause performance degradation:

H1: An interrupted task will require more time to complete than a non-interrupted task within the same task category.

Additional hypotheses and analysis relating to the effects of an interruption on a user's task performance can be found in (Bailey et al, 2000a).

Annoyance. Although not previously investigated, interrupting a user engaged in a task should cause that user to experience a higher level of annoyance than when not interrupted. Thus, we formulated three hypotheses relating to annoyance:

H2: A user will experience a higher level of annoyance when a peripheral task interrupts her primary task than when it does not.

H3: When a peripheral task interrupts a primary task, the level of annoyance experienced by a user will depend on the category of that primary task.

H4: When a peripheral task is presented just after the completion of a primary task, the level of annoyance experienced by a user will not depend on the category of that primary task.

The first of these three hypotheses compares the level of annoyance experienced by a user in the control vs. experimental group. The second hypothesis compares the levels of annoyance experienced by users only within the experimental group, while the third hypothesis compares the levels of annoyance only within the control group.

Anxiety. Based on (Zijlstra et al, 1999) and other research demonstrating that interruptions cause increased levels of stress (Boucsein, 1987; Johansson & Aronsson, 1984), a user from the experimental group should experience a greater increase in anxiety than a user from the control group:

H5: A user will experience a greater increase in anxiety when her primary task is interrupted by a peripheral task than when it is not.

The experiment designed to test our five hypotheses is described next.

3 Experimental Method

3.1 Subjects

50 subjects (30 male, 20 female) participated in the experiment. Subjects were between the ages of 18 and 40, had at least one year of computer experience, and were a mix of students and local professionals. A subject was compensated for his/her participation with a five-dollar lunch coupon.

3.2 Experimental Design

The experiment consisted of two groups of users, six primary task categories, and two peripheral task categories. A user was randomly assigned to either the control or experimental group with the constraint that each group had an equal number of males (15) and females (10). A user from the control group was presented with a peripheral task *just after* completing two of three primary tasks from each category, while a user from the experimental group was interrupted with a peripheral task *while performing* two of three primary tasks from each category. In both conditions, the user attended to the peripheral task immediately. Within the experimental group, we analyzed the difference in task performance between the two interrupted tasks and the one non-interrupted task within each category. Between the groups, we analyzed the differences in anxiety and subjective level of annoyance measured through a self-evaluation state anxiety form and a pencil and paper questionnaire, respectively.

We wanted to analyze whether the effects of an interruption would depend on the category of the primary task being performed when that interruption occurred. Thus, our experiment used six primary task categories, each varying in difficulty.

Primary Task Categories

The six primary task categories used in the experiment were:

- *Addition*. Four numbers, each consisting of four digits, were presented to the user. The numbers were right aligned in a 4-row x 1-column table. The task was to add the numbers and then enter the correct sum into a text field positioned underneath the last number.
- *Counting*. A set of 40 words was arranged in a 10-row x 4-column table and presented to the user. The 40 words were randomly chosen from a base set of six words, i.e., each of the six words was repeated in the table. The task was to first count the number of words in the table matching a target word chosen from the base set, and to then enter this count into a text field.
- *Image Comprehension*. A completed tournament bracket starting with eight teams was presented to the user. The task was to answer five questions regarding the outcomes of the pairings. An example of this task is shown in Figure 1.
- *Reading Comprehension*. A short passage (~4-5 sentences) was presented to the user. The task was to read the passage and then answer three questions regarding its content.

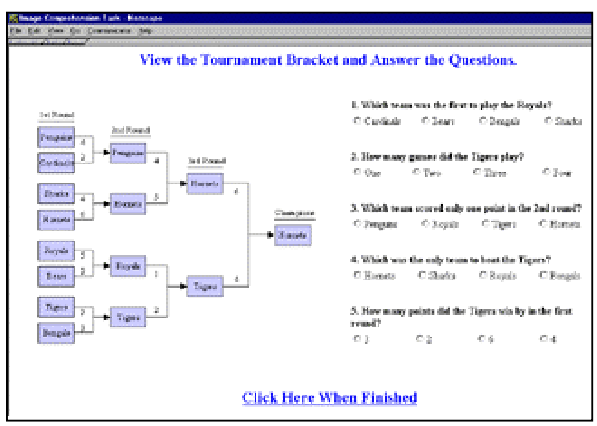

Figure 1: An example of an image comprehension task. A user from the control group would be presented with a peripheral task just after clicking the finished link, whereas a user from the experimental group would be presented with a peripheral task about halfway through this task.

- *Registration*. Eight registration-style questions were presented to the user, e.g., name, age range, and political affiliation. The task was to enter the requested information using three interaction formats; toggle sets, drop-down lists, and free-form text fields. The interaction format was homogenous for each question.
- *Selection*. A set of 40 words along with checkboxes was arranged in a 10-row x 4-column table and presented to the user. The 40 words were randomly chosen from a base set of six words, i.e., each of the six words was repeated in the table. The task was to select each word in the table that correctly matched a target word chosen from the base set.

The task categories were designed to be of varying difficulty and duration (~15-40s). Because a user would need to perform more than one task from each category, multiple sets of similar tasks were designed. The task screens were implemented using HTML and rendered with Netscape Navigator 4.7.

To estimate the completion time for a task within each category, we conducted a pilot study with five users and computed their average completion times.

Peripheral Task Categories

The two peripheral task categories used in the experiment were:

- *Reading comprehension*. A short (3-5 sentence) news summary was presented to the user. The task was to read the summary and then select the most appropriate title from among three choices. Each news summary and its actual title were obtained from an existing news site to enhance realism.
- *Stock decision*. A stock scenario comprised of a fictitious company's name along with the quantity, date, and price of shares previously purchased of that company were presented to the user. The current stock price and a one sentence "news-flash" regarding the company were also presented. The task was to first read and analyze the scenario and to then select one of five actions; do nothing, buy a few more shares, buy many more shares, sell a few shares, or sell all the shares.

The peripheral tasks were designed to last approximately 10-30s. Because a user would receive more than one peripheral task from each category, multiple sets of similar tasks were designed. The peripheral task screens were implemented using HTML. Netscape Navigator was programmed using JavaScript to present the user with a peripheral task either just after the user completed a primary task, signified by clicking the finished link (see Figure 1), or about halfway through a primary task. The choice depended on whether the user had been assigned to the control or experimental group, respectively.

3.3 Hardware/Software

The experiment was conducted on a Pentium III 460 MHz machine with 128MB of RAM running Windows NT. The primary and peripheral tasks were designed using HTML. JavaScript 1.2 was used to implement the dynamics of the experiment, e.g., randomizing, sequencing, and displaying the tasks. Netscape Navigator 4.7 was used to execute the experiment. Each user's screen interaction was recorded for later analysis using Lotus ScreenCam.

3.4 Procedure

After a brief introduction, a user was asked to sign a consent form and then complete the self-evaluation state anxiety form (form Y-1) of the STAI (Spielberger, 1983). After completing these forms, the user moved to the computer to perform the primary and peripheral tasks.

Experimental Group. On the computer, a user completed 18 timed tasks from the six primary task categories (3 tasks per category). Prior to starting each task category, the experimenter gave a verbal description of the category and the user's task, allowed the user to perform a practice task, and answered any questions. A user was instructed to complete both the primary and peripheral task *as quickly as possible while maintaining accuracy on the task*. A user was also instructed to *immediately* attend to a peripheral task whenever it appeared. After any questions were answered, the experimenter left the testing area and the user performed three timed tasks from the current task category. One of the three primary tasks was interrupted using the news task, another with the stock task, and the remaining primary task was not interrupted and served as a control task. If a peripheral task was presented, it was presented approximately halfway through the primary task. This same process was followed for the remaining five task categories. The presentation order of the task categories, tasks within each category, and peripheral tasks was randomized.

Control Group. The procedure for the control group was similar to that of the experimental group. The only difference was that a peripheral task was now presented *just after* the completion of a primary task. Completion of a primary task was indicated by having the user select a "finished" link as shown in Figure 1. As with the experimental group, a peripheral task was presented for two of the three primary tasks performed within each category.

After completing the computer-based tasks, a user was asked to complete the following two forms in order:

- *Another self-evaluation state anxiety form of the STAI*. A user was instructed to complete the form according to how s/he felt, *on average*, while performing a peripheral task. The "on average" instruction was required because the user performed peripheral tasks across several task categories.
- *Pencil and paper questionnaire*. The questionnaire asked a user to rate three items. First, a user was asked to rate the relative difficulty level of each primary task category. Second, a user was asked to rate the level of annoyance experienced when attending to a peripheral task for each of the primary task categories. Finally, a user was asked to rate the level of annoyance experienced when attending to the news and stock tasks independent of the primary task categories.

The entire experimental procedure lasted no more than 60 minutes for a user.

3.5 Measurements

We measured a user's task performance, perceived level of task difficulty, subjective level of annoyance experienced, and anxiety.

Task Performance

For each primary task, the system recorded two performance measurements:

- *Time on Primary Task (TOT)*. The amount of time a user spent performing a primary task. This measurement did not include the time spent on a peripheral task, if presented.
- *Time on Peripheral Task (TOI)*. The amount of time a user spent on a peripheral task, if presented.

The performance measurements, along with the primary and peripheral task categories, were logged to a data file for analysis.

Perceived Level of Task Difficulty and Annoyance

The perceived level of difficulty for each primary task category and the level of annoyance experienced by a user due to the peripheral tasks were measured using a pencil and paper questionnaire. The questionnaire asked a user to rate the:

- *Relative difficulty level for each primary task category*. The difficulty scale ranged from 1 (easiest) to 6 (hardest). A user indicated the level of difficulty by placing an 'X' in the appropriate location and was allowed to rate two or more categories as being equally difficult.
- *The level of annoyance experienced for each category of primary tasks when a peripheral task was presented*. A user placed six labels representing the primary task categories along a single, continuous scale ranging from 1 (Not Annoying) to 25 (Intolerable).
- *The level of annoyance experienced for each category of peripheral tasks, independent of a primary task category*. A user placed two labels representing the peripheral task categories along a single, continuous scale ranging from 1 (Not Annoying) to 25 (Intolerable).

The annoyance scale was pre-marked in nine equidistant locations identifying different levels of annoyance such as slightly annoying, somewhat annoying, and moderately annoying. Our annoyance scale was adapted from scales used to measure annoyance caused by aircraft (Gunn et al, 1981) and transportation noise (Miedema & Vos, 1999).

Anxiety

A user's anxiety was measured just before and after performing the computer-based tasks. Both anxiety measurements were measured using the self-evaluation state anxiety form (Y-1) of the STAI.

4 Analysis

In this section, we provide an analysis of the collected data, structured in terms of task performance, annoyance, and anxiety. Because gender did not show a main effect in the analysis of the data, we do not include it here.

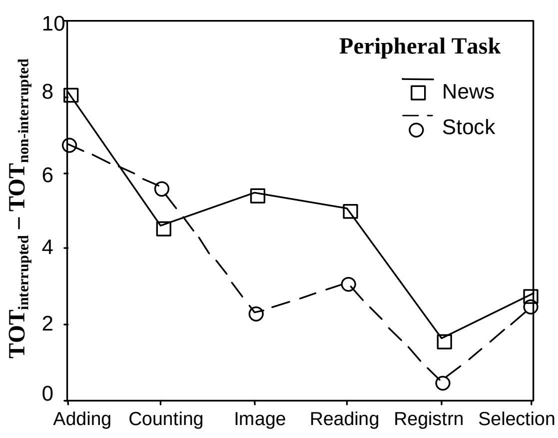

Figure 2: The differences in task completion time (TOT) for interrupted vs. non-interrupted tasks of the experimental group.

4.1 H1: The Effect of Interruptions on Task Performance

Before analyzing the task performance data, we reviewed each user's screen interaction and removed performance measurements associated with a task having substantial error. This was done to ensure the performance times analyzed were from users completing the tasks in a similar manner. Errors were not analyzed, because most were procedural errors such as not meeting the stated objective of a task.

For each subject in the experimental group, the differences in TOT between the two interrupted tasks and the one non-interrupted task were calculated for each category. The TOT differences are graphed as a function of both primary and peripheral task category in Figure 2.

The TOT differences were analyzed using a full-factorial ANOVA with the primary and peripheral task category as factors. The primary task category had a main effect on the TOT differences $F(5,260)=4.08$, $p<.01$), while the peripheral task category did not ($F(1,260)=1.68$, $p<.20$). No interaction was detected $F(5,260)=.39$, $p<.86$). These results indicate that the disruptive effect of an interruption in terms of task performance depends on the category of task being performed (memory load) when that interruption occurs (Bailey et al, 2000a).

To compare whether an interrupted task requires more time to complete than a non-interrupted task, we performed one-tailed t-tests comparing the TOT differences in each category with 0. Mostly consistent with our first hypothesis, a user required more time to complete an interrupted task than a non-interrupted task for all categories except Registration (Adding, $t=6.34$, $p<.00$; Counting, $t=8.90$, $p<.00$; Image, $t=3.43$, $p<.01$; Reading, $t=2.56$, $p<.01$; Selection, $t=8.42$, $p<.00$; Registration, $t=1.25$, $p<.11$). Although the mean TOT differences for Registration did not reach a significant level, the mean was in the positive direction. Furthermore, this result seems reasonable as the registration tasks ostensibly required the lowest memory load at the point of interruption.

Finally, the category of primary task had no effect on the amount of time a user spent on an interruption (Bailey et al, 2000a). Together, the analysis of the performance data indicates that a user can switch easily from a primary task to a peripheral task, but has difficulty switching back to the previously suspended primary task. And the level of difficulty experienced depends on the memory load of the primary task at the point of interruption.

4.2 H2-H4: The Effect of Interruptions on Annoyance

H2 - Level of annoyance experienced by a user in the control vs. experimental group

The peripheral task category annoyance ratings were analyzed using a full-factorial ANOVA with peripheral task category and group as factors. The ratings are graphed in Figure 3.

The results are consistent with our second hypothesis. Whether a user was presented with a peripheral task during or just after a primary task (group) had a main effect on the level of annoyance experienced when performing that peripheral task ($F(1,96)=20.226$, $p<.00$). The category of peripheral task did not have a main effect on the annoyance

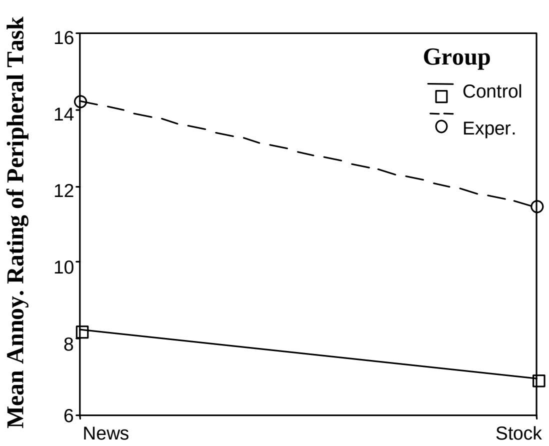

Figure 3: The mean level of annoyance experienced by a user when performing a peripheral task. The annoyance ratings are graphed as a function of peripheral task category and group.

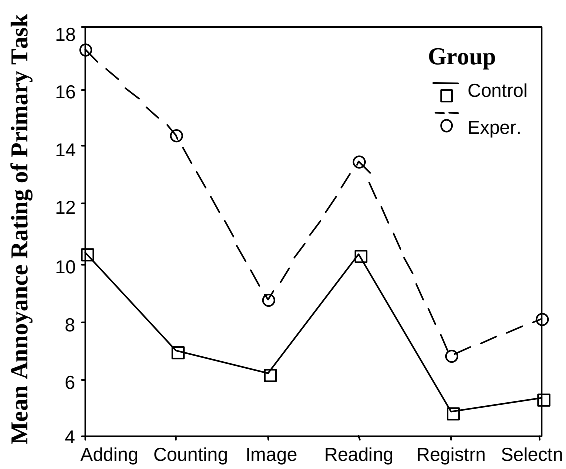

Figure 4: The mean level of annoyance experienced by a user when attending to a peripheral task during (or just after completing) the primary tasks from each category.

rating (F(1,96)=2.98, p<.09) nor were there any interactions (F(1,96)=.42, p<.52).

The primary task category annoyance ratings were analyzed using a full-factorial ANOVA with primary task category and group as factors. These ratings are graphed in Figure 4.

The results further support our second hypothesis. Whether a user was presented with a peripheral task during or just after a primary task (group) had a main effect on the level of annoyance experienced by a user for performing that primary task (F(1,288)=36.49, p<.00). The category of primary task being performed also had a main effect on the level annoyance experienced by a user (F(5,288)=14.55, p<.00). No interactions were present in the data (F(5,288)=2.10, p<.07).

H3 – Level of annoyance experienced by a user in the experimental group

The primary task annoyance ratings from the experimental group were analyzed using an ANOVA with primary task category as the factor. These annoyance ratings are graphed in Figure 4 as the upper line. The results are consistent with our third hypothesis. The category of primary task had a main effect on the level of annoyance experienced by a user due to an interruption (F(5,144)=11.35, p<.00).

H4 – Level of annoyance experienced by a user in the control group

The primary task annoyance ratings from the control group were analyzed using an ANOVA with primary task category as the factor. These annoyance ratings are graphed in Figure 4 as the lower line. However, the results are inconsistent with our fourth hypothesis. The primary task category *did* have a main effect on the level of annoyance experienced by a user even when the peripheral task was presented just after the completion of the primary task (F(5,144)=4.61, p<.01). This result suggests that the mental stress caused by a task is not immediately released upon the completion of that task.

Together, the analysis of the annoyance ratings demonstrate that the level of annoyance experienced by a user due to an interruption depends on both the category of task being performed and the time at which that interruption is presented.

4.3 H5: The Effect of Interruptions on Anxiety

The anxiety difference for a user was computed by subtracting the "before" measure from the "after" measure. The mean anxiety differences for each group are graphed in Figure 5.

The anxiety differences were analyzed using a t-test and the results are consistent with our fifth hypothesis. The mean increase in anxiety for a user in the experimental group was significantly greater than the mean increase in anxiety for a user in the control group (F(1,48)=5.12, p<.03). The results demonstrate that a peripheral task causes a greater increase in anxiety when it is presented during a primary task than when it is presented just after the completion of that task.

4.4 The Effect of Interruptions on Perceived Task Difficulty

The primary task difficulty ratings were analyzed using a full-factorial ANOVA with primary task category and group as factors. These ratings are graphed in Figure 6.

A user did perceive the primary task categories to be of varying difficulties (F(5,288)=44.68, p<.00).

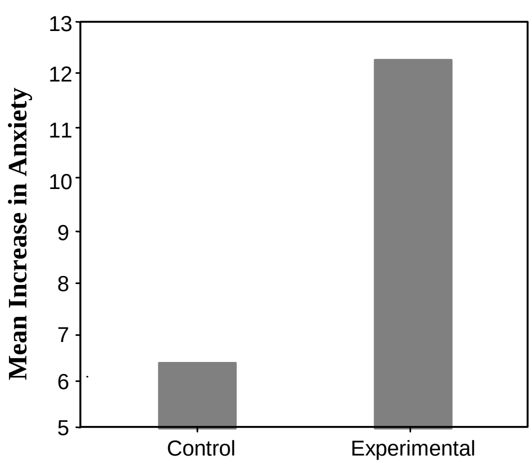

Figure 5: The mean increase in anxiety for users within the control and experimental groups.

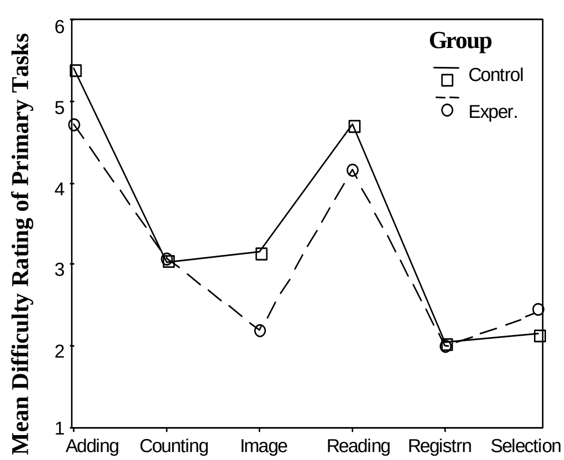

Figure 6: The mean perceived difficulty rating for a primary task, graphed as a function of primary task category and group.

Whether a user was presented with a peripheral task during or just after the primary task (group) also had a main effect on the perceived level of task difficulty ($F(1,288)=4.76$, $p<.03$). No interactions were present in the data ($F(5,288)=1.66$, $p<.14$).

These findings indicate that interrupting a user already engaged in a task causes him to perceive that task as being more difficult to complete than when that same task is not interrupted. A comparison of Figures 4 and 6 suggests a possible correlation between a user's rating of task difficulty and his level of annoyance experienced due to an interruption. A regression analysis using annoyance as the dependent variable and perceived difficulty as the independent variable shows a linear relationship between them ($F(1,298)=62.96$, $R=.42$, $p<.00$). This indicates that the more difficult a user perceives a task to be, the more annoying it is to be interrupted during that task.

5 Discussion and Future Work

Just as one person tries not to unnecessarily interrupt another, an application should not unnecessarily interrupt a user already engaged in a task. Not only is interrupting a user rude behavior, but as this work demonstrates, it also has a disruptive effect on a user's task performance and emotional state.

The primary implication of this work is that we need to build a system, an *attention manager*, which manages user attention among applications that are competing for it. An attention manager would first observe or predict an opportune moment for gaining user attention and then notify the next waiting application. An opportune moment can be defined as a period of low memory load occurring within a user's task sequence such as at a task boundary or during a delayed system response.

By delaying its interaction with the user until a more opportune moment arises, an application bypasses, or at least mitigates, the disruptive effects of an interruption. Although building such a system poses many new research challenges, it is not infeasible. For example, (Horvitz, 1999) built a system that learned the temporal patterns of a user interacting with an email system. The learned behavior was then used to predict when the user would finish reading the current email message and delayed any agent interaction with the user until that time passed.

In addition to implementing an attention manager in the future, we would also like to extend our experiment to examine the effects of:

- Multiple task interruptions (frequency) on a user's task performance and emotional state.
- An audio display interruption on a user's task performance and emotional state.
- An interruption on other tasks such as viewing a video or listening to an audio segment.

6 Conclusion

In this paper, we described an experiment measuring the effects of an interruption on a user's task performance, annoyance, and anxiety in the user interface. The key findings of this work are that (i) a user performs slower on an interrupted task than a non-interrupted task, (ii) the level of annoyance experienced by a user depends on both the category of primary task being performed and the time at which a peripheral task is displayed, (iii) a user experiences a greater increase in anxiety when a peripheral task interrupts her primary task than when it does not, and (iv) a user perceives an interrupted task to be more difficult to complete than a non-interrupted task. The implication of these results is that we need to build systems, such as an attention manager, which help manage user attention among competing applications, thus mitigating the effects of unnecessarily interrupting a user.

Acknowledgments

We would like to thank Tracy VanDemark for implementing the primary and peripheral tasks, and for quickly responding to our numerous requests for changes.

References

Bailey, B.P., J.A. Konstan, and J.V. Carlis. Measuring the Effects of Interruptions on Task Performance in the User Interface. *IEEE Conference on Systems, Man, and Cybernetics,* 2000.

Bailey, B.P., J.A. Konstan, and J.V. Carlis. Adjusting Windows: Balancing Information Awareness with Intrusion. *Proc. HFWeb*, 2000.

Boucsein, W. Psychophysiological Investigation of Stress Induced by Temporal Factors in Human-Computer Interaction in the Workplace. *Amsterdam: North-Holland*, pp. 163-182, 1987.

Gunn, W.J., T. Shigehisa, J.L. Fletcher, and W.T. Shepherd. Annoyance Response to Aircraft Noise as a Function of Contextual Effects and Personality Characteristics. *Journal of Auditory Research.* 21 (1): 51-83, 1981.

Horvitz, E. Principles of Mixed-Initiative User Interfaces. *Conference Proceedings on Human Factors in Computing Systems*, 1999.

Johansson, G. and G. Aronsson. Stress Reactions in Computerized Administrative Work. *Journal of Occupational Behaviour* 5 (1984): 159-181.

Kreifeldt, J.G. and M.E. McCarthy. Interruption as a Test of the User-computer Interface. *Proc. 17th Annual Conference on Manual Control.* JPL Publication 81-95, pp. 655-667, Jet Propulsion Laboratory, California Institute of Technology, 1981.

Maes, P. Agents that Reduce Work and Information Overload. *Communications of the ACM* 37 (7): 30-40, 1994.

Maglio, P. and C.S. Campbell. Tradeoffs in Displaying Peripheral Information. *Conference Proceedings on Human Factors in Computing Systems*, 2000.

Miedema, H.M.E and H. Vos. Demographic and Attitudinal Factors that Modify Annoyance from Transportation Noise. *Journal of the Acoustical Society of America* 105 (6): 3336-3344, 1999.

Nass, C., J. Steuer and E.R. Tauber. Computers Are Social Actors. *Conference Proceedings on Human Factors in Computing Systems*, 1994.

Spielberger, C.D. Manual for the State-Trait Anxiety Inventory (Form Y). Mind Garden, 1983.

Zijlstra, F.R.H., R.A. Roe, A.B. Leonora, and I. Krediet. Temporal Factors in Mental Work: Effects of Interrupted Activities. *Journal of Occupational and Organizational Psychology* 72 (1999): 163-185.