

IM Waiting: Timing and Responsiveness in Semi-Synchronous Communication

Daniel Avrahami

Intel Research Seattle
1100 NE 45th St., Seattle Wa 98105, USA
daniel.avrahami@intel.com

Susan R. Fussell, Scott E. Hudson

Carnegie Mellon University
5000 Forbes Ave., Pittsburgh, PA 15213, USA
{sfussell, scott.hudson}@cs.cmu.edu

ABSTRACT

Responsiveness, or the time until a person responds to communication, can affect the dynamics of a conversation as well as participants' perceptions of one another. In this paper, we present a careful examination of responsiveness to instant messaging communication, showing, for example, that work-fragmentation significantly correlates with faster responsiveness. We show also that the presentation of the incoming communication significantly affects responsiveness (even more so than indicators that the communication was ongoing), suggesting the potential for dynamically influencing responsiveness. This work contributes to a better understanding of computer-mediated communication and to the design of new tools for computer-mediated communication.

Author Keywords

Interpersonal communication, responsiveness, availability.

ACM Classification Keywords

H5.3. Group and Organization Interfaces: CSCW; H5.2. Information interfaces and presentation: User Interfaces.

INTRODUCTION

Each day, the average office worker is faced with an abundance of incoming communication, some of which is dealt with immediately, some answered as time permits, and some ignored altogether. To decide whether to engage in communication, people quickly weigh a multitude of factors, such as the cost of postponing their primary task, the perceived benefit of the communication to themselves and to the initiator, and their ongoing relationship with the initiator. When such communication is mediated by technology, a recipient can choose not only whether to engage in communication, but also *when* to engage in the communication. This *responsiveness*, or the time until a

person responds to communication, can affect the ongoing dynamics of a conversation as well as participants' perceptions of one another.

In interviews presented in [28], for example, email users described modifying their responsiveness in order to project different "responsiveness images", in an attempt to convey to the sender their availability, as well as their perception of the importance of the email. However, previous analyses of people's ways of dealing with incoming communication have typically drawn coarse distinctions in responsiveness (e.g., immediate response, delayed response, and non-response). For example, Dabbish et al. [10] examined whether or not participants reported responding to an incoming email as a function of their assessment of the messages' importance and their social relationship with the sender.

In email, broad distinctions between categories of responsiveness make sense. Email is generally used asynchronously and an instantaneous response is rarely expected. Face-to-face conversation is altogether different—social norms demand responsiveness to any attempts to initiate communication; even minimal delays in responding can lead to unwanted attributes of rudeness or inattention [22]. In synchronous and semi-synchronous computer-mediated communication, people may likewise possess fine-grained expectations about when a response should arrive, and draw unwanted conclusions if the response is delayed.

Responsiveness to communication is particularly interesting in semi-synchronous text-based media, such as chat and instant messaging (*IM*). *IM*, typically run as a computer program, allows users to exchange short textual messages ("instant messages") with their list of contacts (or "buddies"). The lightweight nature of these media allows conversation to range from synchronous, with rapid exchanges of messages, to asynchronous, with hours and even days passing between messages [24] (earning *IM* its description as "semi-synchronous"). Furthermore, this semi-synchronous nature of *IM* allows users to multitask while engaged in communication. In Isaacs' study [16], users multitasked during 86% of the conversations, and they switched in and out of the *IM* window approximately 4 times per conversation (to do other work, or to hold other *IM* conversations).

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

CSCW'08, November 8–12, 2008, San Diego, California, USA.
Copyright 2008 ACM 978-1-60558-007-4/08/11...\$5.00.

Unlike other synchronous communication media (such as face-to-face or telephone), responsiveness in semi-synchronous communication allows real-time decision-making with every turn in the conversation (e.g., whether to click on the window to open the message, whether to respond). Indeed, IM users can choose whether or not, and how rapidly, to respond to each and every message in the conversation (regardless of whether it initiates a new conversation or continues an existing one). To understand responsiveness in semi-synchronous communication, it should thus be examined not only at the beginning of a conversation (as would responsiveness to an incoming phone call), but with each exchange in the conversation.

In our own prior work on responsiveness in IM [3], we looked only at responsiveness to the start of IM sessions – similar to responsiveness to an incoming phone call. While the predictive models described in [3] offer interesting opportunities for the design of communication tools, the focus on the first messages in new sessions offer qualitatively different insights than we provide here.

In this paper, we analyze a large IM dataset to provide a deeper understanding of responsiveness and the factors that govern it: How does the user’s ongoing activity (or activities) affect his or her responsiveness to incoming (and potentially interrupting) communication? Will responsiveness vary based on who sent the message? Will people respond at different speeds during different parts of the day? How does the content of the communication affect the user’s responsiveness to it? Will responsiveness, when the communication is already ongoing, differ from responsiveness to attempts to initiate new communication?

We present an analysis of responsiveness to IM at precisely this level of granularity. Specifically, we attempt to answer the following research question:

How do context, message characteristics, and demographic characteristics affect a user’s responsiveness to incoming communication?

We present, for example, findings showing that faster responsiveness is significantly correlated with increased mouse movement, keyboard activity, and frequent transition between computer-applications (possibly due to work-fragmentation). We also show that the presentation of the incoming communication significantly affects responsiveness (even more so than indicators that the communication was ongoing), suggesting the potential for tools to dynamically influence responsiveness.

In the remainder of this paper, we first present a preliminary survey that guided our study design. We then review the factors we predict will influence responsiveness to IM messages, followed by a description of the data collection mechanism and the collected data. Next, we present the full list of measures that were extracted and investigated. We then describe the steps taken to prepare the data for analysis, followed by a detailed description of




Figure 1. Responsiveness to incoming communication.

the analysis method and our findings. We conclude by discussing how this work contributes to a better understanding of computer-mediated communication and to the design of new tools for computer-mediated communication.

Formative Survey

In order to guide our study, we conducted a short formative survey. We asked 22 professionals (4 females and 18 males) of different professions (programmers, designers, a web developer, researchers, a usability expert, and a homemaker), to complete a short survey in which they described their use of IM. Nineteen of the respondents had used IM for over three years, and the remaining three had used it for over a year. Fifteen of the respondents used IM both at home and at work, three used IM at work but not at home, and four used IM at home but not at work.

Among other questions, we asked respondents to tell us whether or not they would read an incoming message and whether they would reply to it if they were working towards a deadline on a work-related electronic document. In their answers, respondents brought up various elements that would affect their behavior, such as their engagement in, and the state and importance of their ongoing task. They also cited the identity of the sender and their relationship with them. Finally, they reported that, once having read the message, its content and importance played a role in their decision whether to respond. The respondents’ descriptions of their decisions whether to engage in communication guided our data collection and analysis. Some of their responses are brought in the next section.

Responsiveness and Context

“It depends on...how engaged I am with the task I’m doing. If I’m really engaged, they get to wait a while.”

Similar to an incoming phone call, an incoming instant message finds the user in some particular context that may affect their responsiveness to the communication. Furthermore, context may affect responsiveness such that it changes from message to message within the same conversation. (In phone calls, responsiveness is most interesting in the time until a call is accepted.) Multitasking when engaged in a phone call or face-to-face conversation can be difficult or inappropriate since high responsiveness is usually expected. Delays in responses are quickly noticed and negatively interpreted (e.g., [22], [26]). Unlike phone calls, the semi-synchronous nature of IM allows users to

multitask more easily while engaged in communication by using breaks between conversation turns to resume other tasks or attend to other communication. (This isn't to say that all delays in IM responses go unnoticed.)

In the work described here, we examine how responsiveness to incoming messages is affected by the user's context *prior* to the arrival of the communication (see Figure 1), looking specifically at the user's other ongoing computer activities, their recent and ongoing IM activity, and global context including the day of the week and time of day. While we focus on the context into which a message arrives, related work examined a user's behavior *after* the arrival of an incoming email or instant message, showing an increase in behaviors associated with task suspension (such as document-saves and paragraph completions) [15].

Responsiveness and Computer Activity

When looking at a user's context as represented by the user's activities on the computer, one may find the user engaged in greatly varied activities. For example, the incoming message may find the user working on a complex programming or design task that requires their undivided attention, or may find them using the computer for messaging and other communication, or simply for listening to music. Prior field studies ([8][14][20]) reported that information workers spend only a limited time in a specific task context before switching to another task (either due to external or internal interruptions).

Responsiveness and General IM Context

"I usually reply right away – but if am busy I will try not to have a long conversation."

Due to the semi-synchronous nature of IM, users may find themselves engaged in more than one conversation in parallel. However, high levels of responsiveness to simultaneous communication may be difficult to sustain and may result in the user feeling overwhelmed (potentially since these communications compete for similar mental resources [23] and computer resources). Note that while the presence of other ongoing communication may reduce one's responsiveness, the *recency* of communication with others may be an indication of the user's receptiveness to communication.

Responsiveness and the Communication Partner

"...not all buddies are created equal. Some are my boss, others are really good friends and one is my wife. There are social consequences to not being responsive with certain people."

A number of factors that are associated with the sender of the incoming message may affect a user's responsiveness to that message. For example, the specific identity of the sender or the type of relationship the user has with them (e.g., co-worker, friend, etc.) may affect responsiveness to

IMs, as it does for email [10]. The time that has passed since the previous communication with this particular person may also affect responsiveness. On one hand, recent communication may suggest that the user will be fast to respond to further communication. On the other hand, users may be interested and curious about incoming communication from buddies with whom they have not communicated for a while.

The State of the Message Window (and User Preference)

"[my responsiveness] depends on how I am notified..."

When an incoming message arrives, it can be presented on a user's screen in a number of ways. A message appearing in a message window that is already open is likely to indicate that some communication with the buddy has previously started. (Note that it is possible that this communication thus far includes only outgoing or only incoming messages.) An open window can be either "in focus" as the currently active application or it can be "out of focus", that is, not the active application, in which case its taskbar icon will flash to alert the user.

If the message window was not already open, its method of appearance depends on the user's preference. In our study, users selected among three presentation options. In the first, a message window is automatically created and displayed on the desktop in front of all other applications (the system's default presentation method). In the second method, a message window is automatically created but appears minimized on the user's taskbar. Flashing of the window's taskbar icon notifies the user of the arrival of the message. The user must click on the taskbar icon in order to bring the message to the foreground. In the third method, the user is notified of the incoming message through a small (16x16 pixel) blinking icon in the corner of their screen. The user must then click on the small icon in order to make the message window appear on the desktop.

It is reasonable to assume different properties of each of these possible states of the message window will affect responsiveness. Prior research has shown that the presentation and timing of a secondary interrupting task significantly affects a user's performance on both the primary and secondary tasks (e.g., [21]). One would expect, since a message in an existing window likely indicates that communication is ongoing, that responsiveness would be faster (whether the window is in focus or not). On the other hand, the visibility, or *saliency*, of the message window is also likely to affect responsiveness. One would expect that messages in windows displayed on top of all other applications will receive faster responses than messages in windows that are "out of sight". Note that responsiveness may be slower to messages in windows that are not in focus (whether or not they are minimized) since they require additional user action to bring into focus (clicking on the taskbar icon or on the systray icon, using the keyboard to bring the message to the foreground, etc.).

Responsiveness and Message Characteristics

“...if it’s a quick question I would reply, if it’s just to talk, I’d wait until I was done...”

Finally, the content of the message and the content of the conversation to which it belongs are sure to have an effect on responsiveness. While a detailed examination of the content of messages and their relation to responsiveness is left for future work, we extracted and examined a number of simple and potentially relevant attributes including the length of the incoming message, whether it contained a question, whether it contained a URL, and whether it contained an emoticon. (Emoticons are combinations of characters, such as the famous :-) smiley face, that are often used in chat to express emotion.)

In related work, different linguistic features, extracted from Usenet messages were shown to significantly correlate with the likelihood of these messages receiving a response [6]. Similarly, elements of the content of email messages were shown to significantly correlate with the importance that users attributed to these messages, and in turn, the likelihood that they would respond to them [10]. Prior research also suggests that the length of an incoming message will affect responsiveness. This is not to suggest that it is the length of the message *per se* that will have this effect, rather that other factors that are manifested in message length (such as the complexity of the content or the courtesy of the communication) have an effect on communication. Isaacs et al. showed that the length of messages is affected by user’s frequency of use of IM [16]. Similarly, the relationship between IM users was shown to have a significant effect on the length of the messages exchanged, with significantly longer messages, on average, exchanged between work-contacts [2].

DATA COLLECTION

Our data were collected using a custom plug-in for a commercial IM client called Trillian Pro. This client enables users to connect to any of the major IM services from within a single application, allowing us to recruit participants without concern for the specific IM service they were using. Our plugin recorded IM events (e.g., messages sent/received, message window opened/closed) and computer events (e.g., keyboard, mouse, and window activity). The text of messages was not recorded unless we received specific permission from the participants. Otherwise, messages were masked so that we could still perform counts of characters and words (punctuation was left intact).

A copy of Trillian Pro was purchased for each participant and the data recording plug-in was installed. Our plug-in was written in C and implemented as a Dynamically-Linked-Library (DLL) that is run from inside Trillian Pro (the plug-in automatically starts and stops whenever Trillian Pro is started or stopped by the participant.) Participants were instructed to use Trillian Pro for all their

| Group | N | % male | Mean age | Messages |
|-------------|----|--------|----------|----------|
| Researchers | 6 | 50% | 40 | 7,290 |
| Interns | 2 | 50% | 34.5 | 10,343 |
| Startup | 3 | 33% | 32 | 34,670 |
| Students | 8 | 25% | 24.5 | 73,906 |
| Total | 19 | 37% | 31.6 | 126,209 |

Table 1. Participant groups used in the study with demographic characteristics, and total messages.

IM interactions for a period of at least four weeks. Participants received a small monetary compensation for their participation.

Participants

Using the mechanism described above, we collected a total of approximately 6,600 hours of recorded data, observing over 125,000 incoming and outgoing instant messages from 19 participants who communicated with a total of nearly 500 buddies (see [2] and [3] for other analyses of this data).

Our participants included eight employees at a large industrial research laboratory, three employees at a local high-tech startup, and eight graduate students. Of the researchers, six were full time employees (three first-line managers and three full-time researchers) and two were summer interns. All 19 participants used IM in the course of their everyday work (during their participation, each of the students was engaged in a number of group projects as part of their studies). . For confidentiality reasons, we were not allowed to record the text of messages from the researchers nor interns. We will refer to our groups of participants as the *Researchers* group (including the six full-time research employees), the *Interns* group, the *Startup* group, and the *Students* group (See Table 1).

Since we had complete knowledge of desktop and IM state of the participants but not of their buddies’, the analyses described next examined only responsiveness to incoming messages (rather than also examining buddies’ responsiveness to participants’ outgoing messages).

MEASURES

The set of measures examined in this work were computed from participants’ logs. The measures are grouped into three high level categories: context, message characteristics, and demographic characteristics.

Responsiveness Measure

For the purpose of our analysis, we define a user’s responsiveness to an incoming message from a buddy *as the time (in seconds) until the user sent an outgoing message to that same buddy* (Figure 1). This formalization allows us to examine responsiveness devoid of the exact content of the exchange since this was not always made available to us.

Context Measures

These measures represent the context into which an incoming message arrives and include global context (such as the time of day), the participant's other ongoing desktop activities and measures of their IM activity:

Global Context

- Day of the Week (Monday through Sunday)
- Part of Day (Morning: 6:00-11:30, Lunch: 11:30-14:30, Afternoon: 14:30-18:00, and Night: 18:00-6:00)

IM Context (Specific)

These measures describe elements pertaining to the conversation and the conversation partner.

- The identity of the buddy (Buddy ID)
- The relationship with this buddy as indicated by the user (Co-worker (senior), Co-worker (peer), Co-worker (junior), Co-worker (other), Friend & Co-worker, Acquaintance, Friend, Family, Significant-other, and Spouse)
- Time since the last outgoing message to this buddy (log-transformed)
- Time since the last incoming message from this buddy (log-transformed)
- Whether the message window already existed before the message arrived (New Window vs. Existing Window)
- Whether the message window was in focus, as the currently active application, or out of focus (Window In Focus vs. Window Out of Focus)

IM Context (General)

These measures describe the general status of participants' use of IM and their communication with other buddies.

- Online Status (Online, Idle, Be Right Back, Away)
- Length of time in current online status (log-transformed)
- Whether there are other IM windows open (Single Window vs. Multiple Windows)
- Time since the last outgoing message to a different buddy (log-transformed)

Desktop Context

These measures describe desktop activity (*prior* to the arrival of the message).

- Number of Window-Title Switches including both switching between different applications as well as changing between documents or web-pages in the same application (Principal Component Analysis (PCA) on log-transformation – see below)
- The amount of Keyboard activity (PCA on log-transformation – see below)
- The amount of mouse activity (PCA on log-transformation – see below)

- The type of the application that was most in focus in the two minutes prior to the arrival of the message (browser, email, word processing, IM client, presentation, etc.)¹

Message Characteristics

- The length of the message, in characters (log-transformed)
- Whether the message contains a question (0 vs. 1)
- Whether the message contains a URL (0 vs. 1)
- Whether the message contains an emoticon (0 vs. 1)

Demographic Characteristics

These measures represent elements that were fixed for each of the participants during their participation period.

- Group (Researchers, Startup, Interns, Students)
- Participant ID
- Gender (Female vs. Male)
- Age

PREPARING THE DATA FOR ANALYSIS

Before beginning the analysis, a number of steps were necessary to ensure that the analysis provided the most informative and accurate results.

Accounting for Differences in Duration of Participation

Since the data collected represent naturally occurring IM interaction, different message volumes were recorded from different participants. Furthermore, some participants voluntarily continued their participation beyond the required four weeks, again, resulting in differences in the amount of data logged from different participants. To avoid having the data of a small number of participants excessively influence the results of the analysis, we used only the data recorded from each participant's first 45 days of participation. This resulted in an average of 36 days per participant (Min=17, Max=45, SD=8.96) with 73,571 total messages from all the participants (37,547 incoming and 36,024 outgoing).

Handling Non-Response

240 incoming messages for which the participant did not send an outgoing message to the same buddy before completing the study were removed.

Normalizing Measures

The time until a response, which is the primary measure in our analysis, as well as a number of the explanatory measures (for example, the time since receiving a message from a different buddy), exhibit a distribution with a peak and a long tail. (Similar responsiveness distributions were

¹ Similar measures for other time periods were computed. However, in order to avoid singularity, it was necessary to include only one of these measures in the analysis.

reported in [17] in their investigation of responsiveness to email and online discussion boards.) To address this issue, we used a base 10 log-transformation on these measures.

Accounting for Dependence between Messages

It was important to consider that messages arriving in close proximity are not independent of one another. That is, two different messages arriving one after the other (even if they were sent by different buddies) are likely to find the user in a very similar state and to result in similar responsiveness.

Indeed, in our data, the responsiveness to a message was highly correlated with responsiveness to the previous message. The correlation between two consecutive messages from the same buddy was $r=.454$. Thus, in order to account for this lack of independence between the consecutive data points, we included in the analysis the user's responsiveness to the previous incoming message from the same buddy (often referred to as the "lag-1") as a control measure.

Reducing Measure Covariance with Principal Component Analysis

Measures of computer activity were computed for a set of time-periods prior to the arrival of a message, specifically window-title switches, mouse movement, keyboard activity, and the most used application in the 0.5, 1, 2, 5, and 10 minutes immediately preceding the arrival of the message. As expected, however, the correlation between measures computed for different periods is very high. For example, the correlation between keyboard activity in the 30 and 60 seconds prior to the arrival of a message is $r=.82$.

In order to prevent the covariance of these individual measures from adversely affecting the analysis, we created three new measures "summarizing" title switches (WinTitleSwitchesPCA), mouse events (log transformed) (MouseEventsPCA), and keyboard activity (log transformed) (KBCountPCA). This was done by conducting Principal Component Analysis (PCA) three times, keeping the first component from each.

Online Status and Responsiveness

Indicators of presence and explicit indications of availability are one of the unique and most important features of Instant Messaging. Through the online status of a buddy, users can tell, before initiating communication, whether a buddy is online and present at his/her computer, whether the buddy has been inactive for some time, or has indicated that he or she is occupied or busy (as noted previously, the distinction between presence and availability is too often blurred and ignored [4, 12]).

Our field data included messages arriving when participants were in different states of presence and availability, such as *Online*, *Idle*, *Away*, and *Be-Right-Back* (*brb*). Since our interest is in responsiveness in the absence of indications of inactivity or unavailability that are "known" to the system and buddies, we excluded from our main analysis data for

which the user either explicitly indicated unavailability or was indicated by the system to be inactive, leaving only data for which the user was in an *Online* status (32203 messages, or 86% of all the incoming messages). A separate analysis confirmed the significant relationship ($p<.001$) between a user's status and his/her responsiveness.

THE ANALYSIS

The analysis was done as a single mixed model analysis. Responsiveness, or the time until a response is sent (log-transformed) was the dependent measure. The full set of context, message characteristics, and demographic measures listed earlier were included as independent measures. The state of the message window was described with the Window (Existing vs. New) and the Focus (In Focus vs. Out of Focus) measures, and the 2-way interaction between them Window*Focus. We also included two 2-way interactions looking at the state of the message window and the presence of other message windows: MultipleIMWindows * Window, MultipleIMWindows * Focus.

ParticipantID and BuddyID were modeled as random effects. Further, since each participant belonged to only one participation group, ParticipantID was nested in Group. Similarly, since in most cases buddies appeared in only a single participant's buddy list, BuddyID was nested first in ParticipantID, then in Group. This analysis allowed controlling for differences that originate from participation groups or that originate from individual (or dyadic) differences.

RESULTS

The analysis found a large number of significant effects on responsiveness. To make the results easier to follow, we've grouped them by the different measure categories (context, message characteristics, and demographic characteristics).

Context

The analysis showed that a user's other ongoing activities had a significant correlation with their level of responsiveness to incoming messages.

Desktop Context

The main application that was used on the computer in the two minutes prior to the arrival of a message had significant effect on responsiveness to that message ($F[21,31486]=5.3$, $p<.001$). For example, using a development tool (e.g., Microsoft Visual Studio or the Eclipse IDE), a word processor, or PowerPoint, was correlated with significantly slower responsiveness to incoming messages. Surprisingly, while also considered a productivity tool, the use of a statistics tool (e.g., SPSS or JMP) was associated with significantly faster responsiveness.

The amount of keyboard activity prior to the arrival of the message also had a significant effect on responsiveness ($F[1,30237]=116.3$, $p<.001$), as did the amount of mouse activity ($F[1,29331]=19.9$, $p<.001$). The amount of

window-title switches had a significant effect on responsiveness ($F[1,29270]=5.3, p<.05$). In all three cases, greater work-fragmentation (i.e., longer mouse movements, more keyboard activity, or more title switches) was correlated with faster responsiveness (this finding is discussed further in the discussion section). One should keep in mind that these levels of computer activity occurred *prior* to the arrival of the message, not after its arrival.

IM Context (Specific)

In general, elements of the communication to which the incoming message belonged had significant effect on responsiveness. While significant differences were found in responsiveness to different specific individuals (shown through predictions of the random effect of BuddyID and confirmed through a statistically significant increase in adjusted r-square with the inclusion of BuddyID as a random effect in the model), surprisingly, the type of relationship did not have a significant effect on responsiveness ($F[9,177]=1.1, n.s.$).

The time since the previous message sent to the buddy didn't have a significant effect on responsiveness ($F[1,31750]=.19, n.s.$), however the time since the previous message received from the buddy ($F[1,31626]=482.3, p<.001$) did. Longer time since the previous message was associated with faster responsiveness.

The state of the message window when the message arrives had a large and significant effect on responsiveness. Participants were faster to respond to messages appearing in a window that already existed ($M=32$ seconds²; $F[1,31834]=148.8, p<.001$) than to those appearing in a new window ($M=69$ seconds), presumably because of ongoing communication in the existing window. Similarly, the salience of the message window resulted in significantly faster responsiveness when the window was in focus ($M=29$ seconds; $F[1,31807]=394, p<.001$) than when it was out of focus ($M=79$ seconds). The interaction of Window (Existing vs. New) by Focus (In Focus vs. Out of Focus) was significant ($F[1,31363]=16.3, p<.001$), with a smaller difference in responsiveness between existing vs. new window when the window was in focus. A pair-wise comparison showed that the salience of the message window has a stronger effect on responsiveness than whether the window was new or already existed.

Responsiveness to the previous message ("lag-1" control measure) was indeed highly correlated with responsiveness to the current message ($F[1,31847]=4190, p<.001$).

IM Context (General)

A user's IM context had significant effect on responsiveness. The length of time (log-transformed) that the user was in an online status had a significant effect




Figure 2. The effect of the interaction between Other IM Windows (Single vs. Multiple) and Focus (In Focus vs. Out of Focus) on responsiveness.

($F[1,22457]=118.2, p<.001$) with quicker responsiveness when the user hadn't been online for long. Similarly, the time (log-transformed) since the user sent a message to a different buddy had a significant effect on responsiveness ($F[1,28120]=10.5, p<.001$). Responsiveness was faster when communication with others was more recent.

The analysis found no main effect of the presence (or absence) of other IM windows on responsiveness ($F[1,31782]=.22, n.s.$). There was, however, a significant interaction between the presence of other IM windows and Focus (In Focus vs. Out of Focus) ($F[1,30802]=17.5, p<.001$; see Figure 2) and between the presence of other IM windows and Window (New vs. Existing) ($F[1,31830]=32.2, p<.001$). A planned comparison showed that when an incoming message arrived in a window that was out of focus, the presence of other message windows had a significant effect on responsiveness. Messages received a slower response when other IM windows were present than not ($M=85$ vs. $M=74$; $t(31748)=2.78, p<.01$). However, when the message arrived in a window that was already in focus, responsiveness was much faster and the presence of other IM windows did not show a significant effect ($M=27$ vs. $M=30$ seconds; $t(31589)=1.4, n.s.$). This finding is interesting and stresses the significant role of the salience of an incoming message on the user's responsiveness. Similarly, the presence of other IM windows did not show a significant effect on responsiveness when the message arrived in a new window, however it did show a significant difference when the window already existed. This finding and its implication for design are discussed later.

Global Context

As expected, Day of Week had a significant effect on responsiveness ($F[6,22215]=9.8, p<.001$) and so did the Part of Day ($F[3,22199]=13.3, p<.001$). (Similar results were reported in [5].) In our data, responsiveness was significantly faster during the morning hours and at night (both $M=44$ seconds) compared to responsiveness during lunch and afternoon (both $M=52$; $t(21462)=6.15, p<.001$).

² We report Least-Square Means (back-transformed to seconds).




Figure 3. Content and Responsiveness:
The significant effects of the presence of (a) a question, (b) a URL, and (c) an emoticon on responsiveness.

Message Characteristics

Message characteristics influenced how quickly participants responded to messages. On average, the longer the message, the faster it was responded to ($F[1,31615]=234.4$, $p<.001$). The time to respond was 36% shorter with every 10-fold increase in the length of the message. Participants were also faster to respond to messages containing questions ($M=38$ seconds) than messages that did not ($M=60$ seconds; $F[1,31799]=185.5$, $p<.001$; Figure 3a). In contrast, participants were slower to respond to messages containing a URL ($M=70$ seconds) than to messages without a URL ($M=32$ seconds; $F[1,31838]=41.4$, $p<.001$; Figure 3b). Finally, participants were only marginally significantly slower to respond to messages containing an emoticon ($M=50$ seconds) than messages that did not ($M=46$ seconds; $F[1,31840]=3.6$, $p=.06$; Figure 3c).

Demographic Characteristics

Responsiveness was also influenced by several of our demographic measures. Participation group had a significant effect ($F[3,14]=5.5$, $p=.01$). A pair-wise comparison showed that participants in the Students group were significantly faster to respond on average ($M=32$ seconds) than participants in the Startup group ($M=105$ seconds; $t(15)=3.3$, $p<.01$). Neither the responsiveness of the Students nor the Startup participants was significantly different from that of participants in the Interns group ($M=37$ seconds) or the Researchers group ($M=41$ seconds). While the age of the participant showed no effect on responsiveness ($F[1,15]=.07$, n.s.), gender had a marginal difference; women in our study responded to messages faster, on average than men ($M=39$ vs. $M=59$ seconds; $F[1,13]=4.1$; $p=.06$).

DISCUSSION

We have presented an in-depth analysis of factors affecting users' responsiveness to incoming instant messages. This analysis was performed on data collected from participants' computers in an unobtrusive fashion and without user intervention over extended periods. The findings show that many, although not all, of the measures examined had significant effect on responsiveness. By observing communication at the beginnings, middles and ends of conversations, this work enhances our understanding of

responsiveness beyond previous research examining responsiveness only at the beginning of communication (e.g., [1] and [9]).

Effects of Responsiveness on Ongoing Communication

The effects of each of the factors we have presented is often relatively small (seconds, not minutes). To what extent then do these effects matter? Because the semi-synchronous nature of IM allows for easy multi-tasking, we might anticipate that delays in responding would motivate senders to switch to another task, thereby creating a chain of delayed responsiveness (because the IM window is no longer in focus). As Figure 4 shows, this is indeed the case; slower responsiveness reduces the likelihood that the message window of the person waiting will still be in focus when the response arrives. Recall now our finding that messages that arrive in a window that is out of focus receive, on average, slower responsiveness. Taken together, we may conclude that even minimal delays in responding can influence subsequent delayed responding, thereby influencing the overall chronological pattern of a conversation.




Figure 4. The relationship between the time the user is waiting for a response and the probability that the message window is still the application in focus.

Work-Fragmentation

Users' ongoing computer activities prior to the arrival of a message significantly affected their responsiveness. This finding is in agreement with previous work on the

interaction between people's primary task and their performance (and choices) when attending to an interrupting secondary task. In the real world, incoming messages are none other than such interrupting secondary tasks (unless, of course, the IM communication was itself the user's primary activity). Our analysis also showed a significant effect of the type of application used by participants on their responsiveness (e.g., slower responsiveness when using a programming environment). This finding is consistent with Fogerty et al.'s findings that showed that features describing the computer applications recently used were significant predictors of self-reported interruptibility [11].

One of our most interesting findings is the significant inverse correlation between responsiveness and the user's *work-fragmentation*, reflected by amounts of mouse activity and frequency of application switching. When users display increased levels of mouse movement and switch between applications frequently *prior* to the arrival of a message, they are likely to respond faster to incoming messages. This finding suggests that users who are engaged in a task or tasks that involve frequent switching between applications are more receptive to incoming communication. Borrowing the terminology from Gonzales and Mark [14], it is possible that users who are in-between work spheres are more willing to engage in additional external tasks such as incoming communication. In general, we suspect that infrequent switching between applications is associated with a user devoting their attention to a single task, resulting in unwillingness to be interrupted and in slower responsiveness.

The Effect of Message Characteristics

Our results showed that message characteristics also influenced responsiveness. Messages containing questions are typically associated with the expectation of faster responsiveness as the asker is likely to be waiting for a reply. Not surprisingly then, our participants responded significantly faster to incoming messages that contain questions. In contrast, participants were slower to respond to messages containing a URL, possibly because these messages required the receiver to follow that URL to some website before responding. The analysis also showed a small relationship between the presence of an emoticon in a message and slower responsiveness. The common use of emoticons to indicate that the meaning of a message was correctly understood, typically does not assume the conversation-floor (just as when providing non-verbal feedback in face-to-face conversations). Responses to such acknowledgments may thus be slower (It is also possible that emoticons indicate a more relaxed style of conversation affording slower responsiveness).

Investigating the effect of other linguistic features on responsiveness, similar to that presented in [6] would be valuable.

Responsiveness, Norms, and Culture

Responsiveness to communication is very likely to be influenced by cultural and normative elements; a person who is busy and does not desire to communicate may still be coerced by organization norms and other forces to engage in the communication [13][25]. Socio-cultural effects can be seen in our analyses of demographic characteristics, which showed that both gender and participant group influenced responsiveness. However, we measured only a subset of possible influences. Indeed, cultural differences have been shown by prior research to result in differences in communication and the use of communication technology (e.g., [7][18][27]). Organizational norms, too, can have great impact on the use and adoption of communication technology [19]. Future work will need to look more closely at how these factors influence responsiveness.

Design Opportunities

From a practical perspective, our findings enable designers of communication tools to foresee potential problems and highlight design opportunities.

Consider our finding that an ongoing conversation in the background will receive slower responses than a salient message in a new conversation, and the use of this finding, for example, in the design of an IM-based expert help-desk system. For such a system, our result points to a potential problem; open requests for help may be neglected in favor of new ones merely for being in the background. This suggests the need to ensure the prominence of ongoing help requests when new ones arrive. Our finding regarding responsiveness to messages containing URLs suggests designing a seamless link between the conversation, a reference (the URL), and the external application (the browser), to improve users' ability to discuss external resources. Finally, our findings on work-fragmentation suggest the need to distinguish task from sub-task boundaries *across* applications to avoid messages disrupting high-level tasks.

CONCLUSION

In this paper we described results from an in-depth analysis of factors that affect responsiveness to incoming instant messages and discuss practical implications of our findings. While this work describes investigation of responsiveness in a single medium (IM), the general classes of measures that were investigated – context, message characteristics, and demographic characteristics – are not at all unique to IM, but generalize to other forms of interpersonal communication. We propose that it would thus be beneficial to investigate responsiveness as it is manifested in other media (and as different media interact). It would further be interesting to examine the change in the effect of measures of context and content when an incoming message is machine generated rather than as part of interpersonal communication.

ACKNOWLEDGMENTS

We would like to thank Robert Kraut, Judith Avrahami, Yaakov Kareev, Thi Avrahami, Aaron Bauer, Darren Gergle, and Laura Dabbish for their valuable comments and suggestions. This material is based upon work supported by DARPA under Contract No. NBCHD030010, and by the NSF under grants IIS-0121560 and IIS-0325351.

REFERENCES

1. Avrahami, D., Gergle, D., Hudson, S.E., & Kiesler, S. Improving the match between callers and receivers: A study on the effect of contextual information on cell phone interruptions. *Behaviour & Information Technology*, 26 (2007).
2. Avrahami, D., & Hudson, S.E. Communication characteristics of instant messaging: Effects and predictions of interpersonal relationship. In *Proceedings of CSCW 2006*. 505-514.
3. Avrahami, D., & Hudson, S.E. Responsiveness in Instant Messaging: Predictive models supporting interpersonal communication. In *Proceedings of CHI 2006*.
4. Begole, J.B., Matsakis, N.E., and Tang, J.C. Lilsys: Sensing Unavailability. In *Proceedings of CSCW 2004*.
5. Begole, J.B., Tang, J.C., Smith, R.B., and Yankelovich, N. Work Rhythms: Analyzing visualizations of awareness histories of distributed groups. In *Proceedings of CSCW 2002*. 334-343.
6. Burke, M., Joyce, E., Kim, T., Anand, V., & Kraut, R.E. Introductions and requests: Rhetorical strategies that elicit response in online communities. In *Proceedings of C&T 2007*.
7. Choi, B., Lee, I., Kim, J., & Jeon, Y. A qualitative cross-national study of cultural influences on mobile data service design. . In *Proceedings of CHI 2005*.
8. Czerwinski, M., Horvitz, E., & Wilhite, S. A diary study of task switching and interruptions. In *Proceedings of CHI 2004*. 175-182.
9. Dabbish, L., & Kraut, R.E. Controlling interruptions: Awareness displays and social motivation for coordination. In *Proceedings of CSCW 2004*. 182-191.
10. Dabbish, L., Kraut, R.E., Fussell, S.R., & Kiesler, S. Understanding email use: Predicting action on a message. In *Proceedings of CHI 2005*. 691-700.
11. Fogarty, J., Hudson, S.E., & Lai, J. Examining the robustness of sensor-based statistical models of human interruptibility. In *Proceedings of CHI 2004*.
12. Fogarty, J., Lai, J., & Christensen, J. Presence versus availability: The design and evaluation of a context-aware communication client. *IJHCS*, 61, 3 (2004).
13. Ghosh, T., Yates, J., & Orlikowski, W.J. Using communication norms for coordination: Evidence from a distributed team. In *Proceedings of ICIS 2004*.
14. Gonzalez, V., & Mark, G. "Constant, constant, multi-tasking craziness": Managing multiple working spheres. In *Proceedings of CHI 2004*.
15. Iqbal, S.T., & Horvitz, E. Disruption and recovery of computing tasks: Field study, analysis and directions. In *Proceedings of CHI 2007*.
16. Isaacs, E., Walendowski, A., Whittaker, S., Schiano, D.J., & Kamm, C. The character, functions, and styles of instant messaging in the workplace. In *Proceedings of CSCW 2002*. 11-20.
17. Kalman, Y.M., Ravid, G., Raban, D.R., and Rafaeli, S. Pauses and response latencies: A chronemic analysis of asynchronous CMC. *Journal of Computer-Mediated Communication*, 12 (2006).
18. Kayan, S., Fussell, S.R., & Setlock, L.D. Cultural differences in the use of instant messaging in Asia and North America *Proceedings of CSCW 2006*. 525-528
19. Kraut, R.E., Rice, R.E., Cool, C., & Fish, R.S. Varieties of social influence: The role of utility and norms in the success of a communication medium. *Organizational Science*, 9 (1998), 437-453.
20. Mark, G., Gonzalez, V., and Harris, J. No task left behind? Examining the nature of fragmented work. In *Proceedings of CHI 2005*. 321-330.
21. McFarlane, D.C. Comparison of four primary methods for coordinating the interruption of people in human-computer interaction. *Human-Computer Interaction*, 17, (2002), 63-139.
22. McLaughlin, M.L., & Cody, M.J. Awkward silences: Behavioral antecedents and consequences of the conversational lapse. *Human Communication Research*, 8, (1982), 299-316.
23. Miyata, Y., & Norman, D.A. (1986). Psychological issues in support of multiple activities. In Norman, D.A., and Draper, S.W. eds. *User Centered Systems Design: New Perspectives on Human-Computer Interaction*, 265-284. Erlbaum, Hillsdale.
24. Nardi, B.A., Whittaker, S., & Bradner, E. Interaction and outeraction: Instant messaging in action. In *Proceedings of CSCW 2000*. 79-88.
25. Rennecker, J., & Godwin, L. Delays and interruptions: A self-perpetuating paradox of communication technology use. *Information and Organization*, 15 (2005), 247-266.
26. Schegloff, E.A., & Sacks, H. Opening and closing. *Semiotica*, 8 (1973), 289-327.
27. Setlock, L.D., Fussell, S.R., & Neuwirth, C. Taking it out of context: collaborating within and across cultures in face-to-face settings and via instant messaging. In *Proceedings of CSCW 2004*. 604-613.
28. Tyler, J.R., & Tang, J.C. When can I expect an email response? A study of rhythms in email usage. In *Proceedings of ECSCW 2003*. 238-258.