

When Can I Expect an Email Response? A Study of Rhythms in Email Usage

Joshua R. Tyler, John C. Tang

Hewlett-Packard Labs, 1501 Page Mill Road, Palo Alto CA, 94304

Sun Microsystems Laboratories, 2600 Casey Avenue, Mountain View CA, 94043

joshua.tyler@hp.com, john.tang@sun.com

Abstract. A study of email responsiveness was conducted to understand how the timing of email responses conveys important information. Interviews and observations explored users' perceptions of how they responded to email and formed expectations of others' responses to them. We identified ways in which users maintain and cultivate a *responsiveness image* for projecting expectations about their email response. We also discuss other contextual cues people use to discover email responsiveness, which include using other tools such as the calendar and phone, accounting for the amount of work time overlap available, and establishing a pacing between email correspondents. These cues help users develop a sense of when to expect a response and when breakdown has occurred, requiring further action.

Anyone who uses email regularly has sent a message and wondered, "When will I get a response to this email?" Or, "How long should I wait for a response to this message before taking further action?" Beyond the content of email messages, the timing of when email is sent, when it is read, and when a response is received are all examples of rhythms of email activity that help users coordinate their email correspondence.

Previous work has demonstrated that people have rhythmic temporal patterns of activity in the workplace, and that these rhythms can help coordinate interaction (Begole et al., 2002). We wanted to extend this work by exploring what meaningful temporal patterns occur in the usage of email. Email is clearly a crucial and ubiquitous tool for office workers, and understanding the types of

rhythms that govern its use will help identify design implications for improving the effectiveness of email services.

Specifically, we explored the following questions about email rhythms:

- How do individuals decide when to read and respond to a message?
- How do individuals form expectations about how long it will take others to respond to an email?
- How do these expectations affect email behaviors?

Email Background

While email usage has been studied since at least the 1980s, email usage has continued to evolve, suggesting an ongoing need to study and understand this evolution. Over the years, email has evolved into a multi-purpose tool used for more than just sending messages. Some people use email to track their work and manage their current responsibilities (Mackay, 1988), and some people have “overloaded” email with additional functionalities, such as managing an address book, task list, or conversation archive (Whittaker and Sidner, 1996).

People have also come to develop different, personally unique, and well-established strategies for dealing with the organization and storage of email (Bälter, 2000; Ducheneaut and Bellotti, 2001; Whittaker and Sidner, 1996). Similarly, one might expect individuals to have particular strategies for deciding when to respond to email messages, when to expect responses, and how to make those decisions. One recent study explored these differences by comparing the ability to check email continuously versus only three times a day, and found that checking email less often was actually more distracting (Patterson, 2000).

Context also plays a role in email rhythms. An early study of email (Sproull and Kiesler, 1986) found that the relative lack of social context cues in email had implications for what kind of information was exchanged and how that was perceived. Advanced communications systems have used contextual hints and information to better facilitate mutual understanding in work practice (Greenberg, 1996; Milewski and Smith, 2000; Tang et al., 2001). The PRIORITIES project uses classification learning algorithms to help assign the relative expected importance of incoming email messages, and incorporates this information into the user interface (Horvitz et al., 1999; Horvitz et al., 2002). In our research, we wanted to focus on the temporal patterns of email usage to see if they suggest tools to help users manage their email correspondence.

Study Methodology

We used interviews and qualitative observations to study rhythms in email usage. This study was conducted in two phases using semi-structured, observation-based interviews as a way to elicit users’ perceptions and attitudes about email usage.

Phase I was an initial exploration that broadly surveyed e-mail rhythms and identified more focused issues on responsiveness to study in Phase II.

The subjects came primarily (37 out of 40 total individuals) from two large technology corporations (Sun Microsystems, Inc. and Hewlett-Packard). While both companies are headquartered in the San Francisco Bay area, they include tens of thousands of employees spread around the world. Hewlett-Packard (HP) had recently completed a merger with another company (Compaq), which sometimes came up in the interviews, as we noted differences in corporate culture regarding email practices. Three subjects were from the Stanford Graduate School of Business (GSB), and reflected email experiences from their prior work experience (4.5 years on average). Through the merger and through the business school students, we got a glimpse of email behaviors beyond the two main companies from which the subjects were drawn.

The subjects were regular email users and spanned a wide range of positions—research, development, sales, marketing, communications, administration, management, legal, etc. We note when our observations differ due to the organizational context of the subject's email use.

Interviews lasted from 40-60 minutes and were generally conducted in the subject's workspace. For most (31) of our subjects, the workspace was a typical office setting—in front of a computer, on a desk, in either an office or a cubicle. The three interviews with business students were conducted in a university computer lab, and six interviews with remote subjects were conducted over the phone with the subjects in their typical work environment.

While interviews are a good way to elicit users' perceptions on email usage, we were not able to verify to what extent these perceptions were substantiated by actual practice (i.e., we did not verify the user's responsiveness perceptions with actual time stamps on email messages). However, we believe these perceptions are valuable as they govern the users' email behavior.

A variety of email clients were represented, but most of the users primarily used Microsoft Outlook (20), CDE Mailer (9) or Netscape Messenger (8). While most of our observations concern features that are common to all email clients, we note cases in which users reacted to features specific to a client. The number of messages in the user's inbox ranged from zero to over 17,000, which reflected different strategies for managing email. Those with larger inbox sizes relied more on search rather than filing messages into folders. This wide range of inbox sizes was also found in previous studies (Mackay, 1988; Whittaker and Sidner, 1996).

Note that a significant proportion of email does not require a response. Many messages are broadcast or cc:'ed to users for informational purposes without requiring a response. However, such messages are usually dealt with quite easily. It is managing or tracking the messages that do require responses that demand user attention and cognitive effort. Our work focuses on these messages that elicit responses and how users manage the response process.

Phase I—Exploring Email Rhythms

Phase I was conducted with 16 subjects from Sun (7), HP (6), and the Stanford GSB (3). These interviews and observations were designed to broadly explore the concept of email rhythms—how users convey information through the timing of their email, and what they learn from the timing of email from their coworkers. We asked about:

- where, when, and how often email is checked
- how the handling of email is prioritized
- usage of advanced email features
- usage of other communication media, and how this interacts with email
- how absences or communication delays are handled
- perceptions of others' responsiveness, and how these are formed

Inbox Walkthrough

One technique we used was an “inbox walkthrough.” With the subject's email client open, we would step through his or her inbox, asking him or her to describe the relationship with the sender of each message. In some cases this walkthrough was done on the email outbox or a mail folder.

At first, we conducted the walkthrough on a per-message basis. As we discovered that email rhythms are based much more on relationships than isolated messages, we began asking users to sort the inbox by sender (to see many or all messages from a particular sender). We then focused the interview on the relationships represented by the messages—it became more of a “relationship walkthrough.”

Phase II—Focusing on Responsiveness Rhythms

In Phase II, we delved deeper into some of the topics identified in Phase I surrounding the rhythms of responding to messages. In particular, we focused on:

- how users decide when to reply to a particular message
- how long they expect the respondent to take to reply

To help explore these topics, we included people who work from home or at multiple locations, to see if they pay particular attention to email responsiveness. We also interviewed administrative assistants, because their role puts them in contact with a wide variety of people with different email behaviors. These subjects came from Sun (13) and HP (11) without any overlap with Phase I subjects.

Another modification we made to our walkthrough process was to ask the subject to choose relationships representing different levels of responsiveness (e.g., “very responsive” or “very unresponsive”). Finally, we constructed more focused questions on the topics identified in Phase I, which is why most of the quantitative data is only available for Phase II subjects.

Presentation of Quotes and Anecdotes

In both phases, interviews were audio-recorded and analyzed, along with our interview notes, to identify recurring issues. These issues are illustrated in this paper by representative anecdotes and quotes taken from both phases of the study. Where there was diversity in the responses, we offer multiple examples representing the various perspectives articulated. When possible, we also provide quantitative descriptions of our results, created by reviewing and coding our interview notes. The names used to refer to the subjects in this paper are fictitious, though we preserved gender with the substituted names.

Email Responsiveness and Rhythms

Our interviews showed that people had a clear sense of when to expect email responses from people based on how quickly they had responded in the past, and that they could form this expectation after just a few interactions. The time that users reported expecting an email response ranged from fifteen minutes to a few days. This varied depending on the urgency of the message, the correspondents involved, and the work culture. Expectations for quick responses (under an hour) reflected a high level of knowledge about the recipient's availability and behavior. In cases where little context about the recipient was known, the expectation was usually one day (mentioned by 16 out of 24 Phase II subjects).

... I would typically think it should be 24 hours. Even if they don't have the answer they should at least say, 'I got your email, it will require a little bit of research, and I'll get back to you by whenever.'

From our interviews, we identified seven main observations of email usage that we will discuss in this paper. We found that users:

- display typical patterns of response behaviors
- maintain a *responsiveness image*
- take advantage of contextual cues to explain responsiveness
- use email with other media
- use email *peri-synchronously* when quick replies are expected
- reciprocate the email behavior of others
- often experience apprehension when contacting a new email correspondent

Taken together, these observations demonstrate that, despite being considered an asynchronous communication medium, the use of email actually involves negotiation and coordination between senders and receivers. Current email clients do not support this negotiation that surrounds the use of email. By identifying and describing these observations in detail, we can begin to consider how to design email services that better support the email behaviors that have evolved in current use.

Response Behaviors

Most of our subjects (21 out of 24 in Phase II) check their email “constantly”—they have an email client open all day long (when at their computers), and are notified immediately of new email. One exception we found was Karen, an engineer who works from home, who turned off automatic fetching of new mail in her client:

... because it’s too distracting... So I said, ‘OK, forget it,’ when I decide that I want to read email, I’ll check it.

Some of the subjects who check email from home work in “offline” mode, where they explicitly download new email, rather than being automatically notified. This behavior is often the result of a technical limitation of not having an always-on Internet connection at home. But we also found that some people prefer to avoid the “distraction” of email when working outside the office:

- I’m always on when I’m here... Constantly. Except when I’m at home... I’m in an offline mode... I’m going to deal with email very differently...

- I don’t want to be bothered by email, because the reason why I’m working at home is because I want to concentrate on something specific that should not have anything to do with email.

We initially expected individuals to exhibit different types of basal email rhythms (e.g. Alice is a “fast emailer” whereas Bob is “slow”). While those differences exist, we generally found more significant timing themes on a per-organization, per-relationship, and per-conversation basis. For example, our subjects typically responded more quickly to:

- messages from people within their workgroup (per-organization)
- messages from people with whom they have a history of quick communication (per-relationship)
- messages in a continuing conversation thread (per-conversation)

Responsiveness Image

In the course of our interviews, we found a family of behaviors that reflect a user’s desire to project a specific image of the time between receiving and reading or replying to email. We call this projection the user’s *responsiveness image*.

For example, we observed:

- users sending short messages signaling their intent to reply before they had read a message thoroughly
- users citing the perception of availability that they desired others to have of them as one of the most significant factors in prioritizing their responses
- a user who avoids opening read-receipt¹ marked messages until she has a reply prepared

¹ A “read-receipt” is a feature in some email clients that automatically notifies the sender of an email with a receipt when the message is read. The recipient typically is not aware that this receipt has been generated.

- users who became quite upset upon learning that read-receipts might be giving out information on their email reading habits of which they were not aware

These observations provide evidence of how individuals manage their responsiveness image. They evoke Goffman's (1959) notion of interaction as a performance, designed to leave specific impressions on its observers, consistent with the objective of the actor. In our interviews, we observed that people actively craft their image through the way they respond to email. People go to special lengths to project a certain type of responsiveness, in certain situations. While most "projecting" reflects an attempt to appear more responsive, there are also cases where people projected a less responsive image, as shown in the following examples from our interviews.

Bob, a lawyer, is meticulous about always acknowledging email messages immediately. If he does not have the time or resources to respond properly, he will at least send back a message signaling his receipt of the message, and intent to work on it soon. In this way, he projects a high level of responsiveness.

Carol, an administrative assistant for a high-level executive, will sometimes deliberately delay responding to an email to project a degree of inaccessibility. In response to requests of her manager that are low priority, she says:

I don't want the people to think that they can get an immediate, I'll drop anything for you. ... people tend to think they can get Jerry [her manager] any time. I won't respond usually, not until the next day.

Diane is an office worker who typically responds very quickly to her email (usually within an hour). She had just begun an email relationship with her 11-year old niece where she was very deliberate about the pacing of her messages—she responds within a day.

It's something I want to keep open, and so I may, on purpose, spread out my responses a little bit, so that she doesn't think it's a chore to respond to me, and then we can have more of an interaction, actually, than we've been able to have in the past.

People attempt to project an image in order to counteract assumptions that people might otherwise make about responsiveness. For example, Cindy, an engineer who works at home twice a week explained:

I think I check mail more often when I'm at home. Because I feel I'm out of the group, say if they suddenly set up a meeting, I think I should know, I should keep checking.

Many of our interviewees (11 out of 13 at Sun) claimed to use a simple responsiveness-maintenance technique: If they are unable to respond in a "reasonable" amount of time (usually about 24 hours), they send a message saying they need more time. In this way, they are able to preserve a responsive image.

I usually will get back to email or people the same day, ... even if it's just to say 'I need to look into that and I'll get back to you as soon as I can.' Because I kind of think it's rude if you ask someone for information and they just totally ignore you. It's one thing if they need time

to find out the answer, but it's nice to hear that, that they're in the office, that they received your message but just need some time to think about it.

Since users invest substantial effort to maintain a certain responsiveness image, they were frustrated to learn that their image could be undermined by features in their email client such as read-receipts. The default settings for Microsoft Outlook among the users we interviewed were that receivers were unaware that they were generating a read-receipt when they read e-mail messages. Thus, they did not realize that the senders could have more revealing information about when they read the messages than they were trying to project. This lack of control of what information others have of your activities can frustrate the maintenance of a responsiveness image and create privacy concerns.

We later learned of a related example (outside of our formal interviews) of a user who timed her email sending activity to manage when people could learn of her availability (an "availability image"). Rhonda, a high-level executive, likes to start her work early in the morning, in part to have undisturbed time before everyone else comes in to the office. Some of her work, however, involves collaborating with remote people who are in a time zone "ahead" of her. She discovered that if she sent email to those people when she first came in to work, she would often get incoming requests (by phone) since they realized that she was now available for contact. These requests disrupted her window of undisturbed work activity. Her solution was to compose responses as she went through her email at the beginning of the day, but to not send the messages until she was willing to be discovered as available.

Contextual Cues for Email Responsiveness

During the course of our interviews, we noticed that contextual cues were frequently used to help explain email responsiveness. These cues were often articulated when we asked why they thought people responded to email more slowly than expected:

... this week she happens to be in class. So she'll check her email on a break.

My manager ... he's very busy, at meetings all day.

I'm guessing, doesn't respond as quickly as I wish because, they're so busy they can't get to their email.

I think it depends upon what he is doing when he gets the email.

Comments in the interviews helped identify what contextual cues were being used to help interpret email responsiveness. Many of these observations suggest ways it might be helpful to convey these cues through email or by integrating with other tools.

Contextual Cues From Within Email

Besides tracking people's historical response rates, other contextual cues for response rate were conveyed through email itself. An important resource is the use of "auto-reply" or "vacation" messages. This feature creates an automatic response to any email received with a message that usually explains that the user is absent and sets an expectation of when they will be able to respond to the message. The value of these messages is described by an administrative assistant who complained about people who failed to use them:

You ask someone for information that they know, and you sit around waiting and waiting for them to get back to you, and you find out that they've been out of town.

In Phase II, we asked the subjects "How long of a work absence is required to cause you to use an auto-reply message?" The most common response was one day (8 out of 24), followed by two days (5) and one week (4). The other 7 respondents did not use an auto-reply mechanism.

One interesting use of the auto-reply feature was described by Ellen, a supply buyer. She typically takes a day to catch up on email after returning from an extended trip. During this time, she may be less responsive to new incoming mail. Her solution is to leave the "Out-of-office" auto-reply feature activated during this day after she returns (even though she is in her office), conveying to senders that they should not expect her typical responsiveness. We later learned of an example (outside of our interview study) of another user, Sam, who always leaves an auto-reply message on. The message explains that email is an unreliable way to reach him and his ability to respond to email is erratic, and directs people to contact him over the phone.

While auto-reply messages give an explicit explanation of why an email response might be delayed, these messages are only sent in response to receiving a message from someone. Thus, the sender would not know in advance about the delayed response, and might have composed the message differently if the delayed response could have been anticipated.

Another way context can be conveyed through email is through an email message itself. Several users mentioned that they would notify their work group via email when they expected to be out of the office (e.g., vacation, training). Others also mentioned that they could sometimes tell when an email correspondent was busy or in a rush through the writing style of the message. If the message had typographical errors, no capitalization or punctuation, incomplete sentences, or other indications of being hurried, that could convey that the email sender was busy, especially if that was not their usual writing style.

In a sense, a lack of an expected email response, if it becomes a recurring pattern, is also a contextual cue on email responsiveness. Many people (especially at Sun) were quick to suggest that they were not getting a response because their correspondent must be overwhelmed by email at the time. As the experience of struggling to manage email becomes more common, it becomes a

more likely reason to project upon others to explain a lack of a timely response. By contrast, the subjects at HP were less likely to report feeling overwhelmed by email, demonstrating how the corporate culture of email usage provides context for expectations about email responses.

To the extent that our ability to keep up with email varies over time, it might be valuable to convey some indicator of the current size of the recipient's email backlog. This might be measured in terms of the number of unread messages in the user's inbox relative to what is typical for that user. Knowing if the user currently has a large email backlog may help prospective email senders consider the most effective way of contacting that person at the moment. For example, it may suggest a different way of crafting a message to get the attention needed, a different medium of contacting the person (e.g., using the phone instead of email), or waiting for a different time to send the message.

Contextual Cues From Outside Email

Besides cues from email itself, users cited a number of contextual cues from other sources that often helped explain delays in email response. Shared online calendars have been shown to be a valuable resource for coordinating group work (Palen, 1999), and several users mentioned them as a way of explaining email delays. Browsing someone else's calendar is one way of discovering that they are on vacation (especially if they do not use auto-reply vacation messages). At Sun, 9 of the 13 subjects described relationships and situations in which they expected a response within an hour, and sometimes within 15-20 minutes. Failure to receive a response within that time would prompt them to check the calendar of the correspondent. If her calendar showed that she had a meeting scheduled at that time, it would often explain the delay in responding, and suggest when she would return and have a chance to respond to the message.

Others mentioned that the knowledge of their correspondent's job role factored in to their expectation of response time. Certain jobs (e.g., programmers, administrative assistants) were typically associated with being at the desk most of the time, which would tend to indicate faster response times. Other jobs (e.g., managers) implied being in meetings a lot of the time and away from the desk (and thus unavailable for responding to email).

The recent emergence of using wirelessly connected laptops and mobile devices to read email during meetings represents another contextual cue that would modify the effect of meetings and other activities outside the office on email responsiveness. People who worked in the field were considered to be traveling or meeting with customers much of the time, thus limiting their access to email (and consequently their responsiveness).

The social context of email usage also provides cues about responsiveness. A few people mentioned reputations of email responsiveness that were shared among others:

I'm known by my bosses and by my employees to follow up through email probably, faster than anybody else around here.

... people would say, 'Send him email, he responds quickly.'

Thus, people not only formed expectations of when they would get email responses from those they corresponded with, they also shared their expectations with others to help form social reputations of email responsiveness.

Email behavior was also noted as a topic of social conversation that conveys expectations of response time. Two people, in particular, mentioned how others' complaints about "he's not responding to my email" would come up in conversation. In one case, however, the subject being interviewed was getting normal email responses from the person being complained about, reinforcing the observation that email response rates can differ for each pair-wise combination.

In another example, sharing among a team that no one was getting an email response from a particular person helped piece together a story that something was happening:

Once [our manager] was traveling ... but we didn't hear from him for a day. And we thought, you know, something's going on. There must be some bad meetings, something bad must have happened. It was a day. None of us, we checked, none of us had heard from him.

Email users currently draw upon contextual cues from email, other tools such as a calendar, other people's experiences, and other knowledge to help interpret email responsiveness patterns. Understanding how these cues are used could help identify what cues would be useful to convey through an interface to someone who is preparing to compose an email message. Supporting this kind of awareness between sender and receiver builds on the notion of social translucence (Erickson and Kellogg, 2000) by providing cues that allow social negotiation of appropriate action.

Email with Other Media

Email is used in conjunction with other communication media, such as the telephone, instant messaging (IM), and face-to-face contact. A common tactic we observed when a message was important or urgent was to send a voicemail in conjunction with an email. Typically, the email contained detailed data or attached documents, and the voicemail signaled the urgency of the message. Of Phase II subjects, 10 out of the 11 at HP, and 4 of 11 at Sun claimed to use this tactic (at HP, voicemail is used more frequently than at Sun, where email is predominant). Typically, the email was sent as the "primary" message, and the voicemail served as a pointer to the email, but a few subjects mentioned the converse.

One subject even mentioned leaving voicemail at both office and cell phone numbers. Two more subjects reported combining an email with an electronic page to convey the message's urgency. This tactic appeared to address the concern that the recipient may not notice or have access to the email message

quickly enough. However, using the phone or pager as an alert notification not only creates more work for the sender, but also adds to the load of messages (voice or page) that the receiver must deal with. This practice suggests the need for better ways of delivering or notifying recipients of urgent messages, perhaps through more integration with other communication media.

IM use was not common for our subjects. Only 4 of the 24 subjects in Phase II claimed to be current IM users. Two of them, however, described how they use IM in place of the telephone, not to replace email. They used IM in situations where they need an immediate answer (within a few minutes) or expect multiple rounds of dialogue, situations in which people have traditionally used the phone. Consistent with other studies of IM use (Nardi et al., 2000; Isaacs et al., 2002) our subjects preferred IM because it enables them to multitask—to continue with their other work in between dialogue exchanges, and to carry on multiple conversations at one time. For them, email is not as easily replaced by IM, since it is used for sending longer, detailed messages, file attachments, and messages with many recipients, which are difficult to accomplish over IM.

Asynchronous and Peri-synchronous Email

Email is often described as an asynchronous medium, and indeed, our subjects valued the ability communicate without actually having to catch the recipient at the moment. However, we also detected a distinct style of writing email *peri-synchronously*. By peri-synchronous, we mean sending messages nearly synchronously such that they are expected to receive a response shortly. Email correspondents who share a significant overlap of work time can exchange email peri-synchronously, allowing several exchanges to occur within a day. This style contrasts with completely asynchronous email usage, which might occur when emailing with international correspondents who are shifted in time by about seven hours or more. Such email will probably be read and acted upon while the sender is offline.

Most of the work-related email that our subjects talked about occurred within a peri-synchronous context. There was enough time overlap that correspondents could expect a response within a day, and would even afford iterations and discussion throughout the day. In this sense, any individual message could leave room for discussion or be incomplete, but the thread could resolve to a conclusion within a day.

However, email correspondents with international colleagues who were time shifted by seven hours or more exhibited a different pattern. Since they could not expect a response within a day, people talked about thinking more carefully about those messages and taking the time to be more clear, explicit, and complete. Gary, a California-based member of a field services team, commented on correspondence with people in Europe or Asia,

You have to get everything that you need down in one shot, in order to get the turnaround time... So instead of saying ‘I think it could be this, look here and tell me what that says,’ I will list out four or five possibilities, ‘Here’s what I think it probably is, if it’s not that it might be this and if it’s not that it might be this other thing.’

Furthermore, international email correspondence was more likely to encounter cultural differences in styles of usage. Several of the American email users commented that messages to their colleagues in India or Asia did not prompt as quick or complete a response as they had hoped. They were quick to recognize that this might be due to a different work culture.

Also, personal email has a different asynchronous style associated with it. People often do not expect responses to personal email within a day. They do not seem to expect such a high priority in responding to personal email, and think of it more like traditional letter writing where a response could be more leisurely.

In this way, we found that the context of using email asynchronously or peri-synchronously affected the way they composed their messages. Peri-synchronous use did not require catching the recipient at the moment, like using IM or the phone, but did afford the flexibility of interaction and discussion through exchanges during the day. Being aware of this context and the underlying reason for it (e.g., different time zones, work culture, or relative priority) is another cue that people use to set expectations of when they might get a response (which in turn affects how they compose their messages).

Reciprocity in Email Behaviors

The influence of reciprocity in governing human behavior is the subject of much work in social and evolutionary psychology. For example, Trivers (1971) found that the reciprocation of seemingly insignificant favors and courtesies is crucial to the “social contract” fabric of human society, and others have described the biological and economic foundations of such actions (Sigmund et al., 2002). We found that there is also a notion of reciprocity between participants in an email conversation. This reciprocity appeared strongly in exchanges that evolved into IM-like exchanges of 1-2 line messages every few minutes, as well as in longer-term interactions—when one of the participants replied more slowly than the other, subsequent messages back and forth typically proceeded at the slower pace.

Kate, a business development leader, described a former manager who was very responsive—no matter where he was, he would send a clear, useful response to your message within 24 hours. This behavior inspired special efforts to reciprocate:

His responsiveness made people very responsive to him... What I can tell you is what he did for the people who worked with him—incredible loyalty... You would work around the clock for this guy... You could not let him down. This was a problem. Because when you would ask him something, he would come back to you. So when he would ask you something ... you had to get back to him.

Kate also described a former peer who was “very unresponsive.” He would rarely respond to email, instead requiring catching him on the phone. As a result, Kate “calibrated” her email relationship to his, becoming less responsive. Reciprocity is important to Kate: “People react to the way you treat them.” If a peer is unresponsive, then “he’s going to get the same deal.”

I’ve discovered that people do adjust their responsiveness... I’ve seen that happen. Because I try to be very responsive to people, and I expect that same responsiveness. So if they don’t match up, then I’m going to change my responsiveness level. It’s a real tit for tat kind of a mechanism.

Taking stock of how well you have reciprocated in responding to others can lead to a notion of watching out for a “responsiveness debt.” For example, Patty, a solutions manager, will frequently scan the last week’s worth of email in her inbox to assess whether or not she’s maintained her responsiveness image on a per-person basis:

I also take a look a lot at what’s happened in the last seven days, and did I respond to everything that’s been going on... I do it on a daily basis... In the last seven days, I’ve gotten three messages from Annie, seven messages from John, that kind of thing, what were those messages about, and am I still keeping up with all that.

We also noticed several instances of the same two people carrying on two or more overlapping conversations, each with distinct timing patterns. These observations suggest that reciprocity can be a *per-conversation* as well as a *per-relationship* phenomenon.

Another type of reciprocation we observed was the choice of media. Several subjects try to learn which medium their correspondents prefer, and will contact them in that way. Lois, a marketing manager, personally dislikes the phone and strongly prefers email or IM. However, she says:

The first thing I ever find out about anybody I work with is how they communicate. That’s just really important to me.

By adapting and reciprocating communication preferences in this way, she increases her chances of getting timely responses.

Anxiety in New Email Relationships

We found that when writing to a new email correspondent, the sender will often have anxiety about whether and when the email will be read, and if and when a reply can be expected. People typically develop these expectations over the course of an email relationship, but when there is no such foundation, they do not know what to expect. Even if the recipient is a personal acquaintance, there is a degree of anxiety because he or she is not an *email acquaintance*.

Conclusions

By treating email as an asynchronous communication medium, current email tools provide little opportunity for users to exchange cues about when to expect a response or when are good times to be reached by email. Yet, our observations demonstrate that there are subtle ways in which email is a mutually negotiated medium: People project responsiveness images to each other, they use email perisynchronously, and they calibrate their email behaviors to mirror the rhythms of their correspondents. Finding ways to convey these contextual cues could provide better support for the mutual negotiation between senders and receivers, which would in turn enable the design of richer, more effective email services.

Mutually Setting Expectations

One perspective from which we can summarize our findings is in terms of the *recipient's burden*. Without any cues to negotiate between senders and recipients, much of the burden of managing email currently falls on the recipient. Recipients need to track whether or not they have responded to email messages in a timely way, and may also need to integrate notifications from email, voicemail, pagers, and other media. Without any mechanism for negotiating how much or what kind of email a recipient would like to handle, they are further saddled with the task of screening unwanted spam and managing incoming mail with filters, agents, tedious filing strategies, etc.

From the sender's perspective, our observations show that previous experiences in corresponding with a person enable the sender to formulate expectations about when a message he or she sends will be read, and more importantly, responded to (in cases where a response is desired). Distinct from this *response expectation* is a *breakdown perception*—when the sender believes that something has gone wrong, and will take further action. We observed a number of factors that influence these expectations.

We asked people about both expectations—when subjects expected an email response and how long they would wait before following up—and found that they expressed distinct time frames in their responses. For example, some expect a response within a half hour or an hour, but will wait a day until walking to the recipient's office to find him or her. Thus, they expected a response within an hour, but would wait until a day to consider it a breakdown that required follow-up. We saw a similar pattern in users who expected a response in fifteen minutes; after waiting for an hour with no response, they followed up with a telephone call or by checking the recipient's calendar to see if he or she was unavailable due to a meeting. Again, the response expectation (fifteen minutes) was distinct from the breakdown perception (one hour). Figure 1 shows a visualization of these expectations, and puts them on an abstract timeline.

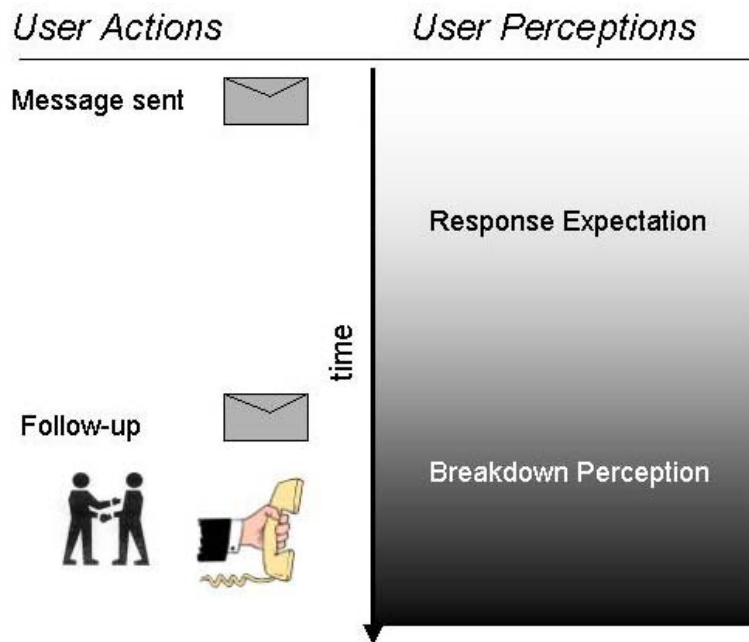


Figure 1. A visualization of the “expectation-to-breakdown timeline.” When the breakdown perception is formed, the sender will initiate a follow-up action, such as a phone call, another email, checking their calendar, or a face-to-face visit.

This distinction between response perception and breakdown also governed how people maintained their responsiveness image. For example, most of our subjects said that they typically expect a response within 24 hours, and thus most people set auto-reply vacation messages if they were absent for a day. But some would only activate an out-of-office auto-reply message for absences of *two days* or more. This longer latency is calibrated to the breakdown perception, rather than the response expectation. In this case, they are not interested in maintaining a typical response time, but one that is just good enough to prevent the sender from worrying about not receiving a response. This distinction also governed when users will send an email alerting the sender that they need more time. This message is sent after the response expectation point (e.g., one day) has passed, but before the breakdown perception (e.g., two days).

Clearly, the expectation and breakdown points are not the same for every email, or even for every person. We observed the following influences on these expectations:

- The recipient—expectations of responsiveness vary from person to person
- Location of recipient—the time shift between sender and recipient may increase the time expected for a response
- Urgency of topic—people expect attention to be paid to urgent messages, which are often signaled with markers such as “URGENT:” in the subject field or priority flags

- Combination with voicemail—as described previously, important email messages are often accompanied by a voicemail, to convey the urgency of the request

If an email reaches the breakdown point without a response, the sender will either give up on the request or follow it up. The most common follow-up methods we observed were:

- contacting the recipient through different means, typically a telephone call or in-person visit (when possible)
- sending another email, with more urgency or adamancy in the subject line
- checking on the calendar or other context of the recipient, to see if there are explanations for the lack of responsiveness

The expectation-to-breakdown timeline is a useful framework for understanding how individuals track the progress of their incoming and outgoing email requests and correspondence.

Design Implications

Besides being a useful analytic framework, this timeline could also be a helpful design concept for future email services that help users manage their email responsiveness. Perhaps an *email responsiveness service* could analyze a variety of measures of email behavior and abstract a representation of this timeline for individual messages, and this timeline could be visualized in email clients. Such a timeline could remind recipients of messages that they have not yet responded to within their typical response times. The timeline could also alert senders of messages for which they have not received a response within their typical expectations.

Such a timeline representation should be designed so that it integrates these various measures of email behavior into a simple graphic that users can easily use as a basis for a responsiveness expectation (rather than having to examine many different metrics and sources to make that judgement themselves). The timeline portrayal should also appropriately reflect that these response expectations are inexact, so as not to mislead users into a false sense of precise predictions of when to expect a response. Furthermore, the service should be designed to aggregate and abstract the data in such a way as to protect the privacy of the users. A useful depiction of the timeline that gives users a sense of when to expect an email response does not need to reveal the detailed measures used to create that representation.

Many of the factors needed to create such a timeline representation for email messages could be drawn from observations about an individual's current email usage. Logs of how quickly people respond to particular email senders over time could be analyzed to develop a response time prediction for a message and identify a range after which breakdown has probably occurred. These predictions

based on historical rhythms could also be modified by other inputs tracked by a dynamically updated awareness service. Just as current IM systems detect when users are logged on and how recently they have been active on their computer keyboard, this information could be tracked over time and used to develop a sense of how much opportunity a user has had to access and thus respond to email. That is, if an email recipient has been away from his keyboard more than usual, it might lengthen the time before a response should be expected without triggering an indication that a breakdown has occurred.

Another factor that could be tracked is the recipient's current backlog of email. If an email recipient currently has a larger than usual number of unread messages in her inbox, she may not be able to maintain her usual responsiveness. Aggregating and analyzing these and other existing data sources could provide a basis for constructing an expectation-to-breakdown timeline.

Taking current conditions into account to modify the expectation-to-breakdown timeline would be useful even to email correspondents that have expectations based on previous experience. For example, you may know someone to be a quick email responder, but if he is currently away from his keyboard and experiencing a large email backlog, the longer response expectation time depicted in the timeline would modify your default expectation. Thus, the timeline may suggest a different course of action (e.g., call instead of wait for an email reply or look for someone else to address the issue).

Another example application of such a timeline is in helping address the anxiety that many senders experience when emailing someone for the first time. Without any prior experience to establish an expectation, senders can face a dilemma of how long to wait for a response before deciding that follow-up actions are required. Appropriately conveying cues to set responsiveness expectations for someone to whom you are emailing for the first time could help smooth the initial contact through email. Of course, sharing cues about email responsiveness, especially with strangers who are emailing you for the first time, raises privacy issues. Users need to be aware of and maintain control of the responsiveness image that they present. More exploration is needed to understand the delicate balance between sharing useful contextual cues about email responsiveness and preserving users' control over their responsiveness image.

Our research identified some of the cues people use to form expectations of getting an email response. Faced with an impoverished communication medium such as email, people have developed many techniques for adding some of the richness of negotiation back to the medium through the use of timing and other responsiveness-related behaviors. Applying this understanding to the design of new email services that share the awareness of these behaviors among email correspondents, with proper respect to privacy concerns, could create a richer, more effectively negotiated communication channel.

Acknowledgements

We would like to thank Kyle Forster, Jason Heidema, and Randy Schwemmin for their involvement in this work, Jeanette Blomberg, Randy Trigg, and Pam Hinds for their guidance, Marie-Jo Fremont for calling in some favors, and all of our interview subject volunteers.

References

- Bälter, O. (2000): 'Keystroke Level Analysis of Email Message Organization', in *Proceedings of CHI 2000*, ACM Press, 105-112.
- Begole, J., Tang, J., Smith, R., and Yankelovich, N. (2002): 'Work Rhythms: Analyzing Visualizations of Awareness Histories of Distributed Groups', in *Proceedings of CSCW 2002*, New Orleans LA, November 2002, in press.
- Ducheneaut, N. and Bellotti, V. (2001): 'Email as Habitat: An Exploration of Embedded Personal Information Management', in *Interactions*, 8 (5), (Sept/Oct 2001), 30-38.
- Erickson, T. and Kellogg, W. (2000): 'Social Translucence: An Approach to Designing Systems that Mesh with Social Processes', in *Transactions on Computer-Human Interaction*, Vol. 7, No. 1, 2000, 59-83.
- Goffman, E. (1959): *The Presentation of Self in Everyday Life*, Doubleday, Garden City NY, 1959.
- Greenberg, S. (1996): 'Peepholes: Low Cost Awareness of One's Community', in *Companion Proceedings of CHI '96*, 138-145.
- Horvitz, E., Jacobs, A., and Hovel, C. (1999): 'Attention-Sensitive Alerting', in *Proceedings of UAI '99, Conference on Uncertainty and Artificial Intelligence*, Stockholm, Sweden, July 1999, Morgan Kaufmann: San Francisco, 305-313.
- Horvitz, E., Koch, P., Kadie, C., and Jacobs, A. (2002): 'Coordinate: Probabilistic Forecasting of Presence and Availability', in *Proceedings of the Eighteenth National Conference on Uncertainty and Artificial Intelligence*, Edmonton, Alberta, July 2002, AAAI Press, 224-233.
- Isaacs, E., Walendowski, A., Whittaker, S., Schiano, D., and Kamm, C. (2002): 'The Character, Functions, and Styles of Instant Messaging in the Workplace', in *Proceedings of CSCW 2002*, New Orleans LA, November 2002, ACM Press, 11-20.
- Mackay, W. (1988): 'More Than Just a Communication System: Diversity in the Use of Electronic Mail', in *Proceedings of CSCW '88*, Portland OR, Sep 26-29 1988, ACM Press, 344-353.
- Milewski, Allen E., and Smith, Thomas M. (2000): 'Providing Presence Cues to Telephone Users', in *Proceedings of CSCW 2000*, Philadelphia PA, December 2000, ACM Press, 89-96.
- Nardi, B., Whittaker, S., and Bradner, E. (2000): 'Interaction and Outeraction: Instant Messaging in Action', in *Proceedings of CSCW 2000*, Philadelphia PA, December 2000, ACM Press, 79-88.
- Palen, L. (1999): 'Social, Individual and Technological Issues for Groupware Calendar Systems', in *Proceedings of the Conference on Computer Human Interaction (CHI) 1999*, Pittsburgh PA, May 1999, pp. 17-24.

- Patterson, M. (2000): *Email in Daily Work*, draft available at (http://www.andrew.cmu.edu/user/mp72/email_in_daily_work.html).
- Sigmund, K., Fehr, E., and Nowak, M. (2002): 'The Economics of Fair Play', in *Scientific American*, January 2002, 81-85.
- Sproull, L., and Kiesler, S. (1986). 'Reducing Social Context Cues: Electronic Mail in Organizational Communication', in *Management Science*, 32, 1492- 1512.
- Tang, J., Yankelovich, N., Begole, J., Van Kleek, M., Li, F., Bhalodia, J. (2001): 'ConNexus to Awarenex: Extending Awareness to Mobile Users', in *Proceedings of CHI '01*, ACM Press, 2001.
- Trivers, R. (1971): 'The Evolution of Reciprocal Altruism', in *Quarterly Review of Biology*, 46, 35-37.
- Whittaker, S., and Sidner, C. (1996): 'Email Overload: Exploring Personal Information Management of Email', in *Proceedings of CHI '96*, ACM Press, 276-283.