

Going with the Flow: Email Awareness and Task Management

Nelson Siu Lee Iverson Anthony Tang

Human Communications Technology Laboratory
University of British Columbia
2356 Main Mall, Vancouver, BC, Canada V5T 1Z4
{nsiu, leei, tonyt}@ece.ubc.ca

ABSTRACT

Email use in the context of everyday work practices, or *email flow*, has not been heavily studied. We present the results of a pair of studies examining how users *interlace* email with their day-to-day, ongoing work processes. We demonstrate that our subjects use email as a tool for managing moment-to-moment attention and task focus. We also provide a model of this workflow that builds upon an existing model by Venolia et al. Finally, we provide specific design recommendations to enhance the usability of email clients in support of these modes of interaction.

Author Keywords

Email deferral, email flow, task management

ACM Classification Keywords

H.5.3 Group and Organization Interfaces – Computer-supported cooperative work

INTRODUCTION

There is a consensus amongst researchers that task management is an important activity supported by email. Mackay was the first to examine how email was being used for more than just communication [13]. Since then Whittaker & Sidner [21], Bellotti et al. [3,4], Ducheneaut & Bellotti [8], and Gwizdka [11] have all studied the close tie between people's tasks and their email practices. Venolia et al. then consolidated these findings into five areas of email activity [20]: *flow*, *triage*, *task management*, *archive*, and *retrieve*.

Of the five categories, email flow has been the least studied because the investigative methods employed by researchers do not do a good job of capturing email usage over extended periods. While commonly used techniques for studying email users such as short interviews and

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'06, November 4–8, 2006, Banff, Alberta, Canada.
Copyright 2006 ACM 1-59593-249-6/06/0011...\$5.00.

questionnaires are suitable for capturing direct observations of triage or task management, they *do not situate people's email use with respect to their other ongoing work activities*. For example, how do users allocate attention between email and all the other work they are preoccupied with during the day?

Email for many people is the predominant communication medium in support of moment-to-moment productivity, and a habitat within which they conduct their much of their job duties [9]. If email is primarily seen as a tool in support of other activities [8], then a basic goal for understanding email usage should be to investigate its relationship to these surrounding tasks, and the strategies users have evolved to manage these relationships.

Understanding the balance between maintaining new mail awareness *vs.* focus on non-email work would greatly benefit the design of next generation of email clients. In this paper we describe a pair of studies that have been designed specifically to begin bridging this gap. Using a combination of user shadowing and software usage logs, we have been able to document and analyze our users' email interactions over an extended period of time. We seek to develop a model of *situated email interaction* that includes the user's views of email-oriented tasks, and the transitions between these tasks. Such a model should answer the following questions:

- What do our users do with email during flow?
- Why do they do it?
- How do tools help users achieve their goals?
- How do they decide what to do next? And why?

Our answers to these questions reveal that *triage*, *flow*, and *task management* are not compartmentalized silos but are instead closely interwoven actions. In particular, we demonstrate that at least for our subject pool of high-volume email users, the majority of interactions with email in the middle of the day are part of what Venolia et al. [20] labels as *flow*. Moreover, virtually all of these interactions can be well described in terms of a user's short-term attention and task management objectives.

One of the clearest implications of this finding is an understanding of why email is so important in the day-to-day lives of many workers: *the asynchrony of email allows them to balance between concentrated, task-oriented activities and the need to communicate with others to organize those activities*. This paper will first present an overview of the difficulties surrounding email flow research, which became motivating factors for our qualitative shadowing study. Out of this study came our classification of email flow as comprising of *glance, scan, and defer* actions. These behaviours were further examined using a quantitative software logging study that deepened our understanding of some of these flow actions. We conclude by describing a number of simple, direct modifications that can be made to email clients to better support the needs of users in managing their short-term attention, task focus, planned interruptions and deferral activities.

UNDERSTANDING EMAIL FLOW

Since timely email communication is often task critical [8], users maintain an awareness of their inbox even when processing email is not the primary activity at hand [20]. Consequently, many users no longer periodically “check email” as they did in the past; instead, they leave their email clients open throughout the workday [9]. How do these users continuously maintain awareness of new messages? How do users decide to move from simply having an awareness of new messages to taking action on one? Building email clients to support modern email usage will require a more in-depth understanding of flow than we currently possess.

The study methods used by email researchers to date have made it difficult to expose flow related activities, and fail to capture the contextual dynamics of email use during the workday. In-situ interviews conducted at the start of the user’s day (e.g. [16] or [21]) isolate email activity from the broader work context. While good for studying the specific activities they were interested in such as triage or folder strategies, the data would not have been as revealing about people’s *ongoing* email management.

Researchers have also made use of questionnaires and inbox files in order to probe email related activities that cannot be captured by interviews [20]. This type of self-report inquiry however assumes the user is retrospectively cognizant of all their email activities. While questionnaire surveys are appropriate for assessing perceptions of email use (e.g. email filing strategies [20]), they are unlikely to capture the subtle activities that people take for granted in the course of their transient email monitoring. Finally, some researchers have analyzed archived inbox files offline [21]. This approach would miss altogether the dynamics of the user’s interactions with their email client.

Our workplace shadowing study complements these prior approaches by allowing us to record and identify contextual

patterns of email use as they vary throughout the day. While our presence in the office may have influenced our participants’ general behaviour, we do not believe our presence systematically affected any email related activity. This was confirmed in the post-study questionnaire, where participants reported the investigators did not affect their general workflow.

Day-long in-situ shadowing

Our aim was to gain an understanding of email use with the intuition that users will exhibit different strategies for handling email flow throughout the day. Each shadowing session began with a twenty minute interview collecting general demographic information about the user and descriptions of their job function. An investigator then positioned himself behind the user with a view of their workstation and work surfaces. Detailed minute-by-minute field notes were taken to build a picture of the user’s email usage, and any relevant work tasks that seemed to trigger email reading or generation. We limited asking clarification questions to once an hour in order to minimize intrusion. Also we always asked these questions after the user had transitioned into a period of inactivity. Some of these questions included, “Why did you read that email first?”, or “How come you decided to ignore that new mail notification popup?”

We noted all of the participants’ email activity for the entirety of a workday, save for a brief time when one participant went to the restroom with his SmartPhone. At the end of the day, users were given a questionnaire to assess their perceptions of their email use and the study.

Participants

Our four participants (one female, three male) were recruited from industry and academia using a targeted email broadcast. Three participants used Microsoft Outlook as their email client, while one used Mozilla Mail. They all used email as the core communications medium for their work, with each receiving over 50 non-spam emails per day. To provide anonymity while referring to our participants, we will use gender-appropriate pseudonyms.

Flora is an administrative assistant who performs a wide variety of tasks within her department. She may be doing reimbursements for a handful of graduate students at one moment, and then handling the department head’s meeting schedule the next. She spends most of her time at her desk, and does all of her work related email at her office desktop computer.

Larry is a lead program manager at a large software firm. He spends his entire day in and out of meetings, but constantly stays within reach of his email using his laptop while on the move. At home he uses both the web email access his company provides, as well as his SmartPhone to check new messages.

Owen is the IT administrator of a medium size software company and leads a small department of four people. His email setup is very interesting as he has an email client visible on at least one of two monitors at all times. Most of his time is spent at his desk during the day. Owen also has a Blackberry email device that handles both company emails and some personal emails from a side consulting job. He uses this early in the morning while on the bus to work and again after work on the way home.

Will is a research lab manager at a university whose role includes troubleshooting lab IT issues, tracking the large amount of research equipment being lent to graduate students, and handling any miscellaneous tasks that requires attention. Much of his time therefore is spent traveling to various locations on campus, coupled with brief 15-45 minute stays at his desk in between.

RESULTS: FLOW HANDLING ACTIONS

Similar patterns emerged in all of the participant's strategies for handling email flow. We used an open coding technique to analyze our field notes [12], and to understand our users' email activities in the context of their work. An iterative clustering of the low level actions from our notes culminated in three distinct groups of actions that pertain to managing email flow (summarized in Table 1). The common thread through these three handling strategies is that they are all employed with the aim of minimizing email disruptions to the work task at hand. The following sections present a detailed description of each strategy along with a timeframe for when each is used. Even though we discuss each strategy as a distinct category, keep in mind that users do transition between these strategies fluidly.

Glance: How many new messages are there?

Users *glance* at their inbox to maintain an awareness of the volume of incoming unread email. This interaction often lasts for less than a second. We saw this strategy employed at least once an hour even as users were deeply focused on non-email work tasks. Users only took further action if there was a surprising amount of email discovered in the glance.

Glances are a lightweight form of email awareness that takes place opportunistically during momentary breaks in the primary task, or while transitioning from one desktop application to another. Because they are so quick, it is unlikely the glance provides anything but a minimal awareness of the inbox. *Glances* provide a sense for *how much email is in the inbox* and perhaps an iconic memory for words.

11:21 am: Flora is working on a paper task. As she reaches for the "Sign Here" sticky notes, she glances at her email client, which has been left open and visible. The email client has 7 unread messages. Flora mutters "Holy

Type	Characteristics & Goals
Glance	<ul style="list-style-type: none"> Brief and opportunistic during primary task What is the rate of my incoming email? I care about: # of unread emails in my inbox
Scan	<ul style="list-style-type: none"> Short break to gain richer awareness of email inbox Is there email to be handled immediately? I care about: who sent me this and what is it about
Defer	<ul style="list-style-type: none"> Explicitly or implicitly flag emails to be handled at a later time Can I handle this email later? I care about: who sent me this and what is it about

Table 1. Summary of email flow handling strategies

smokes," and stops working on her paper task to scan through her inbox.

Glances may or may not provide an exact count of unread emails depending on the user, but we suspect that most use *glances* to gain information on the *rate* of incoming messages. Incoming email rate is important because sudden fluctuations are often associated with "emergencies."

11:30 am: Larry's inbox suddenly "hiccups," scrolling down with 10 new unread emails. Within 5 seconds, Larry minimizes his current window, and opens up the newest email, which is an issue that needs to be resolved within the hour.

Users will sometimes escalate their email-related activities based on a *glance*, with the most common transition being to a *scan* of the inbox.

Scan: Is there anything to handle now?

Users *scan* their inbox in search of new emails requiring immediate attention. These *scans* occur at task transition boundaries when the user takes a break, or when a glance detects an unusual influx of messages. Users primarily attend to the author and subject line of the messages when scanning the inbox, actively searching for both expected and unexpected emails that would be considered important (e.g. mail from manager, a message about an urgent pending issue, etc). In our sample of participants, *scans* were brief, user-initiated "interruptions" lasting no more than 5s-30s, occurring as many as two to three times an hour. Users acted on a message if it was important or if the barriers to replying were minimal; otherwise, our participants returned to their primary task after a scan.

3:21 pm: Earlier in the afternoon, Flora spoke with Bill, who was to prepare a document for her. She expected him to have it complete since she is otherwise blocked on a task. In the meantime, she has been working on another spreadsheet work item. Flora looks bored, and suddenly decides to check her inbox to see whether Bill has sent her the email. Flora recalls later, "Sometimes, when I'm waiting for someone to send me something, I don't really

notice anyone else—I was just looking to see if they’ve sent it, because usually it’s important.”

3:35 pm: Flora checks her email by bringing up her email window. This time, Bill’s email has arrived. Flora immediately opens up the email, and deals with it.

During a scan, users are interested in only two properties in each new message: *who sent me this email*, and *what is this email about?* Efficiency is paramount since *scans* occur during relatively short breaks in the primary task. In only a minority of cases would participants skim an email using the preview pane. Most of the time scans are conducted using the thread pane alone. Participants frequently re-flagged examined messages as “unread” to *defer* it for revisit [21], though none used the built-in flagging or categorizing capabilities found in their email clients.

Scans are distinct from *glances*: scans are longer in duration, and are explicitly used to identify new important emails. As we mentioned, *glances* sometimes transition to *scans*. If many items from a scan were dealt with in succession, then users seemed to transition into a *triage*-like mode, and take the opportunity to complete an entire pass on unread mail. The difference of this behaviour from a full triage session is that the user’s mindset is to return to their non-email task. In full triage (e.g. when having just arrived at the office), the user is dedicated to just handing email for an extended duration.

Defer: Should I do this now?

Users *defer* emails until later to manage *overflow*: emails that cannot or should not be dealt with now. While our users were generally good at keeping up with their incoming email stream, we observed many emails were explicitly flagged for later action (e.g. messages left in the inbox remarked “unread”, message windows that were left opened on the desktop, or half-written replies that were visible on the desktop or saved in the “Drafts” folder). Users frequently *deferred* messages for revisitation as a strategy for managing their attention budget since only the most important emails get handled during *scans*.

The fact that a user defers an email does not imply that the message is less important. Emails can be deferred for many reasons: it could be very important and therefore requires careful examination and a well crafted reply, or it may be unimportant and not deserving time immediately (if ever), or it may contain mostly reference information, and therefore filed away until it is needed [21].

Deferred emails are handled *after* more pressing activities are completed. In at least one case, a deferred email contained a task that could not or should not have been handled when it was received.

12:30 pm: Michael got a request this morning on his Blackberry to setup an intranet website. He had deferred

it since he could not handle the request earlier, but now has the time to action on the request.

Closely coupled with the *deferral* of emails is how users revisit these deferred messages. Users revisit emails in two main ways. It occurs when users finish an activity and are waiting to move to their next task, or towards the end of the day when the pace of incoming mail is likely low and the user’s time is likely free. Some of our participants were seen running through all their unread messages at least once toward the end of the day to catch issues or tasks that should be read or responded to.

The revisit activity (and hence the defer activity) is often task or issue oriented. In the following example, Larry takes several steps to find some emails related to a specific issue in a late-day revisit pass.

3:16 pm: Larry is done with his “today” email (a few remaining unread emails are deferred). He groups his inbox by sender, finding one particular contact. Finding a subject line, he groups emails by subject, opening the most recent message in the thread. Satisfied the issue is resolved, Larry deletes the entire email thread. Larry groups again by sender, scrolls back to the original sender to check there are no more emails from that sender. Larry later recalls, “By reading the newest one, I don’t have to read each email in the thread.”

The revisit activity is distinct from the retrieve activity that Venolia et al. [20] described because revisit deals with *deferred* emails that may contain tasks, whereas the retrieve activity refers to archived emails—emails that have been stored and already handled. The revisit activity is also closely tied with handling email flow, whereas the retrieve activity is a form of information retrieval.

Connecting email flow and task management

The set of actions described so far suggests a pattern of integration between email flow awareness and people’s moment-to-moment task management. Our participants prioritized on the fly between both email activities and their core job activities, switching between the two as needed and as importance dictated. This task management was very ad-hoc in nature and priorities changed fluidly as new information arrived.

Although our participants did make use of calendar applications such as Microsoft Outlook to schedule meetings and group engagements, none of them relied on their calendars to plan out ephemeral tasks. These were usually prioritized on an ad hoc basis taking into account what is in the inbox, and then updated as new messages arrive using lightweight actions (i.e. glance, scan, and defer). We know that people prefer to keep messages in their inbox as a reminder that subsequent action is required [21]. When asked why they did not use the flagging capability available in their email client, most cited the fear of the flagged list growing into an unmanageable size and

Category	Recorded Parameters
Interaction Data	<ul style="list-style-type: none"> • Time of 1st read/review/send events • Message read/review duration
Message Information	<ul style="list-style-type: none"> • Author • Subject Line • Message ID • Thread ID • Number of lines in message body • Timestamp of message • Time of download
Archive Information	<ul style="list-style-type: none"> • Send / receive frequencies with all contacts • Author, subject line, timestamps of archived messages
Inbox Status	<ul style="list-style-type: none"> • Number of unread messages at • Total number of messages during each read

Table 2. Summary of recorded interaction parameters

having some tasks go unnoticed. Having the items in the inbox forces the user to notice it each time they returned to look for new messages.

The shadowing study also highlighted the prominent role of deferral as a frequently used email handling technique. Circumstances often necessitate the need to delay email handling, leaving messages in a half read state or replies left partially composed. The existing email literature has little information on how users postpone required actions within their inbox during the day, choosing instead to focus on other aspects of use such as triage or task management. Understanding how users defer messages and the conditions surrounding its use is crucial because it represents an entry point into the user's future task list. Obtaining data on this mode of use therefore was the focus in the second half of our research, which tried to gather quantitative measures of email interactions in hopes of capturing a long term picture of deferral in action.

PASSIVE INTERFACE MONITORING

In order to delve deeper into the qualitative study results, we wanted to gather email usage data over not only a longer duration, but also in a less intrusive manner. We chose to use a special data-logging email client that recorded our users' actions over a three-week period. An analysis of glance and scan would have been very interesting, but it was not a part of our study because of the need for eye tracking data. Instead, we decided to concentrate on deferral alone because of its importance in flow, and it will be the focus from this point forth.

The open source client Mozilla Thunderbird was used as the platform for our data capture application. Using Thunderbird's plug-in extension facilities, we augmented the interface with a Javascript application that passively recorded a user's interactions with no visible changes to the client front end. The logging application recorded the user's interactions in real-time to a local RDF data file,

Name	Characteristic
Steven	<ul style="list-style-type: none"> • Associate professor • Receives 20 emails per day average • <i>Email client:</i> Netscape Communicator
Ian	<ul style="list-style-type: none"> • Department technician • Receives 40 emails per day average • <i>Email client:</i> Microsoft Outlook
Will	<ul style="list-style-type: none"> • University research lab manager • Receives 30 emails per day average • <i>Email client:</i> Netscape Communicator

Table 3. Descriptions of logging study participants

noting when messages were opened, reviewed, how long it was opened for, the number of unread messages in the inbox at the time, etc. (Table 2). The logs were automatically emailed back weekly to the investigators.

Privacy concerns

We encountered a great deal of difficulty recruiting subjects for this phase of the study due to people's reluctance to have their email interactions monitored. A decision was made early on to capture the subject line of each message (without the message body) in addition to the interaction data. This was done in case there was a need to understand the contextual relationship between several messages at a later date. Many potential participants were very worried however that even their subject lines would prove too great a disclosure of potentially confidential communication. After exhausting our options soliciting subjects through general department mailing lists, a targeted email broadcast was sent to potential subjects who trusted the investigators, and would be likely to agree to the study. In the end we were able to attract three volunteers to participate in our study. All were linked to the university in their job functions and received a large volume of messages (Table 3). They all used email as their main medium for communication, which ensured a significant amount of their work related communication would be captured during the study.

Initial setup

Each subject received a 30-minute orientation at the start of the study during which they were informed of the general goals of the project, the type of data that was collected by the logging tool, and the potential privacy concerns. Subjects were never told exactly what aspects of their email

Category	Steven	Ian	Will
Days Monitored	37	35	20
# of Handling Events	1299	1725	1021
Messages Received	644	1005	500
Messages Sent	336	389	298

Table 4. Log Data Summary

interactions we were interested in examining. We installed a copy of Mozilla Thunderbird on their computer and augmented it with our data logging extension. Attention was paid to ensure the subjects were acclimatized to the interface changes. We helped them import their existing mail archives and folder structures to the new installation, as well as any custom filters, address books, and email signatures. All the users found the layout of Thunderbird (i.e. folders, thread pane, preview pane) to be satisfactory and a familiar setup compared with their existing clients.

RESULTS: INTERFACE MONITORING

We focused our analysis of the log data around actions on new messages arriving throughout the day. Since we were interested in flow related activities, we specifically examined how often users delayed their email handling. While deferred email handling had been previously noted by other researchers [8,21], there have been never before attempts to study it directly in a quantitative manner.

Our real world data was comparable to an earlier Dabbish study [2] where participants were asked to estimate how many messages they responded to (Table 6). Most of our users’ messages were never replied to as in the Dabbish study. The vast majority of the messages that did receive a reply were composed immediately after opening. This leaves then deferred messages that still needed a reply, but could not be tended to right this moment. Within this set of messages, we noticed that users deferred their emails over two different timescales.

Users often delayed their responses until the end of an email session in what we have termed an *intra-session deferral*. By examining the subject lines in our logs, we were able to observe what the users were doing in the interim. We found several common actions including:

- Reviewing old message before replying
- Reading other messages in a group discussion to see what others have written
- Reading other unrelated messages before replying
- Initiating a message to a related 3rd party first

In the other cases, user delayed their responses over a longer time period on the order of hours to several days in what we have called an *inter-session deferral*. Even though we could not directly observe why people performed this type of deferral from the subject lines, we postulate that the

Subject	Intra-session Defer	Inter-session Defer
Steven	70%	30%
Ian	80%	20%
Will	72%	28%

Table 5. Comparison of deferral types

following are plausible reasons for engaging in an inter-session deferral:

- A need to perform an extensive information search before a reply can be crafted
- A desire to manage one’s responsiveness image [18]
- A higher priority task trumps the need to respond at this time

DISCUSSION

We have highlighted some previously undocumented components of email flow through the use of passive observational techniques. What we have found is that there exists a close interconnectedness between flow activities and task management. We also found that deferred handling is a major component of email use. Our study participants relied on their inbox as an ad-hoc tool for directing their attention throughout the day, using glance, scan, and defer actions to support their situational awareness and dynamic task prioritization needs.

Users gravitate towards email as a medium for organizing short-term tasks because of its asynchrony [5]. Email is unique amongst electronic mediums in that it gives users the freedom to handle incoming tasks when they see fit. This type of flexibility is lacking in other communication modalities (e.g. phone, IM), where users have to decide in real-time whether they wish to preempt their existing workflow.

Our users demonstrated a close integration of their email awareness activities with their engaged task based activities. This suggests to us that the model by Venolia et al. [20] showing email as five distinct and isolated activities is incomplete. We can see from our observational study that email flow is not just a background activity for maintaining awareness, but is also a part of a user’s low-level task management strategy.

Subject	No Reply	Immediate Reply	Deferred Reply
<i>Dabbish estimates [2]</i>	64%	23.0%	13.0%
Steven	72%	19.6%	8.4%
Ian	84%	12.8%	3.2%
Will	65%	25.2%	9.8%

Table 6. Comparison of reply percentages

Instead, we propose a new workflow model of email in which flow and task management activities are closely linked (Figure 1). Users make use of actions such as the ones we described (i.e. glance, scan, defer) to balance their attention between focusing on primary work tasks and maintaining an awareness of new messages. During the free periods in their work, users may take glances of the inbox to determine if any new messages have arrived. They rarely divert their attention away from ongoing tasks unless the number of new messages goes above a certain “interest” threshold. Users may also skip the glance all together and proceed to a scan if they are expecting an important message or if they have completed their task and wish to examine the inbox for any new ones. After interesting messages are identified, the user will reply to it immediately, ignore it, or defer it until later. In some cases however, a new message will cause users to abandon their current activity in deference to a more important one that just arrived.

The type of moment-to-moment task management we observed and have discussed here is a not often addressed in the email literature. Even though some of our users did use an electronic calendar for managing meetings and long-term events, the majority of their daily task management was related to their own personal organization. People used their inbox as an informal tool in support of these low level tasks because of its lightweight nature and low overhead. The types of tasks tracked in this manner are usually ephemeral and often too trivial to organize formally. Users also employ lightweight actions such as leaving messages

visible on their desktop or remarking messages as unread to assist themselves in managing their flow of information [21]. As a result, we believe that users can clearly benefit from improved lightweight tools that assist them in these activities.

The quantitative data clearly reveals an ongoing practice of delayed email handling as a central part of people’s email strategy. Even though this observation could have been made from the qualitative data alone, the interaction logs provided us with much more clarity in the understanding of this behaviour. The use of software logging as a method of investigating email usage holds a great deal of promise, of which this study has only begun to explore. Future studies can for example look for a possible relationship between perceived message importance and the reading order employed by the user. One may also try to probe the motivations behind other aspects of email use by linking the interface monitoring to additional contextual data such as eye tracker information.

Design implications

We have highlighted in our study how messages cannot simply be categorized using a binary state as most clients do (i.e. read or unread), and that deferred handling is a common usage strategy. Messages left in the inbox as task reminders can each exist in any number of indeterminate states. Email clients therefore should provide flexible mechanisms to help users track these pending messages. These features will have to be minimally intrusive to the user’s workflow, and have to be lightweight in the amount of attention required to use them.

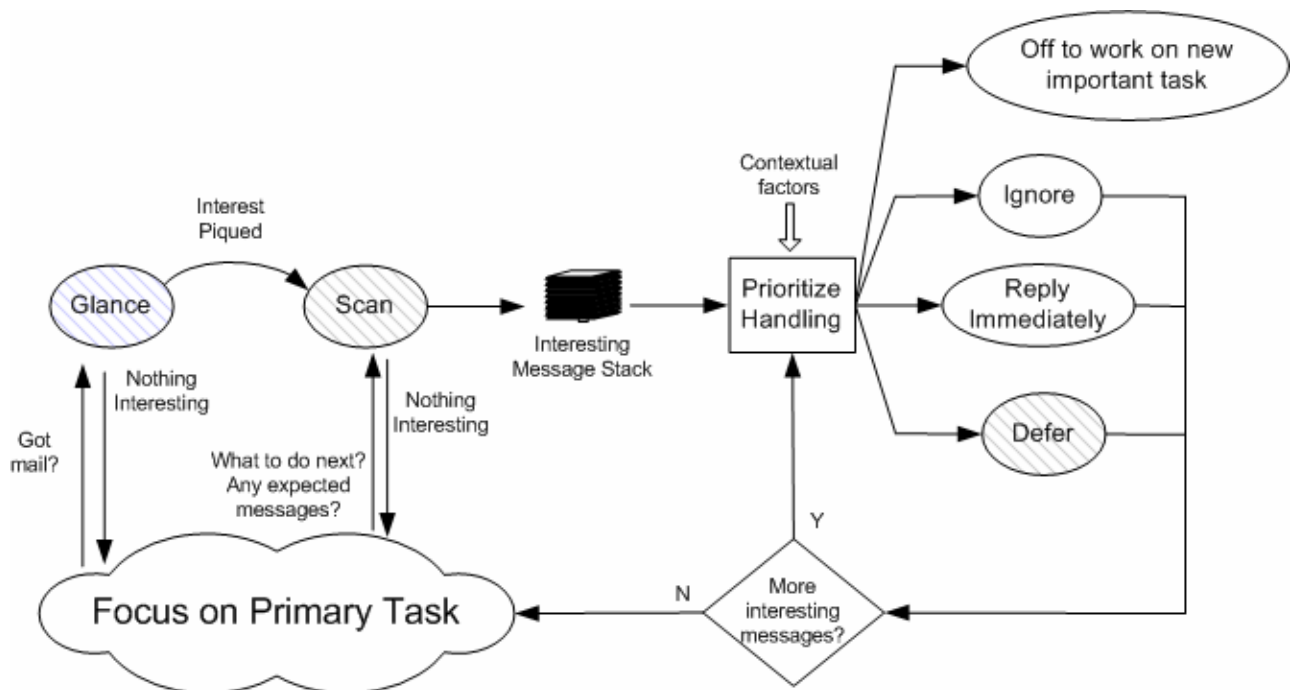


Figure 1 - Workflow model of Email Flow

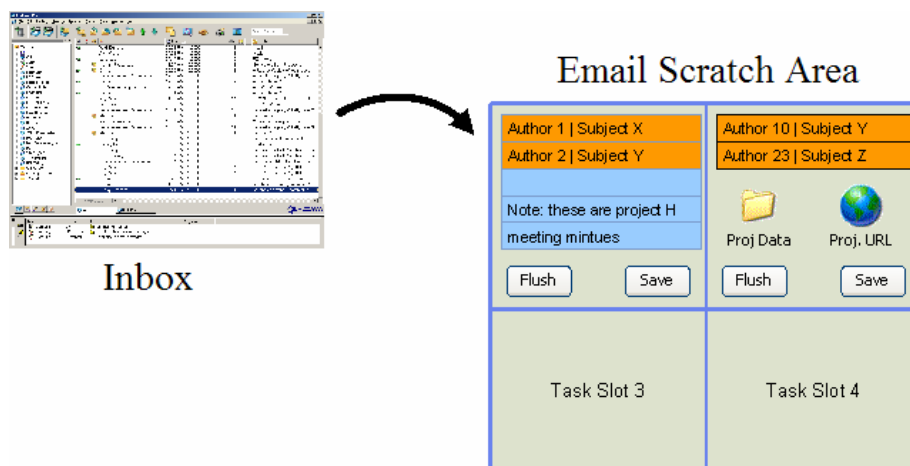


Figure 2 - Email scratch area with separate task slots

Recent client such as Chandler [1], Microsoft Outlook[14], and Google’s Gmail [10] have tried to address this issue by providing the ability to flag deferred emails and tag important messages respectively. Chandler allows users to tag messages as incomplete and offers a separate view to isolate these messages. Microsoft Outlook 2003 has long provided the ability to flag messages as important, and Gmail also offers a similar functionality with its “star” feature. However, users are unable to regroup, annotate, or reorder the messages within these consolidated views. This is important because they are the same type of actions people use to manage tasks and physical documents in the real world.

One can envision that users who juggle many deferred items can benefit from a general scratch area where flagged messages can be grouped, ordered, and annotated with minimal overhead. With the increasing affordability of large screen LCD’s, one can imagine setting aside even a separate sub-area of the desktop for the purpose of email and task organization. When a message is deferred for later reading, or if the task it pertains to is blocked, the user can simply move the message from the inbox into a generic “task slot” on the proposed *email scratch area* (Figure 2). A task slot is analogous to a digital Post-it note, but is more powerful in that users can add annotations to the task or message, group and order multiple messages pertaining to a single task, or add pointers to relevant files, URLs, and more.

The email scratch area provides a highly flexible surface for dealing with the wide range of deferred messages that can arise. The user is never restricted to narrow pre-defined tags when describing each message. It is a lightweight solution because messages are only a mouse drag or click away. There is no need for high overhead constraints such as defining the expected deadline as in the Taskmaster system when entering a new item. Unlike in Gmail, users would be able to group messages together as they see fit,

reordering them to create a sorted to-do list if necessary, and annotating the groupings to customize them. The visibility of the task slots on the desktop means there is no need to define a schedule for bringing deferred messages back into view such as with Chandler. The flagged messages automatically refresh the user’s memory each time they come into view.

The idea of combining email with other digital documents on the desktop is not a new idea, but the proposed solution is novel because it is targeting task related issues in email and deferred message management in particular. Presto was one of the earliest systems to have both email and other documents co-existing on the desktop [8]. The ReMail system by Rohall et al. showed that annotations on messages can be a useful addition to the email interface [17], and Whittaker et al.’s ContactMap showed how giving users flexible email organization on a desktop space is appealing [22]. The email scratch area embodies all of these concepts into one coherent email function.

Another area where there is potential for design improvements is in support of the scan action we have described in this paper. The motivation for performing scans of the inbox can be traced back to the coarse granularity of email notifications, which trigger on any new message whether or not it is relevant to the user. The scan action is a partial remedy employed by users who have turned off the alert feature in their client, but still have the need to maintain ongoing awareness of new messages. Users will likely increase the frequency of their scans if they are anticipating an important message, which in turn can cause unnecessary disruption to their work in progress.

Clients can partially improve upon this situation in a specific use case, such as when a user is waiting for a reply to a message they sent out. This is accomplished by having the user flag important outgoing messages at send time using a special “send with alert” button (Figure 3). Such a feature would be a one-touch lightweight action, yet

powerful because it empowers the user with a level of control over the client's notification mechanism. When the reply comes back from the other person, the client will know it is appropriate to alert the user, and that it is contextually relevant. The human involvement in this proposed solution avoids the need for any intelligent agents or natural language processing of emails.

CONCLUSIONS AND FUTURE WORK

As the most significant communication medium of the Internet age, email is unique amongst its IM and voice cousins in that it gives the receiver almost all of the power in determining how information will be processed and when a conversation will be continued. By deeply studying the engaged work practices of a number of high volume email users, we have demonstrated that these users have taken advantage of this affordance. They have adapted their email interaction strategy to support a deeply engaged process of moment-to-moment attention and task management, and have done this within the limitations of current client interaction modes.

In designing future email clients (or adapting existing ones), we recommend that the workflow model shown in Figure 1 be used to guide new interface concepts that can display the information users need to engage in *glance* and *scan* actions without having to "open" an email client. We also suggest that better support be provided for the kinds of lightweight deferral that seem to be endemic to these interactions. Finally, we suggest a number of other ways of presenting information that might help users to better control interruptions and emails closely related to tasks.

Certainly with the small number of subjects we were able to follow, more validation and broadening of our results is in order. In the next phase of the work, we hope to expand the generalizability of our findings with a much larger subject pool. Our recruiting difficulties could be minimized if we scrubbed all personally identifiable data from the software logs, or if we enlisted the support of large organizations opened to assisting us with this type of research.

We will validate our strategies with even more quantitative data in parallel with our qualitative results. These may include using eye tracking data, improving the level of detail in the software interaction logs, and performing field studies on many more participants. As the daily barrage of email continues to grow, users require interfaces that understand and respect how they manage incoming email. We hope that with further research into email flow, eventually designers will be able to produce interface solutions that directly supporting email usage habits.

REFERENCES

1. Bälter, O. and C. L. Sidner (2002). Bifrost inbox organizer: giving users control over the inbox. NordiCHI '02: Proceedings of the second Nordic



Figure 3 - Send with Alert Button

- conference on Human-computer interaction, ACM Press, 111-118.
2. Balter, Olle (2000): Keystroke Level Analysis of Email Message Organization. In Proceedings of the ACM CHI 2000 Human Factors in Computing Systems Conference. April 1-6, 2000, The Hague, The Netherlands. p.105-112.
3. Bellotti, V.; Ducheneaut, N.; Howard, M. A.; Smith, I. E. Taskmaster: recasting email as task management. CSCW 2002 Workshop on Re-designing E-mail for the 21st Century; 2002 November 16; New Orleans, LA.
4. Bellotti, V., N. Ducheneaut, et al. (2003). Taking email to task: the design and evaluation of a task management centered email tool. CHI '03: Proceedings of the SIGCHI conference on Human factors in computing systems, ACM Press, 345-352.
5. Berghel, H. "E-mail: The good, the bad, and the ugly," Communications of the ACM, 40:4(April 1997), pp. 11-15.
6. Chandler in a Nutshell. <http://chandler.osafoundation.org>
7. Dabbish, L.A., Kraut, R.E., Fussell, S.R. and Kiesler, S.B. Understanding email use: predicting action on a message, *Proc CHI '05*, 691-700.
8. Dourish, P., Edwards, W. K., LaMarca, A., and Salisbury, M. 1999. Presto: an experimental architecture for fluid interactive document spaces. *ACM Trans. Comput.-Hum. Interact.* 6, 2 (Jun. 1999), 133-161.
9. Ducheneaut, N., and Bellotti, V. E-mail as habitat: an exploration of embedded personal information management. *Interactions*, 8(5), 30-38.
10. Google Gmail <http://gmail.google.com>
11. Gwizdka, J. (2002). Reinventing the Inbox: Supporting Task Management of Pending Tasks in Email. In Proceedings of CHI 2002 Human Factors in Computing Systems, ACM, NY, 550-551.

12. Hughes, J., King, V., Rodden, T., and Anderson, H. The role of ethnography in interactive systems design. *Interactions*, 2(2), 56-65.
13. Mackay, W.E. More than just a communication system: diversity in the use of electronic mail, *Proc CSCW '98*, 344-353.
14. Microsoft Outlook
<http://www.microsoft.com/Outlook>
15. Miller, T. and Stasko, J. The InfoCanvas: information conveyance through personalized, expressive art. *Proc CHI '01*, 305-306.
16. Neustaedter, C., Brush, A., and Smith, M. Beyond "from" and "received": exploring the dynamics of email triage, *Proc CHI '05*, 1977-1980.
17. Rohall, S. L., Gruen, D., Moody, P., & Kellerman, S.(2001). Email Visualizations to Aid Communications. In Proceedings of InfoVis 2001 The IEEE Symposium on Information Visualization, IEEE. 12-15.
18. Tyler, Joshua R., Tang, John C. When Can I Expect an Email Response? A Study of Rhythms in Email Usage. *Proc. ECSCW 2003*
19. Tyler, J. R., Wilkinson, D. M., and Huberman, B. A. 2003. Email as spectroscopy: automated discovery of community structure within organizations. In *Communities and Technologies*, M. Huysman, E. Wenger, and V. Wulf, Eds. Kluwer B.V., Deventer, The Netherlands, 81-96.
20. Venolia, G. Dabbish, L., Cadiz, J., and Gupta, A. Supporting email workflow, MSR-TR-2001-88 (2001), Microsoft Research, USA.
21. Whittaker, S., and Sidner, C. Email overload: exploring personal information management of email, *Proc CHI '96*, 276-283.
22. Whittaker, S., Jones, Q., Nardi, B., Creech, M., Terveen, L., Isaacs, E., and Hainsworth, J. 2004. ContactMap: Organizing communication in a social desktop. *ACM Trans. Comput.-Hum. Interact.* 11, 4 (Dec. 2004), 445-471.