Beyond "From" and "Received": Social Sorting for Email Triage

Carman Neustaedter[1]
University of Calgary
2500 University Drive NW
Calgary, AB, Canada
carman@cpsc.ucalgary.ca

A.J. Bernheim Brush and Marc A. Smith
Microsoft Research
One Microsoft Way
Redmond, WA
[ajbrush *or* masmith]@microsoft.com

ABSTRACT
Email triage is the process of going through unhandled email and deciding what to do with it. Email triage can quickly become a serious problem for users as the amount of unhandled email grows. We investigate the problem of email triage by first presenting interview and survey results that articulate user needs for email triage. We then use these results to outline a set of implications for the design of software to support email triage and present a prototype design which illustrates these principles. Our prototype aggregates social meta-data about email correspondents and places email triage within the user's social context by providing social sorting of email items using additional meta-data elements.

Author Keywords
Email triage, email user interfaces, social accounting metrics, social computing, data mining

ACM Classification Keywords
H.5.2. Information interfaces and presentation (e.g., HCI): graphical user interfaces, prototyping, user-centred design; H.2.8. Database Applications: data mining.

INTRODUCTION
The volume and use of email has changed dramatically over the past decade while user interfaces to support email have changed relatively little. While email has been the focus of many research efforts, we are particularly interested in understanding and supporting *email triage:* the process by which one goes through unhandled email and decides what to do with it [14]. Email triage varies depending on the person and the amount of time since last checking one's email. There are those people who may have to triage only five or ten messages at a time, and others who have to triage hundreds or even a thousand new messages. Our informal observations suggest that triage does not present a problem for those who must triage a small amount of email; the real problem exists when unhandled email rises above a certain level. Thus, we focus on *high volume triagers* who may not even have the time to triage all of the email which they feel are important. Email management for high volume triagers *could* be a full-time job, but is not because of time constraints.

We feel existing email user interfaces do not provide high volume triagers with an effective means for performing email triage. Existing interfaces usually provide only the most basic information about new email, such as who it is from, when it was received, and the subject line. This information may be enough for people triaging a small amount of email; however, we feel it does little to aid those who must face hundreds of new emails on a frequent basis.

Although many researchers have studied the use of email (e.g., 10, 17, 4, 6, 8), few have focused specifically on email triage. As a result, it is not clear what specific problems users experience when triaging their email, besides having simply too much email. Thus, our goal is to better understand the process of email triage in order to provide insight into designing software to support it.

First, we discuss existing research specific to email triage to ground our work. Second, we present findings from contextual interviews and a survey that clearly articulate the problems facing users when triaging email. Third, we discuss a prototype system that aggregates and visualizes social meta-data to enhance the email triage experience by sorting email along multiple dimensions of the social history of communication between the sender and receiver. We do this by displaying people along with associated social metrics about their email, e.g., the number of emails sent to the person, the number of unread emails from the person. High volume triagers are then empowered with *social sorting*: people and their emails can be sorted based

[1] Work done while at Microsoft Research.

on a variety of social metrics to bring the messages from people most likely to fit a desired social relationship to the top.

PREVIOUS RESEARCH ON EMAIL TRIAGE

Venolia et al. [14] found that the triage activity is primarily performed when people arrive at work first thing in the morning, return from a meeting, or receive an important email. They discovered that people use one of two strategies for triaging email: serial or prioritized. The serial technique involves a sequential viewing of each unread message, while in the prioritized technique users view messages in order of the messages' perceived importance.

Venolia et al. [14] also found a variety of factors that users felt made a message important including: reply to my message, from manager, I'm on TO line, high importance flag, from project member, and from direct report. While these simple attributes may help users determine the importance of each email and thus aid triage, it is not clear if they are as effective for high volume triagers. Our research extends Venolia et al.'s work by articulating the problems users face when triaging their email and how email volume affects people who receive large volumes of email, which was not their focus.

Tyler and Tang [13] report on social pressures when triaging email. They found that people desire to maintain a "responsiveness image" when handling email. That is, the time to respond to an email conveys important social information and people will make a concerted effort to maintain particular social images depending on the email recipient. This has direct effects on email triage as the social responsiveness factor will play a role when people are determining how to handle each new email. Moreover, this work suggests that social meta-data about email will be valuable to users.

Several research prototypes have presented design ideas for supporting email triage, despite the fact that not all were originally designed to support triage.

Horvitz et al.'s [7] Priorities system uses several meta-level attributes to determine the criticality of new unread emails based on people, time, the type of email, and various message attributes. Priorities was focused on email triage, however, because this meta-data is not exposed to the user, it is not possible for people to use this meta-data to easily select the emails they feel are most important depending on their current context.

Lockerd's [9] DriftCatcher provides social meta-data about emails including relationship tie strength between senders and recipients, the average time to respond to an email, and indicators of message content. However, similar to the Priorities system, users are unable to sort by social attributes to bring out emails that are important given the current context.

Another general strategy for handling email triage is to reduce the number of items to triage by grouping emails into conversational threads, categories of interests, or collections of tasks [6, 1, 2, 15, 8]. While this is helpful, people who receive large volumes of emails typically still have a large number of entities to triage within the groups. Moreover, most of these systems provide little in terms of meta-information for users to find the important entities, regardless of how they are grouped.

Smith et al. [12] present alternative spatial design metaphors for visualizing email where people and their emails are placed spatially within the interface according to social groups, e.g., work, project, or personal. While the visualization of social contexts is a valuable idea, the design metaphors provide no means to help the user decide which context is currently of importance.

Viégas et al.'s [16] PostHistory lets users see the dynamic nature of email interaction with particular people over time. While the prototype presents meta-data to show how social relationships in email have changed over time, because it was not designed to aid email triage, it fails to provide access to actual emails.

We now build on this previous research to better understand email triage and inform our design efforts.

STUDY METHODOLOGY

Our study of email triage took part in two stages. First, we conducted a series of contextual interviews to find out firsthand what problems users face when triaging their email. Second, we created and deployed a survey to validate our interview findings and provide greater insight into email triage. Our goal for both stages was to understand the strategies people employ when triaging their email, what works well, what doesn't work well, and what users feel would help them with email triage. All participants in both stages used Microsoft Outlook to handle email as it is the standard email software in our company.

Contextual Interviews

We interviewed ten individuals (4 female, 6 male) in our company with a variety of job roles, e.g., software developers, managers, researchers, administration. We deliberately selected people who received what we considered to be a high volume of email: greater than 100 emails per day. We felt that users receiving less than this would experience fewer problems when triaging email and that any problems they experienced would also occur with high volume email triagers All participants were rewarded with a coupon for a free coffee.

Interviews were scheduled for a time when the participant would have a sizeable amount of email to triage, e.g., first thing in the morning, or after lunch. Participants were also asked to avoid triaging their email before their interview.

Each interview was held in the participant's office and lasted no more than one hour.

Participants were first asked some introductory questions about their email environment, such as how many emails were currently in their inbox, and if they used folders and how. Next, participants were asked to triage their email while thinking-aloud. During the triage process, participants were probed about their techniques for dealing with unhandled email. The interview concluded with a series of follow-up questions, such as what the user felt worked well for triage, what didn't work well, and what would help the user better perform email triage.

Survey

We created a survey that largely paralleled our contextual interviews in order to validate our interview findings with a larger population. The survey asked questions about the user's email environment, triage technique, and satisfaction with current email software. We distributed the survey to 2000 randomly selected employees within our company. Unlike our interviews, our participant selection for the survey did not discriminate against the volume of email received by users. We wanted to see if any differences existed in the techniques used by low *vs.* high volume email users.

After distributing the survey via email, we received 233 (10.1%) completed surveys over the course of one week. Our respondents comprised a variety of jobs including software developers, managers, testers, product support, research, sales, and administration. For portions of our analysis, we classified our 233 respondents into three categories:

1. low volume triagers (LV): respondents who receive fewer than 50 new emails each day—we had 69 (29%) LV respondents;
2. medium volume triagers (MV): respondents who receive between 50 and 100 new emails each day—we had 84 (36%) MV respondents; and,
3. high volume triagers (HV): respondents who receive more than 100 new emails each day—we had 81 (34%) HV respondents.

A potential downside of distributing our survey via email was that people who receive lots of email are less likely to respond. Despite this, we still received a large number of responses from high volume triagers.

RESULTS

We now describe the results from both our interviews and the survey, contrasting two methodologies and respondent types.

Pre-Processing to Reduce Triage Volumes

The first stage of email triage that we found to be common amongst most participants was a pre-processing phase. Here, rules along with folders are often used to reduce the number of items the user needs to triage.

Type of Email	Respondents
Sent to mailing lists with topic of interest	117
Sent to group mailing lists	94
Sent to project mailing lists	70
Sent from a particular person	63
Sent directly to you	50
Project-related	44
From family or friends	35
CC'd to you	28

Table 1: The number of survey respondents (n=233) who use rules to move each type of email to a folder.

How many rules are used to pre-process email before triaging? Ten (4%) of our survey respondents didn't know what rules were. Of those who knew about rules, we asked how many rules they used on a scale from 1 to 5 (1: 0 rules, 2: 1-4, 3: 5-9, 4: 10-14, 5: 15+). The median response for LV was 3 or 5-9 rules (mean, 3.1 ± 1.0), MV was 3 (mean, 3.7 ± 1.3), and HV was 4 (mean, 4.2 ± 1.4). Significant differences were found between LV/MV and HV users ($p < 0.05$).[2] Thus, rules were more heavily employed as email volumes increased.

What types of emails are pre-processed using rules? Table 1 shows our results in order of popularity and indicates that rules are primarily used to folder emails sent to mailing lists. Interview respondents said that rules were generally used to move emails that were less important into folders so that important emails would stand out more in the inbox. Folder usage ranged greatly with interview participants from the use of no folders to the use of almost a thousand folders. Folders were organized according to time, project, person, or interests; similar results were found by Ducheneaut and Bellotti [4].

Setting Up the Email Environment

We found people setup their email clients and utilize features in a number of different ways when it comes to the actual process of triaging email.

Do people use search folders during email triage? One possible mechanism for triaging email is to use a feature provided by Microsoft Outlook called search folders: emails matching a set of search criteria are displayed together in a virtual folder. For example, a search folder could be created to list all emails sent directly to the user. However, we found that nearly 40% of survey respondents didn't even know what search folders were, while another

[2] Kruskal-Wallis tests along with Mann-Whitney post-hoc analyses were performed.

Type of Email	All Groups
Unread emails in my inbox	4.4 ± 0.7
Emails received today	4.4 ± 0.7
Emails received in the last hour	4.0 ± 1.0
Emails received in the last 15 minutes	3.8 ± 1.1
Emails received yesterday	3.7 ± 0.8
Read emails in my inbox	3.6 ± 1.0
Unread emails in my folders	3.5 ± 1.1
Read emails in my folders	3.0 ± 1.0
Emails received last week	3.0 ± 1.0

Table 2: Mean responses to survey questions of the form "It is important to include the following when triaging email..."

31% knew about them but didn't use them. About 24% of respondents only had 1-4 search folders and less than 5% had 5 or more. Usage did not differ based on our user categories. Only one of our interview participants used search folders.

How do people sort emails when performing email triage? We were also interested in knowing how people setup sorting when triaging their email. In the inbox, an overwhelming majority of respondents, about 80%, used Outlook's default sort of newest to oldest most often to view emails. About 9% sorted from oldest to newest most often in the inbox, while 7% sorted by sender and 2% sorted by subject line. In folders, 63% used the default sort of newest to oldest, about 20% sorted email in folders by sender most often, 8% by oldest to newest, and another 8% by subject line. These did not differ largely by user group.

Do people group emails by conversation? Microsoft Outlook also provides users with the ability to group emails by conversation, which has the potential to reduce the number of total items to be triaged by grouping messages with reply relationships. We found that 27% of respondents didn't know about this feature and another 26% knew about it, but didn't use it on their inbox or folders. We took a deeper look at those who said they actually grouped by conversation. For all user groups, we found people only grouped by conversation occasionally, be it in their inbox or folders. We found similar evidence of users not typically grouping by conversation in our interviews; only one participant used this feature. While others didn't use the feature, some did occasionally sort emails by subject line producing a similar effect.

Choosing What to Triage

We wanted to know what emails people check when triaging and if triage extended beyond the inbox to folders.

What emails do people include when triaging? As discussed, email triage involves deciding what to do with unhandled email. We use the word "unhandled" specifically because we did not know prior to the study which types of emails users were including in their triage process, be it read or unread emails in the inbox, or emails in folders which were moved there manually or with rules.

We found in our interviews with HV triagers that they do not often look at the contents of folders when triaging; rather, triaging is primarily performed on unread emails in the inbox, with some people also looking at read items. We asked survey respondents to rate how important it is to triage each of the types of emails listed in Table 2 using a 5-point Likert scale (1-Strongly Disagree to 5-Strongly Agree). Table 2 shows the mean user responses for all groups. Clearly the most important emails to handle are unread emails in the inbox and those which arrived today. People do, however, find it important to include other items when triaging and none rated any as unimportant. Our interviews showed that this often depends on the amount of time the user has to triage. For example, a common practice we saw in our interviews was to triage unread items in folders during spare moments throughout the day, but not include them during regular triage sessions.

Email Triage Approaches

Our survey and interview results show a number of common approaches that users employ for triaging email.

How are emails handled during triage? During our interviews we saw that emails are usually handled on a per email basis where the user looks at the email, acts on it, and then either leaves it in its current location or moves it to a folder. Acting on an email meant either simply reading it (or a portion of it), choosing not to read it, sending a reply, forwarding the email, or performing some other task associated with the email (e.g., writing another email, working on a document). When triaging, emails that users felt would be needed at a future date were moved to folders. Other emails were left in the inbox or moved to folders as tasks to do throughout the day. Occasionally users would handle multiple emails at the same time, e.g., moving a batch of emails to a folder. This most often occurred when the emails were part of the same conversation or thread, determined by the subject line.

Do people triage email sequentially or by priority? We found one of two methods are used for handling a given pass of one's email: sequential or by priority. Venolia et al. [14] also report these two approaches. We asked survey respondents if they used each method with two 5-point Likert scale questions. The mean response for using the sequential approach was 3.4 ± 1.1, while the priority approach was 3.0 ± 1.0. We did not find significant differences between user groups for both of these questions. We also analyzed these questions together: about 19% of all participants triaged by priority only (agreeing with priority and disagreeing with sequential), 30% only sequentially (agreeing with sequential and disagreeing with priority), and 15% used both (agreeing with both). The remaining participants rated their use of one of the two approaches as neutral.

Do people triage email with a single pass or with multiple passes? In our interviews we found people use either a single pass or multiple passes to triage their email. In the single pass case, the user starts at the top or bottom of the unhandled list of emails and handles each email as the user comes across it. Half of our interview participants used the single pass approach for handling email triage and all of them handled email sequentially during this pass. The alternative, using a priority approach with a single pass, means emails will simply get missed. The important aspect of the single pass approach is that each email is only visually scanned once during triage.

In the multiple pass case, the user performs multiple single passes handling a certain type of email during each pass. People still check to see if an email is one they should handle in the current pass by glancing at it or scanning details like the From or Subject line, but we do not consider this to be handling the email. The important aspect about using the multi-pass process is that users will often scan the same email multiple times before actually handling it.

Half of our interview participants used the multi-pass approach for handling email triage. Three of these five performed both sequential and priority passes. One multi-passer strictly used a sequential approach and one only used a priority approach.

We verified our findings with the survey results. Survey respondents rated their use of each strategy on a 5-point Likert scale (1-Strongly Disagree to 5-Strongly Agree). The mean response for using a single-pass strategy was 2.8 ± 1.1, while the mean response for using a multi-pass strategy was 3.6 ± 1.0. We did not find significant differences between user groups for either of the two strategies.

We also analyzed these questions together to see if one strategy was more dominant than the other. About 17% of all participants used a single-pass strategy (agreeing with single-pass and disagreeing with multi-pass), 47% used a multi-pass strategy (agreeing with multi-pass and disagreeing with single-pass), and only 9% used a combination of both (agreeing with both single and multi-pass). The remaining participants rated one of the strategies as being neutral or felt they used neither strategy. Clearly the multi-pass strategy is dominant, despite the need to scans emails more than once.

What emails do people try to handle first? What we found most interesting was that interview participants using the multi-pass strategy would routinely use a first pass to handle emails they considered to be *not* important or junk. This pass would involve finding emails they could quickly delete or get rid of. This seemed counterintuitive to us initially because we thought people would try to handle email that was most important to them first when confronted with only a short amount of time to triage email. Contrarily, we found it was easiest for people to handle emails of little importance (they could quickly delete them or file them) and often once the unimportant emails were gone, it was easier to find the important emails. Subsequent passes would then include handling the more important emails from specific people or about specific projects. During the interviews, we found that the importance of an email largely depends on the current social context of the person, e.g., their schedule for the day, the people they worked closely with, the projects they were working on, and their project role.

Time of Day	LV	MV	HV
First thing when I get to work	4.0 ± 0.9	4.1 ± 0.8	4.4 ± 0.6
Throughout the day as email arrives	4.0 ± 0.7	3.9 ± 0.7	4.1 ± 0.7
After lunch	3.4 ± 0.8	3.4 ± 0.9	3.7 ± 0.9
After meetings	3.3 ± 1.0	3.5 ± 0.9	3.5 ± 1.0
Right before I leave work	3.2 ± 1.1	3.2 ± 1.0	3.8 ± 1.0
In the evening at home	2.6 ± 1.1	2.9 ± 1.2	3.5 ± 1.3
Before I come to work	2.5 ± 1.0	2.8 ± 1.1	3.0 ± 1.2
During meetings	2.2 ± 1.0	2.5 ± 1.0	2.8 ± 1.1

Table 3: Mean responses to survey questions of the form "I spend time triaging my email..."

We asked survey respondents what type of emails they tried to handle first, important or unimportant, using two questions rated on the same 5-point Likert scales. The mean response for trying to handle non-important emails first was 3.5 ± 1.3, while the mean response for trying to handle important emails first was 3.7 ± 1.2. We did not find significant differences between user types for both questions. We also analyzed these questions together: 27% of all participants try to handle non-important emails first (agreeing with non-important and disagreeing with important), 21% try to handle important emails first (agreeing with important and disagreeing with non-important), and 29% (agreeing with both) try to handle both first (likely these people use a sequential approach). The remaining participants rated one of the two questions as neutral. These results show that users are mostly split when handling important *vs.* non-important email first.

Triage Times

We wanted to understand at what times during the day people performed email triage and how much of their time it was consuming.

When do people triage their email? We asked survey respondents when they triage their email using a series of questions based on a 5-point Likert scale (1-Strongly Disagree to 5-Strongly Agree). Table 3 shows the mean responses for each time grouped by user type. Clearly users triage their email at the beginning of the day with subsequent triage times occurring throughout the day as

time permits. We found significant differences ($p < 0.05$) between LV and HV for triaging first thing in the morning, before work, during meetings, before leaving work, and in the evening at home.[3] We also found a significant difference ($p < 0.05$) between MV and HV for triaging email in the evening.[3] That is, HV users triage their email more times throughout the day than LV users, and are more likely to triage their email in the evening than both LV and MV users. This leads us to believe that HV triagers are forced to triage their email more simply because they receive so much of it. In our interviews, we found that HV triagers often feel overwhelmed with email triage when they are unable to check their email frequently.

How much time do people spend triaging email? Our interviews showed that most HV triagers spend 1-2 hours triaging their email. We asked survey respondents about how much time they spend triaging their email, rated from 1-4 (1: 0-19 minutes, 2: 20-59 minutes, 3: 1-2 hours, 4: More than 2 hours). The median response for LV was 2 (mean, 1.7 ± 0.7), MV was 2 (mean, 2.2 ± 0.8), and HV was 2 (mean, 2.5 ± 0.8). While the medians were all the same, variations in the responses between groups led to significant differences between all user groups ($p < 0.05$).[2] Thus, as one would expect, when email volume increases so does the amount of time needed to perform email triage. One interview participant even said that if his email volume reached a point where he felt overwhelmed, he would begin to use personal time to bring things back "under control."

User Satisfaction

Understanding email triage from the user's perspective is of vital importance. We asked participants what they thought about email triage and how triage tools could be improved.

How satisfied are people with email triage? Regardless of the user type, we found that most people felt their strategy was pretty good, but realized there were likely other, more efficient, strategies. People who were in front of their computer for most of the day generally had an easier time triaging their email; most email could simply be handled during the incoming flow. Participants found emails built-up when they were away from their desk in meetings or away from the office making subsequent triage much more difficult. Particular grievances were found when people returned from vacation. When asked what they would like to see in email clients to help them perform email triage answers from interview respondents included mechanisms to keep important things visible, tools to easily find things, and additional task management features.

DISCUSSION

The results of our interviews and survey point out that the real issue facing people for email triage is not whether their current strategy works or does not work. Strategies will vary depending on how much time each person has at the given moment. The real issue is simply time. People need a means to more efficiently handle email triage and this need rises most for people receiving larger volumes of email. People need tools for supporting email triage that take advantage of the strategies already being employed. Using our results we have come up with a set of recommendations for designing systems to support email triage.

First, people do not group by conversation nor do they use search queries. In the case of email triage, it is difficult to know what emails to search for. Therefore, we feel that email triage user interfaces should *provide additional sorting options rather than searching features.*

Second, an imperative task for email triage is finding the important emails so that they can be dealt with. Often these emails are interspersed with emails of lesser importance making it very difficult to find them. We also found that importance is largely dependent on the user's current social context, e.g., the people with whom they are currently working. Thus, email triage user interfaces should *provide a mechanism to bring the important emails to the forefront using socially salient information.*

Third, we found that email triage is typically performed on recent unread items in the inbox, but can include other items depending on the amount of time the user has. For this reason, user interfaces designed to support email triage should *provide mechanisms to view emails from varying time periods, with varying attributes, e.g., read or unread.*

Fourth, people look at their email during triage in different ways. Some triage using a single pass, while others use multiple passes where different types of emails are handled in each pass. This raises a corollary to the second implication, user interfaces should *provide a means to change the social information that is being presented to bring different classes of content into visibility.*

It is important to realize that these implications should not stand alone. Rather, they are highly dependent on one another and should be utilized as such in order to design effective tools for supporting email triage.

DESIGNING FOR EMAIL TRIAGE

Based on the results of our interviews and survey, we have created a prototype application for email triage that supports our design implications by providing social sorting on a variety of social metrics and time periods.

Social accounting meta-data includes a range of information related to socially salient properties about conversational spaces. One aspect of meta-data is quantifiable measures we call social metrics. Our prototype, Email Snarf (**S**ocial **N**etwork **A**nd **R**elationship **F**inder), collects and aggregates social metrics for each of a user's email correspondents. Using this aggregated data, Snarf presents interfaces displaying people and their

[3] ANOVAs along with follow-up T-tests using Bonferroni correction were performed.

		Metric Name	Metric Attributes	Social Information Provided by the Metric
Sent	1	ToFrom	emails sent to each person from the user	who receives the most/least emails from the user
	2	ToFromReply	replies to each person from the user	who receives the most/least replies from the user
	3	CCFrom	emails CC'd to each person from the user	who is included in the most/least email conversations by the user
	4	ToNotFrom	emails sent to each person from others (not the user)	with whom is the user included in conversations (the user appears on the CC line)
Received	5	FromTo	emails sent to the user from each person	who sends the most/least emails to the user
	6	FromToReply	replies sent to the user	who replies the most/least to the user's emails
	7	FromToRead	emails sent to the user and marked *read*	who sends the most/least emails to the user that are read
	8	FromToUnread	emails to the user and marked *unread*	who sends the most/least emails to the user that are not yet read
	9	FromCC	emails CC'd to the user	who includes the user in the most/least email conversations
	10	FromCCRead	emails CC'd to the user and marked *read*	who includes the user in the most/least conversations that the user had read
	11	FromCCUnread	emails CC'd to the user and marked *unread*	who includes the user in the most/least conversations that the user has not read

Table 4: Sent and received metrics aggregated for each correspondent.

corresponding social metrics to present the user with mechanisms for triaging email. We discuss the collection of person-centric social metrics, our visualization of these metrics, and email triage scenarios using our prototype.

Collecting Person-Centric Meta-Data

Email Snarf collects social metrics for email. The collecting portion of Snarf is based on an existing search project called Stuff I've Seen (SIS) [5]. SIS continuously scans a user's local computer and updates a local database index with information about the user's documents and emails. Snarf collects and aggregates data from this database to generate social metrics for a given user's email correspondents. Correspondents include those individuals appearing in the To or CC lines of any emails sent to or from the user. We currently aggregate two types of metrics, for each correspondent, *sent metrics* (Table 4: rows 1-4) and *received metrics* (Table 4: rows 5-11). Sent metrics provide social information about email sent by the user to a correspondent, while received metrics provide information about email received by the user from a correspondent.

We use a labeling scheme to describe each metric: the first word is the role of the corresondent in the metric, the second word is the role of the user in the metric, and any final words are special attributes about the metric. For example, the metric **FromToUnread** shows the number of emails ***From*** the correspondent, sent ***To*** the user, and marked ***Unread.***

All of our metrics are collected for each of the user's correspondents, aggregated for a set of time periods (daily, weekly, monthly, and over all of one's email), and stored in an Access database. Users can adjust the time frame at which the aggregations are updated to include new emails.

While this is only our initial set of metrics, we feel they offer a compelling set of social information to enhance email triage and support our design implications.

Performing Email Triage

Email Snarf presents visualizations of person-centric social meta-data to support email triage. Our emphasis is to provide easy mechanisms for users to view and sort people according to social meta-data in order to bring social relationships to the forefront. We begin by describing two scenarios that illustrate how the social metrics and sorting features found in Email Snarf can support email triage.

Finding Important Emails Using Social Relationships

Our first and most important scenario illustrates how to sort people to bring emails from socially important people to the forefront. Social importance can be determined by many of the metrics we aggregate. For example, the number of emails you send to someone, the number of times you reply to a person's emails, the number of emails that you receive from someone that get marked read, or the number of times someone replies to your emails all indicate social relationships. Using our prototype, users can sort correspondents by one of these social relationship indicators and use a secondary metric to display, say, the number of unread emails from each person. People of current but recent importance can be sorted to the top by selecting a current time period, while people of past importance could be sorted up by selecting a prior time period.

Figure 1 illustrates this type of social sorting with a sample user's email correspondents and a default set of social metrics. People are displayed in a vertical list along with their associated social metrics. We have changed people's names in the figure to protect privacy and more clearly illustrate each person's social relationship to the user. Next to each contact, two social metrics are visualized with red and blue bars.

In Figure 1, people are sorted by the red metric bar which shows the number of messages sent from the user to each person thus selecting the correspondents with whom the user most frequently initiates interaction. The blue metric bar shows the number of unread emails from each person. If a user has zero emails for a metric, the corresponding bar is not shown. For example, the user has not sent the last eight correspondents listed in Figure 1 any emails; therefore, only a blue bar shows for each of them.

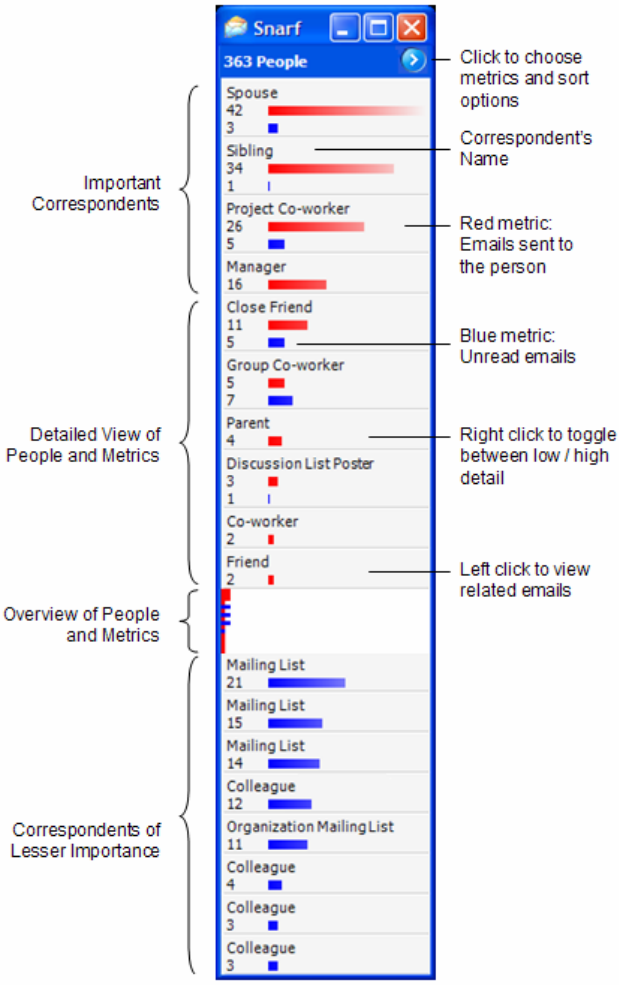

Figure 1: Email Snarf visualizes people and corresponding social metrics.

Alternatively, the Parent, Co-worker, and Friend listed only have red bars indicating they the user has sent them email, but currently has no unread email from them.

People such as the user's Spouse, Sibling, Project Co-worker, and Manager are high on this list and, therefore, have an important relationship with the user simply because the user takes the time to send them the most email. Emails from these people easily stand out with the current sort; thus, it is easy to find the important unread emails. If one were to alternatively only sort by the number of unread emails from each person, emails from these important people would not stand out. Rather, emails from several mailing lists (21, 15, and 14 unread emails respectively) and a colleague who the user rarely corresponds with (12 unread emails) would appear at the top of the list.

This type of sort could also be seen as a mechanism to move unsolicited bulk email to the background; unsolicited email is typically sent by those people with whom you do not have an existing relationship.

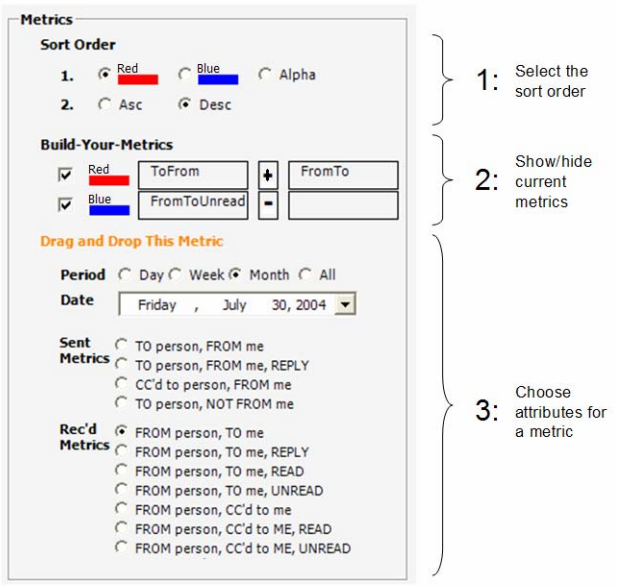

Figure 2: Social metric selection for Email Snarf.

Finding Emerging Topics Using Changes

Our second scenario illustrates how to sort people to bring emails containing emerging topics to the forefront during email triage. Emerging topics typically contain a larger than usual number of emails from a particular person or group of people and signal an issue or problem. In the user interface, the first metric could be used to produce a relationship indicator, similar to the previous scenario. The second metric could show the difference in the amount of email received from people between yesterday and today. Large changes may indicate an emerging discussion or problem topic. If people were sorted by the relationship indicator metric, it would be easy to spot large changes in the amount of email since yesterday for those correspondents with strong social relationships to the user. On the other hand, if people were sorted by the change in the number of emails received, people involved in hot topics would rise to the top of the list even if they weren't people with strong relationships to the user.

Selecting and Viewing Person-Centric Meta-Data

We now describe the multiple ways in which users can select and display social meta-data about their email correspondents. We do not claim that the sorting options we have enabled are the best for everyone and every triage situation. Instead, we argue that when using our interface, users will be able to use social sorting options that are not readily available in existing user interfaces to gain alternative views on their email that will aid email triage.

Social Sorting and Selection of Metrics

Clicking the arrow shown in the top right corner of Figure 1 provides a slide-out window (Figure 2) where users are able to change the social metrics being displayed. The visualization can show up to two social metrics per person:

a red metric and blue metric. Section 1 of Figure 2 allows users to sort people by the red metric, blue metric, or alphabetically by name in either ascending or descending order. Sorting can also be performed in Figure 1 through a right-click context menu.

Section 2 of Figure 2 shows the currently selected metrics in one of four rectangles. Hovering over a given metric name provides a tool tip showing that metric's attributes in more detail (period, date, and metric type). The red and blue metrics can be used to display either a single metric or an arithmetic operation (addition, subtraction, multiplication, or division) on any two metrics.

Section 2 shows us that the blue metric bar will display only the metric **FromToUnread** (the number of unread emails from each person). On the other hand, next to the red metric bar, a plus sign is shown between the **ToFrom** metric and the **FromTo** metric. With this selection, the red bar in Figure 1 would show the sum of the number of emails sent to the contact and the number of emails received from the contact. Clicking the current arithmetic operator causes the operator to cycle through each possible operator. By allowing arithmetic operations on metrics along with the selection of dates and time periods for each metric, it is possible to show the change over time for any of the available metrics. The checkboxes next to the red and blue bars in Section 2 toggle the visibility of each metric in Figure 1. This lets the user sort people by their social relationship and show a single metric of interest like FromToUnread; potentially reducing visual clutter.

The current metrics shown in Section 2 can be changed using a simple drag and drop metaphor. Dragging a metric out of one of the four metric boxes and dropping it will remove the metric. Dragging metrics between boxes moves the metric, overwriting existing metrics in the drop box. New metrics are selected by first choosing the desired metric attributes in Section 3 of Figure 2. Users pick the time period (day, week, month, or all of one's email), a date for which the time period is centered around, and one of the eleven metric types. For example, today's date and the month time period would show a metric for the current month. Once all of the attributes are selected, users drag the orange "Drag and Drop This Metric" label into one of the four metric boxes.

Viewing Individual Emails

Users are also able to view actual emails for each person. Left clicking a person in Figure 1's list opens a message window for the given correspondent. The message window shows a list of all emails matching the contact's metrics along with summary information of the selected social metrics. Double clicking an email will open it in an Outlook message window where full email operations are available (e.g., delete, reply, forward).

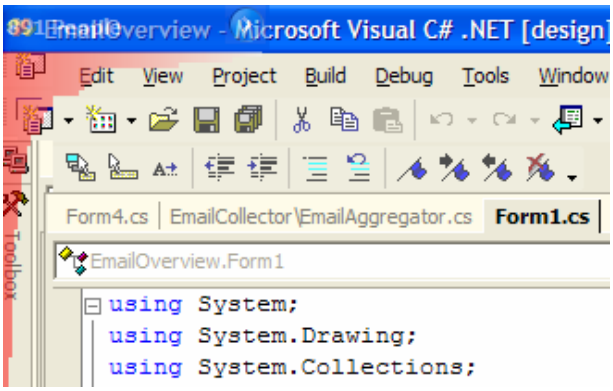

Figure 3: Using Email Snarf as a peripheral awareness tool.

Multiple Levels of Detail

Users are able to interact with the visualization in Figure 1 to view multiple levels of detail, similar to Rao and Card's Table Lens [11]. A *high detail* view is shown in Figure 1 for the first ten correspondents. Here users see the name of the correspondent, bars representing each metric, and a textual count for each metric. A *low detail* view or overview is shown in Figure 1 for a group of correspondents in the middle of the list, as well as at the end of the list. Here users see just the bars representing each metric. Using the low detail view it is possible to get a general sense of metric levels for a large number of correspondents at once. Right clicking a row in the list toggles that correspondent's view between high and low detail levels. A right-click context menu provides options to change the detail level of the entire contact list.

Peripheral Awareness

We designed Email Snarf to provide either direct information about one's email (as already discussed) or a peripheral awareness of one's email environment. To support peripheral awareness, similar to existing applications like Cadiz et al.'s SideShow [3], we provide transparency options along with a position-locking feature. Figure 3 shows Email Snarf locked in the upper left corner of the user's desktop showing unread email counts at a low detail level. A low opacity level allows the user to see her current work and receive a peripheral awareness of current unread emails. This feature provides users with a mechanism to see at-a-glance how "inflamed" their email has become while remaining in their current context.

FUTURE WORK

We plan to continue to explore two main avenues in this project. First, our design work will include a usability study as well as an increased number of interface features. Natural interface extensions would include automatic and user controlled grouping of people into social contexts such as current projects or activities. This idea is similar to using folders and associated rules, yet would use dynamic social filters based on the changing nature of relationships identified by our social metrics. We also plan to include

user pinning of favorite people for whom metrics and emails would always be visible.

Second, we wish to extend our current set of metrics and investigate the role social metrics play in real email usage through a field deployment of our prototype. While we have informally observed which metrics generate interesting and useful social indicators, we plan to validate these hypotheses with actual use. Our user interface also presents interesting scenarios for email usage beyond email triage which we wish to pursue. For example, our current set of social metrics can provide users with information about who they have not sent email to in a long time, or disparity levels in email exchange (who sends them a lot of email that they do not reciprocate).

CONCLUSION

We have presented an investigation of the user needs for email triage in order to inform the redesign of email applications. Our research makes two main contributions. First, we present a thorough understanding of email triage with the results of contextual interviews and a survey. This work takes a significant step beyond previous work on email triage by articulating the problems people face when triaging their email with existing applications and identifying the role that email volume plays in triage. Moreover, we link these findings to design with a set of implications for the design of applications to support email triage.

Second, we present a prototype design to support email triage based on the implications from our initial investigations. Email is a social activity which is embedded in each user's social context, yet many systems fail to provide social information about one's email that can be leveraged by the user. Our design presents sortable social meta-data which people can use to view and triage their email. We make no claim that our user interface is without its usability flaws, nor that it is a full fledged email client. This is our first prototype visualization for viewing social meta-data for email triage and, as such, still has its share of flaws. What we do provide, however, is a look at how one could design a tool to support email triage where social information brings salient social relationships and emails to the forefront.

ACKNOWLEDGMENTS

We thank our survey and interview participants, Gina Venolia, Anthony Tang, members of the MSR Community Technologies group, and members of the SIS project.

REFERENCES

1. Bälter, O., and Sidner, C., Bifrost Inbox Organizer: Giving users control over the inbox, *Proc. Nordic CHI 2002*, ACM Press (2002).
2. Bellotti, V., Ducheneaut, N., Howard, M., and Smith, I., Taking Email to Task: The Design and Evaluation of a Task Management Centered Email Tool, *Proc. CHI 2003*, CHI Letters 5(1), ACM Press (2003), 345-352.
3. Cadiz, JJ, Venolia, G.D., Jancke, G., and Gupta, A., Designing and Deploying and Information Awareness Interface, *Proc. CSCW 2003*, ACM Press (2002), 314-323.
4. Ducheneaut, N., and Bellotti, V., E-mail as Habitat: An Exploration of Embedded Personal Information Management, *Interactions*, 8(5), ACM Press (2001), 30-38.
5. Dumais, S.T., Cutrell, E., Cadiz, J.J., Jacke, G., Sarin, R., and Robbins, D.D., Stuff I've Seen: A system for personal information retrieval and re-use, *Proc. SIGIR 2003*, ACM Press (2003).
6. Gwizdka, J., Email Land – An Exploration of Email User Interfaces Supporting Pending Tasks, *Conference Companion of UIST 2002*, ACM Press (2002), 9-10.
7. Horvitz, E., Jacobs, A., and Hovel, D., Attention-Sensitive Alerting, *Proc. UAI '99*, (1999), 305-313.
8. Kerr, B., and Wilcox, E., Designing Remail: Reinventing the Email Client Through Innovation and Integration, *Proc. Of CHI 2004, Design Case Study*, ACM Press (2004), 837-852.
9. Lockerd, A.L., Understanding Implicit Social Context in Electronic Communication, MIT Master's Thesis, 2002.
10. Mackay, W., More Than Just a Communication System: Diversity in the Use of Electronic Mail, *ACM Transactions on Office Information Systems*, 6 (4), ACM Press (1988), 344-352.
11. Rao, R., and Card, S.K., Exploring large tables with the table lens, *Proc. CHI 1995*, ACM Press (1995), 403-4.
12. Smith, H., Rogers, Y., and Underwood, M., Managing Personal and Work Email in the Same Box: Overcoming the Tensions through New Metaphors, *Proc. HOIT 2003*, Irvine, California.
13. Tyler, J., and Tang, J.C., When Can I Expect an Email Response? A Study of Rhythms in Email Usage, *Proc. ECSCW 2003*, ACM Press (2003).
14. Venolia, G.D., Dabbish, L., Cadiz, J.J., and Gupta, A., Supporting Email Workflow, *Microsoft Research Technical Report*, MSR-TR-2001-88 (2001).
15. Venolia, G.D., and Neustaedter, C., Understanding Sequence and Reply Relationships within Email Conversations: A Mixed-Model Visualization, *Proc., CHI 2003*, ACM Press (2003).
16. Viégas, F., boyd, d., Nguyen, D., Potter, J., and Donath, J., Digital Artifacts for Remembering and Storytelling, *Proc. HICSS-37 2004.*
17. Whittaker, S., and Sidner, C., Email overload: exploring personal information management of email, *Proc. CHI '96*, ACM Press (1996), 276-283.