
The Social Network and Relationship Finder: Social Sorting for Email Triage

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

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

Abstract

Email triage is the process of going through unhandled email and deciding what to do with it. This process can quickly become a serious problem for users with large volumes of email. Studies have found that people use a variety of approaches to triage their email, many of which have a social component. We believe that email clients can better support email triage by providing users with additional sorting features based on socially salient information. We present a prototype email client, SNARF (the **S**ocial **N**etwork and **R**elationship **F**inder), that aggregates social meta-data about email correspondents to aid email triage. Users can then sort their correspondents based on this meta-data bringing emails from socially important people to the foreground.

1 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 one aspect of email management, *email triage*: the process of viewing unhandled email and deciding what to do with it. Most people spend a considerable portion of time triaging their email. Those who receive large volumes of email are forced to triage their email more frequently and spend more time triaging throughout the day; failure to do so can result in an overwhelming feeling of information overload (Neustaedter, Brush, and Smith 2005).

While most people can tell the difference between strangers and the people they know well, unfortunately tools like email clients lack this basic capability. Most email clients provide only a limited set of tools to help people efficiently triage their email using information such as who it is from, when it was received, and the subject. The most important emails given one's current context can be easily buried in the inbox and hard to find. This limits the email client's utility for helping users deal with the increasingly challenging task of managing growing volumes of content along with spam and related email fraud problems.

For this reason, we have designed a prototype email tool called the **S**ocial **N**etwork and **R**elationship **F**inder (SNARF). SNARF is a data mining engine for personal communications that calculates a range of *social metrics* for every email author found in a collection of email, e.g., the number of emails sent to a particular person, the number of unread emails from a correspondent. Our design then visualizes people along with their corresponding metrics and provides a feature for *social sorting*: people and their emails can be sorted using the social metrics to bring emails of likely importance to the top. We feel our prototype takes the necessary first steps at empowering users with tools they can use to efficiently and effectively triage large volumes of email.

We first discuss existing research on email triage and designs that have also looked at augmenting email with social information. Next, we briefly discuss findings from our own studies of email triage that articulate the strategies people use when triaging their email and point towards design directions. Following this, we discuss in detail how SNARF aggregates, visualizes, and sorts email using social meta-data to enhance the email triage experience. We conclude with a discussion of future avenues of research suggested by our work.

2 Related Work

There has been a great deal of work in redesigning email interfaces for various tasks. Email overload has been a recognized problem within the HCI community for a comparatively long period (Whittaker and Sidner, 1996). Ducheneaut and Bellotti (2001) highlight the important task aspects of email, which emphasizes the urgency of helping users gain control over aspects of their inbox. Similarly, Tyler and Tang (2003) report 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.

Whittaker and Sidner point out that there is no “one-size fits all” solution: users have a variety of strategies for handling their mail. One general strategy for handling email triage is to reduce the number of items to triage by grouping emails into conversational threads, IBM’s REMAIL (Kerr and Wilcox, 2004) collects message threads together, and highlights the participants in their messages. The REMAIL interface also allows users to quickly scan for new messages by organization and by presence within a contact list. Venolia and Neustaedter’s Grand Central (2003) and, more recently, Google’s “GMail” also provide a direct interface to message threads.

While a threading approach is valuable, a great deal of email comes in over many threads making additional metadata useful. Dourish et al.’s Presto (1999) made metadata about documents a searchable and examinable aspect of the operating system. Email was treated as a series of individual documents; thus, each could be individually searched by that metadata.

The importance of social metadata in particular was noted by Venolia, Dabbish, Cadiz, and Gupta (2001), who found a variety of factors that users felt made a message important, including whether it was a reply to the users’ message; whether it was from the users’ manager, project members, or subordinates; and whether the message was directly addressed to the user.

Several research prototypes use social information to present new interfaces to email. Lockerd’s (2002) 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. DriftCatcher does not, however, provide any ability to sort or filter email

based on these attributes, and so emphasizes background awareness over an active triage approach.

Balter and Sidner’s (2002) Bifrost similarly groups clusters of email messages by whether they come from “VIP” people (a pre-selected list), and then whether they are sent to the single user or to a small or large distribution list.

Perhaps the most closely related approach to social metadata was that in the Priorities system (Horvitz, Jacobs, and Hovel, 1999). Priorities uses a machine learning technique over several meta-level attributes to determine the importance of new unread emails. Priorities is also focused on email triage, and chooses a single ordering to present to the user based on their previous activity. Neustaedter, Brush, and Smith (2005) argue that users may want a variety of different orderings depending on their social context.

These approaches to social metadata all try to highlight important people. However, they each provide either a fixed ordering of messages (as in Bifrost) or no ordering at all. Our project is meant to explore other approaches to making social metadata available to the user.

3 User Needs for Email Triage

Our design avenues for email user interfaces are largely influenced by our own studies of email triage. For this reason, we outline our main findings from a series of interviews and a survey where the goal was to understand the strategies people employ when triaging their email and how email interfaces could be improved to better support triage. Additional details can be found in Neustaedter, Brush, and Smith (2005).

We found that, depending on one’s current context, people mainly triage their email using either a *single pass* strategy, or a *multi-pass strategy*. In the single pass strategy, each email is visually scanned once in the inbox (top-down, or bottom-up). While efficient, the problem with this approach is that people often are not handling the most important emails first. With a large amount of unhandled email, this could cause time-sensitive or “important” emails to be handled too late or not at all if triage times are limited. In the multi-pass strategy, the user performs multiple single passes over their email, handling a certain type of email during each pass. This strategy is typically less efficient as users must scan the same email multiple times. It does, however, allow users to deal with “important” emails first.

Contrary to our original expectations, we found that rather than handling important emails first, many people handle “unimportant” emails first because they can quickly delete or file them. Often once these emails are handled it is easier for users to find the important emails. We also found that social information was vital for determining the importance of an email. For example, recent emails or emails from someone with a close personal relationship (e.g., close colleagues, direct managers, or significant others) were typically quite important. Just the same, emails from new social contacts working on similar projects were important. Less-recent emails, news-related items, or email from people with a lesser relationship (e.g., someone for whom the user did not typically send replies to) were typically not as important to users. These results verify and extend findings by Venolia et al (Venolia, Dabbish, Cadiz, and Gupta, 2001).

These findings suggest that email interfaces currently lack methods for users to easily find important messages causing them to choose inefficient triage strategies. This also points to the need for email interfaces to incorporate additional socially salient information about one’s emails and correspondents. Such information could empower the user with searching or sorting features that bring emails of likely importance to the forefront.

4 Collecting Social Metrics for Email

The Social Network and Relationship Finder (SNARF) for email is a data mining engine and visualization tool for personal communications. In this section, we discuss how SNARF collects and calculates socially salient information about email.

SNARF collects and aggregates *social accounting meta-data*: social properties for one’s email correspondents. One aspect of meta-data is quantifiable measures we call *social metrics*. These measurements capture multiple dimensions of the relationship between the user and their correspondents and among the correspondents themselves. Measures like the number of times an author sent mail over a time period, the number of those messages that were replies, and the number of those messages that remain unread can be used for supporting email management.

SNARF’s social metric collector is based on an existing search project called Stuff I’ve Seen or SIS (Dumais et al 2003). 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 1: rows 1-4) and *received metrics* (Table 1: 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 correspondent 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**.

| | 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 <i>read</i> | who sends the most/least emails to the user that are read |
| | 8 FromToUnread | emails to the user and marked <i>unread</i> | 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 <i>read</i> | who includes the user in the most/least conversations that the user had read |
| | 11 FromCCUnread | emails CC’d to the user and marked <i>unread</i> | who includes the user in the most/least conversations that the user has not read |

Table 1: Sent and received metrics aggregated for each correspondent.

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 social metrics, we feel they offer a compelling set of social information to enhance email triage. Next, we show how SNARF's social metrics can be used for email triage.

5 Performing Email Triage

The social accounting metrics that SNARF aggregates provides resources for *social sorting*—the re-ranking of collections of email based on attributes like “the number of times that this author has replied to me and I have read their message” or “the number of messages I send to this author that are replies.” These sorting metrics allow for more effective email triage by sorting messages to the top of lists based on the nature and strength of the relationship between the user and each correspondent present in their email. For example, a common user challenge is returning to a large collection of messages after being away from email for some time. Often faced with limited time for triage, users are forced to scan, select, and often initially move or delete camouflaging content that obscures the “good” email.

For this reason, SNARF visualizes the person-centric social meta-data that it aggregates. We begin by describing several scenarios that illustrate a variety of social metrics and sorting features found in SNARF that can support email triage.

5.1 Emails from Socially Important People

Our first scenario illustrates how to sort people to bring emails from socially important people to the forefront. Here we are considering people of social importance over all of one's email, not just people of recent importance. 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.

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




Figure 1: SNARF visualizes people and their corresponding social metrics.

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 (red bar on top and blue on the bottom)

In Figure 1, people are sorted by the red metric bar which shows the number of messages sent from the user to each person (ToFrom) 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 (FromToUnread). 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. 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 unread emails from socially important people. 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.

5.2 Emails from People of Recent Importance

Our second scenario illustrates how to sort people to bring emails from *recent* socially important people to the forefront. Social importance is not static. Rather, it changes over time given one's social context, e.g., work projects, social events. Thus, it is often desirable to easily find emails relating to current projects or upcoming meetings. For example, the same sorting metric from the first scenario—the number of messages sent by the user to a person (ToFrom)—could be used but only considering a recent time period, say the last month. The selection of a current time period will bring people of *recent* importance to the top. A second metric, like the previous scenario, could show the number of unread emails from each person (FromToUnread). Often people will exchange a series of emails several days before an upcoming meeting. If one needed to quickly triage their email for new emails of this type, they could perform the previous sort but for the current week.

5.3 Finding Strangers

Our third scenario illustrates how to sort people to bring emails from strangers to the forefront. These may be individuals who are interested in collaborating on a current project, want to simply establish a relationship, or people sending spam. Here a reverse sort on the number of emails received from a person (FromTo) would bring those people who have only sent the user a single email to the top of the list. Again, a second metric for unread email would ensure only new unread email from strangers was visible at the top of the list. With this display, users could quickly view, handle, and remove emails from strangers, many of which might potentially be unimportant or spam.

5.4 Finding Emerging Topics Using Changes

Our fourth 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 often signal an issue or

problem. In the user interface, the first metric could be used to produce a relationship indicator, similar to the previous scenarios. 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.

5.5 Ensuring Email Reciprocity

Our fifth scenario illustrates how to sort people to bring emails from people where a social imbalance exists to the forefront. At times, people may receive more email from an individual than they sent that person. This type of situation may occur between a manager and her team members or even between people working on a project together. When large imbalances occur there is often a need to rectify the imbalance. For example, a problem may have occurred or a particular aspect of a project may need additional attention. Users can find relationship imbalances by sorting people according to the difference in the number of emails sent from each person (FromTo) and the number of emails sent to each person (ToFrom) for a current time period. This will bring the people who have sent more mail to the user than they have received to the top of the list.

6 Visualizing Social Metrics for Email

We now describe the multiple ways in which users can select and display social meta-data about their email correspondents using SNARF. 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.

6.1 Social Sorting and Selection of Metrics

We provide users with a default set of selectable sorts based on the scenarios in the previous section. In addition, users are able to explore their email on many social dimensions by selecting their own metrics and time periods. 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.

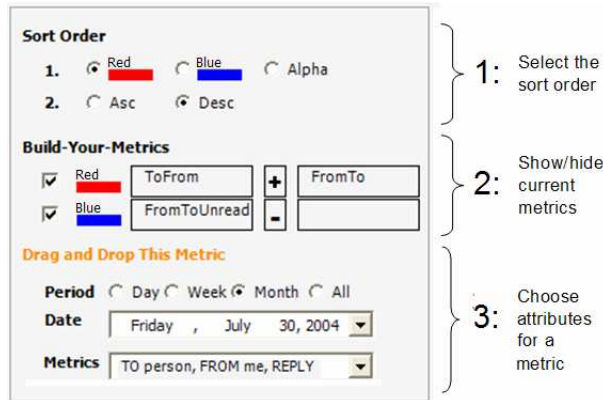


Figure 2: Social metric selection for SNARF.

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. While our visualization appears similar to Bellotti et al's TaskVista (2003), we are providing social sorting options whereas they are clustering emails into tasks. They also do not provide the various social attributes that our work is centered around.

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 a metric type. 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.

6.2 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 (not shown). 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).

6.3 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 (1995). 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 and also in Figure 3 (red bars on left side). 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 of contacts (Figure 1) toggles that correspondent's view between high and low detail levels. A right-click context menu provides

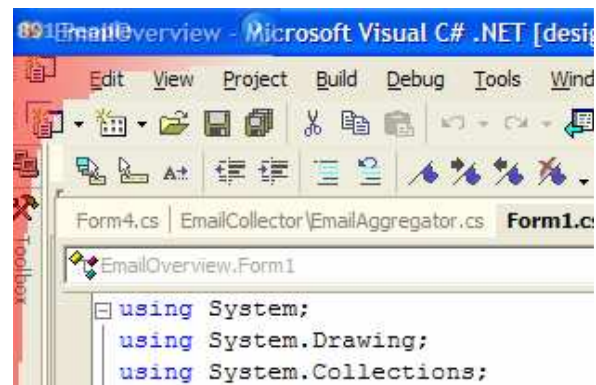


Figure 3: Using SNARF for peripheral awareness.

options to change the detail level of the entire contact list.

6.4 Peripheral Awareness

We designed 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 (2002), we provide transparency options along with a position-locking feature. Figure 3 shows 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.

7 Discussion

While our prototype has yet to be formally evaluated, we feel *social sorting* of email meta-data provides a novel mechanism to aid users when triaging their email. One could even imagine using some of the social accounting metrics that we provide as novel features for the relevance functions of search engines over threaded conversations (and other socially generated content). Search might bring messages with certain keywords to the top of a relevance ranked list because a keyword appeared frequently. Using SNARF-generated metrics, messages can first be selected because they contain a keyword or string and then further re-ranked based on the strength of the relationship (either strong or weak) between the author and the correspondent.

Our work has also opened up avenues for additional visualizations and enhanced user interfaces that make use of the extended attributes and connective structures among email correspondents. For example, a huge volume of our communication events take place in email and yet there are almost no means for gaining a holistic overview of these patterns. Using the social accounting metadata SNARF generates, a range of information visualizations is possible. These could provide insight into the connections between correspondents and their temporal patterns of activity. Moreover, they could facilitate relationship management by allowing people to easily find past colleagues who may be important people to ask for information about new work opportunities, finding new candidates for positions, or marketing new businesses or products.

Email interfaces could also provide improved modes of interacting with large collections of communication

events. Currently the dominant model for email management is the “message browser”—the basic unit of data is an individual message, not conversation threads and certainly not relationships. Using SNARF-generated metrics, user interface components could make a more holistic unit of manipulation available to the user: threads and collections of threads from shifting populations of correspondents become tractable and enable user interfaces to highlight latent but important dimensions of the conversational data. These data include *audience awareness* (Who is on this thread? Who has just joined?), *time sensitivity* (How long do I usually take to reply to this person or group? How long do they normally take to get back to me? Should I resend this request?), *relationship awareness* (What is my relationship to this person? Do we ever appear on the same email together even if they have never directly sent me email?).

8 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).

9 Conclusion

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 to efficiently triage email. The Social Network and Relationship Finder (SNARF) presents a first step towards addressing this problem. SNARF provides sortable social meta-data which

people can use to bring important social relationships and corresponding emails to the forefront. These social sorting options are not readily available in existing email user interfaces.

We make no claim that our user interface is without its usability flaws, nor that it is a full fledged email client that deals with all of one's email problems. This is our first prototype visualization for viewing social meta-data for email triage and, as such, it still has its share of flaws. What we do provide 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. Moreover, our work has unveiled future design possibilities for email applications that utilize social accounting meta-data.

Acknowledgements

We thank Gina Venolia, Anthony Tang, members of the MSR Community Technologies group, and members of the Stuff I've Seen project.

References

1. Bälter, O., and Sidner, C. (2002). Bifrost Inbox Organizer: Giving users control over the inbox. *Proceedings of Nordic CHI 2002*, ACM Press.
2. Bellotti, V., Ducheneaut, N., Howard, M., and Smith, I. (2003). Taking Email to Task: The Design and Evaluation of a Task Management Centered Email Tool. *Proceedings of the ACM Conference of Human Factors in Computing Systems (CHI 2003)*. CHI Letters 5(1), ACM Press, 345-352.
3. Cadiz, JJ, Venolia, G.D., Jancke, G., and Gupta, A. (2002). Designing and Deploying and Information Awareness Interface. *Proceedings of the ACM Conference of Computer-Supported Cooperative Work (CSCW 2003)*. ACM Press, 314-323.
4. Dourish, P, Edwards, W. K, LaMarca, A., and Salisbury, M. (1999) Presto: an Experimental Architecture for Fluid Interactive Document Spaces. *ACM Transactions on Computer Human Interaction*. 6(2), p133-161. New York: ACM Press.
5. Ducheneaut, N., and Bellotti, V. (2001). E-mail as Habitat: An Exploration of Embedded Personal Information Management. *Interactions*, 8(5), ACM Press, 30-38.
6. Dumais, S.T., Cutrell, E., Cadiz, J.J., Jacke, G., Sarin, R., and Robbins, D.D. (2003). Stuff I've Seen: A system for personal information retrieval and re-use. *Proceedings of SIGIR 2003*, ACM Press.
7. Horvitz, E., Jacobs, A., and Hovel, D. (1999). Attention-Sensitive Alerting. *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI 99)*, 305-313.
8. Kerr, B., and Wilcox, E. (2004). Designing Remail: Reinventing the Email Client Through Innovation and Integration. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2004)*, Design Case Study. ACM Press, 837-852.
9. Lockerd, A.L. (2002). Understanding Implicit Social Context in Electronic Communication, MIT Master's Thesis.
10. Neustaedter, C., Bernheim Brush, A.J., and Smith, M. (2005). Beyond "From" and "Received": Exploring the Dynamics of Email Triage. *Extended Abstracts of the Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2005)*, ACM Press, to appear.
11. Rao, R., and Card, S.K. (1995). Exploring large tables with the table lens. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 1995)*, ACM Press, 403-4.
12. Tyler, J., and Tang, J.C. (2003). When Can I Expect an Email Response? A Study of Rhythms in Email Usage. *Proceedings of the European Conference on Computer Supported Cooperative Work (ECSCW 2003)*, ACM Press.
13. Venolia, G.D., Dabbish, L., Cadiz, J.J., and Gupta, A. (2001). Supporting Email Workflow, *Microsoft Research Technical Report*, MSR-TR-2001-88.
14. Venolia, G.D., and Neustaedter, C. (2003). Understanding Sequence and Reply Relationships within Email Conversations: A Mixed-Model Visualization. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2003)*, ACM Press.
15. Whittaker, S., and Sidner, C. (1996). Email overload: exploring personal information management of email. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 96)*, ACM Press, 276-283.