Bifrost Inbox Organizer: Giving users control over the inbox

Olle Bälter

Royal Institute of Technology

Stockholm, Sweden

balter@nada.kth.se

Candace L. Sidner

MERL, a Mitsubishi Electric Research Lab[1]

Cambridge, MA, USA

sidner@merl.com

Abstract

Many email users, especially managers, receive too many email messages to read in the time available to them. The solutions available today often require programming skills on the part of the user to define rules for prioritizing messages or moving messages to folders. We propose a different approach: provide an inbox with several distinct categories of interest. These categories result from the use of a set of general rules that are pre-defined and which can be customized by individual users. We also report on 2.5-D inbox designs to offer users better overviews of their inboxes.

Introduction

Many email users, especially managers, receive too many email messages to read in the time available to them. The solutions available today often require programming skills on the part of the user to define rules for prioritizing messages or moving messages to folders. Constructing and maintaining rules is time consuming, as well as difficult for many users who do not have sufficient programming skills. Furthermore, the prioritization rules cannot meet the needs of users because user priorities can change significantly in a short period of time. For example, after a person talks with a colleague at the watercooler, she could have a new priority for checking her email, but when she meets her manager, who has a new task for her, the priorities can change again. A remedy for these problems might be to use agents that learn how the user would like to read messages, but this solution causes other difficult tasks for users: they must give feedback to the agent to change its behavior, and to do so, they must understand how the

1 [1]This work was undertaken while the author was employed at Lotus Development Corporation.

agent does what it does. Also, the people who have the greatest need for a prioritizing system are also the ones that have least time to spend on giving feedback. Finally, regardless of how the prioritization rules are constructed, moving unread messages to folders automatically causes an "out of sight, out of mind" email situation, that is, email that is not in the inbox tends to be ignored and become forgotten.

We propose a different approach: provide an inbox with several distinct categories of interest. These categories result from the use of a set of general rules that are pre-defined and which can be customized by individual users. These general rules liberate users from reinventing and maintaining situationally relevant rules for handling their email because the user does not need to develop the rules; the user only tailors them via a form to his or her use. The rules we propose are derived from user studies and experiments with a prototype named the Bifrost[2] InboxOrganizer. Our experience with Bifrost, which we will discuss in this paper, indicates that this inbox not only liberates users from rule construction but provides them with a breakdown of their email into categories of interest. These categories allow users to quickly zero in on the email that is most relevant to a particular work task and not to be distracted by email that concerns other tasks. We also believe this inbox will be valuable to users accessing their email over the telephone.

Previous research

Information overflow, or overload, can be described as "information received at such a rapid rate that it cannot be assimilated" (Sheridan & Ferrell 1974). The phenomenon of information overflow is not new. Yates (1989) quotes a manager reacting to information overload in 1920:

I do not think it necessary to send these reports to me in the future unless to draw attention to some peculiar or abnormal condition. I shall depend on you to keep the inspections going, but the clerical work of making out the reports can be saved. If there are any other similar reports which you think can be cut out please give me an expression of your opinion with reference to the same (p 191).

When it comes to email overflow, Hiltz and Turoff (1985) noted that it causes users to answer only parts of the incoming mail, to ignore incoming information systematically, and even to stop using the email system. Mackay (1988) found that the feeling of being overloaded varies widely, independent of the number of messages they send and receive. Lantz (1996) describes how messages from customers and error reports are prioritized by the management, while other messages had to wait, especially if they needed a longer informative answer. This process sometimes resulted in messages that were forgotten.

Whittaker & Sidner (1996) defined three main groups of email users based on folder usage and cleaning frequency: Frequent Filers, Spring Cleaners, and No Filers. While Frequent Filers still managed to maintain order among their messages, Spring Cleaners did so only on rare occasions and No Filers gave up completely. Bälter (1997, 1998) suggested that these groups are a result of a natural development from Beginner to Frequent Filer to Spring Cleaner and finally to No Filer.

---

2 [2]Bifrost is the name of the rainbow colored bridge between this world and Valhalla, the Viking conception of Heaven. Only the righteous may pass this bridge, in this case referring to the messages that should be brought to the user's attention.

We conclude from these observations that the different user groups need different types of support to maintain control over their inboxes.

Managers have specific overload problems. Several studies have shown that managers are under time pressure (Carlson 1951; Stewart 1967; McCall, Morrison & Hannan 1978). Other studies show that managers received more email messages than others (Whittaker & Sidner 1996, Bälter 1997) and commonly had problems handling email (Lantz 1996).

A remedy for overloaded managers may be to delegate more to their subordinates, but delegation is a complicated task (Milewski and Lewis 1997). Managers fear that quality will decrease, that they will lose control and get less credit, that their subordinates will fail, or that they may be perceived as tyrants by the subordinates. Also, some tasks are a pleasure to perform, and managers may be reluctant to delegate them. All these characteristics work against delegation.

(Savendy 1987) suggested reducing email overflow using filters to store messages in folders or to prioritize them in the inbox. This change could reduce the stress involved in decision making tasks that are lower for sorted email than unsorted

Many of the ideas for filtering email originated in the Information Lens systems developed at the Massachusetts Institute of Technology. The filtering system was designed both to save users from junk mail and to find messages of interest, even though the messages were not directed to the user originally (Malone, Grant, Turbak, Brobst & Cohen 1987, Bannon 1993).

Filtering rules have to be defined in some way. This may be simple for programmers, but not for ordinary users (Boone 1998). The help system Advisor in the Andrew Messaging System exemplifies the difficulty in creating filtering rules. Advisor was designed to handle questions from users about the system itself. An attempt was made to use a filter to automatically redirect these user messages to the persons responsible for different parts of the help system (for example, email messages about “mail" went to someone who only answered questions about mail). However, more than 50% of the messages ended up as unsortable, because most users did not specify the problem enough in the subject line and often simply wrote e.g. "Help!" (Borenstein and Thyberg 1991).

Automatically moving messages into folders before the user has viewed them often makes the messages invisible. An analogy to the office desktop explains this claim. There are perfectly good reasons why office workers create piles of papers on their desks instead of filing them into folders. The purpose of these piles is not only to store information for later retrieval, but also to remind the worker to do something (Malone 1983). Two quotes from Malone's subjects illustrate the behavior:

If I don't put it here where I can visually see it, I won't do it (p 107).

You don't want to put it [a pile on the desk] away because that way you never come across it again (p 107).

One way to avoid users writing programs is to design the filtering system in such a way that the system can learn from the user how to sort messages. Losee (1989) developed a formal model based on economical and statistical decision theory to rank messages on a scale of interest. The model assumes that each message contains certain features such as author, origination time, subject, keywords, category, and recipients of the message. The user provides feedback by

classifying messages as relevant or non-relevant. The features and the feedback are used in a Bayesian artificial neural network that considers prior and new knowledge in order to provide the user with a more accurate ranking in the future. Segal and Kephart (1999) developed an adaptive classifier that predicted most likely folder based on a collection of messages previously placed in folders by the user. This classifier worked well for archived messages, but helped only to archive messages after reading them, not to prioritize reading.

Similar suggestions have been made by Shet & Maes (1993) and Maes (1993, 1997), who propose intelligent agents that learn from users by training, imitating users' actions, and receiving negative feedback when it takes the wrong actions. These agents work as a complement to user defined filtering rules.

A difficulty with learning filtering systems is that the user may feel out of control if they do not understand why the system acts the way it does. Also, the system cannot learn without feedback from the user and this may consume the time that the user is trying to save. A system that prioritizes messages will not know if the prioritization is right until the user has read all messages, which is exactly what the user is trying to avoid.

Given the preceeding criticism of filtering, what conclusions can we draw about filtering and prioritizing? Would it be possible to define general rules that would be applicable to most or all users?

Palme (1984) proposed structuring the messages in order to achieve control. Email and newsgroups messages could be structured based on *conferences, comment trees, keywords, subject, selection by others, author,* and *abstract writing* according to Palme.

Arensburger & Rosenfeld (1995) suggest the categories *personal, listserv, ccs*, and *others*. Any message with more than 10 recipients is considered a listserv message. Any personal addressed message that is not from one of the specified friends (VIPs) or from within the own organization goes into the category others. Their system is publicly available, but the authors claim:

The rules ... will almost certainly not apply to anyone else.

Marx (1995) defines *timely* messages as messages containing calendar information (in the header) or as responses to recent messages. Boone (1998) describes how users organized their messages into *high priority, low priority, social* and *announcements*.

Current Approach

In the current work rules based on a combination of the Arensberg & Rosenfeld and Marx methods were created and associated with a user's inbox. The user could fill out a form to customize the rules for senders that were of special importance to them. The contents of the inbox was re-ordered using the rules so that messages of different priorities were listed in the inbox by category. No messages were removed from the inbox to other folders unless the user explicitly requested such a rule. Before changing a user's inbox to use these rules, we first undertook a study of their daily email use.

Methods

Eight email users at five different departments of a software company were interviewed and observed during their morning reading of email. Some of the subjects were interviewed when they returned from a business trip or a vacation. The subjects were selected based upon an expected large volume of incoming messages, preferably more than 30 messages a day. The interviews were taped and lasted between 30 and 120 minutes.

Based on the results of these interviews and observations a prototype was developed to categorize email messages in the inbox. This prototype was then used by eleven users for a period of two to seven months. During that time period the prototype was refined by requests from the users several times. Finally the prototype users were interviewed. Six of the initially interviewed subjects were also users of the prototype.

Results of the first user study

From the initial interviews the following observations were made:

- All users but one scanned new messages several times in order to read the most important messages and at the same time delete the least interesting.
- Messages related to events in the calendar for the day are important, regardless of arrival date
- Users marked opened messages as unread as a reminder that they were unfinished tasks.
- In general, the less recipients there are of a message, the more important it is.
- Replies are important as they often contain a solution to a problem posted by the recipient.
- More than half of the subjects did not read all of their messages.
- Some subjects feared that they would miss important messages during their scanning of the new messages.
- Subjects who mentioned filtering feared that filtering would move messages out of sight.
- For most users, carbon copies were judged by their receivers as less interesting than other messages.
- Four of the subjects did not delete any messages with the exception of failed delivery reports.

Details about individual users' number of messages are displayed in table 1.

Table 1. User data from the first round of interviews

Position	Estimated number of new messages per day	New messages at time of interview	Total number of messages	% in inbox	Number of folders	Scans	Deletes messages	Reads all messages
Researcher	30	140	1,257	62%	24	3	Yes	Yes
Researcher	30	30	10,000	95%	8	2	No	No
AA	10-20	10	3,200	10%	50	2	Yes	Yes
Researcher	20-30	35	3,300	28%	150	3	Yes	No
Researcher	200	470	6,340	13%	70	3	Yes	No

Manager	30	19	5,400	100 %	0	2	Yes	Yes
Manager	20-40	15-20	7,000	100 %	0	1	Few	Yes
Manager	30-40	0	5,000	96%	10	2	Yes	No
Manager	100	120	20,000	90%	200	3	No	No

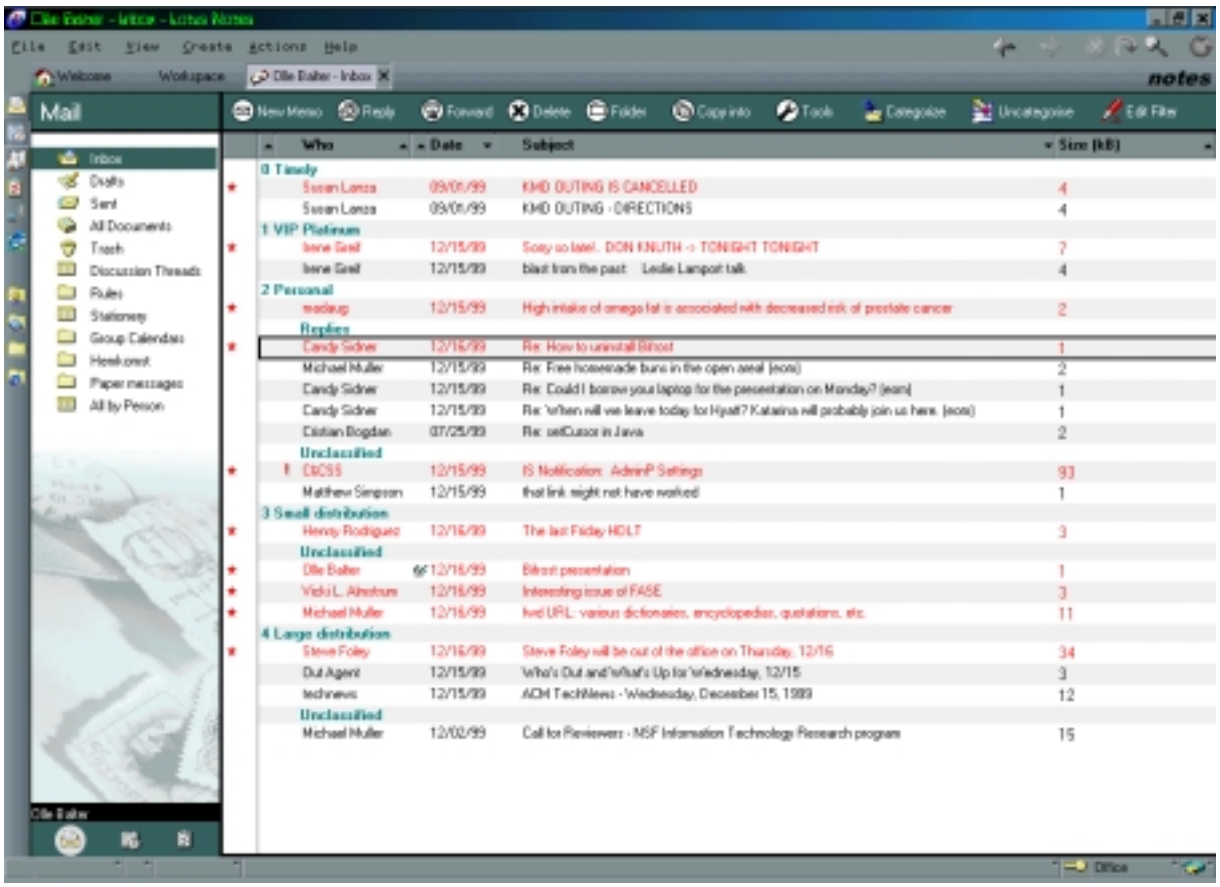

Researcher	30	120	20,000	100 %	6	2	Few	No

Prototype

Messages are categorized by Bifrost in five main categories, as shown in figure 1 where category numbers and headings are presented in addition to the single line descriptions of individual emails.

Figure 1. Categorized inbox.

The first category, *Timely,* contains messages that have a word in the subject or sender in common with an entry in the calendar for today. In the example in figure 1, this category comprises the two messages from Susan Lanza. The message about directions to an outing may be extremely useful on the day of the meeting, but any other time it would be categorized as a *Large distribution* message because the number of recipients is more than seven. The message from Lanza (first in the list) is obviously critical since it indicates that the outing is canceled.

The second category is *VIP Platinum,* and can be used for people whose messages the user always considers as potentially important or urgent regardless of the number of recipients. Users choose such senders, who are dubbed "VIPs."

The *Personal* category contains messages sent exclusively to the user by name. In the example in figure 1 it is subdivided into an unnamed list, a list of "replies," and a list entited "unclassified." If the user does not enter any people in the VIP Gold category (see figure 3), all personal messages will appear as unclassified. The personal category is intended for messages that only the user can handle (hence only one recipient and by name, not some message list). The replies are often answers to questions posted by the mailbox user, or at least messages from people with whom the user has communicated before and whose messages are therefore probably of higher interest than messages sent by others. A subcategory not used in the inbox presented in figure 1 is the calendar related messages. Half of the users used an online calendar and a category was added to collect all calendar related messages. For users that did not use the calendar, this category was small or non-existant.

The *Small distribution* category is intended for group messages. It is divided into named groups and unclassified. Named groups are provided by the user.

The *Large distribution* category is intended for listserv messages and messages with more than seven recipients. It has the same subdivision as Small distribution.

This structure is intended to support users in identifying important messages and in weeding out the less interesting ones. The ordering of the categories implies an order of importance, but in fact any of the categories can become a priority depending on the ever changing events in the user's life.

If the inbox contains a large number of messages, the number of messages in each category will of course be very large, and it will be difficult to get a snap shot of all of them, unless the user deletes messages or store messages in folders. For many users, deleting messages is out of the question since they want to have a record of their communication. Also, deleting messages requires that the user decide that he or she never wants those messages again, a decision many users are not willing to take (Bälter 1998). Folder usage requires naming and maintaining folders, which is a time consuming task, and also requires a decision for each message to store which folder it should go to. Some messages fit several folders and some do not fit anyone, which make storing decisions complicated or even impossible.

However, by categorizing only the unread messages, the number of categorized messages can be kept low and easy to view; see figure 2. Marking messages unread as a reminder of tasks was a common habit among the users in the first study, and this sorting gave these users a to-do list, while the already read messages sunk away to the bottom of the inbox in the "Not Categorized" section.

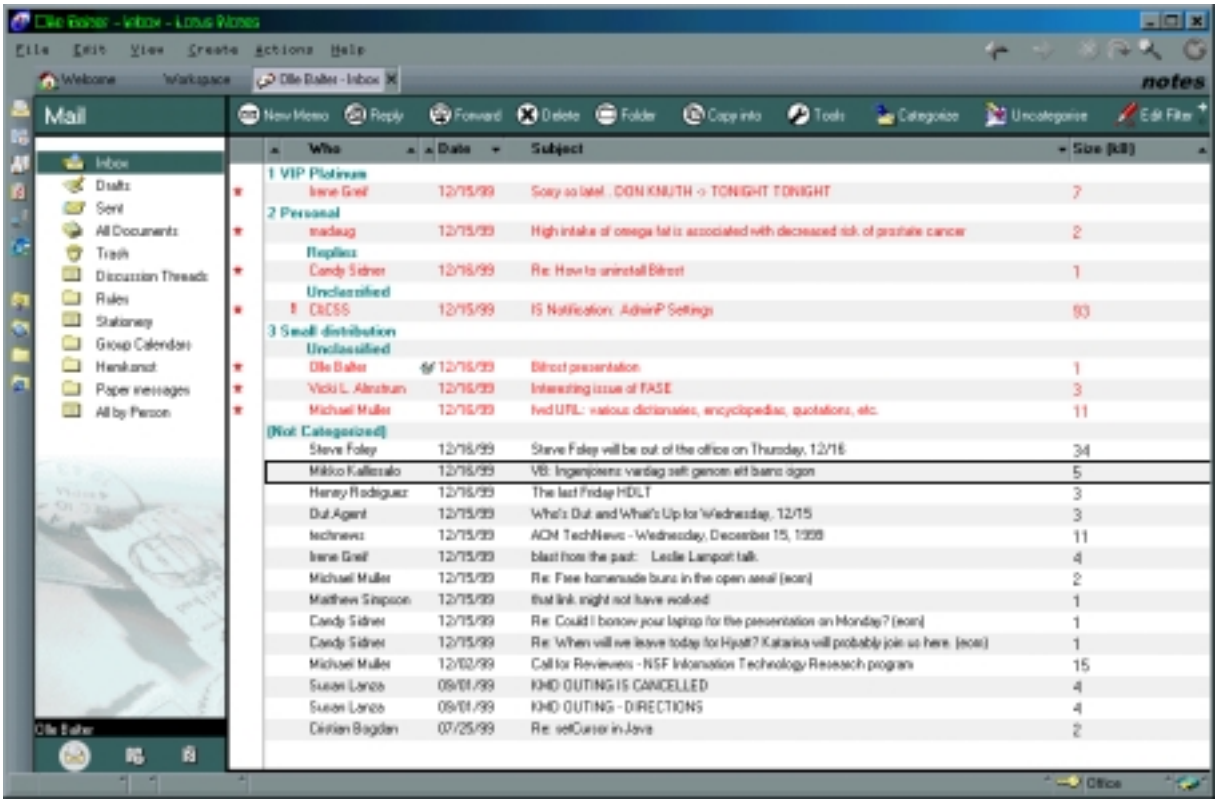

Messages sent as carbon copies are marked with a little icon to help the user identify them (see the messages from Olle Balter in figures 1 and 2). This icon was created because some subjects reported in the first study that cc's were less interesting than other messages.

Figure 2. Inbox with unread messages categorized.

In figure 3, the fields are displayed that are used to improve the precision of the categorizing rules.

Figure 3. User defined fields.

The only mandatory field is *MyNames* where the name of the user must be entered. The *VIP Platinum* field is used for senders whose messages should go into the *VIP Platinum* category. The *VIP Gold* field is used for senders whose messages should be displayed first among the *Personal* or *Small distribution* messages, depending on the number of recipients. The senders and receivers entered in the *Small* and *Large distribution* fields will be displayed first in their category, respectively, while other messages will be displayed as *unclassified* within each category.

Users who found that they trusted the categorization capabilities of Bifrost asked for an *Unwanted* field. Messages from unwanted senders are moved to the trash folder automatically by a rule that triggers on the presence of entries in this field (see rule #1 in figure 4). The Senders # fields in figure 3 can be used to form rules that would group messages from certain senders into a category. This feature was added late during the user testing and was not used by most subjects.

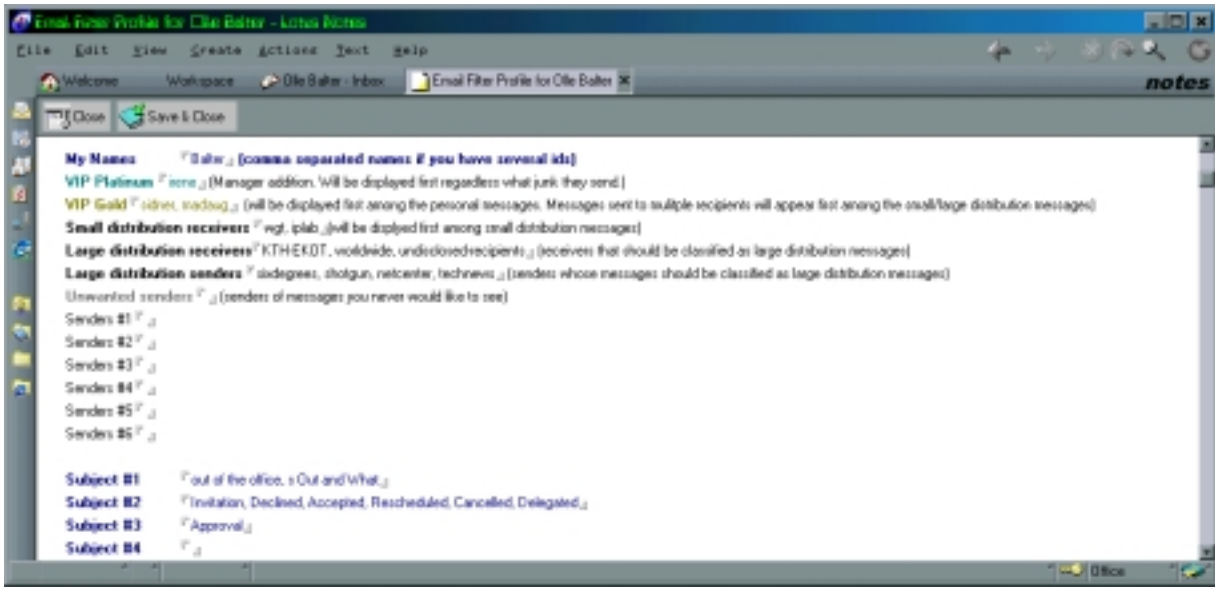

Figure 4. Predefined rules.

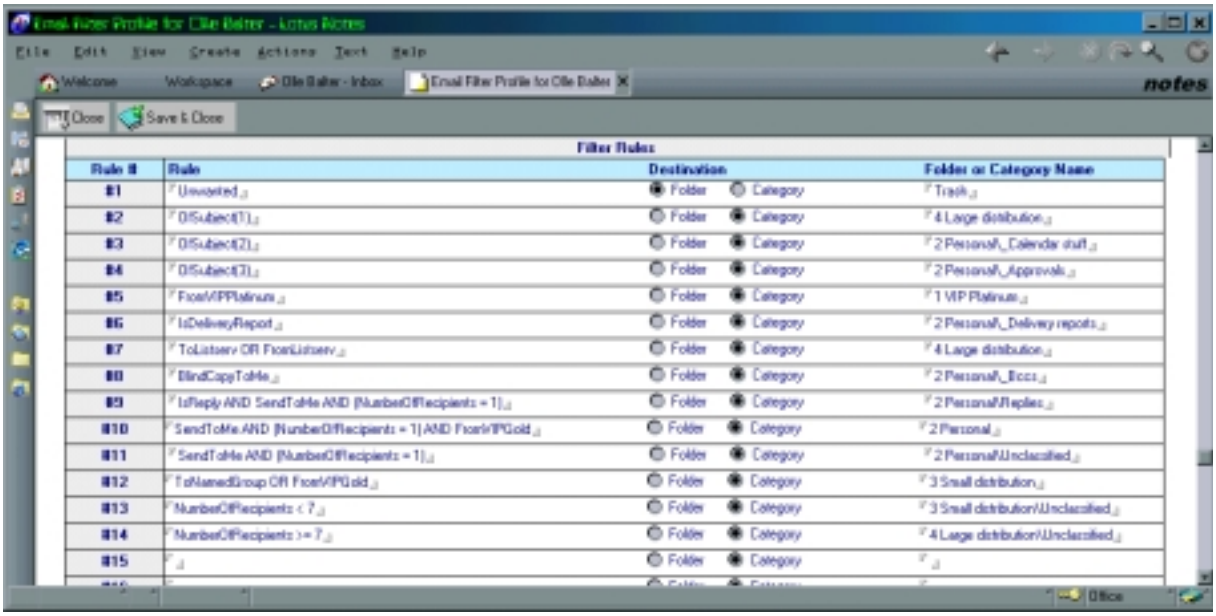

Results of the second user study

Ten of the Bifrost users were interviewed after having access to Bifrost in their ordinary mail system for at least a month. All but one found Bifrost useful. All users were experienced, the median starting year to use email was 1982. Basic user data are displayed in Appendix 1. Several of these users reported being overloaded with email.

> *I look at email negatively. I look at it as a burden, I look at is as a necessary evil for my job, its usefulness has outweighed itself for me. For really truly important things I have to ask people to use voice mail. So I see Bifrost as a way to help me tackle this beast that I* ***depreciate.*** *[sic--means no longer appreciate--the authors]*

Usage

The subjects used Bifrost in two different ways A few used Bifrost as a replacement for the standard inbox and categorized messages daily. A majority used Bifrost categorization when their new email messages had built up to more than a full screen. One user described the effects of email build up. "*It took me two-three weeks to catch up from one weekend of not reading"*. The frequency of categorization varied a lot between users, from twice a week to once a month.

Bifrost is really a tool to help me attack my email.

A reminder that email overload is not caused by the media, but by the people using it was given by a manager when he was asked what he wanted from Bifrost in the future:

A feature that allows me to electronically hunt down and kill people that pointlessly copies [sic] me.

Advantages

Bifrost has two main advantages: it provides users with a reasonable prioritization of their messages, and it simplifies maintenance of a to-do list in the inbox. As a prioritization method, Bifrost changes user behavior: several of the subjects reported that they sifted less through messages or through fewer messages when they sifted.

If I am just running through an inbox, I might be tempted to read a title and get sucked in because it is interesting. Whereas if it is in a pile of listserv stuff I just ignore it altogether. That was a nice thing when I was busy, to not get distracted by unimportant mail.

I can find things that (my manager) sends easily. He usually sends things that I need to pay attention to more completely. I scan certain people's memos while I read other people's memos.

It (Bifrost) really helps me like a machete in a forest to clear a way.

One of the subjects described how a group message (categorized as small distribution) about a meeting was ignored until his manager sent a message (categorized as VIP Platinum) stating that he would like at least one of the people in the group to go to the meeting. The prioritization is of course especially of value when the number of new messages is large and the time to read messages is short.

Bifrost allows me to take things out of the chronology order and deal with them in their importance level.

I can keep a certain amount of information around. For instance, ahm, I have to do the project evaluations, and people are sending me the project evaluations or the links, and I am able to keep them, without putting them into a folder I can put them into a category, and keep them in one place, just as I keep things from my boss in one place, I can keep the things that are necessary for me, which is a group of project evaluations regardless of how many days it crosses over, and clearly visible so I do not have to go and open up, you know, folders to find out where they are, and then of course, inside that folder they are spread all over, so this is a way of, so I find that useable.

Ultimately, I want to keep up with my email. I feel this silly sense of accomplishment if I'm actually on top of it. And I only stay on top of it for a day or two. For two days, maybe. One thing that is kind of neat is, if I am reasonable, a couple of days a go I categorized, by the last week unread, I have gone through the important stuff, I sift through the important stuff, and I keep the not categorized all black, that means for a short period of time when I come in I can deal with them, the more read that are not categorized that's my definition of behind. So I have totally adopted the tool (Bifrost).

The support for maintaining a to-do list works in two different ways. For people who maintained unread marks as a way to remind themselves that a message represents an unfinished task, the Bifrost categorization provided a means for keeping the messages in the priority categories, because read messages (i.e. ones already attended to) move to the bottom of the inbox. One of the users did not bother to mark messages unread, but instead used the Bifrost categorization itself to keep the interesting messages at the top of his inbox. He kept messages that he wanted to have at the top as categorized, and uncategorized the messages manually when he wanted them to fall down.

I sort of keep <messages in the inbox> as a to-do list. Something I always felt I would like to have.

It is a way to give me a "To-do" kind of feature right in my mail file, without having to turn my mail messages into ToDo documents. It helps me to group important messages at the top of my mail file so I can easily find them.

It holds together a stream of thought.

Instead of bouncing all over my inbox, …, I read them a little more consecutively… I am not sure it is good or bad, but it seems to help.

Weeding out the listserv messages simplified deletion of unwanted messages:

For all the spam mail that I get, it kind of appears in one place, and the titles sort of give it away anyway, but since it is all sort of there together I find it easier to chuck, chuck, chuck, chuck, chuck. <deleting messages>

The other thing is if I get behind and don't have time to read some of the like "top stories of the day"-stuff. It's nice that is all in one place and I can just say OK there is three days worth of this and I just kill two days and then read the most recent ones.

The Bifrost organization of the inbox improves the user's ability to see an overview of all the messages. Categories that are not of interest currently can be collapsed to reduce clutter and release screen real estate:

Once the categories stops fitting on a page, I am much less happy,

Collapsing uninteresting categories helps getting an overview. When the scroll bar is gone I immediately feel that I am in more control over my life.

One of the managers who used to schedule the first hour every day to read email, stopped doing so. Bifrost helped him to control his own time since he no longer needed a full hour to identify the most urgent messages.

Bifrost can also help people to not open messages:

I certainly hit insert a lot more to get them unread now. (without opening them, this will make the messages fall down to Not categorized at the next categorization)

Disadvantages

The main disadvantages of Bifrost reported by the users result from its current implementation. Bifrost is slow to perform categorization when the inbox is located on a server rather than a local machine. Delays are especially common for users accessing their mail over a slow moden. A second disadvantage concerns the user interface: users must initiate the categorization of the inbox manually. However, these two problems could easily be addressed by changes in the implementation.

The only user that did not like Bifrost had an unrelated problem with her unread marks. Occasionally her unread marks became corrupted in her base mail system, and she could not use the color red [3]as a means of indentifying new, unread messages. Bifrost made this problem worse by splitting the new messages into several "buckets".

All the other users found the disadvantages of Bifrost to be of minor importance and related to their own maintainance of the user defined fields in figure 3:

---

3 [3]In Notes(TM) email, messages that have not been read are indicated by the color red. Read messages are colored black.

Flaws, there is changes in what is important and who is important, but that is OK, because it (Bifrost) is still more, at a greater percentage, than not helpful, it is just something I have to remember. It is only as smart as I can keep it smart.

Need to change rules

Three of the Bifrost users did not add any names to the field. The remaining six users added between 6 and 47 names with a median of 16. None of the users changed the predefined rules that interpret the user defined fields.

I haven't changed them, Half out of laziness, but also half out of not needing to. I mean, they are at an OK granularity for me.

Folder usage effects

Two users claimed that Bifrost reduced their usage of folders to archive messages while one user started to use folders after the installation of Bifrost.

It moved me into folders.

Now I can use the Bifrost categories, and use folders only for long-term storage.

These contradictory behaviors can have a simple explanation: Bifrost brings order to the inbox. For some users, order gives them the organization they need to stop using folders, or reduce folder usage. For others who previously found folder usage too cumbersome, Bifrost provided an organization for the messages that are difficult to put in a folder. This automatic organization liberates users from the burden of archiving messages that are too complicated to put in folders; for example, messages that could go into several folders at the same time or require a new folder but for which the user has no name. This organization saves the user time. The difficult messages can remain in the inbox with the Bifrost categorization.

Differences between users

Is it possible to create a set of rules that are so general that they can be applied to all users? Perhaps not, but Bifrost comes close. The user studies indicate only one problem: blind carbon copies (as known as Bccs).

Sending blind carbon copies is a common trick among spam senders to make the message appear as an important message. For most users, these messages should be considered as less important. However, especially for managers who often receive real messages as Bccs, these messages may be more important. The current solution is to place Bccs in a sub-category of Personal messages. Bccs from VIP Platinum will end up in the VIP Platinum category.

Another difference observed among users was the number of days over which to categorize messages. Most users only used categorize "selected" or "all unread". However, one manager frequently categorized the last day and week when he was behind in his reading. This only influenced the number of menu alternatives, not the rules per se.

Discussion

Covey (1992) has made the distinction between *urgent* and *important* tasks and claims that it is essential to handle tasks in importance order, not urgency. Applied to the domain of email messages, a good prioritizing system should therefore be able to distinguish important from urgent messages. However, from the user studies described above, it is clear that many users have difficulties separating *important* from *interesting* messages. One of the users described how Bifrost helped him to focus on the important messages first without becoming distracted by the distribution list messages that caught his interest.

An ideal categorization system should be able to distinguish between urgent, important, and interesting messages. However, due to the fast changes in what is urgent and the dependence on off-line events such as random meetings it is questionable whether this will ever be possible.

New ways to support message handling

Bifrost provides support only for identifying which new messages to read. The next step, how to handle these messages, needs a different solution. Many of the studied users marked messages that were not completely handled as unread. The reason was that the unread appearance worked as a reminder to do something with the message. This is a rather poor support for a task with a rich set of varieties. In the real world we can sort incoming magazines, letters, papers, ads, etc. into different piles. Similar possibilities in the email interface might be an advantage. The problem so far has been the lack of differences between email messages - they all look the same. However, with Bifrost there is a way to create initial piles and after that the user can proceed with sorting message manually, if needed.

A pile could show a shortened version of one message and indicate the total number of messages in the pile. The difference towards a folder is that a pile does not have to be named and to prevent the user from forgetting the messages the top message in the pile could automatically shift (at a slow pace).

We will give two examples of such a design. First, a 2.5-D interface where messages behave like ordinary paper messages in that sense that they are opaque: a message put on top of another message will block the view of the lower message (see figure 5). The third message in the pile *Small distribution\Unclassified* is hovering as a reminder (messages not completely visible will do this at a slow pace). Numbers indicate the number of unread and total number of messages in the pile. When there are too many messages in a pile to display the sender, the bottommost messages will overlap (but still hover at a slow pace).

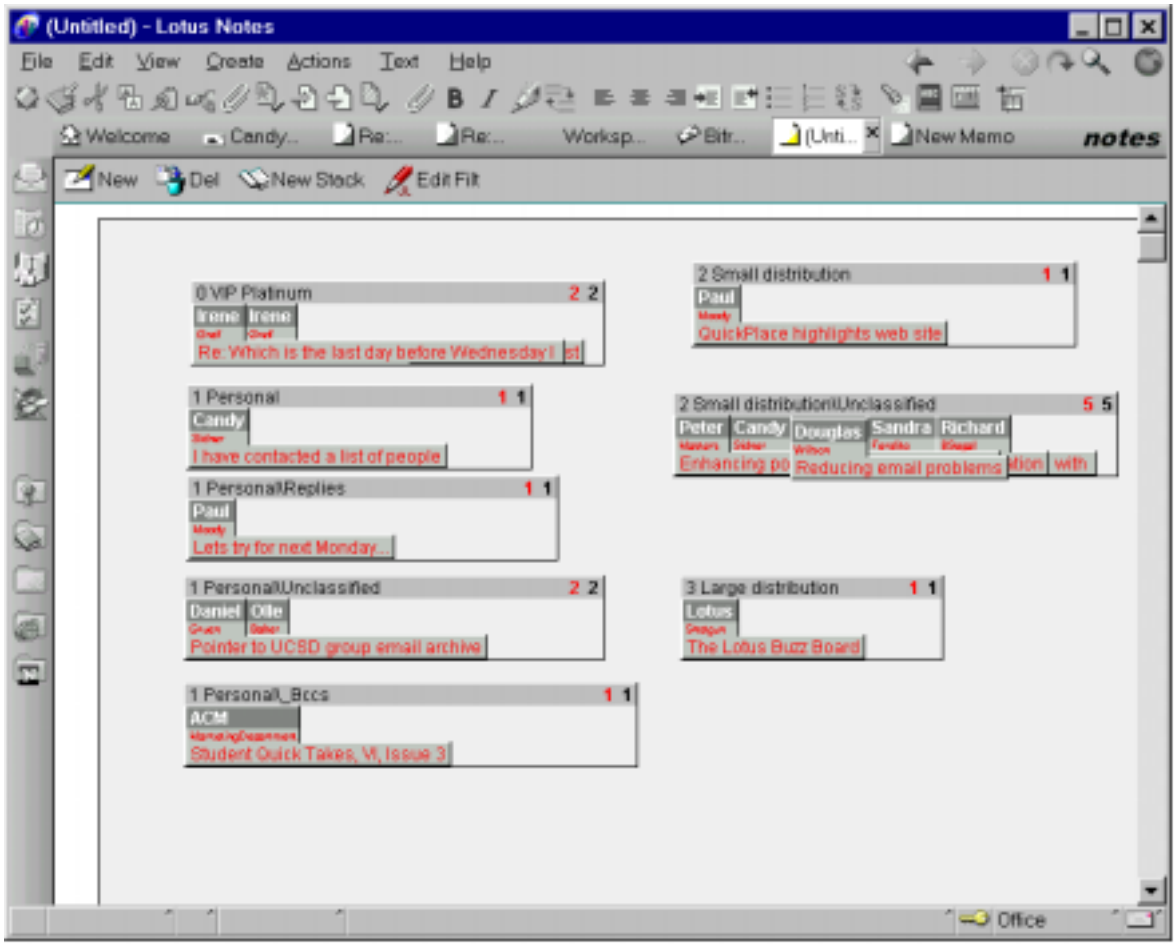

Figure 5. Bifrost categories represented as piles.

We also experimented with a second design which makes messages form piles in form of lists automatically (see figure 6). In figure 6 one message (from Sidner) is outside the piles. A user formed pile is about to be named “To do”. Font size is used as an indicator of time - smaller font represents an older message. If a pile is resized and no longer can display all messages, an indicator of the number of hidden messages will be displayedand the bottommost message will be replaced by one hidden message at a time in a slow pace to work as a reminder. A collapsed pile will only display the category name (if any) and the number of read and unread messages. For more ideas on piles, see Mander, Salomon & Wong (1992).

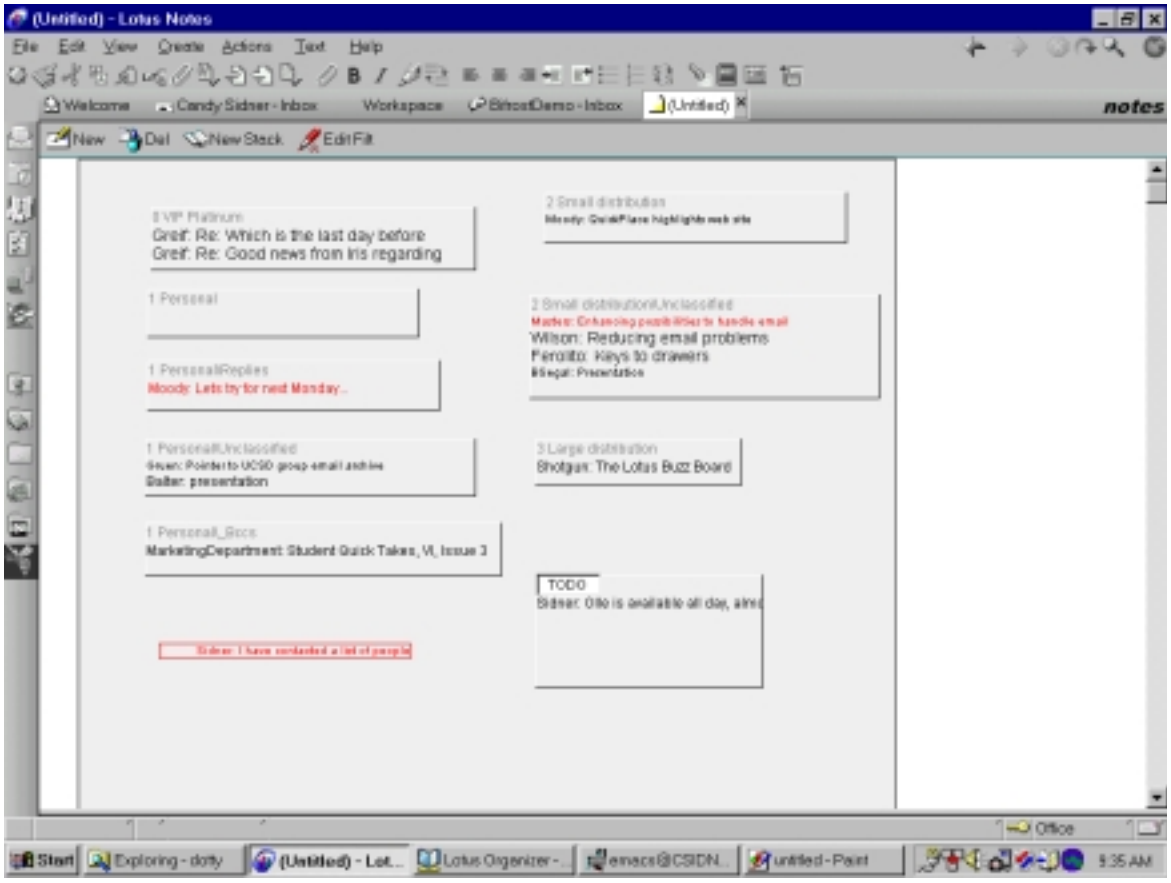

Figure 6. Bifrost categories represented as piles in form of lists.

This kind of visualization gives the user a bird's eye view of the messages and supports implicit task handling -- the user can create piles without naming them or deciding what the pile is for. The interface is natural, this is what we all do with messages in the real world. The problem with “out of sight- out of mind” is reduced by the slow sifting though the messages.

Conclusions

We have described the development of the Bifrost InboxOrganizer and its warm reception by users. Bifrost consists of a rule set that can be used to categorize email messages, without burdening users with creating or maintaining rules. The intention was to support users who received more than 30 messages per day to prioritize their reading.

Similar techniques could be applied to other messaging systems, such as hand held devices, cell phones, and voice mail. When it comes to phone messages, the notion of group and distribution messages may be inadequate, but the notion of VIP's, replies, and unclassified callers are vital to users' phone inboxes.

According to nine of the ten interviewed users, Bifrost was an excellent tool for prioritizing email reading, and it also provided assistance in maintaining a to-do list in the inbox and simplified folder usage.

Our conclusion is that it is extremely difficult for an automatic agent system to handle prioritization tasks well due to the rapid changes in the off-line world. Also, a user giving feedback to an agent how a message should be prioritized compared to the other messages would have to read all messages to know that, and the user might not be aware of how to distinguish between important and interesting messages.

Even though Bifrost seems to be a step in the right direction, there is no support in email systems today for implicit task handling. Our suggestion is to take advantage of a 2.5-D interface design in email systems as well. Thanks to the categorization made possible by Bifrost, this could give the user a fast and natural way to overview messages.

Email is one of the most widely used communication systems but the inbox has changed little since the beginning of the seventies. We believe our approach could make a difference.

Acknowledgments

This study was funded by The Swedish Transportation and Communications Research Board and Lotus Development Corporation. The underlying programming in Bifrost was initially made by Salvatore Mazzotta at Lotus Professional Services. A large part of the Lotus Research group has contributed in different ways: Carolyn Boettner, Li Carbera, Dan Gruen, Paul Moody, John Patterson, Steve Rohall, and Bob Stachel. Also, many thanks go to the anonymous users that participated in the studies.

References

Arensburger & Rosenfeld (1995): *To take arms against a sea of email.* Communications of the ACM 38(3) pp 108-109.

Bannon L. (1993): *CSCW: An Initial Exploration.* Scandinavian Journal of Information Systems, vol. 5, pp 3-24.

Boone G. (1998): *Concept features in Re:Agent, an Intelligent email agent.* Proceedings of the second Annual Conference on Autonomous Agents, pp. 141-148.

Borenstein N. & Thyberg C. (1991): *Power, ease of use and cooperative work in a practical multimedia message system* . International Journal of Man-Machine Studies, vol. 34, pp 229-259.

Bälter O. (1997): *Strategies for organising email messages*. In proceedings of HCI'97, Springer, London, United Kingdom, pp 21-38.

Bälter O. (1998): *Electronic Mail in a Working Context.* Doctoral thesis, TRITA-NA-9820, Royal Institute of Technology, Stockholm, Sweden.

Carlsson S. (1951): *Executive Behaviour: A study of the Work Load and the Working Methods of Managing Directors.* Strömbergs, Stockholm, Sweden.

Covey S. (1992): *The Seven Habits of Highly Effective People. Powerful Lessons in Personal Change*. Simon & Schuster, London, United Kingdom. ISBN-0-671-71117-2.

Hiltz S.R. & Turoff M. (1985): *Structuring computer mediated communication systems to avoid information overload* . In Communications of the ACM, vol. 28, no 7, pp 680-689.

Lansdale M. (1988): *The Psychology of personal information management*. Applied Ergonomics, vol. 19, no 1, pp 55-66.

Lantz A. (1996): *Computer Mediated Communication in a Work Context: an Interdisciplinary Approach*. Doctoral Thesis, Stockholm University, Department of Psychology, Stockholm, Sweden. ISBN 91-7153-519-5.

Losee R. (1989): *Minimising information overload: the ranking of electronic messages*. Journal of Information Science, vol. 15, pp 179-189.

Mackay W. (1988): *More Than Just a Communication System: Diversity in the use of Electronic Mail* . In Proceedings of CSCW'88, pp 215-218.

Maes P. (1993): *Learning Interface Agents.* ACM-SIGCHI International Workshop on Intelligent User Interfaces.

Maes P. (1997): *Agents that Reduce Work and Information Overload*. In Bradshaw, J. (Ed.) Software Agents, ISBN 0262522349. AAAI/MIT Press, pp 145-164.

Malone T. (1983): *How do people organise their desks? Implications for the design of office information systems.* ACM Transactions on Office Information Systems, vol. 1, no 1, pp 99-112.

Malone T., Grant K., Turbak F., Brobst S. & Cohen M. (1987): *Intelligent Information sharing systems.* Communications of the ACM, vol. 30, no 5, pp 390-402.

Mander R., Saloman G. & Wong Y.Y. (1992): *A 'Pile' Metaphor for Supporting Casual Organization of Information.* In Proceedings of CHI'92, pp 627-634.

Marx M: (1995). *Toward Effective Conversational Messaging*. Master Thesis, Media Art and Sciences MIT.

McCall M. W., Morrison A. M. & Hannan R. L. (1978): *Studies of Managerial Work: Results and Methods*, Technical Report no 9, Centre for Creative Leadership, Greensboro, North Carolina, USA.

Milewski A:E: & Lewis S.H. (1997): *Delegating to software agents*. International Journal of Human-Computer studies, vol. 46, pp 485-500.

Palme J. (1984): *You have 134 Unread Mail! Do you want to read them know?* Proceedings of IFIP Wg 6.5 Working Conference on Computer-Based Message Services.

Savendy G. (1987): *Handbook of Human Factors*, Wiley NY 1987, pp 368, 405.

Segal, R. And Kephart J.O. (1999) *MailCat: An Intelligent Assistant for Oganizing Email*, Proceedings of the third Annual Conference on Autonomous Agents, May, pp. 276-282.

Sheridan T. & Ferrell W. (1974): *Man-Machine Systems: Information, Control, and Decision Models of Human Performance.* MIT Press, Cambridge, Massachusetts, USA.

Sheth B. & Maes P. (1993): *Evolving Agents for Personalized Information Filtering* . In Proceedings of the Ninth IEEE Conference on Artificial Intelligence for Applications, pp 345-352.

Stewart R. (1967): *Managers and Their Jobs*, Macmillan, London, United Kingdom.

Whittaker S. & Sidner C. (1996): *Email overload: exploring personal information management of email*. In Proceedings of CHI'96, pp 276-283.

Yates J. (1989): *Control though Communication.* The Rise of System in American Management. John Hopkins University Press, Baltimore, Maryland, USA, ISBN 0-8018-4613-7.

Appendix 1. User data from the second round of interviews

Manager	No	Not anymore	No	No	No	No	Yes	No	Yes	Yes
Estimated number of new messages per day	30	20-30	15-20	30-40	30	200-300	20-50	3 here + 40-80 at her university	100-200	80
Sent messages	25	5	5-10	10-20	15	10-40	12-25	5-10 (trying to cut down)	100	20
Total number of stored messages	1,705	3,800	3261 (+35000 on CD)	17,119	1,635	39,882	7,402	699	2477 (+8000)	25,099
% thereof in inbox	64%	2%	54%	0.4%	38%	2%	47%	92%	2%	1%
# folders	22	30-40	10	240	13	82	11	4	14	130
Deletes messages	Yes	Yes	No	Few	Yes	Yes	Yes	No	Yes	Yes, those of clearly no use
Reads all messages	Yes	Yes	No	Yes	Yes	No. Maybe 10%	Yes	No	No	No
Found Bifrost useful	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes
# people reporting	1	0	1	0	0	0	4	0	50	70
Used email since	75	82	84	80	95	79-80	93	94	77	82
Used online calendar	No	Yes	No	No	Yes	Sometimes	Yes	No	Yes	Yes
User #	1	2	4	3	7	8	5	6	9	10
Added entries in filter	10	7	0	18	17	0	0	0	6	47
Folder effect	Less	Started	None	None	Less	None	None	Stopped	None	None

Appendix 2. Additional implementation ideas

Besides improving efficiency and automatically categorizing messages, users asked for help in adding people to the user defined fields. For example, users requested adding menus to messages to allow them to add the sender or subject to a certain field in the user defined fields (as in, for example, Hotmail).

Some users requested a way to separate Internet messages from internal ones, as these users communicated mostly internally and therefore considered Internet mail less important.

Messages that are part of the same conversation thread may end up in different categories in Bifrost. In some cases, this exactly the effect users want:

> *A reply that I got from my thesis adviser somehow got lost in my inbox, and I clicked to categorize the last two months, and saw that there were some stuff unread in replies, and that was important so it was good that I found it again, and I would have lost it if it was not for that (Bifrost).*

However, some subjects were surprised by this effect. An improved way to navigate between messages in a thread and an overall view of threads could reduce their confusion.

It is possible that the names in the personal address book as well as any person to whom the user has sent a reply should be considered as VIP Gold without entering them manually in the VIP Gold field.

If a message is falsely categorized, it should be possible for the user to drag it to the correct category, after which the system would request a field definition from the user to make it possible to categorize other similar messages the same way.

The algorithm for identifying replies should be refined to handle SMTP. Currently it can only identify Notes replies. Internet mail is recognized as replies by the Re: in the subject line (which will not work in communication with non-English speaking countries).

Bifrost could be introduced to users by opening a dialogue window when the user becomes overloaded, for example after a vacation when messages stack up. If the user would like to get help from Bifrost to categorize messages, a wizard could be used to fill in the fields in the filter.

For users that get more than, say 60 messages a day, there is a risk that Bifrost reduce interest in reading messages from new senders (as these would not turn up in VIP Gold or Platinum until the receiver has acknowledged their importance). Therefore, functionality to identify senders that have many unread messages, but are not in the rules might become necessary.