

Lessons from the ReMail Prototypes

Daniel Gruen, Steven L. Rohall, Suzanne Minassian
Bernard Kerr, Paul Moody, Bob Stachel, Martin Wattenberg, Eric Wilcox

IBM T.J. Watson Research Center
One Rogers Street
Cambridge, MA 02142 USA
+1 617 693 5786
{firstname_lastname}@us.ibm.com

ABSTRACT

Electronic mail has become the most widely-used application for business productivity and communication, yet many people are frustrated with their email. Though email usage has changed, our email clients largely have not. In this paper, we describe a prototype email client developed out of a multi-year iterative design process aimed at providing those who “live in their email” with an improved, integrated email experience. We highlight innovative features and describe the user trials for each version of the prototype with resulting modifications. Finally, we discuss how these studies have recast our understanding of the email “habitat” and user needs.

Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information search and retrieval – *clustering, information filtering, search process*;
H.4.3 [Information Systems Applications]: Communications applications – *electronic mail*; H.5.2 [Information Interfaces and Presentations]: User interfaces – *evaluation/methodology, graphical user interfaces, interaction styles, prototyping, screen design*; I.2.7 [Artificial Intelligence]: Natural language processing – *text analysis*

General Terms

Design, Experimentation, Human Factors.

Keywords

Electronic mail, information visualization, threads, text analysis, attention management, summarization, instant messaging.

1. INTRODUCTION

Electronic mail has become the primary business productivity application. It has emerged as the most-used communications tool in the US and Canada, and, according to an Institute for the Future study, 97% of workers report using email daily or at least several times each week [14]. In fact, US workers average 49 minutes a day managing email, and 25% spend more than one hour per day on that task [15].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CSCW'04, November 6-10, 2004, Chicago, Illinois, USA.
Copyright 2004 ACM 1-58113-810-5/04/0011...\$5.00.

Tools for handling email have failed to keep up with our growing reliance on electronic mail for a variety of uses in the workplace [2]. First, workers feel overwhelmed by their email. While the average user gets approximately 24 messages per day, “high-volume” users can easily get several hundred [4] [14]. (Ironically, 34% of internal business messages were deemed “unnecessary” [15].) Second, people use their email inboxes to manage their tasks, yet they complain that “things fall through the cracks” as messages scroll out of sight and are lost among newly arriving mail. Current organizational structures within email clients, such as folders, often prove inadequate, especially for high-volume email users [13]. Finally, people feel pressure to be more responsive in replying to email messages, reporting that 27% of messages sent “require” immediate attention [15].

Email clients have changed little since they were first invented. Most email clients today continue to use the same basic set of features and organizational structures: a textual listing of messages, the ability to preview and act on a selected message, and multiple folders in which messages can be filed. However, studies have shown that folder systems quickly degrade with the number of messages people receive [13]. Most people end up keeping all of their email in one large folder [1]. The content and use of email has also changed. In addition to traditional letters, email now consists of invitations, receipts, transactions, discussions, conversations, tasks, and newsletters, to name a few variations [2].

Our group has been investigating how people use email and how we might design and build a better email system. Our studies include user observations, interviews, design mockups, iterative prototype development, and user evaluations. These efforts included a significant prototyping effort in an attempt to “reinvent” the email client experience.

During the past three years, we have modified the prototype based on iterative user testing and new design explorations. Our prototype, known as ReMail, evolved into an integrated mail and messaging client framework that lets us explore new features in different combinations. Throughout this time, new research and products were released that supported or extended our research. This paper describes the evolution of our prototypes based on our studies and user tests. We also present overall lessons from the project and thoughts on where email should be heading in the future.

2. INITIAL PROTOTYPE

Our first working ReMail prototype (Figure 1) was constructed during an intensive 3-intern project during the summer of 2001 [9]. Our goal in building this prototype was to implement features

we had explored in storyboards and design prototypes in a system robust enough for people to use it on a daily basis with their own email. Email usage is quite personal and the best user feedback can be obtained by allowing people to use a system on their own messages in the context of their real work. In our case, that meant being able to handle mail databases with tens of thousands of messages and more than one gigabyte of data

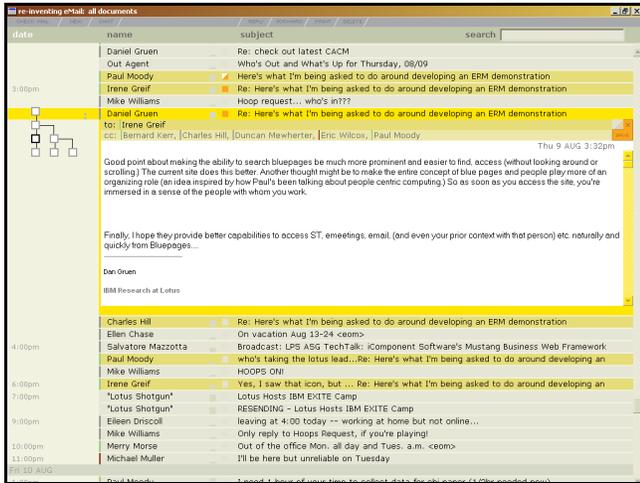


Figure 1. The Initial ReMail Prototype

2.1 Features

Key features introduced in this prototype were a thread map, gathered threads, thread highlighting, inline message previewing, date separators, instant messaging integration, and annotations.

2.1.1 Thread Map

Because people spend a great deal of time searching their inboxes for related messages, many features in this prototype dealt with message conversations, or threads. The thread map was an interactive visualization which represented messages with a reply-to relationship. The intended implementation, as described by Rohall [10], was a “mixed model” [12] which fully supported both sequence and reply relationships. However, the project’s short duration allowed for only the hierarchical thread map to be completed (Figure 2). This map displayed the reply structure of the conversation with each message represented by a square. Selected messages were shown with thick borders, and message headers could be displayed in a rollover pop-up. The map was rendered next to the selected message in the list view to help people see and navigate to other messages in the thread, cutting the amount of time they wasted finding information. Clicking on the node would scroll the message list to the relevant message.

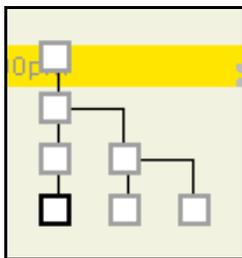


Figure 2. The Thread Map

2.1.2 Gathered Threads

The user could also choose to view individual threads in a “gathered” format, which displayed only the original message and most recent message in the list (Figure 3a). Showing only these two messages significantly reduced the number of items within the inbox view. The user could expand the thread to view all its messages (Figure 3b).

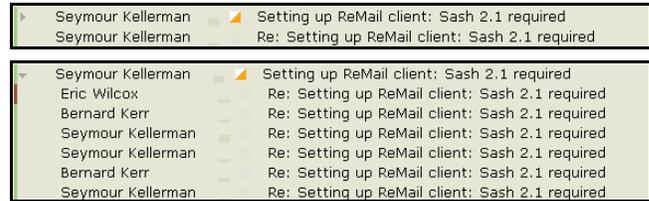


Figure 3a. Collapsed Gathered Threads (top)

Figure 3b. Expanded Gathered Threads (bottom)

2.1.3 Thread Highlighting

The interface provided users with a list of messages, similar to a standard inbox. When a message was selected, other messages in that thread were highlighted in a secondary color showing the user where other messages in that thread existed (see the dark yellow list items in Figure 1). This allowed users to quickly find related messages in the message list. Users could also perform actions on an entire thread, such as forwarding it.

2.1.4 Instant Messaging Integration

IBM Lotus Sametime, an instant messaging application, was integrated into this prototype to allow users to chat directly from the email interface rather than having to switch to a separate instant messaging application. Online status indicators appeared next to the author name in the message list. Therefore chatting was not restricted to members already on the individual’s buddy list; status indicators were displayed for any person in the “To”, “CC”, or “BCC” fields of the message, and a chat could be initiated with anyone online. Chats could also be saved, extending the current thread.

2.1.5 Inline Message Previewing

A message was previewed inline where it was selected, as displayed in Figure 1. It was possible to have more than one message preview open at a time, making it easier to work with multiple messages simultaneously.

2.1.6 Date Separators

The message list was divided into sections by the date the message was received or sent, with a bar that extended across the entire width as a heading. This provided the ability to scroll through volumes of mail more easily by day and see all the messages from the same day as a group. We removed the redundant date columns from individual entries, leaving more room for other information.

2.1.7 Message Annotations

An annotation was the digital equivalent of a Post-It® note. It allowed users to mark a message with a small colored icon of their choice that could also contain text, flagging a message for later reference. The clearly visible icon helped manage users’ attention and differentiated one piece of mail from another. The message

annotation widget was integrated with the close box, so users could easily annotate a message with a simple gesture.

2.2 Implementation

To access users' mail databases, we wrote a custom COM object in C++ using the C++ API to IBM Lotus Notes, the standard messaging client in our company. The user interface for the prototype was written in Dynamic HTML and JavaScript using IBM Sash, an extension of the JavaScript language that allows web applications to extend beyond the browser onto the Windows desktop.

2.3 User Trial

2.3.1 Procedure

We tested the prototype with 8 business email users within our company who used it against their own email for 2 to 5 weeks. Users included members of our extended group who had not worked closely on this project. Our goal was to collect general impressions of the new features and how useful they were, gather feedback to improve the user interface, and fine tune the algorithms used to thread messages.

After an initial interview and demo, we left the subjects to use the test client over a period of at least two weeks. We instructed them to use the client whenever they could comfortably, but to always feel free to use their existing Notes client and not let the testing impede work. We also gave the subjects a small logbook to record any thoughts or problems they had. Each page of this logbook contained a place to check off the type of comment (Usefulness; UI; Problem; Threading; or Other) and places to record the comment, what they were doing at the time, and the messages it involved.

At the end of the test period, we interviewed the subjects again, asking for their general impressions of the client, what they did and did not like, and their thoughts about each of the features. We also went through the logbooks with the subjects item by item.

2.3.2 Results

All subjects saw email as an important business tool in their work. Subjects were asked to describe the role email played in their daily activities, and the following responses were typical:

“Critical to my job”

“One of the most important communication tools... It's your work.”

“Critical role. I use it to communicate with just about everyone I know. In fact, people who don't use email, I've become more distant.”

In general, users liked the client and specifically mentioned threading, Sametime integration, and date separators as features they found most useful. Because of its limited functionality, none of the subjects used the ReMail client as their main email tool.

Threads were seen as a useful feature, both for finding information and for collecting related documents. Not all users found the thread map intuitive, however. In the prototype, the nodes in the thread map were spaced equally; one subject suggested it would be clearer if the visualization indicated the time the messages were received, an idea (though not implemented) which we had also discussed in other contexts

[8][10]. Another subject commented that the structure of the map did not give enough clues on which messages were important. People liked to find a single message in a thread and then gather the other messages in the thread together. One subject commented that this was particularly useful for locating an earlier message in a thread (e.g., the one with a link or an attachment) without hunting through the inbox.

As expected, all users encountered situations in which threads were not constructed correctly. For example, messages from the same sender years apart with the same subject (e.g. “Hello” or “Meeting”) were all grouped together by the algorithm we used. While errors often resulted because the system was “too aggressive”, one subject expressed a desire to be able to tell the system to be “more aggressive” for a specific thread. There was also a suggestion that improvements to the thread map visualization could compensate for imperfect threading. For example, adding a representation of time to the visualization would show that two sets of messages threaded together had, in fact, been received years apart.

Users liked the ability to gather threads, both to simplify the inbox view in general and to make it easier to deal with a specific thread. Two users (and one researcher on the project who used the prototype extensively for his own mail work) requested the ability to have threads gathered by default for all messages; this feature was subsequently added to this prototype and incorporated in later ones.

Users did not annotate messages often, partly because annotations were available only in ReMail and not in Notes. Several users complained that operating the annotation widget was not initially intuitive, but once they figured it out, they liked the ability to mark messages to return to later.

Subjects found it useful to have awareness information about who was online displayed near the messages. Several users commented that this was often more useful than a buddy list because the list of email senders showed the people with whom they were currently working, even if they had not been explicitly added to their buddy list. One person suggested that this helped them decide which messages to reply to first, prioritizing for people who were online and would be more likely to see the response quickly.

People generally liked the way mail was divided by day, and this feature was often mentioned unprompted as a nice aspect of the design. This was especially true for subjects who did not receive large amounts of mail because several day breaks would be visible on the screen at the same time and weekends stood out clearly. One user commented that this display was particularly useful after returning from vacation, since it gave him a clearer view of what had come in.

The inline preview feature was not well liked. Although the intent was to allow people to see the details of messages within the context of the entire message list, in fact, when more than one preview was open, little of the message list was visible.

We demonstrated this first prototype to executives and product groups within IBM as well as to select customers. Based on their positive feedback and interest from the product groups, our group decided to devote significant effort to a follow-on project for 2002 [6]. This led to the construction of our second prototype; although successful for the initial testing and proof of concept, the technologies used in our initial prototype did not provide the

scalability, robustness, or extensibility we needed to pursue the project further

3. SECOND PROTOTYPE

Our second prototyping effort began with an intensive design period leading to a comprehensive design specification document. (A detailed discussion of the design rationale can be found in [6]) The goal of the design was to create an integrated email experience reflecting feedback from our tests and providing a platform for further innovation. Figure 4 shows an overall screenshot of this design.

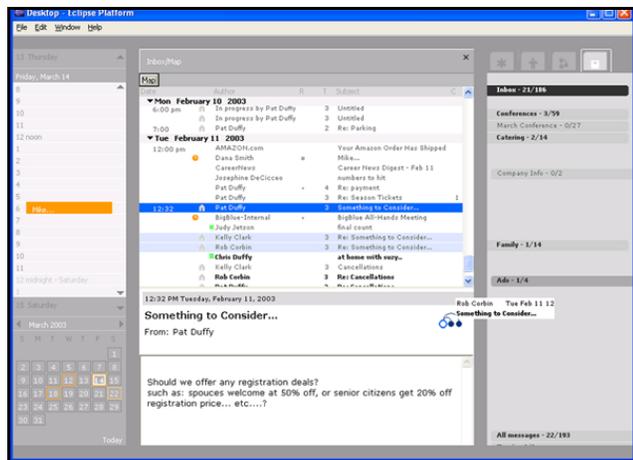


Figure 4. The Second ReMail Prototype

3.1 Features

A number of features were modified for this version, including the thread map and annotations. Online awareness and date separators were favored features that stayed in the revised prototype. Numerous features were added as well.

3.1.1 Thread Map

Feedback on the thread map confirmed that chronology was important to users because email represents events over time, as our research had predicted [10]. Implementing a chronological version of the thread map became a high priority. We designed a scheme, called thread arcs, to show message chronology along with the reply relationships of a threaded email conversation [5] (Figure 5).

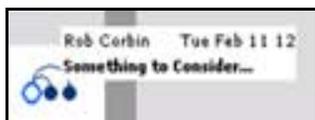


Figure 5. The Thread Map with Hover

This design maintained compactness with both deep and wide trees, while clearly indicating chronological order. As in the original prototype, hovering over a node provided message information and clicking on a node opened that message in the preview pane.

3.1.2 Correspondent Map

A new visualization, the correspondent map, was added to help users view messages from people they corresponded with most frequently. This map grouped messages within a collection by

sender, domain, and number of messages sent (Figure 6). The color of a rectangle indicated the age of the most recent message from that sender. Selecting the “unanswered” check box grayed out the rectangles of those to whom the user had sent mail more recently than they had received it.

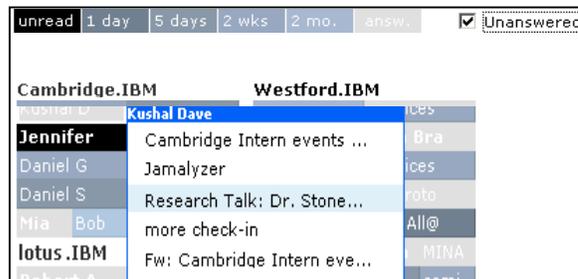


Figure 6. The Correspondent Map

3.1.3 “In-Sight” and “Out-of-Sight” Collections

The revised ReMail prototype incorporated an advanced concept of folders called “Collections”, which blended the features of standard folders with dynamic, rule-based views. Like traditional folders, the user could manually move messages into or out of a collection, and a message could belong to multiple collections. Also, as with traditional filtering rules, messages could be automatically sorted into collections using queries. These queries could be based on content, subject, author, and other information (e.g., “is ‘CSCW’ mentioned anywhere in the message?”).

An important capability of a collection was that it could be “in-sight” or “out-of-sight.” Marking a collection out-of-sight removed messages in that collection from the inbox view. Less important messages could be removed from the inbox so that the remaining messages commanded more of the user’s attention.

3.1.4 Multiple Sources

The revised prototype could access documents from additional sources and treat those documents like email messages. Initially, one source other than Notes was implemented for this prototype with the intention of adding a wider variety at a later date. This first additional source was IBM Lotus Quickplaces, a web-based collaborative team space. Documents from that source were pulled into ReMail as messages. This allowed users to forward Quickplace documents as regular mail, file them into an appropriate collection, or make calendar entries from them.

3.1.5 Calendar

Because calendaring plays an important role in the email experience, we added a calendar where messages could be used to create appointments. Having the calendar on the same screen as the inbox list allowed dragging and dropping of messages: dropping a message onto a day or time slot caused a dialog box to open in which additional event details could be entered. This capability was developed with the recognition that email often talks about meetings and appointments, but in a less formal way than a traditional “meeting invitation.”

3.2 Implementation

After exploring a number of technologies, we decided to build the prototype using Eclipse for both application development and deployment. This open-source platform provided a stable and scaleable solution for our development needs. The platform also

enabled users to easily add new features as we developed them, and position specific components on their workspace as they wished.

3.3 User Trial

Midway through the design and development of this prototype we conducted a formal usability test, running the prototype in its current state with a mail database crafted specifically for the tests. As we were not yet at a point where the prototype could be used reliably against people's own mail, this test focused on the usability and perceived value of the features. Using the same set of mail for each subject allowed us to guide the test to explore key usability areas.

3.3.1 Procedure

We ran sessions with five users who had varying email client experience. Tests focused on the overall layout of the ReMail prototype, threads and thread maps, collections, calendar actions, and multiple sources. Each test took between two and three hours. Tests were videotaped, and full screen captures were taken. The recordings were then transcribed, and the key observations merged and grouped by feature area.

Sessions began with a discussion of subjects' current email practices and issues. Subjects were then placed in situations that gave them an opportunity to discover features and visual elements on their own. If they did not find them, we drew their attention to the features through increasingly directed prompting. Ultimately, if needed, the feature was explained and demonstrated to them. This process gave us an opportunity to see if subjects would notice the features on their own, understand them unaided, understand them once explained, and then to get their impressions of the features. Subjects were asked about the perceived value of the features in their own work, situations in which they might use them, how the features should operate, and suggestions for improving the features.

3.3.2 Results

Overall, subjects found value in each of the features; however, some found the thread map and the way collections were presented and operated confusing at first. Subjects understood the overall layout and liked having the calendar open and accessible at all times. Interestingly, two subjects found our default layout of the calendar on the left and the collections list on the right confusing; they compared our collections to the folders in their current email program and expected the collections to appear on the left. Increased contrast and readability of some of the elements was suggested. ReMail included sent messages in the default message list. This enabled all the messages involved in a thread – including those sent by the user – to be viewed together. Not surprisingly, subjects thought the term "Inbox" was a misnomer for something which included the messages they sent. (Users could create a collection for sent messages and mark it "out of sight" if they preferred that such messages not appear in the

main inbox view; a few users would have liked such a collection to be created for them by default.)

Features retained from the original prototype tested well. All subjects found thread support useful. Subjects liked the ability to navigate easily among related items using the thread map; they also liked the way related items were highlighted in the list view. Three particularly found performing operations on an entire thread, such as deleting it or moving it to the trash, very valuable. There was some initial confusion regarding the mental model they formed of the thread map and what it represented; many subjects thought it had something to do with the flow of the message among individual recipients, with each node representing a different person to whom the initial message had been sent. It was clear to us that the hover-overs on the maps were not informative enough and that the thread map lacked visual cues to indicate that each node referred to a specific message.

All subjects tended to think of collections as folders once they understood their operation. They liked the ability to set automatic rules to bring items into collections and found value in setting collections in-sight/out-of-sight, especially when they were new. But there was confusion in the way in-sight/out-of-sight worked, particularly with items that appeared in multiple collections.

Subjects understood and navigated the calendar unaided, and they found the ability to drag messages to the calendar very valuable and easy to use. Suggestions for improvement included adding more of the expected event behaviors such as alarms, clearer indications on a calendar entry that there was a related message, and different icons for the different types of events.

Four subjects were (or had been in) situations in which having items from multiple sources appear in their inbox would be helpful, though one said it would be important to do so selectively. Subjects were asked if replies they generated to messages from outside sources should be sent from their personal mail or as replies from within the source. Subjects said that they would like to have replies appear as if coming from the same source as the original message by default, though two mentioned situations in which they would want the option to reply with a personal mail message. The house icon, an icon in the message list meant to represent a message's source, was not understood without explanation.

4. CURRENT PROTOTYPE

Feedback from this intermediate round of testing informed the continued development of the prototype. Many existing features were developed further, and new features were added, including date extraction, more choices for external sources, and several search capabilities. Our current prototype now includes a range of visualizations, advanced text analysis, and attention management features (Figure 7).

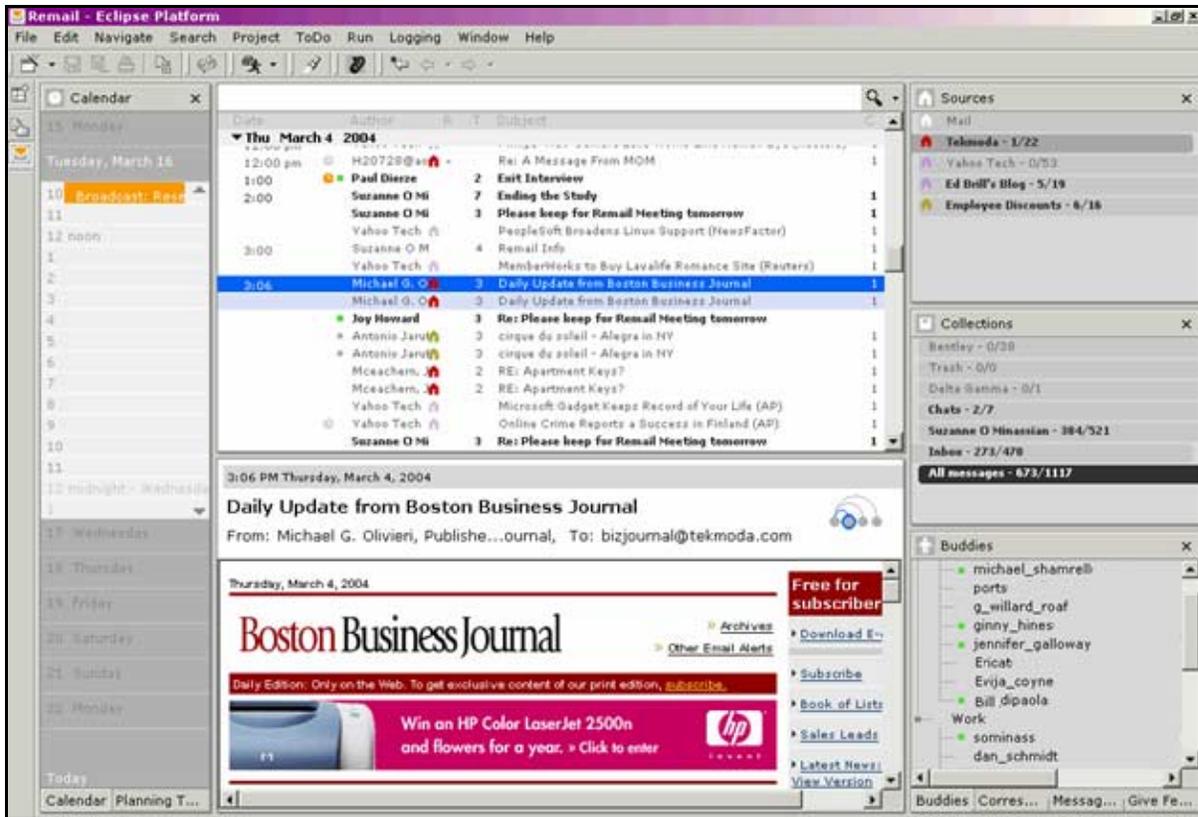


Figure 7. The Current ReMail Prototype

4.1 Features

4.1.1 Thread Map Study

We conducted a focused study to probe more deeply into the effectiveness of thread visualizations and tune their design for the final prototype [5]. In particular, we investigated a number of key qualities we considered important in thread visualizations.

One of the challenges in displaying these conversational threads is that they have two conflicting properties: the arrival sequence of messages and their reply relationship. The focused study confirmed that users find both of these qualities important. Showing the chronology as part of the visualization came through as a strong theme.

As part of the study, we let users explore the email threads in their own mail using a simulation built with Macromedia Director. Users were able to switch between three thread visualizations: thread arcs, tree diagrams, and tree tables during the test. Users were encouraged to use all three visualizations to assess which one they found most useful. Thread arcs proved to be a good technique for displaying the key qualities users considered important.

4.1.2 Improved Thread Summaries

The message list view displayed the number of messages in the thread for each message. Hovering over this thread count displayed a textual summary of the thread (Figure 8). The summary was created by extracting the first meaningful line of each message, and removing spaces and certain key salutations.

While we tried more complex algorithms with limited success [3], this very simple technique has proven remarkably good.

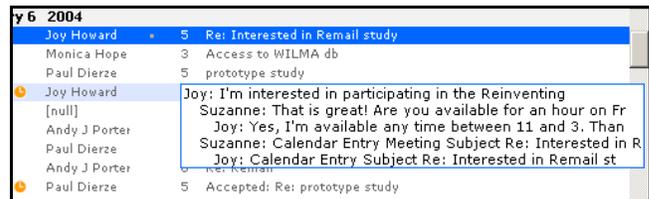


Figure 8. Thread Summaries

4.1.3 Gathered Threads

The option to gather all threads by default was implemented based on feedback from our earlier prototype. Users could select the “gather threads” option to display only the most current message in a thread. The entire thread could be opened in a separate pane, helping the users focus their attention on that thread.

4.1.4 Contextual Message List

We added pop-up message lists for collections, the correspondent map, and sources, letting users view and select messages on the side without having to switch the overall context provided by the main message list (Figure 9).

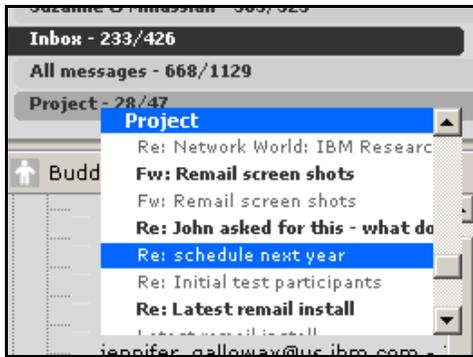


Figure 9. Contextual Message List

4.1.5 Date Extraction

Recognizing that many email messages discuss dates and times in an informal way (e.g., “let’s meet next Monday”), we implemented a system to identify date and time phrases that appear in the bodies of email messages. We parse these phrases to determine their meaning using the context of the message. We surface this information to the user, both automatically when looking at relevant dates and by allowing date-based search. Any dates or times found in the message body will become highlighted (Figure 10) and can be used to easily open a calendar item entry form with the correct information entered. An icon in the message list also indicates the date and time phrases found when hovering.

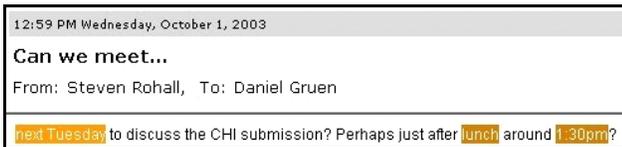


Figure 10. Highlighting Dates and Times

We conducted a series of focused studies to evaluate and improve the accuracy of our date and time extraction algorithms when run against typical email messages. We also wanted to see how users categorized the different dates and times that were found. This study is described in [11].

The date extraction feature was well-received when demoed to over 50 customers in a trade-show lab setting. Most visitors reported experiencing the problem of forgotten events and expressed strong desire for the feature.

4.1.6 Searching

Search capabilities were developed, including an instant search (a character-by-character filtering of the current message list based upon the information in the message’s headers), full-text searching, and date searching using date extraction.

4.1.7 Multiple Data Sources

The information people need to do their jobs comes from a variety of sources. Our current prototype reads Notes, POP3 and IMAP email, NNTP news groups, RSS feeds, and Lotus Quickplaces. Messages from these sources can be in any collection and accessed using the visualizations and tools described above. As before, these sources can be treated as email messages, allowing users to forward or create calendar items.

4.1.8 Message Map

The message map is a new visualization that allows the user to see relationships among messages in a folder quickly [7]. Messages are drawn as rectangles in chronological order (Figure 11). Visual attributes surface information of interest to the user, such as thread highlighting. In the figure, messages with an orange border have the same author as the selection and messages with a “dog ear” are unread (these settings are customizable). If a message is not selected by a user’s search, its rectangle is drawn as a lighter “ghost.”

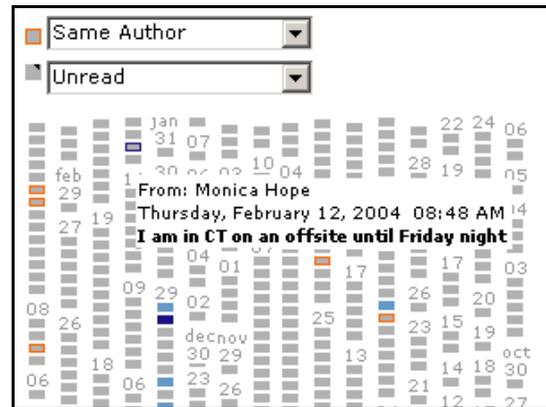


Figure 11. The Message Map Visualization

4.1.9 Improved Hovers

Some users were not sure what the icons in the message list were meant to represent. We added identification hovers to the source icons and the number-of-times-read icons in the message list. In addition, we made email hovers more descriptive. We changed all message hovers in the visualizations to include the first line of message text and more clearly label names as authors or recipients.

4.2 User Trial

4.2.1 Procedure

User studies were conducted with nine individuals who had a variety of jobs and technical backgrounds within our corporation. Subjects included two representatives from the same technical support team, and three from a public relations group. The four other users included an educational software manager, a project manager, a financial analyst, and a “technology advocate.” Breaking down the sample this way allowed us to view the results in terms of job functions as well as individual preferences. All subjects used Lotus Notes as their primary email client. Subjects had been using Lotus Notes for an average of nine years and received between 15 and 100 email messages per day (mean 39).

The studies were conducted individually and began with a one-hour interview in which subjects described their current email usage and their difficulties with email in general. During the interview, ReMail was downloaded onto the subjects’ machines and set up to access their mail databases. Next, users were given a brief overview of ReMail and its features, and an online reference guide for assistance. Subjects were asked to use ReMail for a minimum of one week but were allowed to keep Notes open to use when they wished. Subjects could contact members of the ReMail team for help at any time during the trial. Subjects used ReMail for 5 to 15 working days (mean 8.4).

At the end of the trial period, users were interviewed about their experience. The experimenter asked them to locate various types of messages and timed them as they looked for them. They filled out a survey which asked about their use of specific features in detail. They rated the usefulness of the features using a 7-point Likert scale, 7 being the highest, which were totaled and averaged. We also collected two log files that had been maintained – one recording user actions and the other with any comments the subjects created during the trial using a feedback mechanism we added to the interface for this test.

4.2.2 Results

4.2.2.1 Exploratory Interviews

When asked to identify the problems they face with email, subjects' responses ranged from issues of folder management to technical issues, such as server capacity. Interestingly, aspects of email tools that some users cited as problems were listed as positive features by others. For example, one user saw the integration of instant messaging into the latest release of Lotus Notes as a valuable addition, while another subject complained about the number of instant messaging interruptions he already received. One user found the amount of time required to set up rules on folders was not justified by their value, but another user considered rules important to her organization scheme. This finding highlights the great variability in how people use email, and validates our component approach, which allows users to select from among a set of integrated features and determine how those featured should be arrayed on screen.

4.2.2.2 The Individual Features

The results of the user study were positive. In general, users saw value in the features, and as expected, valued some features more than others. Four of the seven features ranked most useful (from a total list of 14 which included general email features) were new ReMail features. These included the quick search bar, thread map, collections, and additional sources (Table 1). Of the three visualizations, the thread map was rated highest for ease of use and helping people find messages. Subjects reported confusion using the message map and correspondent map visualization, though several stated that this was due to unfamiliarity. One stated that a tutorial would have been useful. Collections were a favored feature, as many subjects relied on folders for organization. Those who did not take the time to create rule-based collections on their own because of perceived complexity said they would have liked collections meaningful to their company or job automatically created for them.

Feature	Average	α	n
Sort-By Headers	6.13	0.99	8
Quick Search Bar	6.00	1.77	8
Thread Map	5.75	1.49	8
Collections	5.57	1.81	7
Message List	5.00	1.60	8
Sources	4.86	1.86	7
Preview Pane	4.33	2.24	9

Table 1. Subject Ratings of ReMail's Features on a 7 point scale (7=Very Useful, 1 = Not Useful)
(α = Standard Deviation, n = number respondents)

These features also proved to be among the most used by subjects. Our log files recorded 44 different user actions, such as selecting a message from the thread map, clicking on a message in the list,

editing a collection, or typing characters into the instant search bar. Viewing information in the three inbox hover areas were the most frequently performed actions. Together, hovering over the thread summary, the date extraction clock, and the number of times read made up 40% of the actions users performed with thread summaries leading at 26%.

The next most frequent actions performed were, respectively, viewing a collection, searching, navigating with the thread map, and selecting a message from one of the contextual message lists. Instant searching made up 79% of the searches subjects used, while date searching and full text searching accounted for 13% and 8% of the searches, respectively. Some comments from the closing interviews were as follows:

"I like the fact that [external source] messages are created which makes them easily searchable."

"Searching in ReMail is astounding."

"[On collections] Great feature. I like the fact that I could ... point documents meeting a search value ... in the collection. I also like being able to decide whether or not the documents in the collections should be 'in-sight.'"

4.2.2.3 Finding Information

Fifty-five percent of the survey respondents perceived ReMail to be faster for finding specific messages than their current email client, while the other 45% said it was no different. None found it slower. Eighty-eight percent of the respondents felt ReMail helped them find related messages faster; again, the others found it no different. More than half of the survey respondents agreed that the thread map helped them find messages faster, was easy to understand, and made their mailbox easier to use. Seventy-eight percent stated that the thread map made it easier to find related messages.

The preliminary timed tests we conducted did not show significant differences in locating messages within the subjects' mail box (though the raw numbers on average were better for ReMail). Because these differences do not correspond to subjects' reports of their own experiences, we hypothesize that this is due either to our test not reflecting the kind of searches that subjects naturally encountered in their daily use of ReMail, or that various aspects of the UI made the searching experience feel faster and more efficient to the subjects. This is something to probe further.

4.2.2.4 ReMail in the Workplace

We asked the respondents to rate the usefulness of ReMail in searching for information, organizing information, managing attention, and scheduling within their respective roles. Table 2 shows how respondents rated ReMail in these areas

Some of the features of ReMail were useful to particular groups of users. The RSS reader was deemed a necessity among the PR team, but seen as a less critical feature to the others. "I have to keep track of what is going on...this (RSS) source was invaluable to me," and "I used info updates from tech publications and business magazines daily," were two of the remarks made by subjects in this group. The team that dealt with internal technical troubleshooting said that saving chats was very valuable because a great deal of their work is accomplished through instant messaging. Being able to save a chat, forward it, and perform searching as a regular email were important features they wanted.

Function	Average	α	n
Searching	6.00	1.00	7
Organizing	5.29	1.38	7
Attention Management	4.60	1.34	6
Scheduling	3.50	1.64	6

Table 2. Subject Rating of ReMail in Particular Job Functions
(α = Standard Deviation, n = number respondents)

4.2.2.5 General Comments

Five users expressed interest in using ReMail beyond the trial period. One subject, who stored large volumes of customer data in his email archive, remarked that the ability to create collections combined with our search capabilities made ReMail an optimal environment to store and search his data. Four subjects said they enjoyed having a single interface for the RSS and NNTP sources and may continue using ReMail as a way to read and forward this information.

Users saw value in some of the features that they did not use frequently, like date extraction and dragging mail to the calendar. Subjects were asked why they did not use these features and pointed to ReMail not transferring data back to Lotus Notes as the cause. Likewise, some users complained of their Notes folders not transferring to the ReMail prototype.

4.2.2.6 Going Beyond Email

Interestingly, only 33% of the respondents pointed to the amount of email they received as a core problem. This was in contrast to our prior study in which subjects complained during interviews mostly about the volume of mail they received. Forty-four percent of the subjects in the current study reported that they lose track of important items and feel pressure to respond quickly to mail. For them, a bigger problem than sheer volume was finding ways to respond to mail more easily and the need to work with multiple messages and other items external to their mail to do so. For example, almost all subjects used databases, collaborative spaces, and the web for work on a daily basis. They also used tools specific to their roles, like spreadsheets, query tools, source code editors, and data repositories.

5. LESSONS

Several overall lessons emerged from our body of studies, interviews, prototypes and tests.

5.1 Connections and context

The sense of overload people feel report when working with their mail is more than simply a result of the sheer volume of messages. The scattered nature of the typical unaided inbox—in which items from multiple activities are interleaved based on the order in which they happened to arrive—provides a challenge for attention management, increases the likelihood that items will get lost, and makes it harder for people to respond quickly. The feedback we received consistently pointed to the value of connecting related items both through conversational threads and collections. Users liked the ability to work with these items together, both through features like the thread map and pop-up views, and by performing actions (move, delete, print, etc.) on an entire thread. The correspondent map similarly allowed users to quickly access messages grouped by person; the date extraction and calendaring

features allowed messages related to a specific date to be viewed together. For people who's work routinely included monitoring non-email feeds, the ability to include items from multiple sources in collections and threads was important. Users in our last study expressed the importance of connecting messages beyond email and information feeds to include the other items involved in an activity.

5.2 The importance of search

Not surprisingly, search was seen as very important and people particularly liked the quick search capability in our final prototype. The ability to easily search for dates within messages was also seen as a valuable feature, suggesting the value of similar text analytics to enhance search capabilities.

5.3 The role of visualizations

The thread map visualization provided a useful way to see structure and background information at a glance, allowing users to notice, for example, if they had not yet read all the messages in a thread or if a message currently being viewed had already been answered. It was widely seen as useful. On the other hand, our more complex visualizations such as the message and correspondent maps were seen to be interesting but not frequently used in practice. This suggests a value for simple visualizations that enrich the context of items already being seen by displaying non-critical, easily understood information.

5.4 Exploiting organizational structure

Structures exist within large enterprises and organizations that could be used to enhance the email experience. Knowledge of groups and organizational domains can be used to group people in meaningful ways. Enterprise-wide collections such as for "My Manager", "Action Required", "FYI", or "Expense Reports" could be predefined based on knowledge of internal systems and mail accounts. Our studies showed that while people often did not create rule-based collections themselves, they found them valuable once they were created. Users reported that they would have liked meaningful collections, such as those suggested above, automatically created for them.

5.5 The challenge of studying email

As other researchers have noted [2] studying email can be quite challenging. Because we wanted people to use the prototype with their actual mail for an extended period, we had to implement many features, such as printing and rich text support, that were not specific topics of research. We also included an "Open in Notes" button to open the current message in the user's existing email client. Despite this, most users were not able to use the prototype as much as we would have liked because of occasional missing features or performance limitations. For this reason we found it important to combine lessons from a number of sources, including traditional user tests, interviews, and feedback to our demos and storyboards.

6. CONCLUSION

The fundamental goal of this project was to reduce problems people face when managing their email, particularly keeping track of information, feeling pressure to respond quickly, and feeling overwhelmed by mail volume. We addressed these problems through mechanisms such as threads and collections to bring together related items, summaries and hover-overs to allow people

to quickly see information about messages and locate specific items, and quick-searches to make it easier to retrieve and locate messages. We added drag and drop and the use of automatic date extraction to make it easier for people to integrate messages involving events with their calendar and visualizations to give people various overviews and ways to navigate their mail.

As expected, there were individual differences in how valuable features were perceived to be, arising in part from differences in job needs and work settings. This validates the value of a componentized architecture, both for research and for deployment, to enable users to select the features that fit the way they work and configure them as desired.

Current social practices and business expectations have made the email client vital to work forcing users to live in their mail clients. As one subject reported, "From the moment I wake up until the moment I go to bed my email is on and alerting me." Our approach to email has been to make improvements to the email client to make it a better place in which to work.

While we have achieved much of that goal, it is also clear that we can only go so far with this approach. Improving the efficiency of working within an email tool is only one part of the problem. The biggest problem users dealt with in our most recent study was efficiently handling activities in which email played an important role, but often involved other items as well.

An effective system should support work that involves email rather than focusing on supporting email work. This entails structuring the email experience so it is organized by aspects of the work setting including business activities, tasks, and events, and it is integrated with relevant, non-email items in a cohesive experience. Our group is currently researching a unified approach to activities that would support this true reinvention of the email experience. People have been living in their email; the real goal is to find ways for people to live in their work.

7. ACKNOWLEDGMENTS

We would like to acknowledge the many people who have contributed to ReMail studies and the various versions of the prototypes. We thank Mia Stern for her development of the date extraction system, Kushal Dave for development expertise, Robert Armes for his work on the installation and logging, intern Jennifer Liu for her development of multiple sources, and intern Ka-Ping Yee for his work on thread summarization. We thank Steve Foley and Seymour Kellerman for their contributions to the system implementation, Andy Lafleur and Deb Maurer for help with testing, interns Kayvon Fatahalian and Evan Jones for work on the first prototype, and Maida Eisenberg for editorial assistance.

8. REFERENCES

- [1] Bälter, O. *Electronic Mail in a Working Context*. Doctoral dissertation, Royal Institute of Technology, Stockholm, Sweden, 1998.

- [2] Ducheneaut, N. and V. Bellotti. E-mail as Habitat. *Interactions*, 8(5), Sept.-Oct. 2001, ACM, pp. 30-38.
- [3] Lam, D., S.L. Rohall, C. Schmandt, and M. Stern. Exploiting Email Structure to Improve Summarization. Poster session (CSCW'02), New Orleans, LA, November 16-20, 2002.
- [4] Levitt, M. *Email Usage Forecast and Analysis, 2000-2005*. IDC Report # W23011, September 2000.
- [5] Kerr, B. THREAD ARCS: An Email Thread Visualization. In *Proceedings of IEEE InfoVis*, Seattle, WA, October 19-21, 2003.
- [6] Kerr, Bernard and E. Wilcox. Designing Remail: Reinventing the Email Client Through Innovation and Integration. *Design Case Study CHI'04*, Vienna, Austria, April 24-29, 2004.
- [7] Rohall, S.L., Interactive Poster: Visualizations in the ReMail Prototype. *IEEE InfoVis*, Seattle, WA, October 19-21, 2003.
- [8] Rohall, S.L. Redesigning Email for the 21st Century. Century Workshop Position Paper (CSCW'02), New Orleans, LA, November 16-20, 2002.
- [9] Rohall, S.L., and D. Gruen. Remail: A Reinvented Email Prototype. Demonstration (CSCW'02), New Orleans, LA, November 16-20, 2002.
- [10] Rohall, S.L., D. Gruen, P. Moody, and S. Kellerman. Email Visualizations to Aid Communications. In *Late Breaking Hot Topics – Proceedings of IEEE InfoVis*, San Diego, CA, October 22-23, 2001.
- [11] Stern, Mia K. Dates and Times in Email Messages. *Intelligent User Interfaces*, Madeira, Portugal, January 13-16, 2004.
- [12] Venolia, Gina and Neustaedter, Carman. Understanding Sequence and Reply Relationships within Email Conversations: A Mixed-Model Visualization. In *Proceedings of CHI'03*, Ft. Lauderdale, FL, April 5-10, 2003.
- [13] Whittaker, S. and C. Sidner. Email Overload: Exploring Personal Information Management of Email. In *Proceedings of CHI'96*, Vancouver, B.C., April 13-18, 1996.
- [14] *Pitney Bowes Study Reveals Increased Use of Electronic Communications Tools Among North American and European Workers*. Pitney Bowes press release, August 7, 2000.
- [15] *The Spam Within: Gartner Says One-Third of Business Email is 'Occupational Spam.'* Gartner, Inc. press release, April 19, 2001.