



## Visualizing Email Conversations and Related Web Resources

Craig Sayers, Malena Mesarina, Jhilmil Jain and John Recker  
HP Laboratories  
HPL-2008-138

### **Keyword(s):**

conversation visualization, information visualization, email visualization, conversational thumbnail, email, web browser, related links, related web pages, relevance

### **Abstract:**

In this note we explore the design of small visual images to represent email conversations about a particular web resource. A working system has been integrated within an Internet browser. While users navigate the web, it searches and displays relevant conversations from their personal email store, providing a visual summary via a thumbnail which depicts conversational participants and message flow along with related web pages and their relevance. Preliminary user studies show our visualizations of participants and related resources were more easily understood than the messaging display. Small image thumbnails have frequently proven to be convenient surrogates for larger pictures and we hope small conversational thumbnails may prove similarly beneficial for email conversations and related web resources.

External Posting Date: October 21, 2008 [Fulltext] - Approved for External Publication

Internal Posting Date: October 21, 2008 [Fulltext]

Submitted to CHI 2009

© Copyright 2008 Hewlett-Packard Development Company, L.P.



# Visualizing Email Conversations and Related Web Resources

Craig Sayers, Malena Mesarina, Jhilmil Jain and John Recker  
{first.last@hp.com}

## ABSTRACT

In this note we explore the design of small visual images to represent email conversations about a particular web resource. A working system has been integrated within an Internet browser. While users navigate the web, it searches and displays relevant conversations from their personal email store, providing a visual summary via a thumbnail which depicts conversational participants and message flow along with related web pages and their relevance. Preliminary user studies show our visualizations of participants and related resources were more easily understood than the messaging display. Small image thumbnails have frequently proven to be convenient surrogates for larger pictures and we hope small conversational thumbnails may prove similarly beneficial for email conversations and related web resources.

## Author Keywords

Conversation visualization, information visualization, email visualization, conversational thumbnail, email, web browser, related links, related web pages, relevance.

## ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

## INTRODUCTION

We have been examining better integration between email communication and the World Wide Web. An experimental web browser plugin has been developed which links messaging and the Internet by automatically finding and displaying email conversations which are relevant to the current web page. It works by searching for the web page URL in the user's personal email store and then sorting and displaying matching messages. A screen shot of the user interface is shown in Figure 1 and more details are available [1].

While showing summary text of every relevant message is adequate for short conversations it does not scale well. This is particularly true for our application since the email conversations are peripheral to the browsing task and so are given only limited display area. Accordingly, we have been exploring the design of small visual images to represent relevant email conversations. An example of these conversational thumbnails is shown in Figure 2. In this note we describe the design and evaluation of these thumbnails.

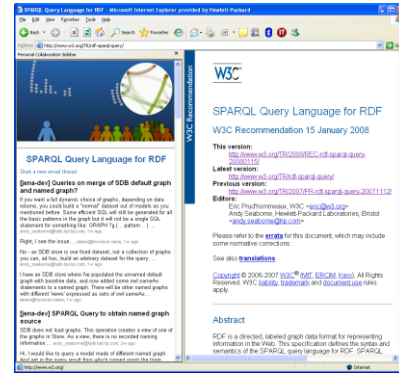


Figure 1 – A user has browsed to a web page and the sidebar on the left displays conversations from the user's personal email store which are relevant to that page. Our conversational thumbnail in the upper left corner is a visual summary of those messages.



Figure 2 - A closeup of the conversational thumbnail. In the upper right corner we show the communication paths, along the bottom are all the participants, and the spheres orbiting the conversation are pointers to related web pages.

## CONVERSATION VISUALIZATION

A number of researchers have looked at visualizing communication. Text-based visualizations are a natural fit for textual content and a particularly nice example is Viegas, Golder and Donnath's system for visualizing keywords from email content [5].

For showing conversational thumbnails a common approach is to focus on the threaded structure. For example, Wattenberg and Millen's visualization [6] represents each message as a rectangle, indents them based on the thread structure, and colors based on perceived value (using votes from readers). This is similar to the thumbnail visualization used in Dave et al's FlashForum reader [2], and the Message Map used in Rohall's visualization for the ReMail system [4]. Rohall's visualization system for email also included a thread map (showing links between threaded messages).

More abstract visualizations are also possible. In the Babble system by Erickson et. al. [3] each person talking about a topic is represented by a colored dot and the dots move toward or away from the center based on how recently they have been involved in the conversation. Xiong and Donath even made a system to visualize newsgroup postings as a field of flowers [7]. Each flower was one subscriber, the height of each stem represented the time that person had been subscribed, and the size of the flower indicated the number of postings.

### OUR APPROACH

We are exploring the visualization of conversations which are relevant to a visited web resource. A reliance on URLs contained within email messages conveniently avoided the need for any manual topic selection or natural language processing. Since there can be many messages about a single topic, a summary of the people and message flows was desirable, and since conversational topics are web resources, it was particularly useful to display not only all conversations related to that page, but also links to other pages and their relevance both to the current topic and to the current user.

The experimental conversational thumbnail merges three different information facets: people, messages, and related web resources. Each of those will be described in turn:

#### Representing people

We represent each person involved in the conversation (specifically each person who was either the sender or another recipient of an email message in the user's inbox which contains a link synonymous with that of the visited web resource) using an iconic representation (see Figure 3).

Showing only the upper body conserved space, and also gave some gender neutrality, while adding drop shadows and alternating the front/back ordering of adjacent people enhanced definition. The current user is shown in blue, while the color of every other participant is chosen consistently based on an alphabetical ordering of their email address. Each person's position is a function of how relevant they are to the current user. For example, someone with whom he or she just communicated is much closer than someone who was CC'd on an email from a third party six months ago. Users can hover over any person to see details (for example their email address) and in the future will be able to click on any person to find additional information.




Figure 3 – Representing conversational participants.

The layout algorithm goes to some lengths to make sure people are visible and appropriately-positioned. In the current implementation it is like virtual seats at a stadium. We place the current user a few seats in from the left-hand edge (chosen to gain a more balanced visual composition). The seats closest to him/her are the most expensive. The money each person can spend on a ticket is a function of the number of relevant conversations and whether they were directly to or from the current user. People are seated in the most expensive seat they can afford starting at the front row and working backwards. Seating is prioritized based on the time of the email which each person most-recently created or received. So, someone with whom you communicated a few minutes ago is almost guaranteed a seat, while someone from a conversation last month might miss out if all the seats are already taken. The size of each person decreases subtly as they move further from the target user. This reinforces the distance and conveniently provides more space for visualizing messages.

#### Representing messages

In the large white circle at the top right of the thumbnail (see Figure 2) we show each direct message from each conversation about the web page. As shown in Figure 4, the current user is placed at the center and spots around the edge represent other email senders or recipients. The location of each person is constant across multiple conversations and determined based on an alphabetical ordering of their email address. Each curved line represents a single email message from sender to recipient and is colored to match the color of the sender. To provide a temporal representation the lines fade away with time, so a thin sharp line is recent, whereas a faded wide line is older. This is intuitive (think of ink on paper fading in the sun) and conveniently allows multiple messages between the same people to be visible – the older message is wider and in the back, while the newer message is thinner and to the front. By using Bezier curves rather than straight lines, we're able to avoid overlap between sent and received messages along the same path. Since people are shown in consistent places (and with consistent colors) it is possible to recognize common conversational patterns and over time we have started to notice recurring patterns. We intend that with a quick glance experienced users should immediately know if it's a conversation with their employees, their work peers, or any other commonly-recurring group.




Figure 4 - Representing messages between participants.

## Representing related resources

When browsing the web, one technique is to navigate using links coded in the visited web page. Those links were considered appropriate by the page author – they are likely accurate, but not at all personal. Another source of navigation comes from your email – if someone sends you a single email containing a link to two different web resources, then we can infer there is some relationship between them and, perhaps more importantly, we can infer that's relevant to you. Thus, in addition to showing messages containing a link to the current web resource we also search those messages for links to other web resources. In displaying those links we aim to show how relevant it is to the current web page, how relevant it is to me, and how important it is (see Figure 6).

We visually represent links as planet-like objects in the “sky” orbiting the current conversation. Users can hover over a planet to see details of the site, or click on a planet to navigate to that site. The size and vertical location of the planets depends on their relevance to the current user – was he or she the sender or recipient, how frequently was the link mentioned in email, and how long ago did the communication happen.

Web pages are closer to the conversation if their URL is closer to the current web page URL. For example, if the visited page is from hp.com, then other pages from hp.com will be closer than pages from say amazon.com. To accomplish this we compute a similarity metric by comparing different portions of the web address.

Referring back to Figure 2, a user has visited the W3C SPARQL standards page and the conversations are from W3C mailing lists relevant to that standard. The planets in the closest orbit are from other generic W3C standards documents, those in the next-closest orbit are specific versions of the W3C standards, while planets in the more distant orbits are other non-W3C web pages that were mentioned (for example blogs and personal pages mentioned in email signatures).




Figure 6 - Representing related resources.

## EXAMPLES

In this section we present some examples of the thumbnail visualizations.

A common email pattern, where one user sends a link out to a number of recipients, is shown in Figure 5a. In this case the current user was one of many to receive this message,




Figure 5 - example thumbnails

the sender is the person rendered closest to the current user, all the other recipients are on the right, and the message was recently sent because the lines are thin and dark.

Another pattern is shown in Figure 5b, where the current user had a discussion with three other people and one participant then forwarded to a much larger group. Notice here how more people are involved in the conversations, that it's older (because the lines have faded) and that there are also two relevant web sites.

Figure 5c shows communication with a project group. Here the pattern of communication is very different – now the current user is much more in the middle of the conversation and over time may come to recognize that particular star-shape as being communication with his or her closest team members.

## TESTING ASSUMPTIONS

Our interface development is still in progress however we can present some preliminary results from a user study to test our design assumptions.

Participants in the user study were five employees at our company, each selected from a different job function: research manager, scientist, writer, administrative assistant and research intern. Two were female, three male, and they ranged in age from 30 to 45.

We installed the visualization as a browser plugin within Internet Explorer which would show related email messages in a sidebar as the user browsed the web. We did not explain the visualization, but instead asked them to explore it on their own. We interviewed them after 20 minutes and then again after one week and asked them to explain the various facets of the visualization.

### *Representing people*

All participants correctly surmised that the people represented those who had been involved in email conversations. However the subtlety of the seating positions was less-obvious and several users assumed the size of each person was based on the frequency of messages and that therefore those on the upper rows were less important. Having people on the left or right of the current user also caused some confusion as users attempted to ascribe meaning.

### *Representing messages*

Everyone figured out that they were at the center of the circle, but the abstraction of messages as lines was not apparent to either the writer or the admin, one thought that perhaps the circle represented a clock. Interestingly in both cases they had tried hovering over and clicking on the message lines and people spots and so if we had supported an appropriate action in each case then the meaning may have been more readily apparent.

### *Representing related resources*

All users noted that the planets represented links to web pages and by the second interview all had decided that the size of the planets depended on the number of times the URL was mentioned in an email. (The latter is not exactly correct since we had actually computed the size based on a weighted function so that, for example, links I send carry more weight than those I receive.)

No users correctly interpreted the planet positions. Even though several noticed that pages closer to the current page were more to the right and that similar web pages appeared close together, they were then confused when dissimilar pages also appeared close together on the outer orbits.

## **DISCUSSION**

Displaying people and related links were both immediately appreciated by all our test users. Interestingly, while we'd computed multi-dimensional weighted functions to determine position and size, our test users tended to assume a much simpler and more democratic relationship: *size of object is proportional to the number of messages*.

The visualization has room for improvement, especially in the messaging display which caused the most difficulty for our users. Improved interactivity should help.

In addition to formal study participants we have had a number of internal beta testers for the browser sidebar. During development we've moved from pure text, to text with a static icon, and now to text with the conversational thumbnail. Anecdotal comments from testers have become noticeably more positive with each step. For this application the aesthetics and novelty of the display may be even more important than the functional aspects.

In future work we plan to explore alternative size and positioning algorithms and alternative messaging displays along with more quantitative efficacy studies (using

instrumented software). We also wish the thumbnail to serve more as a navigational aid into the users email store (for example clicking on any person should show their communication related to that topic). And lastly, we are very interested in objective measures for qualitative aspects as we seek to build interfaces which are not only functional but that also delight our users.

## **CONCLUSIONS**

The graphical conversation visualization shows messaging patterns, conversation participants, and other related web pages. In addition it shows how relevant people/pages are and how long ago conversations occurred.

Preliminary user studies show the display of people and related links was more easily understood than the messaging display. In addition we note that users ascribed simple one-dimensional functions even though we had in fact implemented something much more complex.

The graphical visualization has a strong aesthetic component and introduces novelty and interest to the display. While these are difficult to quantify, we note they do provide perceived user benefit for at least some users, and their value should not be underestimated.

## **ACKNOWLEDGEMENTS**

Special thanks to Tyler Close whose inspiration and work helped create Sidebar. Thanks also to Vlad Bolshakov for his programming work, to Scott Golder, Erik Geelhoed, Badrinath Ramamurthy, and Venugopal Srinivasmurthy for their earlier advice and assistance, and to all those developers who built the open source libraries on which our solution depends. Thanks also to our colleagues at HP for their advice and encouragement, especially our test users.

## **REFERENCES**

1. Close, T., Recker, J., Sayers, C., and Badrinath, R. Sidebar: Ad-hoc, yet organized Personal Collaboration. HP Labs Technical Report HPL-2008-17, 2008
2. Dave, K., Wattenberg M., and Muller, M., Flash Forums and ForumReader: Navigating a New Kind of Large-Scale Online Discussion, In *Proc. CSSW 2004*, ACM Press (2004) 232-241.
3. Erickson, T., et.al. Socially Translucent Systems: Social Proxies, Persistent Conversation, and the Design of "Babble", In *Proc. CHI 1999*, ACM Press(1999), 72-79.
4. Rohall, S.L., Interactive Poster: Visualizations in the ReMail Prototype, In *Proc. IEEE InfoVis 2003*, 104-105.
5. Viegas, F.B., Golder, S. and Donath, J., Visualizing Email Content: Portraying Relationships from Conversational Histories, In *Proc. CHI 2006*, ACM Press(2006), 979-988
6. Wattenberg, M.M., and Millen, D.R., Conversation Thumbnails for Large-scale Discussions, In *Proc. CHI 2003*, ACM Press(2003), 742-743.
7. Xiong, R., and Donath, J., PeopleGarden: Creating Data Portraits for Users, In *Proc. UIUST 1999*, 37-44.