

Who Have I Been Talking To?

Tiago Garcia
Instituto Superior Técnico
Av. Rovisco Pais, 49
1050-001 Lisboa, Portugal
tiago.garcia@ist.utl.pt

João Aires
Instituto Superior Técnico
Av. Rovisco Pais, 49
1050-001 Lisboa, Portugal
joao.aires@ist.utl.pt

Daniel Gonçalves
INESC-ID
Rua Alves Redol, 9
1000-029 Lisboa, Portugal
daniel.goncalves@inesc-id.pt

ABSTRACT

One of the most frequently used ways to communicate with others is email. We have created "Who Have I Been Talking To?", a visualization of the users' email that can provide insights about their relationships with others: how one relation has evolved over time (reflected by the number of messages exchanged); the nature of the relationship (comparing the subjects of messages and how asymmetric is the ratio between sent and received email); and allowing users to compare their relationships with different people. We found that a variant of a stacked bar chart over an interactive timeline allowed all of those goals to be achieved. User studies showed that the users were able to understand the visualization and gain meaningful insights about their electronic communications with others.

Author Keywords

Visualization, Relationships, Personal Information Management, Email

ACM Classification Keywords

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

General Terms

Design, Experimentation.

1. INTRODUCTION

Email remains the most used way of communicating with others. According to numbers released by Royal Pidgin for 2010, excluding spam, there were over 32 billion email messages sent every day by an estimated 1.88 billion users. It can be argued that email is the primary form in which many of our relationships are maintained. There is, however, little support to help us understand those relationships as a whole. By design, email is composed of individual messages, with little relation with each other. At best, they are organized as threads or conversations, but that still doesn't give us any way to understand more general patterns and relevant features about our relationships, how they have evolved over time, and the people we relate with.

Works such as ReMail [3] give us a way to see how a particular thread or conversation has evolved over time. However, they lack

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, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AVI '12, May 21-25, 2012, Capri Island, Italy

Copyright © 2012 ACM 978-1-4503-1287-5/12/05... \$10.00

any means for a more complete view of the evolving relationships to be inspected. Digital Artifacts Remembering Storytelling [4], Visualizing Email Content [5] and Contrasting Portraits of Email Practices [2] take a broader view. However, while they give us overall information about whom have we been exchanging emails with, little context is given, making it difficult for the users to answer the "why" of such communications: we don't have an overall view of who we are contacting over time. The main problem to be overcome when dealing with email is the sheer amount of information that we must deal with. Solutions must scale in order for the system to be usable for both the user that receives 10 emails per week and the one that receives 10 emails per hour.

We argue that by using an appropriate information visualization technique it should be possible to display a user's email in such a way as to make implicit, contextual, information apparent. We developed WHIBTTo ("Who Have I Been Talking To?"), a visualization that displays a visual summary of how the users' email communications with others have evolved over time. In particular, it allows users to: understand how their communication with a particular person has evolved over time; understand the nature of a relationship (Is it more a unidirectional or a bidirectional thing? What subjects were addressed?); compare relationships with different people (does someone communicate with them more than others? Is the number of messages sent to different people correlated?)

Time was without a question a very important dimension to this problem, so the visualization is timeline-based. To cope with large numbers of messages and still have a meaningful visualization, we used a variant of stacked bars, where bars are not actually stacked but rather *superimposed*, thus accommodating larger numbers of bars, and being more amenable to large differences in bar size in the same visualization. A user study showed that our visualization is easy to understand and use. The users were able to gain meaningful insights from their email collections and were very pleased with the visualization.

2. Visualizing "Who HAVE I been talking to"

Our visualization requires data from email messages to be indexed in order to then be visualized. We used Gmail as the source since it is very popular (making it easy to find test subjects) and because it provides access contact lists that can be cross-referenced email senders and recipients, presenting them in more meaningful ways to users.

Figure 1 shows the overall look of this application. There are three main visualization areas. At the top is the "Main Timeline", the most important part of the visualization. Below, on the left, we find a "Secondary Timeline", and the "Topics Cloud" on the right.

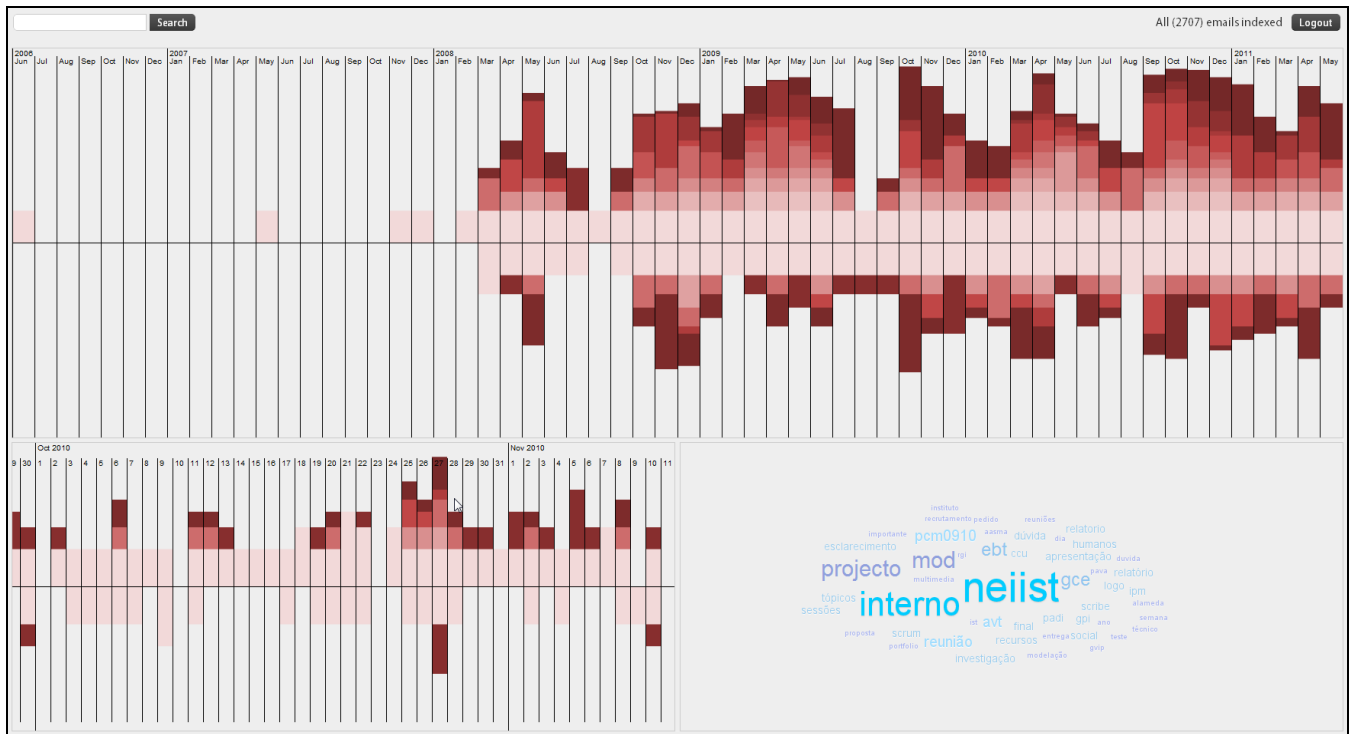


Figure 1. Overall View of WHIBTTo

The Main Timeline provides information about e-mail relationships on a monthly basis. It uses superimposed bars to encode different email facets. A horizontal rule divides the visualization in two. Bars on top represent received e-mails, whilst the bars below represent sent emails. The Secondary Timeline depicts daily information. We decided to keep the two timelines separate instead of letting the user zoom in and zoom out on the same visualization because we considered it to be important to keep the monthly context visible as a whole while watching the daily context and vice versa. They are coupled, and selecting someone in one of them also selects that contact in the other.

The topic cloud is the bottom-right partition of the visualization area. It shows the most discussed/distinctive topics. The keywords are chosen using the tfidf algorithm. When no contact is selected, this cloud shows the most discussed topics from every message. However, when a contact is selected the cloud shows the topics most used when in communication with that person. When a word is selected, every bar that represents a time and person where the topic was mentioned is highlighted providing immediate answers to questions such as "Who mentioned what, and when?"

2.1.1 Superimposed Bars

Each individual bar stands for a different month, appropriately labeled at the top. The entire visualization can, thus, be read as a timeline of the user's email exchanges. Surprisingly, though, while taller bars represent more emails, the height of a bar is *not* directly proportional to the absolute number of emails sent/received in that month. Many tools allow users to find that information already. What they seldom do is provide insights as to *how concentrated around a particular subset of people is our communication*. We achieved that goal by sacrificing the visualization of absolute

numbers, replacing it instead with the relative numbers of emails for different people.

Imagine a user talks with three different persons. Depicting each by its own bar (Figure 2, left), with height proportional to the number of messages, would be prohibitively expensive in terms of screen real estate, and impair our ability to gain an overall picture of our communications. So, we represent the data for each time period as a single bar, by combining the bars for the different people involved. We sort them by height and superimpose them, with shorter bars in front (Figure 2, right). The traditional solution, stacked bars would significantly increase the visualization's height, or require the bars to be so short they would become hard to discern and compare. Also, superimposed bars make comparing different people easier. For instance, it is simpler to realize that we communicated twice as much with Person A than with Person B by seeing a bar for Person B and, extending upwards from it, the bar for Person A. A stacked bar solution would require users to compare absolute heights of bars placed on top of each other, which is not an easy task.




Figure 2. The meaning of WHIBTTo's bars.



Figure 3. Use of Luminance.

The use of luminance allows to easily glance information other than the number of messages. In Figure 3, we can see that in the first bar there is a continuous gradient in luminance, meaning that the amount of emails sent to different people varies gradually between the person that got less emails and the one that received the most. On the second bar, there is clearly someone that dominated the users' emails for that month, as there is a visible gap in luminance. This allows us to identify asymmetries in the communication, and people with disproportionate importance in our communication.

Each sub-bar can stand for a set of people, to whom we have sent/received the same number of emails. This minimizes the vertical space required for the visualization. While the number of people we've communicated with becomes unclear, our goals were related to with the relative amount of emails exchanged than. As such, this was an acceptable trade-off. Another drawback of not stacking the bars is losing the sense of global communication magnitude: a taller bar for January than for February does not necessarily mean that more communication took place on January. Rather, it depicts there was someone in January we talked more to than the person we talked the most to in February. While confusing at first, it goes towards our goal of making the relative importance of different people we communicate with readily visible.

To get information about the absolute number of messages sent and received during a particular month, the user can hover the mouse cursor over the month's name (Figure 4). Hovering over a bar produces information about the users it represents and number of exchanged messages (Figure 5).



Figure 4. Number of messages in a month




Figure 5. Tooltip with information about data in a bar.

It is possible to study how the communication with a particular person has evolved over time by highlighting all the corresponding bars, either by clicking on a bar or using the search box. The search box allows us to enter an e-mail or contact name and select them. A contact can contain several e-mail addresses, and if a contact name is chosen, every e-mail of that contact will be selected with the same color. It also provides an auto-complete feature.



Figure 6. Selected Contacts Bar

The selected contacts view provides the information of what contacts are selected and what color is assigned to them, as depicted in Figure 6. Next to each name there is an 'X' that, when pressed, removes that selection

It is possible to select more than one person and compare different relationships and find correlations between them. In Figure 7 we see that the highlighted person is someone the user didn't communicate with until March 2010, and that has been growing in importance since, becoming the top person in communication volume.




Figure 7. Representation of a selected contact.

3. Evaluation

We gathered 10 users and had them perform a set of six tasks followed by a satisfaction questionnaire. Evaluating this kind of system only makes sense if the users are able to visualize their own data. All users averaged over 5 emails exchanged per week, for at least one year. Before performing the tasks, after a short demo, the users had 5-10 minutes to freely use the prototype. The set of tasks we had the users perform is representative of the questions that users pose themselves when managing their emails and reflecting on their communication patterns:

1. Who have you been talking too the most, recently?
2. Select a person you communicate with and analyze that relationship. Did the rate of conversation grow or decrease? Can you tell why?

3. Select two persons and verify if you can detect some correlation between those two electronic relationships.
4. Look at your Main Timeline. Can you see a clear increase and decrease of communication volume? Why?
5. Can you tell, if you talk with a lot of people or more with a more restricted group?
6. Can you describe the most discussed topics across the time since you've had the email account?

By necessity, these tasks are open-ended and, to some extent, vague. Each user has a different set of information, making it impossible to predict what will be found.

We logged the time spent and errors made as well as any relevant detail or comment from or about the user. After performing the tasks, the users were asked to answer a satisfaction questionnaire. This questionnaire was divided in two parts. The first was the standard System Usability Scale [1]. The second consisted in three essay questions where users could express their opinions.

Only two of our 6 tasks registered any errors (Tasks 2 and 5), and even then, only one error occurred for each one of these tasks. This tends to indicate that our system is easy to use and to understand (suspicion confirmed by the user's feedback described below). Regarding the time spent for each task, Task 2 was the one that took the longer. This can be explained by that task requiring the user not only to select some contact but also to analyze the patterns of email exchange with that person and explain them. As expected, such task is not only long but also very subjective and dependent on each persons' email archive: some users have very long relations with their contact while others have relations of no more than a couple of months. Lastly, some users were also more thorough than others on the analysis required for this task. This last fact is also applicable for most of the tasks with relatively big standard deviation (Tasks like 4 and 5 that required the user to analyze and describe something). Overall, all tasks had fairly low times, showing the visualization is efficient. Also, all users were able to finish all tasks, showing its effectiveness.

	Avg	SD
Task 1	12.9	9.4
Task 2	35	21.1
Task 3	20.6	11.6
Task 4	19.6	12.1
Task 5	18.1	10.4
Task 6	8.4	4.6

Table 1. Task times.

We calculated the SUS score for each user. The average SUS score of 84.3 indicates a very good usability degree perceived by the users during the tests and the small (6.1) standard deviation indicates that each user's SUS's scores are clustered closely around the mean meaning that there's no big discrepancy between user's scores.

In the three essay questions, the most cited positive aspect of our prototype was its usability as users found the system quite easy to use and understand. The visualization was considered to be clean and attractive and the overall subject of the project considered very interesting. The presence of a daily timeline was also considered a very positive aspect. People were pleased with the amount of information the prototype was able to present in an interesting way.

As for negative aspects, most of the users claimed the need to actually see the content of the emails as the topic cloud was considered insufficient for some cases. Another mentioned negative aspect was the fact that it's only possible to search by contact. Some users would like to have the ability to also search by word and see all emails containing those words.

4. Conclusions

Our visualization allows users to get that at-a-glance view of their email relationships. By using our system the user can see the evolution of his/her relationships as a whole or focus on a particular relationship or even set of relationships for comparison. In order to help users understand the 'why' we provide two context indicators: the time and topics for those relationships. Our user tests and satisfaction questionnaire revealed that the users found the system easy and interesting to use, and proved that the visualization is effective and efficient. Future developments would include the ability filter the timeline by keyword and not just contact person, allowing for tasks such as finding the person that has talked more about a particular subject.

5. Acknowledgements

This work was partly supported by FCT (INESC-ID multiannual funding) through the PIDDAC Program funds, and by FCT through the PAELife project AAL/0014/2009.

6. REFERENCES

- [1] J. Brooke. SUS: A quick and dirty usability scale. In P. W. Jordan, B. Weerdmeester, A. Thomas, and I. L. Mclelland, editors, Usability evaluation in industry. Taylor and Francis, London, 1996.
- [2] Perer and M. A. Smith. Contrasting portraits of email practices: visual approaches to reection and analysis. In Proceedings of the working conference on Advanced visual interfaces, AVI '06, pages 389-395, New York, NY, USA, 2006. ACM.
- [3] Steven L. Rohall, et al.. 2004. ReMail: a reinvented email prototype. In CHI '04 extended abstracts. ACM, New York, USA, 791-792.
- [4] F. B. Viégas et al.. Digital Artifacts for Remembering and Storytelling: PostHistory and Social Network Fragments. In Proc. of the 37th Annual Hawaii International Conference on System Sciences (HICSS'04) - Track 4, page 40109.1, Washington, DC, USA, 2004. IEEE Computer Society.
- [5] F. B. Viégas, S. Golder, and J. Donath. Visualizing email content: portraying relationships from conversational histories. In Proceedings of CHI'06, pp. 979-988, New York, NY, USA, 2006. ACM.