

Detecting and Understanding Temporal Patterns for Everyday Collaboration

Danyel Fisher
Department of Informatics
University of California, Irvine

My research group has been involved in attempting to visualize and construct tools to support what we refer to as the “social workscape”—the space of documents, people, and tasks within which each individual is engaged as part of their work. I read these tasks broadly as being intrinsically shared: a document written alone is intended for an audience, and is likely to be presented before others. The day-to-day work that individuals accomplish, both separately and collectively, can be spoken of as parts of “everyday collaboration.” The word “everyday” emphasizes that the collaboration is not the domain of particular collaborative tools, but rather merely the collection of actions around common tasks.

In order to analyze this “workscape”, I exploit recurrent structures that relate the details of specific activity to broader patterns. With others, I have been investigating two sorts of structures. *Social structures* describe the patterns of contact and collaboration that emerge between people. They relate individuals to groups and collaborative activities. Relevantly for this workshop, we also examine *temporal structures*. These describe how patterns of interaction change over time. They highlight the rhythms and trajectories of collaboration, as group members, activities, and topics of concern come and go.

While a more complete discussion of both the social networks and temporal patterns is being presented at the papers track (in a paper co-authored with Paul Dourish, from which this paper is partially adapted), this workshop is a good opportunity to explore my temporal findings more thoroughly.

Some CSCW research has begun to examine the importance of temporal issues. Reddy and Dourish report on an ethnographic investigation of information seeking in a surgical intensive care unit (SICU). Their research finds that the various workgroups in the organization made use of their understandings of the temporal structure of work in order to help coordinate their different activities. Temporal *rhythms* characterize different aspects of the work such as patient flow, medication administration, and shift changes. Similarly, Hudson et al.’s study of research managers also shows that their days have a structure around which their actions are organized; so, the pattern for the day sets expectations and desires about availability and interruptability, for example. In a series of studies, Begole, Tang and colleagues at Sun Labs have examined how some of these temporal structures can be automatically uncovered and incorporated into awareness tools.

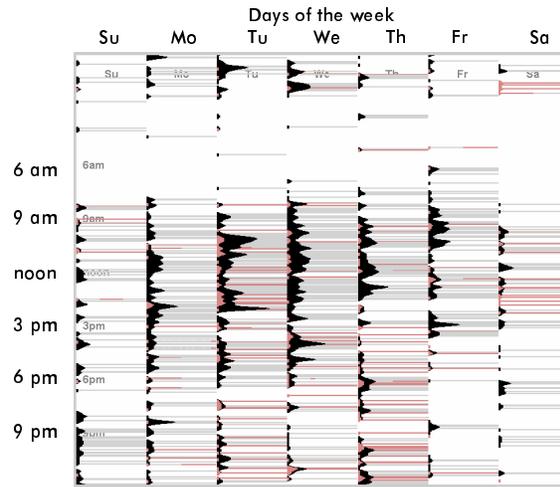
In order to analyze these temporal patterns, I have developed the tool “Soylent”. Soylent is an infrastructure for examining and visualizing the electronic traces of contact and collaboration between users. It collects outgoing email messages into an archive, and displays them to the sender. Because the archive is stored locally and is maintained privately for each user, the Soylent system is a system for visualizing only a single users’ correspondence.

This discussion highlights the results from several Soylent visualizations, each of which shows different temporal structures. One is an actogram, not unlike the one used by Begole et al to trace activity over the course of days and find repeated patterns in action. A second is a “top ten” list,

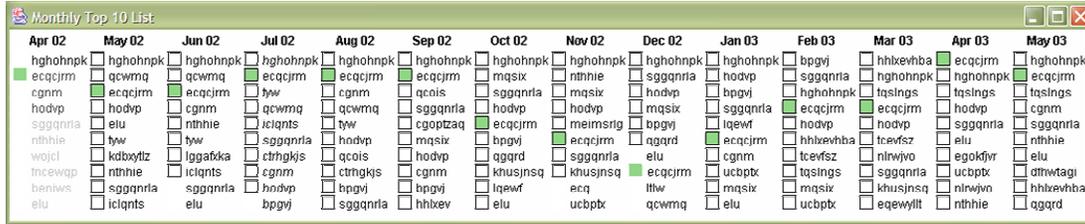
which shows long-term changes in interaction, and a bar chart, which shows patterns of burstiness in interaction.. Last, we discuss a combined network-and-temporal perspective. In passing, this discussion will mention some of the repeatedly seen patterns both from fieldwork and from reviewing our own records.

The Actogram

Our actogram view shows the scope of a daily patten, extended over time. In this diagram, the authors' outgoing messages are each represented by grey bars. The black lines on the edge show the cumulative count of gray bars (averaged over a small time span), and so give a general sense of how many messages are sent during that period. In red ink are visible messages sent to one particular person—in this case, the authors' advisor. Some features, such as the night time rest and the dominance of the daily schedule are visible, as is a Friday evening interruption. Note the clustering on afternoons, especially Thursday afternoons (after a regularly scheduled meeting) and—perhaps more worrisomely—very early Saturday mornings.



While this image is more typical of an academic calendar, we observed much more regular schedules during our field study at a large corporation. Most group members maintained a much more regular schedule, running predictably from 9-5 during weekdays. For their calendars, what stood out most dramatically was the interruptions and exceptions to the normal temporal patterns: the days that didn't match the usual routine. During our interviews, the respondents were able to tell stories about these exceptional occasions, and could identify the few occasions when their schedule digressed. This reinforces our belief that these patterns are meaningful and salient to their users.



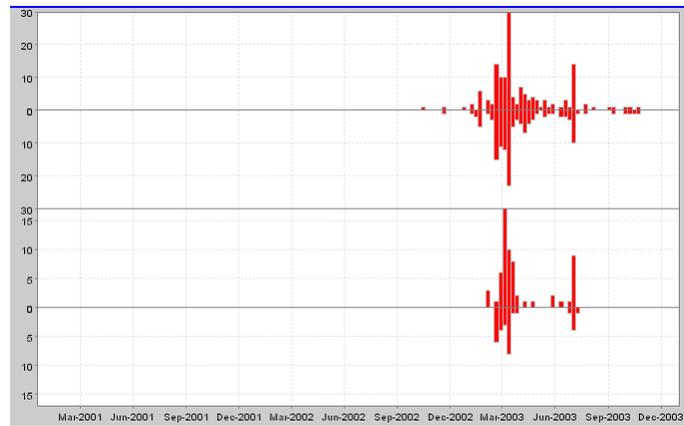
The Top Ten List

The top-ten list was suggested by a user, and has turned into an invaluable tool for understanding social transitions. It portrays, on a monthly or weekly basis, the top ten outgoing correspondents. The figure above highlights in green boxes one name appearing monthly from one April until the following May. The green-labelled name appears in the top few each month, until October—when that correspondent began a period of working remotely. When the respondent returned in January, the level of mail began to rise again.

From the top ten list, we have derived a few regularities. For example, many users showed us the “intern” pattern, in which several new names appeared in the top couple of names for a small period of time before entirely disappearing. Users were also able to show particular times when a single correspondent was very prominent for a special project.

The Bar Chart

While the top-ten list compares a list of users, the bar chart examines one individual. The four sections of the bar chart (from top to bottom) are date-indexed views of the number of weekly outgoing and then incoming email messages, and the number of weekly outgoing and incoming attachments.



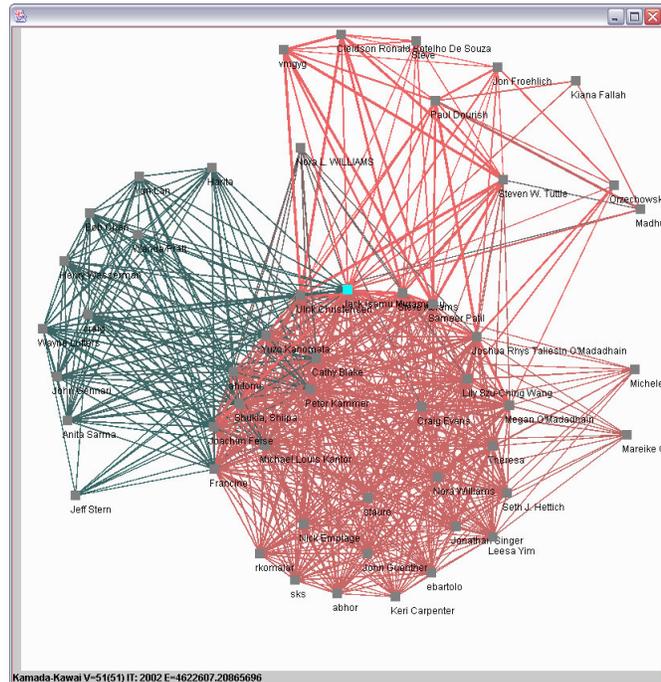
This particular image shows a writer corresponding with his co-author; the axes of the chart run over a period from March 2001 to December 2003. The paper they worked on had a deadline (in March 2003, the end of the first large burst), but they continued to keep in touch afterward; there was a second spike when they submitted revisions (in August 2003) for that paper. It is informative to examine some of the asymmetries: note, for example, that the co-author sent more attachments (presumably, drafts) earlier in the editing cycle, while the respondent here sent the most later on. This turns out to be a not-atypical example of a group building toward a deadline; deadlines often look like this rapid buildup period.

Combined Social and Temporal Perspective

Last, I examine a view with a combined social and temporal perspective. This image depends on Soylent's network visualization techniques.

In this color image, the user is examining the person highlighted in blue ("Jack"), near the center of the screen. An edge between a pair of names means that the user of the system had carbon-copied a message to both of those names in at least one message. In addition this view is restricted to people who are tied to the name under focus. Thus, every person listed here has been carbon-copied on a message to Jack.

In addition, this diagram has been colored by the age of edges. Green ties are old ones; they indicate messages from a year or more back. Bright red edges are new—within the last few weeks. Last, the dull red edges at the bottom are a few months old.



Several interesting combination social- and temporal characteristics are visible from this viewpoint. The large cluster at bottom and the green cluster to its left are closely related; they are a large social group to which Jack belonged (as did the interviewee whose mailbox this is); the green ones are people who had graduated, and thus had fallen out of recent contact. The social group had been on hiatus recently, hence the darkening cluster; however, Jack was still active, interacting with many people in several contexts (at top).

Conclusion

This paper has collected a series of temporal visualizations on social interaction from the Soylent tool. The examples presented here show some of the varieties of information; the discussion has connected the visualizations with the real user stories that they elicited. Through these stories, a number of recurrent patterns have surfaced which seem to be at least somewhat reliable in discussing how individuals structure their interactions.

These lead to design possibilities for tools for everyday collaboration. Tools that are aware of the temporal structure of interactions—if nothing else, paying attention to those names that have long since faded from view—can lead to new ways of thinking about our software design.

Bibliography

1. Begole, J., Tang, J., and Hill, R. 2003. Rhythm Modeling, Visualizations and Applications. Proc. ACM Symp. User Interface Software and Technology UIST 2003. New York: ACM.
2. Hudson, J., Christensen, J., Kellogg, W., and Erickson, T. 2002. 'I'd be overwhelmed, but it's just one more thing to do': Availability and interruption in research management, Proc. ACM Conf. Human Factors in Computing Systems CHI 2002 (Minneapolis, MN), 97-104. New York: ACM.
3. Reddy, M. and Dourish, P. 2002. A Finger on the Pulse: Temporal Rhythms and Information Seeking in Medical Work. Proc. ACM Conf. Computer-Supported Cooperative Work CSCW 2002 (New Orleans, LA). New York: ACM.