Social and Temporal Structures in Everyday Collaboration

Danyel Fisher and Paul Dourish
School of Information and Computer Science
University of California, Irvine
Irvine, CA 92697-3425
danyelf@acm.org, jpd@ics.uci.edu

ABSTRACT
Everyday work frequently involves coordinating and collaborating with others, but the structure of collaboration is largely invisible to conventional desktop applications. We are exploring ways to support everyday collaboration by allowing applications access to the social, organizational, and temporal settings within which work is conducted. In this paper, we present two generations of systems supporting everyday collaboration, focusing on ways to recover and represent the temporal and social structures of online activity.

Categories and Subject Descriptors: H.1.2 [Models and Principles]: User/Machine Systems – human factors; H.5.3 [Information Interfaces and Presentation]: Group and Organizational Interfaces – theory and models.

Keywords: awareness, social networks, collaboration patterns.

NTRODUCTION
Although CSCW research distinguishes between single-user and multi-user technologies, the boundary between individual and collaborative work is much less clear. many tasks carried out with ostensibly single-user tools are, on closer examination, collaborative in nature. Documents are frequently assembled out of pieces provided by others; presentations may be crafted as parts of larger projects, and are designed to suit the needs of both presenters and audiences; and spreadsheets are often used to coordinate collective activity. We refer to this as "everyday collaboration."

While CSCW research has developed a range of tools to support explicitly collaborative work, everyday collaboration is supported only poorly. The essence of everyday collaboration, as it occurs in the real world, is that it is carried out using conventional tools; most collaborative documents are produced using Word rather than multi-user text editors. Features such as change management, versioning, and templates incorporated into some applications provide basic support for collaboration, but little of the insights gained from CSCW research have been applied to everyday collaborative tasks. The barrier between single-user and collaborative applications makes everyday collaboration challenging. Our current research is directed towards this problem, and asks, how can the single-user experience be reconfigured to support the collaborative tasks being carried out?

CHI 2004, April 24–29, 2004, Vienna, Austria.

We propose that an effective solution is not to turn single-user applications into groupware tools, but rather to reveal the collective activity that is already being carried out through those tools. We want to help people coordinate their work by providing them with ways to see how their work is connected to that of their colleagues. To do this, we are exploring the potential for using *single-user tools as technologies supporting awareness*.

Awareness is a widely noted aspect of collaborative practice. Much research into collaboration in practice has shown how the explicit, task-focused aspects of activity are complemented by and coordinated through passive mutual monitoring which provides people with an ongoing, informal awareness of each other's activity. Heath and Luff [6], for example, showed how operators in London Underground control rooms delicately coordinated their separate activities by informally monitoring each others' actions and arranging their own work to match. Others have noted similar processes at work in electronically mediated collaboration [5, 17]. They allow people to coordinate their own work with that of others. Can we incorporate support for informal awareness of everyday collaboration into conventional applications?

We tackle this by exploiting recurrent structures that relate the details of specific activity to broader patterns. We 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. *Temporal structures* 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.

By making social and temporal structures of collaboration apparent in interaction, our goal is to help people make sense of the activities around them, and so help them to coordinate their work with others. Rather than seeing interaction lists as statically stuck in "now," a structural

perspective can derive a history of interactions; rather than seeing colleagues purely as individuals, a structural perspective helps to identify roles and relationships in social groups.

Our research is driven by two hypotheses:

First, that there is structure in interaction, which can be found in electronic traces of activity.

Second, that this structure can be used to build contextualized awareness tools that successfully present an appropriate selection of information.

To test the first hypothesis, we have developed Soylent, an infrastructure and exploratory tool to determine what elements of social and temporal organization we might be able to identify in electronic records of everyday activity. We outline the design of Soylent, and discuss an initial user engagement, carried out at an external site, to test the utility of this tool. This user engagement generated a series of findings about social and temporal structure which we have incorporated into further development activities.

To test the second hypothesis, we have developed awareness tools based on structural information. We present TellMeAbout, an initial client that uses the Soylent infrastructure to provide end users with an understanding of the structures within which their work is embedded. TellMeAbout can also be used as a dynamic awareness tool which reveals shifts in attention and focus as everyday collaboration is carried out.

MAKING STRUCTURE VISIBLE

The essence of our approach is to make the social and temporal settings of everyday collaboration visible in traditional, single-user interaction.

Scenario

To describe something of what we mean by temporal and social settings, consider how our proposed approach could help resolve some contemporary problems.

> *Joe is a software salesman operating in the Western region. One morning, he gets a call from John, a potential client whose name he doesn t immediately recognize. As they talk, he realizes they had been in touch by email several months ago. He wants to locate that information to remind himself how they know each other, so he quickly starts to skim through his email records. The name doesn t show in his address book; the contact had been fleeting, and hadn t seemed like a useful entry. A quick search on his name doesn t come up with anything; many mail clients are poor at checking carbon-copied names.*

With a network and temporal awareness system, however, the history of interaction would become quickly apparent:

> **TellMeAbout -person johnsmith@bigco.com**
> **John Smith** appeared three times between June and July, 2001, always with **James Doe.** (You have sent him three messages; **James Doe** has cc'd him on one).

> *With this information, Joe is quickly able to pull up a history of this otherwise-obscure person, and is reminded of their past working context.*

Approach

As suggested here, our approach is based on three principles. First, we are interested in exploring the use of automated mechanisms for determining structural information. Second, we use structures to explain activity, rather than presenting the structures themselves, which may be less immediately meaningful. Third, since we believe that the information can only be adequately interpreted by the end user, our goal is not to have systems take actions on the user's behalf, but rather simply present relevant information for examination.

Another approach to making structure visible, akin to that explored by ContactMap [24] or related systems, would be to allow people to describe their social groups and daily patterns to an application, which could then use this information to present relevant information about ongoing activities and as an interface to other tasks (e.g. email.) For our purposes, this approach has a number of drawbacks. It imposes significant costs of time and effort on the part of the user, it relies on effective recall and anticipation for accurate results, and it must be repeated when the patterns change. However, in the explicit approach, one can at least be sure that the networks that are generated have salience to the end user. We will discuss this trade-off in more detail later.

BACKGROUND

Both social networks and temporal studies have developing histories in studies of group interaction.

Social Networks

The concept of social networks has long been used as an analytic tool in social science [21]. Perhaps the earliest popular manifestation of this line of research was Milgram's pioneering investigation of the so-called "small world phenomenon," leading to the observation that two randomly selected individuals in the United States could be related to each other through a chain of six or fewer first-name acquaintances [13]. Broadly, social networks describe the ways in which sets of people are linked together by some relationship. One can create social networks based on kinship patterns (a family tree), communication patterns such as Christmas card exchange, or working relationships, such as the graphs describing the film appearances of Kevin Bacon. As a tool in the social sciences, social networks allow comparative analysis between different social settings. Analyzing the structural properties of social groups reveals properties and consequences of particular social organizations.

Social network analysis has been used occasionally within HCI and CSCW research to uncover the social structure of electronically-mediated activities. Mackay [12], for example, used a social network approach when studying the spread of customizations in UNIX software; social network analysis identified the critical role of specific individuals in linking together other groups and supporting the widespread diffusion of innovation. Wellman and colleagues [22, 23] have demonstrated the applicability of

social network models to virtual communities and other settings of computer-mediated interaction. Pickering and King [16] study the use of computer-mediated communication in organizations and highlight the role of CMC in maintaining weak links, especially between "dispersed occupational communities."

We use social network analysis as a mechanism for discovering and understanding group structure. Individuals' roles at particular times can be partially understood by their network positions relative to each other, and relative to the others with whom they collaborate.

Temporal Patterns

In addition to social structure as a source of information about collaborative work, we have also found that temporal structure of activities provide individuals with resources to manage their coordination. A wide range of sociological investigations have noted the importance of temporality in social affairs [19, 26]. Similar concerns have surfaced in empirical CSCW research. Reddy and Dourish [18] 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 [7]. In a series of studies, Begole, Tang and colleagues at Sun Labs [2, 3] have examined how some of these temporal structures can be automatically uncovered and incorporated into awareness tools

Temporal patterns emerge at different scales. While the Sun work focuses on daily and weekly patterns, longer-term patterns are also a focus of research attention. Analysis of interaction between individuals and in online discussion spaces have revealed long-term patterns in collaboration that can characterize the nature of collaboration and communication in electronic settings [8, 25]. Studies such as these suggest that automated analysis may be able to serve as the basis for individual solutions.

SOYLENT INFRASTRUCTURE & VISUAL TOOLS

Our first hypothesis is that patterns of contact and collaboration leave electronic traces. This led to the design of an infrastructure for examining, visualizing and inspecting these traces. This infrastructure and its associated visual tools are collectively called *Soylent*.

Infrastructure

Soylent analyzes aspects of user activity to uncover social and temporal information. Our primary source of information is electronic mail activity, but Soylent is designed to access other sources of information flexibly, and to integrate across them. Soylent gathers information from archived email record headers: both incoming and outgoing messages are examined, indicating all participants in messages, their sending times and dates, and any attached files. This infrastructure builds a database of records of interaction.

Clearly, automatically gathering and analyzing traces of individual activity poses significant potential for invasion of privacy. It is important to note two things here. First, our system (and the field trial we shall shortly discuss) is constructed so that each user runs an application which gathers information purely on their own behalf and purely for their own use; the database in which records are stored is generally a private database for each individual, and information about each person's activity is shared with no others. Second, we reflect back to each user purely information about their own activities; information about one user is not combined with information about any other (and, indeed, to do so would be to defeat the object here.)

Visualization Tools

The database is then queried by the visual tools. The visual tools are designed to highlight temporal and social structure in the stored communication. In this paper, we will primarily discuss two of these views, the *network diagram*, which shows social structure; and the *top ten list*, which shows longer-term changes in interaction.

The network view is an ego-centric view of social structure, explained in figure 1. Pairs of correspondents are tied together if they were co-recipients of a message sent by the user of the tool; once the data for many messages is aggregated, tight clusters suggest groups who are closely related, while disconnected entities suggest sets of people who have little or any relation. The user can filter the view in a number of ways, removing groups of correspondents who are sparsely connected, color-coding individual groups, and limiting the ties to those that occurred at particular times.

Figure 2, for example, shows one authors' email connections. Very infrequent connections have been filtered out; pairs of people must have been sent three or more messages to be shown here. The large cluster at the lower right are participants in a workshop, for example, while the group at top middle is the set of people involved in a publishing collaboration: a co-author connects to

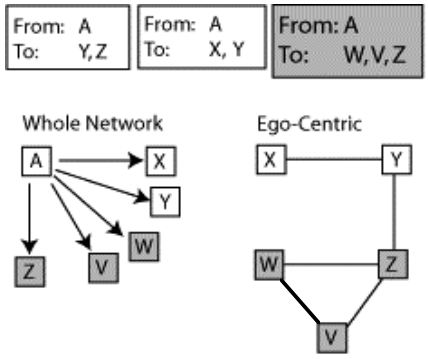

Figure 1: Ego-Centric Network Views. In this diagram, A has just sent a message to W, V, and Z, As a result, W, V, and Z are connected to each other in the ego-centric view.

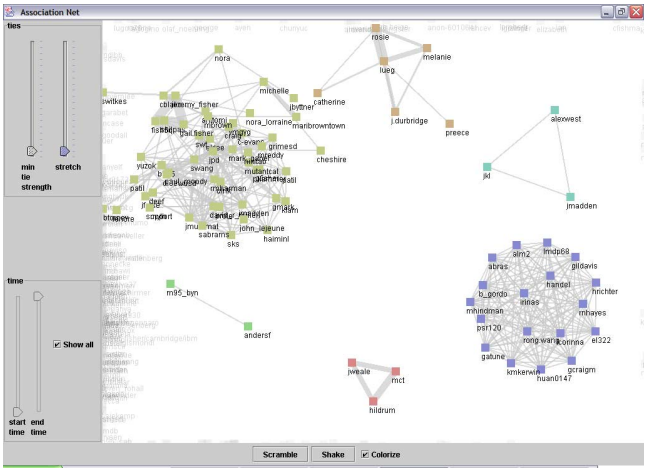

Figure 2: The Soylent network view.

The *monthly and weekly top ten lists* show the top few correspondents separated across time. The view shows a list of the top ten message recipients for each month or week; changes in these lists show shifting patterns of correspondence. Other views not discussed here present different "slices" through the data, such as group overlap and daily patterns of contact.contributors, an editor, and a publisher; the editor and the publisher are strongly tied to each other. The large cluster at the left is a tightly connected cluster of work and social interactions.

The data recorded by Soylent, especially over long periods, is both rich and complex. Each view provides a range of ways of exploring the data from some particular perspective. The views are designed to be used together; they operate in concert and allow an analyst to triangulate on the data and highlight correlations between different perspectives.

USER ENGAGEMENT

Soylent was not designed as an end-user tool, but rather as a tool to allow us to verify our first hypothesis. This hypothesis has two parts – the first concerning the presence of *structure*, and the second that this structure leaves *electronic traces*. Accordingly, we set about determining whether Soylent could be used to find structures and roles in electronically-gathered information.

A group of fifteen users from a large software development firm used Soylent. Their job roles varied, we interviewed managers, software developers, designers, and administrative assistants. We used the system as a way of directing an inbox walkthrough [20] organized not by sending order but by social context; the features of the system triggered conversations about user's social interactions and identified important events.

First, we helped the users install the software on their PCs, and introduced them to the functionality. Next, we gave them a workbook to guide them as they used the software to explore their own data. After a few days to explore the system, we sat down with them again to go through what they had learned, sitting with them at their computers so that we could explore and discuss the data with them.

Results: Finding Structure

Our subjects were proficient at identifying the relevance and connection of the different patterns they saw. Structures and patterns emerged and were recognized. Most importantly, perhaps, during our conversations the users would tie the connections between correspondents, and shifts in participation over time, to stories and events. These events – "the arrival of the summer interns"; "the fall patent negotiation" – would be used to explain how people were connected with each other. In other words, the patterns that they found were meaningfully related to their work. All users found structures that made sense to them, both in social and temporal dimensions of the data.

Results: Detecting Recurrent Patterns

As we discussed the Soylent display with our subjects, we found that certain structures and relationships arose repeatedly. Some people had important roles: they were involved in a variety of contexts, or had a specialized position. Others acted as bridges between two different groups. Groups had cores and peripheries. Over time, people would become more or less involved in groups; those groups themselves would appear and fade away. People would work on regular schedules – a usual morning email check, or a typical end for their day. In other words, our users showed us recurrent structures within their data.

We characterize these recurrent structures as "patterns" – sets of relationships that occur across different individuals, different sets of users, different periods of time, and so forth. These patterns identify characteristic structures in either the social or temporal aspects of interaction. These patterns are associated with user stories – distinctive moments and particular issues illustrating social relationships between sets of people. Yet these patterns occurred in a variety of interviews, and despite the differing details, point to underlying similarities in social relationships. We present a number of these below.

The Onion Pattern ("Core and Periphery")

One common form of collaboration pattern is that of a core and periphery: a project has a central team, for example, and a number of consultants who are peripherally involved. The *Onion* pattern identified a central core, often a tight clique, surrounded by an outer less-tightly-connected periphery.

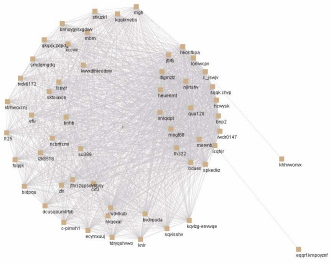

Figure 3: Onion pattern.

The Nexus Pattern ("Multiple Roles").

Certain individuals (such as team managers secretaries) had a distinctive appearance in the network. They would often be at the center of a hub-and-spoke pattern. This "nexus" pattern could be interpreted as a series of multiple collaborations or projects. An employee, reporting back to his manager on a series of projects, would have the manager tied into each of these different contexts.

Figure 4 highlights two examples of the "nexus pattern" in the upper section. The left-hand nexus center, for example, is a joint researcher with the informant on a series of different projects, and shares both work and social ties. The circle of connected people describe some of the different contexts that they share.

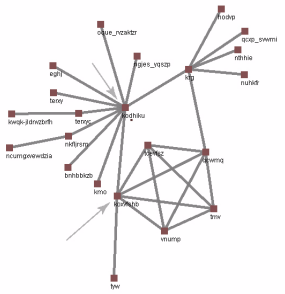

Figure 4: Nexus pattern.

The Butterfly Pattern ("Dual Roles")

The node pointed by the arrow in figure 5 illustrates these two roles. The central node designates a student who was involved in both a colloquium (the clique on the left) and a research team (the clique on the right). The sender of the mail both coordinated the colloquium and was an active member of that team, and so frequently sent mail to both lists of members.The butterfly pattern — named for its two large "wings" surrounding a single center — was a pattern visible when an actor could be found in two roles. A member of a design team would also be a member of a research group, for example. Visually, this looks like two disconnected groups joined by a single actor.

We have suggested that temporal rhythms and social networks can shed light on each other; this example is a good one to understand how they mutually interact. The colloquium was on a very specific date; while there were a number of email messages sent to it, they stretched only for the several months before the colloquium. In contrast, the research group involvement was longer-tem, and involved interactions both before and after the colloquium.

Interestingly, the butterfly pattern is not just a social structure but also, often, a temporal structure. The Soylent network visualization allows users to adjust their window of interest in order to look at network changes over time. Using that tool, it becomes clear that the two "wings" of this butterfly have different temporal extents.

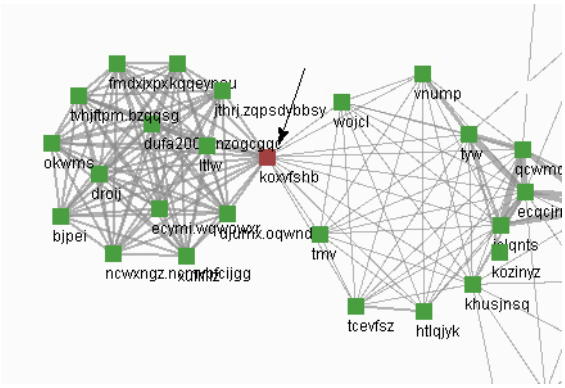

Figure 5: Butterfly pattern.

Shifting Involvement (Long-Term Change)

Both gradual and sudden changes in project teams were a recurrent theme among most workers. The top-ten list made it clear that, in addition to recurrent sets of correspondents, most participants had a series of changing team members: people would become prominent for a period of time, then slip away. These changes happened with notable events: the starts or ends of collaborations; joining a new team; the arrival and departure of summer interns.

Figure 6 shows a list for one authors' outbox. This image has been highlighted to show the changes in the list. Note that the highlighted person is in the top few names during the first months of the visualization; in October, it drops to a much lower ranking, and only rises again in January. This corresponds to that person traveling across the country, working at another site; it has been frequently observed [e.g. 10] that increasing distance tends to decrease the frequency of interaction and communication.

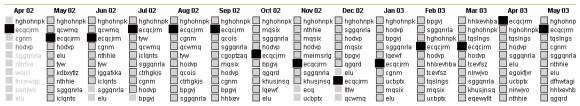

Figure 6: Monthly top ten display.

Figure 7 illustrates a similar concept, that of sporadic interactions. This view shows messages flowing in both directions between two people who worked together on two specific projects, separated by almost a year.

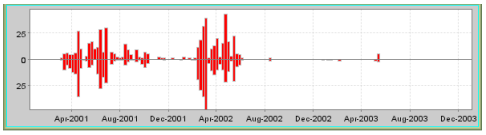

Figure 7: Week-by-week view of one person.

INTERPRETING AND USING PATTERNS

Although we have presented the patterns here as distinct structures, this is not how they occur in real data. Instead, patterns occur in concert with others, overlapping and interacting with each other. We understand these different patterns to mean different sorts of things. At high strength, patterns reflect repeated and significant connections; at low strength, patterns highlight a general topology of interactions. For example, one informant's immediate work group looks like an onion pattern when looking at all her mail, but shows a nexus pattern when looking at the most frequent recurrences. While she had sometimes mailed her whole team (forming the onion's outer layer), she also worked specifically with different sets of people in the group, rotating between tasks with a designer, the software team, and so on.

Nonetheless, the occurrence of these patterns throughout the study suggests that they have important properties for analysis.

They are *recurrent*, in a number of ways For a single user, they arise regularly through time and with different groups; in organizations, they recur for different individuals; and

they recur across organizations (in our own group's data as well as our field organization.) Although the precise details change, the general structures operate across multiple examples.

The are *findable*. The patterns that we have shown here are both uncoverable in traces of user action, and quantifiable so that they can be extracted by a software system and presented to users for analysis.

They serve to *render work meaningful*; that is, the structures provide a framework for interpreting and understanding what is going on. People can make sense of particular events or particular objects by relating them to a larger whole through the structures that time them together.

They operate *across different activities and applications*; they denote broader patterns of contact and collaboration.

These properties suggest, first, that these structures might be a useful basis for interactive technologies that can help mesh separate and individual activities into a broader setting. Studies of social networks and temporal patterns show not only that these structures exist, but also that people actively employ them in the course of their work. This further suggests, then, that they can provide an effective foundation for designing novel awareness technologies. Accordingly, the second challenge for our investigation is to harness this information as part of a tool for presenting information about everyday collaboration.

AWARENESS IN EVERYDAY COLLABORATION

To recap, our first hypothesis was that there is detectable structure in patterns of contact. Our user engagement was designed to test this, and, indeed, this discussion of the patterns has shown that this structure is visible and relevant to our users.

The second hypothesis is that we can use this structure in situated awareness tools. In order to demonstrate this, we have implemented systems that capitalize upon the Soylent infrastructure to provide end users with information about the social and temporal structures within which everyday collaboration is enmeshed.

The primary role for the social and temporal patterns in awareness systems is to tie specific people or moments of time to broader trends. That is, when engaged in activity connected to one individual, social patterns can be used to place that person in a context of collaborators and peers, while temporal patterns can provide greater specificity by placing current activities within a temporal context. Our original goal was to make the regularities within everyday collaboration clear; the social and temporal patterns that Soylent can uncover are just these sorts of regularities.

Regularities, of course, can take different forms, and the relationship of specific instances to broad patterns can, likewise, manifest itself in different ways:

- *Structures may bridge between different activities*. Single tasks often spread out between multiple applications. By uncovering relationships between individuals and artifacts, collaboration structures may be able to show how the separate activities located in different applications are, in fact, part of the same task. This is not dissimilar to the approach taken by Kaptelinin [9], although in this case we depend on understandings of social interaction rather than computer interaction.
- *Structures may help disambiguate different activities*. By placing individual activities in a broader context, the collaboration structures may help not only relate but also separate activities by linking them to different higher-level patterns.
- *Structures may throw current activities into relief*. Frequently, the interesting relationship between current activities and patterns of past action are that current activities do not fit the pattern of the past. Clearly, we do not always follow the same path. Activity might be meaningful not because it fits the pattern, but because it does not.

Of course, while a software system might be able to note commonalities and differences between specific activities and broader patterns, the meaningfulness of those relationships is a matter purely for users to determine. Our systems do not take any action on behalf of a user, but rather suggest how activities are related to recurrent structures of contact and collaboration.

Awareness tools are generally designed as an adjunct to traditional tools, and so must be easily inspectable without detracting from the primary task. Soylent's network view can be intimidatingly complex, and is clearly inappropriate for this. Patterns provide a way to reduce this complexity. Just as user interviews centered on patterns as the basics of stories and events, so too we use patterns to summarize and simplify user interactions. An analysis based around the patterns helps find important events and group dynamics; we collect the information we gain from that analysis into a simplified interface that shows minimal information but allows exploration. We have developed an API to expose this refined network information, and are developing various awareness applications using it.

TellMeAbout: A Sample Tool

To exercise this interface, we have built a system called TellMeAbout to use these insights for an end-user tool. TellMeAbout is the simplest client to the Soylent infrastructure; it describes how particular individuals, objects, or groups are situated, socially and temporally.

```
> TellMeAbout -person bmarkham
79 messages since Mar 23 '01,
  most recently May 12 '02
  especially
    Mar 26 '01-May 7 '01,
    May 28 '01-Jun 11 '01
closest connections include (gayle)
```

As in the sample output printed above, TellMeAbout gives temporal information about the user, including the first time the user was ever emailed; the most recent time they were emailed, and, if relevant, particular periods during which they were most closely involved. In addition, it shows network information about what other users they are closely connected to. Looking at temporal structures

provides characterizations of the rhythms of contact; looking at social structures helps tie this person to others. This quickly solves dilemma faced by Joe in our opening scenario: his client's context would be visible, the range of their past interaction highlighted, and person who introduced them would be visible.

While TellMeAbout is an extremely simple tool, it is the basis for more elaborate tools that provide a richer experience. For example, when connected to an email server, TellMeAbout can generate information about the people from whom the user has most recently received email. Connecting this to a Tickertape-style display [15] allows a continually-updated display of information relevant to the current situation, updated continually as new email arrives or is sent.

Linking TellMeAbout into Other Tools

In both "standalone" mode and "tickertape" mode, TellMeAbout stands separate from other applications, although it is part of the conventional single-user desktop. We are currently exploring opportunities to integrate TellMeAbout and other clients for Soylent services into existing applications; augmenting an email client, for example, with information about the temporal structures of message arrival and delivery, and the social structures surrounding particular individuals. We are also examining calendar applications as both displays and sensors [14].

DISCUSSION AND FURTHER WORK

TellMeAbout is one of a range of tools that we are developing which make information gathered by the Soylent infrastructure. Our ongoing work aims to refine the details of presentation, to smooth the interoperation between tools, and to verify their effectiveness across domains. However, the experiments presented here have borne out both hypotheses we set out at the start.

A number of researchers investigating collaborative work through social networks have expressed skepticism concerning the efficacy of automated approaches to building social networks [11, 24]. Automatic techniques are likely to suffer from two problems – first, that they may miss certain important data, and second, that they lack external validation which would ensure that collected data are actually meaningful. Electronic data fails to capture some modes of interaction; an automatically-generated audit of communication will inevitably be a limited representation of work. However, the approach we have been exploring does not attempt to gather formal social networks for the purpose of sociological analysis. Although we use social network research to motivate and inform our approach, the data we gather are more informal. They are not used for formal analysis but rather are offered to users for interpretation. This mitigates the problems of inaccuracy.

One other area that we are currently exploring is further sources of information about activities. The filesystem is a potentially very rich source. A great many user files are related to others, having either come from or been sent to others via the web, email, or other transfer mechanisms. Further, conventional filesystem structures tend to cluster these relationships with others (one folder contains drafts of a paper exchanged with a few others, while another contains presentation slides collected from a different set of colleagues.) We have developed tools to integrate information about filesystem activity and continue to explore their potential.

CONCLUSION

John Donne observed that "no man is an island," and social scientists such as Becker and Bourdieu have observed that even such a quintessentially solitary activity as artistic creation is enmeshed in a complex web of social relations [1, 4]. Similarly, we have noted that much of the activity carried out through standard office applications is, in essence, collaborative; conventional tools are used to coordinated and conduct a wide range of interactions and everyday collaborations. Everyday collaboration is poorly supported by conventional collaborative tool development.

We have proposed an approach to supporting everyday collaboration. We are developing applications to help people understand, coordinate and manage the collaboration that they achieve through conventional "single-user" applications. Our approach is to make people's "social workscapes" visible – to let them see the structure of their collaborative interactions. We have focused in particular on two sorts of structure – social structure (that is, the patterns of contact and collaboration between people) and temporal structure (that is, how those contacts and collaborations are distributed through days, weeks, and years.)

The critical question to be answered is, to what extent is it possible to uncover usable temporal and social structures from traces of electronic activity?

We have been addressing this question through a combination of software development and user engagements. Our investigations show that, first, it is possible to automatically derive information about recurrent patterns of contact and collaboration in everyday tool use, second, these patterns are recognizable and meaningful to people in terms of their everyday work, and third, that they can be used to develop awareness tools for either standalone use or for augmenting traditionally "single-user" applications with information about the collaboration that is being carried out through them. Our goal, through the use of mechanisms such as this, is to break down the traditional barrier between single-user and collaborative activity and to provide users with mechanisms that let them coordinate their work across the multiple applications and activity contexts that characterize everyday collaboration.

ACKNOWLEDGEMENTS

We are grateful to Katie Faust, Paul Moody, Bonnie Nardi, David Redmiles and Andre van der Hoek for their assistance in the conduct and presentation of this research. We are especially grateful to the volunteers who participated in our user study. This work was supported in part by the National Science Foundation under awards IIS-0133749, IIS-0205724 and IIS-0326105.

REFERENCES

1. Becker, H. 1982. Art Worlds. Berkeley, CA: University of California Press.
2. 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.
3. Begole, J., Tang, J., Smith, R., and Yankelovich, N. 2002. Work Rhythms: Analyzing Visualizations of Awareness Histories of Distributed Groups. Proc. ACM Conf. Computer-Supported Cooperative Work CSCW 2002 (New Orleans, LA), 334-343. New York: ACM.
4. Bourdieu, P. 1993. The Field of Cultural Production. New York: Columbia University Press.
5. Dourish, P. and Bellotti, V. 1992. Awareness and Coordination in Shared Workspaces. Proc. ACM Conf. Computer-Supported Cooperative Work CSCW 92 (Toronto, ON). New York: ACM.
6. Heath, C. and Luff, P. 1992. Collaboration and control; crisis management and multimedia technology in London Underground line control rooms. Computer Supported Cooperative Work, 1(1), 69-94.
7. 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.
8. Jones, Q. 2003. Applying Cyber-Archeology. Proc. European Conf. Computer-Supported Cooperative Work ECSCW 2003 (Helsinki, Finland). Dordrecht: Kluwer.
9. Kaptelinin, V. 2003. UMEA: Translating Interaction Histories into Project Contexts. Proc. ACM Conf. Human Factors in Computing Systems CHI 2003 (Ft. Lauderdale, FL)., 353-360. New York: ACM.
10. Kraut, R., Egido, C., and Galegher, J. 1990. Patterns of Contact and Communication in Scientific Research Collaboration. In Galegher, Kraut, and Egido (eds), Intellectual Teamwork: Social and Technological Foundations of Cooperative Work. Nillsdale, NJ: Lawrence Erlbaum.
11. McDonald, D. 2003. Recommending Collaboration with Social Networks: A Comparative Evaluation. Proc. ACM Conf. Human Factors in Computing Systems CHI'03 (Ft. Lauderdale, FL). New York: ACM.
12. Mackay, W. 1990. Patterns of Sharing Customizable Software. Proc. ACM Conf. Computer-Supported Cooperative Work CSCW 1990 (Los Angeles, CA). New York: ACM.
13. Milgram, S. 1967. The Small World Problem. Psychology Today, 2, 60-67.
14. Mynatt, E. and Tullio, J. 2001. Inferring Calendar Event Attendence. Proc. ACM Conf. Intelligent User Interfaces IUI 2001. New York: ACM.
15. Parsowith, S., Fitzpatrick, G., Segall, B., and Kaplan, S. 1998. Tickertape: Notification and Communication in a Single Line. Proc. Asia Pacific Computer-Human Interaction APCHI 98 (Kangawa, Japan).
16. Pickering, J. and King, J. 1992. Hardwiring Weak Ties: Individual and Institutional Issues in Computer-Mediated Communication. Proc. ACM Conf. Computer-Supported Cooperative Work CSCW 1992 (Toronto, ON), 356-361. New York: ACM.
17. Prinz, W., Mark, G., and Pankoke-Babatz, U. 1998. Designing Groupware for Congruency-in-Use. Proc. ACM Conf. Computer-Supported Cooperative Work CSCW 1998 (Seattle, WA), 373-382. New York: ACM.
18. 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.
19. Sorokin, P. and Merton, R. 1937. Social Time: A Methodological and Functional Analysis. American Journal of Sociology, 42(5), 615-629.
20. Tyler, J. and Tang, J. 2003. When Can I Expect an Email Response? A Study of Rhythms in Email Usage. Proc. European Conf. Computer-Supported Cooperative Work ECSCW 2003 (Helsinki, Finland). Dordrecht: Kluwer.
21. Wasserman, S. and Faust, K. 1994. Social Network Analysis. Cambridge: Cambridge University Press.
22. Wellman, B., Boase, J. and Chen, W. 2002. The Networked Nature of Community: Online and Offline. IT & Society, 1(1), 151-165.
23. Wellman, B, Salaff, J., Dimitrova, D., Garton, L., Gulia, M., and Haythornthwaite, C. 1996. Computer Networks as Social Networks: Collaborative Work, Telework and Virtual Community. Annual Review of Sociology, 22, 213-238.
24. Whittaker, S., Jones, Q., Nardi, B., Terveen, L., Creech, M., Isaacs, E., and Hainsworth, J. 2002. ContactMap: Using Personal Social Networks to Organize Communication in a Social Desktop. Proc. ACM Conf. Computer-Supported Cooperative Work CSCW 2002 (New Orleans, LA). New York: ACM.
25. Whittaker, S., Jones, Q., and Terveen, L. 2002. Contact Management: Identifying Contacts to Support Long-Term Communication. Proc. ACM Conf. Computer-Supported Cooperative Work CSCW 2002 (New Orleans, LA). New York: ACM.
26. Zerubavel, E. 1985. The Seven-Day Circle: The History and Meaning of the Week. New York: The Free Press.