ContactMap: using personal social networks to organize communication in a social desktop

STEVE WHITTAKER[1], QUENTIN JONES[2], BONNIE NARDI[3], MIKE CREECH[4], LOREN TERVEEN[5], ELLEN ISAACS[6], JOHN HAINSWORTH[7]

A shared physical workplace is a rich social and informational environment. Tasks such as managing communication commitments, keeping track of collaborators and friends, and "social data mining" of local expertise for advice and information are supported naturally by a shared physical workplace. However, many people now collaborate remotely using tools such as email and voicemail. Our field studies show that these tools do not support processes such as social reminding or social data mining. In part, this is because these tools are organized around *messages*, rather than *people*. In response to this problem, and informed by our field studies, we created ContactMap, a system that makes *people* the primary unit of interaction. ContactMap provides a structured visual representation of the important people in the user's social network; this representation directly supports social reminding and social data mining. We conducted an empirical evaluation of ContactMap, comparing it with traditional email systems on tasks motivated by our fieldwork. Users performed better with ContactMap than their usual email system, and they strongly preferred ContactMap for these tasks. Analysis suggests that ContactMap's visual interface supports rapid scanning, allowing users to quickly identify relevant contacts and information. It also affords associative reminding about important people and relations between people. We discuss the implications of these results for future communication interfaces and for theories of mediated communication.

Categories and Subject Descriptors: H5.2 [**Information Interfaces and Presentation**]: User Interfaces - User-centered design; Theory and methods; Graphical user interfaces, Interaction Styles; K.4.3 [**Organizational Impacts**]: Computer-supported collaborative work H.4.3 [**Communications Applications**] Electronic mail

General Terms: Design, Experimentation, Human Factors

Additional Key Words and Phrases: Human-computer interaction, interpersonal communication, personal information management, personal social networks, visualization, social data mining, social reminding, iterative user-centric design, email, instant messaging.

1. INTRODUCTION

1.1. Motivation

Physical workplaces are often configured so that coworkers are in close physical proximity. Many studies have documented the benefits gained from physical co-location. For example, proximity promotes social interactions and casual social encounters (Allen, 1977, Isaacs et al., 1997, Kraut et al., 1990a, Kraut et al., 1990b, Whittaker et al., 1994). As people move around their workplace, they opportunistically encounter coworkers and this can remind people about conversations they intended to engage in. This lets people discharge commitments they may otherwise have forgotten (Isaacs et al., 1997, Kraut et al., 1990a, Kraut et al., 1990b, Whittaker et al., 1994). Casual encounters also result in unplanned conversations that allow people to keep in touch and maintain social relationships (Bly et al., 1993, Dourish and Bly, 1992, Heath and Luff, 1991, Fish et al.,

Authors' addresses: Steve Whittaker, AT&T Labs-Research, Florham Park, NJ, 07932, USA; Quentin Jones, New Jersey Institute of Technology, Newark, NJ, 07102, USA; Bonnie A. Nardi, Agilent Laboratories Palo Alto, CA 94303, USA; Michael Creech, BlueOak Software, Los Altos, CA, 94024, USA; Loren Terveen, University of Minnesota, Minneapolis, MN, 55455, USA; Ellen Isaacs, Izix.com Consulting, CA, 94002, USA; John Hainsworth, Princeton University, NJ, 08544, USA.

1991, 1993, Tang et al., 1994).

Furthermore, a shared physical work environment also affords social data mining. As shown by research on social networks and recommender systems, people prefer obtaining information and advice directly from other people, rather than official sources or documents. This makes physically co-located colleagues an important information resource (Allen, 1997, Ackerman and McDonald, 1996, Granovetter, 1973, McDonald and Ackerman, 1998, Resnick and Varian, 1997, Wellman, 2001b).

To summarize, a shared physical workplace functions as a *social interface*, providing ready access to important colleagues, which facilitates both communication management and information access. However, the prevalence of remote work and distributed, cross-organizational teams means that many office workers no longer share a physical environment with their colleagues. Instead, they rely on electronic communication environments, typically email and voicemail. While these are very effective communications media (Bälter, 1997, Mackay, 1988, Sproull and Kiesler, 1991, Whittaker et al., 1998, 2000a, 2002a), we report fieldwork showing that these applications do not provide the social reminding and information access that are a natural byproduct of physical proximity.

One reason for this is that email and voicemail applications focus on *message processing* rather than people or social relationships. These applications typically do not show users who their important contacts are, or when they last were in touch. They consequently provide little direct support for managing communication commitments or keeping in touch with important contacts. Further exacerbating the problem, busy people are overwhelmed by the volume of incoming messages, making it more likely that they will lose track of deferred communication tasks involving important contacts. As a result, they often fail to honor commitments they've taken on (Bälter, 1997, Mackay, 1988, Whittaker and Sidner, 1996, Whittaker et al., 1998, 2000a, 2002a).

Likewise, these applications provide little support for social data mining. Our field studies showed that people were constantly frustrated in their attempts to interrogate their email and voicemail archives to find previous contacts with expertise that they now required. One major problem was remembering the identity of prior contacts. Users often tried to use the relations among people as an associative retrieval cue to remind them of the identity of a forgotten contact ('I can't remember the person's name but I *can* remember the other people she worked with'). However social relations are not well represented in current message-centric communication interfaces. This also led to problems in tracking project progress ('there were a lot of messages exchanged about pricing, but what did we actually decide?'). Again, users wanted to use relationship information to monitor the status of email and voicemail conversations. They wanted to track projects, and they thought of projects in terms of the people who were involved ('I know that John, Mary and Jim were involved in the pricing discussion, let me access the messages they exchanged'), but current tools don't make it easy to track conversations in this way.

In short, users of current asynchronous communications applications have little support for the social reminding and social data mining that are a natural byproduct of a face-to-face environment. We attempt to redress this problem by providing a social interface to communications applications that replicates many of the social functions of a shared physical workplace. We created *ContactMap* a novel visual interface to asynchronous communication systems, one that puts people at the center. ContactMap is centered on the notion of *personal social networks*, consisting of those people who are central to a user's work and social life. These contacts are depicted in a structured visual representation that shows important contacts and their relationships. From this representation, users can straightforwardly access their personal communication data and initiate new communications. The goal of this interface is to support social reminding and social data mining, thus allowing users to focus on relationships with important contacts instead of processing incoming messages. Furthermore the interface is intended to support *personal* communication rather than access to public data sources.

1.2 Related Work

In designing a social network UI, we capitalize on foundational research into social networks (Freeman, 1998, Granovetter, 1973, Wasserman and Faust, 1994, Wellman, 2001a, 2000b). This work has developed methods for analyzing communication records to identify important social contacts and represent the relationships among those contacts. According to this approach, people are viewed as interdependent and connected by relational ties. These ties serve as channels for the transfer of information and communication (Wasserman and Faust, 1994). Our goals are somewhat different from classical social network research, however. We are concerned with supporting users in constructing their own personal social networks as *tools* to aid social reminding and social data mining. This contrasts with the goal of classical social network research, which aims to construct veridical representations of social relationships among a set of people, to serve as tools for analysts.

Other research has begun to explore the concept of social interfaces. Much of this work has focused on communications applications such as real time messaging or Usenet access. Instant Messaging (IM) is a highly successful technology supporting real-time text communication (Isaacs et al., 2002a, Nardi et al., 2000a, Milewski and Smith, 2000, Tang et al., 2001, Whittaker et al., 1997). IM systems feature a *buddy list*, i.e. a list of the people with whom the user frequently exchanges instant messages. Users initiate communications with important contacts by clicking on entries in the buddy list. More importantly, the buddy list can serve to remind users about keeping in touch with those contacts (Nardi et al., 2000a, Isaacs et al., 2002a). However, there are several limitations of the IM buddy list. It is used only to initiate communication, not to search for archived communications from important contacts, so it does not support social data mining. (One notable exception here is ICQ, which has a searchable archive function). In addition, buddy lists often are limited to a small number of contacts. Furthermore, while buddy lists usually allow users to organize entries by group, they do not show the detailed *relations* among various contacts that are critical for social data mining.

The Babble system takes a person-centric approach similar to that proposed here (Bradner et al., 1999, Erickson and Kellogg, 2000). It supports synchronous and asynchronous textual communication between small groups, loosely organized into topics or 'rooms', similar to chat rooms or MUDs. The visual focus is a social representation that shows active users in a room, their communication activity, and the relations between their recent interactions. Unlike the buddy list, the Babble interface attempts to extract and represent social relations derived dynamically from prior communications: by visually browsing the interface, a new user can see at a glance which Babble 'rooms' are active, who is currently active or talking, and whom they are talking to. A similar approach of representing dynamic conversational activity is taken in ChatCircles (Viergas and Donath, 1999). However, while the Babble and ChatCircles work shares our goal of providing a structured social interface to communication, they aim to do this for public data. Like email and voicemail UIs, these interfaces show activity for all participants and not just those that are important to the user. However, since social data mining exploits one's personal social network, such public data is less useful for this purpose (Allen, 1997, Granovetter, 1973, McDonald and Ackerman, 1998). ContactMap in contrast offers a system that uses one's *personal social network* as an interface to support social reminding and social data mining.

Other research has also explored social interfaces to support conversational data mining in public data such as Usenet. Smith and Fiore (2001) developed visualizations that highlight conversational structure and social information for Usenet conversations. Social and interpersonal information about past conversations is presented using two representations: a 'piano roll', showing messages for a given thread sorted by poster, ranked by the number of contributions made by that poster; and a sociogram showing reciprocity patterns – such as who replies to the user, and whom the user replies to. Donath et al (1999) designed the Loom interface to Usenet data. This presents the intersection between authors, activity, and time. The interface shows the connection between sequential posts in a thread, allowing users to determine active threads and identify ignored posts. However, both these social interfaces to Usenet are largely intended to support access to all prior conversational history, rather than communications from important contacts, so they do not provide good support for social reminding. Again the focus of these interfaces is on public, rather than personal, views of data which also makes social data mining less effective (Allen, 1997, Granovetter, 1973, McDonald and Ackerman, 1998).

Work on email filtering research also has some similarity to our research. The goal of email filtering is to address communication overload (Bälter, 1997, Jones, 1997, Mackay, 2000, Maes, 1997, Malone et al., 1987, Marx and Schmandt, 1996, Thorngate, 1990, Whittaker and Sidner 1996, Whittaker et al., 1998, 2000a, 2002a) by identifying important incoming messages, thus letting users prioritize their message processing. Important messages are often those from people central to one's social network (Bälter, 1997, Marx and Schmandt, 1996). Email filtering lets users write rules that find messages from important people or about important topics (Balter, 1997, Maes, 1994, Malone et al., 1987, Marx and Schmandt, 1996). The filtering program then makes these more visually salient by reordering messages, color coding them, or providing alternative views on the email inbox. The focus on messages from important contacts is related to our goals; however, this work has more limited objectives than does ours. Its goal is strictly to streamline the processing of incoming messages. In consequence, email filtering provides no support for important social processes such as keeping in touch, managing communication commitments, or social data mining.

1.3 Method / Structure of this Article

Our methodological approach is requirements-driven iterative system design, combined with quantitative and qualitative evaluation of working prototypes (Whittaker et al., 2000b). We proceed by first gathering requirements from interviews and observations aimed at identifying people's problems with current communication systems. Based on those interviews, we identify early design requirements, develop initial system prototypes and critique these by iterative user feedback. We conclude by conducting laboratory based user studies and user interviews to evaluate the utility of our novel system, to determine whether it supports the initial user and design requirements. The structure of this paper mimics that method. Section 2 first summarizes two field studies indicating the importance of social reminding and social data mining and the lack of support that current communications applications provide for these tasks. Section 2 also presents the user and design requirements for interfaces to support social reminding and social data mining, derived from those studies. Section 3 describes our system ContactMap. We outline the basic system features and motivate the design by explaining how it supports user requirements and key social reminding and data mining tasks. Section 4 reports an empirical evaluation of the social network interface. We evaluate how effectively ContactMap supports the social reminding and data mining tasks derived from our interviews, when compared with people's regular email program. We present quantitative and qualitative data showing how and why ContactMap outperforms people's regular email system on these tasks. We also present users' qualitative feedback about our UI design, along with suggestions about how the UI might be redesigned in the light of this feedback. Finally in Section 6, we discuss the implications of our work for theories of asynchronous communication, and consider future issues concerning user interfaces to communication systems.

2. FIELD STUDIES TO IDENTIFY DESIGN REQUIREMENTS FOR CONTACTMAP

The design requirements for ContactMap were derived from two sets of user studies (Nardi et al., 2000b, Whittaker et al., 2002b). We summarize the key findings here. The studies had two related goals: (1) to identify the key problems that users experienced with current communication applications, and (2) to document the main strategies they have evolved to address these problems. The studies analyzed semi-structured interviews and observations of 44 business professionals. We asked them to identify the tools they used to communicate with others, explain how they used those tools, describe the main problems with those tools, and identify the strategies they used to address those problems. We observed and interviewed people using: email, voicemail, IM, fax, phone, and written documents, along with various adjunct 'applications' such as address books, PDAs, 'todo' lists, organization charts and sticky notes. We refer the reader to those earlier studies for the specific questions addressed, methods and detailed results. Here we summarize four main problems users experienced with current asynchronous communication applications.

2.1 Honoring communication commitments.

Users experienced considerable difficulties in honoring outstanding communication commitments; responding to important messages was a recurrent problem. Our users frequently described their interactions as a series of social commitments, that is, communications they owe, or are owed. Consistent with other work on email and voicemail processing (Bälter, 1997, Mackay, 1988, Whittaker and Sidner, 1996, Whittaker et al., 1998, 2000a, 2002a) we found that users often failed to discharge such communication commitments. Although users tried to respond immediately to important contacts, they reported that the sheer volume of messages they received meant that outstanding undischarged messages were sometimes forgotten. In common with earlier research (Bälter, 1997, Mackay, 1998, Whittaker and Sidner, 1996, Whittaker et al., 1998, 2000a, 2002a), we found that users often defer responding to important incoming messages because they do not have the time or the information needed to reply at once. As a workaround strategy, users often left undischarged messages in the inbox as reminders, and then regularly scanned back through the inbox to remind themselves of these commitments. However users who receive large amounts of mail found that these outstanding messages were quickly displaced from view in the inbox by new messages, thus subverting their reminding strategy. High volume email users encountered these reminders only if they happened to scroll back up the inbox. A more successful reminding strategy was to make paper 'todo' lists of outstanding commitments, but many users considered this strategy too time-consuming.

2.2 Keeping in touch

Another critical task users reported was keeping in touch with important contacts. Again, current communications applications did not provide good support for this task. Communication with long-term contacts (whether by email or voicemail) often is sporadic, and people complained that it was very hard to keep these important contacts in mind (Nardi et al., 2000a, 2000b). Users suffering from communication overload spent all their time responding to incoming messages rather than maintaining important contacts. As with commitment tracking, messages from less important people displaced messages from significant people making those messages less visible in the inbox. One strategy that people employed for keeping in touch was to have a paper 'hotlist' of important contacts. They kept this hotlist close to a computer or phone, so that they were reminded about contacting these important people when using these devices.

2.3 Social recommendation

Users also experienced problems trying to access information from their communication archives for social data mining. They repeatedly searched their email and voicemail archives for information about personal contacts who might give them social recommendations or assistance with current projects. They would attempt to answer questions such as 'who was that person I worked with two years ago who knew about network protocols?' or 'who would know about how to file an international patent?'. They regularly searched email, voicemail, and associated address books for such contact information. People commonly relied on the relations between contacts to regenerate contact information. Often they couldn't remember the identity or contact information of the target person, so they would resort to *associative reminding*. They might try to recall other people who worked on the same project, the organization in which the work took place, or other projects they were working on at the same time. From these cues they were often able to track down the target contact.

Neither messaging systems nor address books are designed to represent structured social information – i.e., relationships between people – making them inefficient tools for this type of search. Address books store contacts as alphabetic lists. They do not directly represent information such as projects people worked on or organizations they belonged to. They therefore fail to provide the structured social information needed for associative reminding. Another reason users relied on searching archives was that it was too much effort to maintain an up-to-date address book for all potentially relevant contacts.

2.4 Tracking project status

Finally, users had problems accessing social information from communication archives to track project status. Most of our users participated in collaborative projects that relied heavily on email and voicemail, but they experienced difficulties in monitoring these streams. Some users relied on message subject lines for project tracking, but inconsistent user behaviors can undermine the utility of subject lines. One problem is 'topic drift', which leads to messages about different topics having the same subject line. Alternatively, messages about the same topic can have different subject lines (Erickson and Kellogg, 2000, Herring, 1999, Jones et. al. 2002). One common way users addressed these problems was to retrieve relevant project information by ad hoc social groupings: attempting to remember which set of people were involved in a given project task, then using that social information as a cue to determine which conversations to access ('I know that Julia, Mary and Phil were all involved in the new equipment purchase, so let me access the messages that they exchanged'). However, neither email nor voicemail made it easy to access sets of messages from ad hoc groups of people; it was too unwieldy to set up a system alias for each potential subtask.

2.5 Summary and Design Requirements

One general conclusion of these studies was that not all contacts and communications were treated alike. Users wanted to focus on a small subset of all the people they communicate with. They tried to keep these important contacts in mind and respond quickly to communications from them. They also exploited this subset of important contacts for data mining for work and social purposes. When keeping in touch and honoring communicative commitments, users experienced a common problem with current systems - that communications with important people often suffered from being 'out of sight' and hence 'out of mind'. Current applications focus on processing incoming messages, rather than managing important contacts and communicative commitments. In consequence these applications did not provide good support for social data mining from communication archives. The absence of structured social information meant that people experienced problems in trying to regenerate information about important contacts and in using social information to track project status.

In conclusion, current communication applications did not adequately support the following four communication tasks involving a user's subset of important contacts:

- Honoring communication commitments by tracking and responding to messages from important contacts;
- Keeping in touch with important contacts in order to maintain social relations with them;
- Eliciting social recommendations by exploiting the network of contacts implicit in communication archives;
- Tracking project status using social information within communication archives by accessing messages from relevant clusters of contacts.

This generates a set of design requirements for novel communications systems intended to support these tasks. Users need:

- Tools that help identify and represent contacts that are important to the user;
- Tools that remind users about those important contacts so they can keep in touch and honor communication commitments;
- Tools to support both project tracking and the associative reminding needed for social data mining;
- A structured representation of important contacts; it must show relations between contacts and important information about them to facilitate social data mining.

3. THE CONTACTMAP SYSTEM

Our *ContactMap* system addresses these design requirements. It has the following properties:

- It shows a visual representation of the set of contacts that are important to the user;
- This visual representation is structured to show relations between contacts. Relations are indicated by color coding and spatial proximity;
- Tools remind users about important contacts and outstanding commitments;
- The structured visual representation can be used both to initiate communication and to access communication records, for either a single contact or a set of contacts;
- Tools are provided to help users set up the representation. The tools help users identify important contacts and arrange them spatially on the screen.

Our design was influenced by our field studies and prior system designs. One strong influence was the set of workaround strategies that our subjects developed to address their problems with current email and voicemail systems. In particular, we were struck by the extensive use of hotlists. A hotlist is a list of the contacts with whom a user wants to maintain regular contact. The hotlist is directly associated with communications devices and is placed close to a phone or computer to function as a reminder. We wanted our design to provide the functionality of a hotlist. A related observation concerns the success of buddy lists in IM systems. The buddy list also highlights people who are important to the user. Together these observations suggest the utility of basing interaction around a list of significant contacts. However, as our field studies indicated, representing important contacts as a linear list is insufficient. For social data mining, users need access not only to a list of important contacts, but also to information about the relations between them. Our design presents a *network* representation of important contacts.

We now describe the features of the system in more detail and then explain how these features address the requirements described above.

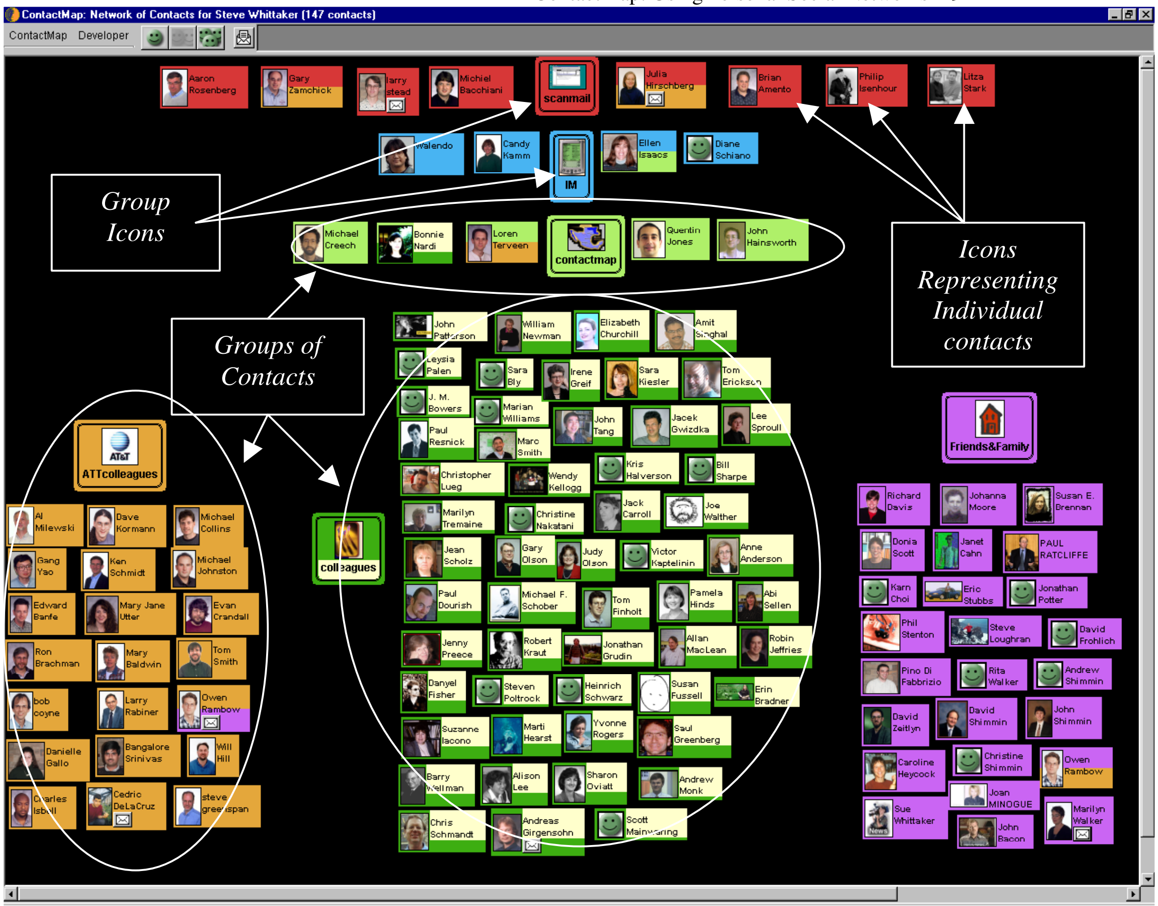

Figure 1 - ContactMap user interface showing 147 contacts in 6 groups

The central element of ContactMap (see Figure 1) is a visual representation of the set of contacts that are critical to a user's work or social life. The visualization models the user's personal social network, showing important contacts and the relations between them. Each contact is represented by a label and a picture, depicted on an icon representing a business card. We used pictures both to accentuate the social nature of the UI, and also because of the ease with which users can scan and recognize familiar faces (Bruce, 1988). Spatial layout and color-coding are used to indicate relations between contacts. Users spatially arrange and categorize contacts to reflect their relationships with each other and to themselves in a 'messy desktop' organization (Barreau and Nardi, 1995, Lansdale, 1988, Malone, 1983, Whittaker and Hirschberg, 2001). Contacts may also be assigned to one or more groups, with the icons for a given group (e.g. scanmail, IM, Friends&Family, colleagues, contactmap, ATTcolleagues) assigned a common color. Groups typically constitute social categories, such as friends and family (Friends&Family), work projects (e.g. scanmail, IM, contactmap), or organizational affiliations (ATT, ATTcolleagues). Group icons may also be given a logo (see ATTcolleagues which has the blue AT&T globe logo). Feedback on an initial prototype indicated that users wanted to have contacts belong to multiple groups. Multiple group membership is shown by striping, so that a single contact can depict the colors of the multiple groups to which that contact belongs.

This structured visual representation lets users see at a glance who is in their social network and perceive relationships between network members. We also experimented with more complex visual representations including network diagrams showing relational links (Freeman, 1998), visual outputs from hierarchical clustering (Backer, 1995), and concentric circles organized around a representation of the user. User feedback about these designs suggested they were less comprehensible than the visual desktop design we present here.

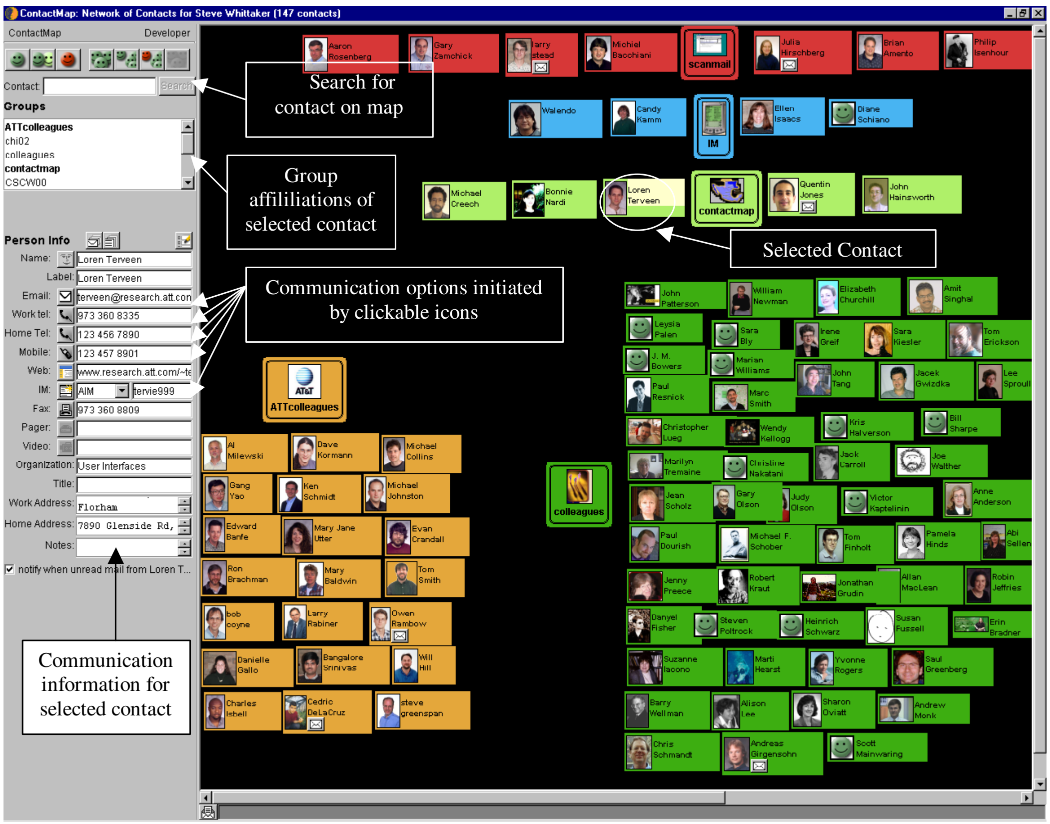

Figure 2 – Contact Information Panel showing communication initiation functions

For each contact, ContactMap provides information for communicating with that person using various modes. Clicking on a contact (e.g. Loren Terveen) causes information for that contact to be shown in the left display panel (see Figure 2). The information includes group affiliations (in this example, the contact has dual affiliations of ATTcolleagues and contactmap), indicated with bold font in the Groups panel. The left panel also contains contact information such as email address, phone numbers, web page, fax, pager. The communication icons (Email, Work Tel., Mobile, IM, FAX) are active; when selected, they invoke communication functions such as addressing an email message or initiating a click-to-dial phone call. Since the system maintains contact details, the user doesn't have to remember these when initiating a communication. Furthermore, these communication functions work not just for individuals, but also groups, including ad hoc sets of contacts defined by the user. For example, selecting a group icon from the map or the panel (e.g. the IM group) highlights the group (as shown in Figure 3). Clicking on the email icon addresses an email to all the relevant contacts. In the same way, users can manually select multiple contacts from anywhere on the map and address an email to all of them. This is analogous to setting up ad hoc email aliases – a process that is highly laborious in most email programs. This selection method also makes it less likely that users will forget relevant contacts, because users scan the map to determine who should be included.

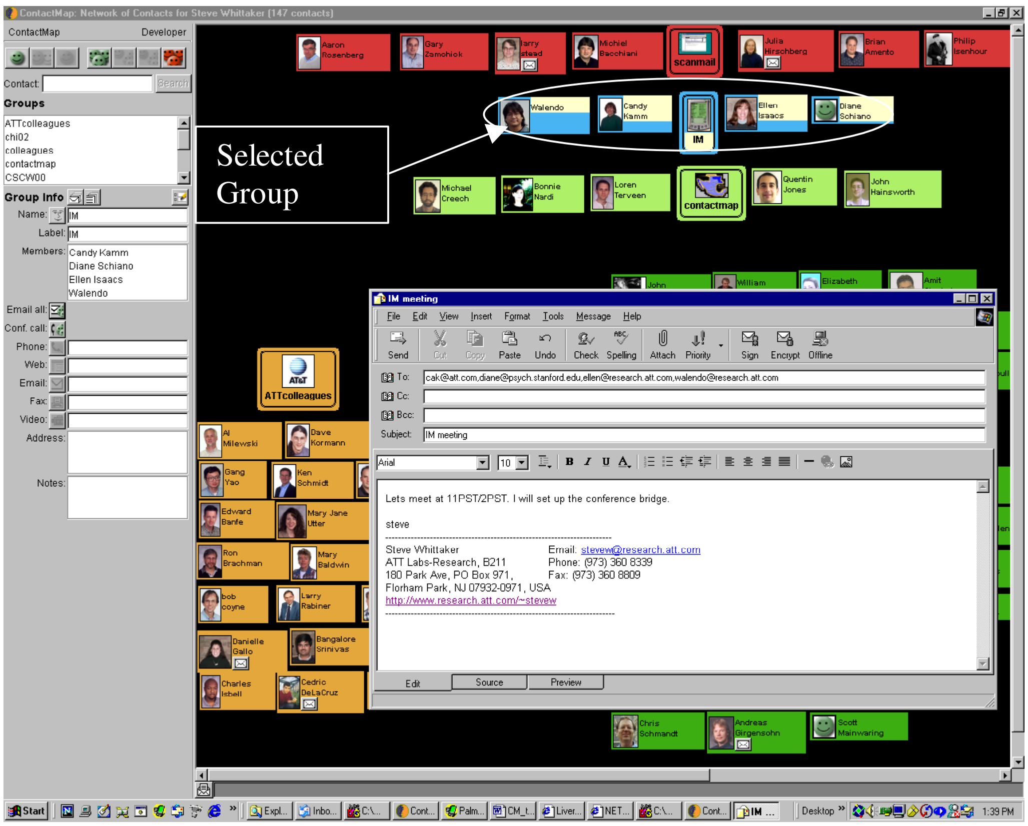

Figure 3 – Sending email to the IM group in ContactMap

ContactMap also enables users to rapidly access communication archives for social data mining. Selecting a contact and then clicking on the Search icon accesses all prior emails exchanged between user and contact(s). The Search functionality also can be applied to defined groups or ad hoc sets of contacts selected by the user. As Figure 4, ContactMap pops up a viewer summarizing emails exchanged by the selected contacts (in this example, those in the highlighted IM group). The viewer allows the user to sort messages by sender; subject; and date. It also presents the body of the selected message, and indicates attachments. Email search operates over all email folders and includes messages sent by the user; this contrasts with current email systems, which segregate incoming and outgoing messages. Note that access to communication history is only currently supported for email. However, we now have access to searchable voicemail corpora, (Hirschberg et al., 2001, Whittaker et al., 2002a) and plan to extend search to transcripts of voicemail messages.

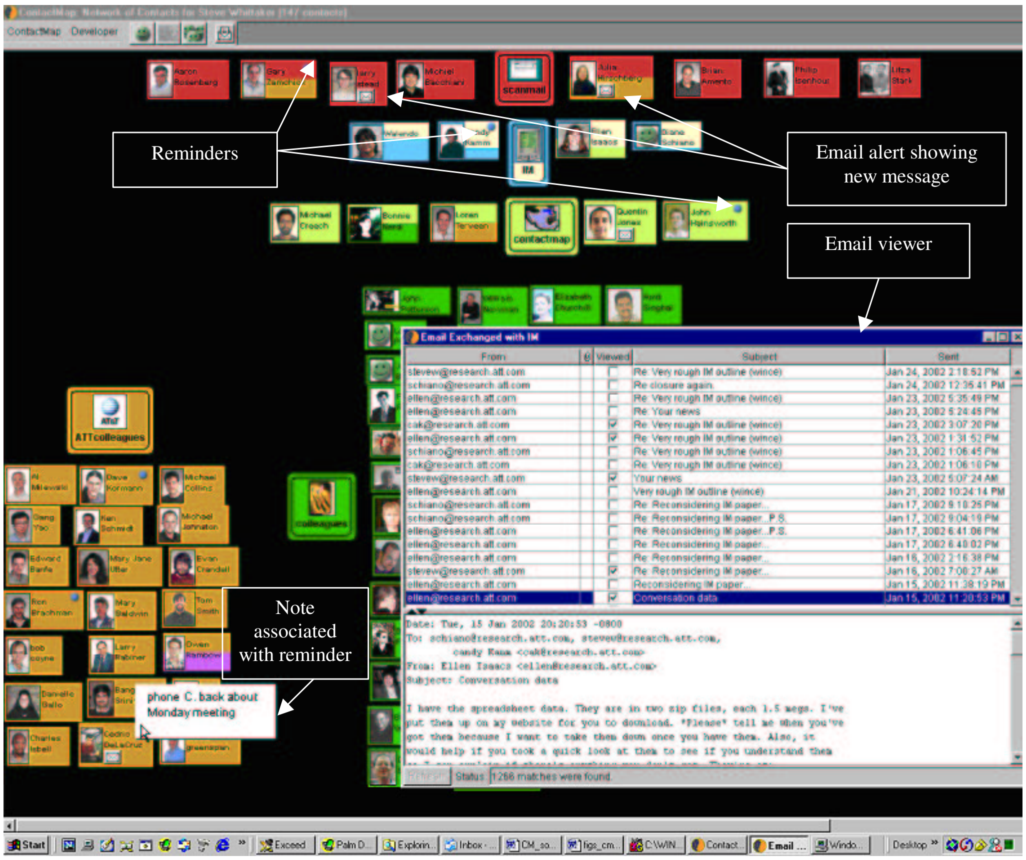

Figure 4 – Accessing emails exchanged with selected contacts, reminders and email alerts in ContactMap

ContactMap also supports social reminding through an alerting mechanism. The presence of an unread incoming message is signaled by a small envelope icon appearing on the relevant contact. In Figure 4, for example, there are unread messages from Cedric DelaCruz of the ATTcolleagues group, Larry Stead and Julia Hirschberg of scanmail, and Quentin Jones of contactmap. Clicking on the envelope icon shows header information for the new message. Users select the contacts and groups for which alerts should be posted. They can also place explicit reminder notes on contacts (signaled by a blue dot in Figure 4) to indicate some outstanding action, such as the need to phone them back. Rolling over the contact with the mouse displays the relevant note. For example, the contact Cedric DelaCruz has an associated reminder to 'phone C. back about Monday meeting'. Leaving reminders on contacts is an explicit way to remember commitments. More subtly, constant visual access to one's contacts may keep important contacts in one's mind and implicitly lead one to recall outstanding commitments to those contacts.

Building the Map

Our field study participants noted that creating and maintaining information about important contacts is laborious (Whittaker et al., 2002b). Feedback about initial versions of our system reinforced this point, with users complaining about the "start-up" cost of identifying contacts and organizing them on the map. We therefore designed tools that help users with these tasks. These tools extract information from email archives. Initially we experimented with fully automatic techniques for extracting important contacts from email and clustering them on the map. We identified important contacts based on their frequency and reciprocity of communication with the user. They were clustered on the map according to co-mentions, i.e. which contacts were mentioned together in message headers. These automatic techniques met with limited success, however. Users pointed out examples of important contacts (such as friends or family) whom they communicated with intermittently in email but whom were nevertheless felt to be central to their personal social network. Similarly, users felt the need to actively reorganize contacts on the map to represent perceived relations between them. For these reasons, we abandoned attempts at automatic contact selection and layout, deciding instead to build tools to support manual selection and layout.

These new tools extract all potential contacts from email. There are often thousands of potential contacts, so users need tools to help select and organize them on the map. We therefore compute various features associated with each contact, and present the results in a table interface (see Figure 5). The table is a tool to help the user select and organize contacts, but in contrast to our initial approach, users determine which contacts get added to the map, and how the map is organized.

User feedback suggested relevant features to include in the table to help with these decisions. These features include a contact's organization (operationalized as the domain from the email address, such as att.com), geographical location (e.g. uk, fr, de), frequency of communication, how long since the last communication, how long since the first communication, and reciprocity (how likely the contact is to respond to one's email and vice versa). These features can be extracted directly from email message headers. Users also stated that email folder structure was an important cue for organizing contacts. For each contact, we computed the most frequent folder in which messages from the contact were filed. Finally, we computed an overall *importance* metric for each contact that combines frequency and reciprocity. This metric has been shown to be useful in initial feedback sessions. The assumption is that these are the two key aspects of what it means for someone to be an important contact. The algorithm for computing importance was based on extensive early user feedback. We presented users orderings based on various candidate metrics, asking them to comment on the orderings and reorder them to better reflect importance. The algorithm was refined on the basis of this feedback. One beneficial effect of this metric is that even if someone sends you a lot of email, as long you never mail them (or do so infrequently), they are not scored highly; this metric thus is an effective spam filter.

The table can be sorted by any of the features, in ascending or descending order, making it easy to identify (for example) the people one communicates with the most, for the longest time, from a particular organization, etc. The initial sort order is based on the importance metric. The 'status' column indicates whether contacts already have been added to the map or whether the user has judged them irrelevant. Contacts that have been "processed" (i.e., added to the map or marked as irrelevant) are moved to the bottom of the sorting table. This lets users focus on the agenda of pending work, i.e., contacts whose fate they have yet to determine. Contacts already added to the map are highlighted in green in the table, and irrelevant contacts are highlighted in red.

Figure 5 – ContactMap sorting table, showing analysis of communication history with different contacts

Users add contacts (and groups of contacts) to the map by dragging them from the sorting table and dropping them onto the map, an interaction paradigm based on that of Amento et al (1999, in press). When a single contact is dropped onto the map, a business card object for that contact is created and populated with all available contact information. When a set of contacts is dropped onto the map, a group icon is created (with icons for the individual contacts also created and arranged around the group icon), and the user is prompted to supply a name for the group. Users are free to reorganize contacts and groups manually on the map. They also can add photos or other images to contact and group icons. Users can also rerun the email analysis tool whenever they wish to update the map and they can also add people or groups manually.

To summarize, ContactMap lets users construct a visual representation of their personal social networks. The representation is structured to show relations between contacts, indicated by spatial organization and color coding. The interface can be used to initiate communication and access communication records. ContactMap helps users keep contacts in mind through explicit alerts and implicit reminding. It also helps users track commitments. Finally ContactMap provides email analysis tools that greatly ease the process of constructing a personal social network.

We now explain how ContactMap meets our design requirements, in particular, how it addresses the 4 social communication tasks identified in our interviews.

3.1 Honoring communication commitments

The key task here is tracking and responding to messages from important contacts. The problem arises because users don't discharge messages immediately. ContactMap's visual representation provides implicit and explicit support for remembering these outstanding communication commitments. Users can place explicit reminders on contacts, and rolling over a contact shows any reminder notes. Normal use of the map leads to implicit reminding, as users opportunistically notice contacts and recall commitments associated with that contact as they use the map for accessing and initiating communications. This implicit prompting is designed to serve the same reminding function as bumping into people in a shared physical environment (Kraut et al., 1990b, Whittaker et al., 1994). For example, to answer a message from a particular contact, one must access the relevant location on the map. This should remind users of other contacts who are spatially (and hence conceptually) close to the accessed contact. As we pointed out before, this contrasts to email and voicemail systems, where there is no guarantee that adjacent messages are conceptually related.

3.2 Keeping in touch

The key task here is to stay aware of important contacts in order to remember to communicate with them periodically. Again, normal use of ContactMap results in users repeatedly seeing important contacts. Like bumping into someone in the hallway at work, traversing the map leads one to encounter images of different contacts, keeping one aware of them. Again, the map's visual representation of relationships between contacts plays a useful role. Contacts are usually spatially clustered on the map according to work projects, organizations, or social groupings. Thus, when accessing a particular contact on the map (e.g., to send a message), a user will notice and be reminded of conceptually related contacts because of their spatial proximity to the target contact. And if a user then wants to initiate communication with one of these contacts, the various communication channels supported by ContactMap make it very easy to do so. As before, this contrasts with email and voicemail, where associative reminding is unlikely since adjacent messages are often unrelated.

3.3 Social recommendations

The task here is exploiting the network of contacts to elicit recommendations for tasks such as information access and labor recruitment. ContactMap supports this process directly. First, a contact is visually represented with pictures of that person's face. It is well known that people have excellent abilities to scan and locate faces (Bruce, 1988), thus making the map an effective way to find appropriate people, even when they can't remember their names. Furthermore, the structure of the map supports associative reminding: even if one can't remember a particular person's identity, one may remember other people who worked on the same project, at the same place, or in the same organization, and these people may help users recall or find the target contact. In email and voicemail, on the other hand, laborious search and browsing is necessary to find forgotten individuals and their messages, and there is no direct support for associative reminding.

3.4 Project tracking

The task here is to use social structure and communication archives to monitor project progress. First, ContactMap users may (and typically do) decide to group contacts in terms of projects. In that case, simply scanning the map to locate a relevant project makes it straightforward to access all communications relating to the project. However, as our interviewees noted, they often needed to access a flow of messages among an ad hoc set of people carrying out a subtask within a project or across formal project lines. The map makes it easy to scan for ad hoc groups of people, select them, and access all messages they sent or received. In contrast, while threading is available in email or voicemail, the subject line is often misleading because of topic drift (Erickson and Kellogg, 2000, Herring, 1999). People also create folders for project tracking, but these have been shown to be ineffective for this purpose (Whittaker and Sidner, 1996).

4. EVALUATING CONTACTMAP

4.1 Method

We conducted a laboratory experiment to evaluate the core functions of ContactMap. We comparing (a) ContactMap with (b) people's regular emailer along with its associated address books, and calendars. We wanted to determine how well they supported the four social communication tasks identified in our interviews. We also knew from our interviews that people typically used their email systems for the types of tasks we intended to examine here. Since users have a great deal of experience with those tools, this was a challenging test for ContactMap. Fifteen subjects participated in the experiment. Twelve used Netscape Communicator and 3 used Microsoft Outlook as their regular email program. Users were researchers, managers, secretaries and marketing staff from a large industrial research laboratory. They were volunteers, and we gave them a nominal reward on completing the study. They carried out all tasks after setting up their own personal social network in ContactMap.

The experiment consisted of 3 phases: (a) map construction, (b) task execution with task specific UI comparisons, and (c) general UI comparisons.

4.1.1 Map Construction

Users took about 45 minutes on average setting up their map. They ran the email analysis program on their email archive to extract potential contacts. They then used the sorting table to explore these potential contacts, adding to the map the ones they judged important. We logged the spatial layout and the structure of all contacts and groups on the map.

4.1.2 Task execution and task specific UI comparisons

The next two phases of the experiment took place a day later. Users first carried out five brief practice tasks to familiarize themselves with ContactMap functionality and the experimental procedure. Study participants had used their current emailer for an average of 3.2 years, and so were highly familiar with its operation. Nevertheless we confirmed that users were familiar with relevant features of their emailer that might be helpful for the communication tasks, e.g. search, view by person, and view by thread.

Users then carried out 8 experimental tasks. All tasks were done in both ContactMap and the user's regular emailer. The within subjects design was used to control for variability in users' email archives and contacts. The order was randomized: half the users carried out a given task first with their emailer, and half started with ContactMap. We logged key strokes, time to solution, and various success measures relevant to each task such as the number of messages accessed or the number of contacts correctly identified. Users were allowed a maximum of 2 minutes for each task. There were two examples of four types of tasks:

- *Communicative Commitment Tracking* - e.g., "You have become ill and have to go into quarantine for the next couple of days, send an email message to relevant people canceling all relevant meetings and social engagements".
- *Keeping in Touch* - e.g., "Congratulations! You have decided to get married! Send an email to all friends to let them know about this happy event".
- *Exploiting One's Personal Network for Social Recommendations* - e.g., "You are looking for a new job. Send an email to as many people as you can who could write you a suitable reference for a new job or organizational role".
- *Project Tracking* - e.g., "You are trying to track the status of activity X[1]: find recent 5 messages sent and 5 messages received about that activity".

These tasks were derived directly from our field studies, and all could clearly be executed with a standard emailer. We expected that ContactMap's structured visual presentation of the user's personal social network supporting social reminding and social data mining would lead to better performance for each task than with a standard emailer. After each task, users made *task specific UI comparisons*. For the task they had just completed, we asked them to express their preference for either ContactMap or their emailer. We then asked them to give a reason for their choice. Again, we expected that ContactMap's better support for social reminding and social data mining would lead users to prefer it for these tasks.

4.1.3 General UI comparisons

After finishing the tasks, users answered 5 general questions comparing the suitability of ContactMap and their emailer for social reminding and social data mining tasks. These questions addressed keeping in touch, honoring communication commitments, social recommendations, project tracking and following up on email. Again these were derived from our initial user interviews.

4.1.4 Qualitative User Feedback about ContactMap Design

We also conducted follow up interviews with six ContactMap users to obtain general reactions to the technology and suggestions for design improvements. We also noted spontaneous user comments made during the experiment that concerned interface design.

4.2 RESULTS

---

[1] We had previously asked users to generate a list of their current projects and activities which we selected from here.

4.2.1 The Structure of User's Maps

Before describing our experimental comparisons, we briefly discuss map construction and the differences between users' maps. More detailed analysis of map structure is described in Nardi et al., 2002. Map structure and complexity were highly variable. The number of map contacts varied between 15 and 184 with a mean of 95.1, and a standard deviation of 61.4. The number of groups varied between 2 and 23 with a mean of 11.1, and a standard deviation of 6.2. The average number of contacts per group (allowing for the fact that contacts could potentially belong to multiple groups) ranged between 3.7 and 14.7 with a mean of 8.5, and a standard deviation of 2.9. Most contacts were included in groups: only 7% of contacts were not part of a group. Most users had some contacts in multiple groups – on average 10% of contacts were in more that one group.

The layout of contacts and clusters was important (Malone, 1983). For example, users often placed frequently needed important contacts and clusters at the top of the map, with less frequently accessed contacts and clusters below them. Several users talked about placing "current projects" near the top of their maps and more archival information lower down. Users also often employed a visual 'seeding' process when using the sorting table to set up the map. They would use the table to rapidly identify a subset of contacts that were placeholders for a larger set of contacts that they later intended to add to the map. They would place these 'seeds' on the map and incrementally add other related contacts to the placeholder set as they worked through the sorting table. We also noted the common types of groups that people constructed, and found some consistency in the types of groups constructed. Most users organized their social network into *workgroups, work projects, friends,* and *family.* People's networks also represented those affiliated with their *special interests* such as a stock club, the PTA or small businesses they were running on the side.

4.2.2 Experimental comparison of ContactMap and users' emailer

For each task, we measured performance using efficiency measures such as time to complete the task, or success measures. The success measures depended on the task, with some tasks requiring users to generate a list of contacts and others requiring them to access a set of messages. We used the measures to compare ContactMap with the user's usual email software, testing the hypothesis that a better interface would let users (a) find more relevant messages or contacts; (b) complete these tasks in a shorter time. We compared performance using two separate analyses of variance (ANOVA) where the independent measures were Interface Type (ContactMap, emailer), Task (honoring communication commitments, keeping in touch, project tracking, social recommendation), Order (whether users carried out the task first using ContactMap and then their emailer or vice versa). The dependent measures in the two analyses were success or time. As Figure 7 shows, for the success measure, people performed better using ContactMap than with their normal emailer ($F_{(1,224)} = 23.52$, $p < 0.0001$), with the ContactMap interface performing better for all tasks on post hoc tests (All $p < 0.05$). There were also differences between tasks ($F_{(3,224)} = 5.99$, $p < 0.001$), with posthoc tests showing that users performed better on keeping in touch and honoring commitments than project tracking, although there were no interactions between interface and task ($F_{(3,224)} = 0.90$, $p > 0.10$). There were also no order effects: users performed no better the second time they carried out a particular task ($F_{(1,224)} = 0.08$, $p > 0.10$).

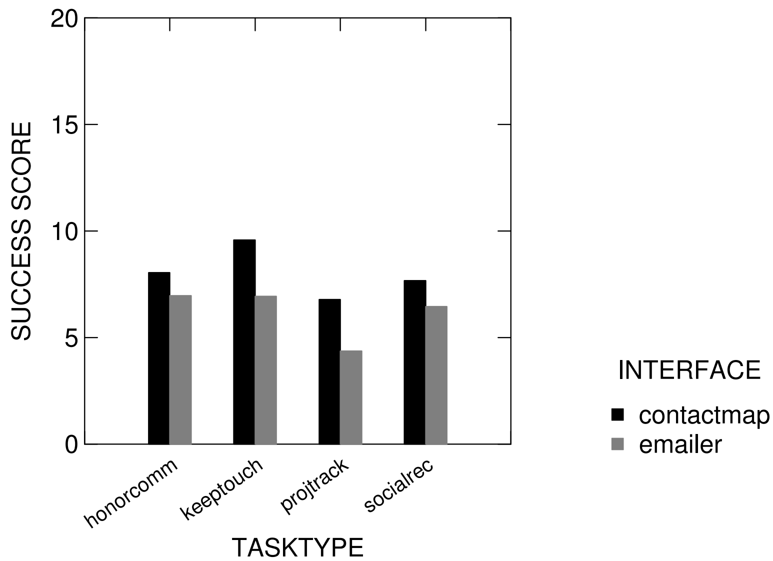

Figure 7 –Success Scores on Four Communication Tasks for ContactMap and email interfaces

The results were similar for the time data. As Figure 8 shows, for the time data, people performed tasks more quickly with ContactMap than their normal emailer ($F_{(1,224)}$ = 11.07, $p < 0.001$), with post hoc tests showing that the ContactMap interface was faster for all tasks, except keeping in touch, where ContactMap and emailer performance were equal. Again, there were differences between tasks ($F_{(3,224)}$ = 3.47, $p < 0.02$), with post hoc tests showing that users performed better on keeping in touch than social recommendation, although there were no interactions between interface and task ($F_{(3,224)}$ = 0.15, $p > 0.10$). Again, there were no order effects: users performed no better the second time that they carried out a particular task ($F_{(1,224)}$ = 1.95, $p > 0.10$). In a separate analysis we looked at whether there were differences between the emailers that people used. There were no differences between Netscape and Outlook emailers for either success ($F_{(1,238)}$ = 1.46, $p > 0.10$), or time $F_{(1,238)}$ = 0.15, $p > 0.10$).

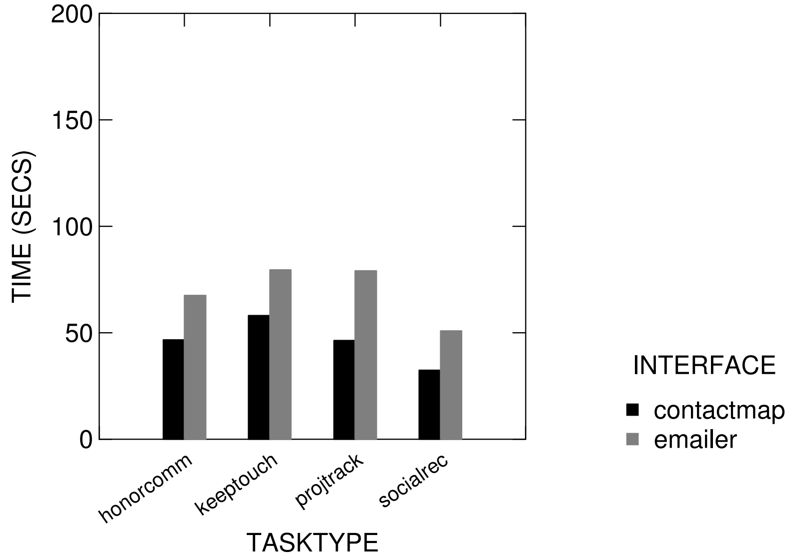

Figure 8 – Task Completion Time on Four Communication Tasks for ContactMap and Email Interfaces

We now turn to the subjective data for users' judgments about which interface they felt was better suited to each task. After each task, users made a single comparison with scores ranging from -2 (strongly preferred emailer) to +2 (strongly preferred ContactMap). The boxplot in Figure 9 shows that all task types had overall positive preference scores favoring ContactMap: the respective means were 1.1 for keeping in touch, 1.0 for outstanding communications tracking, 0.8 for project tracking and 1.0 for social recommendations. (The boxes in the boxplots show the middle half of the data for each task and the whiskers extending from the box reach to the most extreme non-outlier. Outlying points are plotted individually). To test whether there was an overall preference for ContactMap, we carried out one sample t tests for each individual task. On all 4 tasks, users were significantly more likely to rate ContactMap as more suitable than their emailers for carrying out that task ($t_{(29)} = 7.06$, $p < 0.001$, for honoring communication commitments, $t_{(29)} = 8.27$, $p < 0.001$, for keeping in touch, $t_{(29)} = 2.98$, $p < 0.01$, for project tracking and $t_{(29)} = 9.87$, $p < 0.0001$, for social recommendation tasks).

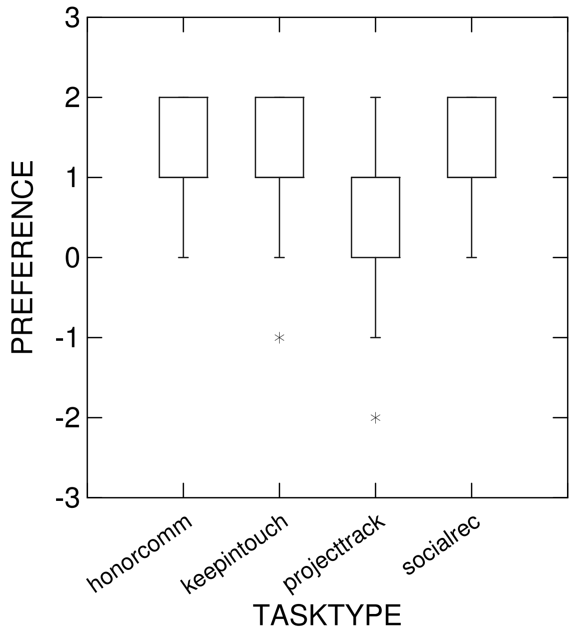

Figure 9 – Subjective preferences for ContactMap or Emailer for Four Communication Task Types

Results were similar on the 5 general UI comparisons that users made at the end of the experiment. One sample t tests again showed ContactMap was preferred for 3 of these 5 questions: keeping in touch (mean difference is 1.27, $t_{(14)} = 6.97$, $p < 0.0001$), social recommendations (mean difference is again 1.27, $t_{(14)} = 6.97$, $p < 0.0001$) and finding people associated with current projects (mean difference is 1.4, $t_{(14)} = 8.57$, $p < 0.0001$). However, ContactMap was rated as equivalent to their emailer for following up email (mean difference is 0.13, $t_{(14)} = 0.49$, $p > 0.10$) and for honoring communication commitments (mean difference is 0.33, $t_{(14)} = 1.05$, $p > 0.10$).

To summarize our findings, the general superiority of ContactMap for social communication tasks for both objective and subjective measures was striking, bearing in mind that people had very brief experience with ContactMap and on average 3.2 years with their emailer.

4.2.4 User comments about interface preferences

We next analyzed the comments users made after each task to explain their interface preference. These comments suggested four main reasons for the greater utility of ContactMap. First, the ContactMap interface supports *visual scanning* that provided rapid access to important contacts. This could be *local scanning,* where users quickly found related contacts by exploiting structural information provided by colors or spatial positions. It could also be *general visual scanning* across the entire map. Scanning was useful for tasks such as keeping in touch, tracking outstanding actions, and social recommendation, where relevant contacts might either be clustered together on the map or spread across different groups or organizations: *"ContactMap helped me to find the relevant people easily - I just looked in the relevant clusters to find them. I also got ideas by just scanning rather than searching for individual people."* Users also commented that seeing people's faces assisted in the scanning process: *"seeing faces really helps you to find people quickly"*. In email, general scanning was hard: relevant contacts were often spread across multiple folders, making it difficult to scan a large set of contacts quickly.

Second, the fact that contacts are constantly visible in ContactMap *reminded* users about important people, obviating the need to *remember* their identity. Even though these names and their aliases were often in the user's email address book or email archives, people felt it was often too laborious to access data from those sources, so people sometimes tried to remember the names of contacts. In contrast, ContactMap users scanned the map to remind themselves of relevant people: *"ContactMap reminds me of who my friends are – in Netscape I have to remember myself."* Reminding was important in tasks such as keeping in touch and social recommendation. In both tasks, relevant contacts may be accessed infrequently – increasing the possibility of forgetting them. As we had intended, ContactMap supported *associative reminding*. Seeing one relevant contact seemed to suggest another: "*it's easy to see a quick overview of relevant people. They're across many different groups but it's easy to pick them out spatially. Seeing a name often generates ideas of other people to pick, and finding them is easy."* Reminding was also useful for honoring communication commitments. Because users are often engaged in multiple communication activities, the set of people associated with their current activities is often ad hoc, increasing the likelihood that contacts associated with those activities may be forgotten – "*with ContactMap I could see all of my contacts at once and select them quickly. With Outlook, I had to scroll through the contact list to make sure I wasn't missing anyone. In the end I missed one person."* Reminding was less useful for frequent and routine tasks such as project tracking. Here, users sometimes already had email aliases and email folders containing sets of relevant contact names, or they were able to directly remember all contacts.

Another critical advantage for ContactMap is that *contacts are structured,* making it possible to efficiently access groups of contacts. "*I can pick everyone in one click. It's easier and faster to click on a few clusters rather than typing and thinking of them all."* This structure also serves a reminding function - accessing one contact may remind the user about another relevant person within the same group and help users access a more complete set of contacts. *"It's easier in ContactMap to see the groups. In email when you need to send to a lot of people it's easy to forget someone."* Overall, ContactMap's structure seemed to be especially useful for project tracking and keeping in touch, as long as users had taken the trouble to construct the relevant groups. While email folders provide some form of structure, they do not allow archival access or sending to specific subgroups of people.

A final advantage for ContactMap was the ability to select and send email to predefined or ad hoc groups of contacts. Several users commented that this support for multiple addressing was equivalent to being able to create impromptu email aliases. This seemed to be an advantage for all tasks, except when contacts were highly familiar and well remembered.

4.2.4 Effects of Map Structure on Performance

We also investigated which aspects of map structure best predicted successful map use. Was the sheer number of contacts on the map critical for success using ContactMap? Or were groups and structure more important for successful usage? We tested several hypotheses about the relationship between map structure and our performance measures of success and time.

We found that having more contacts or more groups on the map did not affect either objective performance as assessed by success scores ($r_{(14)}$=0.30, p>0.05, $r_{(14)}$=0.22, p>0.05) or task completion times ($r_{(14)}$=0.18, p>0.05, $r_{(14)}$=0.27, p>0.05). However, the structure of the map was important, in particular the extent to which contacts were organized into groups. Users with a larger proportion of contacts that were integrated into groups had higher success scores ($r_{(14)}$=0.54, p<0.05), and increased preference for ContactMap over their emailer $r_{(14)}$=0.53, p<0.05). Having more contacts integrated into groups may improve task performance because more structured maps directly support retrieval and associative reminding. Related contacts can be found in the same location, reducing the need for scanning the map, and this may be beneficial for both social reminding and social data mining tasks.

4.2.5 User comments for redesigning ContactMap

Finally, we conducted semi-structured interviews with six users between one and three days after the experiment, asking them more general questions about the interface design, and soliciting their suggestions for improvement. We also noted spontaneous comments about redesign that users made while carrying out the experiment.

One set of user comments addressed the scalability of the interface, particularly for large numbers of contacts. Users noted the visual complexity of the network when executing specific tasks. Several users with large maps proposed mechanisms for “hiding” more peripheral contacts. They suggested constructing task-specific views where relevant parts of the network are highlighted and others hidden when they executed certain tasks. For example, the user might switch to a view showing the IM group when s/he was engaged in work associated with that project. Another common suggestion was that infrequently accessed groups be collapsible, so that only the group icon is generally visible. Clicking on that icon would expand the group to show all members. However, while these suggestions have the merit of reducing visual complexity, they have the disadvantage of reducing the likelihood of impromptu ‘sightings’ of relevant contacts, thus compromising the important visual reminding that follows from seeing a contact.

Users also commented about network change, pointing out that they need to continually update their network as new contacts come into their lives. However, they suggested that such updates be incremental to preserve the spatial relations holding between pre-existing contacts. Consistent with our hypotheses, ContactMap was seen as a *visual workspace*, where spatial position is an important memory cue making radical updates confusing. Currently ContactMap allows the user to run the email analysis tool to find and add new contacts whenever desired. In the future, closer integration with email programs might allow users to incrementally add new contacts, as they receive email from them. People also talked about the importance of phone-based contacts, which were not extracted automatically by the bootstrapping program. Accordingly, we plan to extend the current contact seeding process to include telephone or voicemail logs (Hirschberg et al., 2001, Whittaker et al., 2002a). Automatically extracting information from existing online address books would also be useful.

Users also noted the importance of photographs. These images not only assisted scanning and reminding, but also gave users a strong sense of the people in their network – which was reduced with names and labels alone. Nevertheless, acquiring images is laborious: People collected images by downloading pictures of colleagues from an internal web site, using an image finding search engine, requesting on-line photos from those contacts, or scanning in personal photos. These methods are quite time-consuming (although diverting), and we would like to find an approach that would make picture sharing easier. Finally we noted that most users seemed to enjoy the process of extracting and organizing their contacts. They commented that it was illuminating to see who their main contacts were and to try to impose an organization on them, and this in turn gave them insight into their communication behavior.

While users were generally positive about the ContactMap user interface and its support for social reminding and social data mining, they pointed out there were occasions when they needed access to message-centric information. One scenario they mentioned was when an important communication arrived from an unknown or unexpected person, for example if a colleague's secretary or a new coworker sent email. If the secretary or colleague was not on the user's ContactMap, then there would be no alerting about the arrival of the new message. They also pointed out the general utility of having messages ordered by time, which potentially helps with task management (Whittaker and Sidner, Whittaker et al., 1998, 2000a, 2002). Nevertheless, users argued that person-centric and message-centric approaches are complementary, supporting different types of communication tasks, and that the ideal interface would contain both sets of views allowing users to switch between them, based on their current activity. Finally, people pointed out that much of their networking is done while they are away from their computers in mobile settings. This indicates a need to exploring future versions of ContactMap operating on small devices. This raises the significant visualization challenge of presenting complex visual information on a very small display.

5. DISCUSSION

Following a requirements-driven iterative design approach, we first conducted interviews to investigate current user interfaces to asynchronous communication. These interviews revealed that by organizing communication systems around messages, as opposed to people, current interfaces fail to provide the support for social reminding and social data mining that a shared physical workplace provides. These observations motivated our prototype system, ContactMap, which provides a social rather than a message-centric interface to communication to support these social processes. ContactMap is centered on a structured visual representation of the user's important contacts. This visual representation is intended to support reminding about outstanding communication commitments and allows users to keep significant contacts persistently in mind to support keeping in touch with them. The social representation also allows for associative reminding when accessing communication archives, to help in mining archives for social recommendations or when project tracking. We conducted a laboratory study comparing ContactMap with people's usual emailer for social reminding and social data mining tasks. People performed significantly better with ContactMap on almost all tasks, and they expressed a preference for using ContactMap for the majority of social communication tasks. These results are striking given users' limited experience with

ContactMap compared with their extensive daily use of their emailers. Analysis of user comments suggests that ContactMap's visual social network interface fulfilled our original design objectives in supporting scanning and associative reminding, allowing users to identify relevant contacts and messages more quickly.

Our observations and user comments suggest a number of outstanding system and empirical issues. Future research questions for ContactMap include: How will people use this social interface for everyday work to manage communications with their contacts? If networks increase in size, how can we support network growth to preserve spatial consistency and the *visual workspace* that our users wanted? Two possible solutions to the scaling problem include hiding data associated with more peripheral contacts, or using task-specific views. For example, users often constructed clusters of contacts that corresponded to particular tasks, and these structures might be used to specify task-specific views. This approach is similar to the notion of context-specific spaces, corresponding to 'rooms' (Henderson and Card, 1986) or conversations (Whittaker et al., 1997). A different approach using data hiding might allow more peripheral groups to be collapsible, so that such groups could be represented by their group icon. One limitation of these suggestions, however, is that they restrict the set of contacts that is continuously visible. In doing so, they potentially compromise the opportunistic reminding about contacts that the interface is intended to support. A third possible solution might be fisheye views, or other hyperbolic geometries that would allow peripheral contacts to be visible except smaller (Furnas and Bederson, 1995).

Other research questions involve discovering how these personal social networks change over time. Do networks continue to grow, or do users remove older contacts as they add new ones? Can we automatically recommend contacts to users to include in their networks (McDonald and Ackerman, 1998, Resnick and Varian, 1997, Terveen and Hill, 2001)? Can we enable sharing of contact information so that users can easily share sets of contacts with each other, with all the privacy considerations this entails? Finally, how can the interface to ContactMap be redesigned so that it can be used with small devices such as cell phones and personal digital assistants, as many of our interviewees commented that much of their communication occurs while they are on the move (Whittaker et al., 2002b)?

Another issue concerns integration with existing message-centric user interfaces. While we have shown demonstrable benefits to providing social network interfaces for social reminding and social data mining, our users were clear about the utility of current message-centric interfaces. Future versions of ContactMap will attempt to integrate both approaches into a single UI, so that messages and social networks present different views of the same underlying data, allowing users to choose a representation that best suits their current communication task.

Our findings extend other recent work on social interfaces. Work on IM (Isaacs et al., 2002a, Nardi et al., 2000a, Milewski and Smith, 2000, Tang et al., 2001, Whittaker et al., 1997) and on various proprietary messaging systems (Babble, ChatCircles) has used people as the focus of the user interface. ContactMap extends these ideas by expanding the social interface beyond real-time messaging and using it as a method to access all asynchronous communications. Other related systems have used social interfaces to support social data mining in public conversations such as UseNet (Donath et al., 1999, Smith and Fiore, 2001). ContactMap's approach addresses a slightly different and more general problem than either of these systems, because it focuses on *managing personal communications.* This requires the user interface to support not only social data mining, but also social reminding tasks such as honoring communication commitments and keeping in touch. Support for social reminding is crucial when managing personal communications, but is of less importance when using public data, where there is a reduced need to be responsive or maintain relations. Our work also differs from many of these systems in providing data about the utility of the proposed social interface. With some notable exceptions (Bradner et al., 1999, Isaacs et al., 2002) much of this work has focused on the design of novel social interfaces and paid less attention to the tasks that the interfaces are intended to support or whether they prove useful. A related point is our users' preference for simple representations and for active participation in the construction of the representation. Our early attempts to automatically extract and represent contacts without user involvement were unsuccessful.

This shift from message-centric to people-centric technologies has implications for user's control of their work (Bälter, 1997, Whittaker and Sidner, 1996). The fact that email and voicemail applications are designed around messages and not people makes these applications into reactive technologies. They enable users to process and archive the messages that they receive, but do not highlight who sent them. ContactMap in contrast allows participants to be more proactive. By providing social interfaces that represent people who are important to the user, we allow users to focus on establishing communicative relationships and honoring communication commitments to those people who are important to them.

This work has implications for theories of mediated communication, in particular accounts of asynchronous communication. Elsewhere we have argued that theories of mediated communication are derivative of face-to-face communication theories (Nardi and Whittaker, in press, Whittaker, in press). As a result they have tended to focus on the act of communication itself, i.e. *interaction*, as opposed to the work that is needed to make such communication take place at all, i.e. *outeraction* (Nardi et al: 2000a, 2000b, Nardi and Whittaker, in press). Social reminding and social data mining are both examples of outeraction tasks that current computer mediated communication theories fail to explain. They are both prerequisites for the act of communication: one has to remember to get back in touch with someone, or to remember their contact information before any form of communication can take place. Future empirical and theoretical work needs to better elaborate these outeraction phenomena, and refine theories that better account for them.

A similar focus on interaction as opposed to outeraction may explain the limits of message-centric UIs. Message-centric interfaces focus on the interaction event itself, i.e. composing or replying to a message. Such UIs do not help users with outeraction tasks, such as remembering that they have to reply to a message or helping them find information about the message recipient. Again, refining our notion of outeraction should help with the design of interfaces that better support all aspects of asynchronous communication, including outeraction. Overall, our results argue for the usefulness of personal social networks in organizing communication on the computer desktop. By providing users with straightforward methods to extract and visualize their networks, we hope to provide better support for the outeraction tasks of social reminding or social data mining that are crucial for communication in today's workplaces.

ACKNOWLEDGMENTS

Many thanks to many colleagues at AT&T Labs-Research for discussions about this work, and to the interviewees and experimental subjects for their invaluable assistance.

REFERENCES

ACKERMAN, M. AND MCDONALD, D. 1996 Answer Garden 2: Merging organizational memory with collaborative help, In *Proceedings of Computer Supported Cooperative Work*, 97-105, ACM Press, New York.

ALLEN, T. J. 1977. *Managing the flow of technology*. MIT Press, Cambridge, MA.

AMENTO, B., HILL, W., TERVEEN, L., HIX, D., AND JU, P. 1999. An Empirical Evaluation of User Interfaces for Topic Management of Web Sites, In *Proceedings of Conference on Computer Human Interaction*, 552-559. ACM Press, New York.

AMENTO, B., TERVEEN, L., HILL, W., HIX, D. IN PRESS. Experiments in Social Data Mining: The TopicShop System. To appear in *ACM Transactions on Computer-Human Interaction.*

BARREAU D. AND NARDI, B. 1995. Finding and Reminding: File Organization from the Desktop, *SIGCHI Bulletin*, **27** 3, 39-43.

BACKER, E. 1995. *Computer Assisted Reasoning in Cluster Analysis*. New York: Prentice Hall.

BLY, S. HARRISON, S., AND IRWIN, S. 1993. Media spaces: Bringing people together in a video, audio, and computing environment. Communications of the ACM, 36,:28-47.

OLLE BÄLTER. 1997. Strategies for organising e-mail. In H. THIMBLEBY, B. O'CONAILL, AND P. THOMAS Eds, *Proceedings of Human Computer Interaction*, 21-38, Springer, London, UK.

BRADNER, E., KELLOGG, W, AND ERICKSON, T. 1999. The Adoption and Use of Babble: A Field Study of Chat in the Workplace. The Proceedings of the European Computer Supported Cooperative Work conference.

BRUCE, V. 1988. *Recognising Faces*. Lawrence Erlbaum Associates. Mahwah, N.J.

DONATH, J., KARAHALIOS, K., AND VIEGAS, F. 1999. Visualizing conversations. *In Proceedings of HICSS-32*, Maui, HI.

DOURISH, P. AND BLY, S. 1992. Portholes: Supporting Awareness in a Distributed Work Group. In *Proc. ACM Conference on Human Factors in Computing Systems CHI'92* (Monterey, Calif.). New York: ACM.

ERICKSON, T. AND KELLOGG, W. 2000. Social Translucence: An Approach to Designing Systems that Mesh with Social Processes. *In Transactions on Computer-Human Interaction.* 7, 59-83. New York: ACM Press, 2000.

FISH R.S., KRAUT, R.E. AND CHALFONTE, B.L. 1992. The Video Window System in Informal Communications. In *Proceedings of Computer Supported Cooperative Work*, PAGE NUMBERS, ACM Press, New York.

FISH, R. S., KRAUT, R. E., ROOT, R. W., RICE, R. 1993. Evaluating video as a technology for informal communication. Communications of the ACM, 36, 48-61.

FREEMAN, L. 1998. Computer Programs in Social Network Analysis. *Connections* 11, 26-31.

FURNAS, G. AND BEDERSON, B. 1995. Space-scale diagrams: Understanding multiscale interfaces. In *Proceedings of the ACM Conference on Computer Human Interaction*.

GRANOVETTER, M. 1973 The Strength of Weak Ties. *American Journal of Sociology* 78, 6, 1360-1380.

HAYTHORNTHWAITE, C., WELLMAN, B., AND MANTEI, M. 1995. Work Relationships and Media Use: A Social Network Analysis. *Group Decision and Negotiation,* 4(3), 193-211.

HEATH, C. AND LUFF, P. 1991. Collaborative Activity and Technological Design: Task Coordination in London Underground Control Rooms. In Proc. of ECSCW'91 Kluwer Academic Publishers, Amsterdam.

HENDERSON, A., AND CARD, S. 1986. Rooms: the use of multiple virtual workspaces to reduce space contention in a window-based graphical user interface, *ACM Transactions on Graphics*, 5, 211-243.

HERRING, S., 1999, Interactional Coherence in CMC, *Journal of Computer Mediated Communication, 4*

HIRSCHBERG, J., BACCHIANI, M., HINDLE, D., ISENHOUR, P., ROSENBERG, A., STARK, L., STEAD, L., ZAMCHICK, G., AND WHITTAKER, S. 2001. SCANMail: Browsing and Searching Speech Data by Content, *Proceedings of Eurospeech 2001*, Aalborg.

ISAACS, E., WHITTAKER, S., O'CONAILL, B., AND FROHLICH, D. 1997. Informal communication re-examined: new functions for video in supporting opportunistic encounters. In K. Finn, A. Sellen, S. Wilbur (Eds.), *Video mediated communication*. LEA: NJ.

ISAACS, E., WALENDOWSKI, A., AND RANGANATHAN, D. 2002. Hubbub: A sound-enhanced mobile instant messenger that supports awareness and opportunistic interactions. In the *Proceedings of the Conference Computer-Human Interaction*, ACM Press, New York, 179-186.

JONES, Q (1997). Virtual-Communities, Virtual Settlements & Cyber-Archaeology: A Theoretical Outline, *Journal of Mediated Communication*, 3(3).

JONES Q., RAVID G., AND RAFAELI S. 2002. "An Empirical Exploration of Mass Interaction System Dynamics: Individual Information Overload and Usenet Discourse." In: Proceedings of the 35rd Annual Hawaii International Conference on System Sciences, IEEE, Big Island, Hawaii.

KRAUT, R., C. EGIDO, AND J. GALEGHER. 1990A. Patterns of contact and communication in scientific research collaboration, 149-171 In J. GALEGHER, R. KRAUT, AND C. EGIDO, Eds*., Intellectual Teamwork: Social and Technological Foundations of Cooperative Work*. Hillsdale, N.J.: Lawrence Erlbaum Associates.

KRAUT, R., FISH, R., ROOT, B., AND CHALFONTE, B. 1990B. Informal communication in organizations: Form, function and technology. In S. OSKAMP AND S. SPACAPAN, (Eds.), People's reactions to technology in factories, offices and aerospace, 145-199, Sage Publications

KRAUT, R., AND L. STREETER. 1994. Coordination in large scale software development. *Communications of the ACM*, 1996.

LANSDALE, M. 1988. The psychology of personal information management. In *Applied Ergonomics,* **19,** 55-66.

MCDONALD, D. AND ACKERMAN, M. 1998. Just talk to me: a field study of expertise location. In *Proceedings of Computer Supported Cooperative Work*, 315-324, ACM Press, New York.

MACKAY, W.E. 1988. More than Just a Communication System: Diversity in the Use of Electronic Mail In *Proceedings of Computer Supported Cooperative Work,* ACM Press, New York.

MACKAY, W.E. 2000 Responding to cognitive overload: Co-adaptation between users and technology. *Intellectica,* 30, 177-193.

MAES, P 1994. Agents that Reduce Work and Information Overload. *Communications of the ACM,* 37(7), 31-40, 146, ACM Press, New York.

MALONE, T. 1983. How do people organize their desktops? Implications for the design of office systems. *ACM Transactions on Office Information Systems,* 1, 99-112.

MALONE, T. W., GRANT, K. R., LAI, K.-Y., RAO, R. and ROSENBLITT, D. 1987. Semistructured messages are surprisingly useful for computer-supported coordination. ACM TOOIS. 5, 115-131.

MARX, M. SCHMANDT, C. 1996. CLUES: Dynamic Personalized Message Filtering. *Proceedings of Computer Supported Cooperative Work,* ACM Press, New York, 113-121

MILEWSKI, A. AND SMITH, T. M. 2000 Providing Presence Cues to Telephone Users. *Proceedings of Computer Supported Cooperative Work,* ACM Press, New York.

MILLEN, D. AND HENDERSON, D. 2000. PhoneMan – The Benefits of Personal Call Histories, In *Proceedings of Conference on Computer Human Interaction,* 155-156, ACM Press, New York.

NARDI, B. AND WHITTAKER, S. in press. The role of face-to-face communication in distributed work. In Hinds, P. and Kiesler, S, (Eds.), *Distributed work: New ways of working across distance using technology*. MIT Press.

NARDI, B., WHITTAKER, S., BRADNER, E. 2000A. Interaction and Outeraction: Instant Messaging in

Action. In *Proceedings of Conference on Computer Supported Cooperative Work,* 79-88, ACM Press, New York.

NARDI, B., WHITTAKER, S., AND SCHWARZ, H. 2000b. It's Not What You Know, It's Who You Know: Work in the Information Age. *First Monday*, www.firstmonday.org.

NARDI, B., WHITTAKER, S., ISAACS, E., CREECH, M., JOHNSON, J., AND HAINSWORTH, J. 2002. Integrating Communication and Information Through ContactMap. *Communications of the ACM,***45**, 89-95.

RESNICK, P. AND VARIAN, H. 1997. Recommender Systems, *Communications of the ACM*, 40(3), 106-108.

SACK, W. 2001. Conversation Map: An Interface for Very Large- Scale Conversations, *Journal of Management Information Systems*.

SMITH, M. AND FIORE, A. 2001. Visualization components for persistent conversations. In *Proceedings of Conference on Computer Human Interaction*, 136–143, ACM Press, New York.

SPROULL, L. AND KIESLER, S. 1991. *Connections*. MIT Press, Cambridge, MA.

TANG, J.C., ISSACS, E.A. AND RUA, M. 1994. Supporting Distributed Group with a Montage of Lightweight Interactions. In *Proceedings of Conference on Computer Supported Cooperative Work,* ACM Press, New York

TANG, J., YANKELOVICH, N., BEGOLE, J., VAN KLEEK, LI, F., AND BHALODIA J., 2001, ConNexus to Awarenex: Extending awareness to mobile users, *Proceedings of Conference on Computer Human Interaction*, ACM Press, New York.

TERVEEN, L. AND HILL, W. 2001. Human Computer Collaboration in Recommender Systems. In J. CARROLL Ed. HCI in the New Millennium, Addison Wesley.

THORNGATE, W. 1990. The economy of attention and the development of psychology. *Canadian Psychology*, 21, 62-70.

theBrain: www.thebrain.com.

VIEGAS, F.B. AND DONATH, J.S. 1999. Chat Circles, In *Proceedings of Conference on Computer Human Interaction* 9-16, ACM Press, New York.

WASSERMAN, S. AND FAUST, K. 1994. *Social Network Analysis*. Cambridge University Press, Cambridge.

WELLMAN, B. 2001A. Physical Place and Cyber Place: The Rise of Networked Individualism. *International Journal of Urban and Regional Research*, 25(2), 227-52.

WELLMAN, B. 2001B. Computer Networks as Social Networks. *Science*, 293, 2031-2034.

WHITTAKER, S. TO APPEAR. Computer mediated communication: a review. In Graesser, A. (Ed.) *The Handbook of Discourse Processes*. MIT Press, Cambridge, MA.

WHITTAKER, S., DAVIES, R., HIRSCHBERG, J., AND MULLER, U. 2000A. Jotmail: a voicemail interface that enables you to see what was said. In *Proceedings of Conference on Human Computer Interaction*, 89-96, ACM Press, New York.

WHITTAKER, S., FROHLICH, D., AND DALY-JONES, O. 1994. Informal communication: what is it like and how might we support it? In *Proceedings of Conference on Computer Human Interaction*, ACM Press, New York, 130-137.

WHITTAKER, S., AND HIRSCHBERG, J. 2001. The character, value and management of paper archives. *Transactions on Computer Human Interaction,* **8,** 150-170.

WHITTAKER, S., HIRSCHBERG, J., AMENTO, B., STARK, L., BACCHIANI, M., ISENHOUR, P., STEAD, L., ZAMCHICK G., & ROSENBERG, A. 2002A SCANMail: a voicemail interface that makes speech browsable, readable and searchable. To appear in *Proceedings of Conference on Human Computer Interaction,* ACM Press, New York.

WHITTAKER, S., HIRSCHBERG, J., AND NAKATANI, C. H. 1998. All talk and all action: strategies for managing voicemail messages, In *Proceedings of Conference on Computer Human Interaction*, 1998, ACM Press, New York.

WHITTAKER, S. JONES, Q., AND TERVEEN, L. 2002B. Persistence and Conversation Stream Management: Conversation and Contact Management. In *HICCS'02.*

WHITTAKER, S. AND SIDNER, C. 1996. Email overload: exploring personal information management of email. In Proceedings of Conference on Computer Human Interaction, 276-283, ACM Press, New York.

WHITTAKER, S., SWANSON, G., KUCAN, J., & SIDNER, C., 1997. Telenotes: managing lightweight interactions in the desktop. *Transactions on Computer Human Interaction*, 4, 137-168.

WHITTAKER, S., TERVEEN, L., AND NARDI, B. 2000B. Let's stop pushing the envelope and start addressing it: a reference task agenda for HCI. Human Computer Interaction, 15, 75-106.