ContactMap: Organizing Communication in a Social Desktop

STEVE WHITTAKER
Sheffield University
QUENTIN JONES
New Jersey Institute of Technology
BONNIE NARDI
University of California, Irvine
MIKE CREECH
BlueOak Software
LOREN TERVEEN
University of Minnesota
ELLEN ISAACS
Izix.com
and
JOHN HAINSWORTH
Princeton University

Modern work is a highly social process, offering many cues for people to organize communication and access information. Shared physical workplaces provide natural support for tasks such as (a) *social reminding* about communication commitments and keeping track of collaborators and friends, and (b) *social data mining* of local expertise for advice and information. However, many people now collaborate remotely using tools such as email and voicemail. Our field studies show that these tools do not provide the social cues needed for group work processes. In part, this is because the tools are organized around *messages*, rather than *people*. In response to this problem, we created ContactMap, a system that makes *people* the primary unit of interaction. ContactMap provides a structured social desktop representation of users' important contacts that directly supports social reminding and social data mining. We conducted an empirical evaluation of ContactMap,

Authors' addresses: S. Whittaker, Sheffield University, Sheffield, S1 2TN, UK; email: swhittaker@shef.ac.uk; Q. Jones, New Jersey Institute of Technology, Newark, NJ 07102; email: ggjones@acm.org; B. Nardi, University of California, Irvine, CA 92697; email: nardi@ics.uci.edu; M. Creech, BlueOak Software, Los Altos, CA 94024; email: mike@creeches.com; L. Terveen, University of Minnesota, Minneapolis, MN 55455; email: terveen@cs.umn.edu; E. Isaacs, Izix.com, Consulting, CA 94002; email: ellen@izix.com; J. Hainsworth, Princeton University, Princeton, NJ 08544; email: hains@cs.princeton.edu.
Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from the Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.
© 2004 ACM 1073-0616/04/1200-0445 $5.00

comparing it with traditional email systems, on tasks suggested by our fieldwork. Users performed better with ContactMap and preferred ContactMap for the majority of these tasks. We discuss future enhancements of our system and the implications of these results for future communication interfaces and for theories of mediated communication.

Categories and Subject Descriptors: H.4.3 [**Communications Applications**]: Electronic Mail; H.5.2 [**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

General Terms: Design, Experimentation, Human Factors

Additional Key Words and Phrases: Email, human-computer interaction, instant messaging, interpersonal communication, iterative user-centered design, personal information management, personal social desktop, social data mining, social reminding, visualization

1. INTRODUCTION: SOCIAL WORKSPACES AND TECHNOLOGICAL SUPPORT

Modern work is a highly social process, and social cues are critical both for managing communications and for accessing information. Modern teamwork is communication-centered [Kraut et al. 1990a; Mortensen and Hinds 2002], but workers experience major difficulties in managing their communication commitments [Duchenaut and Bellotti 2001; Mackay 1988; Whittaker and Sidner 1996] and keeping in touch with others [Whittaker et al. 2002a]. Shared physical workplaces reduce these problems. They support *social reminding*, in which casual encounters with co-workers (a) remind individuals about outstanding work tasks and (b) help people keep in touch [Bly et al. 1993; Kraut et al. 1990b; Tang et al. 1994].

Modern teamwork is also information rich, and social processes are critical in obtaining necessary information. A shared physical workplace provides ready support for this process. Often information is not available from official sources but can only be accessed informally from colleagues [Ackerman and McDonald 1996; Granovetter 1973; Kautz et al. 1997; Wellman 2001b]. We call the process of using other people to obtain information *social data mining.* This includes (a) access to expertise, and (b) using social information to track the progress of project tasks.

A shared physical workplace, therefore, functions as a *structured social interface*, configured to provide ready access to co-workers, and reflecting both team affiliations and social relationships [Kraut et al. 1990a]. However, the prevalence of remote work and distributed, cross-organizational teams means that many modern teams no longer share a physical environment. Instead, they increasingly rely on electronic communication environments, typically email and voicemail. While these are effective communication tools [Sproull and Kiesler 1991; Whittaker et al. 1998, 2002a], they do not support the social reminding and social data mining processes that are a natural byproduct of physical proximity.

One reason for this is that email and voicemail are organized around *messages* rather than people or social relationships. These applications do not present a social interface to support *social reminding*: they do not show users

who their important contacts are, or when they last were in touch. Users also cannot distinguish important messages from unimportant ones. As a result, they often fail to honor commitments from important contacts [Whittaker and Sidner 1996; Whittaker et al. 1998, 2000, 2002a]. Likewise, the absence of an explicit social representation means that message-centric applications provide little support for *social data mining*. They do not represent the *relationships* between co-workers that are useful in the associative retrieval processes characteristic of social data mining ("I can't remember the name of the person who knows about X, but I *can* remember the other people she worked with").

More recent interfaces and tools have been developed that address social processes more directly. We now review the effectiveness of these tools.

Foundational research into social networks has developed methods to identify important social contacts and represent the relations ("ties") between them [Freeman 1998; Granovetter 1973; Wasserman and Faust 1994; Wellman 2001a]. The representation of important contacts and the relations between them are clearly important for social reminding and social data mining. However, while social network research aims to construct veridical social representations of complex, diverse networks, it does not directly address our aims here. To support social reminding and social data mining, we must provide an alternative type of social representation that is an end-user tool with explicit mechanisms for personal organization, reminding, and searching.

Other research into social interfaces has focused on communications applications such as real time messaging. Instant Messaging (IM) is centered on a social representation, the *buddy list*. A buddy list consists of people a user regularly communicates with [Isaacs et al. 2002a; Nardi et al. 2000a; Milewski and Smith 2000; Tang et al. 2001]. The buddy list supports social reminding, prompting users to keep in touch with their buddies [Nardi et al. 2000a; Isaacs et al. 2002a], but buddy lists are less effective for supporting social data mining. Although some IM clients such as ICQ and Yahoo IM have searchable archive functions, the buddy list is not used to access these. 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.

Other social interfaces tackle group communication. Babble [Bradner et al. 1999; Erickson and Kellogg 2000] supports synchronous and asynchronous textual communication between small groups with a visualization that shows active users, their communication activity, and relations between their recent interactions. ChatCircles [Viergas and Donath 1999] explored a similar approach to representing dynamic conversational activity. Other research has explored social visualizations to provide access to Usenet conversations [Donath 1995; Donath et al. 1999; Smith and Fiore 2001]. Other recent social network tools such as Orkut and Friendster use a social representation to support making friends through public networks. However the focus of all these interfaces is on *public* rather than *personal* data. But social reminding involves contacts and commitments that are specific to particular users, and social data mining exploits one's personal contacts, making interfaces to public data less useful for these activities.

In summary, existing communications applications do not directly support the social reminding and social data mining that are a natural aspect of a face-to-face team environment. We address this problem by creating *ContactMap*, a novel visual interface that replicates many functions of a shared physical workplace. ContactMap serves as a *personal social desktop* that represents people who are central to a user's work and social life. The representation shows the user's important contacts and the relationships between them. This social representation also allows users to see their outstanding communicative commitments at a glance and to access their personal communication data. In this way, ContactMap's visual representation directly addresses social reminding and social data mining.

Our methodological approach combines requirements-driven iterative design with quantitative and qualitative evaluation of working prototypes. The structure of this article mirrors our method. Section 2 summarizes field studies showing the importance of social reminding and social data mining. It also shows the lack of support current communication applications provide for these tasks, and derives design requirements to support them. Section 3 describes ContactMap, outlining the design rationale and system features and explaining how ContactMap supports social reminding and data mining tasks. Section 4 reports an empirical evaluation of ContactMap that shows subjects perform social reminding and data mining tasks more effectively with ContactMap than with their regular email system. We also present user feedback about our interface design and suggestions for redesign based on this feedback. Section 5 discusses the implications of our work for theories of asynchronous communication and future design issues concerning social interfaces to communication systems.

2. FIELD STUDIES: DESIGN REQUIREMENTS FOR CONTACTMAP

The design requirements for ContactMap were derived from three user studies: Nardi et al. [2000b]; Whittaker et al. [2002a, b]. The studies aimed to identify unmet user needs, specifically to identify tasks that are not well supported by current communication applications. We conducted semi-structured interviews and observations of 44 business professionals. We investigated their use of many communication tools and "adjunct" applications, including: email, voicemail, IM, fax, phone, written documents, address books, PDAs, 'to do' lists, organization charts, and sticky notes. We refer the reader to the earlier studies for details.

Consistent with other research [Duchenaut and Bellotti 2001; Whittaker and Sidner 1996], we found that current communication applications are effective for message processing, that is reading, addressing, and replying to messages. However, they performed less well for tasks that had a strong social component. These technologies do not explicitly represent important contacts. As a result, informants reported a common problem—communications with important contacts often suffered from being "out of sight" and hence "out of mind". The lack of structured social information in these applications also led to problems in trying to access information associated with these important contacts.

We summarize four social communication tasks that users did not execute well because of the lack of explicit representation of social information.

2.1 Social Reminding

2.1.1 *Honoring Communication Commitments.* Informants characterized many work interactions as a series of social commitments, that is, communications they owe or are owed. Consistent with previous work, our informants reported difficulty in managing communication commitments involving important contacts [Duchenaut and Bellotti 2001; Whittaker and Sidner 1996; Whittaker et al. 1998]. Sometimes the sheer number of messages they received kept them from dealing with them immediately. As a workaround strategy, they would leave messages in the inbox as reminders, and then regularly scan the inbox to remind themselves of these commitments. However informants who receive large amounts of mail found that important messages disappeared from view as new messages arrived, subverting their reminding strategy. A more successful reminding strategy was to make paper "to do" lists of outstanding commitments, but many users considered this strategy too time-consuming.

2.1.2 *Keeping in Touch.* Another problematic social reminding task was keeping in touch with important contacts. Informants reported that current tools did not provide good support for this task. Communication with long-term contacts is often sporadic, and people complained that it was hard to keep these important contacts in mind [Nardi et al. 2000a, b; Whittaker et al. 2002a, b]. Users suffering from communication overload spent all their time processing incoming messages, rather than maintaining important contacts, with the result that messages from less important contacts displaced messages from significant contacts. Again, some informants resorted to a paper-based workaround strategy, maintaining a paper "hotlist" of important contacts. Informants kept their hotlists close to a computer or phone, so that they were reminded about these important contacts when they used these devices.

2.2 Social Data Mining

2.2.1 *Social Recommendation.* Informants often accessed their email and voicemail archives and associated address books to search for personal contacts who might give them relevant information. People exploited relations between contacts to locate a desired contact. When they couldn't remember the identity of the target person, they resorted 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. Neither messaging systems nor email address books provide the structured social information needed for associative reminding. Address books store contacts as alphabetic lists, but do not directly represent information about people such as projects worked on, or organizations belonged to. And relations between contacts may be multi-faceted, with one contact simultaneously related to multiple other groups of contacts. Such complex relations are hard to represent using current tools. Informants relied on memory or access to email,

paper or electronic address lists when seeking social recommendations, but this was often a laborious trial–and-error process.

2.2.2 *Tracking Project Status.* Most of our informants participated in collaborative projects that relied heavily on email and voicemail, but reported trouble accessing information from these archives to track project status. Some used message subject lines, but inconsistent user behavior can undermine this strategy. "Topic drift" leads to messages about different topics having the same subject line; conversely, messages about the same topic can have different subject lines [Erickson and Kellogg 2000; Herring 1999; Jones et al. 2002]. A common workaround strategy was to retrieve information by *ad hoc* social groups. Trying to remember which people were involved in a given task, then using that information as a retrieval cue ("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 systems made it easy to access messages in terms of *ad hoc* groups of people.

In sum, the field studies showed the utility of social information in managing communication and accessing information. Informants used a small subset of their contacts as resources for information and recommendations. They sought to keep these important contacts in mind and respond quickly to communications from them, but current technologies do not provide good support for exploiting such social information.

3. THE CONTACTMAP SYSTEM

3.1 Design Goals

The lack of support for social reminding and social data mining generates a set of design goals. Users need tools that:

—Identify and represent important contacts;
—Remind them about important contacts and associated tasks, supporting the social reminding tasks of keeping in touch and honoring communication commitments;
—Provide a structured social interface to communication archives to support the social data mining tasks of social recommendation and project tracking.

Note that our design goal is not to *replace* message-centered communication systems, which are highly effective for reading, replying to, and composing messages. Rather, we explore interfaces that present an alternative, social view of message data, to support important tasks that are not well supported by current systems.

3.2 Design Metaphor: The Social Desktop

Our final design was the conclusion of four phases and 18 months of iterative prototyping involving both paper and software prototypes. At each stage we

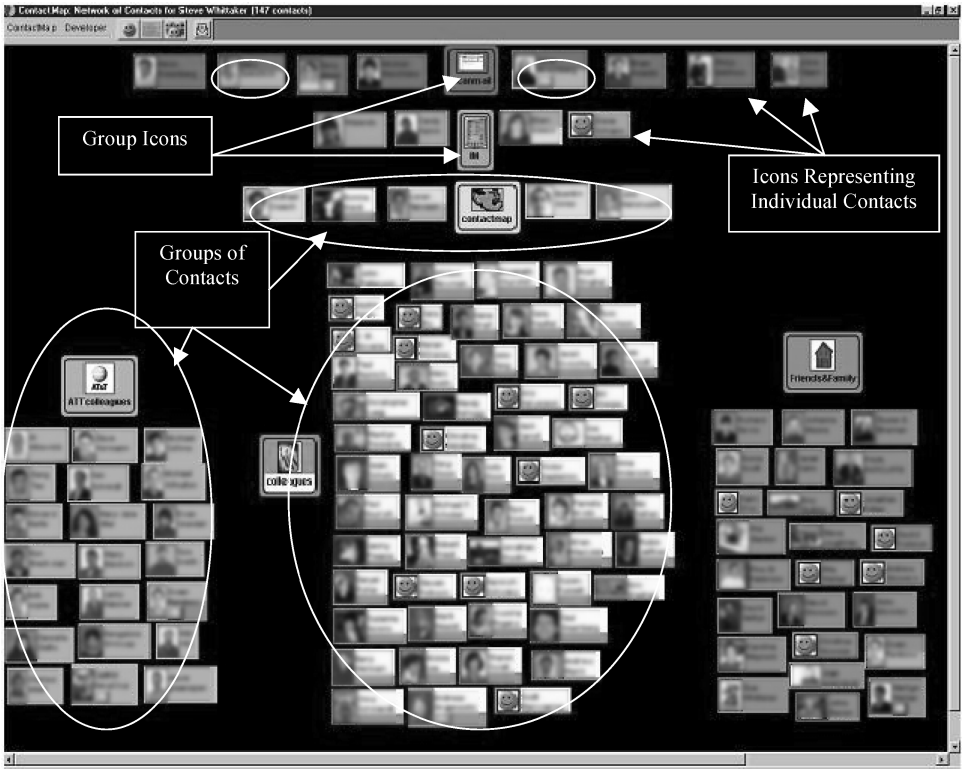

Fig. 1. ContactMap user interface showing social desktop representation of contacts in six groups.

collected and incorporated feedback from representative users. The following sections present the central design metaphor, the *social desktop,* contrasting it with designs of related systems. Next we talk about how the social desktop is integrated with communications data. Finally we describe how ContactMap provides direct support for the social reminding and social data mining tasks we identified in the field studies.

The central design element is the *social desktop* [see Figure 1]. The social desktop is a user-constructed arrangement of important contacts, represented by photographs or icons, using spatial and color cues to show social, affiliate, or project-based relationships. Images and names in all figures are blurred to preserve contact anonymity. The social desktop is analogous to a shared workplace, in affording easy, structured access to work colleagues. A traditional desktop organizes information-related tasks, providing ready access to important data and pending tasks [Malone 1983]. ContactMap aims to play the same role for communication-oriented tasks.

Our field study informants emphasized that social relations are complex and that contacts may be related to multiple groups. ContactMap provides this functionality. The icons in a group [e.g. scanmail, IM, Friends&Family, colleagues, contactmap, ATTcolleagues in Figure 1) are assigned a common color. Membership in

multiple groups is shown by striping, so a single contact can depict the colors of all groups to which the contact belongs [e.g. two contacts circled at the top of the figure are striped in the relevant colors to indicate membership of both scanmail and ATTcolleagues). Groups typically constitute work projects (e.g. scanmail, IM, contactmap), organizational affiliations (ATT, ATTcolleagues), or social categories (Friends&Family). Group icons may also be given a logo to make the group more salient (ATTcolleagues uses the blue AT&T globe).

The social desktop contrasts with related systems in the way in which it structures social information. Email and Instant Messaging are organized around lists of names, such as buddy lists, or names in message headers. Experiences with simple paper prototypes suggested, however, that name lists are ineffective for social reminding and social data mining. Names are not as evocative as photographs, because human faces are a much more effective memory cue [Bruce 1988]. In addition, as a one-dimensional representation, a list is ill-suited to represent the multi-faceted relations between different contacts that are needed in social data mining.

Our field study participants also wanted *personal* representations of social information, noting that creating and maintaining detailed contact information is laborious [Whittaker et al. 2002a, b]. Feedback on early prototypes of our system reinforced this point, with users complaining about the startup cost of identifying and organizing contacts. We therefore designed tools that help users with these tasks, by processing email archives to identify and help organize important contacts on the social desktop.

Initially we tried fully automatic techniques for extracting important contacts and clustering them on the desktop. In one working prototype, importance and organization of contacts was based on frequency and reciprocity of email communication with the user. In other prototypes, we experimented with email-derived social network diagrams showing 2D-representations of relations between contacts, and hierarchical clustering derived from communication patterns [Backer 1995; Wasserman and Faust 1994]. However, consistent with other studies [McDonald 2003], users were not satisfied with these automatic techniques, arguing that they were neither intuitive nor useful for social communication tasks. Rather than ties of greater or lesser strength, users wanted to group contacts based on their affiliation, work project, or social category. Further, they wanted to define their social maps *themselves* rather than having this done automatically. They also identified contacts (such as friends or family) whom they exchanged email with only intermittently, but still considered important. And they wanted to organize contacts on the map to represent *perceived* relations between them. Therefore, we put aside our effort to completely automate the process of contact selection and layout using standard social representations, such as social networks or hierarchical clustering. We turned instead to building tools to guide and support users in constructing their own social representations.

The need for *personally constructed* representations also argues against the use of other automatic methods to construct representations of public social data [Donath et al. 1999; Erickson and Kellogg 2000; Smith and Fiore 2001]. While these systems offer rich and interesting social visualizations, they are

not specific to particular users, and thus are not well suited to supporting social reminding and social data mining.

Two other important design decisions concern *scalability* and *maintenance*. While our photograph-based representation works for up to about 200 contacts, display limitations make it hard to show more. This may not be a practical problem, however, since our early prototyping (and other research) indicates limits on the number of active contacts that people maintain. Although there is considerable individual variability, most people actively communicate with fewer than 150 contacts [Dunbar 1998; Whittaker et al. 2002b]. Over time, people replace older, inactive contacts with new ones, leaving the total number of contacts roughly constant. ContactMap also provides some support for the task of maintaining one's contacts. As a social desktop, it is intended to be always visible as people execute their everyday work tasks. This constant exposure should remind users to add important new contacts and remove older inactive ones, just as people who use folders in their email system periodically archive old folders and create new ones. We also provided users with tools to identify important new contacts and quickly add them to their maps.

Finally, any software, ContactMap included, differs from shared physical workplaces. While shared physical environments provide effective social access to colleagues, they are *public,* and, are therefore configured according to management dictates or collective needs, not to the needs of individual users. In contrast, ContactMap lets users organize their social desktop to represent just those colleagues they consider important. There are other differences between ContactMap and shared physical workplaces: recent studies have shown that physical environments can be *interruptive*, distracting workers from more important activities [Hudson et al. 2002; Rouncefield et al. 1994]. In contrast, ContactMap manages the problem of availability by providing personalized alerting that can be configured depending on how busy users are, or who they need to pay attention to. So while ContactMap offers some (but obviously not all) of the advantages of shared physical workplaces, it also removes some of the disadvantages.

3.3 Integrating the Social Desktop with Communication Information

In addition to providing a structured social representation, the social desktop needs to provide tools that directly support social reminding and social data mining tasks. Users not only need to see their important contacts, but to be reminded of the communication commitments and access information associated with those contacts. We enhanced the social desktop with implicit and explicit reminder tools, integrating it with email data to allow structured access to communication archives. Emails could also be sent and received through the social desktop. By adding this social representation to email, we directly supported both social reminding and social data mining.

Social Reminding, Implicit Reminding, Alerts and Notes. Like a physical desktop, ContactMap supports reminding. Seeing important contacts while engaged in other tasks can remind users of outstanding communication commitments. ContactMap also provides explicit support for social reminding through

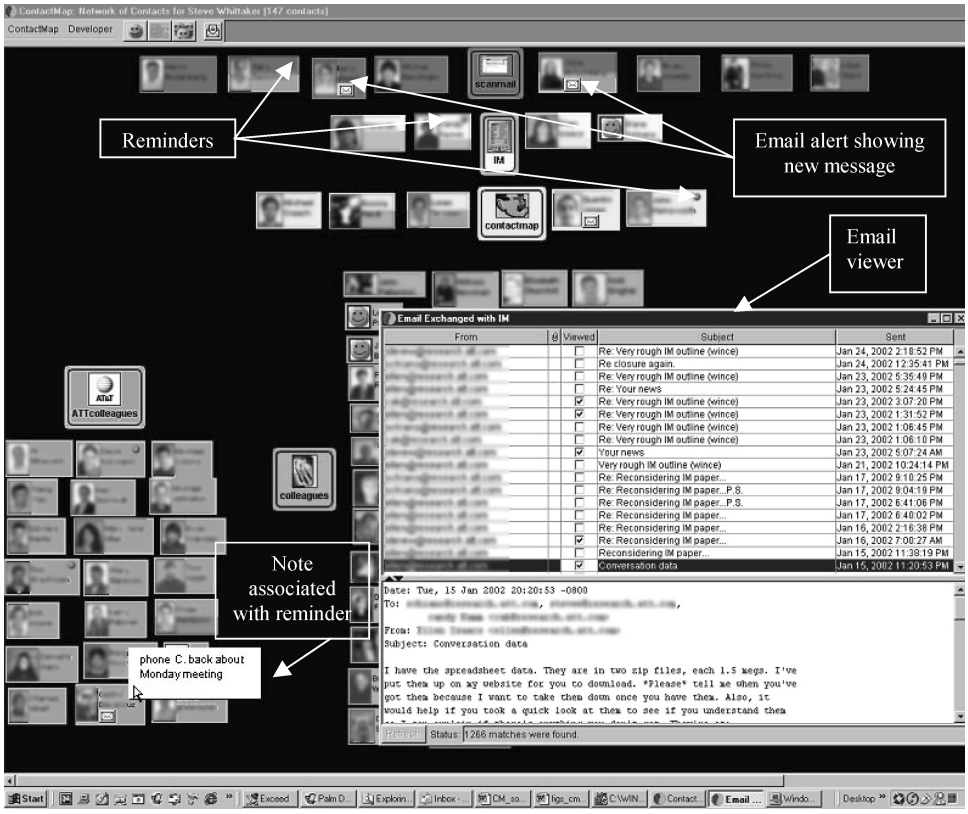

Fig. 2. Accessing emails exchanged with selected contacts, reminders and email alerts in ContactMap.

alerting mechanisms and reminder notes. To provide this support, ContactMap is integrated with the user's email system: new email messages from important contacts are signaled by small envelope icons on the relevant contact. In Figure 2, for example, there are unread messages in contacts in both the ATTcolleagues, scanmail, and contactmap groups. Clicking on an envelope icon shows header information for the new message. Users define contacts and groups for which alerts should be posted, thus focusing on critical contacts. They can also place notes on contacts (signaled by blue dots depicted in the upper right of the relevant contact icons in Figure 2) to indicate outstanding actions, such as the need to call someone. Rolling over the contact with the mouse displays the relevant note. For example, one contact in the lower part of the ATTcolleagues group has an associated reminder to "phone C. back about Monday meeting".

Social Data Mining: Structured Access to Social Data. ContactMap also enables socially structured access to communication archives to support social data mining. A search function retrieves all email exchanged with a specified contact, defined group, or *ad hoc* set of contacts. ContactMap displays a viewer summarizing email exchanged with the selected contacts (the example in Figure 2 shows email exchanged with the IM group). Retrieved messages can

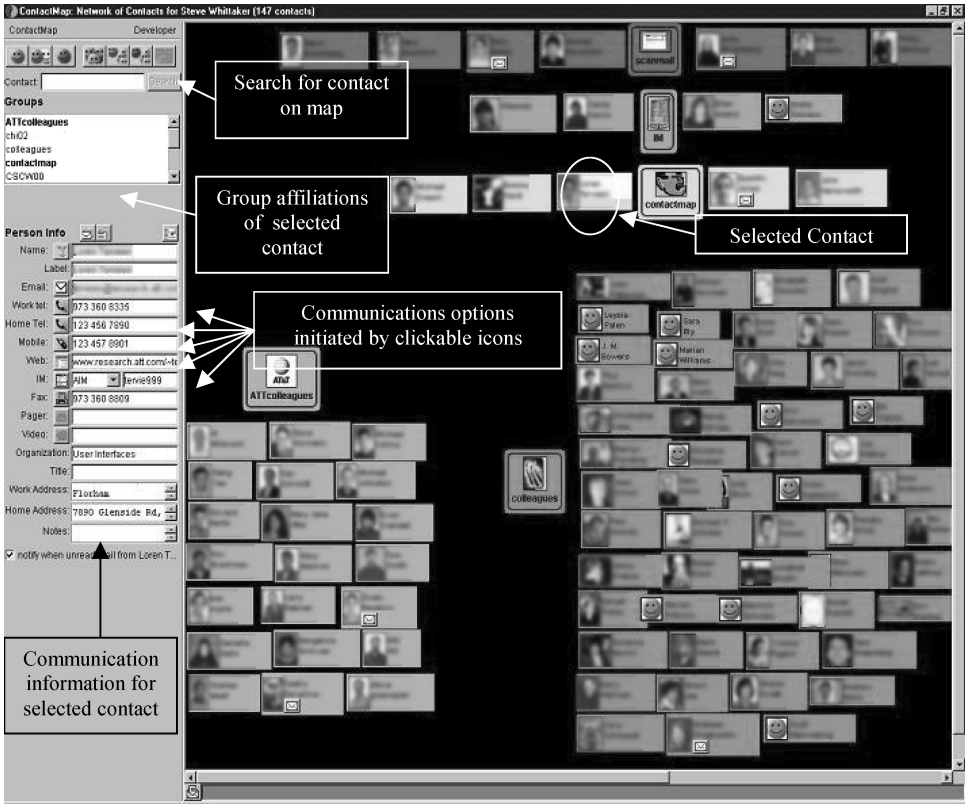

Fig. 3. Contact information panel showing communication initiation functions.

be sorted by sender, subject, or date. Message body and attachments are also displayed. The default search operates over all email folders and includes messages sent by the user. This contrasts with most current email systems that segregate incoming and outgoing messages, requiring the user to explicitly stipulate exactly the folders they want searched. Seeing one's own messages with related messages from others also gives a more complete picture of the communication record, and facilitates project and commitment tracking.

Addressing Emails or Other Communications. ContactMap can also be used to initiate communication with contacts using various modes. Clicking on a contact displays information for that contact in the left display panel (see Figure 3). This information includes group affiliations (in this example, the contact has dual affiliations of ATTcolleagues and contactmap), as well as contact information such as email address, phone numbers, web page, fax, pager. The communication icons (Email, Work Tel., Mobile, IM, FAX) are active; selecting them invokes the appropriate communication functions, for example, addressing an email message or initiating a click-to-dial phone call. Furthermore, these functions work not just for individuals, but also groups, including *ad hoc* sets of contacts defined by the user. For example, one could select a group icon (e.g. the IM group), then address an email to all contacts in the group with a single

Fig. 4. ContactMap sorting table, showing analysis of communication history with different contacts.

click, rather like a predefined email alias. Likewise, one could select contacts one-by-one, from anywhere on the map, and address an email to the whole set. This is a very easy way to create an *ad hoc* email alias, a process that is laborious in most email programs. This selection method also reduces the chance of forgetting relevant contacts, since users scan the map to determine who should be included in a specific message.

Tools for Building the Personal Social Desktop. As we noted above, users reacted to early prototypes by requesting that we provide tools to help construct and maintain their social desktop. We did this by analyzing user email archives to identify *potential contacts*, that is, anyone a user had sent a message to, or received a message from. Active email users have thousands of potential contacts, so additional information is needed to decide which ones to select and how to organize them on their maps. The system computes various features for each contact and presents the results in an interactive table interface (see Figure 4). The table was designed after extensive prototyping [Whittaker et al. 2002b]. Features for each contact include user name, messages sent and received, folder where messages are most frequently filed, and an overall measure of contact importance.

The table-sorting mechanisms are based on Amento et al. [2003]. The table can be sorted by any of the features, making it easy to identify people one communicates with most, for the longest time, from a particular organization, and so on. The "status" column indicates whether contacts have been added to the map or judged irrelevant. Users add contacts (and groups of contacts) to the map by dragging them from the sorting table and dropping them onto the map, and the system automatically constructs labeled icons. Users can spatially reorganize contacts and groups as they desire. They also can add photos or other images to icons. Users can rerun the email analysis tool at any time to update the map and can also add contacts or groups manually.

To summarize, ContactMap lets users construct a visual representation of their social desktop. The representation is structured to show relations between contacts, which are indicated by spatial organization and color coding. ContactMap helps users keep contacts in mind through explicit alerts and implicit visual reminding. It also helps users track commitments and access communication archives using social relationships and associative reminding. Finally ContactMap provides email analysis tools that ease the process of constructing and maintaining a personal social desktop.

3.4 Supporting Social Reminding and Social Data Mining

We now explain how ContactMap meets our design requirements, in particular, how it addresses the 4 social communication tasks identified in our interviews. Along the way, we contrast ContactMap features with those provided by email and voicemail, which are still the prevalent communication tools.

3.4.1 *Honoring Communication Commitments.* Since users cannot deal with all messages immediately, they face the task of tracking and responding to messages with implicit obligations. ContactMap provides implicit and explicit support for remembering such commitments. Users can place *explicit* reminders on contacts; moving the mouse over a contact then shows any reminders. Normal use of the map for accessing and initiating communications also leads to *implicit* reminding. When accessing a particular contact on the map, a user is likely to notice and be reminded of related contacts because of their spatial proximity to, or color match with, the target contact. This implicit prompting is modeled after the reminding function played by bumping into people in a shared physical environment [Kraut et al. 1990b; Whittaker et al. 1994]. As mentioned previously, this contrasts with email and voicemail systems, where there is no guarantee that adjacent messages are conceptually related.

3.4.2 *Keeping In Touch.* People want to stay aware of important contacts in order to remember to communicate with them periodically. Again, normal use of ContactMap results in repeatedly encountering important contacts. Like bumping into someone at work, traversing the map leads one to see images of different contacts, thus maintaining awareness of them. The map's visual representation of relationships between contacts again plays a useful role. Contacts usually are arranged or color coded on the map according to work projects, organizations, or social groupings, so that accessing one contact prompts awareness of another

related contact. Again this contrasts with email and voicemail, where associative reminding is unlikely since adjacent messages are usually unrelated.

3.4.3 *Social Recommendations.* People use their set of contacts to elicit recommendations and advice. ContactMap supports this process directly. First, a contact is represented with a picture of a person's face. People are very good at scanning and locating faces, making the map an effective way to find appropriate people, even when one 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 be able to exploit ContactMap's organization to 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 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.

3.4.4 *Project Tracking.* People also use social structure and communication records to monitor project progress. ContactMap supports this activity, too. Users typically group contacts in terms of projects. When this is done, scanning the map to locate the relevant project makes it easy to access all communications relating to the project. However, our informants noted that they often needed to access 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 and access all messages they sent or received. In contrast, while threading is available in email, topic drift means it is an unreliable technique for tracking tasks. People also create folders for project tracking, but these have been shown to be ineffective for this purpose because project information is often inconsistently filed [Whittaker and Sidner 1996].

4. EVALUATING CONTACTMAP

We conducted a laboratory experiment combined with semi-structured interviews to evaluate the core functions of ContactMap. In the laboratory study, we compared ContactMap with subjects' regular email programs, associated address books and calendars. Our goal was to determine how well each system supported the four social communication tasks identified in our interviews. We compared ContactMap with email, both because of the prevalence of email use [Duchenaut and Bellotti 2001; Whittaker and Sidner 1996], and because our interviews showed that people typically used their email systems for these tasks. This was a challenging test for ContactMap, since subjects had a lot of experience using email for social reminding and social data mining tasks. Note, however, that the study was not intended to compare ContactMap and email for *all* tasks that email currently is used for, such as responding to and composing messages. Our objective with ContactMap was not to replace email, but rather to supplement it with explicit support for social communication tasks that users find important, and that email does not support effectively.

We used a laboratory study for several reasons. First, it provides control over the tasks carried out with ContactMap. Second, we did not want to conduct a

long-term field trial with ContactMap, as this would have involved developing and supporting a fully featured, robust email client. Nevertheless, a laboratory study makes it hard to investigate extended aspects of map usage, such as map evolution or the utility for reminding of infrequently accessed contacts. Therefore, our study looks at a snapshot of map use at a specific time. We used follow-up, semi-structured interviews carried out a few days after the study to probe these other issues.

4.1 Method

Fifteen subjects participated in the experiment. Twelve used Netscape Communicator and three used Microsoft Outlook as their regular email program. Subjects were researchers, managers, secretaries, and marketing staff at a large industrial research laboratory. Subjects were volunteers, and received a nominal reward when they completed the study. Using experimental subjects from the same organization as the research team raises the potential for bias. However, we minimized this risk; none of the subjects were involved with the project, and all were blind to the hypotheses of the experiment and the goals of the research.

The experiment consisted of two phases. Subjects first constructed their maps by running the email analysis program, selecting important contacts, and organizing them on the map. Then, a day later, they carried out the experimental tasks including task-specific and general comparisons between ContactMap an their emailer.

4.1.1 *Map Construction.* Subjects averaged about 45 minutes to set up their maps. They ran the analysis tool on their email archive to extract potential contacts. They then used the sorting table to identify important contacts and add them to the map. We logged the spatial position and group structure of all contacts on the map.

4.1.2 *Task Execution and Task-Specific UI Comparisons.* Subjects first did five brief practice tasks to learn about ContactMap functionality and the experimental procedure. Subjects were very familiar with their email tools, having used them an average of 3.2 years. Nevertheless, we still had them carry out the same practice tasks with their email program to confirm that they were aware of features that might be helpful for the experimental tasks. We did not allow people to continue until they had correctly completed the tasks.

Subjects then carried out eight experimental tasks twice, once with ContactMap and once with their regular emailer. The within-subjects design controlled for variability in subjects' email archives and contacts. The order was randomized: half the subjects did a given task first with their emailer, and half with ContactMap. Subjects were given a maximum of 2 minutes for each task. There were two examples of each of the main social communication tasks identified in the field studies. Tasks were completed by either addressing an email to a set of people or identifying a set of messages. We had subjects send group emails as this was a convenient way to log how easy subjects found it to identify different sets of contacts. We logged key strokes, time to solution, and

task-specific success measures such as the number of messages accessed or the number of contacts correctly identified. Recall that the evaluation did not aim to compare ContactMap and email for *all* tasks that email is currently used for. Rather, we focused on tasks where our field studies revealed unmet user needs.

—*Honoring Communication Commitments*. This social reminding task tests subjects' awareness of their outstanding tasks and activities. Subjects first had to identify all their outstanding communication tasks and send a message to all relevant individuals involved with these tasks to postpone them. "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". Identifying people who are involved with outstanding tasks and then communicating with them, is exactly consistent with our field study informants' view of honoring communication commitments [Whittaker et al. 2002a, 2002b].

—*Keeping in Touch*. Keeping in touch is another example of social reminding. In everyday life, it requires people to remember others they haven't been in contact with for some time. However, we had no direct access to subjects' prior communication history, making it hard to determine who they had not been in touch with recently. Instead, we asked subjects to identify people such as friends whom our field studies showed were contacted on a sporadic basis, and whom users often needed to be reminded about [Whittaker et al. 2002a, 2002b]. We tested the extent to which each system allowed subjects to access these infrequent contacts: "Congratulations! You have decided to get married! Send an email to all friends to let them know about this happy event". We asked people to send a message to their friends as this provided a convenient way to log their contact choices.

—*Exploiting One's Contacts for Social Recommendations*. Social recommendation is a form of social data mining where one exploits contacts for their advice or expertise. This task required subjects to identify people who knew them well to write a job recommendation. Our field study suggested this was a typical social data mining task leading people to browse their communication archives to access information about prior contacts and conversations: "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". Again we asked people to send a message as this enabled us to log the contacts they had chosen.

—*Project Tracking*. Project tracking may involve social data mining, where people use information about *ad hoc* groups of colleagues to access information about ongoing projects. In this task, we required subjects to identify messages associated with a particular project activity, a typical social data mining task: "You are trying to track the status of activity X[1]: find five recent messages sent and five messages received about that activity".

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

After completing each task, subjects made *task-specific* UI comparisons: they were asked to express a preference for either ContactMap or their regular email system, and then explain their choice. Since ContactMap provides direct support for social reminding and social data mining tasks, we expected subjects both to perform better on the objective success measures with ContactMap and to prefer it to their regular email software. After finishing the tasks, subjects answered five *general* questions comparing contactmap and their email software for the four types of tasks in the experiment.

Finally, we conducted follow up interviews with six people who used ContactMap for three days after the experiment to obtain reactions to using the system and suggestions for design improvements. We also noted spontaneous comments subjects made during the experiment relating to interface design.

4.2 Results

4.2.1 *The Structure of Subject's Maps.* We first briefly discuss map construction and variation among subject's maps (see Nardi et al. [2002] for more detailed analysis). Map structure and complexity were highly variable. The number of 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 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 subjects had some contacts in multiple groups—on average 10% of contacts were in more that one group.

Just as with physical desktops, the spatial arrangement of contacts and groups was significant [Malone 1983]. For example, nine subjects placed important, frequently-used contacts and groups at the top of the map, with less frequently accessed contacts and groups below. Seven spoke of placing current projects near the top and more archival information lower down. Seven subjects also employed a visual "seeding" process when setting up their maps. They would use the sorting table to rapidly identify a subset of contacts that were placeholders for a larger set of contacts that they intended to add to the map. They would place these "seeds" on the map and incrementally add other contacts to the growing group as they worked down the table. We also noted some common patterns across subjects' groups. Eight organized their social desktop into *workgroups, work projects, friends,* and *family,* with thirteen having at least two of these categories. In some cases, workgroups and projects corresponded to listservs or email aliases, but no user incorporated listservs as contacts on the map. We examine the effects of map structure on performance in Section 4.2.4.

4.2.2 *Experimental Comparison of ContactMap and Subjects' Emailer.* We measured task performance with objective success and efficiency (i.e., task completion time) measures. For some tasks, the success measure was a set of retrieved contacts, for others, it was a set of retrieved messages. We used the measures to compare ContactMap and subjects' email software, testing

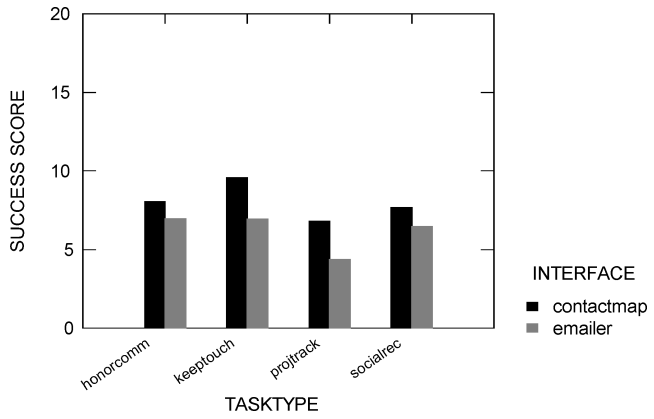

Fig. 5. Success scores for ContactMap and email interfaces.

the hypothesis that a better interface would let subjects (a) find more relevant messages or contacts, and (b) complete these tasks more quickly. We compared performance using two separate Analyses of Variance (ANOVAs), where the independent measures were **Interface Type** (ContactMap, emailer), **Task** (honoring communication commitments, keeping in touch, project tracking, social recommendation), and **Order** (whether subjects used ContactMap or their emailer first). The dependent measures in the two analyses were success and time. Figure 5 shows that subjects performed better with ContactMap than with their normal emailer ($F_{(1,224)} = 23.52$, $p < 0.0001$), with ContactMap better for all tasks on post-hoc tests (Tukey tests, all $p < 0.05$). There were also differences between tasks ($F_{(3,224)} = 5.99$, $p < 0.001$), with post-hoc tests showing that subjects 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 also were no order effects: subjects performed no better the second time they carried out a particular task ($F_{(1,224)} = 0.08$, $p > 0.10$).

Results were similar for completion time. Figure 6 shows that subjects completed tasks more quickly with ContactMap than their normal emailer ($F_{(1,224)} = 11.07$, $p < 0.001$). Post-hoc tests showed that ContactMap was faster for all tasks except keeping in touch, where this difference approached significance (Tukey tests, $0.10 > p > 0.05$). Again, there were differences between tasks ($F_{(3,224)} = 3.47$, $p < 0.02$), with post-hoc tests showing that subjects performed better overall 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: subjects were no faster the second time they carried out a given task ($F_{(1,224)} = 1.95$, $p > 0.10$). In a separate analysis, we looked for differences between the email tools that subjects used. There were no differences between Netscape and Outlook for either success ($F_{(1,238)} = 1.46$, $p > 0.10$), or time $F_{(1,238)} = 0.15$, $p > 0.10$).

Subjective preferences were consistent with the objective results. After each task, subjects gave a preference score from −2 (strongly preferred emailer)

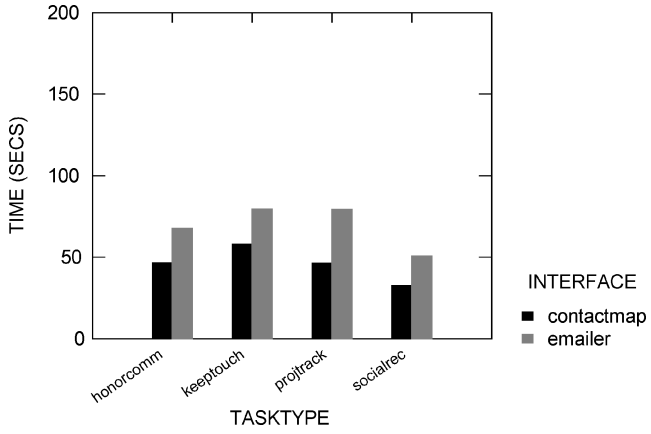

Fig. 6. Task completion time for ContactMap and email interfaces.

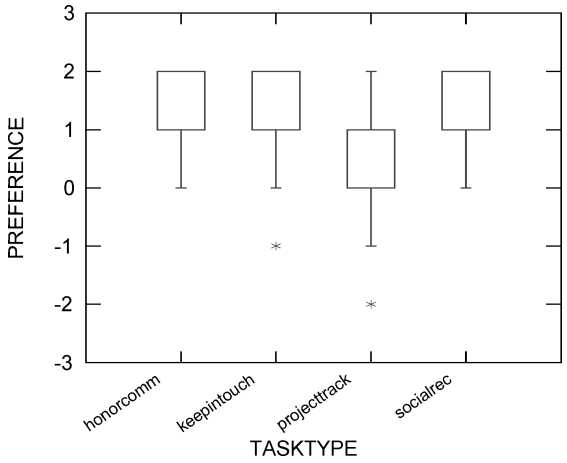

Fig. 7. Subjective preferences for ContactMap or email interface.

to +2 (strongly preferred ContactMap). The box plot in Figure 7 shows that subjects preferred ContactMap for all tasks: the 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 plots show the middle half of the data for each task and the whiskers extending from the box reach to the most extreme nonoutlier. Outlying points are plotted individually). To test for overall preference, we carried out one-sample t tests for each task. On all four tasks, subjects 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).

Results were similar for the five general questions subjects answered at the end. One-sample t tests showed ContactMap was preferred for three of these five

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, $t_{(14)} = 8.57$, $p < 0.0001$). ContactMap and the emailer were rated equivalent for following up email (mean difference is 0.13, $t_{(14)} = 0.49$, $p > 0.10$) and honoring communication commitments (mean difference is 0.33, $t_{(14)} = 1.05$, $p > 0.10$).

To summarize our findings, subjects performed all four tasks better, and three of the four tasks more quickly with ContactMap. Subjects also generally preferred ContactMap to their standard emailer for this set of tasks. Further investigation clearly is warranted, however, to clarify inherent differences between the tasks and understand more deeply the situations where ContactMap offers the greatest benefits.

4.2.3 *Subjects' Comments About Interface Preferences.* We analyzed the comments subjects made after each task to explain their preference. Their comments showed how, in contrast to email, ContactMap supported social reminding and provided structured access to social information.

As expected, ContactMap was effective for social reminding. Important contacts are constantly visible, which meant there was no need to remember their identity: "*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.*" Even though important names and aliases were usually in the subjects' email archives or address books, subjects considered it laborious to access data from those sources, instead they often tried to *remember* contacts. *"ContactMap reminds me of who my friends are—in Netscape I have to remember myself."* Remembering is known to be less accurate than recognition [Baddeley 1999].

ContactMap also provided structural information that directed visual scanning when locating contacts: *"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."* As we had intended, ContactMap also 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.*" The same structure was useful when trying to access archival information. In email, scanning was hard: relevant contacts were often spread across multiple folders, making it difficult to scan a large set of contacts quickly. And email folders do not allow using groups of contacts as retrieval cues, that is there was no way to view sets of messages involving defined groups of contacts.

4.2.4 *Effects of Map Structure on Performance.* We also investigated how map structure influenced successful task outcomes. Was the sheer number of contacts on the map critical for success using ContactMap? Or were groups and structure more important? Was complexity a problem—did more contacts or groups lead to reduced performance? We tested these hypotheses about the

relationship between map structure and the performance measures of success and time.

We found that having more contacts or more groups on the map did not affect either objective performance ($r_{(14)} = 0.30$, $p > 0.05$, $r_{(14)} = 0.22$, $p > 0.05$), or task completion times ($r_{(14)} = 0.18$, $p > 0.05$, $r_{(14)} = 0.27$, $p > 0.05$). This indicates that performance was independent of the overall complexity of the map.

The structure of the map was important, however, specifically the extent to which contacts were organized into groups. Subjects with a larger proportion of contacts in groups had higher success scores ($r_{(14)} = 0.54$, $p < 0.05$), and increased preference for ContactMap $r_{(14)} = 0.53$, $p < 0.05$). Having more contacts in groups may improve performance, because more structured maps directly support retrieval and associative reminding. Related contacts can be found in the same location, reducing the need for visual scanning, which may help both social reminding and social data mining.

4.2.5 *Limitations of the Experimental Study.* One weakness of the experiment is that three of the four experimental tasks had subjects send email to groups rather than individuals—a task that ContactMap is specifically designed to do well. Future studies need to control for this, by having subjects respond to individuals.

Subjects also used ContactMap to construct personal maps the day before they carried out the experimental tasks. We do not think, however, that this experience gave ContactMap an advantage for the experimental tasks. The task of map construction is very different from the experimental tasks, and when subjects set up their maps, they had no idea what the experimental tasks might be.

However, it is possible that recent exposure might make subjects more familiar with contacts and their organization in ContactMap than in email. We intend ContactMap to be used on a daily basis as a "social desktop", being repeatedly accessed for many communication tasks. We should, therefore, expect people to be familiar with the layout of their maps in the same way that they are familiar with the structure and functionality of their current email systems. Nevertheless, one area where there may be bias in our procedure is in accessing *infrequent* contacts who may not have been accessed in email for periods of months or years, in contrast to ContactMap where they had been accessed on the prior day. Although subjects were allowed to update and reorganize their email address books, PDAs, and aliases as they set up the map, only seven chose to do so.

4.2.6 *Semi-Structured Interviews and Comments for Redesigning ContactMap.* We conducted semi-structured interviews with six subjects who used the interface for three days after the experiment. We asked general questions about the interface design and solicited suggestions for improvement. Users were generally highly positive about the system. They enjoyed setting up their maps and using the system, although they made comments about possible design improvements.

All subjects were positive about the email alerting facility and the ability to leave notes. They particularly liked the ability to customize alerting so that

alerts were only received for specified contacts or groups. One subject felt that typing reminder notes was cumbersome: he wanted the ability to dictate these notes and have the audio attached to the relevant icon, or transcribed using speech recognition.

All six subjects were positive about the use of photographs. These not only assisted scanning and reminding, but also gave users a strong sense of their important contacts, which was reduced when only names and labels were available. 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 online photos from contacts, or scanning in photos. These methods are time-consuming (although fun), and we would like to find an approach that would make picture sharing easier.

Another set of comments addressed map complexity. Three subjects felt that having too many contacts or groups on the map could be distracting or might increase the complexity of using the interface. These concerns may turn out to be unfounded, however; recall that our results showed that performance was independent of the number of contacts or groups on the map. If complexity did turn out to be a problem, we could apply visualization techniques to address it such as hyperbolic methods that keep peripheral contacts constantly visible, though smaller [Furnas and Bederson 1995]. A related concern mentioned by two subjects was how the interface might scale for larger numbers of contacts. But as we have already pointed out, one's total number of active contacts remains fairly constant, with older inactive contacts continually replaced by newer ones [Whittaker et al. 2002b]. So, if people inherently limit their number of contacts (for cognitive and social reasons), then scalability of the visualization may not prove to be a significant issue.

This leads to the next issue, maintenance. Three subjects commented on the need to continually update their contacts. Consistent with our design goals, however, ContactMap was seen as a *visual workspace*; people felt that using it as a backdrop for work would implicitly prompt them to add new contacts. We might also supplement this implicit reminding with automatic methods. The system could analyze recent email behavior to prompt users to add potentially important new contacts based on the user's prior contact selection behavior [Whittaker et al. 2002b]. Users may also be willing to spend time maintaining their maps as their comments indicated they found this process rewarding. They commented that it was illuminating to see their contacts and try to impose organization on them, because of the insight this gave into their communication behavior.

Two subjects also brought up the significance of phone-based contacts, which are not extracted automatically by our analysis tools. We need to extend our tools to include telephone or voicemail logs [Whittaker et al. 2002]. Automatically extracting contact information from existing online address books would also be useful.

While all six subjects were positive about the ContactMap user interface and its support for social reminding and social data mining, they reinforced the continued importance of message-centric information. They saw person-centered and message-centered tools as complementary, supporting different

types of communication tasks. The following comments point to these trade-offs:

"Seeing contacts is really useful for doing things related to projects or tasks, but if I want to see recent messages or just quickly respond to a message, then Netscape is better."

"It's great to see messages related to people, but I still need the time-based Inbox view to give me an overall sense of what's happened in the last hour or two."

The ideal interface should contain both sets of views, with users able to switch between them based on their current activity. Finally, two subjects pointed out that much of their communication is done while they are on the go, and away from their computers. This indicates a need to explore future versions of ContactMap operating on mobile devices. It raises the significant visualization challenge of presenting complex visual information on a small display.

5. DISCUSSION

Our field study identified social communication tasks that are not well supported by message-centric communication interfaces. We addressed these unmet requirements by designing and evaluating a *social desktop*—a structured, visual representation of the user's important contacts. We directly compared our interface with email for these social communication tasks. Consistent with our original design goals, the social desktop outperforms email in allowing people to keep their contacts in mind, respond quickly to their communications, and to access information relating to those contacts. Although there are weaknesses in our experimental design, these are promising results.

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 advanced prototypes [Bradner et al. 1999; Smith and Fiore 2001; Viergas and Donath 1999] have put people at the center of the user interface. Other successful social networking tools such as Orkut (http://www.orkut.com/) and Friendster (http://www.friendster.com/index.jsp) are focused on the utility of public networks for making friends. ContactMap extends this work by expanding the social interface beyond real-time messaging and using it as a method to potentially access all asynchronous communications. We also extend the social representation to depict the structured personal organizations people use for everyday personal work tasks.

ContactMap may also generalize to other important user tasks and applications. A structured social representation might be used as an alternative interface to the operating system, allowing documents and files to be accessed using social information [Dourish et al. 1999]. It could also be used for filtering in communication applications, allowing important contacts' messages to receive priority, or controlling social access in location-based applications. Finally, a structured social interface might be used in more collaborative contexts for the sharing of documents and expertise.

Our work also differs from earlier systems in providing empirical data about the utility of the novel interface. With a few exceptions [Bradner et al. 1999; Isaacs et al. 2002] much prior work has focused on novel interface designs and

paid less attention to the tasks that the interfaces are intended to support and whether they prove useful. Our users' preference for simple representations and for creating their social maps themselves illustrated the benefit of our empirical approach. Careful user evaluation of our early prototypes that automatically extracted and organized contacts showed the problems with this strategy and led us to abandon it. Our experiment also showed the benefits of our final design.

Nevertheless, one outstanding design issue concerns integration with existing message-centered user interfaces. While we have demonstrated benefits of a social desktop interface for social reminding and social data mining, our subjects were clear about the utility of current interfaces for standard message processing tasks. An effective way to combine the two approaches might be to have ContactMap as one of several views onto email data. Users would then have both people and message-centric views of their communication records, and be able to choose the one that best suits their current task.

Other research questions involve discovering how these social desktop representations evolve over time. It is clear that contacts change over time [Wellman 2001b; Whittaker et al. 2002b], but how can we best support this process? Can we automatically recommend contacts be added or removed [Resnick and Varian 1997; Terveen and Hill 2001]? And since people in close-knit workgroups often have largely overlapping sets of contacts, could we support contact sharing, with all the privacy considerations this entails?

Our results also have implications for theories of mediated communication, particularly asynchronous communication. Elsewhere we have argued that theories of mediated communication are derivative of face-to-face communication theories [Nardi and Whittaker 2002; Whittaker et al. 2002, 2002a, b; Whittaker 2003]. As a result, they have tended to focus on the act of communication itself, that is *interaction*, as opposed to the processes needed to make such communication take place at all, that is *outeraction* [Nardi et al. 2000a, b; Whittaker 2003]. 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 in touch with someone, or 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 to account for them.

A similar focus on interaction may explain the limits of message-centered UIs. Message-centered interfaces focus on the interaction event itself, that is, 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 finding information about the message recipient. Again, refining our notion of outeraction should lead to interfaces that better support all aspects of asynchronous communication. Overall, our results argue for the usefulness of personal social desktops. By providing users with straightforward methods to extract and visualize their social contacts, we hope to provide better support for the outeraction tasks of social reminding and 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*. ACM Press, New York. 97–105.

AMENTO, B., TERVEEN, L., HILL, W., HIX, D., AND SCHULMAN, R. 2003. Experiments in social data mining: The topicShop system. *ACM Trans. Comput.-Hum. Interact. 10*, 1, 54–85.

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

BADDELEY, A. 1999. *Essentials of Human Memory*. Taylor & Francis, London.

BLY, S., HARRISON, S., AND IRWIN, S. 1993. Media spaces: Bringing people together in a video, audio, and computing environment. Commun. ACM *36*, 28–47.

BRADNER, E., KELLOGG, W., AND ERICKSON, T. 1999. The adoption and use of babble: A field study of chat in the workplace. In *Proceedings of the European Computer Supported Cooperative Work Conference*.

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

DONATH, J. 1995. Visual who. In *Proceedings of ACM Multimedia '95* (Nov 5–9). San Franciso, CA.

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

DOURISH, P., EDWARDS, W. K., LAMARCA, A., AND SALISBURY, M. 1999. Presto: An experimental architecture for fluid interactive document spaces. *ACM Trans. Comput.-Hum. Interact. 6*, 2.

DUCHENAUT, N. AND BELLOTTI, V. 2001. Email as habitat: An exploration of embedded personal information management. *Interact. 8*, 5, 30–38.

DUNBAR, R. 1998. *Gossip, Grooming and the Evolution of Language*. Harvard University Press, Boston.

ERICKSON, T. AND KELLOGG, W. 2000. Social translucence: An approach to designing systems that mesh with social processes. *ACM Trans. Comput.-Human Interact. 7*, 59–83.

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*. ACM Press, New York.

GRANOVETTER, M. 1973. The strength of weak ties. *Amer. J. Sociol. 78*, 6, 1360–1380.

HERRING, S. 1999. Interactional coherence in CMC, *J. Comput. Mediat. Commun.* 4.

HINDS, P. AND KIESLER, S. 2002. *Distributed Work*. MIT Press, Cambridge, MA.

HUDSON, J. M., CHRISTENSEN, J., KELLOGG, W. A., AND ERICKSON, T. 2002. "I'd be overwhelmed, but it's just one more thing to do:" Availability and interruption in research management. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI'02)*. ACM Press, New York.

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

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.

KAUTZ, H., SELMAN, B., AND SHAH, M. 1997. The Hidden Web. *AI Magazine 18*, 2, 27–36.

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

KRAUT, R., FISH, R., ROOT, B., AND CHALFONTE, B. 1990b. Informal communication in organizations: Form, function and technology. In *People's Reactions to Technology in Factories, Offices and Aerospace*, S. Oskamp and S. Spacapan, Eds. Sage Publications. 145–199.

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.

MALONE, T. 1983. How do people organize their desktops? Implications for the design of office systems. *ACM Trans. Off. Inform. Syst. 1*, 99–112.

MILEWSKI, A. AND SMITH, T. M. 2000. Providing presence cues to telephone users. In *Proceedings of Computer Supported Cooperative Work*. ACM Press, New York.

MORTENSEN, M. AND HINDS, P. 2002. Fuzzy teams. In *Distributed Work: New Ways of Working Across Distance Using Technology*, P. Hinds and S. Kiesler, Eds. MIT Press, Cambridge, MA. 283–309.

NARDI, B. AND WHITTAKER, S. 2002. The role of face-to-face communication in distributed work. In *Distributed Work: New Ways of Working Across Distance Using Technology*, P. Hinds and S. Kiesler, Eds. MIT Press, Cambridge, MAs.

NARDI, B., WHITTAKER, S., AND SCHWARZ, H. 2000a. 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., AND BRADNER, E. 2000b. Interaction and outeraction: Instant messaging in action. In *Proceedings of Conference on Computer Supported Cooperative Work*. ACM Press, New York. 79–88.

NARDI, B., WHITTAKER, S., ISAACS, E., CREECH, M., JOHNSON, J., AND HAINSWORTH, J. 2002. Integrating communication and information through ContactMap. *Commun. ACM 45*, 89–95.

RESNICK, P. AND VARIAN, H. 1997. Recommender Systems. *Commun. ACM 40*, 3, 106–108.

ROUNCEFIELD, M., HUGHES, J., RODDEN, T., AND VILLER, S. 1994. Working with constant interruption: CSCW and the small office. CSCW, ACM Press, New York. 275–286.

SACK, W. 2001. Conversation map: An interface for very large-scale conversations. *J. Manage. Inform. Syst.*

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

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

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

TANG, J. C., ISAACS, 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.

TERVEEN, L. AND HILL, W. 2001. Human computer collaboration in recommender systems. In *HCI in the New Millennium*, J. Carroll, Ed. Addison Wesley.

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

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. *Internat J. Urban and Region. Resear. 25*, 2, 227–252.

WELLMAN, B. 2001b. Computer Networks as social networks. *Science 293*, 2031–2034.

WHITTAKER, S. 2003. Theories and methods in mediated communication. In *The Handbook of Discourse Processes*, A. Graesser, M. Gernsbacher, and S. Goldman, Eds. LEA Press, Mahwah, NJ.

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

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., HIRSCHBERG, J., AMENTO, B., STARK, L., BACCHIANI, M., ISENHOUR, P., STEAD, L., ZAMCHICK, G., AND ROSENBERG, A. 2002. 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*. ACM Press, New York.

WHITTAKER, S., JONES, Q., AND TERVEEN, L. 2002a. Persistence and conversation stream management: Conversation and contact management. In *Proceedings of HICCS35*. Maui, HI.

WHITTAKER, S., JONES, Q., AND TERVEEN, L. 2002b. Contact management: Identifying contacts to support long term communication. In *Proceedings of Conference on Computer Supported Cooperative Work*. New York: ACM Press. 216–225.

Received May 2002; revised October 2003, July 2004; accepted July 2004 by John M. Carroll