hebb: An Infrastructure for Conveying Shared Interests

Scott Carter[1], Jennifer Mankoff[1], Patrick Goddi[2]

[1] Electrical Engineering and Computer Science Department
University of California at Berkeley
Berkeley, CA 94720-1770
jmankoff,sacarter@cs.berkeley.edu
[2] Hewlett-Packard
1501 Page Mill Road
Palo Alto, CA 94304
goddip@exch.hpl.hp.com

Abstract. Awareness of others' interests can lead to fruitful collaborations, friendships and positive social change. Interviews of groups involved in both research and corporate work revealed a lack of awareness of shared interests among workers sharing an organizational affiliation and collocated in the same building or complex but still physically separated (*e.g.*, by walls or floors). Our study showed that loosely coupled, co-located groups were less likely to discover shared interests in the way that many tightly collocated groups do, such as by overhearing conversations or noticing paraphernalia. Based on these findings we iteratively developed techniques to capture and display shared interests. Our platform includes an e-mail sensor to discover personal interests, a search algorithm to determine shared interests, a public peripheral display and lightweight location-tracking system to convey those interests.

1 Introduction

Mutual relationships can provide springboards for conversations. Such relationships are built on common interests. The hypothesis driving this work is that making shared interests visible will support conversation and help build relationships. We conducted a series of interviews to determine specific communication and collaboration issues that arise among small size working groups. The results of these interviews drove the development of techniques to capture user interests and display shared interests to collocated group members. Our contributions include the identification of the need for improved grounding and communication among loosely coupled, co-located groups, and the development and deployment of a system that can capture and display shared interests to support such grounding.

Previous investigations of small size working groups have shown that when group members are collocated they are more likely to share context and experiences that lead to communication and collaboration. Specifically, collocation has been shown to facilitate the kind of impromptu meetings and exchange of tacit knowledge important in establishing the common ground necessary to conduct meaningful conversations [19]. However, in a series of interviews that we conducted with members of six small working groups we found that the benefits of collocation often do not extend beyond the group. Shared interests with nearby groups or individuals with similar interests go unnoticed, limiting the chance for communication and collaboration.

In his analysis of conversations, Clark suggests that conversation participants frame their discussions in terms of what all parties understand to be common knowledge, context or shared experiences [4]. Common context may include an employer or a common space that all common participants share. However, certain pieces of common ground, such as shared interests, are often not as readily discernible. Our goal is to expose previously unknown shared interests to improve common ground and induce cross-group collaboration and communication.

To this end, we developed a system designed to capture and convey shared interests. From our interviews we discovered that most group members in the work places we studied used e-mail as their primary means of communicating interests to others, and we therefore designed a sensor that culls interests from e-mails. Furthermore, we used iterative design techniques, including paper prototyping, to build a peripheral display to discover and display shared interests. Our display is designed for deployment in one or more public spaces frequented by members of loosely related groups that could benefit from support for building closer relationships. When people with common interests approach one of our displays, those interests are shown on those displays.

In this paper we first describe previous work that has looked at the impact of informal communication in work environments. We then describe a formative evaluation consisting of a series of interviews of research groups and other working groups that revealed a lack of informal communication and collaboration between collocated but physically separated groups. We then describe how this evaluation led to the design and evaluation of a system, *hebb*, capable of sensing and conveying shared interests to users.

2 The Informal Exchange of Knowledge

Informal communication is important to support because it helps people establish common ground necessary for meaningful conversations and relationships. Common ground, as Clark defines it in his book, *Using Language*, is information that two parties share and are aware that they share. According to Clark, "Everything we do is rooted in information we have about our surroundings, activities, perceptions, emotions, plans, interests. Everything we do jointly with others is also rooted in this information, but only in that part we think they share with us" [4]. A key concept developed by Olson and Olson is that collocated workers are better able to establish such common ground. In their work on methods of supporting distributed working groups they have outlined some of the aspects of collocated work that often make it more successful than remote work [20,22,21]. In particular collocated groups are more likely to experience activities collectively, which, according to Clark, is “the most important source of common ground” [16]. Furthermore, other work has used the theory of common ground to model awareness of the users concerns [3] and facilitate collaborative repair tasks [24].

Even though grounding is better established in collocated groups than remote ones, collocated groups still miss many opportunities to interact and collaborate. In most organizational contexts, the exchange of tacit knowledge is critically important and even a collocated group may not have the grounding necessary to hold conversations that could lead to collaboration [14]. Studies that have investigated the nature of informal communication practices in organizational environments show that nearly all intra-workplace transfer of tacit knowledge occurs during unscheduled, brief interactions [32,1]. Furthermore, workers are often mobile within their environments — walking to other offices or buildings within their complex to discover what other people are working on (“walkabouts”) or where they are, to discuss something urgent or to coordinate some other activity [1]. In most of these cases, conversation is spawned by at-hand phenomena, such as a poster on the side of a cubicle or an overheard conversation with a colleague.

While many of these studies concentrated on how tacit knowledge is transferred amongst collocated groups they did not address situations in which such knowledge is *not* transferred, a particular problem for the loosely-coupled groups we studied. As humans are able to make use of opportunistic, situated information [5], technology could potentially help to support grounding in such situations. Many technologies have been created to support media connections between spaces [6,7,9]. Some public displays have been created that attempt to address the issue of providing common ground to inspire conversation. McCarthy’s Groupcast is a peripheral display that recognizes passers-by and posts content of interest to at least one of the users [15]. Snowdon’s and Grasso’s CWall published articles captured from a community web site to an interactive public display [29]. Huang’s *et al.* Awareness Module present items posted by users in a group to a public displays [12]. However, none of those were deployed specifically to increase ties between different weakly connected but collocated groups. The groups we studied share the disadvantages described by Olson and Olson, but are not distributed and thus could benefit from face-to-face contact should common ground be highlighted or established.

The aforementioned studies on the importance of and mechanisms for providing common ground, combined with our interviews (described next), led us to build the application described in this paper, which actively displays information that can function as conversational reference-points.

3 Interviews

We conducted a series of interviews to determine problems and issues with current work communication practices. Specifically, we conducted open-ended interviews with one to two members selected from six different small (5-10 members) working groups that included academic and industrial research labs, a corporate design firm and an academic administrative group. From these interviews we generated a series of sociograms depicting social relationships. Though we were originally interested in evaluating collaboration and communication within each group, we quickly learned that while strictly collocated groups share many resources for sharing tacit knowledge the same phenomena did not extend to groups that were nearby but nonetheless unable to hear or see each other even when they shared common interests.

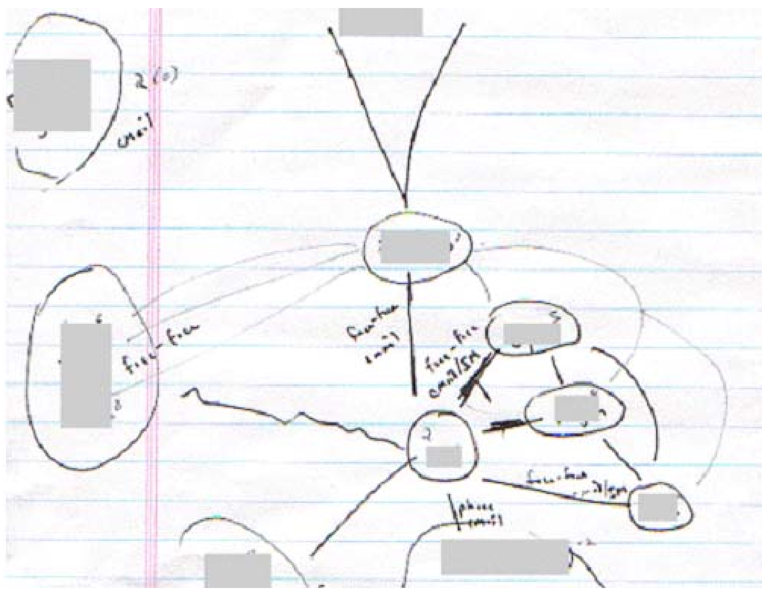

Fig. 1. Part of a drawing made by a participant during an interview. Notice the "information funnel" at the top, grouping and linking. Names have been removed to maintain anonymity.

3.1 Method

We conducted interviews with one to two members of six different groups: three academic research labs, one industry lab, one design firm and one academic administrative group. We e-mailed groups to solicit subjects and chose the members longest associated with the group. We asked participants open-ended questions about their core working group's collaboration and communication habits. Several questions involved their group's relationship with other nearby groups working on similar topics. Surprisingly, we found that participants were quite fond of drawing their own view of their social network (See Figure 1). After the first participant volunteered a graph using a nearby whiteboard including their relationship to members of those other groups, we asked all other subjects to do the same on paper (all of which, save one, drew detailed diagrams). These artifacts were useful as a point-of-references during the interview and also aided in our sociometric analysis.

3.2 Creating a Picture of Interaction: Sociograms

Data gathered from qualitative research, and especially interviews, is often difficult to analyze. Often interview data is transcribed and cataloged but utilized only in ad hoc references [18]. Therefore, we began interviews with the specific goal of being able to generate sociograms, visual representations of the strength of ties for a particular group [31], from the data gathered. This goal not only helped to direct the flow of interviews but also facilitated analysis of interview data. Social network methodology is gaining recognition in many academic communities as a standard means of investigating social structures [26,8,30,17]. Core to social network theory is the belief that individual, group and organizational behavior is affected more by the kinds of relationships in which actors are involved than by their own particular conditions. Therefore, a social network analysis seeks to determine tie strength between individuals in a group. The strength of that tie depends on the number and types of resources, such as verbal information, documents or goods they exchange, the frequency of exchanges, and the extent to which at least one party considers those resources personal [11]. Graphical representations of groups arise naturally out of a social network analysis.

3.3 Results

We found that members of groups tend to have strong connections to other members of the same group, usually born out in e-mail communication, but significantly weaker connections to other groups in their same organization. This holds even when the other groups work on related issues or are related to the core group through an organizational construct. For example, one member said of groups near him working on similar topics that:

"I mean, I do [try to talk to these groups], I try to stick my head in and say, 'Hey, what you guys doing? You know, I'm John, and I work two doors down from you ... We kind of do the same thing; let's try to stick together.' But we have very little interaction. I think it would be kind of cool ... I'd like to promote that ... I know what they do two doors down, but I don't know what they do one door down from the left."

In this case the group "one door down from the left" in fact conducts research that overlaps with this group. Furthermore, articles that group published recently were attached to the main door to the group's lab. When asked about this, the group member cited above said that he considered most artifacts posted outside the door to be "stale." Thus, artifacts that could lead to collaboration are there to be noticed but there exists no catalyst to contextualize that noticing in day-to-day routines. This group member went on to mention some work he is doing to try to foster more relationships between his group and nearby groups by holding joint meetings but mentioned it was difficult to get high attendance rates. A member of a different group said that he is "definitely" interested in collaborating more with nearby groups but noted that as most interaction occurred through inter-group e-mail there was little chance of discovering shared interests. A third called inter-group interactions "vital" to his work but lamented that the difficulty was making sure the "right person saw the right thing." In all three cases mentioned above, artifacts that could lead to collaboration between groups do exist to be noticed but are buried in communications that are too formal and decontextualized to pique interest.

Another common trend our interviews revealed is a heavy reliance on e-mail to communicate new ideas. For example, one group member mentioned that when he often "generate(s) a list of questions that I have about, you know, various aspects of a project ... that I'll just e-mail to [my boss]." Another interviewee reported relying on e-mail as the primary means of deciding new directions to take on upcoming projects. Still another reported replacing informal requests (*e.g.* a request to go to lunch) with e-mail when he had to work in a space separate from his group. Groups that did not depend as heavily on e-mail tended to hold more spontaneous face-to-face meetings to develop ideas.

3.4 Sociogram analysis

We created a sociogram for each group. We assigned each person a node and placed two nodes closer together the more they share in common. Node shade indicates the number of connections that node shares with other nodes. Thus in the sociograms in Figure 2 darker nodes share more connections and lighter nodes less. We calculated edge widths using the following information:

- Whether or not group members shared the same physical space
- The number of projects that the two group members shared in common
- The number of organizational structures that the group members shared in common
- Self-reported importance of other members

Concretely, a node edge was given one point if the two people represented by the edge shared the space, one point for each project and organizational unit in which the two are involved (*e.g.*, people sharing the same department and same lab would get two points) and a point if an interviewee specifically singled out a person as important. We used the Netdraw program [2] to generate the sociograms shown in Figure 2.

Note that after generating the sociograms we grouped organizational units using shaded areas. The darkest area represents the subgroup that we interviewed in each case. The other subgroups were unknown to us originally and emerged out of the interview analysis as related in some way (either organizationally or as a community of practice) to the six subgroups that we interviewed. For example, the dark gray area in Figure 2 (b) is a subgroup that we analyzed, and all other nodes represent non-members who are nonetheless important to the group. Furthermore, we distinguished between two types of organizational units — the core group analyzed in interviews and the larger group of which it is a part. We used a lighter area for the larger group.

In most cases the sociograms provide visual evidence of a lack of collaboration among related groups. For example, Figure 2 (a) depicts two communities of practice within a larger group that largely do not collaborate even though they are working on similar tasks. Here, the lack of collaboration is due to physical separation and a top-down work structure. The subgroup in Figure 2 (b) has almost no interaction with other members of their larger group because of physical separation and unfamiliarity because of an organizational restructuring. The subgroups in Figure 2 (c), (d) and (e) are largely isolated from outside influences. The group in Figure 2 (f) is interesting because the group that we interviewed overlaps with but is not exactly equal to the most connected group (indicated by the darkest nodes) — two of the most connected nodes fell outside the core group we interviewed. These external members had lower rank and spent less time in the group's space than other members.

Reasons for a lack of inter-group connectivity vary from restructurings that place two unfamiliar social networks under the same organizational umbrella, to labs that are physically separated but in the same building, to a top-down work structure that limits the extent to which spontaneous interactions can have an impact.

3.5 Discussion

Sociometric analysis showed that the subgroups we studied have limited awareness of and communication with other subgroups. Sociograms are a good analysis tool for getting an overall picture of the social dynamics of a group. However, they leave out some key characteristics that may be important in analyzing the effects of new technologies introduced to the network. For example, it is difficult to integrate information about the specific media used between nodes or to categorize interactions between nodes [11]. Nevertheless, our sociometric analysis did show that small groups often share only weak connections with other groups regardless of media.

From our interviews we learned that people with similar interests often do not communicate about these interests because they are unaware of the relationship. As a result, many topics that people share an interest in remain unexplored. In all but one group, group members expressed a desire to "know what people are working on" in nearby rooms and office and have more opportunities to work with the people in those spaces. To accomplish these goals, the groups need a lightweight means of starting conversation. People are unlikely to explore some new topic, however, unless they share some common ground [10]. Often this grounding is established by features of a collocated work space not available to the groups studied here. Thus, technologies are needed to encourage cross-group

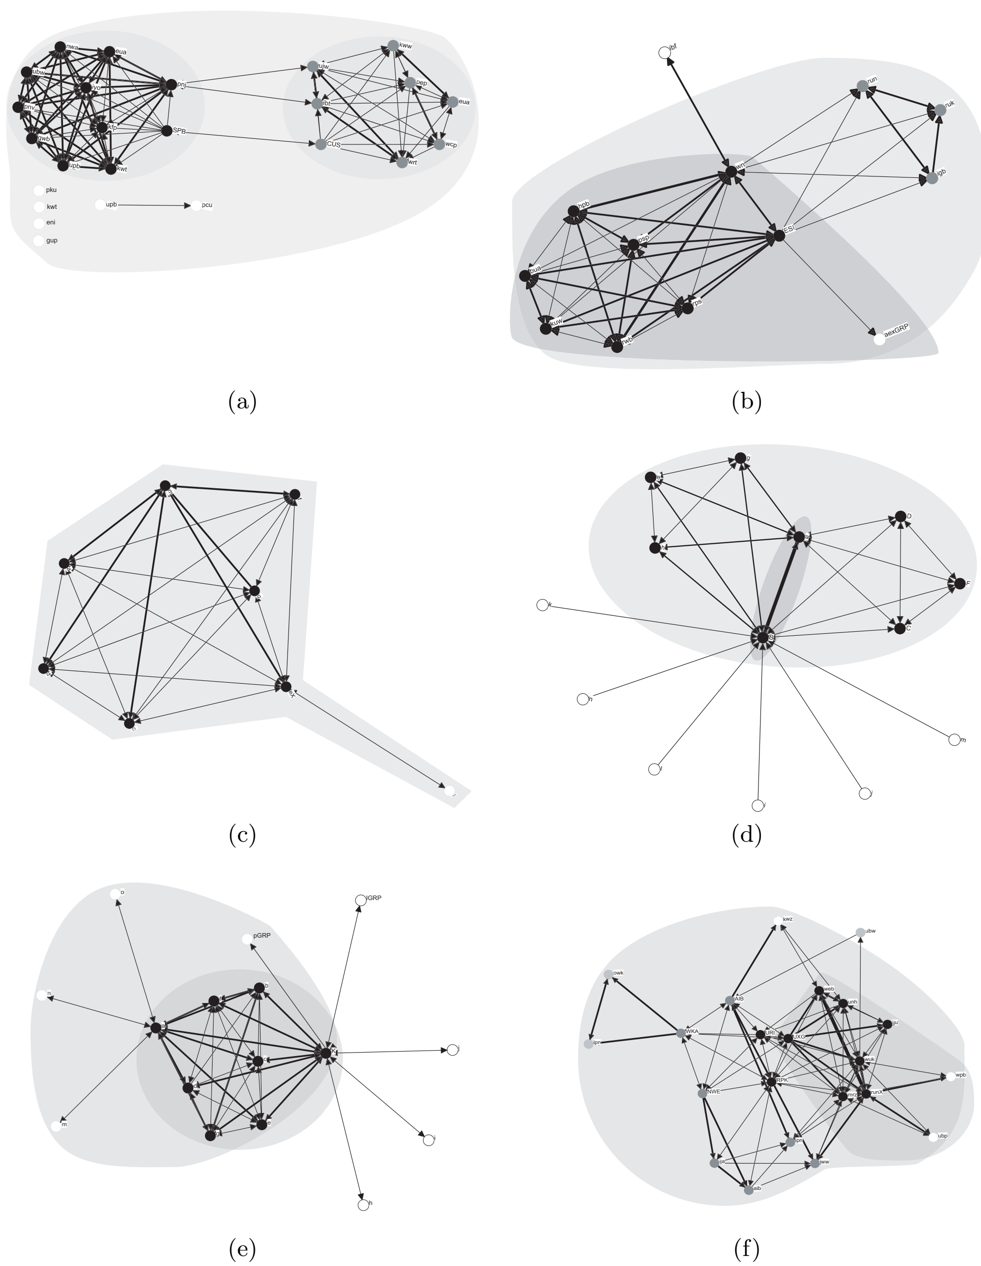

Fig. 2. Sociograms for all groups. Gray areas are demarcations of organizational units, the darkest area in each sociogram indicating the boundary of the group that we interviewed. In all cases save one (graph c), members of the core group identified other influential individuals as being part of one larger organizational unit (nodes in the light gray area) or outside interests (nodes in the white area).

collaboration in collocated settings to provide this common ground through relationships between group members.

4 System Design

Based on our interviews we set out to iteratively develop a system capable of finding and displaying shared interests. In this section, we present a scenario of how this system might be used. Also, we describe the interest sensor we built in more detail, specifically how the interest sensor extracts information from e-mails using part-of-speech tagging and how we dealt with privacy concerns by extracting only high-level topic descriptions. We then describe the design of a public and a personal display to share these topics as well as an evaluation of the public display that clarified the seminal information that it should convey.

4.1 Use Scenario

To illustrate how our interest sensor might be used to build common ground we begin with a scenario based loosely on an actual system we have built: John and Chris are researchers in a cognitive science department. John recently sent an e-mail to Chris regarding a new haptics system that they are considering purchasing for their research in human perception. Mary, who works on cognitive models of users of neurosurgical devices and works in a lab downstairs from John and Chris, recently sent one of her peers an e-mail about a haptic device to simulate surgical situations that her friend uses at another university. One morning while Mary is getting coffee in a breakroom near her office, Chris stops by a different breakroom close to his office to refill his water bottle. A public display in both spaces recognizes their presence, generates a shared interest and posts it. Mary notices that a projected display in her room now displays a graphic window containing a few attached phrases: “shared interest: haptics” “shared by john (not around)”, “mary (in downstairs breakroom)”, “chris (in upstairs breakroom)”. Mary reads the display and pushes a button on the window to maximize it. Another window pops up containing Mary’s, John’s and Chris’s pictures. She sees Chris’ picture, recognizes him, clicks on it, and starts to walk to the breakroom upstairs. Chris, who had not noticed any change in the display at first and was about to leave, now sees that it is blinking a message: “someone visiting Chris from downstairs breakroom regarding haptics.” Instead of leaving, Chris pulls out his PDA and navigates to the web page for the public display where he finds e-mails regarding the haptics system he has sent to John. When Mary arrives, they talk about how they could use the haptics system on both of their projects. Since he knows that the public display web site does not provide Mary access to his e-mails, he decides to forward them to her...

4.2 System Components

In this section we explain how various system components work together to convey shared interests as users carry out their daily routines. We first describe

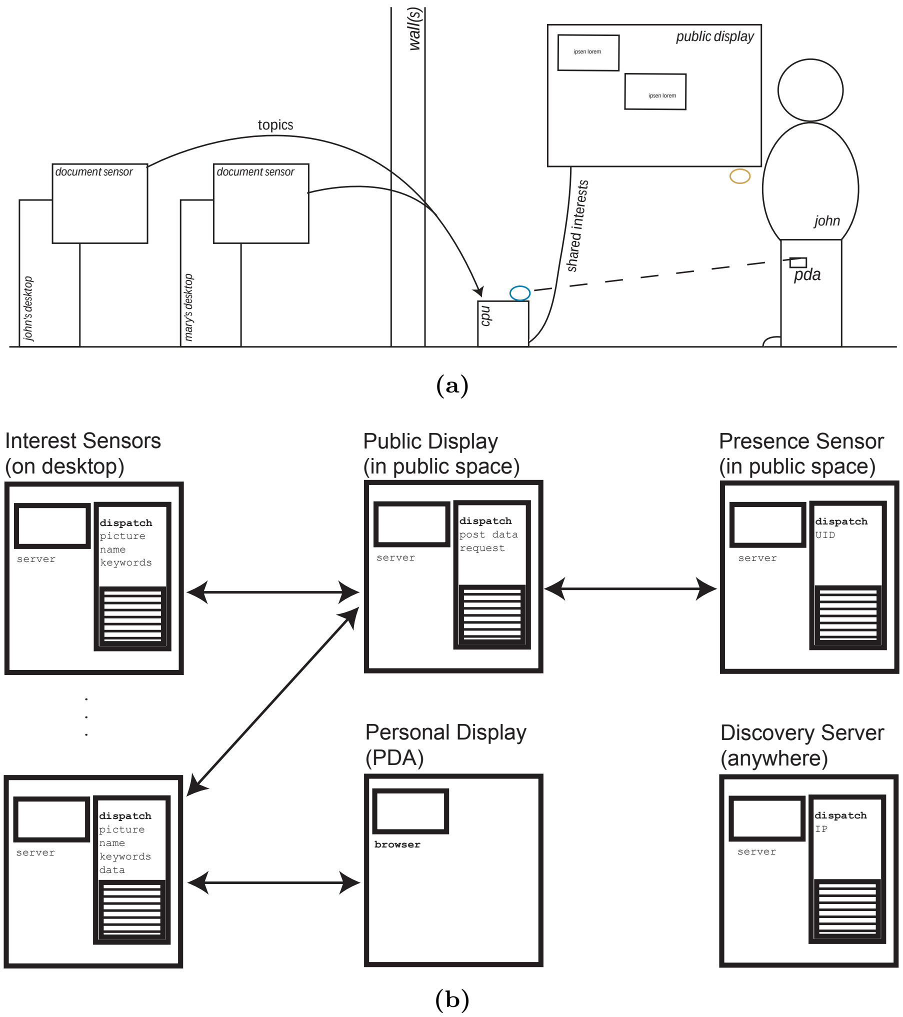

Fig. 3. (a) System use scenario. Two interest sensors, one running on John's desktop and the other on Mary's (Chris's not shown), extract general topics of interest from their e-mails. These topics are forwarded to a nearby public display, which John happens to be near. The display senses that John's PDA is nearby and shows topics of interest he shares with others that he may not know about. **(b)** System architecture. Components send events they generate to other registered components.

how all of the components interact and then detail the design of each individual component.

Overview Components used in this system include interest sensors, public displays and presence sensors (see Figure 3). Each component can register with another component to receive semantically tagged data that component generates. The interest sensor generates picture, name and keyword events, corresponding to a picture of the user of the interest sensor, the user's name, keywords (generated from the algorithm described in the next section) as well as encrypted full document data (for use on personal PDAs). The public display generates events indicating from which documents keywords were recently displayed and the presence sensor generates unique user identifiers (UIDs) based on users sensed in the space via either RFID badging or presence of the user's PDA on the local wireless network. A server on each component handles incoming requests by saving the IP address of the requesting component in an event-specific table. When a component generates an event, it sends it to all components listed in the table. Components, especially document sensors, may go on and off line sporadically and may not be statically addressed. For this reason, a discovery server allows components to update and retrieve location information.

Interest Sensor Design Interest sensing could be done in number of ways, including recording chatter in public spaces, sensing the content of recent printouts or user specification. Our interviews revealed that most groups developed new ideas either in face-to-face, spontaneous meetings in their lab or over e-mail, and it is important to concentrate on the early stage of ideation as that is the time groups are most open to collaboration [14]. We considered two means of using microphones to capture interests: a direct audio link and voice recognition with visual display [13]. However, as our analysis found that information was remarkable only in so far as the receivers perceived it to be relevant, we discarded the former as it lacks a filtering mechanism. Voice recording and recognition is appealing as it allows both easy installation and filtering, but has two central problems both arising from difficulties discovering the identity of who is talking because of environmental noise: 1) without being able to identify exactly who shares the same interests we would have to present that information generally to both groups, again decontextualizing the information and 2) fine, per individual privacy controls would in some cases be impossible to establish.

For those reasons, and because most of the subjects we interviewed indicated that they use e-mail as a primary communication medium with others in their group we chose to sense interests from e-mail. Furthermore, e-mail has been shown to be an effective means of discovering shared interests [28]. We used interviews and a pilot deployment to determine the most salient shared interests and report on our findings below.

The interest sensor is written in Java and is designed to have minimal impact on work practice, requiring setup only once. Thereafter it restarts each time the user reboots her machine. To setup the sensor the user completes the dialog screen, entering configuration information for her primary e-mail account. The e-mail sensor then checks mail once every ninety seconds on this account. For each message received (and sent, if available) the sensor first eliminates commonly used generic terms using keyword matching and then assigns the message an ID number and extracts a list of pertinent nouns and phrases using the part-of-speech (POS) tagger Qtag [23]. The sensor also records all attachments received, parsing only the name of the attachment. The sensor then sends the list of all important phrases and attachment names to all known nearby displays.

One major issue we faced was privacy concerns, because shared interests must be displayed to potentially unknown others to help build relationships. Therefore, our solution involves releasing only one- or two-word descriptions of possible interests determined from an e-mail document. The remainder of the document does not leave the computer on which the interest sensor runs. Only these high-level and decontextualized descriptions leave the user's individual machine to be available to matching algorithms running on other machines. Furthermore, when a keyword is posted, the full content of the document in which it appeared is made available to user's personal PDAs.

When topics sent from individual interest sensors arrive at a public display relevant messages are selected to display as follows: They are associated with the unique user ID of the interest sensor and cataloged. Then, given a particular set of users, the public display runs an algorithm to determine interests to display. The algorithm used is similar to the Term Frequency inverse Document Frequency (TFiDF) algorithm used in other text retrieval systems [27]. Specifically, the algorithm builds a word vector for each document based on words sent from that user's interest sensor. It then assigns a weight to each word in the vector that is directly related to its frequency within the document but indirectly related to its frequency across all sensed documents. The algorithm then multiplies these vectors for each pair of sensed users, displaying the keywords with the highest rank. In this way, the algorithm limits the displayed interests to those that are as yet relatively unexplored.

Public Display Design The primary purpose of the public display component is to convey the gist of shared interests as well as to facilitate collaboration amongst interested parties. To this end, we designed a public display capable of showing topics gleaned from interest sensors and people related to those topics while also providing some mechanisms to allow users to interact with displayed information.

The public display tailors displayed interests to people who would most likely see them using a presence sensor. The public display uses two means of sensing nearby users: implicitly by finding nearby PDAs and explicitly by means of an RFID reader situated next to the display that users can badge in to. To sense nearby PDAs, the presence sensor simply downloads a list of MAC addresses from the wireless access point closest to the display and uses that information to index into a table of MAC addresses and user names. Similarly, the RFID reader senses a unique ID and uses that information to index into a table of

RFID IDs and user names. When the presence sensor discovers a change in the users at a space it notifies the public display and sends it a list of users. The public display then searches for relationships between these users. In addition, when a user badges in explicitly the public display responds by showing their name briefly on the display to provide feedback of successful recognition.

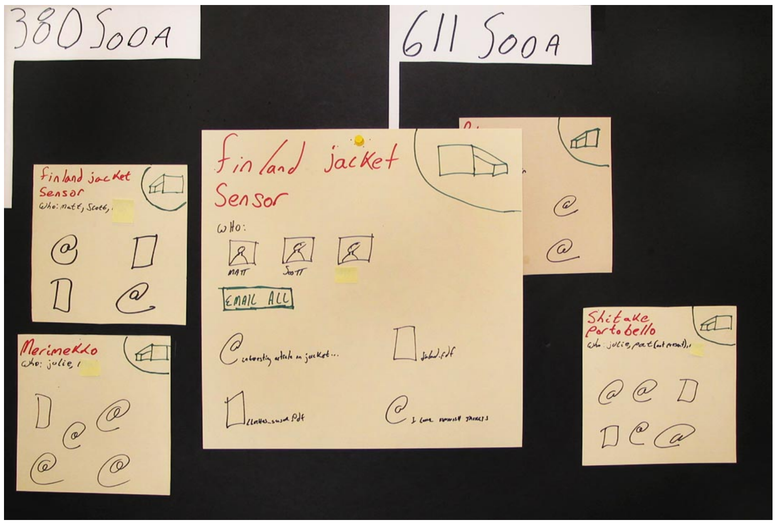

Fig. 4. Public display prototype.

Personal Display Design Personal PDAs allow each user to see the full contents of documents from which displayed keywords were derived. Furthermore we designed the system such that users can view the full contents of documents generated by *their own* interest sensor only. Users can access this data on their personal PDA via a Web site associated with the public display.

4.3 Paper Prototype Study

Before building the final version of the system we evaluated a paper prototype of the display with two users using Rettig's methodology [25]. The paper prototype consisted of a black poster board on which we attached several small cut-outs that represented found interests. The cutouts included: the topic, names of people associated with that topic, icons representing documents (e-mail messages or attachments) related to that topic and a maximize icon.

We recruited two participants with relatively little familiarity with public displays. Our approach was to provide the participants with a scenario in which they were waiting for someone else in a hallway and had noticed the display in an alcove before approaching it. We conducted the study in just such a space – an alcove off of a hallway. Before arriving at the study, we asked users to notify us of some of the topics that they had recently discussed in their e-mails. We used this interest list to generate several topics that they might have in common with fictitious users. When we instructed the user to approach the screen, a person "playing computer" would post relevant found interests. When a user pressed the maximize icon the computer would put up a larger cutout onto the poster board that included pictures of the users sharing this interest as well as a button allowing them to e-mail all of the users and notify them of a visit to their space. We then asked each user to complete a series of specific tasks testing interaction with the display.

We found that our interface, though minimalist, was still too cluttered for lightweight interactions and that people were confused about which spaces other users currently occupied. For example, note in Figure 4 the icons at the bottom of the topic cut-outs. These were meant to convey the general nature of common documents, such as whether they are e-mail messages or PDFs or some other kind of document. But subjects were unable to explain them as anything but "random." Furthermore, topics are positioned according to which area its users occupy. Specifically, if a user is in 611 Soda any relationship with someone in 320 Soda will show up on the left half of the display. We originally designed the display this way to separate topics shared among users in the same space versus nearby spaces. However, subjects were not readily able to determine this. Additionally, this particular arrangement made it difficult for subjects to interpret the placement of topics shared by people across many spaces.

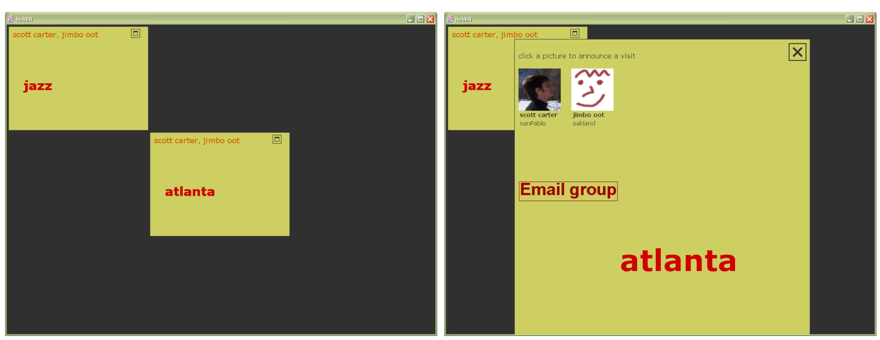

(a) (b)

Fig. 5. **(a)** Topics display **(b)** maximized view of a topic

In our next iteration we took out all references to individual documents and put user location directly next to user names (Figure 5). Also, we added some lightweight interactive features into the interface to start communication about some topic when not all of the users are in the same room. Users are able to both send an e-mail about the topic to all interested users and, as in the scenario, notify users near other displays that they are "on their way" to that space to talk about one of the displayed topics. In both cases, users only need touch one button on the display ("E-mail group" to e-mail and a person's picture to announce a visit).

In summary, we designed and tested a public interest display on paper. Based on user feedback, we then implemented a fully working system that senses and responds to user presence with shared interests, The displays are intended for deployment in one or more public spaces shared by loosely connected groups, and interests are determined by the union of users present near any deployed displays.

5 Early Results

The success of the system we have developed depends on the extent to which it measurably encourages communication between group members about topics that they otherwise would not have explored. To test this, we are in the process of deploying a functioning system to working environments. We initially deployed the interest sensor to one academic and one industrial research group for two weeks for early testing. We are following up this pilot deployment with another deployment of the full system to two academic research labs.

In our initial pilot, we deployed the interest sensor to two groups of five and six members each and found that while users appreciated that the software did not require much maintenance they wanted more feedback and control. Specifically, they wished to be able to see exactly what the sensor was forwarding to public terminals and be able to turn off the sensor at any time. We integrated these changes into the second version of the sensor, creating a new window that allows users to monitor and control outgoing information (Figure 6). This is a significant change because whereas before users could start up their system and not know that the interest sensor is running in the background, the addition of an interface makes it observable. That is, users devote more of their attentional resources to it than we had intended. We expect that this trend may diminish over time, however.

Users also commented that the topics are often too common to be of interest (*e.g.* "notes"). Since the choosing algorithm in fact weights the least common interests most heavily, we believe this is due to the relatively small size of data accumulated by the display prior to our test and will improve with time.

In our current deployment, public displays and presence sensors are situated in rooms central to each group. At this point in the study, we have experience with integrating the hardware into the two environments. In particular, we have found it challenging to find a deployment area within environments that both blends in with daily activity and stands out enough to be distinguished from other equipment. In one space, the public display was situated on a desk frequently used by undergraduate students. For users near that space, the display stood out because it was new and took up space, thus changing an environment

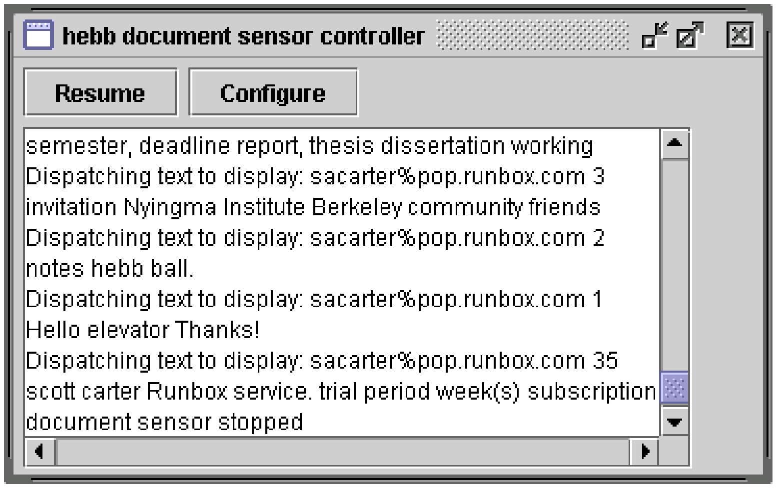

Fig. 6. Interest sensor controller showing history

that those users had grown familiar with. Other users not as familiar with that environment, though, tended not to notice it as much and thus not perceive changes in the display. This suggests that social practices have as much to do with the noticability of displayed items as the perceptual qualities of those items.

Also, users have found full document access via personal PDAs useful because it augments their recall of the context in which they used the displayed keywords. However, use of the display itself has revealed that users would find it useful to "take over" the display itself to access their documents but also to search for related information. This suggests that while the system may encourage awareness, it should provide more controls to support conversation.

6 Conclusions and Future Work

In this paper we reported on a series of interviews investigating collaboration and communication in small groups as well as a system developed in response to our findings from the interviews. Our contributions include the identification of the need for improved grounding and communication among loosely coupled, co-located groups, and the development of a system that can support such grounding by capturing and displaying shared interests. The sensor we designed captures interests from e-mail and displays them via peripheral displays positioned in areas routinely visited by group members. We adjusted the interest sensor to satisfy early use reports indicating that users required more control and feedback of content that the sensor discovered. Furthermore, a paper prototype of the peripheral display indicated that our original designs were too complex and high fidelity for the kind of impromptu interactions we were hoping to inspire. In response to this issue we simplified the layout and controls of the final display. We are in the process of testing this system, and early results show that users find the system useful when they have access to aids that help recontextualize displayed items.

7 Acknowledgments

References

1. V. Bellotti and S. Bly. Walking away from the desktop computer: Distributed collaboration and mobility in a product design team. In *Proceedings of ACM CSCW'96 Conference on Computer-Supported Cooperative Work*, pp. 209–218, 1996.
2. S. Borgatti, M. Everett, and L. Freeman. *Ucinet for Windows: Software for Social Network Analysis*, 2002.
3. D. Brock and J. G. Trafton. Cognitive representation of common ground in user interfaces. In *User Modeling*, 1999.
4. H. H. Clark. *Using Language*. Cambridge University Press, Cambridge, England, 1996.
5. O. de Bruijn and R. Spence. Serendipity within a ubiquitous computing environment: A case for opportunistic browsing. In *Ubicomp 2001: Ubiquitous Computing, Third International Conference*, volume 2201 of *Lecture Notes in Computer Science*, pp. 362–369. Springer, 2001.
6. P. Dourish, A. Adler, V. Bellotti, and A. Henderson. Your place or mine? learning from long-term use of audio-video communication. *Computer Supported Cooperative Work*, 5(1):33–62, 1996.
7. P. Dourish and S. Bly. Portholes: Supporting awareness in a distributed work group. In *Proceedings of ACM CHI'92 Conference on Human Factors in Computing Systems*, pp. 541–547, 1992.
8. L. Garton, C. Haythornthwaite, and B. Wellman. Studying on-line social networks. *Journal of Computer Mediated Communication*, 3(1), 1997.
9. S. Greenberg. Peepholes: Low cost awareness of one's community. In *Extended abstracts of ACM CHI '96 Conference on Human Factors in Computing Systems*, pp. 206–207, 1996.
10. H. Grice. Logic and conversation. In P. Cole and J. Morgan, editors, *Syntax and semtics*, volume 3. Academic Press, 1975.
11. C. Haythornthwaite. A social network theory of tie strength and media use: A framework for evaluating multi-level impacts of new media. Technical report, Graduate School of Library and Information Science, University of Illinois at Urbana-Champaign, Champaign, IL, 1999.
12. E. M. Huang, J. Tuillo, T. J. Costa, and J. F. McCarthy. Promoting awareness of work activities through peripheral displays. In *Extended Abstracts of ACM CHI'02 Conference on Human Factors in Computing Systems*, pp. 648–649, 2002.
13. Hubbub.
14. H. T. Ikujiro Nonaka, Hirotaka Takeuchi. *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation*. Oxford Press, 19945.
15. J. F. McCarthy, T. J. Costa, and E. S. Liongosari. UniCast, OutCast & GroupCast: Three steps toward ubiquitous, peripheral displays. In *Ubicomp 2001: Ubiquitous Computing, Third International Conference*, volume 2201 of *Lecture Notes in Computer Science*, pp. 332–345. Springer, 2001.

16. A. Monk. Common ground in electronically mediated communication: Clark's theory of language use. *To appear*, 2003.
17. B. Nardi, S. Whittaker, E. Isaacs, M. Creech, J. Johnson, and J. Hainsworth. Integrating communication and information through contactmap. *Communications of the ACM*, 45(4):89—95, 2002.
18. G. Olson. Tutorial on the evaluation of cscw systems. In *Conference on Computer Supported Cooperative Work*, 2003.
19. G. M. Olson and J. S. Olson. Distance matters. *Human-Computer Interaction*, 15(2/3):139–178, 2000.
20. G. M. Olson, J. S. Olson, M. R. Carter, and M. Storrosten. Small group design meetings: An analysis of collaboration. *Human-Computer Interaction*, 7(4):347–374, 1992.
21. J. S. Olson, G. M. Olson, and D. K. Meader. What mix of video and audio is useful for small groups doing remote real-time design work? In *Proceedings of ACM CHI'95 Conference on Human Factors in Computing Systems*, pp. 362–368, 1995.
22. J. S. Olson, G. M. Olson, M. Storrosten, and M. Carter. Groupwork close up: A comparison of the group design process with and without a simple group editor. *ACM Transactions on Information Systems*, 11(4):321–348, Oct. 1993. Special Issue on Computer-Supported Cooperative Work (CSCW).
23. Qtag. Web Page. http://web.bham.ac.uk/O.Mason/software/tagger/.
24. . J. S. R.E. Kraut, S. R. Fussell. Visual information as a conversational resource in collaborative physical tasks. *Human-Computer Interaction*, 18:13–18, 2003.
25. M. Rettig. Practical programmer: Prototyping for tiny fingers. *Communications of the ACM*, 37(4):21–27, 1994.
26. W. Sack. Conversation map: A content-based usenet newsgroup browser. In *Proceedings of the 2000 International Conference on Intelligent User Interfaces*, pp. 233–240, 2000.
27. G. Salton. Automatic text processing. In *The Transformation, Analysis and Retrieval of Information by Computer*. Addison-Wesley, 1988.
28. M. Schwartz and D. Wood. Discovering shared interests among people using graph analysis of global electronic mail traffic. *Communications of the ACM*, 36(8):78–89, 1992.
29. D. Snowdon and A. Grasso. Diffusing information in organizational settings: learning from experience. In *Proceedings of ACM CHI'02 Conference on Human Factors in Computing Systems*, pp. 331–338, 2002.
30. J. Tyler, D. Wilkinson, and B. A. Huberman. Email as spectroscopy: Automated discovery of community structure within organizations. In *International Conference on Communities and Technologies*, 2003.
31. S. Wasserman and K. Faust. *Social network analysis: methods and applications.* Cambridge University Press, 1994.
32. S. Whittaker, D. Frohlich, and O. Daly-Jones. Informal workplace communication: What is it like and how might we support it? In *Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systems*, pp. 131–137, 1994.