



Building Connections among Loosely Coupled Groups: Hebb's Rule at Work

S. CARTER¹, J. MANKOFF¹ & P. GODDI²

¹Electrical Engineering and Computer Science Department University of California at Berkeley, Berkeley, CA 94720-1770, USA (E-mail: sacarter@cs.berkeley.edu; jmankoff@cs.berkeley.edu); ²Hewlett-Packard Laboratories, 1501 Page Mill Road, Palo Alto, CA 94304, USA (E-mail: 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 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 a system 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. We deployed the system to two groups for two months and found that the system did lead to increased awareness of shared interests.

Key words: CSCW, social computing, peripheral displays, sensors

1. Introduction

When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased.

D.O. Hebb *in* The Organization of Behavior (Hebb, 1949)

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. We implemented these techniques and deployed them over 2 months. Our contributions include the identification of the need for improved grounding and communication among loosely coupled groups, and the development and analysis of a deployed 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, work-related conversations (Olson and Olson, 2000). 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. In particular, we noticed few collaborations between loosely coupled groups, or those that spend a majority of their working day in separate physical environments, but whose members occasionally share public spaces. Members of such groups may share interests, but because of their infrequent interaction those interests may go unnoticed, limiting opportunities for communication and collaboration.

In his analysis of conversations, Clark (1996) suggests that conversation participants frame their discussions in terms of what all parties understand to be common knowledge, context or shared experiences. 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 awareness 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, work-related relationships. Displays use context information derived from sensors to present interests shared among those near some display.

In this paper we first describe previous work regarding 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. We conclude with a analysis of a field experiment with the deployed system.

2. The informal exchange of knowledge

It is important to support informal communication as it helps people establish common ground necessary for meaningful conversations and relationships. Common ground, as Clark (1996) defines it in his book, *Using Language* (1996), 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”. 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 (Olson et al., 1992, 1993, 1995). In particular, collocated groups are more likely to experience activities collectively, which, according to Clark, is “the most important source of common ground” (Monk, 2003). Furthermore, the theory of common ground has been shown to be valuable in other collaborative settings, such as repair tasks (Kraut and Fussell, 2003).

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 (Nonaka and Takeuchi, 1945). For example, members of two different groups may have entirely different notions about a term or process, such as what it means to conduct a study (e.g., a new member of a working group may show up to meetings on the hour, while seasoned members may be aware that meetings never start until ten past the hour). Or they may not share each other’s informal understandings of information access and elicitation (e.g., in a particular group there may be an official process of obtaining information on a particular topic, but in actuality everyone simply asks a particular informant). 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 (Whittaker et al., 1994; Bellotti and Bly, 1996). 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 (Bellotti and Bly, 1996). 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 (de Bruijn and Spence, 2001), technology could potentially help to support grounding in such situations. Many technologies have been created to support media connections between spaces (Dourish and Bly, 1992; Dourish et al., 1996; Greenberg, 1996). 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 (McCarthy et al., 2001). Snowdon’s and Grasso’s (2002) CWall published articles captured from a community web site to an interactive public display. Huang et al.’s (2002) Awareness Module and Greenberg’s and Rounding’s (2001) Notification Collage present items posted by users in a group to a public displays. Okamoto et al.’s (2002) Silhouettell system connected distributed groups spread across different cultures. However, none of those were deployed specifically to increase ties between different loosely coupled work groups, instead concentrating on strongly collocated groups, broad communities or groups distributed across the world. 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 laboratories, 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.

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 respondents longest associated with the group. We asked participants open-ended questions about their collaboration and communication habits, including the following:

- Describe where you might go, some typical activities you might engage in and who you might meet during a typical day.
- What information do you share with people publicly, e.g., on a communal whiteboard outside of a meeting room? How comfortable do you feel about doing that?
- Do you hold meetings with your advisor/boss? What do you discuss?
- How have you gone about finding other people who share your interests? Please describe some situations in which you have talked face-to-face with another person about a newly discovered shared interest.

Other 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, are often difficult to analyze. Often interview data is transcribed and cataloged but utilized only in *ad hoc* references (Olson, 2003). 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 (Wasserman and Faust, 1994), 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 (Garton et al., 1997;

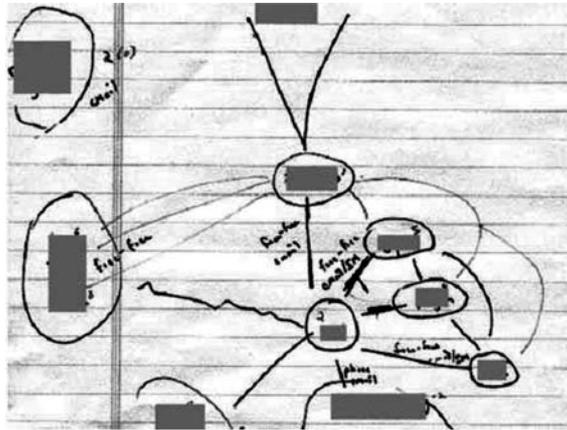


Figure 1. Part of a drawing made by a participant during an interview. Notice the “information funnel” at the top, as well as how the interviewee circled and linked people and groups. Names have been removed to maintain anonymity.

Sack, 2000; Nardi et al., 2002; Tyler et al., 2003). 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 (Haythornthwaite, 1999). 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

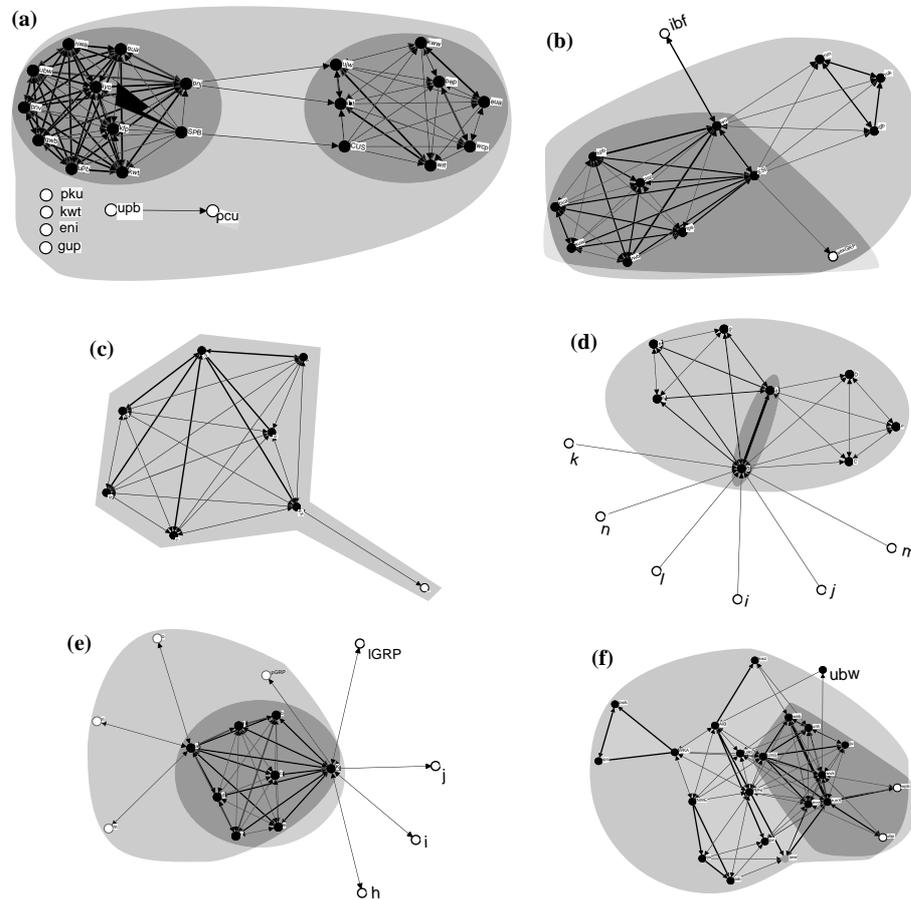


Figure 2. Sociograms for all groups. Gray areas are demarcations of organizational units, the darkest area in each sociogram indicating the boundary of the group that we interviewed. In all cases save one (graph c), members of the core group identified other influential individuals as being part of one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

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 (Borgatti et al., 2002) 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 (Haythornthwaite, 1999). 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 discovering shared interests and starting a conversation about them. People are unlikely to explore some new topic, however, unless they share some common ground (Grice, 1975). 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 collaboration in loosely coupled 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. Our system includes a public, peripheral display, a private, focal display, presence sensors and an implicit interest sensor. Our iterations included paper prototypes of the public and private displays as well as evaluations of interactive versions of the interest sensor. 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 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 (see Figure 3a).

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

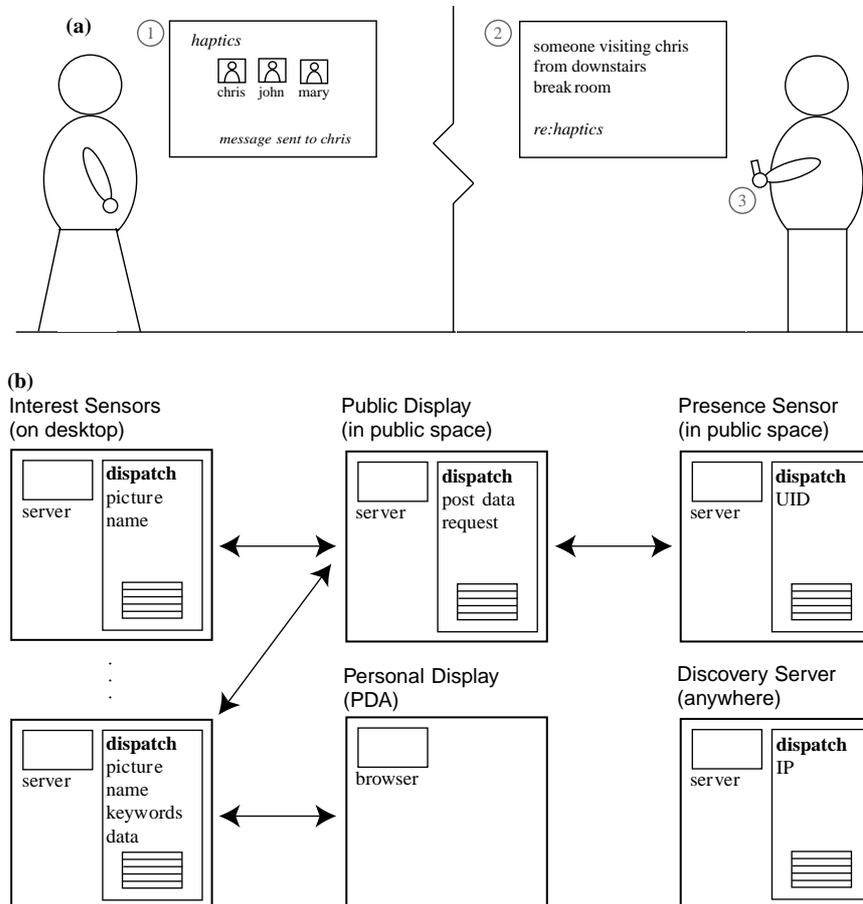


Figure 3. (a) System use scenario. (1) Mary (left), downstairs, uses the public display to notify Chris (right), upstairs, that she would like to discuss haptics. (2) Chris sees a message on the public display near him that someone will be visiting soon. (3) While waiting for Mary, Chris uses his PDA to view the document of his that the hebb system used to identify this shared interest. (b) System architecture. Components send events they generate to other registered components.

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)”. Then (1) 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’s picture, recognizes him, clicks on it, and starts to walk to the breakroom upstairs. (2) 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.” (3) 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 how all of the components interact and then detail the design of each individual component.

4.2.1. *Overview*

Components used in this system include interest sensors, public and private displays and presence sensors (see Figure 3b). Each component can register with another component to receive semantically tagged data that component generates. The interest sensor monitors a user's e-mail and generates keyword events (generated from the algorithm described in the next section) tagged with her name and picture as well as encrypted full document data (for use on personal PDAs). The public display notifies other components indicating from which documents keywords were recently displayed and the presence sensor senses users in the space via either RFID badging or presence of the user's PDA on the local wireless network. When a component generates an event, it sends it to all components registered with a local server. A discovery server allows components to update and retrieve location information.

4.2.2. *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 (Nonaka and Takeuchi, 1945). Thus, the interest sensor should be designed either to capture spontaneous meetings or e-mail or both.

We considered two means of using microphones to capture interests generated in spontaneous meetings: a direct audio link and voice recognition with visual display (Hubbub, 2004). 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 (Schwartz and Wood, 1992; MacArthur and Bruza, 2003). 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 90s on this account. For each message received (and sent, if available) the sensor first eliminates commonly used generic terms and spam 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 (2004). 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 others who are potentially unknown to help build relationships. Therefore, after discussions with interview subjects we derived a solution in which the interest sensor releases 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. 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 (Salton, 1988). 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.

4.2.3 *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 software we developed can run on any standard resolution, interactive display.

The public display (Figure 4) 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 and displays each discovered shared interest for at least five minutes (more if the presence sensor does not report a change in nearby users). The display notifies users of updates by alternating background colors as well as orientations of the displayed interests. 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.

We also developed lightweight interactive features to help start communication about some topic when not all of the involved 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 use scenario in Section 4.1, 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).

4.2.4 *Personal Display Design*

Personal PDAs allow each user to see the full contents of documents from which displayed keywords were derived and to browse through previous topic postings. 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

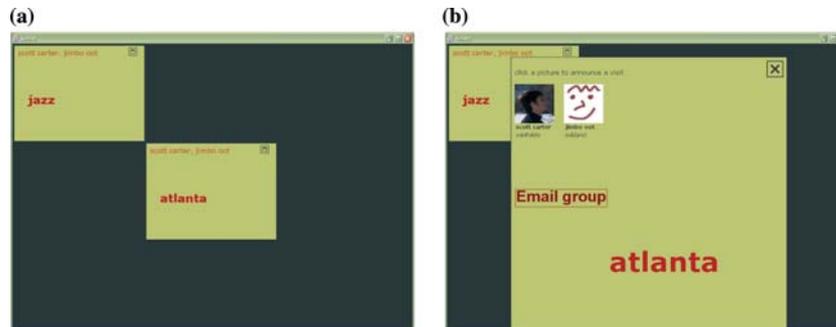


Figure 4. The public display. The display updates with between 1-5 topics (a) when it receives new data from an interest or presence sensor. When a user clicks on one of the topic areas, the display shows a maximized view of that topic (b).

data on their personal PDA via a Web site associated with the public display. The Web site lists links to all documents belonging to a specific user and related to a shared interest that recently appeared on the public display.

5. Field experiment

The success of the system we have developed depends on the extent to which it encourages awareness and communication between group members about topics that they otherwise would not have explored. To test this, we deployed 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 followed up this pilot deployment with another deployment of the full system that connects two academic research labs. The full deployment ran for two months.

5.1. PILOT DEPLOYMENT

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 wanted the ability to see exactly what the sensor was forwarding to public terminals and 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 5). 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. Thus, users want to devote more of their

attentional resources to it than we had intended. We expect that this trend may diminish over time, however.

5.2. FULL DEPLOYMENT

After the initial pilot, we deployed the system to two academic research groups working on similar topics. The groups were spread across three different spaces: one group of three members was collocated (in Table I, space P) while the other group of four was split between two spaces (in Table I, spaces S and J). This particular arrangement allowed us to determine how well the system supports intra- and inter-group awareness and communication. We deployed the interest sensor first for 4 weeks to establish common e-mail patterns and then deployed the rest of the system for another 4 weeks. We deployed a total of three public displays, all in work areas, and a presence sensor along with each display. We logged all user interactions with the public and private display. We also monitored and interviewed user groups to determine use that does not appear in logs. This could occur in two ways. Firstly, our access log data could not ascertain the extent to which participants used the public display to maintain peripherally an awareness of other user's interests. Secondly, we could not record instances in which users verbally discussed something that they noticed on the display without interacting with the system in any way. In addition, the interviews helped us understand how users adapted to the system.

6. Results and discussion

Quantitative measures of interaction with the system as well as qualitative reports from participants showed that the system helped increase awareness

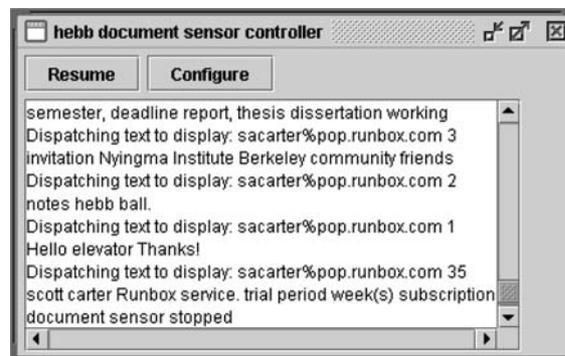


Figure 5. Interest sensor controller. In the main window the sensor lists exactly what words and phrases were sent to public displays. Users can stop or reconfigure the display at any time.

Table I. User interaction data. One group of three members was collocated in space P while the other group of four was split between spaces S and J. All three spaces were in separate buildings.

	Space S	Space J	Space P
Public display topic updates	396	412	350
Public display user accesses	29	45	67
Private display user accesses	19	41	47
Private display user topic requests	12	27	25
Person-hours in space per day	10	14	17

and communication between participants using this system. Our field experiment also led to other key observations: that users adapt to implicit sensing systems, peripheral displays need their own dedicated space and that for peripheral displays object onset captures attention and minimizes annoyance.

6.1. USER INTERACTION WITH DISPLAYS

The public and private displays logged updates and user interactions. In Table I, public display topic updates refers to when the public display refreshed with new information, public display user accesses refers to when users actually used the public display to access more information about a topic, private display user accesses refers to when a user loaded the private display web page, private display user topic requests refers to when a user accessed contextualizing information about a topic on a private display, and person-hours in space per day is equal to the sum of the total number of hours each group member was in the space per day. We found that participants often used the personal display and expanded displayed sections of the public display when they were not familiar with the person involved in a relationship (we derived familiarity from the pre-study logs of e-mail correspondence). Also, note that the extent to which participants used the system depended directly upon the number of person-hours spent in the space.

6.2. USER COMMUNICATION CHANGES

Users tended to e-mail each other more regularly after the deployment than before, indicating increased connections. E-mail has been shown to

approximate social connections in work groups (Tyler et al., 2003). From e-mail recipient information logged by the interest sensors we created sociograms (Figure 6) that show e-mail habits between study participants for a month before versus during the month-long deployment of the displays. In the sociograms, nodes represent users, shaded areas spatial and organizational relationships between users and lines e-mail communication between users. Also, line thickness varies with e-mail correspondence rate: a thin line represents 1-3 e-mails, medium lines 4-7 e-mails and thick lines 8 or more e-mails sent between participants. Figure 6 compares sociograms showing pre-study and within-study e-mail rates between the two groups; there are both more and stronger connections in the later figure, especially between groups in different organizations.

Our data showed that users who contacted others did not necessarily do so directly after expanding a topic. Rather, discussions with users revealed that they often sent e-mails about a topic after having noticed something that interested them *and* after they felt they were at a “stopping point” on whatever task they were working on at the time. The edges in Figure 6(b) not present in Figure 6(a) were largely due to mutual correspondence occurring after such an incident.

We interpret these results to mean that our system increased correspondence between the two groups. However, these effects could be due to natural fluctuation in the rhythms of e-mail use or to the fact that group members had established common-ground merely by jointly participating in the study. The latter issue is one common to field deployments: it is inevitable that a group changes with the introduction of a system. These issues make it harder to evaluate the system’s effect. However, we believe that this effect would not solely cause the extent of the increase in communication represented in the sociograms. We justified this claim by exploring why e-mail rates increased for some participants. In particular, we found that e-mails tended to increase after an e-mail was sent between participants regarding a topic posted to the public and private displays. We explored such a situation involving two participants that began e-mailing more after one participant noticed a topic on the display that he had in common with another. Figure 7 shows that while both participants were accessing content on the hebb system throughout the deployment, they did not begin e-mailing each other until after one of them acted on a common interest shown on a hebb display. Follow-up interviews with the participants revealed that they did indeed discuss this interest via e-mail.

6.3. USER ADAPTION TO IMPLICIT SENSING

When asked to compare the hebb system to other methods of increasing collaboration and awareness between groups such as group meetings, e-mail

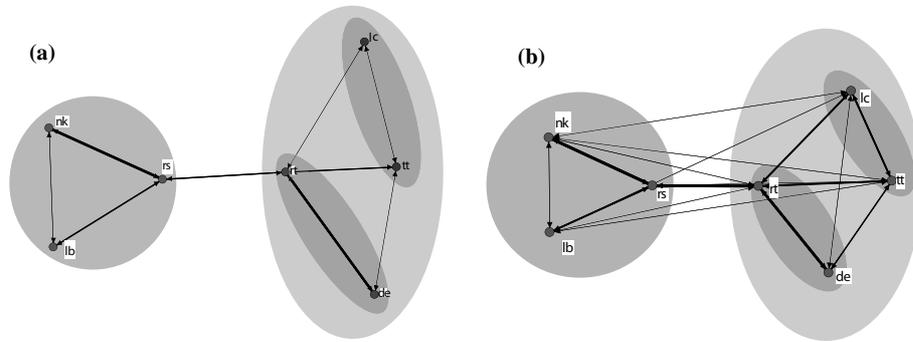


Figure 6. Sociograms of communication between participants based on e-mail (a) pre- and (b) post-deployment. The thickness of the lines indicates amount of e-mail traffic (a thin line represents 1-3 e-mails, medium lines 4-7 and thick lines 8 or more over the deployment period). Areas are shaded according to the process described in Section 3.4.

lists, newsgroups and community displays, participants reported that the benefit of implicit sensing was more important than potential privacy tradeoffs and overload. In particular, participants reported a strong desire to have a system that “displays peripheral information peripherally” and seemed exasperated at the prospect of subscribing to another e-mail list or to post content explicitly. Furthermore, users discussed a desire for the public display to show sentences from personal e-mail rather than only words and phrases. These results were particularly surprising given that early attempts to deploy the system met with ridicule precisely because the implicit sensing component was considered intrusive and the public display component a violation of privacy. Some users did, however, argue for a lightweight control that would allow them to flag particular phrases for capture when composing an e-mail. Thus, it is essential to allow for user adaptation when assessing the value of systems that use implicit sensing.

6.4. SPATIAL LOCATION OF PERIPHERAL DISPLAYS

Our interviews revealed that most participants did not carry PDAs with them because they found them cumbersome to carry and the Internet browser

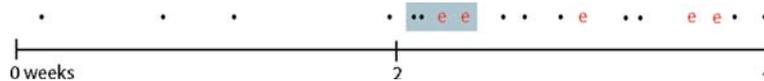


Figure 7. E-mails following content accesses for two users during the second half of the deployment. Circles represent public and private display content accesses by two participants and “e”s represent e-mail sent between the two participants. The shaded area represents hebb postings and e-mails that were related to one another.

“awkward to use”. Because we designed the private display as a web-based UI, a quarter of the way into the deployment we gave participants the option of accessing their personal display via a web page bookmarked in their favorite Internet browser on their regular desktop machine. We found that about half of the users used the desktop version, but some users still used the PDA as a personal display. However, those who used the PDA used it neither as a mobile device, as we had expected, nor as a display that they could monitor peripherally from their desktop. Rather, they set it aside and accessed it when they noticed something on the public display that intrigued them. When we asked one participant why he chose to use the display that way instead of via the desktop, he said that the “desktop is for daily work stuff. . . it doesn’t belong there”. This statement reflected the feelings of other such users. Thus, some users seemed to have developed a mental model of a peripheral system as one that must be physically separate from their daily work environment.

6.5. OBJECT ONSET AND ATTENTION CAPTURE

Early log analysis showed extremely sparse interaction with the displays. Interview questions probing this issue revealed that participants would often not perceive updates on the display even when new phrases and names appeared on the display. To remedy this we modified the public display 2½ weeks into the second half of the deployment to switch positive and negative space on the display after each update (approximately every 15 min). That is, we alternated the blank and filled spaces in the display shown in Figure 4. Follow-up interviews revealed this change to be significant enough to raise participant’s awareness of changes on the display. This result follows from research showing object onset to be among the best means of capturing attention (Yantis, 1993). Interestingly, though, interviewees reported this option to be much preferable to motion, also a powerful means of capturing attention. Specifically, interviewees reported that they could see the update but that the display did not “demand attention like. . . [Microsoft] Clippy”. Thus, visual onset in peripheral displays as a means to capture attention may strike a balance between awareness and annoyance.

6.6. SUMMARY

Use of the hebb system led to increased awareness and communication between participants. Observations and interviews also revealed surprising adaptations to the system, including changing tradeoffs between privacy and implicit sensing as well as unexpected use of personal displays as task-specific data access appliances. Lastly, a revealed tradeoff between peripheral

noticing and interruption led to our adapting the public display to utilize object onset but not animation.

7. Conclusions and future work

In this paper we reported on a series of interviews investigating collaboration and communication in small groups as well as the development and analysis of a system designed in response to our findings from the interviews. Our contributions include the identification of the need for improved grounding and communication among loosely coupled groups, and the development, deployment and analysis 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 public displays positioned in areas routinely seen by group members as well as private displays for each participant. Based on a pilot study, we adjusted the interest sensor to satisfy early use reports indicating that users required more control and feedback of content that the sensor discovered. Analysis of our final deployment showed that the system encouraged more communication between test groups. Also, interviews we conducted further revealed that users in this study adapted their attitudes towards implicit sensing systems, becoming more apt to release private information to gain benefits from the system over time. Our interviews also found that peripheral displays may need their own dedicated space rather than taking up space on a desktop. Finally, we found that one means of balancing attention capture and interruption in peripheral displays is to use object onset rather than animation.

In future work we plan to continue field deployments to explore how users continue to adapt to the system. Also, as we found user's reflections regarding their use of the system in the field experiment to be most valuable, we are also interested in conducting participatory brainstorming sessions to reveal new features for and potential uses of the hebb system.

Acknowledgements

Hewlett-Packard and grants from the NSF (IIS-0205644 and IIS-0209213) funded this work. Thanks also to Dr. Anind Dey, Dr. Peter Lyman, Christine Ro and Mimi Yang.

References

- Bellotti, V. and S. Bly (1996): 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.

- Borgatti, S., M. Everett and L. Freeman (2002): Ucinet for Windows: Software for Social Network Analysis.
- Clark, H.H. (1996): *Using Language*. Cambridge, England: Cambridge University Press.
- de Bruijn, O. and R. Spence (2001): Serendipity within a Ubiquitous Computing Environment: A Case for Opportunistic Browsing. In *UbiComp 2001: Ubiquitous Computing, Third International Conference*, vol. 2201 of *Lecture Notes in Computer Science*. pp. 362–369.
- Dourish, P., A. Adler, V. Bellotti and A. Henderson (1996): Your Place or Mine? Learning from Long-Term Use of Audio-Video Communication. *Computer Supported Cooperative Work* vol. 5, no. 2, 33–62.
- Dourish, P. and S. Bly (1992): Portholes: Supporting Awareness in a Distributed Work Group. In *Proceedings of ACM CHI'92 Conference on Human Factors in Computing Systems*. pp. 541–547.
- Garton, L., C. Haythornthwaite and B. Wellman (1997): Studying On-Line Social Networks. *Journal of Computer Mediated Communication* vol. 3, no. 1.
- Greenberg, S. (1996): 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.
- Greenberg, S. and M. Rounding (2001): The Notification Collage: Posting Information to Public and Personal Displays. In *Proceedings of ACM CHI 2001 Conference on Human Factors in Computing Systems*. pp. 514–521.
- Grice, H. (1975): Logic and Conversation. In P. Cole and J. Morgan (eds.): *Syntax and semantics*, vol. 3. Academic Press, pp. 41–58.
- Haythornthwaite, C. (1999): 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.
- Hebb, D.O. (1949): *The Organization of Behavior: A neuropsychological theory*. Wiley, New-York
- Huang, E.M., J. Tuillo, T.J. Costa, and J.F. McCarthy (2002): 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.
- Hubbub: 2004. Web Page. <http://www.gvu.gatech.edu/people/sha.xinwei/topologicalmedia/hubbub/>.
- McArthur, R. and P. Bruza (2003): Discovery of Implicit and Explicit Connections Between People Using Email Utterance. In *Proceedings of the European Conference on Computer Supported Cooperative Work*. pp. 21–40.
- McCarthy, J.F., T.J. Costa and E.S. Liongosari (2001): UniCast, OutCast & GroupCast: Three Steps Toward Ubiquitous, Peripheral Displays. In *UbiComp 2001: Ubiquitous Computing, Third International Conference*, vol. 2201 of *Lecture Notes in Computer Science*. pp. 332–345.
- Monk, A. (2003): Common Ground in Electronically Mediated Communication: Clark's Theory of Language Use. In J. Carroll (ed.): *HCI models, theories and frameworks: towards a multidisciplinary science*. Morgan Kaufmann Publishers, pp. 265–289.
- Nardi, B., S. Whittaker, E. Isaacs, M. Creech, J. Johnson, and J. Hainsworth (2002): Integrating communication and information through ContactMap. *Communications of the ACM* vol. 45, no. 4, 89–95.
- Nonaka, I. and H. Takeuchi (1995): *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation*. Oxford Press, Oxford.
- Okamoto, M., K. Isbister, H. Nakanishi and T. Ishida (2002): Supporting Cross-Cultural Communication with a Large-Screen System. *New Generation Computing* vol. 20, no. 2, 165–185.

- Olson, G.M. (2003): Tutorial on the Evaluation of CSCW Systems. Presented at the Conference on Computer Supported Cooperative Work.
- Olson, G.M. and J.S. Olson (2000): Distance Matters. *Human-Computer Interaction* vol. 15, no. 2/3, 139–178.
- Olson, G.M., J.S. Olson, M.R. Carter and M. Storrosten (1992): Small Group Design Meetings: An Analysis of Collaboration. *Human-Computer Interaction* vol. 7, no. 4, 347–374.
- Olson, J.S., G.M. Olson and D.K. Meader (1995): 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.
- Olson, J.S., G.M. Olson, M. Storrosten and M. Carter (1993): Groupwork Close Up: A Comparison of the Group Design Process With and Without a Simple Group Editor. *ACM Transactions on Information Systems* vol. 14, no. 4, 321–348. Special Issue on Computer-Supported Cooperative Work (CSCW).
- Qtag: 2004. Web Page. <http://web.bham.ac.uk/O.Mason/software/tagger/>.
- Kraut, R.E., Fussell, S.R.J.S. (2003): Visual information as a conversational resource in collaborative physical tasks. *Human-Computer Interaction* vol. 18, no. 1, 13–18.
- Sack, W. (2000): Conversation Map: A Content-Based Usenet Newsgroup Browser. In *Proceedings of the 2000 International Conference on Intelligent User Interfaces*. pp. 233–240.
- Salton, G. (1988): *Automatic Text Processing: The Transformation, Analysis, and Retrieval of Information by Computer*. Addison-Wesley, Reading, MA.
- Schwartz, M. and D. Wood (1992): Discovering Shared Interests Among People Using Graph Analysis of Global Electronic Mail Traffic. *Communications of the ACM* vol. 36, no. 8, 78–89.
- Snowdon, D. and A. Grasso (2002): Diffusing information in organizational settings: learning from experience. In *Proceedings of ACM CHI'02 Conference on Human Factors in Computing Systems*. pp. 331–338.
- Tyler, J., D. Wilkinson and B.A. Huberman (2003): Email as Spectroscopy: Automated Discovery of Community Structure within Organizations. In *Proceedings of International Conference on Communities and Technologies*.
- Wasserman, S. and K. Faust (1994): *Social Network Analysis: Methods and Applications*. Cambridge University Press, Cambridge.
- Whittaker, S., D. Frohlich and O. Daly-Jones (1994): 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.
- Yantis, S. (1993): Stimulus-Driven Attentional Capture. *Current Directions in Psychological Science* vol. 2, (no. 5) 156–161.