

Social Translucence: An Approach to Designing Systems that Support Social Processes*

[Thomas Erickson](#) and Wendy A. Kellogg

IBM T.J. Watson Research Center

snowfall@acm.org, wkellogg@us.ibm.com

Abstract

We are interested in designing systems that support communication and collaboration among large groups of people over computer networks. We begin by asking what properties of the physical world support graceful human-human communication in face to face situations, and argue that it is possible to design digital systems that support coherent behavior by making participants and their activities visible to one another. We call such systems "socially translucent systems" and suggest that they have three characteristics -- visibility, awareness, and accountability -- which enable people to draw upon their social experience and expertise to structure their interactions with one another. To motivate and focus our ideas we develop a vision of knowledge communities, conversationally-based systems that support the creation, management and re-use of knowledge in a social context. We describe our experience in designing and deploying one layer of functionality for knowledge communities, embodied in a working system called "Babble," and discuss research issues raised by a socially translucent approach to design.

Categories and Subject Descriptors: H.1.2 [Models and Principles]: User/Machine systems: *human factors; human information processing*; H.5.2 [Information Interfaces and Presentation]: User Interfaces: *Graphical User Interfaces (GUI); Theory and Methods*; H.5.3 [Information Interfaces and Presentation]: Group and Organization Interfaces: *Asynchronous Interaction; Collaborative computing; Computer-supported cooperative work; Organizational design; Synchronous interaction; Theory and models*; K.4.3 [Computers and Society]: Organizational Impacts: *Computer-supported Collaborative Work*

General Terms: Design; Human Factors

Additional Keywords and Phrases: Computer Mediated Communication; CMC; CMI; CSCW; Social Computing; Social Navigation; Social Visualization; Visualization

1. Introduction

We are concerned with designing systems to support communication and collaboration among large groups of people over computer networks. We are particularly interested in the question of how to design such systems so that they support communication that is deep, coherent and productive.

It seems evident that our digital tools are primitive and cumbersome. Most of the communication and collaboration tools in daily use could be implemented, with little loss of functionality or usability, on the time-shared mainframe and character-based terminal systems of three decades ago. This is not because we managed to get things right on the first try: email, mailing lists, chat rooms, discussion databases, etc., all have substantial

shortcomings. There are a host of problems: addressing; managing threads; bringing other people into the middle of a conversation; keeping a conversation on track; knowing who (or whether) anyone is listening; getting people to respond in a timely manner; finding old messages with crucial information in them; etc. It is difficult to conduct a long-running, productive conversation through the digital medium, especially if there are more than a few people involved. Yet we do not believe that these shortcomings are inherent in digital systems.

The difficulty of digital communication and collaboration stands in stark contrast to our ability to communicate and collaborate with one another in the physical world. Consider these ordinary situations:

In another town on business, you and a few colleagues are looking for a place to have dinner. You notice a small restaurant: through its window you see a cozy room with waiters bustling about; you hear the murmur of conversation, and the clink of glasses and cutlery. You head for the entrance...

You've arrived at the opening reception for a convention. You look around for someone to talk to and see someone you recognize gesturing excitedly as others listen intently. Curious, you wander over...

You're shopping for wine to bring to dinner. As you browse the racks you hear a muttered 'Aha!', and watch another shopper grab two bottles out of a nearly empty bin. You get a bottle for yourself...

These examples are unremarkable. Every day we make countless decisions based on the activity of those around us. As social creatures, we are immersed in a sea of social information. We've evolved an exquisite sensitivity to the actions and interactions of others. Whether it's wrapping up a talk when the audience starts fidgeting, or deciding to forego the grocery shopping because the parking lot is jammed, social information like this provides the basis for inferences, planning, and coordination of activity.

In all of this, the only thing that is remarkable is how radically things change when we move from the ordinary world into the world of digital systems. Digital systems are generally opaque to social information. Most of our knowledge about people, most of our attunement to their interactions, most of our facility for improvising in a changing situation, goes unused. In the digital world we are socially blind.

In this paper our principal goal is to make some progress in redressing this blindness. We begin by describing the notion of social translucence, an approach to designing digital systems that emphasizes making social information visible within the system. Next, to focus and motivate the concept of social translucence, we apply it to knowledge management and sketch the outlines of what we call knowledge communities. Then we ask what it means to actually implement social translucence in a digital system and describe relevant work within HCI, and our own implementation efforts. Finally, we describe some of the most pressing research issues raised by a socially translucent approach to systems design.

2. Foundations: Social Translucence and The Architecture of interaction

As designers of communication and collaboration systems, we find ourselves taking inspiration from work in the areas of architecture and urban design. This is not surprising since, like architects and urban designers, we are concerned with creating contexts that support various forms of human-human interaction. What architecture and urbanism have to offer is long experience in exploring the interrelationship between physical spaces and social interaction -- the interested reader should see Alexander, et al. [1977], Gehl [1980], Jacobs [1961], Lynch [1990] and Whyte [1988]. In what follows, we primarily draw upon our own examples. Although we have learned much from architecture and urbanism, the fact is that designers in those domains can assume the existence of a consistent and unquestioned physics that underlies social interaction. There is no such constancy in the digital world, and so our goal in this section is to look deeply at social interaction as it is embedded in physical space, and try to extract principles that are sufficiently abstract that they might be transposed to the digital realm.

2.1 Visibility, Awareness, and Accountability

In the building where our group works there is a door that opens from the stairwell into the hallway. This door has a design problem: opened quickly, it is likely to slam into anyone who is about to enter from the other direction. In an attempt to fix this problem, a small sign was placed on the door: it reads, 'Please Open Slowly.' As you might guess, the sign is not a particularly effective solution.

Let's contrast this solution with one of a different sort: putting a glass window in the door. The glass window approach means that the sign is no longer required. As people approach the door they see whether anyone is on the other side and, if so, they modulate their actions appropriately. This is a simple example of what we call a socially translucent system. While it is obvious why this solution works, it is useful to examine the reasons behind it carefully. We see three reasons for the effectiveness of the glass window:

- First, the glass window makes socially significant information *visible*. That is, as humans, we are perceptually attuned to movement and human faces and figures: we notice and react to them more readily than we notice and interpret a printed sign.
- Second, the glass window supports *awareness*: I don't open the door quickly because *I know* that you're on the other side. This awareness brings our social rules into play to govern our actions: we have been raised in a culture in which slamming doors into other people is not sanctioned.
- There is a third, somewhat subtler reason for the efficacy of the glass window. Suppose that I don't care whether I hurt others: nevertheless, I'll open the door slowly because *I know that you know that I know* you're there, and therefore I will be held *accountable* for my actions. (This distinction is useful because, while accountability and awareness usually co-occur in the physical world, they are not necessarily coupled in the digital realm.) It is through such individual feelings of accountability that norms, rules, and customs become effective mechanisms for social control.

We see these three properties of socially translucent systems -- visibility, awareness, and accountability -- as building blocks of social interaction. Notice that, as exemplified by our other examples, social translucence is not *just* about people acting in accordance with social rules. In socially translucent systems we believe it will be easier for users to carry on coherent discussions; to observe and imitate others' actions; to engage in peer pressure; to create, notice, and conform to social conventions. We see social translucence as a fundamental requirement for supporting all types of communication and collaboration.

2.2 Translucence and the Power of Constraints

There is one other aspect of social translucence that deserves mention. Why is it that we speak of socially *translucent* systems rather than socially *transparent* systems? Because there is a vital tension between privacy and visibility. What we say and do with another person depends on who, and how many, are watching. Note that privacy is neither good nor bad on its own--it simply supports certain types of behavior and inhibits others. For example, the perceived validity of an election depends crucially on keeping certain of its aspects very private, and other aspects very public. As before, what we are seeing is the impact of awareness and accountability: in the election, it is desirable that the voters *not* be accountable to others for their votes, but that those who count the votes be accountable to all.

It would be a mistake, however, to think that translucence is only about the tension between privacy and visibility. Rather, translucence stands in more generally for the power of constraints. To see this, let's look at a more complex example:

A group of thirty people -- authors of the chapters of a forthcoming book -- had gathered to mutually critique chapters. For three days, each author worked with a small group of about six others. At the end of the workshop

the organizers decided to try to collectively create an organization for the book. The authors, none of whom had read all the chapters, would decide on what the book sections should be, and how the chapters should be ordered.

Everyone gathered in a room, each author with a copy of his or her chapter. To start the process, pieces of paper with possible section names had been placed on the floor, and authors were asked to put their chapters near appropriate sections. After this, the procedure was simple: anyone could pick up any chapter and move it elsewhere; anyone could change the name of a book section; anyone could propose a new section by writing a name on a new piece of paper.

Although the ensuing process was characterized by a lot of milling about and simultaneous conversations, it was exceptionally effective. In half an hour the group had arrived at an organization for a book of 30 chapters, with everyone participating in the discussion.

What is of interest here is how the spatial nature of the setting enabled what was, in effect, a process of social computation. First, as in the case of the door, the participants could see what was happening, and thus awareness and accountability came into play. For example, when someone went to move a chapter to another area of the room (i.e., move it to another section of the book), there would usually be one or more people around. Although not required by the rules, what happened was that the mover would politely offer a rationale for moving the chapter to those in the vicinity (hereafter the 'on-lookers'), thus triggering a discussion about the purpose of that section and the point of the chapter. The consequence of this discussion was that either:

- the mover and the on-lookers would agree on the move
- the on-lookers would convince the mover that the chapter was indeed in the right place
- the mover and on-lookers would decide to change the name and definition of the section so that the chapter fit the section better

In each of these cases the result was that there was a greater shared understanding of the section names and definitions, the gist of each chapter, and the rationale for the chapter's inclusion.

In addition to the awareness and accountability brought into play by the visibility of the activity of moving chapters around in the room, another spatial property played an important role: physical constraints. Constraints shaped the way in which people could participate in the process of organization. The fact that the chapters and section names were spread all over the room had an important impact: it meant that no one person could dominate the organization of the book. Those who had strong opinions about where their chapters belonged tended to hover near their chapters, ready to 'defend' their chapters' positions against would-be reorganizers. In contrast, those who had ideas about the arrangement of the book as a whole had to flit about from section to section, thus giving up any strong control over where their chapters (or any single chapter) were positioned. Similarly, people who stayed near a single section heading gained, over time, a detailed understanding of the rationale for the section as a result of repeatedly participating in the to-move-or-not-to-move discussions for its component chapters. This regulation of activity came as a side effect of the fact that the ability to hear and see in a crowded room decreases as distance increases; that is, the space is *translucent* (not transparent) to vision, speech, and hearing.

Note that it was not simply the existence of the constraints that were important; in addition, as with the visibility of socially significant information, two other levels are of importance. First, it was important that *people were aware of the existence and nature of the constraints*. This awareness means that the participants were able to anticipate the ways in which the constraints structured the group's interaction and adjust their own actions accordingly. Thus, based on the amount of ambient noise, speakers adjusted the volume of their speech so that they could be heard by those to whom they spoke. Awareness of and experience with the physics of real world interaction enabled smooth interaction among the group. In situations where an awareness of certain constraints is lacking (e.g., if a participant has a hearing aid that doesn't cope well with high levels of ambient noise), the interaction may break down, with people's communicative acts failing unexpectedly, and requiring joint action to

detect and repair failures.

Second, it was important that *participants were aware of the others' awareness of the constraints*. Thus, in the situation described, there was a generally shared awareness that people on one side of the room were unable to see or hear discussions on the other side of the room. Everyone understood, by virtue of their common experience with the physics of human interaction, what was going on. Thus, while I might be held accountable for moving your chair if you were standing nearby and could have easily been consulted, it was a different matter if you were on the other side of the room where you knew I could not have seen you. That is, not only do constraints serve to structure interaction, but the existence of a *shared awareness of constraints* is also a resource for structuring interaction.

These distinctions among the existence of constraints, participants' individual awareness of constraints, and shared awareness of constraints is important because, although these things are usually bundled together in physical environments, they are not necessarily coupled in digital ones. In digital environments it is rarely evident what the constraints are, nor whether the constraints are necessarily shared. The fact that I can hear you speaking over a speaker phone does not necessarily mean that you can hear me if I speak (half duplex lines), nor does my ability to send you email necessarily imply that you can reply to me. The assurances of communicative symmetry that arise out of our experience in unmediated physical reality are not automatically present in the digital realm.

2.3 Summary

We've discussed two dimensions of social translucence. First, a system that makes social information visible enables participants to be both aware of what is happening, and to be held accountable for their actions as a consequence of public knowledge of that awareness. Second, people also have a sophisticated understanding of the physics that underlie the visibility of their social interactions. The fact that physical space is translucent (and not transparent) to socially salient information is an important resource for structuring interactions. Neither of these dimensions of social translucence is a given in the digital domain.

It is important to note that although our examples are architectural, there is considerable work in the HCI tradition that also bears on these issues. In particular, a considerable body of work begins with research in video-mediated communication (e.g., Fish et al. [1990]; see Finn et al. [1997] for a collection of work), which has since been generalized and is often referred to under the rubric awareness (e.g. Dourish and Bellotti [1992], Gutwin et al [1996]). Another significant strand of research is ethnographic studies of physical work places such as transportation control rooms [Heath and Luff 1991; Harper et al. 1991] and office environments [Bellotti and Bly, 1996].

3. Application Domain: Knowledge Management

What might it mean to have social translucence in a digital system? How might making social information more visible actually change the way digital systems are used? Why might this be a desirable thing? To answer these questions, let's look at knowledge management from a socially translucent perspective.

3.1 Knowledge Management as a Social Phenomenon

Knowledge management is a currently popular term for the attempt to provide organizations with tools for capturing, retrieving and disseminating information about their own activities to their own employees. In a sense, it is an attempt to make organizations self-conscious, to enable them to tap their own experience in solving problems rather than having to reinvent solutions to recurring problems. Knowledge management is often seen as a problem of putting useful information into databases and providing schemes for organizing and retrieving the

information. This perspective leads people to think in terms of data mining and text clustering and databases and documents. This isn't wrong, but it's only part of the picture.

The production and use of knowledge is deeply entwined with social phenomena. For example, one of us once interviewed accountants at a large accounting company about how they would use a proposed database of their company's internal documents. A surprising theme emerged: the accountants said that they'd love to access the documents so that they could find out who wrote them. As one explained, 'Well, if I'm putting together a proposal for Exxon, I really want to talk to people who have already worked with them: they'll know the politics, the history, and other information that usually doesn't get into the reports.' Of particular import was the fact that someone who had worked with the prospective client could give referrals, thus saving the accountant from having to make a cold call. The ability to say 'so-and-so said I ought to call,' was of great value to the accountants (and illustrates yet another function of accountability). Having a referral, however tenuous the connection, is a valuable *social resource* that can only be directly conveyed from one person to another: saying 'I found your name in the corporate knowledge database,' is not the same. It was only through the people -- and the social networks they were part of -- that the accountants could get the knowledge and social resources they needed.

Even as documents may serve as indices into social networks, so may social networks serve as pointers to documents. One of the authors follows an area of research known as genre theory. Giving a search engine a query like "+genre +theory" will generate about fifty thousand hits, most of which are not very relevant. In contrast, the query "+genre +yates +bazerman" generates 20 mostly relevant hits. What is happening is that the query effectively selects for papers that mention "genre" and that also have bibliographies that include both "Yates" (a reference frequently cited by those who study genres in digital documents) and "Bazerman" (a rhetorician and scholar of 'classical', i.e. non-digital genre, typically cited by other classical genre theorists). Thus, the query selects for scholarly papers that draw on the work of two largely divergent research communities. Of course, this is really a trick: the search engine knows nothing of people or scholarly communities or even of the difference between a scholarly paper and other texts; the query works because of the bibliographic conventions of research papers and takes advantage of the co-occurrence of two relatively uncommon names. But it is interesting to think about the possibilities of a system that was designed to 'know' about the notions of authorship, citation, and research communities.

This sort of situation -- the production and use of knowledge in a social milieu -- is not the exception; it is the rule. A variety of research programs -- social studies of science, critical theory, the sociology of knowledge, and ethnographies of the work place -- all point to the deep connections between knowledge and social and cultural contexts. Knowledge, whether it be of bugs in the Java Virtual Machine or of how to begin negotiations with an executive from another culture, is discovered, shared, and used in a social context, not just for its own sake, but to construct the identities and advance the agendas of the individuals, groups, and institutions involved. Having the information in a database isn't as useful as we would hope, unless it also provides an *entree* into the social networks that produced the data. It is from the social networks -- not from the information itself -- that social resources can be recruited.

3.2 From Knowledge Management to Knowledge Communities

Imagine a knowledge management system that was designed from a social perspective, a system predicated on the assumption that knowledge is distributed throughout a network of people, and that only a small proportion of it is captured in concrete form. As the above vignettes suggest, such a system would, along with its data and documents, also provide a rich set of connections back to the social network of people who produced the information. But, if we think in terms of making socially significant activity visible, considerably more possibilities suggest themselves. Imagine that the knowledge management system provided access not only to authors, but to people who were accessing and using the knowledge. Suppose that -- just as we look for crowded restaurants, eye fellow shoppers, or look for engaging conversations -- we could see similar traces of those making use of information in a knowledge management system. After all, some of the knowledge users might have to go to considerable work to apply the knowledge to their own ends, thereby developing an understanding

of its shortcomings and particularities, as well as building on it. If we could capture traces of this knowledge work, others with similar needs might find as much value in talking with users as with the original authors. Such a system would not be just a database from which workers retrieved knowledge, it would be a *knowledge community*, a place within which people would discover, use, and manipulate knowledge, and could encounter and interact with others who are doing likewise.

A knowledge community of this sort has a formidable social problem to overcome: Why should those who produce and use knowledge take the time to engage in such interactions? Why should they wish to? What benefits would they gain that might compensate them for their efforts [Grudin 1989]? This is a deep and difficult research problem -- however, we are optimistic that the problem has a solution (after all, such large scale interaction occurs in the physical world in a wide array of communities and social collectivities). We suggest that the solution to the problem arises as a direct result of social translucence: making knowledge work visible (thus allowing people to observe and contact one another), also enables those who are skilled at unearthing, applying, and adapting knowledge to receive credit for what is all too often an invisible form of work. If knowledge work is made visible it can be recognized and rewarded by the organization (awareness and accountability at the organizational level), and it can permit knowledge work to shift from something that takes time away from 'real work' to being 'real work' in and of itself.

This discussion brings us to an important question: Through what mechanism can knowledge work be made visible? How can it leave traces so that not only can others see it occurring, but that those who were not present can gain value from it at a later time? We believe that the answer to this lies in conversation.

3.3 Conversation: Knowledge Work Made Visible

Conversation is essential. We use conversation as a medium for decision making. It is through conversation that we create, develop, validate, and share knowledge. When systems -- computational or bureaucratic -- freeze, or simply prove too rigid, we pick up the phone to figure out appropriate work-arounds. While e-commerce allows transactions without conversations, it's likely that customers will often want to question, clarify, and negotiate during the purchase of anything other than a commodity. And with all our advances in information retrieval, the preferred method for obtaining information is still to ask a colleague.

Why is this? We suggest that conversation has two characteristics that are central to its power and ubiquity. One vital characteristic of talk is that it is a deeply interactive intellectual process (see Clark [1996] for a detailed exposition). As we talk we refer to a common ground of already established understandings, shared experiences, and past history. As the conversation proceeds, we are continuously attempting to interpret what is said, verify that we have been understood, and offer new contributions. Sometimes misunderstandings occur, and so we attempt to fix them by rephrasing our words, or 'debugging' the previous conversation to reveal that what we thought were shared understandings were not, in fact, shared. What all this amounts to is that conversation is a superb method for eliciting, unpacking, articulating, applying, and recontextualizing knowledge.

Conversation is more than simply an intellectual endeavor: it is a fundamentally social process (e.g., see Goffman [1963 and 1967], and Kendon [1990]). Conversation is social in two ways. First, people speak *to* an *audience*. Speakers notice how their audience is reacting and steer their remarks appropriately: nods and eye contact convey one message; questions and furrowed brows another; yawns and fidgeting still another. Second, conversation is social in that people portray themselves through conversation. They advance their personal agendas, project their personal style, take credit, share blame, and accomplish other social ends through their talk, often with a great deal of subtlety. The social nature of talk is not an undesirable side effect, but rather the heart of it: personal motivations fuel conversation and provide the energy for the considerable intellectual work it takes, whether the conversation in question is banter over morning coffee or about the composition of a journal article.

In addition, conversation within the digital medium, has a property of great importance for our purposes: it

persists. Instantiated as text, whether typed in or spoken and recognized, persistence expands conversation beyond those within earshot, rendering it accessible to those in other places and at later times. Thus, digital conversation may be synchronous or asynchronous, and its audience intimate or vast. Its persistence opens the door to a variety of new uses and practices: persistent conversations may be searched, browsed, replayed, annotated, visualized, restructured, and recontextualized, with what are likely to be profound impacts on personal, social, and institutional practices.

3.4 The Vision: Conversationally-Based Knowledge Communities

While considerably more remains to be said, we will stop here and sketch an outline of the sort of system we have in mind. Our sketch will be rough because our purpose is not to define a system in detail, but simply to provide a framework in which to position existing research and to raise new research questions, the subjects of the final two sections of this paper.

We see several types of needs, each of which suggests a layer of functionality:

Activity Support. First, it is necessary to support the initiation and conduct of social activity via digital media. In particular, we want to support long-running, deep, coherent conversations that can be steered by their participants. In such an environment we believe that knowledge can be constructed through reflective conversations among a community engaged in a shared set of tasks.

Conversation Visualization and Restructuring. The knowledge embedded in conversations should be re-useable. That is, we want to move from today's state, where conversation is of value primarily as it occurs, to a situation in which the comments that make up a conversation can serve as a useful work product. It should be possible for members of a knowledge community to search, navigate, and visualize their conversations. Similarly, participants should be able to add structure to conversations: summarizing, highlighting, linking, and otherwise annotating them.

Organizational Knowledge Spaces: Our experience suggests that knowledge production and use will proceed most easily in a semi-private environment -- a relatively small community where knowledge workers feel 'safe' enough to venture tentative interpretations and conjectures. At the same time, the *raison d'être* of a knowledge community is to share knowledge with the organization within which it is embedded. One resolution to this tension between privacy and visibility is to support an organizational space within which semi-autonomous knowledge communities can exist, each exercising control over the ways and means through which its knowledge is shared with the larger organization.

4. Implementation: social translucence in digital systems

So far we have discussed social translucence as a general approach to designing systems for computer-mediated communication and collaboration, and have described how it might map onto the domain of knowledge management, and why that might be desirable. In the remainder of the paper we ask how these ideas might be converted to practice. In this section we review some of the work that is applicable to designing socially translucent systems and describe a prototype system that we have implemented which illustrates a socially translucent approach to supporting conversational activity. In the final section we discuss some of the research issues that remain to be addressed.

4.1 Making Activity Visible

While the perspective we've developed -- social translucence -- is unique, we are certainly by no means the first to be concerned with making the activities of users of digital systems visible to others. A concern for making other users visible dates back to at least the Finger program on UNIX. More recently, a number of investigators have explored ways of portraying socially-salient information in human computer interfaces. Ackerman and Starr [1995] have argued for the importance of social activity indicators, particularly in synchronous CMC systems. Hill and his colleagues [1992] have discussed the creation of persistent traces of human activity. And a considerable number of researchers have constructed systems that attempt, in various ways, to provide cues about the presence and activity of their users (for example, Benford et al. [1994]; Dieberger [1997]; Gutwin, et al. [1996]; Hill et al. [1995]; Isaacs et al. [1996]; O'Day et al. [1996]; Roseman & Greenberg [1996]).

This raises the question of how social cues might best be portrayed in a digital system. We see three design approaches to answering this question: the realist, the mimetic, and the abstract. The realist approach involves trying to project social information from the physical domain into or through the digital domain. This work is exemplified in teleconferencing systems and media space research -- see Finn, Sellen, and Wilbur [1997] for many examples. While the realist approach is a straightforward answer to the question, in that it minimizes the difficulty of producing and interpreting social cues, it has a number of drawbacks: First, current technology places very real limits on the resolution (both temporal and spatial) of teleconferencing and videospace systems. Subtle cues such as facial expressions, gaze, and posture may go unnoticed due to resolution limits or more pragmatic factors such as the direction and size of the camera's field of view, thus creating a mismatch between the ability to produce and detect social cues. Related to this is the fact that most systems of this sort are still very expensive -- in terms of both financial cost and requirements for infrastructure and management support. Finally, even were these problems resolved, realist solutions don't scale well: they can't support the many tens, hundreds, or thousands of people who even today interact over the internet.

The mimetic approach tries to re-represent social cues from the physical world, as literally as possible, in the digital domain. The mimetic approach is exemplified by graphical MUDs and virtual reality systems, and uses virtual environments and avatars of various degrees of realism to mimic the physical world. Work here ranges from attempts to implement a virtual physics (e.g., the work of Benford and his colleagues [1994]), to the considerably looser re-presentations of social information found in the 2-D and 3-D avatars found in various graphical MUDs and 3-D VRML worlds. While the use of mimetic representations decreases some of the bandwidth requirements of such systems, it doesn't scale much better than realist approaches. It also introduces another problem: in this sort of system social cues must be consciously produced via users manipulating their avatars. Because the goal is to mimic the social cues being produced (often unconsciously) in face to face situations through an interface that is explicit and quite indirect, production of these cues can require considerable attention and effort. On the other hand, it is generally the case that cues that can be produced can also be detected, a benefit over the realist approach.

The abstract approach involves portraying social information in ways that are not closely tied to their physical analogs. Exemplars of the abstract approach include AROMA [Pedersen and Sokoler 1997], the Out to Lunch system [Cohen 1994] (which uses abstract sonic cues to indicate socially salient activity), and Chat Circles [Viegas and Donath 1999] (which uses abstract visual representations). This approach also includes the use of text to portray social information. Text has proved surprisingly powerful as a means for conveying social information, as can be seen in studies of MOOs [Bruckman 1997; Curtis 1992; Cherny 1999; Davis and Brewer 1997], textual chat rooms [Danet et al. 1998; Werry 1996], and the Chat-and-Search system of Iwayama et al. [1999]. The abstract approach scales well, both textual and abstract graphical representations being relatively compact.

While we think all these approaches are promising, we're particularly interested in the abstract approach. First, we believe that systems that attempt to leverage social processes need to be developed through a process of creating and deploying working systems, and studying their use in ordinary work contexts. This intent to deploy, in and of itself, is a strike against the realist and mimetic approaches, both of which face substantial pragmatic barriers (e.g., expense, infrastructure, support) to deployment outside of research institutions. Second, and more

importantly, we believe that the abstract approach has not received sufficient attention from designers and researchers, particularly with respect to graphical representations. Text and simple graphics have many powerful characteristics: they are easy to produce and manipulate; they persist over time, leaving interpretable traces (helpful to those trying to learn the representation); and they enable the use of technologies such as search and visualization engines.

4.2 Abstract Representations of Social Information: The Babble Prototype

As the first step towards designing an infrastructure for knowledge communities, our group designed, implemented and deployed a system intended to provide an environment for supporting smooth, long-running, productive conversations among members of small to medium sized groups. Here we describe the prototype, called Babble, focusing on the way in which it uses textual and graphical representations to make socially salient information visible (see Erickson et al. [1999] for more information).

4.2.1 Portraying Social Information in Babble

In designing Babble, we used two tactics to make social information visible: a persistent textual representation of the conversation, and a synchronous representation of the activity of participants.

```
===Friday 12 Dec 97 3:43:44 From: Bill
Hi Steven!
===Friday 12 Dec 97 3:44:49 From: Steven
Hellooo Bill. A little guidance please?
Is the [...] summary we're preparing for [...]
supposed to be an exercise in feeling
good, or are we supposed to be giving
him hard-headed guidance?
===Friday 12Dec 97 3:56:55 From: Bill
yes :-)
```

Figure 1. A segment of conversation displayed as a single, shared, persistent document.

In Babble, a conversation is represented as a single, persistent document (e.g., Figure 1). Each comment is preceded by a header that has a time stamp and the participant's name; new comments are appended to the end of the document; and the entire conversation is shown in a single window. This approach is employed by some types of asynchronous bulletin boards, and we believe it has significant advantages over other forms of textual representation. Unlike chat, where conversation is ephemeral, or like mailing lists, where the past becomes buried in message archives, accessing the conversation's history is just a matter of scrolling. This visibility of the conversation's history can support the emergence and enforcement of conversational norms (e.g., Erickson [1999]).

This type of representation has a number of advantages both for readers and for writers. Readers, whether they are newcomers or simply infrequent participants, can get an overview of social norms that govern a conversation by skimming through it. For example, the length of comments is apparent, as is the informality of the conversation (inferable from the simplified syntax and the absence of punctuation), and degree of humor and politeness. In short, the persistent trace left by textual conversation is a rich source of socially salient information (awareness), and its power is enhanced because participants know that the representation is shared and thus that everyone is privy to the same set of cues (accountability). The persistence of the conversation has consequences for authoring as well. Another important element of this representation is its shared sequential structure: everyone sees the same order of remarks, and can therefore count on that as a resource for structuring their interaction. Thus participants can and do participate with short, indexical utterances like 'Yes!', 'Thank you,' and

'Great idea!'. This type of response is less likely to occur in a mailing list where the absence of shared, sequential structure requires participants to quote the message to which they are responding (and because participants are often annoyed at opening a message and finding only an 'insignificant' comment in it). While such indexical comments may not extend the conceptual bounds of the conversation, they can make it more convivial and inviting by providing an easy way for participants to signal agreement, encouragement and empathy.

A novel aspect of Babble is the *social proxy*, a minimalist graphical representation of users that depicts their presence and their activities (Figure 2). The social proxy portrays the conversation as a large circle, and the participants as colored dots (shown as small numbered circles in the schematic in Figure 2), referred to, hereafter, as marbles. Marbles within the circle are involved in the current conversation; marbles outside the circle represent those who are logged on but are in other conversations. The marbles of those who are active in the current conversation, either talking (i.e., typing) or 'listening' (i.e., interacting via mouse clicks and movements) are shown near the circle's center; with inactivity marbles drift out to the periphery. When people leave the current conversation their marbles move outside the circle; when they enter the conversation, their marbles move into the circle. When people log onto the system it creates virtual wedges for their marbles, adjusting the position of all the marbles in the social proxy; when they depart, the wedges are destroyed, and the remaining marbles adjust to uniformly occupy the space. All marble movements are animated, thus making arrivals, movements, and departures visually salient.




Figure 2. Social proxy schematic. Part (a) shows the layout of the social proxy: dots 1, 2 and 3, inside the circle, are part of the 'current' conversation; dot 4 is in another conversation. Part (b) shows the dot animation: they move abruptly to the center when they are active, and slowly drift to the periphery with inactivity. Thus, a tight cluster of dots represents an active conversation.

Although simple, this social proxy gives a sense of the size of the audience, the amount of conversational activity, as well as indicating whether people are gathering or dispersing, and who it is that is coming and going. Also, because the portrayal is graphical, it has a perceptual directness (like the glass window) that a list of written names lacks. Experientially, the social proxy is interesting because it focuses attention on the group as a whole, and the coherence (or lack thereof) of its activity.

4.2.2 Social Activity in Babble

As of this writing, Babble has been in daily use by its implementers for two years, and has been deployed to about eight other groups who have used it for periods of two to six months. Most (though not all) deployments have been to relatively small groups in which most of the members know one another. Babble has proved to be an effective environment for supporting informal group conversations, on a wide range of topics. Here we describe a few of the observed social phenomena which are most relevant to the knowledge communities vision

(see Bradner et al. [1999] for a broader and deeper discussion of usage and adoption issues observed in six deployments).

One social practice we've observed is waylaying, in which a user watches for a particular person to become active on Babble (signaled by the movement of their marble into the center of the social proxy), and then initiates a conversation by via private chat, telephone, or other means. Because the movement of the marble occurs when the user has just begun an episode of typing or mousing, it indicates a opportune moment for contact (since the user's attention has just shifted to communication with the group). Waylay is used for purposes ranging from asking questions to initiating casual social chat.

Babble also supports the maintenance group awareness. For example, when members of a Babble travel, many report reading through conversations that occurred in their absence to 'find out what happened.' For someone who is a member of the group and understands the context seemingly trivial comments can convey considerable information about what's going on at the individual, group, and organizational levels. Thus, a sign off -- 'I have to go to the [project] meeting now' -- reveals that one participant is still involved in a particular project, and a question -- 'Does anyone know how to do a screen capture' -- indicates that another participant is beginning to write a paper.

Babble was typically regarded as a semi-private place, and conversations were often frank and unguarded since they were produced for a small and well-known audience that shared considerable context. This became apparent when 'strangers' (e.g., visitors; new members; or, in one case, an unannounced conversational software agent) appeared in various Babbles, their presence and activities being displayed in the social proxy (awareness) along with those of the regulars. In each case, the appearance of strangers provoked concern about how unguarded conversations might be interpreted by those from different contexts (accountability). Sometimes this led to the creation of rules for how new members were admitted into a Babble or for what areas of a group's Babble were appropriate demonstrate to outsiders.

While one should be cautious about generalizing from a few experiences, these results are not at variance with what has been observed in other on-line systems (for examples see Bruckman [1997] Bruckman and Resnick [1995]; Cherny [1999]; Churchill and Bly [1999]; Fanderclai [1996]; and Kovalainen et al. [1998]). In general, we regard Babble as an existence proof that textual and graphical representations of social information can provide a foundation for rich social interaction.

5. Some Research Issues

Coming at the design of digital systems from the perspective of social translucence raises a huge number of questions. In this section we describe some of the issues that seem of greatest importance to us with respect to the design of knowledge communities; we do not presume to map out the entire terrain.

5.1 Social Proxies: What Should be Represented?

While the Babble social proxy is simple, we think the concept is promising, and see a variety of ways of extending its scope. The dimensions of this problem include the people and activities represented in the proxy, the spatial scope of the proxy, and the temporal scope of the proxy. For example, in the Babble prototype we show the participants in public conversations, but we do not show who is participating in the private one-to-one chats that the system also supports, or even whether a private chat is taking place. On the one hand, an argument in favor of making private chats visible is that it might increase awareness of the feature, and cause others to take up its use; in addition, it can be argued that -- just as in shared physical spaces -- seeing that two people are chatting (without knowing what is said) can convey useful information without necessarily infringing on their privacy. On the other hand, it is easy to generate negative scenarios. We simply wish to suggest that the tradeoffs are not well understood, and that neither a blindly optimistic approach of making everything visible, nor a

trenchantly pessimistic approach of putting privacy above all else, is particularly useful. One approach might be to explore ways of making it evident to all participants which of their actions are visible: this is something that we tend to understand pretty well for physical space, but which is not clear how to do effectively in digital systems.

The spatial and temporal scope of the social proxy is another issue. Babble currently shows the synchronous activity of only those participants in the conversation being viewed. Obviously, it could be useful to have a view of several conversations, or the entire knowledge community, so that one could notice activity elsewhere. For example, a community proxy (Figure 3) could depict the global structure of a knowledge community, showing the number of topics, and the distribution of activity across them. Community members, with an understanding of the semantics of the structure, could make inferences about the meaning of clusters of participants and the distribution of new information.




Figure 3. A social proxy showing the global structure of a knowledge community: larger circles represent conversation topics, with filled circles indicating new information; smaller dots represent participants.

Similarly, it might be useful to provide a diachronic (longitudinal) view of a conversation. Even in a conversation where people are never simultaneously present, it still makes a difference whether fifty people check in every day, or one person checks in every month. In the former case, there will be a lot of social motivation for people to contribute; in the latter, there will be little. Such a proxy could show arrivals and departures, flurries of conversation, and dialog patterns, even in conversations where two people are never simultaneously present. Figure 4 shows two examples of diachronic social proxies: the one on the left is based on the model of a ticker tape, with each row representing the activity of an individual (a line indicates presence, a blip indicates speaking) that provides a visualization of global activity across the day (notice the system crash at 10:30am); the proxy on the right is a schematic of a Babble-like social proxy in which the marbles leave contrails -- this provides a different view of activity over time within a conversation. In the contrail proxy we can see that although dots 1 and 2 are both active, 2 is a recent arrival. Similarly, we see dot 3 has recently gone idle, whereas dot 4 has been inactive for a longer period. These proxies represent two ways of broadening what Johnson-Lenz [1991] refers to as the 'rolling present,' the time that constitutes the 'now' of a particular interaction. Each of the diachronic proxies makes certain things more or less visible: thus, for example, the tickertape proxy makes it easy to assess who logs on earliest; similarly, the 'contrail' social proxy makes it easy to see who has recently attended to a particular conversation. In both cases, this information could be used to positive or negative ends, and its presence will likely shape interactions in unexpected ways.




Figure 4: Two examples of a diachronic (longitudinal) social proxy. The one on the left is modeled after a ticker tape with lines indicating presence and blips indicating talk; the other is a 'conrail' proxy in which the dots leave ephemeral trails.

5.2 Supporting Coherent Activity

One of the problems with on-line interaction is that it is often characterized by incoherence and drift (e.g., Herring [1999]). However, this is not necessarily the case. Erickson [1999] presents a case study of an on-line limerick-making game among several dozen people that, over the course of a year, succeeded in repeatedly producing coherent results. The coherence of the interaction in question was facilitated by a set of rules that were shared and mutually enforced by the participants (e.g., only one line could be added by a participant at a time). In addition, as the interaction proceeded, new conventions evolved to govern their interactions (e.g., the person who completes a limerick is obligated to start the next one). While the case described is unusual in several respects, one important factor is that the combination of the system's interface and the nature of the interaction resulted in many of the conventions being highly visible. That is, when people violated one of the conventions that structured the conversation (e.g., adding two lines at once), the violation was visible as an interruption in a spatial patterning of text. Because the violation was visible both to the violator and to the other participants, it increased awareness and accountability.

This raises the question of whether social proxies (and other means of portraying social information) might be useful as a way of representing the conventions of an interaction and allowing their enforcement by the participants. For example, the circular form of the Babble version of the social proxy implies an open conversation in which all are allowed to participate. One might redesign the social proxy to make visible the constraints that define different styles of interaction, such as a lecture (figure 5). The 'lecture' social proxy is designed in both form and animation dynamics so when a lecture is proceeding as it ought, it appears well-formed. The idea is that a dot is gradually moved towards the apex of the proxy as a function of cumulative talking, so that normally one dot (the lecturer) is all the way at the front, and all the other dots (the audience) are arrayed at the back. However, if the interaction begins to depart from the norm (for example, with multiple members of the audience talking), the proxy will begin to look out of balance, thus supporting awareness and accountability.




Figure 5. A schematic of a 'lecture' social proxy, in which dots move towards the apex of the wedge with cumulative activity.

5.3 Visualizing Conversation

Given a knowledge community that supports fluent conversation, it will very quickly produce a nearly impenetrable mass of text (as the Babble system has). While experienced users may perhaps have some success in navigating it, newcomers or visitors, are likely to be overwhelmed. This problem is exacerbated by the fact that the nature of on-line conversation includes a more casual approach to spelling, abbreviation, and acronyms, thus limiting the usefulness of conventional search techniques. Thus we see a need for mechanisms that support the visualization of conversation.

Conversations have considerable structure. They vary in their numbers of participants, patterns of participation, dynamics (e.g., degree of asynchrony), and duration. Contributions may vary in length, formality: some conversations may be composed primarily of rapid-fire, near synchronous exchanges of short utterances among many people; others may be deliberate, slowly paced, dialogs consisting of long, carefully composed remarks. This structure can be exposed through visualizations. A technique as simple as 'greeking' the text down to gray lines, and showing boundaries between utterances, and providing a time line can provide an overview of the tempo and rhythms of the conversation. To this visualization can be added information about the participants, commonly used words, the size (and fluctuation thereof) of the audience, etc. Or one might add an overlay for searching: Figure 6 shows a simple textual visualization, with an overlay of hits in response to a query for several phrases, each hit highlighted with a color corresponding to its search term. Such tools could help a participant harvest conversation so that it could be cleaned up and made publicly available. We believe that representations of this sort are likely to be meaningful to those who participated in the conversations, and can provide a useful structure for browsing and navigating old conversations. To date, little work has been done on this topic (see Donath et al.[1999]and Small [1996]).



Figure 6: A 'greeked' conversation visualization with search hits shown.

5.4 Restructuring Conversation

While visualization is likely to be helpful for those who have participated in conversations, it will be less useful for visitors and new members. For a knowledge community we envision a core of original conversation, around which has accreted a body of more ordered text -- a discourse base that supports the re-use of conversationally produced knowledge. There is a need for tools that support what Johnson-Lenz [1991] calls harvesting: the creation of summaries, indices, and other distillations of conversational content. While some such tools exist, most are designed to work on documents and larger textual corpi, whose volume, stricter conventions, and homogeneity of authorship make them quite distinct from conversations; whether such tools can be adapted to process conversation is a non-trivial research issue.

5.5 Organizational Knowledge Spaces

The creation of discourse bases -- productive conversations that have been restructured (or have had new elements like summaries, glossaries, and edited paths added on), is only part of the solution to making the knowledge produced through conversations accessible. There needs to be a way for members of the larger organization, in which the knowledge community is embedded, to browse its discourse bases and discover useful knowledge. This is not just a benefit for those seeking knowledge -- it also provides increased visibility and payoffs for those who do the work of producing, adapting and distilling knowledge. Even so, this leads to the tension between privacy and visibility observed in the Babble prototype. One area to explore is how to create translucent representations of conversations that allow outsiders to see enough to determine if the conversation is relevant, without exposing the jokes, camaraderie, and improprieties that characterize relaxed talk among peers. For example, just as the conversational visualization shown in Figure 6 showed search hits without revealing the text of the conversation, so might a community-scale social proxy (as in Figure 3) be used to reveal the locus and extent of a particular concept being discussed within a knowledge community without revealing the potentially sensitive surround. Given clues that useful knowledge is present, interested parties could *request* summaries of the topic, *petition* for admission to the community, or simply *converse* with some of the community members. Notice how this strategy blends technical and social mechanisms: technology is used to locate hot spots without revealing the particulars, and social mechanisms are used to control access. Such an approach has at least two benefits: explicit requests for information provide a way of documenting the value of the knowledge work; and could also serve as catalysts for restructuring. More generally, this raises the issue of how one might visualize the knowledge space -- contents, structure, activity -- of an entire organization. So far as we know, except for Donath's visualization of a community's use of mailing lists [Donath 1995], this is unexplored territory.

6. Conclusion

Let us return to the tale of the door with the glass window. Although we have focused on designing systems for communication and collaboration, the contrast between the opaque door with its sign and the door with its glass window seems an apt metaphor for a very general problem with technological systems today. In the first case, the system creates a barrier between users; the remedy requires that they do extra work--noticing the sign, interpreting its words, and adjusting their actions just in case. In the second case, the system reveals the presence of those using it, enabling already-established social rules and mechanisms to come smoothly into play.

From our perspective, the digital world appears to be populated by technologies that impose walls between people, rather than by technologies that create windows between them. We suggest that understanding how to design digital systems so that they mesh with human behavior at the individual and collective levels is of immense importance. By allowing users to 'see' one another, to make inferences about the activities of others, to

imitate one another, we believe that digital systems can become environments in which new social forms can be invented, adopted, adapted, and propagated -- eventually supporting the same sort of social innovation and diversity that can be observed in physically based cultures.

Acknowledgements

This work is highly collaborative. Thanks to our colleagues at IBM: David N. Smith for creating the Babble prototype; Mark Laff and Amy Katriel for implementations of the social proxy and Java client; Cal Swart for supporting Babble deployments; and all the members of the Niche Networking and Social Computing groups for conversation, inspiration, and general support. This work also benefited from the efforts of our 1998 summer interns, Erin Bradner who lead the deployment studies, and Jason Ellis who designed a number of provocative prototypes, including the lecture social proxy. Thanks to Amy Bruckman, Saul Greenberg, Michael Muller, Charles Taliaferro, and several anonymous reviewers for comments on this work.

References

- Ackerman, M.S. and Starr, B. (1995). Social activity indicators: Interface components for CSCW systems. In *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST'95)*, New York: ACM, pp. 159-168.
- Alexander, C., Ishikawa, S., Silverstein, M., Jacobson, M., Fiksdahl-King, I., and Angel, S. (1977). *A pattern language*. New York: Oxford University Press.
- Bellotti, V. and Bly, S. (1996). Walking Away from the Desktop Computer: Distributed Collaboration in a Product Design Team. *Proceedings of CSCW '96*.
- Benford, S., Bowers, J., Fahlen, L., Mariani, J., and Rodden, T. (1994). Supporting Co-operative work in Virtual Environments. *The Computer Journal*, Vol 37, No. 8, Oxford University Press.
- Bradner, E., Kellogg, W.A. and Erickson, T. (1999). The Adoption and use of Babble: A field study of chat in the workplace. In preparation: submitted to the European Conference on Computer-Supported Cooperative Work (ECSCW '99).
- Bruckman, A. (1997). MOOSE Crossing: Construction, Community, and Learning in a Networked Virtual World for Kids. PhD dissertation, Massachusetts Institute of Technology.
- Bruckman, A. & Resnick, M. (1995). The MEDIAMOO project: Constructionism and professional community. In *Convergence*, 1:1, Spring 1995.
- Cohen, J. (1994). Monitoring background activities. In G. Kramer (Ed.), *Auditory display*. New York: Addison-Wesley, pp. 439-531.
- Cherny, L. (1999). *Conversation and Community: Chat in a Virtual World*. Stanford, CA: CSLI Publications: 1999.
- Clark, H.H. (1996). *Using Language*. Cambridge: Cambridge University Press, 1996.
- Churchill, E.F. and Bly, S. (1999). Virtual environments at work: Ongoing use of MUDs in the workplace. In *Proceedings of WACC'99*, February 22-25, 1999, San Francisco, CA.
- Curtis, P. (1992). Mudding: Social phenomena in text-based virtual realities. In *Proceedings of DIAC '92*. Available via anonymous FTP from [parcftp.xerox.com, pub/MOO/papers /DIAC92](ftp://parcftp.xerox.com/pub/MOO/papers/DIAC92).

- Davis, B.H. and Brewer, J.P. (1997). *Electronic discourse: Linguistic individuals in virtual space*. Albany, NY: SUNY.
- Danet, B. Ruedenberg, L. and Rosenbaum-Tamari, Y. (1998). Hmm...where's that smoke coming from? Writing, play and performance on Internet Relay Chat. In F. Sudweeks, M. McLaughlin, & Rafaeli (Eds.), *Network & netplay: Virtual groups on the Internet*. Menlo Park, CA: AAAI Press, pp. 41-76.
- Dieberger, A. (1997). Supporting social navigation on the World Wide Web. *International Journal of Human-Computer Studies*, 46(6), pp. 805-825.
- Donath, J. (1995) Visual Who: Animating the affinities and activities of an electronic community. *Proceedings of the Third International Conference on Multimedia '95* (San Francisco, CA: November 5-9, 1995). New York: ACM, pp 99-107.
- Donath, J., Karahalios, K., & Viegas, F. (1999). Visualizing conversation. In J.F. Nunamaker, Jr., & R.H. Sprague, Jr., (Eds.), *Proceedings of the Thirty-Second Hawai'i International Conference on Systems Science*. Los Alamitos, CA: IEEE Computer Society Press.
- Dourish, P. and Bellotti, V. (1992) Awareness and Coordination in Shared Workspaces. *Proceedings of CSCW '92* (Toronto, 1992), pp. 107-114.
- Erickson, T. (1997). Social interaction on the net: Virtual community as participatory genre. In J. F. Nunamaker, Jr., & R.H. Sprague, Jr., (Eds.), *Proceedings of the Thirtieth Hawai'i International Conference on Systems Science*, Vol 6. Los Alamitos, CA: IEEE Computer Society Press, pp. 23-30.
- Erickson, T. (1999) Rhyme and punishment: The creation and enforcement of conventions in an on-line participatory limerick genre. In J. F. Nunamaker, Jr., & R. H. Sprague, Jr. (Eds.), *Proceedings of the Thirty-Second Hawai'i International Conference on Systems Science*. Los Alamitos, CA: IEEE Computer Society Press.
- Erickson, T., Smith, D.N., Kellogg, W.A., Laff, M.R., Richards, J.T., and Bradner, E. (1999). Socially translucent systems: Social proxies, persistent conversation, and the design of Babble. *Human Factors in Computing Systems: The Proceedings of CHI '99*, ACM Press.
- Fanderclai, T. (1996). Like magic, only real. In L. Cherny & E. Weise (Eds.), *Wired women: Gender and new realities in cyberspace*. Seattle, WA: Seal Press.
- Finn, K.E., Sellen, A.J., and Wilbur, S.B. (1997). *Video-mediated communication*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Fish, R.S., Kraut, R.E. and Chalfonte, B.L. (1990) The VideoWindow System in Informal Communications Shared Video Spaces. *Proceedings of ACM CSCW'90 Conference on Computer-Supported Cooperative Work*, p.1-11.
- Gehl, Jan. (1980). *Life between buildings: Using public space*. New York: Van Nostrand Reinhold, 1980/1987.
- Goffman, E. (1963). *Behavior in public places: Notes on the social organization of gatherings*. New York: Macmillan Publishing Co.
- Goffman, E. (1967). *Interaction ritual*. New York: Anchor Books
- Grudin, J. (1989). Why groupware applications fail: Problems in design and evaluation. *Office: Technology and People*, Vol. 4, No. 3, pp. 245-264.
- Gutwin, C., Greenberg, S. and Roseman, M. (1996). Workspace Awareness in Real-Time Distributed Groupware:

- Framework, Widgets, and Evaluation. In Sasse, R.J., A. Cunningham, and R. Winder, (Eds.), *People and Computers XI (Proceedings of the HCI'96)*, pp. 281-298, Springer-Verlag.
- Heath, C. C. and Luff, P. (1991) Collaborative Activity and Technological Design: Task Coordination in London Underground Control Rooms. *Proceedings of E-CSCW 91* (Amsterdam, Sept. 24-27 1991), pp. 65-80.
- Herring, S. (1999). Coherence in CMC. In J. F. Nunamaker, Jr., & R.H. Sprague, Jr., (Eds.), *Proceedings of the Thirty-Second Hawai'i International Conference on Systems Science*, Vol 8. Los Alamitos, CA: IEEE Computer Society Press.
- Hill, W., Hollan, J.D., Wroblewski, D., and McCandless, T. (1992). Edit wear and read wear: Text and hypertext. In *Proceedings of ACM CHI '92 Conference on Human Factors in Computing Systems*, pp. 3-9.
- Hill, W., Stead, L., Rosenstein, M., and Furnas, G. (1995). Recommending and evaluating choices in a virtual community of use. In *Proceedings of ACM CHI '95 Conference*.
- Isaacs, E.A., Tang, J.C., and Morris, T. (1996). Piazza: A desktop environment supporting impromptu and planned interactions. In M.S. Ackerman (Ed.), *Proceedings of the ACM 1996 Conference on Computer Supported Cooperative Work*, pp. 315-324.
- Iwayama, N., Murakami, M. & Matsuda, M. (1999), Consulting search engines as conversation. In the Proceedings of the 7th IFIP Conference on Human-Computer Interaction, Interact'99.
- Jacobs, J. (1961). *The death and life of great American cities*. New York: Random House.
- Johnson-Lenz, P. and Johnson-Lenz, T. (1991). Post-mechanistic groupware primitives: Rhythms, boundaries, and containers. *International Journal of Man-Machine Studies*, Vol. 34, 395-417.
- Kendon, A. (1990). *Conducting interaction: Patterns of behavior in focused encounters*. Cambridge, MA: Cambridge University Press.
- Kovalainen, M., Robinson, M., and Auramki, E. (1998). Diaries at work. In *Proceedings of the ADM 1998 Conference on Computer-Supported Cooperative Work (CSCW'98)*, Seattle, Washington, USA, pp 49-58.
- Luff, P., Heath, C., and Greatbatch, D. (1992) Tasks-in-interaction: Paper and Screen Based Documentation in Collaborative Activity. *CSCW 92 Proceedings* (November, 1992), pp. 163-170.
- Lynch, K. (1990). In T. Banerjee and M. Southworth (Eds.), *City sense and city design: Writings and projects of Kevin Lynch*. Cambridge, MA: The MIT Press.
- O'Day, V.L., Bobrow, D.G., and Shirley, M. (1996). The social-technical design circle. In M.S. Ackerman (Ed.), *Proceedings of the ACM 1996 Conference on Computer Supported Cooperative Work*, pp. 160-169.
- Pedersen, E. R., and Sokoler, T. (1997). AROMA: Abstract representation of presence supporting mutual awareness. In *Human Factors in Computing Systems: CHI'97 Conference Proceedings*. ACM Press, pp. 51-58.
- Roseman, M. and Greenberg, S. (1996). TeamRooms: Network places for collaboration. In M.S. Ackerman (Ed.), *Proceedings of the ACM 1996 Conference on Computer Supported Cooperative Work (CSCW'96)*, pp. 325-333.
- Small, D. (1996) Navigating Large Bodies of Text. *IBM Systems Journal*, Vol. 35, No. 3&4.
- Viegas, F. and Donath, J. (1999). Chat Circles. *Proceedings of the CHI99 Conference on Human Factors in Computing Systems*, ACM Press, pp. 9-16.
- Werry, C. C. (1996). Linguistic and interactional features of Internet Relay Chat. In S. Herring (Ed.), *Computer-*

mediated communication: Linguistic, social and cross-cultural perspectives. Amsterdam: John Benjamins, pp. 47-63.

Whyte, W.H. (1988). *City: Return to the center*. New York: Anchor Books.

* To appear in *Transactions on Computer Human Interface*. In preparation, ACM Press, 2000.