

Social Translucence: Using Minimalist Visualizations of Social Activity to Support Collective Interaction¹

Thomas Erickson and Wendy A. Kellogg

1.1 Introduction: The Ubiquity of the Social

As humans, we are fundamentally social creatures. From birth we orient to other people, and as we develop we acquire abilities for interacting with one another ranging from expression and gesture through spoken and written language. As adults we are exquisitely sensitive to the actions and interactions of those around us. Every day we make countless decisions that are shaped by our social context. 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 provides a basis for inferences, planning, and coordinating activity.

When we move from face to face interaction to digitally-mediated interaction, however, everything changes. The subtle social cues that we use to guide and structure our real world interactions are mostly absent. In the digital world we are socially blind, and our attempts to communicate can be awkward and labor-intensive. Although the web is used by millions of people, reading a web page is usually a solitary experience. Even when others are clearly present—as in a chat room or on a conference call—it is difficult to see who is present, who is paying attention, or who wishes to speak. Things that require little effort in face to face settings—taking turns when speaking; noticing when someone has a question; seeing who is responding to whom—require a lot of effort in online settings, if they are possible at all.

Our aim is to design systems that support deep, coherent and productive 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 allow groups to observe their interactions, and to steer them so as to make progress towards a shared goal. In this chapter we describe our work towards this end. We begin by taking a close look at how people interact in physical, face-to-face settings, with particular attention to the ways in which they make use of socially salient cues to regulate their behavior. Here we develop the notion of social translucence, an approach to designing digital systems that emphasizes making social information visible to all participants. Next we describe how we map the concept of social translucence into digital systems. We do not try to imitate the real world (e.g., via virtual reality or video); instead, we use “social proxies,” minimalist graphical representations of the online presence and activities of people. We illustrate this concept with several prototypes designed for different online contexts, noting how they provide the information that supports the social processes—such as imitation, norming, and peer pressure—that underlie so much human behavior. We then move on to describe

¹ Sections 2, 3.1 and 4.1 of this chapter have previously appeared in “Social Translucence: An Approach to Designing Systems that Mesh with Social Processes,” in *ACM Transactions on Human Computer Interaction*, Vol. 7, No. 1. New York: ACM Press, 2000.

Babble, an implemented system that uses two social proxies to support online conversation, and discuss what we've learned while studying its use over the last four years.

1.2 Foundations: Social Translucence and Interaction Architectures

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).

However, 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 we need to go a bit farther than architectural discourse takes us. 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.

1.2.1 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 supporting coherent social behavior.

We see these three properties of socially translucent systems—visibility, awareness, and accountability—as building blocks of social interaction. That is, we claim that 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. To see this, let’s move on to a more complex example.

1.2.2 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: This, it is not an accident that voting booths are designed so that observers can not see who a person is voting for, but that they can see that the voter is alone in the booth; similarly, the ballot box is positioned in public view so that it is easy for observers to see that a voter puts one and only one ballot in the ballot box. The entire process and apparatus that supports elections is carefully design to ensure that, on the one hand, the individual voter is accountable to no person or group with regards to his or her vote, but, on the hand, that both voters and the election administrators are publicly accountable for following the rules of the election process.

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 complicated 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 chapter 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.

1.2.3 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.

1.3 Design: Embedding Social Translucence in Digital Systems

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. First, in addition to our architectural examples, there are a number of ethnographic studies of transportation control rooms (e.g. Heath and Luff, 1991), offices (Bellotti and Bly, 1996) and other physical work places (see Heath and Luff, 2000) which reveal the crucial role of visibility and mutual awareness in supporting coordinated activity.

A concern for making other users visible in digital systems dates back to at least the Finger program on UNIX. More recently, a considerable body of work begins with research in video-mediated communication (e.g. Finn et al., 1997), which has since been generalized and is often referred to under the rubric “awareness” (e.g. Dourish and Bellotti, 1992; Gutwin et al., 1996). A number of investigators have also 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 (e.g. Benford et al., 1994; Gutwin et al., 1996; Hill et al., 1995; Isaacs et al., 1996).

1.3.1. Making Activity Visible: The Realist, Mimetic and Abstract Approaches

This brings us to 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, et al. (1997) for many examples.

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 (Benford et al., 1994) to the considerably looser re-representations of social information in the 2-D and 3-D avatars found in various graphical MUDs and 3-D VRML worlds.

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 and Loom (Donath et al., 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 MUDs and MOOs (Bruckman, 1997; Cherny, 1999) and textual chat rooms (Danet et al., 1998).

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.

1.3.2 The Concept of a Social Proxy

A social proxy is a minimalist graphical representation that portrays socially salient aspects of an online situation. Typically, a social proxy shows participants in the situation, and some of their activities with respect to it. As should be evident from the previous section, we are not in favor of making all aspects of activity visible: the choice of which aspects of activity are to be visible, and which are to remain private, depend on the particulars of the situation.

Social proxies have four basic characteristics:

- **Figure-Ground.** A social proxy typically consists of two components: a relatively large geometric shape with an inside and an outside and sometimes other features that represent the online situation or context (e.g. a circle), and much smaller shapes positioned relative to the larger shape (e.g. small colored dots) that represent participants.
- **Relative Movement.** The presence and activities of participants in an online context are reflected in the location and movement of the smaller shapes relative to the larger one. Most often, the relationships and movements of the proxy's visual elements have a metaphoric correspondence to the position and movement of people's bodies in face-to-face analogs of the online situation
- **Public Not Personal.** Social proxies are public representations. That is, everyone who looks at a social proxy for a given situation, sees the same thing. It is not possible for participants to customize

their views of a social proxy. This is important because it is what supports mutual awareness and accountability: I know that if I see something in the social proxy, that all other viewers can see it as well.

- **Third Person Perspective.** Social proxies are represented from a third-person point of view. When I look at a social proxy, I see myself represented in it in the same way that other participants are represented. This opens an important avenue for learning. As I act within system, I can see how my actions are reflected in my personal representation, and thus I can begin to make inferences about the activities of others.

This is rather abstract, so let's take a look at an example. Figure 1.1 shows a social proxy from the Babble system (discussed in section 1.4), that we refer to as "the Cookie." The purpose of the Cookie is to reflect the real time presence and activities of participants in a multi-channel chat-like system. The large circle represents the 'current' conversation (i.e. the one being viewed by the user). The small colored dots (called "marbles") represent people who are logged onto the Babble system, including the user from whose viewpoint we are seeing things. Marbles that appear inside the circle depict participants who are looking at the current conversation; marbles outside the periphery represent people who are logged on to the chat system but in a different chat. Finally, when participants are active—meaning they either 'speak' (i.e. type), or 'listen' (i.e. click or scroll)—their marbles move to the inner part of the circle, and then, when activity ceases, drift back out to the edge over the course of about twenty minutes.

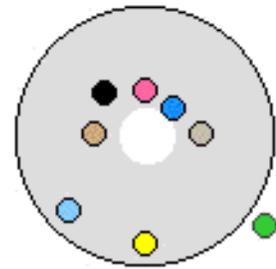


Figure 1.1. The Babble 'Cookie'

Thus, the Cookie shown in Figure 1.1 shows that 'something is happening.' The tight cluster of five marbles around the center core of the circles shows that those participants are engaged (they are either typing, or clicking and scrolling as one often does when participating in the chat). Two other users, depicted by marbles at six and seven o'clock, are also viewing the same conversation, but have not been active. Possibly they are away from their computers, or possibly they are working on other things and ignoring the conversation. The eighth marble (at four o'clock) depicts a user who is logged onto the system but viewing a different conversation; that user may or may not be active—all we can tell is that he or she is connected to the system. Thus, the Cookie shows people are either here or 'around', and if here, how recently they have done something; it also shows when people arrive or depart, in that new marbles appear or disappear (if logging on or off) or move into or out of the bounds of the circle.

We will discuss the Babble Cookie and how people make use of the cues it provides in the next section. However, before we do that, we will present some other examples of social proxies to provide a glimpse of the power and generality of this technique.

1.3.3 The Lecture Proxy

A common situation in the face to face world is that one person will speak to an audience which, by and large, remains quiet. Class room lectures, professional talks, and business presentations are all examples of this. Suppose that we have an online analog of this situation. Perhaps the talk or lecture is being delivered via audio as part of a large conference call; we can imagine that members of the distributed

audience are in their own offices, where they have access to desktop computers, or perhaps they are mobile users who have screen phones.

The social proxy shown in Figure 1.2 assumes that we have some way of identifying who has spoken. As with other social proxies, the large containing shape represents the situation, in this case the lecture, and the small colored marbles represent the participants. The horizontal positions of the marbles reflect a running average of the number or length of comments during the last five minutes. The 'lecturer' starts out on the left, the members of the audience on the right, and the proxy is dynamically updated.



The aim of the Lecture Proxy is to provide a visible representation of the interaction that foregrounds the interactive expectations that define it. As long as the lecture follows its canonical interaction pattern, with the lecturer speaking and the audience being silent, it retains its initial form. However, if a person interrupts with a question or a comment, his or her dot will move a bit to the left, and if the interruptions continue, that person becomes, quite literally, 'out of line' (as shown in Figure 1.2). Because the proxy is seen by everyone, everyone knows (and knows that everyone knows) what is happening. How the group makes use of this information is up to it. The Lecture Proxy may serve to enforce norms about how to behave during a lecture; or it may be used as a signal that people are getting restless and that it is time to shift to a more open discussion (in which case we might imagine that the proxy could shift back to the circular form of the Cookie). The basic point here is that one role that social proxies can play is to make interactive expectations visible, and to highlight how the group is behaving relative to them. We believe that this is an interesting alternative to the technical solution, sometimes employed in conference call situations, of making it impossible for anyone but the speaker or moderator to talk.

1.3.4 The Auction Proxy

Social proxies are not just for conversation, though conversation is certainly one of the most prevalent forms of online interaction. To look at some different cases we turn to the realm of e-commerce and auctions. In the face to face world, auctions are dramatic and intensely social events. People arrive before the auction begins to look at what's for sale; and they not only look at *what* is for sale, they also look at *who else* is looking. Auctions are not just about bidding *for items*, they are also about bidding *against other bidders*. In fact, they serve a wide variety of social ends (Smith, 1989).

However, when we look at online auctions, the social cues that make their face-to-face counterparts such rich and engaging experiences have vanished. The social proxy shown in Figure 1.3 is an attempt to restore some of these cues. The auction itself is represented by the large circle, and participants are represented by colored marbles. The inner core of the circle is a clock, that fills in as the auction's time runs out; the next ring out is the 'bidding room,' where the marbles of people that have placed bids are shown; and those marbles shown around the outside of the circle are those who have showed up to look, but have not (yet) bid. Because online auctions may last for days or even weeks, this proxy is designed to be asynchronous. The way the Auction Proxy works is that

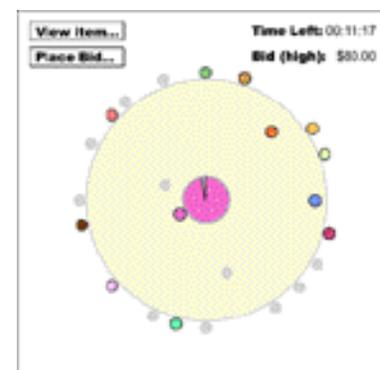


Figure 1.3. The Auction Proxy

participants who have downloaded the web page in the last three minutes have their marbles shown in color; after three minutes pass, their marbles gradually fade to gray, but they remain for the entire duration of the auction. Thus, in Figure 1.3, we see the number of people who have showed up to look or to bid for the entire duration of the auction. The other thing that the Auction Proxy shows is the bidding spread: that is, the closeness of a marble to the center of the circle represents the magnitude of the bid. The innermost marble is the current high bidder, and when another outbids them, the other bidders' marbles are pushed out towards the edge.

Like the other social proxies we've shown, the Auction Proxy creates a public representation of a collective experience: it shows how many people are present, and to what extent they are participating in the activity. Note, as well, that the Auction Proxy does not show certain things (for example, the identity of the bidders, or the number of times an individual has bid). Depending on the circumstances, it may or may not be desirable to show this. As with the Lecture Proxy, the Auction Proxy makes some of the interactive expectations visible: it foregrounds the bidders, particularly the high bidder; and it relegates those who are only watching to the periphery. The Auction Proxy also highlights some new aspects of social proxies. First, social proxies need not represent only synchronous activity, but can represent activities extended over considerable periods of time. Second, the Auction Proxy is being used to create a sense of drama: the clock in the center shows that time is running out; the color of the clock is the same as that of the current high bidder, something which symbolizes the high bidder's current dominance; and the use of color to indicate that someone has connected recently reminds participants that some of the viewers are 'present' and may choose to enter the auction at the last minute with a new high bid.

1.3.5 The Queue Proxy

Auctions are a special case of commercial interaction. For a different example of supporting non-conversational interaction, let's turn to the most ubiquitous hallmark of commerce-oriented interaction: the line or queue.

As experienced users of queues we understand a lot about them. We understand the implications of their length; we make estimates of their speed; we mutter when someone with a problem slows the queue; we become irritated when others 'cut' in front of us; we feel elation if extra personnel show up to handle a lengthy queue. We may decide to postpone a transaction if the queue we are in appears to be moving slowly, only to change our minds if we notice that the growth of the queue behind us has accelerated. However, when we leave face-to-face interaction, the queues have vanished; there appear to be no lines, on-line. But, as anyone who has listened to the 'your call is important to us please remain on the line and your call will be answered in the order in which it was received' message knows, "vanished" doesn't mean that they're gone. They have simply lost most of the cues that transform a really annoying experience into a mildly annoying, or very occasionally, a mildly interesting experience.

Figure 1.4 shows a social proxy designed for online situations where customers are, for example, waiting to chat with a technical support person or a customer representative. The proxy shown in Figure 1.4 conveys information about the length of wait (in this case the color of the marbles encodes estimated wait time), the number of customer service representatives present (shown as triangles), and the length of

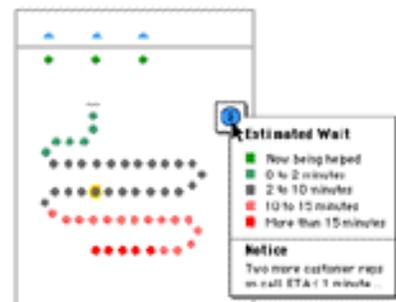


Figure 1.4. The Queue proxy

queue, showing both those in front, and those behind, the user (shown with a gold ‘halo’).

Queues differ from the other situations we’ve discussed—chats, lectures, and auctions—in that we neither seek them out, nor view them as an end in and of themselves. Nevertheless, providing a visible representation of an online queue can support a variety of ends. First, it enables those in the queue, or considering joining it, to engage in the sorts of reasoning about the questions to which we’ve already alluded: Should I join the queue? Should I come back later? How much longer will it be? Second, it provides a potential foundation for social interaction. That is, while face-to-face queues are certainly not necessarily convivial occasions, in some situations, as at a sporting event or a Grateful Dead concert, shared enthusiasms may provide a foundation for opportunistic conversation. Thus, an online queue of people waiting for technical help might benefit if they were able to chat with one another: perhaps some members might have the solutions to others’ problems, or, perhaps the complaints and associated remarks of disgruntled users of a product might serve as a sort of naturalistic focus group for the vendor. Finally, the ability to monitor an online queue can lay the foundations for new types of behavior. Thus, it would become possible to wait in multiple queues at once, and perhaps, if one reached the head of two queues simultaneously, it might be possible to allow someone else to go ahead of one without losing one’s place.

1.3.6 Summary

That’s a look at an array of social proxies. The purpose was to make it clear that the concept of social proxy is quite general. Social proxies can be designed to support a wide range of online interactions, whether they involve conversation or not. They may be synchronous or asynchronous, and they may be associated with activities which are an end in themselves (e.g. auctions), or activities which are simply a means to an end (e.g. waiting in queues). We believe that by providing a shared representation of the activity in which participants are involved, social proxies can help create shared expectations, shared experiences, and can serve as a resource which participants can use to structure their individual and collective interactions. That is, at least, our claim. However, it is important to note that, except for the first, the proxies described so far are concept pieces, meaning that they haven’t been implemented and deployed to real situations. Now, however, we’ll turn to an implemented system, and look at a real example.

1.4. Experience: The Babble System

In the previous section we introduced the concept of social proxies, and discussed examples illustrating the wide range of situations to which social proxies can be applied. In this section, we focus on our experience in designing, implementing and studying a social proxy in the context of an online system called Babble. We begin by describing the vision behind Babble, which provides a further illustration of the importance and utility of making social information publicly available to a system’s users.

1.4.1 The Vision: A Social Approach to 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.

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. (See Erickson and Kellogg (2002) for a more lengthy analysis of the social aspects of producing and using knowledge.)

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.

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 vignette suggests, 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.

Making this sort of activity collectively visible serves one other role. In addition to laying the foundation for a variety of activities that aid in the production and sharing of knowledge, it also has the potential to answer a question that Grudin (1989) has raised with respect to collaborative systems in general: why should those who end up having to do extra work to support collaboration actually do so? The answer is that by making such activities visible, a system of this sort makes it more likely that an organization can see and value such work.

1.4.2 The Babble System

The Babble system (Erickson et al., 1999) represents our first step towards such a knowledge community. Our goal was to be able to support active, conversationally-based communities that could function as part of the business environment. Our hope was to blend the spontaneity and opportunistic nature of ‘water cooler’ or ‘hallway’ encounters and conversations with the possibilities of deeper, more focused talk that might be expected to emerge in a business environment. While there are other requirements for supporting a knowledge community, these aims seemed sufficiently ambitious given that social proxies were untried and the viability of chat-like environments in business situations was unknown (this was before the advent of instant messaging as a workplace phenomenon).

The Babble user interface is shown in Figure 1.5. Starting from the upper left corner, its components are: the user list (all users currently logged on to Babble); the social proxy (as described in section 1.3.2); the topic list (all conversations in the system); and, in the bottom half of the screen, the textual conversation. It should be noted that although we describe Babble as ‘chat-like,’ the conversation model is not quite the same as ordinary chat. It is what we refer to as persistent conversation (Erickson, 1999), or blended synchrony. The basic idea is that Babble conversation may either be synchronous, like chat, or asynchronous, as in bulletin board and netnews systems. Like chat, the conversation accumulates utterance by utterance; like bulletin board systems, the utterances persist across sessions. So that a single conversation in Babble may proceed at a rapid, near real-time pace when many people are around and active, but may then slow and become asynchronous, with people continuing to contribute, but at a much slower pace of a comment per hour, day, week, etc.

A few other features of the Babble user interface are worth noting. First, Babble provides user-specific cues about what has

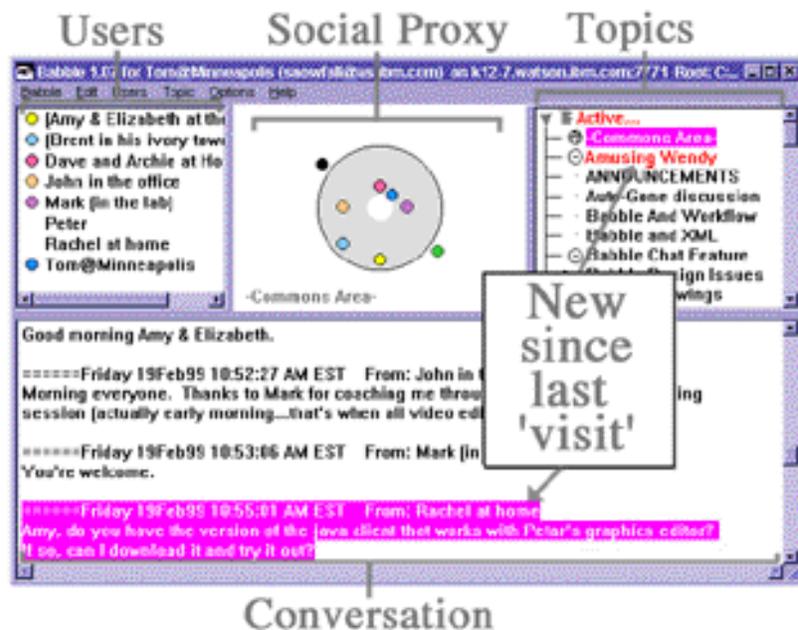


Figure 1.5. The Babble User Interface.

changed recently. Thus, the reverse highlighted text in the Conversation pane indicates that that comment is recent; the topic name (“Amusing Wendy”) shown in red in the Topics list indicates that it has content that has appeared since the user was last there. Also, just to the left of the topic names are spaces where ‘mini-cookies’—miniature social proxies—appear, providing an indication of how people are distributed through the conversation space. Finally, Babble has a number of features that are not on the ‘surface’ of its user interface. Thus, menu commands allow users to create, rename, and modify conversations. Right clicking on a user’s marble provides access to user-specific commands, such as setting preferences and engaging in a private one-to-one chat. And, last, but not least, a menu command also brings up a second social proxy called the Timeline.

The Timeline proxy (Erickson and Laff, 2000) was added because many Babble conversations are asynchronous. The Timeline (Figure 1.6) works as follows: each user is represented by a row in the Timeline; when they are logged on to Babble, they leave a flat trace or line, and when they ‘speak’ they leave a vertical mark or blip on the line. If the line/blip is in color, it means that that user was present/speaking in the conversation currently being viewed by the user of the Timeline; if they were in a different conversation, the line/blip is shown in gray (and the line becomes thinner). As the user mouses over the Timeline proxy the name of the conversation and the user and the time is shown in the upper left corner of the proxy; the user can scroll back through as much as one week of activity. The Timeline proxy also provides access to other functionality via a menu accessed via a right-click on another user’s row (e.g. private chats). For example, in Figure 1.6, we can see that nine people have logged onto Babble (shown by the presence of lines), and that all of them have spent some time in the current conversation (shown by the color/increased thickness of the lines), and that many but not all have ‘spoken’ (shown by the blips). The line being indicated by the cursor shows that the user ‘Peter’ logged on around 11am, made a couple of comments in the “Commons Area” conversation, switched to another topic, and then switched back to the Commons area about 1pm, and then logged off.

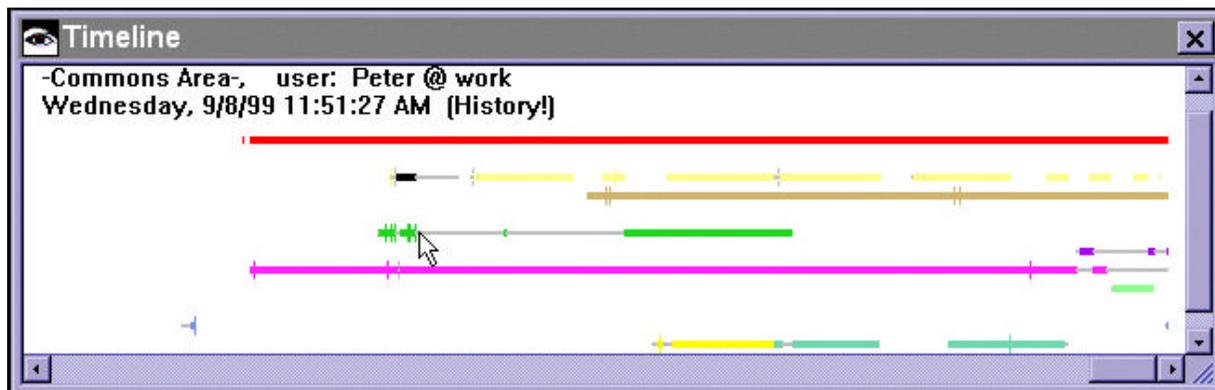


Figure 1.6. The Babble Timeline Proxy.

1.4.3 Experience with Babble

Babble was implemented as a client-server system in SmallTalk. It was first released and used by its developers in the fall of 1997 (Erickson et al., 1999); 1998 saw its first deployments, some of which were studied in depth (Bradner et al., 1999). As of this writing Babble has been deployed to over twenty groups, most though not all of them within IBM (see Thomas et al., 2001 and Erickson and Kellogg, 2002, for recent summaries of deployment experience).

In this section we will focus on our experience with Babble's social proxies. This is not an easy task, because it is difficult to isolate the effect of Babble's social proxies from its other functionality (such as the persistence of its chat, or the use of its private chat as a back channel). So, what we will do is report on our observations of how people make use of the social proxies, and on their comments on the value of the social proxies (in response to several surveys that we conducted).

The Cookie

One phenomenon, originally reported in Bradner, et al. (1999), is called waylay. Waylay refers to the practice in which a user monitors the Cookie for signs of another person's activity (i.e. their marble moving into the center of the circle), and then initiates contact. Contact may be carried out through Babble (entering a comment in a topic or opening a private chat), or through another medium (e.g. the telephone). The Babble Cookie facilitates waylay because, since simply clicking or scrolling triggers the movement of a person's marble, other participants can tell that a person is 'looking at' Babble, even if they 'say' nothing. The indication of 'looking at' is quite valuable; it not only means that the person is at being active on their computer, but it implies that because they are interacting with Babble that they are likely to be receptive to group communication. (After waylay was initially observed, the Babble Cookie was redesigned so that it had a visible center core—this enabled people to see whether someone was touching the center (meaning they'd been active very recently), or whether they had drifted out a little).

Another way in which the Babble Cookie is used is to spot new users. When a person first logs on to Babble, the system assigns their marble a default color: black. As a consequence, black marbles are generally taken to be the sign of a new user, and (depending on the customs of a particular Babble deployment) the user may be greeted, welcomed, or offered help (either in general, or told how to change their marble color). Although we don't know whether the following instance was provoked by a 'black marble,' this user's comment conveys the idea:

“I really like the babble a lot! There is a very friendly and helpful atmosphere.... What triggered me to actually respond at first was however that people had seen me logged in, and welcomed me! The got me over the first hinder of contributing to the chat.”

Another user also observed that Babble seemed to do a good job of supporting first-time users:

“I have been impressed with how easily people have been able to parachute into the community. ... I strongly suspect that elements like the cookie subconsciously help us to relate to each other more easily and make this a more friendly medium than the typical chat.”

More generally, users report that the Cookie is a useful resource for supporting synchronous interaction. Often this has to do with seeing whether there are dots clustered around the center; if so, people often assume that they will get a reply to a question or comment. Here are comments in this vein from two other users:

“The cookie has a definite use for me in indicating how active team members are. It really eases the social interactions knowing whether an immediate reply is likely.”

“The cookie is what makes the difference. ... The cookie makes me feel connected when seeing others read stuff in the same room. Babble without the cookie would be like 'being blind.' Not knowing if it's worth to stay and wait for someone to post a question.”

Finally, users comment that the Cookie helps create a sense of place.

“The cookie is what make Babble useful to me -- it gives me a sense of ‘place’ which I don’t get with [instant messaging, asynchronous discussions, web sites or email].”

“Ah, the cookie... we love the cookie...the cookie is good – our colored dots circulate around to ‘make room’ when someone new joins the conversation – that’s fun. And when someone’s connection dies, they rather dissemble into the ether, angelic like. Which is sort of fun to watch. ... Also, when I’m wondering whether my comments have fallen on deaf ears, I can tell when a response may in fact be on its way when someone’s dot moves back to the center (happens as soon as someone starts typing). So, yes, we like the cookie – it makes me feel like there are actually people in a room with me...”

The Timeline

The Timeline proxy also serves as a resource for users, although it is generally less successful. In part this is due to problems with its user interface, and in it part it may be because the Timeline must be called up with a menu command, rather than being always visible as the Cookie is. Nevertheless, some users find it useful. The most common use is as an indicator of who has been around:

“...I wanted to see if [R__] was on Babble at a particular time [and just not writing anything]”

People also use the Timeline to see if someone has read something that was posted (technically, it is only possible to see that someone has *been* in a topic; however Babble users typically infer that that means someone has *read* the topic’s new content).

“...after a week or so I was able to tell that most people had not read the material I posted...”

Finally, users commented that the Timeline was useful for seeing patterns of usage. It is easy to see when people tend to come on in the morning, and when they tend to depart at the end of the day. Depending on the make up of the community, one may see waves of users arriving, with the Europeans preceding the North Americans by five to eight hours. It is also possible to see whether there are particular patterns of conversation (often people are quite talkative in the morning, and turn their attention to other activities in the afternoon); other events may be visible as well (the most obvious being system crashes, where all traces terminate simultaneously). One user expressed the role of the Timeline with particular eloquence:

“It’s a little like reading an electrocardiogram, the heartbeat of the community. I noticed that I missed [S__] by an hour on Monday morning.... [P__] comes in every so often as a blip. [L__] jumps from space to space....”

While the Timeline is not as generally successful as the Cookie, it nevertheless serves as a resource for those who use it, enabling them to see who has been present or absence, and as the basis for inferences about individual activity (‘who has read what I posted?’) and collective activity.

1.4.4 Summary

Babble is our first step towards designing a knowledge community, an online collaborative system that takes a deeply social approach to supporting the production, elaboration and dissemination of

knowledge. There is a lot that remains to be done. One layer of functionality that is needed has to do with making Babble conversations easy to browse. That is, within an active Babble, it is easy for an energetic group to generate several thousand words a day. This adds up, and the ability to simply scroll and skim or conduct keyword searches of conversations is not adequate. As a consequence we are exploring ways of creating conversation visualizations, where it is easy to see both their content and structure. A second layer of functionality we are investigating has to do with supporting communities of communities. One of the reasons Babble works well is that Babble communities are small enough that people know and trust one another. Thus, Babble participants can offer friendly critiques, off-the-cuff suggestions, and wild ideas that would be less appropriate (and less likely ventured) in a larger, more formal situation. Yet, the *raison d'être* of a knowledge community is to support widespread knowledge sharing. The approach we are currently exploring to resolve this tension is to provide ways in which statistical summaries (but not detailed excerpts) of discussion content might be allowed to 'leak' out of (or perhaps undergo a monitored release from) a Babble, and that an outsider who came across hints of an interesting discussion would have to go through the normal social protocols (e.g. talking to a contact, or perhaps an official 'gatekeeper' within that Babble) to find out more.

Even though Babble is more effective at supporting the exchange of knowledge among a relatively small group than serving as a long term organizational repository of knowledge, we regard it as a successful experiment. Our experience of the last several years has convinced us that social proxies are effective and engaging mechanisms for signaling the presence and activities of participants in online interactions. In particular, it is clear that that the relatively abstract cues produced by social proxies are easily learnable and interpretable by participants, and that we need not be constrained by the requirement to try to mimic familiar cues from face-to-face interactions.

1.5. Summary and Conclusions

In this chapter we've described an approach to designing online systems that we call "social translucence." Its central tenant is that making information about the presence and activity of online participants mutually visible can support social processes—such as peer pressure, imitation, and norming—that help collocated groups interact coherently. However, it is not *just* about making information visible: the word "translucence" signals that we do not propose to reveal all social information; rather, we recognize that there is a subtle interplay between privacy and visibility (as in elections) that shapes interactions in fundamental ways. In face to face situations, individuals and groups exhibit sophisticated understandings of the extent to which the cues they give and give off are visible (as in the "table of contents" example), and use that to adjust their actions. Thus, part of the social translucence research programme involves how to create systems in which participants can understand which of their activities are visible, and the extent and duration of that visibility. To the extent we succeed, we believe that we can design online systems which enable groups to self-organize and otherwise steer their collective interactions.

Babble represents our first steps toward supporting such self-organizing groups. It has proved quite successful in supporting interactions among small communities within IBM (Erickson and Kellogg, 2002; Thomas et al., 2001). Furthermore, our experiences with the Babble Cookie and Timeline indicate that Babble's social proxies serve as useful group resources. The sorts of uses participants report making of the social proxies—seeing when people are around, making inferences about who is listening, noticing and welcoming newcomers, and getting a feel for the overall activity of the group—are just the

sorts of behaviors that we want to support. That is, by making the presence and activities of individuals mutually visible to everyone, we create a foundation on which richer social interaction can be built.

To conclude, 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.

1.6 ACKNOWLEDGMENTS

This work is highly collaborative. Thanks to our colleagues at IBM: David N. Smith for originally creating the Babble prototype; Mark Laff and Peter Malkin for work on the Babble client and server; Cal Swart for supporting Babble deployments; and the members of the Niche Networking and Social Computing groups, and several generations of summer interns, for conversation, inspiration, and general support.

1.7 References

Ackerman MS and Starr B (1995) Social activity indicators: Interface components for CSCW systems. *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST'95)*, New York: ACM, 159-168.

Alexander C, Ishikawa S, Silverstein M, Jacobson M, Fiksdahl-King I, and Angel S (1977) *A pattern language*. Oxford University Press, New York.

Bellotti V and Bly S (1996) Walking Away from the Desktop Computer: Distributed Collaboration in a Product Design Team. *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW '96)*, 209-218.

Benford S, Bowers J, Fahlen L, Mariani J, and Rodden T (1994) Supporting Co-operative work in Virtual Environments. *The Computer Journal*, 37:8.

Bradner E, Kellogg WA, and Erickson T (1999) The Adoption and use of Babble: A field study of chat in the workplace. *Proceedings of the European Conference on Computer-Supported Cooperative Work (ECSCW '99)*, 139-158.

Bruckman A, and Resnick M (1995) The MEDIAMOO project: Constructionism and professional community. *Convergence*, 1:1.

Cohen J (1994) Monitoring background activities. In G Kramer (Ed), *Auditory display*. Addison-Wesley, New York, 439-531.

Cherny L (1999) *Conversation and community: Chat in a virtual world*. CSLI Publications, Stanford, CA.

Danet B, Ruedenberg L, and Rosenbaum-Tamari Y (1998) Hmmm...where's that smoke coming from?—Writing, play and performance on Internet Relay Chat. In F Sudweeks, M McLaughlin and Rafaeli (Eds), *Network & netplay: Virtual groups on the Internet*. AAAI Press, Menlo Park, CA, pp 41-76.

Donath J, Karahalios K, & Viegas F (1999) Visualizing conversation. In: Nunamaker JF Jr and Sprague RH Jr (eds), *Proceedings of the Thirty-Second Hawai'i International Conference on Systems Science*.

Dourish P and Bellotti V (1992) Awareness and coordination in shared workspaces. *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW '92)*, 107-114.

Erickson T (1999) Rhyme and punishment: The creation and enforcement of conventions in an on-line participatory limerick genre. In: Nunamaker JF Jr and Sprague RH Jr (eds), *Proceedings of the Thirty-Second Hawai'i International Conference on Systems Science*.

Erickson T, Smith DN, Kellogg WA, Laff MR, Richards JT, 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*, 72-79.

Erickson T and Kellogg WA (in press, 2002) Knowledge Communities: Online Environments for Supporting Knowledge Management and its Social Context. To appear in: Ackerman MA, Pipek V and Wolf W (eds) *Beyond knowledge management: Sharing expertise*. MIT Press, Cambridge, MA.

Erickson T and Laff MR (2001) The design of the 'Babble' timeline: A social proxy for visualizing group activity over time. *Extended Abstracts: Human Factors in Computing System (CHI 2001)*, 329-220.

Finn KE, Sellen AJ, and Wilbur SB (1997) *Video-mediated communication*. Lawrence Erlbaum Associates, Mahwah, NJ.

Gehl J (1980) *Life between buildings: Using public space*. Van Nostrand Reinhold, New York.

Grudin J (1989) Why groupware applications fail: Problems in design and evaluation. *Office, Technology and People*, Vol. 4.3: 245-264.

Gutwin C, Greenberg S and Roseman M (1996) Workspace awareness in real-time distributed groupware: Framework, widgets, and evaluation. In: Sasse RJ, Cunningham A, and Winder R (eds), *People and Computers XI (Proceedings of the HCI'96)*, 281-298, Springer-Verlag, Berlin.

- Heath C and Luff P (1991) Collaborative activity and technological design: Task coordination in London underground control rooms. *Proceedings of the European Conference on Computer Supported Cooperative Work (ECSCW 91)*, 65-80.
- Heath C and Luff P (2000). *Technology in action*. Cambridge University Press, Cambridge.
- Hill W, Hollan JD, Wroblewski D, and McCandless T (1992) Edit wear and read wear: Text and hypertext. *Human Factors In Computing Systems: The Proceedings of CHI '92*, 3-9.
- Hill W, Stead L, Rosenstein M, and Furnas G (1995) Recommending and evaluating choices in a virtual community of use. *Human Factors in Computing Systems: The Proceedings of CHI '95*.
- Isaacs EA, Tang JC, and Morris T (1996) Piazza: A desktop environment supporting impromptu and planned interactions. *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW '96)*, 315-324.
- Jacobs J (1961) *The death and life of great american cities*. Random House, New York.
- Kendon A (1990) *Conducting interaction: Patterns of behavior in focused encounters*. Cambridge University Press, Cambridge, MA.
- Luff P, Heath C, and Greatbatch D (1992) Tasks-in-interaction: Paper and screen based documentation in collaborative activity. *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW 92)*, 163-170.
- Lynch K (1990) In T Banerjee and M Southworth (Eds), *City sense and city design: Writings and projects of Kevin Lynch*. The MIT Press, Cambridge, MA.
- Pedersen ER, and Sokoler T. (1997). AROMA: Abstract representation of presence supporting mutual awareness. *Human Factors in Computing Systems: The Proceedings of CHI '97*, 51-58.
- Smith CW (1989) *Auctions: The social construction of value*. Free Press, New York.
- Thomas JC, Kellogg WA and Erickson T (2001) The knowledge management puzzle: Human and social factors in knowledge management. *IBM Systems Journal*, 40:4, 863-884.
- Whyte WH (1988) *City: Return to the center*. Anchor Books, New York.