Designing Online Collaborative Environments: Social Visualizations as Shared Resources

Thomas Erickson

IBM T. J. Watson Research Center
snowfall@acm.org

Abstract

How might online collaborative environments be designed so as to better support coherent interaction amongst their users? Drawing from a case study of an example of coherence in an online system, I argue that one way to improve online environments is to provide visualizations that depict the presence and activities of their users. I discuss our approach to creating such visualizations using the concept of the social proxy—a minimalist representation of people and their activities in a particular context—and describe systems we have designed and deployed. I conclude with a series of concept pieces that illustrate the breath of the concept.

1 Introduction

For the last several years, I've been engaged in the study, design and deployment of digital systems that endeavour to support smooth, focused interactions within and among online groups of various sizes. I am particularly interested in conversational interaction that occurs in organizational contexts, and that enables

Proceedings of the 9th International Working Conference on the Language-Action Perspective on Communication Modelling (LAP 2004)

Rutgers University, The State University of New Jersey, New Brunswick, NJ, USA, June 2-3, 2004
(M. Aakhus, M. Lind, eds.)
www.scils.rutgers.edu/lap04/lap04.htm

distributed groups to carry out transactions, create joint work products, and conduct other forms of goal-oriented work.

It is clear that people are very skilled at collaborating in face to face situations. As humans, we are fundamentally social creatures, exquisitely sensitive to the actions and interactions of those around us. We pick up the pace of a presentation if our audience begins to fidget. We forego the grocery shopping if we see that the parking lot is jammed. We decide to eat at the crowded restaurant, rather than the suspiciously empty one next door. The physical world is full of socially produced cues, and a large literature testifies to the many ways in which we make use of such information to govern our own behavior and collaborate with others (e.g., [Goffman, 1963], [Whyte, 1988], [Heath and Luff, 2000]).

However, when we move from face to face interaction to digitally mediated interaction, much changes. The subtle cues that we use to guide our face to face interactions are mostly absent. In digitally mediated settings our attempts to communicate are often awkward. Even when the presence of others is obvious—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. Interactive moves that require little effort in face to face settings—interrupting at the right moment, yielding the floor when someone has a question, or 'going around the table' to do introductions—require much effort in digital systems, if they are possible at all.

In this paper I provide an overview of work aimed at redressing this situation. I begin with an example of a digitally mediated situation that succeeded in supporting a coherent, long-running, productive interaction. Based on an analysis of this situation, I suggest that an important element of supporting coherent online interaction has to do with providing a shared visual representation upon which the group can draw to coordinate and carry out its interactions. I then turn to the question of how to operationalize this conjecture and describe the concept of the "social proxy," and its implementation and deployment as part of a persistent chat system called "Babble." Finally, I present a range of concept pieces that explore ways in which social proxies might support various forms of online interaction.

2 The Anomalous Case of Café Utne

Café Utne is an asynchronous conferencing system aimed at providing an online environment where people can "discuss ideas and issues in a thoughtful and respectful manner." Founded by the *Utne Reader* magazine in 1995, Café Utne continues to be a lively and economically viable forum to this day [Utne, 2004].

At the time this study was conducted (circa 1997-98), Cafe Utne had attracted thousands of members, who had produced thousands of conversations consisting of well over a hundred and fifty thousand posted messages. I had analyzed a number of the conversations that occurred in Cafe Utne, when I discovered a conversation with a some anomalous charactistics. Rather than exhibiting the usual

conversational pattern—an energetic start with multiple themes that gradually disipated into incoherence or wandering, sporadic talk—this conversation was long running, well-ordered, and coherent.

2.1 The Limerick Game

The conversation taking place was a game that involved collectively generating limericks. It stood out because of its length, its remarkably steady rate posting (190 posts per month), its repeated production of well-formed results, and its duration—it had been in existence for over a year when I encountered it, and its descendent is alive and well as of this writing, seven years later. It was also interesting because of its stability in the face of disruptions, and because, over the course of the year it had evolved several persistent conventions.

```
Here's a fun game. We write limericks, each person
contributing a line at a time. You'll recall from this
example that limericks rhyme and scan (iambic pentameter,
and all that) a certain way:

           [example limerick omitted ]

Limit your contribution to one line at a time, at
whichever point the limerick is at when you happen by.
```

Figure 1: The start of the limerick game

Figure 1 shows the first post to the conversation (headers and formatting omitted) which explains the rules of the game: basically, participants can add one line at a line, endeavouring to produce a valid limerick. Shortly after this post, the game is taken up by others. Although the setup was minimal, over the next two days, a dozen people collaborated to produce several limericks, one line at a time, and the limerick conversation was off and running.

While writing limericks is obviously done for fun, the process through which it occurs is as rich as that which underlies more substantive conversations. People aren't just making limericks, they are talking with one another, making jokes, commenting on the limericks and one anothers' contributions, and dealing with newcomers who want to disrupt the game or simply don't understand the rules.

2.2 The Side Talk and Last-does-First Conventions

One of the things that made the limerick conversation interesting is that the participants evolved a variety of conventions for carrying on their interaction. Here I focus on two interrelated conventions: how to mark side conversations (i.e. talk other than limerick-making), and turn taking.

Side talk was, of course, a part of the limerick topic. Making limericks is a game, and like most games, while enjoyable in itself, it also serves as an activity through which participants come to know one another. Thus there was a clear need to have separate side conversations. Early in the limerick conversation, participants experimented with various means of marking their side talk as 'not limerick-making' (e.g. the use of brackets, punctuation and white space), but they generally used the simple expedient of waiting until the limerick that was being constructed was completed; at that point, one or two participants would comment, and then someone would start the next limerick.[1]

A second convention in the limerick topic had to do with turn taking. The opening post of the conversation (Figure 1) defined the convention: "Limit your contribution to one line at a time." This was interpreted, in practice, as meaning that participants could post only one line at a time, *and* that they could not make two posts in a row. This convention was never explicitly invoked as a censure against other participants; instead, participants maintained it implicitly: they would occasionally violate it and then apologize, pointing to extenuating circumstances. However, about a week and a half into the game, what I came to call the "Last-does-First" move came into existence. One participant wrote the last line of a limerick, and without starting a new post, typed a dashed line and wrote the first line of the next limerick. No one complained and, quite rapidly, others adopted the move; it soon turned from an option into a rule, as stated in this bit of side talk after a limerick had been completed without a new one being started:

```
You have to do a first line now!! It's an obligation of the person
who does the last line!
```

Over the ensuing months, the Last-does-First rule came to dominate transitions between limericks: from its introduction as an optional move in the first month, it came to be used in the majority of limerick transitions in the second month, and by the fifth and sixth months was used in all transitions between limericks.

2.3 Clashing Conventions

The advent of the Last-does-First convention caused a problem: it clashed with the side talk convention. The Last-does-First rule, because it was almost always carried out in a single post, eliminated the region between limericks where participants had conducted side talk. Whereas before there was a potential pause in limerick-making every five posts, where participants could chat, after the Last-does-First convention was established, that possiblity vanished. In response, the participants developed a new convention to separate their side talk from their limerick-making by appropriating two features of Café Utne's user interface.

[1] In spite of the synchronous tone of this comment, remember that this is an asynchronous conversation in which inter-post frequency typically ranged from half an hour to a couple of hours.

Figure 2 shows Café Utne's user interface for reading and posting to a conversation. The conversation is shown on a single web page as a sequence of posts in temporal order, each consisting of a header that describes the forum, topic and number of the post, the author's nickname and user ID[2], and the date, time and length of the post. At the bottom of the page is a form that readers can use to enter and post their contributions; in particular, readers should note the first two elements: the nickname field, and the "Hide this posting" checkbox.

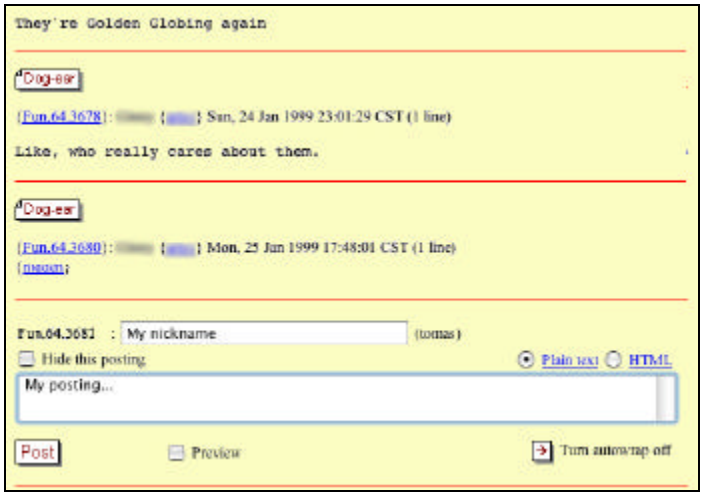

Figure 2: The Café Utne interface for displaying and posting to a conversation.

The nickname field, intended to let users to enter a nickname to supplement their system IDs, was appropriated as a place to make short comments that would be replied to in subsequent 'nicknames': for example, "Groan. Awful pun, stop it Dave!" would be followed by "Sorry about that! ;-)" when Dave next posted. The "Hide this posting" checkbox creates a hyperlink labeled "{hidden}" that allows a post to be displayed on its own page. Although this feature was intended to prevent long postings from dominating the conversation transcript, participants employed this mechanism to segregate their longer bits of side talk from the limerick-making: for example, the hidden link (last post in Figure 2) opens a separate window to display: "What the blazes rhymes with 'oscars'"?

2.4 Some Conjectures

There is considerably more to this account (see [Erickson, 1999]), however this is sufficient for our purposes. The point I wish to make here has to do with the

[2] Here, and throughout the paper, user names and identifiers have been changed or obscured.

apparent ease with which the conversation stayed true to its purpose, and with which the participants developed new conventions to suit their needs.

It seems evident that at least some of this conversation's success has to do with its well-defined nature. It was clearly structured from the start: conventions for the form of the conversation were inherited from the limerick genre, and they, combined with a simple turn taking convention, created a new genre for the collective composition of limericks. It is the conventions of this new genre that are partly responsible for the smoothness of the interaction.

In addition, I believe that it is no coincidence that many of the conventions that structured this interaction were *visible*. Newcomers who have failed to the read the introductory message (a common occurence when asynchronous conversations of this sort grow long), will nevertheless immediately be aware that the conversation is unusual. It will be obvious that the normal post is one line long. Scrolling through the conversation, it is evident that sequences of single-line posts are punctuated with somewhat longer posts (Figure 3 uses a 'greeked' version of part of the conversation to illustrate the visual regularity of its structure). If they pause to read, they will notice that the lines tend to rhyme and scan locally, and even if unfamiliar with limericks, they will still have induced the basic structure of the conversation: one line posts that tend to rhyme and scan in particular ways. Given this (as well as the liklihood that they are familiar with the limerick form, since this probably attracted them to the topic), they are well on their way to making sense of what is going on and being able to participate in a coherent and productive fashion. It is interesting to note that the participants' conventions for keeping side talk out of the flow of limerick-making also served to preserve this visual rhythm.

Figure 3: The series of 1-line posts punctuated by the longer Last-does-First posts produces a strong visual rhythm that is evident to the casual browser

The conjecture I took from this case study was that the visibility of the limerick game's conventions were, in part, responsible for its relative ease of interaction. It made it easier for newcomers to become participants, and also—by making disruptions of conventions publicly visible—made it easier for participants to discuss infractions, thus providing grounds either for their enforcement or their modification and evolution.

3 The Social Proxy

While the limerick game was an interesting and provocative example of a successful and long running conversation, it was clear that much of its success stemmed from its simplicity of content and interaction. How might the visualization of conventions that occurred almost by accident in the limerick game be supported for more usual sorts of conversation?

We use the term "social proxy" to describe our attempt at a general solution. The social proxy is a minimalist graphical representation that portrays socially salient aspects of an online interaction; it is intended to be visible to all users of a system, and updated dynamically. It typically consists of a geometric figure representing an interaction setting, and one or more colored dots that depicts aspects of the presence and activities of participants in that setting. Because, in the general case, the content of a conversation is not as well structured as in the limerick game, social proxies generally take the approach of depicting the structure of the participants' interactions independently of the content.

3.1 Social Proxies in Babble

Figure 4 shows two instances of a social proxy as implemented in a multi-room persistent chat environment called Babble [Erickson, et al., 1999]. The circle represents the chat room the user is currently viewing; dots inside the circle represent others who are viewing the same room, and dots outside the circle represent others logged on to Babble who are viewing other chats. When people in the current room are active (meaning they click or scroll, as when reading, or type, as when 'speaking'), their dots move to the circle's hub; when they cease to be active, their dots gradually drift to the periphery of the circle. Typically, a cluster of dots at the hub of the circle indicates that 'something is going on'—the experience, to a Babble user, is somewhat similar to walking down a street and noticing a crowd: it provokes curiosity and (often) a desire to see what's up. Thus,

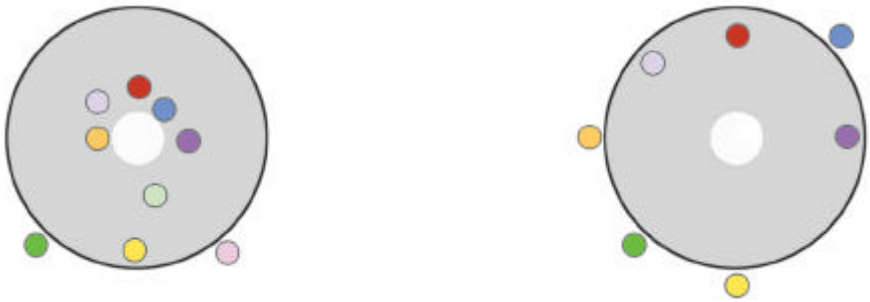

Figure 4: Two instances of the Babble Social Proxy:
(a) a focused interaction and (b) 'not much happening'

in Figure 4a we see an active chat with about half a dozen active participants; in 4b the interaction has ceased and the participants have gone elsewhere or logged off.

Figure 5 shows a second social proxy, also designed for the Babble chat system. In this proxy, each participant's activity is shown in a row of the timeline: if the user is logged on to the system they leave a flat trace, and when they 'speak' (i.e. type) the trace shows a blip. Whereas the circular social proxy in Figure 4 is best suited for showing synchronous or near synchronous patterns of interaction, the timeline proxy of Figure 5 shows patterns of activity over time. Thus, one can see patterns such as what times people log on, whether they tend to stay logged on all day, as well as certain events such as system crashes. Participants who know more about the background of the interaction can recognize other events such as 'morning in Europe' (i.e. European participants log on five or six hours in advance of those in North America), and the convention of saying "hello" when one first logs in to the system (indicated by the blips near the beginning of most traces).

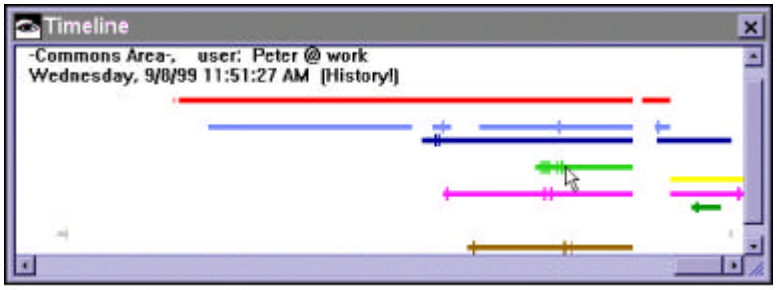

Figure 5: This proxy shows users' presence on the system as flat lines and their posts as blips, thus showing activity over time

Both of these social proxies have been implemented and deployed; intially as part of the Babble application, and subsequently as part its web-based successor, Loops [Halverson, et al., 2003]. As a consequence of observing several dozen deployments, we have quite a bit of experience with how users' of online systems make use of social proxies.

In general, our users report that social proxies are engaging and informative [Erickson and Kellogg, 2002]. In the case of the first proxy, they speak of seeing who is 'in the room,' watching as the dots adjust to 'make room' for a newcomer, noticing a crowd 'gathering' or 'dispersing,' and seeing that people are 'paying attention' to what they say (when other dots move into the center of the proxy after they post). It is also clear that users are able to 'read' Babble proxies, using them to draw inferences about the presence of individuals and the activity of the community as a whole. A user, commenting on the Timeline proxy, remarked:

> 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....

Note that many of the things our users report "seeing" are inferences. The social proxies do not show that people are "paying attention", nor that they are actually talking *to* one another; the proxies show only that someone has clicked or typed. And our users understand this. Making inferential leaps, based on incomplete information, is part and parcel of our life as social beings.

4. Design Explorations

The Babble system, and its successor, Loops, enabled us to explore the reception and usage of social proxies in an online system. It convinced us of the basic viability of the approach, and taught us that users make rich inferences from rather rudimentary information when that information is seen as a product of other participants' activities. This leaves us with the question of what other sorts of online activites might be supported in this way? In this section, I describe a few of the design explorations, presenting mock-ups of interfaces for a variety of online situations. I suggest that, far from being useful just for supporting online chat, social proxies have a wide range of applicability.

4.1 The Lecture Proxy

Since both Babble and its successor Loops dealt with text-based chat, let's turn to a different form of mediated conversation: telephony. Conference calls are a common form of interaction, particularly within large distributed organizations. The advent of digital voice communications (AKA VoIP—Voice over IP) as a viable form of communication [Varshney, et al., 2002], creates new opportunities for supplementing voice with digitally generated visual information.

Imagine a talk or lecture delivered as part of a conference call and accessed by people using phones with screens or or desktop phones adjacent to their computers. The Lecture social proxy, three states of which are shown in figure 6, assumes that we have some way of identifying that someone on a particular connection has spoken, as is possible with Voice Over IP. The background shape represents the lecture 'room;' dots represent people; and the positions of the dots reflect how much they've spoken during the last five minutes. If the lecture is

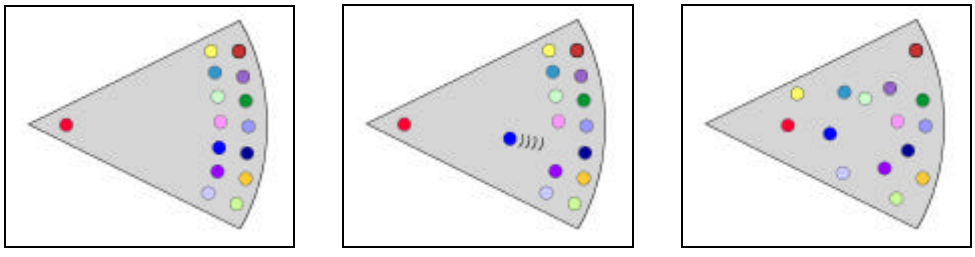

Figure 6: Three instances of a proxy for a telephone lecture: (a) the norm; (b) an audience member speaking; (c) general discussion

going as it 'ought'—with the lecturer speaking and the audience being quiet—the dots in the proxy assume a very regular pattern. 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 middle instance of Figure 6). If multiple audience members speak, their dots move forward as well, imparting a 'raggedness' or incoherence to the visual image, as in the rightmost instance.

What the lecture social proxy is doing is to make the standard convention of lectures—that the lecturer speaks, while the audience remains quiet—visible. The point here is not to prohibit audience members from speaking while the lecturer is talking (indeed, such functionality can and has been implemented in some systems), but rather to eliminate the need for it. In a face to face lecture, it is only convention—and the visibility of adherence to or violation of the convention—that keeps people quiet. In just this manner, the lecture proxy highlights how the interaction is going with respect to the convention, and makes it visible when the interaction is shifting from the norm. By making this shift public, the lecture proxy can serve as an aid in either enforcing a return to the norm, or signaling that perhaps it is time to shift to a different mode of interaction.

4.2 The Conference Call Proxy

In the latter case, it might make sense to shift to a proxy that supports a style of interaction that is more typical of conference calls. Figure 7 shows a proxy that is reminiscent of the first Babble proxy, but which provides some additional functionality that is of particular use in conference call situations. As with the Babble proxy, the dot of the person speaking moves to the center of the circle (in this case the precise center), and then drifts gradually to the periphery after speaking. Thus, in Figure 7, the relative degrees of drift show that the participants in the call are taking turns, and, in fact are 'going around the table' (using the speakers' positions provided by the proxy as a resource), something that is a bit awkward in ordinary conference calls.

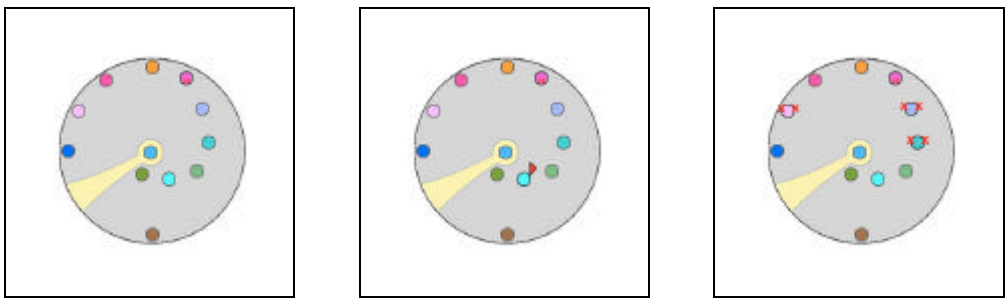

Figure 7: Three uses of a conference call proxy: (a) 'going around the table' (b) asking questions, and (c) signaling that there are difficulties hearing the speaker

In addition, this illustrates another possiblity that social proxies afford: their use as a backchannel. In Figure 7b, the small flag attached to the dot at 5 o'clock indicates that one listener has a question. Similarly, in 7c, participants are able to signal that they are having difficulty hearing (the pairs of "x"s on the dots at 10, 2 and 3 o'clock are "I can't hear" indicators). Note that while one person having difficulty hearing might suggest a problem on that person's line, the simultaneous report of difficulty by several persons suggests that the problem is more likely to be on the speaker's end of the call. Of course, these mechanisms, and others like them, could be used to provide a whole range of non-interruptive communication ranging from getting a show of hands to conduct a vote, to providing a way for someone to indicate that they are stepping out of the call for a few minutes (perhaps by moving the dot just outside of the circle).

4.3 The Auction Social Proxy

So far we've looked at social proxies for supporting various types of conversational interaction. However, the cues provided by social proxies have the potential to contribute to interactions that don't involve conversation. To see an example of this, let's consider the case of auctions.

In the physical world of face to face interaction, auctions are social events [Smith, 1990]. A crowd gathers, inspects the items being offered, and participates in a public bidding process. Participants not only look at what is being auctioned—they also observe who is interested in what, and who bids for what; and they are aware that their own actions and gazes are watched by others. That is, people not only bid *for* items, they also bid *against* other participants. The presence of others contributes to making auctions intensely social and dramatic experiences, as well as enabling them to function as social mechanisms for computing the value of items, asserting the social or professional status of the bidders, and, of course, actually carrying out transactions.

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 auction proxy (Figure 8) is an attempt to express some of the drama of face to face auctions. The large circle represents the auction 'room,' the center circle a clock, and each dot is a participant. People who access information about the to-be-auctioned item are shown around the outside of the circle; if the dot is in color it means that the person has accessed the item information in the last five minutes—after that it turns gray until the person accesses the web page again. If a person bids, his or her dot moves into the circle. The radius of the auction room represents a sliding scale of the bidding to that point in the auction. Thus, in an English style auction (in which bids increase), a new bidder moves to the inner periphery of the

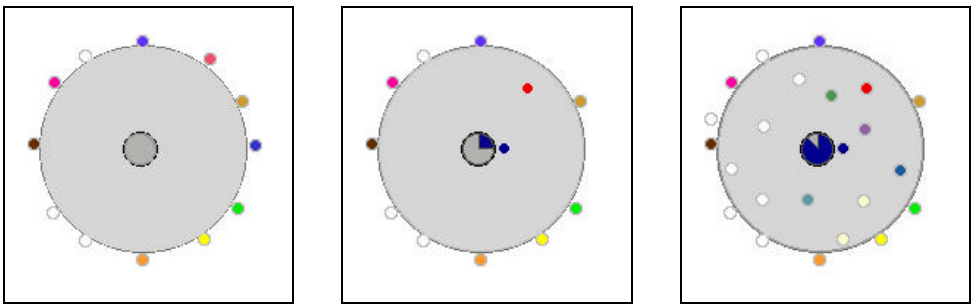

Figure 8: Three instances of the auction proxy: (a) people viewing information about to-be-auctioned item; (b) two bidders; (c) the end game with many bidders

auction room, and other bidders are pushed outwards in proportion to the degree to which they have been outbid.

Figure 8 shows three stages of an auction. In 8a, before the bidding has opened, 12 people have looked at the information describing the to-be-auctioned item. In 8b, part way through the auction, two people have placed bids, moving their dots into the bidding circle. Figure 8c shows the final minutes of the auction, where 11 people have bid, and 11 others are watching—7 of whom, having refreshed the page in the last five minutes (as indicated by the fact that their dots are filled with color), may be waiting until the last few seconds to bid. Although this auction proxy is a conceptual design, a variant of it has been implemented and shown to produce social facilitation effects (Rafaeli et al. 2003).

The auction proxy affords another possibility. Suppose that its final state (resembling that of Figure 8c) is saved as a thumbnail. An online auction system might allow its users to retrieve the thumbnails of auction proxies for all the auctions for a particular type of item that have been carried out over the last year or so. Certainly, such a feature could be useful to both buyers and sellers. Arraying proxy thumbnails in a suitable layout could permit users to get a sense of the nature of the market for that item—the usual numbers of bidders, on-lookers, bids placed, and so on. Besides assisting them in deciding whether it is a useful market to enter into, such a view might show seasonal or other patterns in the market.

4.4 The Building Proxy

Least I leave the impression that social proxies are useful only for online situations, I'll conclude with an example that illustrates how social proxies might be used in any situation where people are spatially distributed, even when the spatial distribution is slight. Figure 9 shows a social proxy for a building: it shows the location and recent movements of people on the floor of a large building (and assumes the availability of very fine-grained location information from a sensor-based tracking system (e.g., [Harter, et al., 2002])). In the building proxy, people

are represented by squares, if they are employees, or triangles, if they are visitors, and the smaller dots associated with each person show their positions over the previous few seconds, thus allowing movement to be detected.

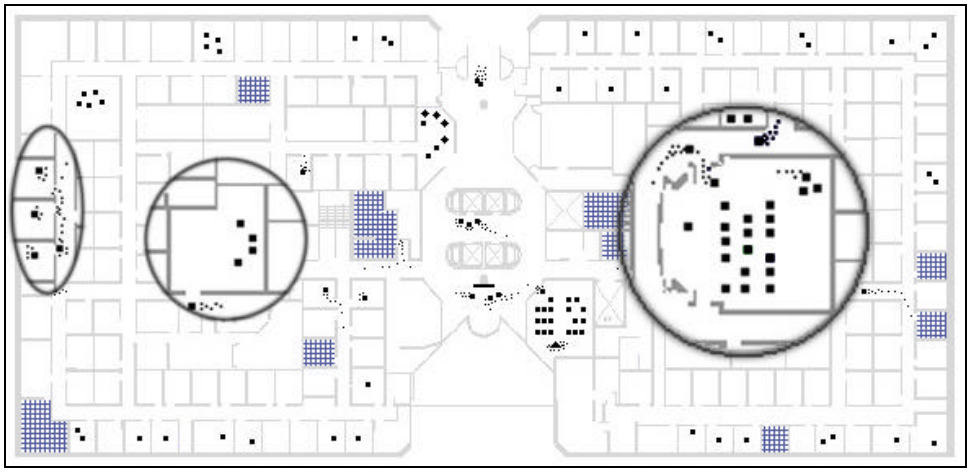

Figure 9: This proxy uses active badge information to show users' locations in a building; dot trails show the last few locations

This proxy raises two issues. First, note that even though no identifying information other than empployee/visitor is provided, the proxy is quite informative. We can see at glance where meetings are being held, which spaces are shared, and which individual offices are occupied. Second, note that someone who is familiar with the context, can read more from the map than a stranger. Thus, as an inhabitant of the building, I can infer that mail is being delivered (in the leftmost circle), make good guesses at who is in the large room indicated by the second circle (it is my laboratory space, and I know who usually sits where), and tell that a meeting (rightmost circle) hasn't started yet (because I know the orientation of the room, and no one is yet at its front). Although location sensing technology in particular, and the sort of visualizations of presence and activity produced by social proxies in general, certainly raise privacy issues, the fact that strangers can not 'see' as much as inhabitants can, seems like a nice property of this sort of representation.

5 Closing Remarks

Humans are social creatures, and as such we've developed a finely honed ability to attend and respond to the actions of those around us. It is this sensitivity that enables us to collaborate gracefully and productively with our fellows. In this paper, I've argued that one way to increase the effectiveness of online collaboration is to design systems that make the presence and activities of their

users visible. Both the Babble and Loops systems, and the design studies I've described, illustrate this approach. By making social cues visible, and allowing visible traces to accumulate over time, we create a public resource that allows people— especially those familiar with the interactive context—to draw inferences about what is happening that can inform the ways in which they participate, and, in turn, may ultimately shape the collective activity of the participants.

This emphasis on visibility raises a number of issues, two critical ones being trustworthiness and privacy. In terms of trust, the role of the social proxy as a collective resource for governing interaction makes it an attractive point of leverage for those who wish to control interactions. It is easy to imagine, for example, unscrupulous online auctioneers who might wish to create counterfeit bidders (just as face-to-face auctions may have their shills). Mechanisms for addressing this sort of concern range from the technical to the social and legal.

With regard to privacy it is important to remember that neither privacy nor visibility are inherently good or bad. Each supports some behaviours, and inhibits others. For example, the perceived validity of elections depends crucially on keeping certain elements of voting behavior private, and others very visible: it is important *both* that a voter be alone in the voting booth, *and* that it be visible that the voter is alone. Likewise, it is important that the act of putting a ballot in the ballot box be visible (so that it can be seen the only one ballot is deposited), but that the content of the ballot be hidden. Similarly, by making careful choices about which cues to reveal or suppress, we can tailor collaborative environments to support particular types of interactions. Privacy and visibility stand in tension with one another, and understanding how to strike a balance appropriate to the situation is one of the critical issues in designing systems that support social interaction.

Acknowledgements

Thanks to my collaborators at IBM for their support, and, in particular, the many conversations that have helped move this work forward. The work on the Babble and Loops applications, and the studies thereof, have been carried out in collaboration with many colleagues, past and present, in the Social Computing and Enhanced Web Experience groups at the IBM T.J. Watson Research Center.

References

[Erickson, 1999] T. Erickson. Rhyme and Punishment: The Creation and Enforcement of Conventions in an On-line Participatory Limerick Genre. *Proceedings of the Thirty-Second Hawaii International Conference on Systems Science.* (eds. J. F. Nunamaker, Jr. R. H. Sprague, Jr.). IEEE Computer Society Press: Los Alamitos, CA, January, 1999.

[Erickson, et al., 1999] T. Erickson, D.N. Smith, W.A. Kellogg, M.R. Laff, J.T. Richards, and E. Bradner. Socially Translucent Systems: Social Proxies, Persistent Conversation, and the Design of Babble. *Human Factors in Computing Systems: The Proceedings of CHI '99*, pp. 72-79. ACM Press, 1999.

[Erickson and Kellogg, 2002] T. Erickson and W.A. Kellogg. Knowledge Communities: Online Environments for Supporting Knowledge Management and its Social Context. *Beyond Knowledge Management: Sharing Expertise* (eds. M.A. Ackerman, V. Pipek and V. Wulf). MIT Press, 2003.

[Goffman, 1963] E. Goffman. *Behavior in Public Places: Notes on the Social Organization of Gatherings.* Doubleday Macmillan, 1963.

[Halverson, et al., 2003] C.A. Halverson, T. Erickson, and J. Sussman. What Counts as Success? Punctuated Patterns of Use in a Persistent chat Environment. *The Proceedings of GROUP 2003*. ACM Press, 2003.

[Harter, et al., 2002] A. Harter, A. Hopper, P. Steggles, A. Ward and P. Webster. The Anatomy of a Context-Aware Application. *Wireless Networks* 8, 187–197. Kluwer, 2002

[Heath and Luff, 2000] C. Heath and P. Luff. *Technology in Action*. Cambridge University Press, 2000.

[Rafaeli, et al., 2003] S. Rafaeli, D. Raban, G. Ravid, and A. Noy. Online Simulations in Management Education about Information and its Uses. *Educating Managers with Tomorrow's Technologies* (eds. C. Wankel and R. Defillipi), pp. 53-80. Information Age Publishing, 2003.

[Smith, 1990] C.W. Smith. *Auctions: The Social Construction of Value.* University of California Press, 1990.

[Utne, 2004] http://www.utne.com/ (as of April 2004)

[Varshney, et al., 2002] U. Varshney, A. Snow, M. McGivern, and C. Howard. Voice Over IP. *Communications of the ACM*, 45(1): 89-96. ACM Press, January 2002.

[Whyte, 1988] W.H. Whyte. *City: Rediscovering the Center*. Doubleday, 1988.