

'Social' Systems: Designing Digital Systems that Support Social Intelligence

THOMAS ERICKSON

Social Computing Group

IBM T. J. Watson Research Center

snowfall@acm.org

Abstract

Large groups of people exhibit social intelligence: coherent behavior directed towards individual or collective goals. This paper examines ways in which such behavior is produced in face to face situations, and discusses how it can be supported in online systems used by geographically distributed groups. It describes the concept of a “social proxy,” a minimalist visualization of the presence and activities of participants in an online interaction that is used to make online social norms visible. It summarizes experience with an implemented system, presents conceptual designs that illustrate the range of situations to which social proxies can be applied, and discusses about how to go about designing these types of visualizations.

Keywords

Coordination, CMC, CSCW, Design, Social Proxy, Social Visualization

Social Intelligence

What do we mean when we speak of “social intelligence?”

There are many possibilities. It might refer to the ways in which a person reasons about the various social collectives – families, teams, communities and so forth – that he or she encounters. Or, conversely, social intelligence might refer to the way in which social collectives ‘reason’ about their individual members – that is, how they operate to determine an individual’s reputation, social status, roles, and other properties of social import. More generally, social intelligence might refer to how social collectives reason about anything, as when participants in an auction collectively determine a price for a rare object (Smith, 1989). Alternatively, we might regard social intelligence as an emergent property of a collection of not-very-intelligent objects such as software agents (e.g., Glasser and Huhns, 1989) or insects (e.g., Gordon, 1999).

I have no interest in trying to discriminate among approaches, for all lead to an array of fascinating and difficult questions. I will simply state that my own interests are in the ways in which human collectives may be said to function in an intelligent manner. That is, for my purposes, I will define social intelligence as *the ways in which groups of people manage to produce coherent behavior directed towards individual or collective ends*. In his recent book, *The Wisdom of Crowds*, James Surowiecki (2004) provides a broad range of examples in which the behavior of social collectives produces seemingly intelligent results.

Why should we be interested in this sort of social intelligence? Two reasons, well illustrated by Surowiecki, are that groups may produce solutions that are qualitatively better than those produced by individuals, and that they may produce solutions more quickly. There is a third reason that this sort of

social intelligence may be of interest that sets it apart from other senses of the term. That is that a solution that is produced by a large collective – think of elections or auctions – may be preferred because, regardless of the quality of the solution or the speed with which it is produced, it may be perceived as fairer or more legitimate. Indeed, the importance of legitimacy is indicated by the fact that a prominent feature of such collective computations are mechanisms designed to safeguard against the subversion of the process by an individual or special interest group.

I am a designer, and the focus of my interest in human social intelligence is how to design technical systems that support it, particularly in situations where people are geographically distributed. To begin with, we will step back and look at how groups function coherently in the world of face to face interaction. Besides seeing more examples of what I mean by human social intelligence, we will also look at some of the properties of face to face situations that support coherent behavior by large groups of people. Next, drawing on these observations, I discuss an approach to designing digital systems that tries to preserve the key properties of face to face situations that support social intelligence. Third, I present a series of examples – some implemented and others purely conceptual sketches – that illustrate the nature and possibilities of the approach. The paper concludes with some claims about how to design systems that support social intelligence.

Social Intelligence in the Face to Face World

A Tale of Two Doors

In the building where I work there is a door that opens from a stairway into a busy hallway. The stairway comes up from the building's cafeteria area, and at lunch time many people carry lunch trays back to their offices. However, this can cause problems. Sometimes, if the door is opened quickly and at just the wrong moment, it can result in a messy accident as the door opens suddenly and a person balancing a tray in one hand and reaching to push the door open with the other, stumbles through.

In recognition of this problem, a small sign was placed on the door: it reads "Open Door Slowly." You can imagine how well this works. The sign may, in fact, be noticed the first few times a person uses the door. However, very quickly, it becomes 'invisible,' in the way signs do, and people return to their usual paces, and accidents occur.

Contrast this solution of putting a sign on the door with a different sort of solution that was subsequently put into practice: inserting a window in the door. The window addresses the problem more effectively, because now a person approaching the door can see if someone is on the other side, and adjust his actions appropriately.

There are three entwined reasons this solution works. First, the window makes the person on the other side visible. Although the sign was also visible, as humans we are predisposed to notice other people: movement and faces engage our attention in a way that signs do not. Second, once we notice that another person is present, that brings our social rules or norms into place. Most of us have been taught from a very young age that we should not slam things into other people; indeed, we are likely to have been taught to assist another person at a door, if they are approaching it with their hands full. Thus, we may very well act in accordance with this norm. Finally, there is a third and more subtle reason that all this works. Suppose that I am not inclined to obey the norm – perhaps I am angry and in a hurry, or perhaps I am a bit of a sociopath and don't care about other people: nevertheless, I am likely to obey the norm because not only can I see you, but you can see me. Thus, you know that I know that you are there, and as a consequence I can be held accountable for my behavior.

There are two lessons that can be taken from this story. First, visibility is important. Visibility enables us to be aware of the context, which in turns allows us to draw upon our social norms to structure our behavior. Second, mutual visibility is important, because it enables people to hold one another accountable for their actions. Once mutual visibility is established – and especially once eye contact occurs – people expect one another to adhere to social norms. Indeed, this is one reason that, in certain circumstances, people may avoid eye contact (as in the old saying that the best strategy for a motorist driving in a foreign country is to avoid making eye contact). There is one other point worth noting about this story: when one person pauses, and allows another holding a tray to pass through the door, it reinforces the norm. That is, while our knowledge of norms may be due to training as a child, norms are constantly reinforced as we interact with one another in a variety of mundane ways.

I suggest that there is nothing special in this story. The world is full of examples of people structuring their behavior according to social norms. As humans, we are fundamentally social creatures, exquisitely sensitive to the actions and interactions of those around us. We notice the social norms in play, act in accordance with them – both because we see others around us doing likewise, and because we know that if we violate the norms, our violations will be visible to others – and by doing so we reinforce the norms.

Collective Cooperation and Social Norms

This use of social information to shape activity is not limited, however, to an individual responding to those around her. If one looks at the behavior of groups, or even the large amorphous entities we call crowds, we see a wide variety of coherent behavior. Let's look at some examples.

Both Erving Goffman (1963) and William Whyte (1988) have remarked on the cooperative skill involved when a pedestrian moves down a crowded sidewalk. Drawing on the analysis of videotapes of pedestrian behavior, Whyte describes what happens as two pedestrians approach one another on a sidewalk and make eye contact:

By their glance they must not only convey the signal, but see if the signal has been acknowledged. A few feet nearer they drop their gaze and make a slight shift in course – to use Michael Wolff's term, the step and slide. The course shift in itself is not enough for a full clearance but it will be enough if the other pedestrian makes a comparable move, as with few exceptions they do." (Whyte, 1988, p 57)

In a crowded situation such as a street crossing or a busy train station at rush hour, these subtle 'dances' are occurring in parallel hundreds of times a minutes. Figure 1, which shows three frames from a video of pedestrians crossing the street in Shanghai, illustrates how masses of pedestrians, about to meet head on, collectively adjust their trajectories to avoid collision. In Figure 1a we see that the crowd at the bottom of the frame is forming columns as they squeeze between stopped taxis. In the next frame (about 5 seconds later), the two masses of pedestrians have formed complementary columns, each having altered their trajectories slightly as they approach one another. Figure 1c shows the situation after about 20 seconds: the four columns of pedestrians have stabilized, two more have formed on the outer edges of the crossing, and a seventh is beginning to form on the far right edge. Although there is no overarching coordination mechanism, other than the change of the traffic light that starts it all off, the members of the crowd have solved the problem of crossing the street *en mass*, without a single pedestrian-pedestrian collision.



Figure 1. Crowds navigating a street crossing: (a) columns begin to form; (b) columns alter directions slightly as they approach one another; (c) at peak flow a series of parallel columns has formed that enables the crowds to efficiently cross.

As in the door story, social norms underlie this coherent behavior. At the individual level, people are trying to achieve their individual goals – most likely getting across the street – while at the same time trying to adhere to the social norm that prohibits shoving or colliding with another person. Most of the coordination of behavior that we see here is due to local adjustments as individuals respond to the activities of their neighbors. Only a few global coordination mechanisms are in evidence: the traffic light that starts and stops the crossing process (reinforced by the assertive Shanghai auto traffic), and the painted crosswalk that demarcates the crossing area (reinforced by the orange barriers at the upper left of picture of the crossing).

Alternatively, consider a somewhat less chaotic example: the behavior of an audience attending a performance of Hamlet at a theatre. When the play is ready to begin, the house lights are lowered: the audience, whose members have been carrying on hundreds of quiet conversations, responds, their collective murmur subsiding into silence, punctuated by the occasional cough. Throughout the performance, silence reigns; audience members who attempt to carry on conversations, or the unfortunate person who has neglected to turn off his cell phone, are subjected to pointed looks from their neighbors. When the play ends, the audience makes an attempt – each individual intentionally acting on his or her own – to give signs of their enthusiasm. Typically the result of this is applause, an individual's hand-claps quickly taken up by others, swelling into a uniform texture of sound. Occasionally, if the play is well received, one or a few individuals may stand up, perhaps leading the rest of the audience to stand as well. On the other hand, if the play is not so well received, a very different situation can result: the attempt at a standing ovation may fail, with a few scatted people standing as the rest remain seated; or worse, even applause may fail to catch on, with distinct isolated claps echoing loudly in the largely silent theatre. Either case – the standing ovation or sporadic applause – is a very powerful expression of the reception of the performance.

Once again, social norms underlie this cooperative behavior. Audience members know when to talk, when to be silent, and when to applaud. This knowledge is reinforced by environmental cues: the lowering of lights signals the start of the performance; other cues, sometimes quite subtle, signal its end; and the brightening of the house lights signal (or at least acknowledge) when the time for applause or ovation has come to an end. Those who are not familiar with the norms of the theatre can manage quite well by imitating those around them. And those less inclined to adhere to the norms, for example a couple this wishes to converse during the performance, may nevertheless feel pressured to comply by the glances, glares or shushes of those around them.

The Role of Behavioral and Environmental Cues

Norms are complex and subtle things. They can be quite flexible: in the theatre, a whispered remark to a companion may be OK, provided it doesn't go on long; and while pedestrian collisions are to be avoided, brushing against another pedestrian is acceptable in crowded situations. Furthermore, in certain situations, norms can be suspended: it is acceptable to shout in a theatre, in the event of a medical emergency; and one would not hesitate to shove a fellow pedestrian out of the path of an oncoming auto. Norms are specific to particular times and cultures: while shouting in the midst of a performance of Hamlet is normally frowned upon, at the time of Shakespeare it was perfectly acceptable for audiences to loudly heckle the performers, comment on the action, and so on. Or, if rather than attending a play, we are at a performance by a stand up comedian, shouted comments from the audience may not just be acceptable but desirable.

How do we, as humans, manage to negotiate these complex and subtle norms? It begins, of course, with early training. An important aspect of growing up and becoming an adult is internalizing the norms of our culture, and developing an increasing sophistication about the ways in which norms are flexible, and the situations in which they may be suspended. But of equal importance is the fact that we learn to take cues from the environment and from the behaviors of those around us. Other people, in particular, are a valuable source of information about the norms in effect. Sometimes this involves watching and imitating others, and at other times the cues may be more direct, as when someone shushes us in the theatre. The social psychologist Stanley Milgram provided a number of demonstrations of the strength of norms, such as his study of the readiness of people to defend a queue they were waiting in against intruders who tried to cut in line (Milgram et al., 1986).

In summary, the coordination evinced in our face to face behavior is in large part governed by social norms, which in turn are supported by physical cues in the environment, and behavioral cues from those around. A large literature, from social psychology (e.g., Milgram, 1977/1992), sociology (e.g., Goffman, 1963) and anthropology (e.g., Whyte, 1988) testifies to the many ways in which we make use of such information to govern our own behavior and collaborate with others. It is this perspective that, in the next section, we will attempt to transfer from face to face interaction to computer mediated interaction among geographically distributed groups.

A Design Approach

As noted earlier, my interest is in how to design online systems that support social intelligence when people are geographically separated. And, as readers will infer from the previous section, the approach I will employ involves invoking social norms to coordinate collective behavior by making physical and behavioral cues visible in online environments.

There are several approaches one might take to making cues about social norms available. Perhaps the most obvious is to design 3D virtual environments that populated by avatars – manipulable visible representations of people and other animate figures (e.g., Churchill et al., 2001). This would allow developers and participants to reconstruct the physical cues present in the material world, and participants to produce and respond to behavioral cues by manipulating their avatars. A second approach is to use video to produce media spaces into which the images of participants are projected and through which they can interact with one another (e.g., Finn et al., 1997). In this case, the ordinary behaviors of participants are made visible and audible, making the production and reception of behavioral cues easy and natural. However, my work is focused on designing systems for use in business environments, and

both of these approaches have drawbacks for this sort of application. They require significant amounts of bandwidth (especially video), and may not scale well to large numbers of people. They require significant amounts of screen space, and thus may not work well on devices with small displays. And they are socially problematic, in different ways: 3D worlds with avatars are associated with games, and are often seen as inappropriate for the workplace; and media spaces, with their transmission of video and audio, may produce problems with privacy, as when uninvolved passersby may unknowingly have their behaviors picked up and broadcast into the media space.

Social Proxies

As a consequence of these limitations, I've been exploring a different line of research. The approach I've developed has to do with creating a shared visualization of people and their activities called a "social proxy." 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 depict aspects of the presence and activities of participants in that setting.

Figure 2 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 users in the that room, all of whom can see what the others type; dots outside the circle represent people in different chat rooms. When those in the current room 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 circle's periphery.




Figure 2. A social proxy for a group chat in the Babble system:
(a) an active chat: (b) after chat has ceased.

Interestingly enough, users find this visualization quite easy to interpret. A cluster of dots at the hub of the circle is taken to mean that 'something is happening' – 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 go find out what is happening. Figure 2a shows a chat with about half a dozen active participants; in the instance of the chat proxy in Figure 2b, we can see that the chat has ceased and the participants have stopped typing (because the dots have drifted to the circle's periphery), gone to other chat rooms (the dots outside of the circle), or logged off.

Figure 3 shows a second social proxy, also designed for the Babble chat system (Erickson and Laff, 2001). 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 type the trace shows a blip. Whereas the circular social proxy in Figure 2 is best suited for showing

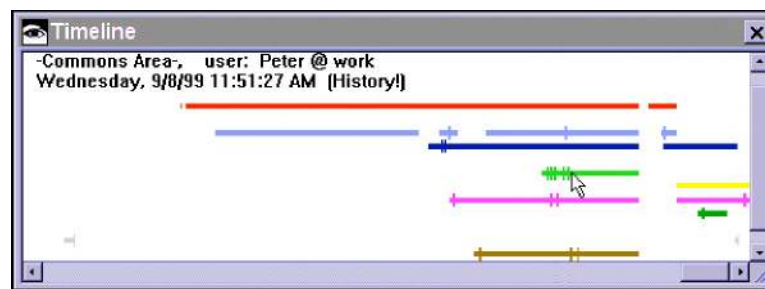


Figure 3. The timeline proxy shows users' presence in the chat room as flat lines and their posts as blips, thus showing activity over time.




Figure 4. Three instances of the lecture proxy: (a) the norm; (b) an audience member interrupting; and (c) many audience members speaking, which violates the norm

synchronous or near synchronous patterns of interaction, the timeline proxy in Figure 3 shows patterns of activity over time. Thus, one can see patterns such as the times of day that 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).

To see how a social proxy can make norms visible and aid in their reinforcement, let’s look at the slightly more complex example shown in Figure 4: the lecture proxy. Imagine a talk or lecture delivered as part of a conference call and accessed by people using phones with screens or phones adjacent to their computers. The lecture proxy, three states of which are shown in Figure 4, 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 dots are positioned according to how much each person has spoken during the last three minutes: the farther to the left a dot is, the more that person has spoken. Thus, if the lecture is proceeding according to the lecture norm – with the lecturer speaking and the audience being quiet – the dots in the proxy assume a very regular pattern (Figure 4a). However, if an audience member interrupts with a question or a comment, his or her dot will move a bit to the left, and if the interruptions turn into something longer than a question, that person becomes, quite literally, ‘out of line’ (Figure 4b). Were this behavior to be taken up by others, their dots would move forward as well, imparting a raggedness or incoherence to the visual image (Figure 4c).

What the lecture proxy is doing is to make the standard norm 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 by reminding participants of the norm. In a face to face lecture, it is only the norm – 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 actual interaction is playing out with respect to the norm, 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 the group that perhaps it is time to shift to a different mode of interaction.

Experience with Social Proxies

The first two social proxies have been implemented and deployed; initially as part of the Babble application (Erickson et al., 1999), and subsequently as part its web-based successor, Loops (Erickson et

al., 2006). Babble and Loops have been deployed to a number of groups (both local workgroups and globally distributed teams) within IBM over a period of about five years. 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 (see Bradner et al., 1999; Erickson and Kellogg, 2003; Halverson et al., 2003).

In general, our users report that social proxies are engaging and informative. In the case of the chat proxy, they speak of seeing who is 'in the room,' watching as the dots adjust to 'make room' for a someone who has just logged on, noticing a crowd 'gathering' or 'dispersing,' and observing that people are 'paying attention' to what they say (e.g., when other dots move to the hub of the proxy after the observer posts a comment). 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. One user compared Babble's timeline proxy to an electrocardiogram, and commented on the differing of movement patterns of various users; others report that they can see conversations between people (e.g. as parallel series of blips on different rows of the timeline proxy).

Note that many of the things our users 'see' in the proxy are not strictly correct; rather, the users are making inferences. The proxies show only that someone has clicked or typed. The inference that someone who begins typing after you have just posted a comment is in the process of replying to you and is thus paying attention is often correct, but it is not guaranteed to be so. Similarly, although parallel series of 'blips' in the timeline proxy are produced when two people are in dialog, it is also possible that they are not talking to one another – yet the common assumption is that parallel sequences of blips represent a conversation. What is important to note is that this uncertainty is not a problem. Making inferential leaps based on incomplete information is part and parcel of our life as social beings, and in most cases mistaken inferences pass without notice or are easily repaired.

Another issue raised by social proxies is the question of learnability. To our surprise, at least in the relatively simple proxies that we've implemented thus far, learnability has not been a problem. One reason for this is that our social proxies are part of communications environments. New users who arrive in a Babble or Loops implementation commonly ask those who are already present about "the dots." Or, alternatively, those present, upon noticing a newcomer (the initial color of a user's dot is black, and those with black dots are generally assumed to be new to the system), may offer to explain what's going on. However, even in cases where there are no experienced users around, novices prove to be quite adept at figuring out what's going on. The dialog shown in Figure 5 occurred when a large number of new users entered a newly created Babble at the same time. In the dialog (which omits other threads of chat interleaved with that shown), we see users speculating about the meaning of the dots, and then gradually figuring it out as they watch the effects of their own actions on their dots.

Nik:	Interesting, why are some dots closer to the middle and some more towards middle/outer radius?
Bob:	Hmmm... need more testing to find out!
Tom:	Nik, it looks like an activity statement... the longer your idle the further from the center you are.
Nik:	Thanks for the info Tom, I'm going to see if I move in closer as a result of typing this message...
Pat:	as soon as u send a message u get closer to the center

Figure 5. A transcript of a Babble conversation in which a group of newly arrived users figure out

This example illustrates one other interesting feature of social proxies: they can serve as 'ice-breakers,' by providing a group of strangers with an immediate topic of conversation. William Whyte, an anthropologist who has studied pedestrian behavior in urban settings, uses the term "triangulation" for

this phenomenon where an object, person or event provokes conversation among strangers (Whyte, 1988). Another example of triangulation comes from the earliest versions of Babble. When first released, Babble did not provide a means for allowing users' to select the colors of their dots; instead, a color was assigned randomly, using the user-changeable nickname as a hash into the color table. That is, as users changed their nicknames, their dots' colors changed, apparently at random. This would occasionally result in someone trying to set their dot to a particular color, trying variation after variation on their nickname, and other users would notice the nickname change and begin talking – making jokes, offering suggestions, with the results that the users' dots gathered together into the center of the circle. Sadly, this convivial behavior vanished when we implemented a color picker for the dots, providing an interesting illustration of the point that supporting ease of use of at the individual level may not be an unvarnished good.

Design Experiments: Some Concept Pieces

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, showed that users attend to the information about the presence and activity of others, and that they make rich inferences from that rather rudimentary information. This leaves us with the question of what other sorts of online activities might be supported in this way? In this section, I describe a few of the design explorations, presenting conceptual designs for 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 ranging from facilitating small group interactions to supporting coordination in very large groups.

The Conference Call Proxy

Figure 6 shows a proxy that is reminiscent of the Babble chat proxy, but which provides some additional functionality that is of particular use in conference call situations. As with the chat 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 the person has stopped speaking. Thus, in Figure 6, 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 quite awkward in an ordinary audio-only conference call.




Figure 6. Three instances of a conference call proxy: (a) 'going around the table' (b) asking questions, and (c) signaling that there are difficulties hearing the speaker

Furthermore, the conference call proxy can serve as a backchannel through which audience members can provide feedback without interrupting the speakers. In Figure 6b, the small flag attached to the dot at 5 o'clock indicates that one listener has a question, and, in 6c, the pairs of x's on the dots at 10, 2 and 3 o'clock indicate that those participants are signaling that they are having difficulty hearing. 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. 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). While the above design is a concept piece, a descendent of this proxy has been implemented as part of IBM's Rendezvous audio conferencing service; similar visualizations have been implemented in other audio conferencing systems (Moors, 2002; Yankelovich et al., 2004). Initial experiences indicate that they support a rich variety of social behaviors (Yankelovich et al., 2004, Kellogg et al., 2006).

The Auction Social Proxy

So far we've looked at social proxies for supporting various types of conversational interaction among relatively small groups. However, the cues provided by social proxies have the potential to contribute to interactions that don't involve conversation, and that can support coherent activity in larger groups. 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. 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. Participants not only bid *for* items, they may also end up bidding *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 (Smith, 1989).

However, in online auctions the social cues that make their face-to-face counterparts such rich and engaging experiences have vanished. The auction proxy (Figure 7) is an attempt to restore some of these cues. The large circle represents the auction 'room,' the center circle a clock, and each dot a participant.




Figure 7. 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.

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

Figure 7 shows three stages of an auction. In 7a, before the bidding has opened, 8 people have looked at the information describing the to-be-auctioned item. In 7b, part way through the auction, two people have placed bids, moving their dots into the bidding circle. Figure 7c shows the final minutes of the auction, where 5 people have bid, and 8 others are watching – 3 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 used to demonstrate that an increased degree of virtual presence cause people to stay longer, do better and exhibit increased preferences for such auctions (Rafaeli and Noy, 2002; Rafaeli and Noy, 2005).

A Task Proxy for Organization-wide Activities

The proxies described so far can support interactions among groups ranging from a few to a few dozen to perhaps, in the case of the auction proxy, a few hundred; in this section, we look at proxy designed to support coordinated action among potentially thousands of people. The case in point involves tasks that require simple but coordinated actions by large numbers of people. Consider the following example.

In June of 2003, a worm appeared on our organization’s internal network. The IT department sent a broadcast email, detailing the measures to be taken – installing a patch, updating anti-virus definitions, and scanning all machines –and stressing the need for rapid compliance. To ensure compliance this email was passed down the management chain – executives emailed their third line managers, third lines emailed their second lines, second lines emailed their first lines, and first lines emailed the employees who reported to them and requested, as well, that each employee report back when he or she had done the task. Upon receiving the email, one of several things happened: some employees did the task and reported back promptly; others did the task but forgot to report back; and still others deferred the task.

Although the ‘worm task’ itself was simple, managing it – and especially ensuring that all employees had carried it out – required a disproportionate amount of time and effort by the organization’s management. There were three problems. First, responses were scattered through managers’ email queues, requiring effort by the manager to locate responses. Second, responses were usually embedded in the email, and not readily apprehensible without opening each message. Third, because this task was just one of many with which employees and managers needed to cope (other tasks included certifying that required training had been completed, verifying that business guidelines had been read; and completing inventories of equipment), the multiple instances of such tasks contributed to general information overload and attention management problems.

This type of task, which we generically refer to as complete-and-acknowledge tasks, is relatively common in large organizations. The worm example described above led us to design a task proxy to facilitate the performance of tasks that require the coordination of the efforts of many individuals (Erickson et al., 2004). Figure 8 shows an example of this task proxy for a single work group at four points in time. Each hexagon represents an individual, and the cluster of hexagons represents a work

group. The hexagon of the user who is viewing the task proxy is marked with an asterisk; the hexagon with the dark chevron (the hexagon at the upper left of the group) represents the group's manager. A hexagon's color represents its user's state with respect to the task: for the proxy shown, white means that no state has been entered, yellow (light grey) signifies that the task is "in progress," and green (dark grey) indicates that it has been "completed." Thus, the progression shown in Figure 8 depicts the gradual completion of the group's task over time as attested to by the participants.




Figure 8. A task proxy for a single group at four points in time, as the group's task status changes from not started to completely finished.

Figure 9 shows a screenshot of an implemented version of the task proxy embedded in its user interface and scaled up to represent a third line organizational. As before, hexagons represent individuals and bundles of hexagons represent groups, but in addition it shows groups of groups (second line departments), and so on: when a user mouses over an element of the task proxy, it dynamically pops up borders, shadows, and (to the right) labels that show the location of the person or group in the organizational hierarchy.

The proxy works as follows. Users can see, at a glance, the overall distribution of task states across the organization. When a user mouses over her own hexagon, her name is displayed in the lower margin, and she can click on it to pop up a dialog that reveals more information about the user, task, and status. She can also use this dialog to change the task status – note that this is preferable to the system automatically reporting completion of the task, because what is desired is that individuals take responsibility (and can thus be held accountable) for the completion of their bit of the task. Because tasks differ in their sensitivity, a "privacy policy" determines whether users can see the names attached to other hexagons, or whether they only see the distribution of task states across the organization. Depending on the privacy policy, a user may be able to see only his name, on the names of those in his immediate work group, only the names of those in his department, and so on.

The task proxy as envisioned here enables two sorts of things. First, it permits the overall status of a task to be visualized. This permits either centralized or decentralized management of the task. That is, it enables a manager to exercise oversight over those components of a task for which he or she is responsible; but also, because the task proxy is visible to everyone, it means that others – other managers, technical leaders without management responsibility, or regular employees – can step in to help a group complete its task. In addition, the visualization makes it possible for other social phenomena to come into play. Thus, one user, noticing that everyone else in her work group has completed the task, may feel some peer pressure to complete it as well. Or, another user may postpone the task, until a local expert has performed it, reasoning that once the expert has completed it, he will have a source of assistance, if

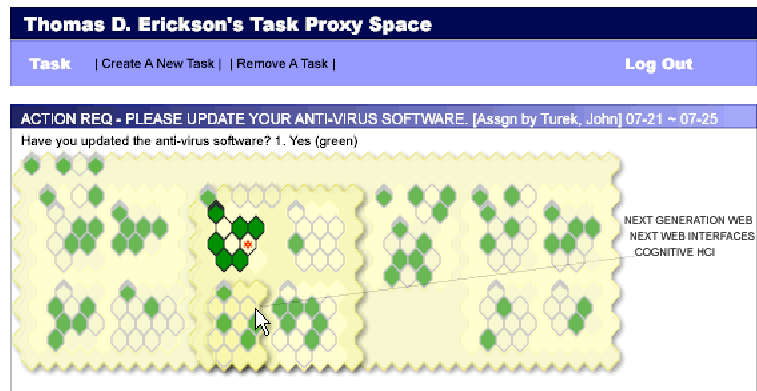


Figure 9. A task proxy for a third line organization of about 140 people divided into four departments and fourteen work groups.

needed. Second, the task proxy provides a contextualized means of communication that is tied to the task and its state: for example, the manager of a task might want to send email to only those who have not yet started it.

Summary

As the examples in this section indicate, social proxies can be applied to a wide variety of online situations and group sizes. Whether facilitating small group conversations, or showing the state of task across an organization, all take the approach of providing information about the presence and activities of participants in a shared visualization. Our hope is that, just as collections of people coordinate their activities in face to face situations, the approach of creating a shared online visualization of the activity of a distributed group can enable the group to engage in the various sorts of coherent behavior that comprise social intelligence. While we have implemented and studied versions of the chat and timeline proxies, and can point to studies of implemented analogs of the conference call proxy conference call proxy (Yankelovich et al., 2004, Kellogg et al., 2006) and the auction proxy (Rafaeli and Noy, 2002; Rafaeli and Noy, 2005), it is important to acknowledge that this is work in progress, and that some of discussion in this section is based on extrapolation or conjecture.

Six Claims about Designing Social Proxies

In this section we offer our current understanding of how to design social proxies for socially intelligent systems. We have learned quite a bit from our deployments, and we have refined and elaborated those lessons through discussions of our conceptual design work. We believe that we understand significantly more about designing systems that support social intelligence than when we embarked on this venture. On the other hand, as we've noted, this is work in progress. We do not feel sufficiently confident of our findings to propose them as guidelines or findings. Instead, we will offer six claims about how to design social proxies for socially intelligent systems. These may be regarded as working hypotheses or starting points, or even as design points for others to react against (we present only claims that a reasonable person might propose doing the opposite of) – regardless, we think that careful consideration of the claims and their underlying rationale can provide others who follow this path a bit of a head start.

Everyone sees the same thing; no user customization

When we show our concept pieces to potential users, or discuss refinements to the design of existing systems, a common suggestion is that they be allowed to customize the social proxy. One example is that a user of Babble might wish to be invisible to other users; another is that one user might wish to make strangers invisible. While this is a reasonable request if one thinks of the chat proxy as a personal tool, it would undermine the function of the proxy as a coordination mechanism. An important aspect of the power of a social proxy is the knowledge that everyone sees the same thing. If I see something, I know that you see it as well *and* that you know that I know. It is this mutuality that supports people being held accountable for their actions, and that leads to useful social phenomena such as feelings of obligation and peer pressure. Thus, we suggest that designers think very carefully about the consequences of tinkering with the shared nature of the visualization, and the accountability it supports.

Portray actions, not interpretation

Social proxies are designed with a particular usage situation in mind, and thus it seems natural to surface the intended meaning of an activity in the visualization. Babble and Loops are for talking, and thus it seems natural to think in terms of conversations and interruptions and paying attention; the auction proxy is for auctions, and it seems natural to think in terms of bidding and buying. However, systems

often end up being used in unexpected ways, and what the designer intended to be a feature for increasing ease of use (for the imagined situation of use) may become an irrelevant and potentially bewildering feature. Instead, we recommend minimizing the amount of interpretation that is built into the system; let the users interpret – they understand the context better than the system ever will. For example, in the Babble chat proxy, a dot moves to the proxy’s hub when the user types or clicks. While the intent was that this indicate that the user is ‘talking’ (typing), or ‘listening’ (clicking to scroll), we have done our best to make it clear that Babble’s social proxy depicts input level activities and not user intentions. Our users have proved much more adept at providing appropriate interpretations of what a key press or click means than we could ever have built into the system.

Social proxies should allow deception

In the course of our face to face interactions, it is often the case that we go to considerable effort to project impressions that don't represent our underlying feelings (Goffman, 1963). We may feign interest, nod understandingly when we are baffled, and act pleased to meet people we loathe. These are vital social skills, and the last thing a social visualization should do is undermine them. Thus, it is useful that one can feign attention in Babble (by clicking on the screen to zoom one's dot into the middle), and it is also useful that one can feign ignorance (‘Sorry, I didn't see your question – I clicked on the screen when switching to another program’). It is also significant that most experienced users of Babble have figured out a way of bringing Babble to the top of their screen without clicking on it, thus avoiding the revealing movement of their dots into the circle’s center. In summary, in my opinion the common impulse among those designing such systems – to strive for as much accuracy as possible – is misguided.

Support micro/macro readings

Whenever possible, a social proxy should be built up out of many small components which persist. Ideally, over time, information will accrete into recognizable patterns at multiple levels, in what Tufte (1990) has called micro/macro readings. For instance, Babble’s timeline proxy, which depicts the activity (presence and talking) of the group over the last week, shows ‘sleep bands’ (at least for groups that are not globally distributed), and other shifts in activity produced by weekends, holidays, and other more global influences. Both these large scale patterns, as well as their fine structures and the perturbations thereof (e.g. activity in what is normally a sleep band) carry information for those users who understand the context.

Ambiguity is useful: suggest rather than inform

When we discuss social visualizations with engineers, a common concern is how well they scale: this works well for a dozen people, they say, but how about thousands? Our response is that accurately presenting information is not the point of a social proxy; its primary role is to provide grist for inferences, and, in fact, it is less important that the inferences are correct. Our users have proved very comfortable with making best guesses from incomplete information. Thus, it is OK to distort activity, to magnify small amounts of activity, and to dampen large amounts of activity; for example, it is much more important for users to be able to tell whether there are 3 or 7 people present, than whether there are 103 and or 107 present. Ideally, the ambiguity of the visualization should be clear to users.

Use a third-person point of view

Although it might be argued that user’s do not need feedback on their own activity since they know what they’re doing, our experience is that this is quite important. People learn what elements of the social visualization mean by watching it over time, and, particularly, by seeing their own behavior reflected in it. For example, in Babble, we have observed groups figuring out the social visualization by group experimentation (‘I clicked and my dot moved to the center!’). Thus, a social visualization should show its users their own activity as others would see it.

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. The door story illustrates the role social norms play in structuring our behavior, and the importance of visibility (and mutual visibility) in enabling those norms to shape behavior. I suggest that many of the technical systems that we've designed to allow groups to interact are similar to the windowless door: they place the burden of coordinating on the individual, rather than allowing groups to use their social intelligence to coordinate their interactions.

In this paper, I've argued that one way to support social intelligence – the process through which groups come to produce coherent behavior – is to design online systems that make the presence and activities of their users visible. Both the Babble and Loops systems, and the design experiments I've described, illustrate this approach. By making behavioral and environmental cues visible, we create a shared resource that reveals some of the social norms that underlie collective behavior. This, in turn, 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 skills). Mechanisms for addressing this sort of concern range from the technical to the social and legal.

Another issue that is often raised with respect to this approach is privacy. Here, it is important to remember that neither privacy nor visibility are inherently good or bad. Each supports some behaviors, 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 online 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 intelligence (Erickson and Kellogg, 2000).

Acknowledgements

Thanks to those involved in the Babble and Loops projects: Erin Bradner, Jason Ellis, Brent Hailpern, Christine Halverson, Wendy Kellogg, Mark Laff, Peter Malkin, John Richards, David N. Smith, Jeremy Sussman, Cal Swart, Tracee Wolf, and several generations of users. And thanks to my IBM colleagues in the Social Computing and Next Generation Web Interface groups for support and inspiration.

References

- Churchill EF, Snowdon DN and Munro AJ (eds) (2001) *Collaborative Virtual Environments: Digital Places and Spaces for Interaction*. Springer-Verlang.
- Erickson T, Huang W, Danis C and Kellogg WA (2004) A Social Proxy for Distributed Tasks: Design and Evaluation of a Working Prototype. In *The Proceedings of CHI 2004*. ACM Press, New York
- Erickson T and Kellogg WA (2000) Social Translucence: An Approach to Designing Systems that Mesh with Social Processes. In *Transactions on Computer-Human Interaction* 7(1):59-83. New York, ACM Press.
- Erickson T, Kellogg WA, Laff M, Sussman J, Wolf TV, Halverson CA and Edwards DA (2006) A Persistent Chat Space for Work Groups: The Design, Evaluation and Deployment of Loops. In *The Proceedings of DIS 2006*. ACM Press, New York.
- Erickson T and Laff M (2001) The Design of the 'Babble' Timeline: A Social Proxy for Visualizing Group Activity over Time. In *The Proceedings of CHI 2001*. ACM Press, New York.
- Erickson T and Kellogg WA (2003) Knowledge Communities: Online Environments for Supporting Knowledge Management and its Social Context. In Ackerman, Pipek and Wulf, V (eds) *Sharing Expertise: Beyond Knowledge Management*, pp. 299-326. MIT Press, Cambridge, MA.
- 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. In *The Proceedings of CHI 1999*. ACM Press, New York.
- Finn EK, Sellen AJ and Wilbur SB (eds) (1997) *Video-Mediated Communication*, Lawrence Erlbaum.
- Glasser L and Huhns (1989) *Distributed Artificial Intelligence: Volume 2* (Research Notes in Artificial Intelligence). Morgan Kaufman.
- Goffman E (1963) *Behavior in Public Places: Notes on the Social Organization of Gatherings*. Doubleday Macmillan, New York.
- Gordon D (1999) *Ants At Work: How An Insect Society Is Organized*. Simon & Schuster, New York.
- Kellogg WA, Erickson T, Wolf TV, Levy S, Christensen J and Bennett WE (2006) Leveraging Digital Backchannels to Enhance User Experience in Electronically Mediated Communication. In *The Proceedings of CSCW 2006*. ACM Press, New York.
- Milgram S, Liberty HJ, Toldeo R and Wackenhut, J. (1986) Response to Intrusion into Waiting Lines. *Journal of Personality and Social Psychology*, 5:4 pp. 683-689. (Reprinted in Sabini and Silver (eds) *The Individual in a Social World: Essays and Experiments*, second edition, McGraw-Hill, New York.)
- Milgram S (1977/1992) *The Individual in a Social World: Essays and Experiments* (second Edition, eds. Sabini and Silver; originally published 1977). McGraw-Hill, New York.
- Moors T (2002) The SmartPhone: Interactive group audio with complementary symbolic control. In *Proceedings of the 4th International Conference on Distributed Communities on the Web* (Sydney, Australia, April 3-5, 2002).
- Rafaeli S and Noy A (2002) Online auctions, messaging, communication and social facilitation: a simulation and experimental evidence. *European Journal of Information Systems*, 11(3):196-207. Palgrave MacMillan.
- Rafaeli S and Noy A (2005) Social Presence: Influence on Bidders in Internet Auctions. *Electronic Markets*, 15(2). Routledge.
- Smith CW (1989) *Auctions: The Social Construction of Value*. Free Press, New York.
- Surowiecki J (2004) *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations*. Random House, New York.
- Tufte ER (1990) *Envisioning Information*. Graphics Press, Cheshire, CT.
- Whyte WH (1988) *City: Rediscovering the Center*. Doubleday, New York.
- Yankelovich N, Walker W, Roberts P, Wessler M, Kaplan J and Provino J (2004). Meeting Central: Making Distributed Meetings More Effective. In *The Proceedings of CSCW 2004*. ACM Press, New York.