

Putting the There There: Visualizing Community Data

Position Paper for the "Dealing with Community Data" CSCW 2000 Workshop

Tom Erickson
snowfall@acm.org
Social Computing Group
IBM T.J. Watson Research Center

September 2000

1. The Community

Since late 1997 my colleagues and I have been working with an ever-evolving program called "Babble." Babble is system that is a cross between a multi-channel textual chat system and an asynchronous textual bulletin board system, with a lightweight graphical awareness mechanism thrown in. The intent of the system is to support everyday, opportunistic interaction among members of a workgroup, though it has not always been used for precisely this purpose. The particulars of the design are unimportant for the moment; the most important aspects of it are that it allows participants to have either synchronous or asynchronous conversations (and shift seamlessly between them), and that it provides a social proxy (a lightweight graphical awareness mechanism) that provides visual cues about the presence and activities of the participants. (Interested parties should see Erickson, et al., 1999 for a description of the system; Bradner, et al, 1999 for a report of its use and adoption (or not) by various groups; and Erickson and Kellogg, 2000, for the 'big picture.')

What is important for the purposes of the workshop is that, over the ensuing three years, we have

- used Babble as a virtual extension of our workgroup,
- deployed it to various groups (about 10) inside of IBM for research purposes, and
- deployed it to one group (or, specifically, five groups of 20 students each, as part of an undergraduate design class at the University level.

That is to say that we have extensive experience in deploying Babble, collecting data from it, and trying to understanding the ways in which various groups have used it.

2. The Community Data

There are three sources of data that are associated with Babble.

- First, because conversation in Babble is persistent (that is, it stays around forever), we have every utterance ever generated within Babble (except for private one-to-one chats, and lacunae due to system failures)
- Second, the Babble Server keeps a log of relevant events including log-ons, log-offs, posts, topic switches, the initiation and termination (but not content) of private one-to-one chats, and so on.
- The third source of data are field studies of the groups using Babble. Field study methods that we've used to study various deployments of Babble have ranged from questionnaires and interviews to participant-observation ethnography.

It is important to emphasize the centrality of field studies. In our experience, this source of data is crucial to making sense of the other data. Knowing, for example, where participants are physically located with respect to

others (same room, same hallway, same building, different states) radically changes the ways in practices in Babble are understood. Similarly, understanding roles, status hierarchies and the social dynamics of the group is also crucial. (See Bradner, et al, 1999, for examples.)

3. Using the Community Data

We want to use the community data in two distinct ways. First, like other researchers, we want to gain a better understanding of what is going on in the communities we are supporting. Second, a more unusual goal is that we are interested in feeding community-generated data back into the community as visualizations of socially salient activity that we call social proxies. We will discuss each of these uses, and some of the issues they raise, in turn.

3.1 Research: Towards a Better Data Logging System

In this section, I summarize some of the research-oriented issues with which we are grappling. The basic point of this section is that, in the case of Babble, the design our data logging system falls far short of our needs. When we began the Babble project I did not carefully think about (and sometimes did not know) the sorts of analytical questions and issues we would face. Here I describe some of the issues that we hope to better deal with in the design of our next generation system.

3.1.1 Logging State, not just Events

A fact of life in collecting data is that accidents happen: files get corrupted, systems crash, backups fail, and humans err. Our logging system for Babble only (for the most part) logged changes of state. Thus we would know that user 1 logged on at time t1 and logged off at time t2, and could therefore compute the duration of their stay -- assuming that the logs were complete. As noted, they sometimes weren't, and so we often had to go to considerable lengths (e.g. searching the log for a "rebooting" log entry) to compute durations of stay, etc. In the next generation system, we hope to do cumulative logging of certain states, rather than rely on trying to compute it from incomplete data. Other examples of states that we currently infer from the logs are the number of participants present in the same topic (i.e. the audience), and the number of people logged onto the system as a whole.

3.1.2 Merging and Replaying Data

While, in some sense, it is nice to have vast amounts of data, trying to make sense of it is difficult. In analyzing activity in a particular Babble deployment, our three sources of data (conversations, the server log, and field study data) are all separate, and must be manually synthesized. In our next generation system, we hope to be able to automatically combine the conversation and log data, essentially enabling ourselves to replay the interactions in the community, either in real time or fast forward modes, from the perspectives of various users.

3.1.3 Protecting User Identity while Supporting Access by Researchers

We deploy Babble to groups with the explicit understanding that we will be collecting and analyzing their conversations and usage, and take pains to explain what data we collect and what we do not. Among other things, we promise to disguise the identity of users whenever we use the data in presentations; in one deployment, we agreed to disguise the identity of users to all but one member of the research team. The latter, stronger promise, seems to us to be a very desirable one to make. However, as currently designed, our logging system makes this a very cumbersome promise to keep. Basically we hand-code large PERL scripts which sift through hundreds of files doing global replaces; one consequence is that this is that we have deferred this analysis until the deployment is over, rather than allowing the researchers to observe the activity of the

community during the course of the deployment. In our next generation system, we would like to provide a way to allow researchers to directly access the data from the operating community, but to essentially view it through a filter which transforms unique identifiers into pseudonyms that are designed to portray certain user characteristics (e.g. gender). A more ambitious (and unworked out) possibility would be to attach information derived from field studies to the pseudonyms (e.g. roles, relationships, social network information).

3.2 Designing Social Proxies

One of the central goals of our research is to explore the consequences of feeding data generated by a community back into the community. Our working assumption is that making the activity of members of a community visible can have a variety of social effects that promote coherent behavior (and, yes, coherent behavior is not necessarily good). There are a variety of means to this end; the particular strategy we have been pursuing is to create minimalist visualizations of social activity that are driven by the community data. We call these visualizations "social proxies."

3.2.1 About Social Proxies

In the social proxy shown in figure 1, the large circle represents the current conversation (with its text shown below), people in that conversation are shown as colored dots inside the circle, and when a person 'participates' in the conversation (either by typing a comment, or simply by clicking or scrolling in the window as one might do when reading) their dot moves quickly to the inner area of the circle. As a person does nothing, their dot slowly drifts to the circle's periphery.

One consequence of the social proxy is that it makes 'lurking' visible. For example, in figure 1, the most recent poster (Rachel, the black dot at 11 o'clock) can infer from the social proxy that Dave and Mark (the magenta and blue dots at 12 and 1 o'clock) have read her question (because their dots indicate very recent interface actions), but that Amy (the yellow dot at 6 o'clock), the person to whom the comment was addressed, has not. (Note: these are inferences and not certainties; any experienced user of Babble can generate reasons why these might be wrong in this case).

Similarly, figure 2 shows the presence and activity of users over the last week (the area visible in the window represents about six hours). In this social proxy, people are represented by rows, and their presence on the system is represented by a line. If they are/were in the conversation being viewed, the line is colored (otherwise it's gray); when people post a comment, their line gets a little spike. Mousing over a line shows the name of the person and the time being pointed at. Thus, we can see that about nine people were on the system Wednesday morning, and that Peter (the green line that the cursor arrow is over) logged on at about 11:30, posted a number of comments, including a possible dialog exchange with Tom (the magenta line below Peter's), and then switched to another topic (when his line becomes gray). The proxy also reveals some of the norms of the group,

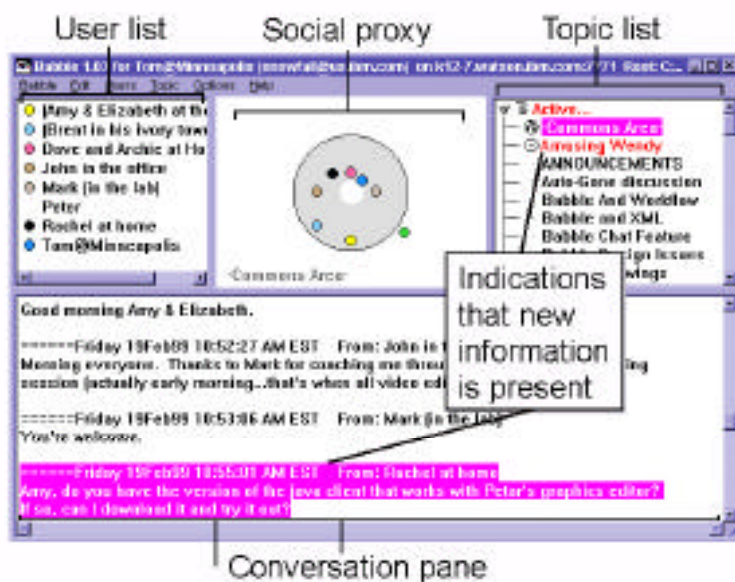


Figure 1. Babble, with a synchronous social proxy.

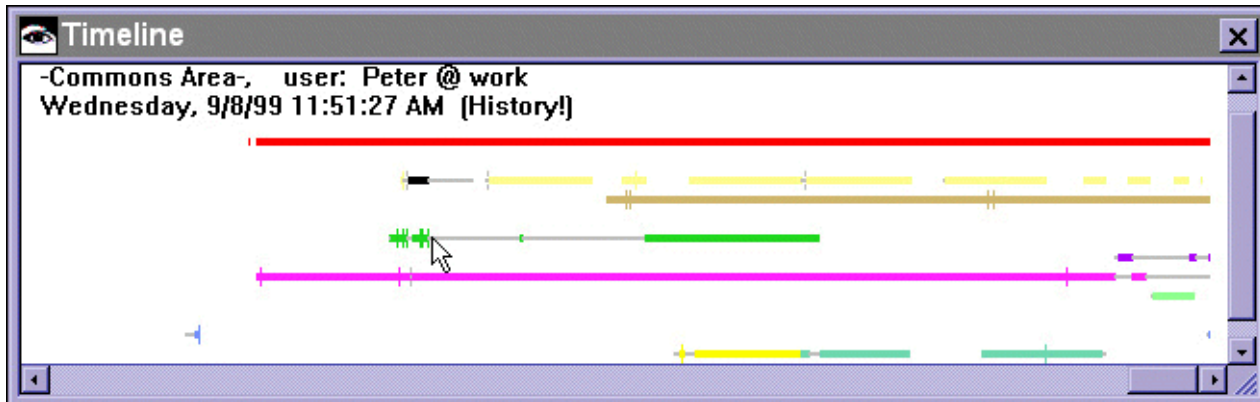


Figure 2. An asynchronous social proxy from Babble.

such as saying something when they first enter the commons area (note, the spikes near the beginning of most of the lines) and the times when people log on in the morning and log out in the evening. One other interesting aspect of the social proxy is that, just as researchers can more readily interpret logged data in the context of field studies, so can the inhabitants of Babble, who understand their social context, read quite a lot from it. Thus, community members might be able to guess the topic of the exchange between Peter and Tom, and what topic Peter was likely to have switched to afterwards. And community members are not surprised by the 'silence' and 'immobility' of the topmost red line: it represents a Babble client used solely to project Babble onto the wall of the semi-public group lab. Its presence has a rather different reading from that of the other lines.

3.2.2 The Tension between Privacy and Visibility

It is quite common for other researchers to express concern about our approach of making community activity visible. While we are certainly aware of the pitfalls of this approach, we also feel strongly that it merits exploration. We suggest that social pressures are important forces for producing coherent behavior within and between groups, and that the approach of making a collaborative system as private as possible is as ill advised as making all activity totally visible. We like to use real-world elections as an example of the desirability of having a mixed model: the perceived validity of an election depends crucially on some of its aspects being very private (what happens within the voting booth), and other aspects being quite visible (e.g. the counting of the votes; or the fact that a voter is alone in the voting booth (the latter being a problematic feature of home-based electronic voting)).

3.2.2 Some Design Guidelines for Social Proxies

While we don't yet feel confident about our approach to designing visualizations of community activity, for the purposes of the workshop we list some of the (often tacit) guidelines we follow. We omit some of the more obvious ones (use rapid animation and other attention capturing effects sparingly) in favor of more controversial guidelines.

Everyone sees the same thing; no user-customization of a social proxy

People frequently suggest that users should be able to customize social proxies so that show only what they want to see. While this is, at one level, a reasonable request, it is deeply contrary to what we are trying to do. A big part of the power of a social proxy is the knowledge that, just as in a physical space, everyone who is 'present' sees the same things everyone else does: I know that you know, and I know that you know that I know. It is this mutuality that supports people being held accountable for their actions. Or to put it differently, it is the

consequences of this mutuality that lie at the root of a noted Cognitive Scientist's strategy for driving in Italy: Don't make eye contact.

Portray actions, not interpretation

Minimize the amount of interpretation that is built into the system; let the users interpret — they understand the context better than the system ever will. Thus, we claim, it is best to reflect actions directly in the visualization: thus, a user's dot moves to the middle of the Babble proxy when they click or drag on the interface. We do not try to establish that they are actually reading the content of Babble, or even paying attention to it. Even if it were possible to somehow reliably make this interpretation, we suggest it is a bad idea because:

Social proxies should allow polite (i.e. deceptive) behavior

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. We may feign interest, nod understandingly when we are baffled, and be very pleased to meet people we loathe. These are vital social skills, and the last thing a social proxy 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 (oh, I'm sorry, I didn't notice your question — I probably just clicked on the Babble screen as I was switching to another program).

Support micro/macro readings

Whenever possible, a social proxy should be built up out of many small components. Ideally, over time, information will accrete into recognizable patterns at multiple levels, what Tufte has called micro/macro readings¹. For instance, in mapping the activity of a mostly co-located group over time, activity will tend to occur during non-sleep hours, and a temporal representation will show 'sleep bands,' and other shifts in activity due to weekends, holidays, and other more global influences. Both these large scale patterns, as well as their fine structures and perturbations thereof (e.g. activity in what is normally a sleep band) carry information.

It's more important to suggest than to inform: ambiguity is useful

Accurately presenting information is not the point of a social proxy. It is more important to provide grist for inferences; it is less important that the inferences are correct. People are 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; ideally, it should be evident to users that such distortions are occurring.

Users should see their own actions reflected in the social proxy (i.e. it should provide a third person view)

People learn what elements of the visualization mean by watching it over time, and, particularly, by seeing their own behavior reflected in it. Thus, the visualization should show them their own activity as others would see it (even though it might be argued that they don't need feedback about what they're doing).

3.3 Conclusions

Conclusions? This is all work in progress, and I have more confidence in the problems encountered than in the solutions suggested.

¹ Edward R. Tufte, *Envisioning Information* (Cheshire, CT: Graphics Press, 1990), pp 37-51.

4. Other Issues

I have struggled quite a bit with ways of characterizing

- a) the activity or 'life' of a community
- b) the activity of life of a single online conversation

I'd be interested in discussing this with others. I find the notion that we could come up with general ways of talking about such things (e.g. a sort of Dow Jones average for community activity) to be quite tempting, though I tend to believe the details of particular situations are usually so important that an attempt at a general metric would either miss all the important stuff about particular communities, or would be gigantic and unwieldy.

5. References

Bradner, E., Kellogg, W., & Erickson, T. The Adoption and Use of Babble: A Field Study of Chat in the Workplace. *The Proceedings of the Sixth European Conference on Computer Supported Cooperative Work (ECSCW '99)*. Dordrecht: Kluwer, 1999.

Erickson, T. & Kellogg, W. "Social Translucence: An Approach to Designing Systems that Mesh with Social Processes." In *Transactions on Computer-Human Interaction*. Vol. 7, No. 1, pp 59-83. New York: ACM Press, 2000.

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

Acknowledgements

The work discussed here is highly collaborative. Thanks to our colleagues at IBM: David N. Smith for creating the Babble prototype; Mark Laff and Amy Katriel for implementations of the social proxy and Java client; Cal Swart for supporting Babble deployments; and all the members of the Niche Networking and Social Computing groups, including two generations of summer interns, for conversation, inspiration, and general support.