

Figuring Out How to Figure Out : Supporting Expertise Sharing in Online Systems *

Thomas Erickson, Christine A. Halverson, Wendy A. Kellogg

IBM T. J. Watson Research Center

May 2005

Abstract

In large organizations, people often take part in processes in which they have no prior experience. In such situations a common problem is figuring out how to begin, and a common solution is the simple expedient of talking to others who have more expertise. However, in large distributed organizations, this expedient is often not so simple. In this chapter, we describe two applications—Babble, and its web-based successor, Loops—which people have turned to this end. We discuss the design of the systems, with particular attention to the ways in which they make people and their activities visible and thus available to one another, and illustrate how people make use of this availability in seeking expertise. We conclude with a discussion of the ways in which the functionality of such systems can aid people in drawing upon one another for assistance and expertise.

1. Figuring Out How to Figure Out

There are two kinds of hard things... there are things that are hard that you know how to do, and things that are hard and you don't have a clue. ...So there you delay delay delay because you don't have a clue. ...I try to figure out how to figure out... who do you ask?

This comment, from an interview with a manager in a large multi-national corporation that we will refer to as Global Corp., illustrates a common problem in large organizations: people often find themselves confronted with problems, situations or tasks that they don't know how to approach. They need to, as the manager put it, "figure out how to figure out." This is the case not only for technical problems that require specialized expertise, but also for more mundane procedural problems. In another interview, a Global Corp manager provides an illustration:

Jil [his manager] walks in here in November and says "good news you've been promoted to STSM." I said "OK, what does that mean?" She said "I don't know, we'll have to figure it out. But congrats!" And she walks out. ...So I read the job library, I asked other people who are STSMs "What do you do?"

Situations like this are not uncommon. Whether it is taking on a new role, filling out a purchase order for an unusual item, or trying to conform to the rules imposed by a rapidly changing legal or regulatory environment, large organizations continually put their employees in the position of doing things they've never done before. This is not arbitrary or capricious behavior on the organization's part; rather it is a natural consequence of their size, the number of internal business processes they support, and their tendency to move employees among positions to broaden their experience and meet the organization's needs.

While organizations provide artifacts such as manuals, checklists, or web-based help pages to assist their employees in navigating business processes, these rarely provide all the knowledge necessary. For example, a process may dictate that something needs to be "approved by an executive from the affected division," without stating how to determine which division is affected, which of several executives in the division might be the appropriate one, or how to go about getting a response from that executive. This is not necessarily a failure to properly document the process. The details of how to proceed are often highly dependent on the situation. Trying to specify all the process details and how

they are mapped onto all possible situations would be a tremendous effort. Furthermore the dynamic nature of the organization means that details of the business process—such as who is appropriate to provide approval—are likely to change over time, increasing the difficulty of keeping the artifact up to date. Thus, the solution alluded to in both examples—talking to others who can help interpret the details as they apply to a particular instance—is a sensible one.

But although talking to others may be sensible, it is not always easy to get started. If a person is situated in a small, locally based organization where most members inhabit the same context, know of one another, and have a shared understanding of everyone's expertise, the problem is simple. But it is another matter when the organization is large or its members are geographically distributed: the issue of who to talk to, or even who to talk to about who to talk to, is often a non-trivial matter.

In this chapter, we take a human centered approach to these issues. Along with a number of other researchers (e.g., see Ackerman, et al, 2003), we see knowledge management as a deeply social process. Humans and the material artifacts in which they embed knowledge are intricately entwined. Even the most thoroughly developed artifact, whether it be a checklist, instruction manual or software system, is embedded in a weave of human practices that helps make it comprehensible and useful to those who draw upon it. It is people, either implicitly or explicitly, who provide the expertise to help unpack, interpret and apply the knowledge embedded in artifacts. Thus, in this chapter, we focus on how people go about obtaining access to others' expertise. In particular, because our research has involved the design, implementation and deployment of systems that support conversation within distributed work groups, we examine how online systems support such expertise sharing.

We begin with a description of the systems in question—Babble, and its web-based successor, Loops. After an account of their design rationale, which is founded on how people orient to one another in face to face situations, we describe the systems' functionality and the general usage patterns we've observed. Then we take a close look at two examples of our system's use, and how they support expertise sharing. We conclude with a discussion of the ways in which the functionality of such systems aids people in drawing upon one another for assistance and expertise.

2. Supporting Interaction at a Distance

In the local realm of face to face interaction, people are exceptionally skilled at interacting with one another. In meetings we wait for a chance to speak, watching carefully for a slight pause that will allow us to enter the conversation without interrupting; failing that, a glance, a gesture, or an 'um' will catch the speaker's attention, who may then cede the floor to us. We interact gracefully with one another even when hardly aware of it: a pedestrian making her way through a rush hour crowd is participating in a series of intricate mutual maneuvers as glancing eye contacts lead to slight alterations of speed or direction that avert collision after collision. Even in the absence of people—as when we decide to forgo grocery shopping when we see that the parking lot is jammed—we draw upon traces of human activity to make inferences that guide our activities. None of this is a surprise: these skills are a basic aspect of being social creatures, and examples pervade the work of sociologists (e.g. Goffman 1963; Kendon 1990), anthropologists (e.g. Whyte 1998; Hall 1983), and other scholars (e.g. Jacobs 1961).

For the purpose of designing systems that support interaction at a distance, it is useful to look more closely at how people coordinate their face to face interactions. Consider the following example. In the building where we work, there is a door that opens onto a busy hallway. On the other side of the door is a stairwell that comes up from near the cafeteria, and around lunch hour there is a steady stream of people carrying lunch trays to their offices or meetings. There is a problem associated with this door: if it is opened quickly from the other side, it can cause rather messy accidents involving those coming up from the cafeteria. To address this problem, a small sign that read "Open door slowly" was posted. As readers may guess, although the sign is effective the first few times a person passes through the door, it quickly becomes 'invisible' in the way that signs do. A better solution—implemented a few years later—is to put a glass window in the door. Now a person approaching from one side can see that someone is on the other side, and adjust their actions appropriately.

Let's examine what makes this solution work. The glass window solution is effective for three reasons. First, and most evidently, it works because it makes what is going on visible. Motion, and human faces, attract attention in a way that signs do not, and thus a person approaching the door is much more likely to become aware of others who are present.

Second, once the person is aware that others are present, our social rules come into play. Most of us have been raised according to standards of courtesy that include not colliding with others, and so we act to avoid that. Third, and somewhat more subtly, even if a person is not inclined to conform to those standards, not only can that person see others through the window, but he knows as well that they can see him, and thus that they know that he knows that they're present: as a consequence, he can be held accountable for his actions. We refer to this triad of effects—visibility producing awareness, awareness invoking social rules, and mutual visibility creating accountability so that social rules are reinforced—as “social translucence” (Erickson and Kellogg 2000), and suggest that it is an important factor in shaping human interactions of all sorts.

However, although this is true for face to face interaction, matters are quite different when people are interacting over a distance via digitally mediated communication systems. When we use instant messaging, email, or even traditional telephony, many of the cues that we use to gracefully coordinate our face to face interactions are absent. Conversational practices that are simple in face to face situations – such as taking turns when talking, or ‘going around the room’ with one person speaking after another – become cumbersome when the conversation takes place over digital media. As designers who are situated in a large organization that is dependent upon mediated communication for its existence, we are interested in addressing this situation.

2.1. Design Rationale

Over the last seven years we have been engaged in designing online conversational environments aimed at supporting small- to medium-sized corporate work groups. Our aim is to design “socially translucent” systems—systems that provide a social context for interaction by providing cues about the presence and activity of participants, and lay the foundation for mutual awareness among participants. We suggest that such systems can, by taking advantage of the human ability to draw inferences from traces of activity, support a range of social processes—imitation, peer pressure, competition—that permit groups to function effectively.




Fig. 1. A social proxy for the Babble system

Our approach to making social information visible employs two tactics: social proxies and persistent conversation. Social proxies are minimalist graphical representations of the presence and activity of participants in an online system. Made available to all users of an online system, their aim is to provide a sense of both the synchronous and asynchronous activity in an online system. Figure 1 shows a simple example, implemented in a multi-user, persistent chat system called Babble (Erickson et al. 1999). The Babble social proxy, which is visible to all users of the system, depicts the current chat room as a large circle, and participants as small colored dots. From the point of view of a particular user, dots shown inside the circle are in the same chat room; dots shown outside the circle are in other rooms. Thus, the social proxy in Figure 1 depicts eight people logged into Babble, seven of whom are in the same room. Dots move to the circle's inner ring when their users type or are otherwise active, and slowly drift to the edge of the circle over the course of about 20 minutes of idleness. Thus, of the seven participants in the chat room, five have been active very recently and the other two have been idle for varying periods. Although this social proxy is very simple, it allows users of the system to gain a rapid impression of how many people are in the same room, and how many of those are active in the chat. Typically, a cluster of dots at the center 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 a desire to find out what's happening.

The other tactic for making social information visible is persistent chat. By this we mean that the transcript of textual chat persists across sessions, with remarks being separated by seconds, as in traditional chat or instant messaging applications, or minutes, days or months, as in bulletin boards or other asynchronous conversation applications. The conversation that takes place is, like more ephemeral chat, a rich source of social information due to both the expressiveness inherent in written language, and the variety of expressive practices that users of text-based communication systems have developed (see Cherny 1999, for many examples). What both of these tactics have in common is that each creates a persistent and public trace that users of the system can draw upon to coordinate their activities.

2.2 Loops and Babble

Our work has been embodied in two systems: a first-generation system called Babble, and a web-based successor called Loops. Loops and Babble were both intended to provide a semi-private online conversation area where members of groups such as work groups, committees, and task forces could have text-based conversations. This section provides an overview of the user interface. We will describe the Loops user interface, as it is a superset of Babble's.

Loops consists of a set of user-definable rooms, each of which can contain a conversation, URLs, static text, and people, as well as user interface elements for seeing who is present, viewing, navigating and modifying the environment. The user experience is that people log in to a Loops server, enter a Loops community, and move from room to room, reading conversations that have changed in their absence, contributing new comments, and encountering other users as they do so. In many cases, users 'hang out' in Loops during the day (monitoring its activity as they do other work on their computers), and taking occasional conversation breaks. The ultimate goal is that Loops should feel like an inhabited place where members of a group can go for social and work related talk.

Figure 2 shows a screenshot of the "SCG" Loop, used by a small but distributed work group. At the top of the screen is the social proxy. It shows that five people are in the room being viewed (the "Commons"), three of whom are active and two of whom are idle. Mousing over the dots that represent people shows their names, and right clicking brings up a menu of person-related functionality, including the ability to have a private, instant message style chat. To the right of the social proxy is a rectangular area called a bulletin board. The bulletin board is a simple text field that anyone can edit, and, in Figure 2, it is shown containing a meeting announcement, followed by a later note that it has been cancelled, and a subsequent reaction to the cancellation. The two lists below the bulletin board show, respectively, the people currently logged onto the system, and the rooms or places that users have created. The places list also enables users to see where people are located in the system (via thumbnail versions of the social proxy next to each room), and where new information has been posted (via highlighting room names in red). To the left of the people and places lists, is the persistent chat pane, where those who are in the room can 'talk.' Since the chat text is time-stamped and persists across sessions, conversations may be synchronous or asynchronous. Finally, peeking out from under the chat pane are slide-out tabs, each of which can contain, like the bulletin board, publicly viewable and editable text and URLs. The two tabs shown in Figure 2 contain a list of contact information for workgroup members, and information about the group's conference call number. All of these elements are public—that is, they are visible to and editable by everyone who has access to the room.

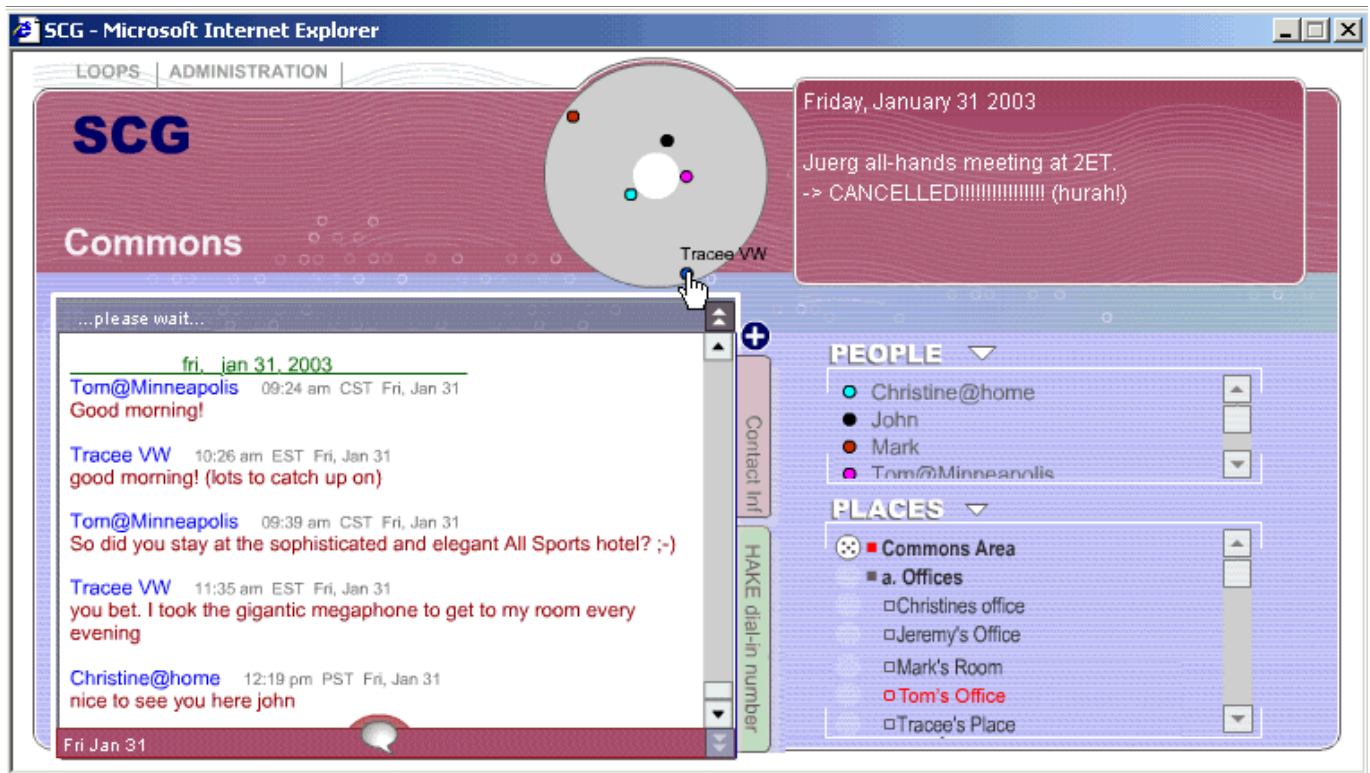


Fig. 2. The Loops system, showing the “Commons Area” of the SCG Loop (screenshot truncated by a third).

The Loops user interface provides access to other functionality as well. There are a variety of options for setting user preferences (e.g., changing one’s dot color, nickname, email address, profile), and indicating that one is “away for a while.” Users can also create and delete rooms, or move them elsewhere in the hierarchy. Both Babble and Loops use sounds to indicate significant activities such as message posting, and the login and departure of participants—this enables users to monitor what is happening, even though the application may not be visible on the screen as they work on something else. Finally, both systems provide facilities for holding one-to-one (or few-to-few) private, ephemeral chats—because the ordinary chat is public and persistent, there is a need for a private and non-persistent back channel.

2.3 Usage Patterns in Loops and Babble

Over the past seven years Babble has been deployed to about two dozen groups, and Loops to about half a dozen, mostly within Global Corp. Deployments have been to two types of groups: relatively small, close-knit but distributed work groups that needed a long-term collaborative space; and larger, cross organizational groups that were either convened for a specific purpose (such as a task force chartered to produce a report), or longer term communities of interest (e.g. members of a particular profession or special interest group) that needed a collaborative space in which members could talk and get to know one another.

Deployments have been studied using a variety of techniques ranging from surveys and interviews to long term participant-observation, as well as analyses of content and system logs (e.g. Bradner et al. 1999; Halverson et al. 2003). As a result of these studies we’ve observed a number of recurring usage patterns that are involved in people drawing upon one another for help. For the most part, these patterns have occurred in both distributed work groups and cross-organizational groups, although some practices are more common in one type of group than the other.

Casual Opportunistic Chat

A frequent use of Babble and Loops is for casual, socially oriented chat. Typically, groups settle on a particular room—usually the “Commons,” the default ‘root’ room that is a part of every deployment—as a place to ‘hang out.’ A typical pattern is that as people log onto the system in the morning, those already present will greet the new arrivals,

and the group will engage in casual chat, banter about the weather, discussion of current events, and so on. Throughout the day, depending on the dynamics of the group using a particular deployment, chat will occur in bursts triggered by arrivals of participants, work related events such as impending meetings, or questions posted by individuals. Because the social proxy shows when particular users have just interacted with the system (generally resulting in an inference that they are paying attention and available to communicate), people often address those who have just become active. Indeed, users have been observed to watch for a person to become active, and to “waylay” them for a chat (Bradner et al. 1999). Although seemingly inconsequential, such practices as greeting someone in the morning, asking about the weather, or engaging in banter about current events provide a lightweight way for strangers to begin to recognize one another, and for acquaintances to maintain and extend existing relationships.

Focused Discussion

Babble and Loops are, of course, used for more focused discussions, often about work-related topics. While this sort of talk sometimes occurs intermixed with casual chat, users often create special rooms for topic or project-centered discussion. For example, users might create a room as a note-taking space and textual back channel for a conference call. Before the call, an agenda may be posted on the bulletin board, and information about the dial-in number, passcode, and URLs for the meeting may be added to one or more tabs. Once the call begins, members may use the persistent chat to take notes, record ‘to do’ items, post questions without interrupting the call, and even to prompt the person representing the group to bring up particular issues during the call. Alternatively, a room might be created to support less synchronous activity, such as reporting and tracking the bugs associated with a development project. In this case, bugs might be discussed in the chat pane, and summarized (with numbers indicating current status) in the tabs or bulletin board. As with casual opportunistic chat, this sort of interaction serves to create new and reinforce existing relations amongst participants.

Catching Up and Leaving Trails

Because the content of Babble and Loops persists, it is quite common for members who have been away on vacation—or have simply missed a meeting or event—to use the system as a way of catching up with what happened in their absence. More generally, participants use the phrase “getting the red out” to refer to the process of catching up—this refers to the fact that the names of rooms containing information the user hasn’t seen are shown in red, and thus on returning from an absence there is often a lot of red in the rooms list. In a more extended example of this, Halverson et al. (2003) describes how a Loops deployed to a distributed software development team was used, in a primarily read only way, by a group of product development managers who were responsible for transforming the application being developed into a full fledged product.

The other side of this usage pattern is that, because users recognize the value of the chat transcript, some adopt the practice of talking to themselves when no one else is logged on. Essentially, they are leaving trails of their activities (“I have to drive to Somers for a meeting now;” “The phone call on Project X went very well!”), knowing that their co-workers will see their words later and possibly want to follow up. Although it is often the case that the residue of the chat—for example, notes taking during a phone meeting—is not thorough enough to convey everything that went on, it usually provides enough hints about what happened to enable latecomers to waylay and question those who were present.

Undirected Queries

One of the most useful aspects of Loops and Babble is their potential to provide their members with access to a vast array of expertise. Yet, especially in cross organizational groups, it is difficult for members to know who to ask their questions of (since they don’t necessarily know who knows what), and they may feel awkward making a direct request of someone they don’t know well. Users typically deal with this tension by posting questions directed to the group as whole, rather than to individuals in particular; this often occurs in the Commons area. Such undirected queries are less awkward because they don’t put anyone on the spot, and because those who answer them may do so at their leisure (since the questions persist over time). Additionally, others may further specify the query by asking questions about the questions, suggest partial answers that others may build on, give pointers and referrals to experts (‘tell Pat that I told you to contact her’), or simply indicate that they, too, are interested in the answer. All of this activity serves to

emphasize that question asking and answering is a legitimate usage of the system, and allows participants to voluntarily enlist others in the process of getting the question answered.

Publishing Personal Information

A common usage pattern in both Loops and Babble is the publication of personal information for the rest of the group. This impulse is most clearly manifested in the creation of types of rooms referred to as “personal places” or “offices.” Although functionally identical to rooms created for group conversation, the simple act of naming a room something like “Asa’s Office” signifies that it is an ‘owned’ place oriented around a particular person. Personal places are often used to provide a profile about the interests and skills of its ‘owner’ (especially in the case of cross-organizational groups). In the first generation Babble system, there was no facility for this, so users resorted to creating ‘chat rooms’ that were intended as information containers rather than as places to talk. Recognizing this usage pattern, Loops added two user interface features—tabs and bulletin boards—to accommodate the publication of information. Thus, in Loops, tabs in people’s “offices” are often used to contain personal resources such as contact information, URL lists, and travel schedules.

2.4 Summary

These usage patterns—although they vary in frequency and nature from situation to situation—are found across the majority of Babble and Loops deployments. They matter because they form the underpinning for the ways in which people draw upon one another’s expertise. Some, like casual opportunistic interaction and publishing personal information, are important because they foster the growth of relationships. Others, such as leaving traces and undirected queries are directly implicated in the processes through which people share expertise, as we shall see in the next section.

3 Finding and Tapping Expertise

Thus far we’ve described the Loops and Babble systems, and discussed a number of common usage patterns that occur across our deployments. We now shift the focus back to our primary theme, which has to do with the way in which people have used these systems to seek assistance and share expertise. This section examines two examples of how the system enables strangers to draw upon one another’s expertise. The first illustrates some of the ways in which systems like Loops and Babble can allow strangers to get to know one another, if just a little bit. The second illustrates how such relationships, even if they are weak, can then be tapped for expertise.

3.1 Ice Breaking: Forming Relations Amongst Strangers

This example takes place in a Babble deployment called ThinkBig. The deployment was a bit unusual in that, rather than being created to serve the needs of a long term community, it was created to support a short term global task force charged with brainstorming new ideas for Global Corp.

Tari [1:15:14]	Are the people outside of the gray circle not in the chat room? For e.g. Sidhu are u not in the chat room?
Gina[1:15:33]	!
Sidhu [1:15:35]	Hello, World!
Tom[1:15:46]	Tari, it looks like they are in a different Babble area
Sari [1:15:59]	The circles move when an individual moves...
Tari [1:17:03]	Sari, you just moved out of the gray circle...what does it mean? Hv you gone to a differnt Babble area?
Nils [1:17:56]	When you enter a room/area, your marble will move to the center and you can tell who else is in the area.
Gerald [1:18:12]	Interesting, is there any significance to the differences in proximity to the center for some dots vs. others? Why are some closer to the middle and some more towards middle/outer radius?

Nils [1:18:29] Hmmm... need more testing to find out!
Gerta [1:18:46] Hi
Tom[1:18:53] Gerald, it looks like an activity statement... the longer your idle the further from the center you are.
Kendra [1:19:10] perhaps the people who came in recently are further away from the core of the commons area?
Gerald [1:19:20] Thanks for the info Tom, I'm going to see if I move in closer as a result of typing this message...
Liam[1:19:53] as soon as u send a message u get closer to the center

Fig. 3 . Segment of Babble chat in which new arrivals get to know one another as they figure out how Babble works. (Reformatted for brevity and readability)

The conversation shown in Figure 3 marks the start of the ThinkBig Babble. The team has just finished their first conference call, and a dozen new users have just logged on to Babble for the first time. This segment—5 minutes of an approximately 20 minute multi-way chat among new arrivals from North America, Europe and Asia—shows the newcomers collectively figuring out how Babble’s social proxy works. In the first part of the segment, Tari** and Tom are exploring what it means when “marbles” (the Babble term for the colored dots) are outside of the circle. That leads Nils to note that entering a room causes a marble to move to the circle’s center, and Gerald to raise the question of what differing proximities to the center means. Tom ventures a conjecture that the distance is proportional to the time spent idle, Kendra suggests the opposite, and Gerald—who has been idle for a bit over a minute and has consequently drifted a tiny ways out—tests Tom’s theory by posting a comment, with Liam confirming the result of his test.

What is interesting here is that the way in which the behavior of Babble’s social proxy forms a focus for discussion amongst strangers. The task of figuring out what is going allows the participants to begin interacting with one another, without requiring them to plunge immediately into content-oriented talk or possibly awkward socializing. Although this example is unusual in the simultaneous convergence of a dozen new users in the absence of experienced users, it is quite typical in that one of the first topics of conversation that a newcomer raises has to do with what the dots and their movements mean. Babble’s proxy, because of its salience and public visibility, becomes a conversation starter.

A variant of this conversation starting behavior is something we call the “black dot” phenomenon. When a new user joins a Babble or Loops community, his or her dot color defaults to black. This has come to be recognized as a hallmark of a new user (since in every deployment of Babble or Loops we have observed, users change their dot colors to something else), and experienced users will welcome the presumed newcomer and point them to the menu item for changing their dot’s color. Furthermore, after the dot’s color changes, those present often engage in banter about the new user’s choice. This is obviously not important conversation, in a substantive sense, but it does provide a way for people to engage in a friendly and non-challenging way, a very desirable property of an environment where strangers encounter one another.

3.2 Tapping (Weak) Relationships

The prelude to this second example is that Bob, a U.S. Global Corp. employee, has received an email from Åke, an unemployed technical person in the United States who is about to move to Sweden for family reasons. Åke is contacting Bob to ask if Bob knows of any Global Corp. job possibilities in Sweden. Unfortunately, Bob doesn’t know about job possibilities there; in fact, he doesn’t even know anyone in Sweden.

It is at this point that the system enters the picture. Bob is a regular participant in a Babble run by the “Netweavers” group, a cross-organizational special interest group whose aim is to support community-based knowledge management throughout Global Corp. In response to Åke’s request, Bob turns to the Netweavers’ Babble, and engages Carlos in the dialog shown in Figure 4.

Bob and Carlos do not know one another well. They have never met face to face, nor do they expect to: Bob is based in the United States, whereas Carlos is in Spain. They only know one another because both are frequent participants in this Babble environment, and have primarily interacted via casual opportunistic chat in the Commons room of Babble. It is also possible that they know a bit more about one another, because both Bob and Carlos have “personal places” in the Netweavers’ Babble in which they’ve published personal information about themselves.

Bob decides to talk with Carlos, because he (erroneously) believes that Carlos used to work in Sweden. Bob begins by enquiring whether Carlos is ‘still around,’ by which he means ‘available for a chat.’ Bob can see from the social proxy that Carlos is logged on, but can also see that Carlos has been idle for a while, and Bob also knows that it is getting quite late in Spain. After waiting a minute for a response, Bob decides to post his query anyway; even if Carlos is not around, it will persist and Carlos will see it later. So, Bob proceeds, explaining the context and asking Carlos if he has contacts in Sweden.

Bob [1:47:11]: Hey Carlos, are you still around?
I know a guy who is a very talented visual/systems designer who has done a lot of online education type work. He is moving to Sweden with his wife and looking for work. I somehow think that you used to work in Sweden (or was it Denmark?)... do you have any leads for him?

Carlos [1:48:54]: Yeah, still here ...
I never lived in Sweden I am afraid nor worked there but I have got a couple of colleagues in Denmark that have got some good connections in Sweden. You may need to contact them: Dana and Dagmar. Feel free to mention my name to them. They may be able to help out.

Bob [1:52:12]: Thanks Carlos... don't know why I had the Sweden bit beside your name turned on! I guess it's because you seem such a man of the world! ;-)

Carlos [1:58:08]: hehehe ... They call that nowadays virtual networking ;-P If those two leads do not work out let me know as I have got another one in Norway who may be willing to help out as well.

Carlos [2:00:42]: Hummm... what the heck ... the Norway contact is Davin. I used to work with her and every now and then we still get to talk to each other. She is really good fun!

Carlos [2:38:32]: Time to go... Last conf. call and then I am out of here... Take care, folks, and will see ya tomorrow! TTFN

Bob [3:15:45]: Thanks Carlos, I've contacted all three, and have dropped your name shamelessly (although I have had a little contact with Dagmar and Dana).

Fig. 4. Segment of Babble chat illustrating how people tap one another for assistance. (Reformatted for brevity and readability)

A minute later Carlos acknowledges that he is present, and a couple of minutes after that replies with the names of two colleagues in Denmark, Dana and Dagmar, and says to mention his name. After a bit of banter (it turns out that Carlos used to work in Denmark, and not Sweden), Carlos volunteers the name of Davin, another colleague located in Norway. Bob sends email to Carlos’ contacts, and posts a comment in Babble thanking Carlos, and indicating that he made use of Carlos’s information and name. Bob is also aware, though nothing comes of it in this case, that other participants in this Babble may read the exchange, and may also volunteer any contacts that they have.

It is important to note that this is not an unusual example—or rather, the example is unusual only in that it is relatively coherent and compact, occurring at a time when it was not interleaved with other talk. It is more common that such interactions occur amidst a rich *mélange* of social and work talk. For example, one instance, in which one participant helped another find a contact from which to obtain project funding, occurred over the course of a multi-threaded, 30 utterance, 17-minute conversation. The conversation consisted of two primary participants, Richard and Sean, and was composed of four distinct threads. Two threads were related to work topics (Richard explaining that he had referred some colleagues to Sean, and a discussion of the use of patterns in knowledge management), and two were more social threads (one an attempt to identify an earlier participant’s real name, another a request by Richard for assistance in developing an Irish accent for an upcoming storytelling performance). The two work related tasks were treated relatively seriously, even as the two interleaved non-work threads were used as an excuse for banter. Yet both the social and work threads developed and played off one other throughout the conversation, which concluded with Richard revealing the names of the colleagues whom he has referred to Sean, and Sean indicating that he would be happy to talk with them.

3.3 Discussion

These examples illustrate how systems like Babble and Loops support people in drawing upon one another for expertise. In the first example, we saw one way in which the social proxy—by serving as a focus of discussion for

strangers—supports the formation of relationships among those who don't know one another. Indeed, a number of the general usage patterns—casual opportunistic chat, focused discussion, undirected queries, and publishing personal information—all serve to promote the formation and strengthening of relationships among those who don't know one another well. Once relationships exist, they provide a means for tapping others for help. Thus, in the second example, Bob uses his relationship with Carlos to enlist him in his goal of finding others who know of job possibilities at Global Corp Swedish locations.

Let's look more closely at this second example. It raises three interesting points. First, it is significant, in these days of interest in knowledge and expertise management, to note that neither Bob nor Carlos nor any of the contacts that Carlos passes on are “experts” on finding jobs in Global Corp's Sweden locations. It is only in the next round of activity, in which Carlos' two Denmark contacts provide the names of Swedish colleagues, Erik and Eva, involved in hiring, that people with expertise on ‘jobs in Sweden’ may be said to have been reached.

The second point of interest is the nature of Carlos's response to Bob: he gives Bob three names, and permission to mention his name. Bob's ability to invoke Carlos's name when emailing Dana, Dagmar and Davin gives his request more power. That is, by mentioning Carlos's name he draws on Carlos's relationships with his contacts, and his request becomes something different from what it would have been if he'd randomly selected the three names from Global's phone directory. Similarly, one of Carlos's contacts, Dagmar, forwards Bob's email to his Swedish colleague Erik, with a personal note asking Erik if he can help, thus acting as an intermediary and presumably strengthening the force of the query. The point is that more than information exchange is happening here: the participants in this chain of activity are using subtle means to draw upon their social networks to support Bob's attempt to find a contact in Sweden. They are not just passing on names, but by virtue of allowing their names to be mentioned—or even serving as intermediaries to pass a message along—they are, in a small way, allowing their social reputations to be enlisted in support of the task.

The third point of interest has to do with the relationships amongst Bob and the others involved in this interaction. As a consequence of the support that Carlos, Dana and Dagmar have provided to Bob's endeavor (Davin was unhelpful), Bob has now accrued small but real obligations to all three. Should any of the three have a problem that Bob might assist them with, they will be more inclined to turn to Bob. There are three reasons for this: first, the interaction has reminded or informed them of Bob's existence; second, they recognize that Bob is now obliged to them in a small way; and third, the interaction has left a residue—Bob's email address and other contact information in their ‘sent email’ queues—which lowers the cost of contacting Bob (and, in fact, should they do so, they can simply reply to his email leaving the original message intact, thus subtly reminding Bob of his obligation).*** So, in summary, the interaction has simultaneously reminded them of the existence of Bob as a possible contact, and lowered their ‘cost’ of exploiting him.

The same is true for the other participants in this interaction: Åke, who made the original request, has now accrued an obligation to Bob (but not to anyone farther down the chain, as Åke knows nothing of them nor they of him); and Carlos, by virtue of allowing his name to be invoked, has accrued small obligations to Dana, Dagmar and Davin. While the original request and the responses to it are neither complicated nor onerous, their importance should not be underestimated. It is the very fact that the request is an easy one to fulfill that makes it an acceptable way of tapping the very weak ties that connect the members of this chain of activity. And, more importantly, it is this non-onerous exercise of responding to the request and allowing one's social reputation to be enlisted in the service of Bob's (and Carlos' (and Dana and Dagmar's)) request that strengthens the connections between people.

It is interesting to reflect on how tapping people for their expertise differs from getting said expertise from a document or other artifact. Setting aside issues of whether the artifact is available, understandable, or complete, there are two important differences. First, tapping a person in this way is a potentially reciprocal process: a person who serves as a source of expertise for another can in turn draw upon the first person, and, as noted, can do so with lower cost. Second, the process is potentially transitive: people are embedded in social networks, and thus tapping one person as a source of expertise can result in a cascade as the tapped person in turn taps into his or her social network, and so on. Finally, the combination of the reciprocal and transitive nature of this process of expertise sharing means as expertise is shared, the social networks that enable it become stronger and more extended.

4 Closing Remarks

At the beginning of this chapter we observed that a fact of life in a large organization—due to its scale and the practice of moving people among organizational units to broaden their experience—is that people frequently find themselves doing things they’ve never done before. A consequence of this is people need to “figure out how to figure out” how to proceed. In such cases, artifacts such as checklists, instruction manuals and written procedures are often (and necessarily) not adequately specified to be helpful. Indeed, artifacts often require the application of considerable expertise before the information they contain can be made useful. Instead, a common strategy is to turn to other people, who, even if they can’t help directly, may be able to provide indirect assistance by helping finding someone else who can help. However, while this is a sensible strategy, it is not so easy to implement in a large, distributed organization.

In this chapter we described Babble and Loops, and the ways in which distributed groups have used them to seek assistance and share expertise. Thus, in response to Åke’s request for help, Bob used Babble to tap Carlos (and indirectly, the whole Netweavers’ community) for assistance. In so doing, Bob took advantage of several things. First, Babble’s social proxy allowed Bob to see that Carlos was present, and thus able to provide an immediate answer to his query. Second, as a consequence of previous casual chats, Bob had a friendly if not terribly deep relationship with Carlos, and thus felt it would be appropriate to solicit his help. Third, Bob thought he knew, either because of the casual chats or because of personal information that Carlos had shared in his online profile, that Carlos had worked in Sweden, and was thus a relevant person from whom to seek help.

This illustrates a sort of paradox for those involved in designing systems that support people helping one another. On the one hand, in our example it was the colleagues who were only weakly connected who were most likely to be able to provide assistance. This is an instance of the general point, made by social network theorists (Granovetter 1983), that those who are most likely to help us are those we know the least. That is, new information about jobs—as well as all sorts of other information—is more likely to come from weakly connected acquaintances who move in different circles than the information seeker and his or her close friends and co-workers. On the other hand, it is precisely those we know the least (or not at all) who have the least incentive to help us, sharing little in the way of common context, experience or acquaintances. Thus, in situations like this we have a paradox: those who are most able to help us are least obligated to do so.

One of the strengths of Babble and Loops is that they support interaction among people who do not know one another well. The strategy of providing persistent chat augmented with the social proxy’s depiction of the presence and activities of participants, adds richness to the system. Many of the practices that have arisen in Babble and Loops—casual opportunistic chat, leaving trails, publishing personal information—make it easier for people to form the relationships and find the information that enables them to draw upon others’ expertise. While Babble and Loops are not a full substitute for face to face interaction, as our examples show they do enable geographically dispersed, lightly connected people to develop weak ties, and to tap one another for help. While Babble and Loops represent first steps, we believe the direction is an important one, at least in the domain of supporting distributed organizations. As Granovetter notes:

... social systems lacking in weak ties will be fragmented and incoherent. New ideas will spread slowly, scientific endeavors will be handicapped, and subgroups separated by race, ethnicity, geography, or other characteristics will have difficulty reaching a modus vivendi.

— A. Granovetter, page 202, 1983.

Acknowledgements

Thanks to David N. Smith and Mark Laff for their work on the design and implementation of Babble, to Mark Laff, Jeremy Sussman and Tracee Wolf for their work on the design and implementation of Loops, and to Cal Swart, who keeps the infrastructure for everything up and running. Thanks is due, as well, to our management, especially John

Richards and Brent Hailpern, for supporting the work in its early days, and to our many users for their patience, enthusiasm and very real contributions.

References

Ackerman MS, Pipek V, Wulf V (2003) *Sharing expertise: beyond knowledge management*. The MIT Press, Cambridge MA

Bradner E, Kellogg WA, Erickson T (1999) The adoption and use of Babble: a field study of chat in the workplace. In: Bødker S, Kyng M, Schmidt K, (eds) *Proceedings of the sixth European conference on computer supported cooperative work*, Kluwer Academic Publishers, Dordrecht/Boston/London, pp 139–158

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

Erickson T, Kellogg WA (2000) Social translucence: an approach to designing systems that mesh with social processes. *Transactions on Computer-Human Interaction*. 7:1:59–83

Erickson T, Kellogg WA (2003) Social translucence: using minimalist visualizations of social activity to support collective interaction. In: Höök K., Benyon D, Munro A (eds) *Designing information spaces: the social navigation approach*. Springer, London, pp 17-42

Erickson T, Smith DN, Kellogg WA, Laff, MR, Richards JT, Bradner E. (1999) Socially translucent systems: social proxies, persistent conversation, and the design of “Babble”. In: Williams MG, Altom MW, Ehrlich K, Newman W (eds) *Human factors in computing systems: the proceedings of CHI 99*. ACM Press, New York, pp 72–79

Goffman E (1963) *Behavior in public places: notes on the social organization of gatherings*. Macmillan Publishing Co., New York

Granovetter MS (1983) The strength of weak ties: a network theory revisited. *Sociological Theory* 1:201–233

Hall ET (1983) *The dance of life: the other dimension of time*. Doubleday, New York

Halverson CA, Erickson T, Sussman J (2003) What counts as success? punctuated patterns of use in a persistent chat environment. In: Schmidt K, Pendergast M, Tremaine M, Simone C (eds) *Proceedings of the 2003 international ACM SIGGROUP conference on supporting group work*. ACM Press, New York, pp 180–189

Jacobs J (1961) *The death and life of great American cities*. Random House, New York

Kendon A (1990) *Conducting interaction: patterns of behavior in focused encounters*. Cambridge University Press, Cambridge MA

Whyte WH (1988) *City: return to the center*. Anchor Books, New York

* To appear in *Resources, Co-Evolution and Artifacts: Theory in CSCW* (eds. MS Ackerman, C Halverson, T Erickson and WA Kellogg). Springer, London. In preparation, 2005.

** Here, and in the rest of the examples, names have been altered to preserve anonymity.

*** This cluster of properties of email—the (usually) automatic archiving of sent messages which preserves both content and email address, enabling both to be reused in subsequent messages—makes email a particularly powerful communication medium for large organizational environments.