

Knowledge Communities: Online Environments for Supporting Knowledge Management and its Social Context*

Thomas Erickson and Wendy A. Kellogg
IBM. T.J. Watson Research Center
snowfall@acm.org; wkellogg@us.ibm.com

1. Introduction

The issue of how to support the re-use of knowledge — under rubrics such as organizational memory, knowledge management and expertise management — has received increasing attention over the last decade. In this chapter we take a strongly social approach to the issue, arguing that knowledge (and expertise) is created, used, and disseminated in ways that are inextricably entwined with the social milieu, and therefore that systems which attempt to support these processes must take social factors into account.

Our approach to managing knowledge or expertise is to do it on-line, via multi-user networked environments that support group communication and collaboration. That is, we are interested in designing on-line environments within which users can engage socially with one another, and, in the process, discover, develop, evolve,

* To appear in *To Appear in Beyond Knowledge Management: Sharing Expertise*. (eds. Ackerman, Mark, Volkmar Pipek, and Volker Wulf). Cambridge, MA, MIT Press, in press 2001.

and explicate knowledge relevant to shared projects and goals. We refer to online multi-user environments used in these ways as “knowledge communities.”

This chapter consists of two sections. In the first section we make the case for a deeply social approach to knowledge management. We begin with an example that depicts a number of ways in which the production and use of knowledge is fundamentally entwined with social phenomena. We note that this socially situated view of knowledge is supported by research in a number of disciplines, and also has made its way into the business discourse that surrounds knowledge management. This view raises a challenge for those designing technology: knowledge management systems must take into account, either explicitly or implicitly, the social context within which knowledge is produced and consumed.

In the second part of the chapter we argue that one way of addressing this challenge is via the sorts of online multi-user systems that we call knowledge communities. We describe some examples of systems that currently function as knowledge communities and then turn to our own work on designing infrastructures for knowledge communities. Our general approach is to design online environments that, by making users and their activities visible to one another, can enable a variety of social phenomena that support social and work-oriented interaction. We describe a system called “Babble,” which we have designed, implemented, and deployed to about twenty workgroups over the last four years. We report on our experience with Babble, and conclude by discussing some of the general issues we see for designing online environments that support a socially-oriented approach to the management of knowledge and expertise.

2. Knowledge Work as Social Work

Knowledge management is often seen as an information problem: how to capture, organize, and retrieve information. Given this perspective, it isn’t surprising that knowledge management evokes notions of data mining and text clustering and databases and documents. This is not wrong, but it is only part of the picture. We suggest that knowledge management is not just an information problem, but that it is, as well, a social problem.

2.1 An Example

One of us once interviewed accountants at a large accounting and consulting firm about their information usage practices. The goal was to find out how they thought they would use a proposed database of their company’s internal documents. In the course of the investigation, an unexpected theme emerged: the accountants said that one of the ways in which they wanted to use the documents was as a means of locating people. The accountants’ claim — that they wanted to use a *document* retrieval system to find *people* — was, at the time, quite surprising. However, in

the course of further interviews, it came to make sense: It was only through the people that the accountants could get some of the knowledge they needed. As one accountant explained, ‘Well, if I’m putting together a proposal for Exxon, I really want to talk to people who’ve already worked with them: they’ll know the politics and the history, and they can introduce me to their contacts. None of that gets into reports!’

For our purposes, there are five important points here. First, as the accountants observed, some types of knowledge tend not to get written down. Sometimes it may be that the knowledge is too politically sensitive: people shy away from recording gossip and innuendo, even though knowledge of it may be very helpful to someone about to do business in that environment. Sometimes knowledge — in the form of comments, opinions, or conjectures — may not be written down because the resulting records can be potentially be subpoenaed. And sometimes knowledge that may seem too trivial to be recorded when first encountered — that the CEO is a teetotaler or a Scotch fancier — can prove quite valuable in the process of establishing a relationship. Because this knowledge is often quite useful for social purposes, we will refer to it by the rubric “social knowledge.”

The second point is that the accountants were not just tapping into social knowledge; they were also getting access to *social resources* such as contacts and referrals. One accountant explained that the worst way to approach a company with a proposal was by making a “cold call”. It is much better if the accountant, let us call him Charles, can begin a call to a new contact by saying ‘My colleague, Jil Smith, suggested I chat with you.’ Being able to say that one has been referred by a mutual acquaintance is a frequent and powerful facilitator for interpersonal interaction — and this is true even if the relationship is only a few hours old. Charles, by virtue of having permission to assert a relationship with Jil, can draw on — to some extent — Jil’s reputation and standing with the person with whom he is trying to open negotiations. Notice, by the way, that social resources can’t be extracted from a person and embedded in a database: opening the conversation by saying ‘I found your name in the corporate knowledge base’ isn’t the same as saying ‘Jil Smith said I should call.’

The third point we take from this example is that people don’t necessarily need access to an “expert.” It may be that Jil Smith has had only one previous engagement with Exxon, and that, in terms of facts, she may have far less expertise than an outside consultant. Nevertheless, Jil’s experience may be sufficient to provide Charles with the social knowledge and social resources necessary to gain entry into the Exxon environment. In fact, it may be preferable for Charles to talk with Jil, because, as a colleague who shares the same work context, she will understand more about what he needs to know, the situations in which he will use the knowledge, and how he is likely to go about using it, than someone traditionally construed as an expert. That is, Jil has social and contextual expertise, in contrast to an outside consultant’s factual expertise.

The fourth point we take from this example is implicit in the previous ones: networks of personal relationships, which are created and reinforced through interpersonal conversation, are critical in supporting knowledge sharing. Let us return to the example of Jil and Charles. Assuming that Jil's assistance was helpful, Charles has now accrued a small debt or obligation to Jil. When Jil needs assistance, she is likely, in turn, to come to Charles with questions or requests for social knowledge that falls within his domain. Even if the required information is outside of his domain, she may seek to obtain access to his social resources — a referral to one of his contacts, for example. Thus are professional relationships established, and thus do social networks grow. In the long run, if not the short, it may be more valuable for an enterprise if its members seek knowledge and social resources from one another — thus building a web of mutual knowledge and trusted relationships — than if, for instance, employees are given instant access to a top-notch external domain expert.

This brings us to our final point, which has to do with the centrality and importance of conversation in knowledge sharing (see Fitzpatrick, this volume). It is no coincidence that both social knowledge and social resources are best shared through talk. It is the time spent discussing apparently trivial social knowledge that suggests that a relationship goes beyond the purely professional — that there is more in play than just a purely instrumental professional exchange. It is the disclosure of politically sensitive information that indicates a degree of trust between two people. It is the ability of one person to take generic information and apply it — on the fly — to the other's problem that increases the reputation of the giver and creates an obligation for the receiver. This sort of talk — and the exchange of knowledge and social resources it involves — both requires and strengthens networks of personal relationships in workplace.

2.2 The Social Construction of Knowledge

This sort of situation is not the exception, it is the rule. A wide variety of research programs — for instance, ethnographies of workplaces, social studies of science, critical theory, organizational memory, the sociology of knowledge — point to the deep connections between knowledge management and social context.

For example, ethnographic studies of workplaces reveal a wide array of social practices implicated in the production and dissemination of knowledge. Lave and Wenger have developed the notion of a community of practice. They note that one way in which people come to master a body of knowledge is through a sort of apprenticeship or "legitimate peripheral participation" in the activities of a group of practitioners (Lave and Wenger, 1991). Wenger (1998) describes the daily work in an insurance claims processing office, and shows how it is entwined with social relationships and processes. Similarly, in an ethnography of copier service technicians, Orr (1996) reveals that technical knowledge is socially distributed across a network of technicians, and that it is tapped into and disseminated through

oral processes such as storytelling.

A similar sense of the social nature of the production and dissemination of knowledge comes from the field of social studies of science — see Latour and Woolgar (1979) and Latour (1987). For example, Traweek's ethnography of particle physicists (1988) examines some of the social phenomena that structure the practice of high energy physics. She notes the impact of social relationships on the placement of graduate students, the evaluation of experiments, and access to equipment and facilities. Her comments on the role of conversation are particularly interesting:

"...talk accomplishes diverse tasks for physicists: it creates, defines, and maintains the boundaries of this dispersed but close-knit community; it is a device for establishing, expressing, and manipulating relationships in networks; it determines the fluctuating reputations of physicists, data, detectors, and ideas; it articulates and affirms the shared moral code about the proper way to conduct scientific inquiry. Acquiring the capacity to gossip, and to gain access to gossip about physicists, data, detectors, and ideas is the final and necessary stage in the training of a high energy physicist." (Traweek, page 122)

At a more general level, Brown and Duguid (1995) note that even documents, which appear to be fixed, immutable public entities whose very purpose is to transcend social contexts, "play an important role, bringing people from different groups together to negotiate and coordinate common practices." Documents, they suggest, in their production, use, and distribution, have their own social life, and function as mediators of and catalysts for social activity.

2.3 Social Capitalism

An awareness of ways in which work is bound up with social factors has assumed a prominent place in business discourse regarding knowledge management. Often referred to as organizational learning in these contexts, knowledge management in the organization is seen as a collective process in which teams create and share knowledge (e.g. Senge, 1990; Nonacka and Takeuchi, 1995; Cohen and Prusak, 2001; Boone, 2001). While proponents typically invoke a systems perspective in thinking about organizational processes, they also emphasize social factors — such as relationships, trust, reputation, and commitment — in their descriptions of how such processes play out. As a Vice President of Strategy puts it:

Expertise location is a big issue in companies today. The goal is not only to provide access to information, but to provide access to people who have the information. ... I don't want raw data, I don't even want information, I want the judgments of people I trust. (Boone, page 22)

Recently the concept of social capital — the "features of social organizations such as networks, norms, and social trust that facilitate coordination and cooperation for mutual benefit" (Putnam, 2000) — and the possible role it may play in the networked organization, has come to the fore. Cohen and Prusak (2001)

explain the connection:

Social capital makes an organization, or any cooperative group, more than a collection of individuals intent on achieving their own private purposes. Social capital bridges the space between people. Its characteristic elements and indicators include high levels of trust, robust personal networks and vibrant communities, shared understandings, and a sense of equitable participation in a joint enterprise—all things that draw individuals together into a group. This kind of connection supports collaboration, commitment, ready access to knowledge and talent, and coherent organizational behavior. (Cohen and Prusak, page 4)

Elaborating on the connection between social capital and knowledge sharing, Cohen and Prusak point out that exchanging knowledge depends on a social connection — "without some degree of mutuality and trust, the knowledge conversations will not get started; without some degree of shared understanding, they will not go very far" (Cohen & Prusak, 2001, p. 86). They also note that the knowledge exchanged in spontaneous conversations "is often social knowledge — shared aims and interests discovered, signals and stories shared that build confidence, trust, and connection—rather than technical or business knowledge that can be directly applied to a product or problem" (Cohen and Prusak, pp. 86-87).

2.4 The Challenge

Thus far we have argued that knowledge management is not just an information problem, but that it is a social problem. That is, we've suggested that effective knowledge management involves networks of people, relationships, and social factors like trust, obligation, and commitment. One can't isolate knowledge from its social context without denaturing it, without stripping it of the social resources and social knowledge that contribute to its utility.

Taking the social nature of knowledge seriously raises a considerable challenge for those interested in designing knowledge management systems. We suggest that the place to start is to stop thinking in terms of knowledge management, and start thinking in terms of supporting the larger social context in which knowledge management is embedded. Our response to this challenge is to explore the role of online multiuser environments. In particular, we are interested in environments within which users can engage socially with one another, and, in the process, discover, develop, evolve, and explicate knowledge. We refer to online multi-user environments used in these ways as "knowledge communities." In what follows we discuss current environments that function as knowledge communities, and then turn to our own work on the topic.

3. Knowledge Communities

Knowledge communities have a long history, albeit not by that name. The idea that

networks of computers might provide a medium within which individuals might come together to share knowledge and expertise dates back to at least 1960. Perhaps the first vision of this nature was offered by Simon Ramo (Ramo, 1961), who wrote of “many millions of human minds ... connected together.” Ramo offered a number of scenarios, including one of an attorney consulting an on-line database that contained more than data:

“Even on the nonroutine legal processes, the attorney, in the coming intellectronic age, will be able to consult with the equivalent of a host of informed fellow attorneys. His request to the system for similar cases will yield an immediate response from the central store, together with questions and advice filed by other attorneys on those similar cases -- even as he will add his facts and guidance into the system for future use by all.” page 10.

— Simon Ramo, The Scientific Extension of the Human Intellect. *Computers and Automation*, Vol 10., No. 2, pp 9--12. February 1961. (Based on a talk given in 1960).

Over the ensuing decades the idea spread and evolved. From its beginning as a vague if exciting vision, it took concrete form in the special purpose DELPHI and EMISARI systems pioneered by Murray Turoff in the early ‘70s (Turoff, 1972; Hiltz and Turoff, 1993) and in the PLATO Notes system in the mid ‘70s (Wooley, 1993). In the late ‘70s and early ‘80s the idea took off, spreading and evolving, under pressures from application domains such as education and gaming, into a variety of genres of software ranging from bulletin board systems to MOOs.

3.1 Some Examples of Knowledge Communities

A complete account of the systems which are used to enable online groups to share knowledge among themselves is well beyond the scope of this chapter. Instead, we will take the tack of looking at some representative examples to give an idea of both the types of systems and the forms of use that are used in managing online knowledge. It is important to note that we are not just interested in the software; we are interested in the combination of the software and the way in which it is put to use by its users — we refer to this combination of technology and usage as a knowledge community.

One genre of software that supports knowledge communities is the MOO. MOOs, originally developed as multi-user text-based gaming environments, have been applied to a number of pedagogical and business ends. Examples include MOOSE Crossing, an educationally-oriented environment for children from eight to thirteen (Bruckman, 1997); Pueblo, a school-centered MOO in Phoenix, Arizona (O'Day, et al, 1996); Tapped In, a distributed community of teachers (Schlager, et al., 1998; Schlager et al., in press); and a MUD used by employees at Argonne National Labs for work-related talk (Churchill and Bly, 1999).

Another genre of system that can support knowledge communities is the electronic mailing list, or listserv. While mailing lists are used for a variety of purposes, the existence of mailing lists used to share knowledge among cohesive,

long-lasting communities is well documented. In one case, a community of about a thousand professional journalists have used a mailing list to help one another with technical problems and to find story-specific information sources for over six years (Millen and Dray, 1999; Millen 2000). Another example, the use of a mailing list to support discourse amongst a scholarly community, is described by Ekeblad (1999). And a third example, the use of a listserv by a community of soap opera fans, to share knowledge ranging from plot summaries to character background information, is described by Baym (Baym 1995; Baym 1997).

In addition to the genres of mailing lists and MOOs, which can be turned to a variety of ends other than knowledge sharing, quite a few systems have been designed with their principal aim being the support of a knowledge community. One example is Answer Garden (Ackerman, 1998), a blending of bulletin board and email systems that makes a network of questions and answers available to its users, and uses email to automatically route new questions to appropriate experts whose answers are then incorporated into the network. The Zephyr Help Instance (Ackerman and Palen, 1996) has a similar purpose — providing online help information — but uses a synchronous chat-like mechanism to broadcast questions and answers to the user community. Another genre of knowledge community system is the collaboratory. Collaboratories are aimed at the needs of the scientific community, and provide real time access to scientific instruments along with synchronous communication channels ranging from textual chat to real time audio and video (Olson and Olson, 2000). Collaboratories are a highly successful class of applications, with many in existence that have supported dozens to hundreds of users for periods of years.

If one examines these systems and the ways in which they're used to share knowledge, an interesting commonality emerges: Virtually all of these systems exhibit a rich array of social phenomena, in spite of the fact that most provide only textual communication mechanisms, typically synchronous chat, asynchronous email, or both (as in MOOs). (Even collaboratories, which are increasingly supporting various forms of high bandwidth synchronous interaction, functioned well when chat was their dominant communication channel.) Examples of the social phenomena found in most knowledge communities range from interpersonal phenomena such as the negotiation of status and reputation or the development of trust, to the emergence of group norms and conventions. While these systems bear eloquent testimony to the ingenuity of their users in using textual representations to support a rich array of social phenomena, we suspect that we can do better.

This brings us to the question which informs our own work. What would it mean to design an infrastructure for a knowledge community from the ground up? That is, if we take seriously the charge that knowledge management is a social problem as well as an information problem, one response is to ask how we can better support social interaction. How do we go about designing a system which supports not just information sharing, but that supports the exchange of social

knowledge and resources, the creation and growth of interpersonal networks and accompanying social phenomena such as trust, obligation, commitment and accountability?

To address this question, we've developed a system called "Babble" which we've used as a testbed for exploring these issues over the last four years. We begin by discussing the rationale that underlies Babble's design: the notion that increasing the visibility of the presence and activity of participants in an online environment can provide a foundation for a variety of social processes and activity. Next we describe the system that we've implemented, and discuss the ways in which we've come to use it as part of our daily work practice. Finally, we discuss our experiences in deploying Babble to other work groups.

3.2 Supporting Online Social Interaction

In the building where our group works there is a door that opens from the stairwell into the hallway. This door has a design problem: opened quickly, it is likely to slam into anyone who is about to enter from the other direction. In an attempt to fix this problem, a small sign was placed on the door: it reads, "Please Open Slowly." As you might guess, the sign is not a particularly effective solution.

Let's contrast this solution with one of a different sort: putting a glass window in the door. The glass window approach means that the sign is no longer required. As people approach the door they see whether anyone is on the other side and, if so, they modulate their actions appropriately. This is a very simple example of what we call a socially translucent system.

While it is obvious why this solution works, it is useful to examine the reasons behind it carefully. We see three reasons for the effectiveness of the glass window: First, the glass window makes socially significant information *visible*. That is, as humans, we are *perceptually attuned* to movement and human faces and figures: we notice and react to them more readily than we notice and interpret a printed sign. Second, the glass window supports *awareness*: I don't open the door quickly because *I know* that you're on the other side. This awareness brings our social rules into play to govern our actions: we have been raised in a culture in which slamming doors into other people is not sanctioned. There is a third, somewhat subtler reason for the efficacy of the glass window. Suppose that I don't care whether I hurt others: nevertheless, I'll open the door slowly because *I believe that you know that I know* you're there, and therefore I will be held *accountable* for my actions. (This distinction is useful because, while accountability and awareness usually co-occur in the physical world, they are not necessarily coupled in the digital realm.) It is through such individual feelings of accountability that norms, rules, and customs become effective mechanisms for social control.

We call systems that exhibit these properties – of perceptual salience, awareness, and accountability – socially translucent systems. But there is one other aspect of social translucence that deserves mention. Why is it that we speak of socially

translucent systems rather than socially *transparent* systems? Because there is a vital tension between privacy and visibility. What we say and do with another person depends on who, and how many, are watching. Note that privacy is neither good nor bad on its own—it simply supports certain types of behavior and inhibits others. For example, the perceived validity of an election depends crucially on keeping certain of its aspects very private, and other aspects very public. As before, what we are seeing is the impact of awareness and accountability: in the election, it is desirable that the voters *not* be accountable to others for their votes, but that those who count the votes be accountable to all.

We see these three properties of socially translucent systems — visibility, awareness, and accountability — as critical building blocks of social interaction. Notice that social translucence is not just about people acting in accordance with social rules (see Erickson & Kellogg, 2000). In socially translucent systems we believe it will be easier for users to carry on coherent discussions; to observe and imitate others' actions; to engage in peer pressure; and to create, notice, and conform to social conventions. We see social translucence as a requirement for supporting online communication and collaboration in general, and knowledge communities in particular.

This brings us to the question of how to support social translucence in online environments. How can we provide the cues that allow our socially based processes to operate — and which are so ubiquitous and lightweight in the physical world — in online systems? Two obvious approaches are to use video or 3D virtual environments. However, these have several drawbacks for our purposes. First, they don't scale well: we would like to support conversations among fairly large numbers of people. Second, both approaches are best suited for supporting synchronous interactions, whereas we would like to support both synchronous and asynchronous interaction. Third, they are both relatively demanding in terms of processing power, bandwidth, and display space and characteristics: we would like to be able to support mobile employees working over sub-56K connections and using devices with smaller displays.

As a consequence, we have taken a more abstract approach to supporting social translucence. The abstract approach involves portraying social information in ways that are not closely tied to its physical analogs. Exemplars of the abstract approach include the Out to Lunch system (Cohen, 1994), which uses abstract sonic cues to indicate socially salient activity, and Chat Circles (Viegas, et al. 1999), which uses abstract visual representations. This approach also includes the use of text to portray social information; as we have already noted, text has proved surprisingly powerful as a means for conveying social information in knowledge communities.

3.3 The Babble System

Babble (Erickson, et al., 1999) is an online environment intended to support both synchronous and asynchronous text-based conversations within small to medium

sized groups. The principle goal of Babble has been to serve as a platform for exploring ideas about the social effects of supporting mutual awareness among online groups. However, to do this effectively, we needed to be able to observe 'real' workgroups using it as part of the daily work process. As a consequence, Babble needed to be sufficiently robust and lightweight to be usable by groups who don't care about the technology itself.

In terms of infrastructure, Babble is a client-server system with both components written in SmallTalk. Babble stores all data, except for user specific preferences and state (e.g., the user's last location, last items read, and so forth) on the server and broadcasts it as needed. Babble clients request the data they need from the server (e.g., when a user switches to a new conversation the client requests the content), and also notify the server of events that it will broadcast to other clients. As this architecture suggests, Babble only works when on a network; when disconnected it has no cache of conversation text. The Babble server runs on a variety of server-class machines; the principle client runs on PCs, though we have had, for varying durations, clients that ran on the Macintosh (in Java) and on the Palm Pilot. Here we discuss only the PC client, since that comprises the vast majority of our experience.

In terms of functionality, Babble resembles a multi-channel, text-based chat system in that many users can connect to it, and select one of a variety of conversations to participate in (or create their own). However, Babble differs from conventional chat in two ways, both of which stem from our interest in supporting knowledge communities. First, the textual conversation that occurs in Babble is persistent: that is, unlike conventional chat where newly arriving users only see what has transpired since they've joined a channel, Babble users can see everything ever typed in any existing conversation. These traces give the system the potential to function as a knowledge store, or what we prefer to call a "discourse base." Second, Babble makes the presence and activity of the participants visible by a variety of means, but principally through what we call the social proxy.

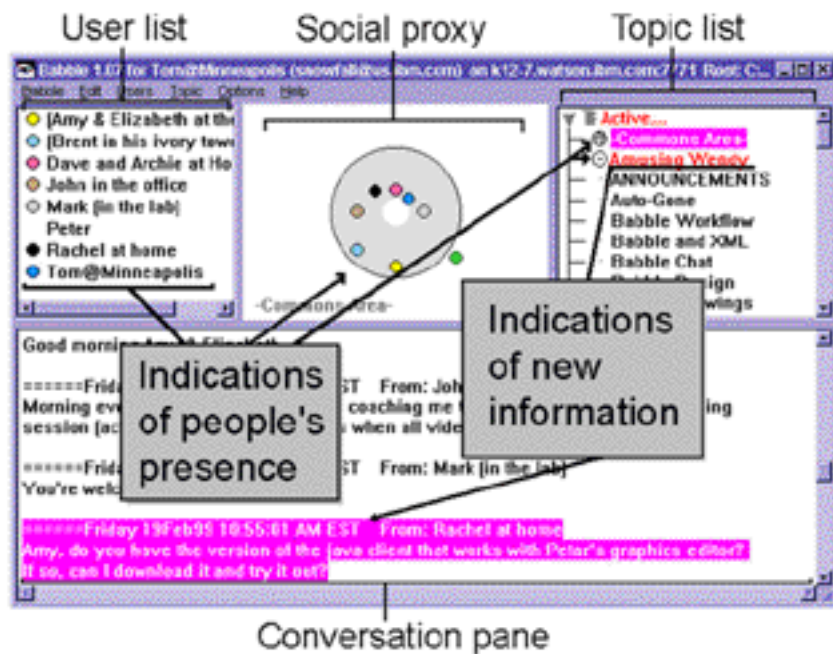


Figure 1. The Babble Interface. From the upper left, clockwise: the list of all users logged on; the social proxy; the topics list; the conversation pane.

Figure 1 shows the Babble user interface. In the upper left hand corner is a list of the names of users currently connected to Babble. In the middle upper pane is the social proxy, which we will describe shortly. In the upper right pane is a hierarchical list of the Babble conversation topics (grouped in categories and subcategories). And the pane that occupies the lower half of the window contains the text of the current conversation (whose topic name is highlighted in the topics list); within the pane, each 'comment' is prefaced with the name of the user, and date and time of its creation (recall that Babble conversations need not be synchronous; indeed, some are asynchronous, with hours, days or weeks separating comments). Babble provides a variety of other types of functionality via the menu bar, context-sensitive menus accessed via right clicks, and keyboard shortcuts. These include functions for creating messages, creating, changing, and deleting topics and categories, conducting private, ephemeral chats, and so forth.

The social proxy, in the upper middle part of the window, represents the current conversation as a large circle, and the participants as colored dots, referred to, hereafter, as marbles. Marbles within the circle are involved in the conversation being viewed; marbles outside the circle represent those who are logged on but are viewing other conversations. What makes the social proxy interesting has to do with the position of the marbles in the circle. When a user becomes active, either 'speaking' (i.e., typing) or 'listening' (i.e., interacting with the conversation window by clicking or scrolling), the user's marble moves rapidly to the center ring of the circle. If the user stops interacting, the marble gradually drifts out to the inner

periphery of the circle over the course of about twenty minutes. Thus, when there is a lot of activity in the conversation, there is a tight cluster of marbles around the center of the circle. The social proxy shown in Figure 1 depicts a situation in which five people have been recently active (i.e., speaking or listening) in the current conversation, and two others have been idle for a while (and an eighth person is off viewing another conversation).

When people leave the current conversation their marbles move outside the circle; when they enter the conversation, their marbles move into the circle. When a person logs onto the system, it creates a virtual wedge for their marble, adjusting the position of all the marbles in the social proxy; when they depart, the wedges are destroyed, and the remaining marbles adjust to uniformly occupy the space. All marble movements are shown with animation, thus making arrivals, movements, and departures visually salient. Although simple, this social proxy gives a sense of the size of the audience, the degree to which the audience is actively listening or contributing, as well as indicating whether people are gathering or dispersing, and who it is that is coming and going.

In addition to the social proxy (which we refer to as 'the cookie'), Babble uses additional mechanisms to reveal the presence and activity of users. In the topic list, to the left of the topic names, are 'mini-cookies', thumbnails of the social proxy for each topic with a participant in it. So, in Figure 1, we can see that there is a single person in the second topic, "Amusing Wendy." Babble also highlights information that the user hasn't yet seen: the names of topics with new material in them are shown in red (e.g., "Amusing Wendy" in Figure 1), and comments that have been added to the current conversation since the user last 'touched' Babble are shown in reverse highlighting.

One of the shortcomings of the cookie is that it only works for synchronous interactions — that is, it shows only the presence and activities of people who are currently logged on to Babble. This is a considerable drawback because the majority of the conversations carried on in Babble are asynchronous, with just a few comments per day (or per week, or per month). As a consequence, we designed a second, asynchronous social proxy for Babble: the Timeline (Figure 2).

The basic goal of the Timeline was to provide a way for a 'speaker' to see that people were 'listening' (or not), even when the listening was offset in time. The Timeline proxy works as follows: each user is represented by a row in the Timeline; when they are logged on to Babble, they leave a flat trace or line, and when they 'speak' they leave a vertical mark or blip on the line. If the line/blip is in color, it means that that user was present/speaking in the conversation currently being viewed by the user of the timeline; if they were in a different conversation, the line/blip is shown in gray (and the line becomes thinner). As the user mouses over the Timeline, the name of the topic, the user, and the time being examined is shown in the upper left corner; the user can scroll back through as much as one week of activity. The Timeline also provides access to other functionality via a menu



Figure 2. The Timeline (showing 3 hours of activity).

accessed via a right-click on another user's row (e.g., private chats).

For example, in Figure 2, we can see that nine people have logged onto Babble (shown by the presence of lines), and that all of them have spent some time in the current conversation (shown by the color/increased thickness of the lines), and that many but not all have 'spoken' (shown by the blips). The line being indicated by the cursor shows that the user 'Peter' logged on around 11am, made a couple of comments in the "Commons Area" conversation, switched to another topic, and then switched back to the Commons area about 1pm, and then logged off.

3.4 How a Babble is Used by a Group

While one must be wary about drawing conclusions concerning the usability of software when it is used by its developers, our aim here is to simply provide a sense for how Babble is actually used by a group. We'll begin by describing the group, and then move on to discuss how Babble is actually used. In the next section we'll discuss our deployments of Babble to other groups and some of the phenomenon that we've observed across different deployments.

Our group has used Babble for about four years. The group is centered around the software development group (AKA "the lab") that designed and implemented the system, and includes a mix of computer scientists and social scientists (including the authors). The size of the group has varied in number over the years from four to nineteen users. Part of the variance is due to the ebb and flow of people characteristic of groups in large organizations; and part is due to current members of the lab inviting "associates" — colleagues with whom they had strong social or professional ties — to join Babble.

Geographically, the group of Babble users is about half co-located in New York, and half distributed. Most of the lab members are located in the same building, although offices tend to be distributed around the building — so actual adjacency is rare. Three members of the lab are telecommuters, and spend the majority of their time tens to hundreds of miles away; other members of the lab frequently work at home. Four of the six associated colleagues (i.e. those not officially members of the lab, but users of Babble) are remotely located.

Socially, the lab is a cohesive group, with considerable camaraderie. In addition to work-based collaboration, the lab members occasionally socialize, although usually within business hours (e.g., going out to lunch) The associates vary in the

strength and number of their ties to the lab members, some known to almost all lab members, and others known only to one or two lab members with whom they have shared interests. Conversation in the Babble system moves fluidly between work and social talk; it is always civil, frequently informal, and joking, teasing, and other ludic behavior is not unusual.

Overall, the Babble system as used by the lab can be characterized as a core of relatively synchronous activity surrounded by a constellation of asynchronous conversations. At the center of activity is a topic called the "Commons Area," a place where collocated and remote members greet one another, share news, engage in banter, and ask general questions. Members of the lab tend 'hang out' in the Commons Area, often remaining logged on for most of the work day. Comments in the Commons Area tend to be short and informal, with relaxed syntax and punctuation, use of paralinguistic expressions ("ummm"), onomatopoeia, emoticons, and playful tropes (for example, the 'tossing of cookies' to 'a dog' who usually 'accompanies' one of the participants — all done via text, of course). The content of conversation in the Commons Area ranges from purely social talk (such as the custom of saying "good morning"), to the posing of general questions, to reminding people of an impending meeting of general interest, to more technical discussions about work projects. (In theory, more topic-oriented discussion is 'supposed' to take place in specific topics; in practice, work talk often grows out of social discussions, and the recognition that a substantive conversation that 'belongs somewhere else' is taking place is often not recognized until after the fact.) Because of the amount of talk that occurs in the Commons, the content of the Commons Area is automatically archived once a week.

In addition to the commons area, there are a variety of other topics or conversation areas in Babble. These have ranged in number from a dozen or so in the early days of Babble to several dozen, the growth being facilitated by the addition of an expandable hierarchical topic list. These topics tend to have asynchronous and mostly sporadic conversation, and they tend to be focused on particular purposes, typically either project-oriented or person oriented. Examples of topics include personal offices (e.g., "Tom's Office"), project-oriented topics (e.g., "Babble Ethnographies—CB Babble"), and occasional non-work topics (e.g., "Bad Jokes").

In general, uses of Babble can be grouped into three general categories: social/ludic; informative; and instrumental. Social/ludic activities are those engaged in for social and entertainment purposes such as the custom of exchanging morning greetings, and the topic devoted to jokes. Informative activities have to do with actions on the system that are addressed to the group as a whole, or to no one in particular, and generally are done without expectation of a reply or responsive action. These activities include posting announcements and other news believed to be of general interest, commenting on project activity, and keeping on-line notebooks or offices. The third type of activity is instrumental, that is, activities

engaged in with a particular end in mind. These include starting or participating in focused discussions, posting bug reports, holding on-line meetings, and asking questions. These activities are often, though not always, addressed to a particular participant or group of participants.

3.5 Adoption and Social Phenomena across Babble Deployments

Over the last four years we've deployed Babbles to about twenty groups. We've studied the deployments using techniques ranging from ethnographic studies — see Bradner, et al. (1999) for a study of six Babbles — to studies based on surveys and analyses of log data and conversation archives.

We have had mixed experiences with the adoption of Babble. Sometimes groups try Babble out, but fail to adopt it (typically about six weeks pass before it is evident whether or not the Babble is going to be adopted by the group). Other times groups use Babble for a period of months, and then cease (either because it was for a particular event or period that has ended, or because the composition or needs of the group change). It isn't clear how to operationally define a successful deployment of Babble: the group uses it for its entire existence? the group uses Babble actively for six months? the group uses Babble to carry out a particular activity? If we take, as a rule of thumb, that a Babble is successful when it is used on a more or less daily basis by several people for more than six weeks, we can say that about half of our Babble deployments have met with success. As of this writing, we have five Babbles running, all of which are well past the six week mark, and all exhibiting robust daily activity.

When a Babble is adopted by a group, it usually supports a variety of communicative purposes and practices (often similar to those described in the previous section). Here, we describe four social phenomena that we've observed in a number (though not all) of successfully adopted deployments that are most relevant to knowledge communities.

One phenomenon is waylay, in which a user watches for a particular person to become active on Babble (signaled by the movement of their marble into the center of the social proxy), and then initiates a conversation (either publicly within Babble, via Babble's private chat mechanism or by some external means such as the telephone). Because the movement of the marble occurs when the user has just begun an episode of typing or mousing, it indicates a opportune moment for contact (since the user's attention has just shifted to communication with the group). Waylay is used for purposes ranging from asking questions to initiating casual social chat. In general, forms of opportunistic interaction such as waylay permit the same sorts of requests for assistance and transfers of social resources that we've observed in face to face knowledge sharing situations, with the accompanying effect of strengthening of interpersonal ties.

Babble also supports the maintenance of group awareness through the exchange of social knowledge. For example, when members of a Babble travel, many report

reading through conversations that occurred in their absence to ‘find out what happened.’ For someone who is a member of the group and understands the context, seemingly trivial comments can convey considerable information about what’s going on at the individual, group, and organizational levels. Thus, a sign off — “I have to go to the [project] meeting now” — reveals that one participant is still involved in a particular project, and a question — “Does anyone know how to do a screen capture” — indicates that someone is beginning to write a paper. Babble also supports group awareness through the Timeline proxy. Babble participants have reported uses such as: looking to see who has visited a topic in which they had posted questions; looking to see whether a colleague who hadn’t posted recently had been online; and using the Timeline to get a sense for the activity of the community as a whole.

Another phenomenon that can be observed across Babbles is the development of social norms. That is, one participant may develop a particular way of doing something, and others will imitate it. Examples of this include what users include in their online nickname (e.g., in some Babbles users append “@mylocation” after their name), the types of online conversations created (e.g., some Babbles have categories for "personal places" or "offices"), and naming conventions (e.g., one Babble uses the term "chit-chat" to signal that a topic is intended for casual conversation. Babble groups also evolve various interactive customs, the most common being to say 'hello' upon logging in (even when no one else is present). Again, the existence of these norms supports social interaction by providing expectations about how to behave.

Finally, we’ve observed that Babbles are typically regarded as semi-private, “trusted” places. This became apparent when ‘strangers’ appeared in various Babble systems. Sometimes the strangers were unannounced new members, sometimes they were visitors provided access by an unreflective manager, and, in one case, the stranger was actually an unannounced conversational software agent. But in all cases, the arrival and presence of the stranger (reflected in the social proxies along with the presence of the regulars) was greeted with considerable consternation. In each case, the appearance of strangers provoked concern about how unguarded conversations might be interpreted by those from different contexts, and led to the creation of visitor and membership policies. We suggest that this concern reflects the success of Babble as an online space that is rich in social context.

One issue that is not clear, so far, is the degree to which Babble’s social proxies contribute to these phenomena. Analytically, it is difficult to isolate the effects of the social proxies, from the effects of purely textual cues. Certainly, there are a number of social practices (such as waylay) which require (or are at least greatly facilitated by) the proxies. It *is* clear that the participants, in general, like the proxies and want them retained as a feature of the system. One user, responding to a question in Babble, writes:

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

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. Another user, commenting on the Timeline proxy, remarks:

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

Nevertheless, although we have compelling anecdotes and a large fund of positive comments by Babble users, analytically separating social benefits conveyed by proxies from those produced by text remains as a challenge for the future.

3.6. Babble as an Infrastructure for Knowledge Communities

Babble clearly succeeds as a multi-user online environment where sustained social interaction takes place. But does it support knowledge communities? Is the social interaction combined with the sharing of information, social knowledge and social resources via personal social networks that, we suggest, is a crucial part of knowledge management? This is indeed what we have observed. In the following, we refer to examples¹ and survey results drawn from a Babble whose membership is composed of a world wide cross-section of people in IBM and Lotus interested in online communities.

Perhaps the first point to make is that participants do feel as though they are part of a community. This is particularly important to those who are remote teleworkers:

“I work remotely and can feel very isolated when I don’t travel regularly (as has been the case for the past six months because of travel restrictions). Babble has provided me with a way to feel connected with a group of people outside my basement walls. It is my portal (so to speak) into IBM.”

Another says:

“As a home office worker, this is perhaps one of the things I miss the most – the ongoing banter I can have with colleagues who are focused on a similar work topic as I am.”

This is not simply a feeling of a vague belonging to a group; participants report feeling as though they are hooked into social networks. One participant reports that participation in Babble strengthened an existing network:

¹ Identifiers have been changed to protect confidentiality, and comments edited for brevity

“Babble has helped me establish a tighter social and professional relationship with all of them – we have much more regular contact with each other, much as we would if we were collocated, via the Babble connection. This in turn has built social capital among us which may be of use in the future.”

And these social networks are not just about talk, they can also be tapped for assistance. The participant continues:

“I have also contacted Vera about getting her input and advice about setting up a knowledge network, which is part of my ‘real work.’ I felt much more comfortable about approaching her with this question as a result of our frequent contacts via Babble than I would have otherwise.”

Another Babble member notes:

“I like the back and forth. ...we have a lot of reflective talk about our own experiences... In at least one case, e.g., a half-joking comment of mine, “anybody want to fund this?” has led to e-mail, phone, and face-to-face meetings and now a serious proposal for funding. I don’t know the final outcome yet,² but we have found out something significant about another part of the business and have made a serious attempt to propose [a] solution to their problems.”

These comments are *prima facie* evidence that knowledge sharing and expertise management are deeply social processes — that people value informal exchanges with colleagues, and may only venture a non-trivial request for information or assistance after a social relationship has been established.

A danger in using the summary remarks of participants to what happens in Babble is that it makes it sound a bit more straightforward and calculated than it is. It is difficult to convey the way in which these effects emerge out of a rich melange of social and work talk. For example, one instance of the transfer of social resources occurred over the course of a multi-threaded, 30 utterance, 17-minute Babble conversation on March 7, 2001. The conversation consisted of two primary participants (‘scienceguy’ and ‘Patrick’), and was composed of four distinct threads. Two threads were related to work topics (Patrick explaining that he had referred some colleagues to scienceguy, 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 scienceguy 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 Patrick revealing the names of the colleagues whom he has referred to scienceguy, and scienceguy indicating that he would be happy to talk with them. (The situation grows more complex when one recognizes that Babble users are remote from one another, and may be simultaneously carrying on other work on their computers, via the telephone, or orally with co-located colleagues.)

² The project was funded.

5. Concluding Remarks

In this chapter we've argued that knowledge management is not just an information problem, but is, as well, a social problem that involves people, relationships, and social factors like trust, obligation, commitment, and accountability. This view raises a considerable challenge for those interested in designing systems to support knowledge management. Our approach has been to explore the creation of infrastructures for knowledge communities: on-line environments within which users can engage socially with one another, and, in the process, discover, develop, evolve, and explicate knowledge.

In our work on Babble, we've begun exploring ways of creating infrastructures that support rich forms of social interaction. We've found that social proxies are a promising development, and continue to be impressed with the power of plain text as a means of supporting interactions that are both complex and subtle. We believe that one of the most important aspects of a knowledge community is that it can be used as a place for unguarded discussion among people who know one another, who share professional interests, and who understand the contexts within which their remarks are being made.

The notion of a knowledge management environment as a 'trusted place' is an interesting and challenging one. How — technically, socially, and organizationally — can we balance the need for a safe and trusting place with the organizational imperative to share information? One decision facing us as designers is how and to what extent we “design in” norms and social conventions. For example, if we build in technical mechanisms to provide privacy, in addition to the usability impact, we also eliminate opportunities for participants to show that they may be trusted, or to rely on others to respect their privacy. The Babble prototype has no technical features for controlling access: anyone who has access to the client could, in theory, enter any Babble space. But, because Babble makes users visible, this results in groups noticing, commenting on, and ultimately discussing how to deal with this issue. We believe that a greater understanding of how to design systems that permit social mechanisms to come into play is of great importance in designing future systems for knowledge management.

Acknowledgments

Thanks to David N. Smith for creating the Babble prototype, to Mark Laff, Peter Malkin, and Amy Katriel for implementation work on the Babble server and clients, and to Cal Swart for critical assistance in the deployment of a multitude of Babbles. Thanks, as well to members of IBM's Social Computing and Pervasive Applications groups for productive conversations, and to the many dozens of 'Babblers' who have shared their insights, responded patiently to surveys, and, most importantly, used Babble in many of productive (and often surprising) ways.

References

- Ackerman, M.S. (1998); Augmenting organizational memory: a field study of Answer Garden, *ACM Transactions on Information Systems*, Vol. **16**, No. 3, July 1998, p203-24.
- Ackerman, M. S. and Palen, L. (1996) The Zephyr Help Instance: promoting ongoing activity in a CSCW system. In *Human Factors in Computing Systems: The Proceedings of CHI '96*. New York: ACM Press, pp. 268 - 275
- Baym, N. K. (1995) The Emergence of Community in Computer-Mediated Communication. *Cybersociety: Computer-Mediated Communication and Community* (ed. S. Jones). Thousand Oaks, CA: Sage Publications, 1995. pp 138-163.
- Baym, N. K., (1997) Interpreting Soap Operas and Creating Community: Inside an Electronic Fan Culture, In *Culture of the Internet* (ed. S. Kiesler). Mahwah, NJ: Lawrence Erlbaum Associates, pp. 103-120.
- Boone, M. E. (2001). *Managing Inter@ctivity*. New York, NY: McGraw-Hill.
- Bradner, E., Kellogg, W.A., and Erickson, T. (1999) The Adoption and use of Babble: A field study of chat in the workplace. *Proceedings of the European Conference on Computer-Supported Cooperative Work (ECSCW '99)*, 139-158, Kluwer Academic Publishers, 1999.
- Brown, J.S. and Duguid, P. (1995) The Social Life of Documents. In *Release 1.0* (ed. E. Dyson), October 1995. New York: EDventure Holdings Inc., pp. 1-18. Also at: www.parc.xerox.com/ops/members/brown/papers/sociallife.html
- Bruckman, A. (1997) MOOSE Goes to School: A Comparison of Three Classrooms Using a CSCL Environment. In the Proceedings of CSCL 97, Toronto, December 1997.
- Churchill, E. F. and Bly, S. (1999) "Virtual Environments at Work: Ongoing use of MUDs in the Workplace." The Proceedings of WACC '99 (San Francisco, CA, 1999). New York: ACM Press.
- Cohen, D. and Prusak, L. (2001) *In good company: How social capital makes organizations work*. Boston, MA: Harvard Business School Press.
- Cohen, J. (1994). Monitoring background activities. In G. Kramer (Ed.), *Auditory display*. New York: Addison-Wesley, pp. 439-531.

- Donath, J., Karahalios, K., & Viegas, F. (1999). Visualizing conversation. In J.F. Nunamaker, Jr., & R.H. Sprague, Jr., (Eds.), *Proceedings of the Thirty-Second Hawai'i International Conference on Systems Science*. Los Alamitos, CA: IEEE Computer Society Press.
- Ekeblad, E. (1999). The Emergence of Multilogue on a Scholarly Mailinglist, Paper presented in the symposium: "Time and coordination in a virtual community of learners," at the 8th European Conference for Research on Learning and Instruction (EARLI), Göteborg, Sweden, August 24-28 1999. Available at: <http://hem.fyrystorg.com/evaek/writings/earli99/multdyn.html>
- Erickson, T., Smith, D.N., Kellogg, W.A., Laff, M.R., Richards, J.T., and Bradner, E. (1999) Socially translucent systems: Social proxies, persistent conversation, and the design of Babble. *CHI 99 Conference Proceedings: Human Factors in Computing System*, 72-79, ACM Press, 1999.
- Erickson, T. and Kellogg, W.A. (2000). Social translucence: An approach to designing systems that mesh with social processes. *ACM Transactions on Computer-Human Interaction*, Vol. 7, No. 1, pp. 59-83.
- Erickson, T. and Laff, M.R. (2001) The design of the 'Babble' timeline: A social proxy for visualizing group activity over time. In *Extended Abstracts: The Proceedings of CHI 2001*. New York: ACM Press.
- Hiltz, S. R. and Turoff, M. (1993). *The Network Nation* (Revised edition), Cambridge, MA: MIT Press.
- Latour, B. (1987) *Science in Action*. Cambridge, MA: Harvard University Press.
- Latour, B. and Woolgar, S. (1979) *Laboratory Life: The Social Construction of Scientific Facts*. London: Sage. (2nd edition; Princeton: Princeton University Press, 1986).
- Lave, J. and Wenger, E. *Situated Learning: Legitimate Peripheral Participation*. Cambridge: Cambridge University Press, 1991.
- Millen, D. (2000) Community Portals and Collective Goods: Conversation Archives as an Information Resource. Proceedings of HICSS - 33rd Annual Hawaii International conference on Systems Sciences, January 4-7, 2000.
- Millen, D. R., & Dray, S. (1999) Information Sharing in an Online Community of Journalists. In Proceedings of Esprit i3 Workshop: Ethnographic Studies in Real and Virtual Environments: Inhabited Information Spaces and Connected Communities. Edinburgh, January 25-26.
- Nonaka, I. and Takeuchi, H. (1995) *The Knowledge-Creating Company : How Japanese Companies Create the Dynamics of Innovation*. New York: Oxford University Press.

- O'Day, V. L., Bobrow, D. G., and Shirley, M. (1996) "The Social-Technical Design Circle. In the Proceedings of CSCW 96. (November 1996, Cambridge, MA). New York: ACM Press.
- Olson, G. and Olson, J. (2000). Distance matters. *Human-Computer Interaction*, Volume **15**, Vol. 2-3.
- Orr, J. E. *Talking About Machines: An Ethnography of a Modern Job*. Ithica, NY: Cornell University Press, 1996.
- Putnam, R. D. (2000) *Bowling Alone: The Collapse and Revival of American Community*. New York: Simon and Schuster.
- Ramo, S. (1961) The Scientific Extension of the Human Intellect. *Computers and Automation*, Vol 10., No. 2, pp 9-12. February 1961. [Based on a talk given in 1960].
- Schlager, M., Fusco, J., & Schank, P. (1998). Cornerstones for an on-line community of education professionals. *IEEE Technology and Society Magazine*, 17(4), 15-21/40.
- Schlager, M., Fusco, J., & Schank, P. (in press). "Evolution of an On-line Education Community of Practice." To appear in K. A. Renninger and W. Shumar (Eds.), *Building virtual communities: Learning and change in cyberspace*. NY: Cambridge University Press.
- Senge, P. M. (1990) *The Fifth Discipline : The Art and Practice of the Learning Organization*. New York: Doubleday.
- Traweek, S. *Beamtimes and Lifetimes: The World of high energy physics*. Cambridge, MA: Harvard University Press, 1988.
- Turoff, M. (1972) Delphi Conferencing: Computer-Based Conferencing with Anonymity, *Technological Forecasting and Social Change*, Vol. **3**, 159-204.
- Viegas, F.B. and Donath, J. Chat Circles. *CHI 99 Conference Proceedings: Human Factors in Computing System*, 9-16, ACM Press, 1999.
- Wenger, E. (1998) *Communities of practice: Learning, meaning, and identity*. Cambridge: Cambridge University Press.
- Wooley, David. (1994). PLATO: The Emergence of Online Community. On -line version at <http://www.thinkofit.com/plato/dwplato.htm>; an earlier version appeared in *Matrix News* (ed. J. Quarterman), January 1994.