The knowledge management puzzle: Human and social factors in knowledge management

by J. C. Thomas
W. A. Kellogg
T. Erickson

Knowledge management is often seen as a problem of capturing, organizing, and retrieving information, evoking notions of data mining, text clustering, databases, and documents. We believe that this view is too simple. Knowledge is inextricably bound up with human cognition, and the management of knowledge occurs within an intricately structured social context. We argue that it is essential for those designing knowledge management systems to consider the human and social factors at play in the production and use of knowledge. We review work—ranging from basic research to applied techniques—that emphasizes cognitive and social factors in knowledge management. We then describe two approaches to designing socially informed knowledge management systems, social computing and knowledge socialization.

Knowledge management (KM)—also known under rubrics such as organizational learning, organizational memory, and expertise management—has received increasing attention over the last decade. Indeed, it is fair to say that knowledge management is well on the way to becoming a distinct field, with its own theories, jargon, practices, tools, skills, and other accoutrements of an independent discipline. This paper is motivated by our concern that the codification of knowledge management is proceeding a little too rapidly, and that we may end up with a conception of knowledge management that is too neat and too simple to survive in the wilds of the workplace.

The dominant conception of knowledge management—particularly that which has spread beyond the circle of researchers and practitioners into the marketplace—is overly tidy. Knowledge management is seen primarily as a problem of capturing, organizing, and retrieving information, evoking notions of databases, documents, query languages, and data mining. Knowledge is seen as passive, analytic, and atomistic: it is composed of facts that can be stored, retrieved, and disseminated, with little concern for the context in which the facts were originally embedded, and little concern for the new and often quite different contexts in which they will be used. In this view, as one widespread advertisement recently claimed, knowledge management is nothing more than getting the right information to the right people at the right time.

This is a nice picture, but one with which we are not comfortable. Whereas there is no denying the importance of factual knowledge and the usefulness of information technologies, we believe that there are many other issues that are of critical import. Our goal, therefore, is to bring forward a set of results—ranging from basic research findings to practical techniques—that we believe to be very relevant to knowledge management, even as they are at risk of being left out of the KM picture. Overall, our strategy in this paper is to back away from a coherent picture of knowledge management. We suggest that it is more valuable to see knowledge management as a puzzle, especially if we focus on the puzzle pieces: our basic approach will be to add a number of new

pieces to the puzzle, and to demonstrate that some very different pictures of KM can be assembled from the richer, if less ordered, set.

In the next section, "Missing pieces: Cognitive and social research and techniques," we begin by surveying the conceptual landscape that informs our work in knowledge management. This involves looking closely at some of the human and social factors that are involved in the creation and communication of knowledge. We discuss both research areas and applied techniques that, we believe, have received insufficient attention in knowledge management. We do not attempt to provide a single, unified framework for knowledge management, an endeavor that we see as premature; rather, our goal is to broaden the reader's view of what is important and relevant for KM. In the following section, "New pictures: Socially informed knowledge management systems," in lieu of offering a unified KM framework, we describe two distinct projects that, each in its own way, draw upon some of the previously described research and techniques to develop socially grounded approaches to knowledge management.

Missing pieces: Cognitive and social research and techniques

One of the reasons we are dissatisfied with the dominant picture of knowledge management is that it pays little attention to human and social factors. In our view, knowledge is bound up with human cognition, and it is created, used, and disseminated in ways that are inextricably entwined with the social milieu. Therefore, we argue that knowledge management systems must take both human and social factors into account. In this section we describe a number of the research findings and applied techniques that motivate our work. We believe that these pieces are vital parts of any picture of knowledge management. At the same time, we acknowledge that there are undoubtedly other missing pieces to the KM puzzle, and that many distinct, but still valid, pictures of KM are possible.

The missing pieces we discuss are quite diverse. They are drawn from a variety of areas ranging from the cognitive and social sciences, to domain-focused disciplines such as social studies of science and computer-supported cooperative work. These pieces reveal the complexity and some of the subtleties underlying the mantra of "the right information to the right people at the right time." To show this more clearly and to provide a bit of structure in what follows, we discuss our pieces in terms of knowledge ("the right information"), presentation and communication of that knowledge (". . . to . . . at the right time"), and social context ("the right people"). Following that, we turn to applied techniques that are relevant to these areas.

Knowledge and intelligence. Until fairly recently, the prevailing view of science held that the world was, in principle, knowable and predictable; that the universe, including human beings, consisted of essentially complex but analytically decomposable machines.[1,2] A common metaphor for knowledge, still quite common in Western society, is that it consists of separate little "beads" or factoids,[3] and that these knowledge "atoms" can be collected, stored, and passed along. Views like this are what underlie the notion that an important part of knowledge management is getting access to the "right knowledge." Although, obviously, it is important to find knowledge that is relevant to whatever problem is at hand, there is quite a lot of research that paints a considerably more complex picture of knowledge.

To begin with, let us take a look at some findings from research in the area of human intelligence. Outgrowths from the endeavor to test "intelligence" over the last century have led to an understanding that there are different types of intelligence that work primarily on different forms of knowledge. Although there are variants on this theme, the most popular recent work, as well as one having a sound empirical base, is probably that of Sternberg.[4,5] Perhaps the most ambitious and elegant theoretical framework was developed by Guilford,[6] who built a three-dimensional model of mental processes. In this work, there were differently sized *Products* of mental operations: Units, Classes, Relations, Systems, Transformations, and Implications. There were different *Operations* (processes) that could be performed: Cognition, Memory, Divergent Thinking, Convergent Thinking, and Evaluation. Finally, there were different types of *Content*: Figural, Symbolic, Semantic, and Behavioral. While this system has largely fallen out of favor as a basis for testing intelligence, it is an interesting framework for KM developers to consider. All too often knowledge management systems are designed with an implicit, unquestioned, and unacknowledged limitation on the varieties of knowledge that are supported.

The field of intelligence testing is relevant to knowledge management in yet another way. Early developers of intelligence testing failed to recognize the

extent to which their tests were measuring, not innate capability, but essentially the degree to which someone had acquired socially sanctioned knowledge. Test developers struggled to develop "culture-free" IQ (intelligence quotient) tests and by and large failed in the attempt, realizing at last that what constitutes intelligence is primarily determined by culture. Perhaps the most telling example in this regard comes not from the field of intelligence testing per se, but from the work of Tom Evans,[7] an AI (artificial intelligence) student of Marvin Minksy, who built a program to solve figure analogies of the form "A is to B as C is to [D1, D2, D3, D4, or D5]." Evans's program worked—too well. It could parse the figures, construct rule components, and find a composed rule that made *every* answer correct! Fully half the work of the dissertation was essentially to get his program to have the same ordering of "elegance" of rules that was socially agreed upon by the test makers. As one example, the authors of the tests thought it more "elegant" to rotate a figure in the plane of the paper as opposed to out into three-dimensional space. In other words, even knowledge that could easily be thought of as factual or mathematical is in fact strongly shaped by social and cultural assumptions.

If even factual knowledge is not quite as objective as we might expect, it is not surprising to find that other forms of knowledge are even more subjective. For example, one important early debate in psychology centered on introspectionism versus empiricism. This debate arose in part due to inconsistencies in subjects' self-reports of experiences of perception and consciousness. At the time, the scientific community reacted by declaring that only objectively observable phenomena should be used in building a reliable understanding of mental processes; today, in the wake of the failure of the behaviorist project, there is greater openness toward subjective forms of knowledge. Although it is clear that some kind of "self-knowledge" is essential for people to behave intelligently (e.g., without knowledge of the limits and capacities of our bodies we might continually be running into things), individuals differ on how such knowledge is best viewed.

In addition, research has shown that there are a number of important cases in which a person's self-knowledge is inaccurate. In the "fundamental attribution fallacy" literature, studies show that the behavior of an individual is highly influenced by context, and yet people give explanations for their behavior based on their own internal values. For example, bystander studies consistently show that people are much more likely to help a person in distress if they are alone rather than if they are with a large group, and yet, when asked whether they would respond differently depending on how many others are present, people claim that it would make no difference.[8] This has important implications for modern knowledge management practices. Not only are people very much influenced by the social context, they may believe that they are not so influenced, when they in fact are. Although some have pointed out that the productivity of both teams[9] and large organizations[10,11] is pervasively influenced by social context, we believe the impact is often underestimated, not only by subjects in social psychology experiments but also in everyday business decisions about knowledge management.

Communication and learning. If knowledge is not so simple as our ways of talking about it assume, neither is the process of communicating it to others. As Brown and Duguid[12,13] note:

> The idea of a document as a carrier is an example of what Michael Reddy calls a "conduit" metaphor. People regularly describe most communication technologies in conduit terms, talking of information as "in" books, files, or databases as if it could just as easily be "out" of them. We ask or are asked to put ideas "down on paper," to "send them along," and so forth.

However, there is quite a lot of research that suggests that it is not just a matter of *getting* the right knowledge *to* people—people need to engage with it and learn it. One of us has argued that a more realistic and useful model of communication is a "design-interpretation" model. In this model, the speaker uses knowledge about the context and the listener to design a communication that, when presented to and interpreted by the listener, will have some desired effect.[14] In the "design-interpretation" model, a knowledge worker would be viewed in an active, constructionist role, consistent with a wide variety of empirical results.

There is quite a lot of research that is relevant to this view. Theorists as disparate as Dewey,[15] Vygotsky,[16] and Piaget and Inhelder[17] have consistently shown that the mere presentation of information does not necessarily result in learning. People have to become actively involved for behavior to change, for insight to occur, for problems to be solved. Vygotsky stressed that this learning and insight had a

significant social component, even if the resulting knowledge was of a type we might classify as mathematical or scientific. Yet, all too often, large organizations come to believe that simply making more information available more widely will "solve" knowledge management problems. By way of contrast,

It is becoming clear that knowledge work is not a solitary occupation, but it involves communication among loosely structured networks and communities of people.

within IBM much of the management training is done via scenario-based training. In this technique, the individual is asked to make choices in realistically portrayed situations such as ones that managers face. These scenarios are based on an analysis of real situations, and assume that when the individual makes a "mistake" in the simulator or is "surprised" by a result, it motivates the person to read and understand the rationale.[18] In the use of such simulators, even if the individual learner is sitting alone in front of a computer console, learning is very much influenced by social context. It is the social context of the scenario that provides much of the motivation and interest as well as guidance on what constitutes a "right answer."

In addition to arranging interactions so that people actively engage with knowledge, there are other considerations from earlier work that are applicable to knowledge management systems. We know, for example, that people are better able to both distinguish and remember knowledge that is encoded on multiple dimensions.[19] However, in contrast to the variety of sensory cues that naturally occur in real-world "paper" systems, many current generation systems provide little in the way of differentiating cues. Given the processing power and memory of today's computers, it would be quite feasible instead to provide sensory "signatures" that are unique to various items. "Folders," for instance, could easily be portrayed not only in different colors, but also by different sizes and textures. Indeed, small musical animations could even hint at the structure or content of a folder or its date of last access. Of course, a challenge in convincing organizations to adopt sensory-rich approaches to laying out a knowledge space is that performance improvements may only be observable after extended usage.

A large number of indicators point toward the reality of an information-processing world moving toward greater fidelity and multimodality. Over the last four decades, user interfaces have evolved from lights and toggle switches to keyboards, mice, icons, and speech I/O. In the entertainment industry, we now see computer-generated full-length movies. Video games strive toward greater responsiveness, more modes of experience, and more detailed images. Research laboratories continue to push the boundaries of multimodal I/O, including virtual reality and augmented reality,[20,21] auditory icons,[22] kinetic typography,[23] and so on. Yet, in a business context, knowledge management writings and practice often seem to focus on the content of systems while ignoring the method of presentation. Beyond considerations of cost, there sometimes seems to be almost a puritanical business-culture ethic toward avoiding presentations that stimulate the senses and utilize the complete human brain.

Social context. Although it is difficult to argue that knowledge management should not be concerned with getting information to the "right people," our common definition of KM provides little insight into which people are the right people. Here we turn to a body of work—drawn primarily from the social sciences and domain-oriented disciplines like Computer-Supported Cooperative Work (CSCW)—that provides some interesting views on the social contexts (both with and without technology) within which knowledge work occurs, and the social factors that seem to be important in supporting knowledge work.

In the field of Computer-Supported Cooperative Work, researchers have often made careful examinations of how people within organizations conduct their work. One point that these studies make is that knowledge work is not a solitary occupation, nor is it sufficient to say that knowledge work involves many people. Rather, in case after case, it becomes clear that knowledge work involves communication among loosely structured networks and communities of people, and that understanding it involves identifying the social practices and relationships that are operative in a particular context. One of the best-known concepts to emerge from such studies is Lave and Wenger's notion of a community of practice. A community of practice is defined by common tasks, methods, goals, or approaches among a group of people. Lave and Wenger show how new workers come to

master a body of knowledge through a sort of apprenticeship or "legitimate peripheral participation" in the activities of a group of experienced workers.[24] Wenger[25] provides a detailed study of an insurance claims processing office, and shows the vital roles that social relationships and processes play in enabling people to meet productivity targets while adhering to corporate policies. Orr's[26] study of photocopier technicians 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.

Similarly, in reviewing ten years of field and laboratory studies of collocated and remote work, Olson and Olson[27] point to a variety of social factors that affect the social context of knowledge management, and how these interact with technologies intended to support remote collaboration. In an interesting discussion of the role of common ground[28] among collaborators, for example, the Olsons describe how greater shared background and awareness of a coworker's activities and mental state contribute to establishing and maintaining common ground. The Olsons also discuss the role of motivation in successful knowledge sharing:

> Motivation has been established as one of the major sources of failure in adoption of groupware in general. In Orlikowski's classic study of the failure to adopt Lotus Notes** in a consultancy, the failure was attributed to the fact that individuals were compensated according to their competitive talents.[29] There was no incentive to share one's best ideas if they were then going to be seen as common, no longer unique. In other organizations, where incentives are aligned with how much others use the knowledge you make available to them, Notes and other jointly authored groupware systems succeed.[27]

Churchill and Bly's[30] study of the use of a MUD (a kind of text-based conversation environment) among a group of scientists at Argonne National Laboratory points to other social factors supporting knowledge sharing and collaboration:

> As Huxor[31] points out, however, "chance" encounters do not occur entirely by chance. Social ties and physical layout can have an effect on who contacts whom and how often. Hillier[32] emphasizes the effect of work environment design (whether virtual or physical) on the establishment and maintenance of "weak ties." These are contacts that one normally would not make through one's central work practices. Arguably these "weak ties" have a strong role to play within an organization's functioning. In this vein, the MUD provides a set of virtual places wherein one can "bump into" others or people can be actively sought. Our interviewees stressed the importance of such planned and unplanned encounters.

Thus, we see that talking about knowledge management as though it involves delivery of information to a person, or to a set of people, is missing quite a lot. Looking at the individual is, by and large, looking at the wrong level of granularity; instead, we need to shift our focus to the social context. That is, most of the phenomena that have been identified as important—relationships, awareness, common ground, incentives, and motivation—are network or social phenomena. In the same way, the other "missing pieces" we discussed earlier are also characterized by a shift in granularity: rather than knowledge as isolated, context-free facts that could be "in" documents or databases, and straightforwardly transferred into people's heads, we see that knowledge is bound up with human intelligence, shaped by social assumptions, and requires active engagement on the part of recipients if it is to be taken up. All of these findings, although just a small sampling of what is contained in the HCI (human-computer interface) and CSCW literatures, demonstrate the pervasiveness of human and social factors in the realm of knowledge management.

Practical techniques for creating and communicating knowledge. As practitioners of knowledge management we have a repertoire of techniques that are aligned with the research we have discussed thus far. And, as designers of systems for managing knowledge, we have conducted a number of explorations on how to provide technological support for some of these techniques. Here we describe the most generally useful techniques, and we examine the ways information technology (IT) could support them. To begin with, we turn to an area that, in our view, has received remarkably little attention: the creation of new knowledge. The lack of attention may be because Western culture has a long tradition of treating creativity as tantamount to magic. Nevertheless, there is now a large body of scientific literature dealing with such issues as creativity, problem solving, and design (see Shneiderman[33] for a nice review). While the state of the science and art is not at the point where we can duplicate the accomplishments of a Shakespeare or Einstein on demand, we can craft

technological and methodological support to increase the creation of new knowledge, both by individuals and by groups. Tools can be designed to help with various processes in the creation of new knowledge. We examine several examples of creative processes that could be feasibly supported by technology including: dialog,[34] the use of metaphors,[35] the use of strategies,[36] and the use of stories.[37]

Bohm Dialogue. In a Bohm Dialogue,[38] a group of people work together to build new knowledge. This process differs from a typical business meeting in several ways. First, continued inquiry is balanced with striving for an answer. We know from studies of human problem solving that people's natural tendency (at least in American culture) is to jump on the first formulation of a problem and try to solve it, rather than exploring alternative formulations. Second, the dialog is noncompetitive. People ask questions and make observations but are asked to "suspend" their thoughts; that is, not to own or push for their specific idea to be adopted by the group as the correct one. Third, the actual conversation is not bound by an analytic agenda that deals with pieces of a problem one by one. Rather, a "container" binds the conversation; that is, there is some focus to the conversation—it does not wander aimlessly over any topic—but everything said is related to the overall system that the group is attempting to understand. Fourth, the rhythm of the Bohm Dialogue is different from the typical meeting, in which everyone begins mentally critiquing the speaker and rehearsing a counterargument before the speaker is even finished speaking. Instead, one person speaks and others listen. After listening, everyone reflects on what was just said before someone else speaks. Bohm, a physicist, likened this "cooler" pace to superconductivity. Whether we accept such a metaphor or not, some remarkable breakthroughs have come from using the Bohm Dialogue.[38]

Tools have been developed to support Bohm Dialogue, in face-to-face settings as well as in remote collaborations. Fischer and his colleagues at the University of Colorado have built a collaboration space that includes both a horizontal, computerized action space—where collaborators literally design and build "in the center"—and a large-screen, large-scale reflection space[39] where material designed to help participants reflect on their activity is displayed. It is clear that for Bohm Dialogue to work as a process, however, both formal organizational mechanisms and informal cultural aspects of a situation must lend themselves to the process. If people are operating in a highly competitive culture in which there are higher intrinsic and extrinsic rewards for having their own idea adopted rather than for, as a group, reaching a breakthrough, dialog will just be seen as a slow, inefficient way to run a meeting. Similarly, in order to be effective, the participants will probably need to have at least some minimal familiarity with systems thinking; otherwise, there will be little motivation to attempt to understand the whole rather than attack the parts piecemeal.

Systematic use of metaphor. Another creative process that can be supported with information technology is the systematic use of metaphor, as suggested by Gordon,[35] Mattimore,[40] and others. Earlier research[41] described a technique for improving performance in problem solving and design tasks. In one experiment, a group of subjects spent an hour creating a design for converting an abandoned church into a restaurant. Another group spent 20 minutes designing, 20 minutes looking at a word list designed to evoke concepts from a wide range of domains, and another 20 minutes designing. The latter group produced designs that fulfilled significantly more of the functions of a restaurant. In another experiment, subjects given an arbitrary word list exhibited more insightful solutions to problems, despite constant time on task. These results suggest that word lists specifically crafted to evoke different metaphors in the mind of a problem solver can aid open-ended design and problem-solving tasks, and that such metaphor techniques could be extended and improved with today's interactive software. Preliminary explorations of what such tools might be like can be found in the demonstration prototypes available at www.research.ibm.com/knowsoc/.

How such a tool might work is illustrated in the following fictitious scenario.

> Joe is stuck. He has what he thinks is a great idea for making the next version of the ThinkPad* easier to use. But he is getting nowhere with the project manager who refuses to consider hardware changes at this late date. He logs on to the automated metaphor tool and engages in a structured interaction that encourages him to break down the elements of his problem, rearrange them, make his implicit assumptions explicit, and play "what-if" games with some of these assumptions. He gets a "funny" feeling when he turns the definitions of hardware and software upside down, but still has no solution. However, later that day, as he is driving home, he realizes that what he

wanted to include in modified hardware could also be provided as software. Fulfillment would not have to be done through the PC division, per se, but could be downloaded via the company Web site or sent as a CD in response to a call to an 800 number. As a result, PC sales increase dramatically once word gets out that IBM has a very easy-to-use system.

Strategy mapping. Recent work at IBM Research has focused on developing a comprehensive compendium of strategies for enhancing creativity and knowledge creation. Although there have always been books about strategies for problem solving in particular domains (e.g., war, diplomacy, chess, business), we believe that people within a particular community of practice often develop a common set of strategies that are then taken more or less for granted within that community. There may be strategies, however, from a totally different domain that might be successfully applied, if only the practitioners knew about the existence of these strategies. Gordon[36] has collected strategies from a large number of disparate domains and developed an abstract planning language in which all the strategies can be described. Since it is generative, the abstract planning language might be used to articulate potentially novel strategies as well as to help practitioners in one domain find unusual strategies from some other domain.

For example, in traditional "Western" medicine, germs have been considered the "enemy" and strategies of care for someone who has an infection typically have involved attacking and killing these enemies. However, other strategies might include boosting and supporting the body's own immune system, altering the characteristics of the germs so that they become harmless, "leading them out" of the body by providing a more hospitable environment elsewhere, altering the germs so they attack each other, coating the germs so they can no longer affect the body, etc. Our conjecture is that providing alternative strategies from other domains might enable doctors to engage in breakthrough thinking. In fact, various drug companies are considering the use of alternative strategies to facilitate innovative thinking.

Stories and storytelling. Stories and storytelling provide another possible way to foster creativity in individuals and groups, and they also provide a valuable way of presenting and communicating knowledge. In some cases, particular stories can illustrate a specific point. One fairly common yet difficult point to make in teaching the concepts of systems thinking is the kind of mutual impact that people have on each other. For example, a marketing department may feel that the engineering department is unresponsive and takes too long to make changes. To counter this, the marketing department may develop a whole suite of requirements and ask for them earlier than is actually necessary, hoping to "speed up" development so that enough features will be provided for a timely, competitive product. Of course, such behavior makes the engineering department feel less like being responsive to marketing. Breaking out of such "vicious circles" is difficult. Direct communication can often backfire under these circumstances, because it can trigger defensiveness and defensive countermoves. An alternative suggested here is to provide a story to both groups about another situation in which the same principles apply. Snowden[42,43] reports several business cases in which the use of various types of stories has helped to produce breakthroughs.

Finding appropriate stories for the situation at hand, however, is nontrivial. In our laboratory, we are developing tools to help. In one such tool, Gordon[44] describes a "script-based browser" that allows a user to find stories based on the type of activity they contain. This approach has been applied to a very large story collection called the "American Heritage Project"—stories commissioned in the 1930s by the Works Progress Administration, many of which are available on line. As work progresses on the abstract planning strategy language described previously, the browser can be used to find stories about analogous activities as well.

In some cases there are other characteristics of a story that may be important in selection. Our laboratory has begun developing a Story Markup Language for describing the various aspects of a story. We plan to develop software for either adding meta-data to stories automatically or helping a user do it in a straightforward fashion. Such meta-data might be used to search for specific kinds of stories or could be used as the basis for visualizations of the set of stories that users can quickly scan to find likely candidates.

The Story Markup Language not only deals with the internal content of the story but also with the social context. Storytelling is fundamentally social: in everyday events, people tell stories to specific other people (who are usually physically present) in particular social contexts (at dinner, in a meeting, etc.).

Social factors influence who tells what stories to whom and when. In designing effective ways to collect and provide access to stories, we think it is important to attend to some of the basic social dynamics that affect everyday storytelling, such as reasons for telling stories, the teller's knowledge of the audience, and the role the audience takes in the telling.

Expressive communication is communication in which individuals or teams are primarily motivated by personal or social aims such as sharing experiences, indicating agreement, or being humorous.

As one example of how the social context of storytelling can influence its teaching effectiveness, we must recall that, in a business context, the audience of a story does not simply "take in" the story. In the case of fictional stories (e.g., stories told in an entertainment context), readers and listeners will "buy" the story as long as it is internally consistent. But in the context of using stories to foster change in the real world, the audience must not only see the story as internally consistent, but also as consistent with external reality. An elaboration of such social factors in storytelling and their implications can be found in Lawrence and Thomas.[45]

Expressive communication. Communication is central to any complex, modern organization. We find it useful to draw a distinction between what we might call instrumental versus expressive communication. Instrumental communication is that which is necessary for accomplishing tasks related to the immediate organizational goals. It is typically supported by specific forms and media, for example, job offers, requisitions, ratings, invoices, RFPs (requests for proposals), contracts, formal award certificates, and so on. Such communications are typically created in well-defined formats (e.g., forms), delivered in specific contexts, and need to include specific information. In contrast, expressive communication is communication in which individuals or teams are primarily motivated by personal or social aims such as sharing experiences, indicating agreement, being humorous, etc. Expressive communication often occurs in informal settings including hallway conversations, informal meetings, stories, notes scrawled on congratulatory cards or napkins, and e-mail about nonbusiness issues.

Organizations clearly rely on instrumental communications, but the role of expressive communications is less clearly understood, with perceived value varying from irrelevant, to disruptive, to vital to the accomplishment of work. Recently some have argued that such communications are important in supporting innovative thinking[46] and the building of social capital[11] within organizations. In any case, it is clear that the effective use of expressive communication within organizations requires certain conditions, e.g., appropriate social and cultural norms. Technology can enhance the use of expressive communication. One idea is to simply modify existing forms to include a field for comments and contact information for an ombudsman, a person responsible for solving problems associated with a mismatch between the assumptions of the built systems and the workday realities. Providing a place for such commentary and the ensuing conversation could be vital for supporting the instrumentality of the system.

An instructive real-world example is provided by Harris and Henderson.[47] In a shipping firm, paper forms were being replaced with a more "efficient" computer system. In one case, a shipment was to be delivered to a boat. In the paper version, in the field marked "Ship to Address," a worker wrote "Call Mr. X at number Y to see where the ship will be at time of delivery." Entering the same data in the computer version yielded an error message. The worker would then enter a bogus, but syntactically correct address in the computer form. As a result of many such errors, the workers ended up using the computerized system *and* the old paper system.

A two-fold approach can be taken for the prevention of such absurdities. First, as a general principle, knowledge systems should provide for expressive communication. This will result in fewer errors and more effective operations; it will also enhance social capital. Second, systems should be designed with an understanding of how work is actually done, and not how IT developers think the work is done. Processes for developing such understanding can be found, for example, in Beyer and Holtzblatt.[48]

Expressive communication as a means of building trust. In Arie De Geus's[49] classic study on the longevity of large organizations, he found that mutual trust was a strong characteristic of long-lived organiza-

tions. In Robert Putnam's[50] study of various regions and local governments in postwar Italy, he also found that mutual trust, facilitated by various informal groups, clubs, and associations, was an excellent predictor of future economic growth. Instrumental communication may inform about a person's competence and reassure us that the person follows the organizational rules. Character, however, is revealed by choices under pressure.[51,52] In a well-structured organization, instrumental communication minimizes choice; hence, it is difficult to learn about someone from purely instrumental communication. On the other hand, if a person tells a story about him or herself, has a social conversation, or participates in creative design sessions, he or she will inevitably reveal something personal. Over time, we learn something about another person and may come to trust them. In fact, there is some evidence that people prefer people who use more expressive means of communication, even in organizational settings.[53] This may be one reason why effective leaders turn to story.[54,55] The importance of mutual trust in coworkers may not be evident if people in an organization are all following a procedure that works. However, in times of change or breakdown, mutual trust will allow collaborative effort to proceed toward organizational (as opposed to individualistic, locally optimized) goals. Looked at another way, what expressive communications allow people to do is build up a more complete and complex model of others so that there will be a basis for behavioral prediction in novel situations requiring conjoint but not prechoreographed action.

Of course, while stories and other forms of expressive communication have the capability of building mutual trust, they also have the capability of reducing mutual trust. As mentioned above, stories in an entertainment context create their own frame. But stories told in the context of an organization will not simply be "accepted"; they will be viewed against the backdrop of the current context and if the story told is too discrepant from actual behavior, one result will be less trust, not more.[37]

Expressive communication is not a sufficient condition to build mutual trust; however, it may be necessary. Hence, the design of knowledge management systems would do well to support expressive communications as well as instrumental ones if they are to help organizations thrive in times of change and adversity. Some recent designs seem to be moving in this direction.[37,56]

Conversation. We conclude our review of practical techniques for supporting a socially informed approach to knowledge management by turning to a technique that is so common as to go unremarked: conversation.

We view conversation as essential. We use it as a medium for decision-making. It is through conversation that we create, develop, validate, and share knowledge. When systems—automated or bureaucratic—freeze, or simply prove too rigid, we pick up the phone to figure out appropriate "workarounds." And with all our advances in information retrieval, the preferred method for obtaining information is still to ask a colleague.

Why is this? We suggest that conversation has two characteristics that are central to its power and ubiquity. One vital characteristic of talk is that it is a deeply interactive intellectual process (see Clark[28] for a detailed exposition). As we talk we refer to a common ground of already established understandings, shared experiences, and past history. As the conversation proceeds, we are continuously attempting to interpret what is said, verify that we have been understood, and offer new contributions. Sometimes misunderstandings occur, and so we attempt to fix them by rephrasing our words, or "debugging" the previous conversation to reveal that what we thought were shared understandings were not, in fact, shared. What all this amounts to is that conversation is a superb method for eliciting, unpacking, articulating, applying, and recontextualizing knowledge.

Conversation is more than simply an intellectual endeavor: it is a fundamentally social process.[57–59] Conversation is social in two ways. First, people speak to an audience. Speakers notice how their audience is reacting and steer their remarks appropriately: nods and eye contact convey one message; questions and furrowed brows another; yawns and fidgeting still another. Second, conversation is social in that people portray themselves through conversation. They advance their personal agendas, project their personal style, take credit, share blame, and accomplish other social ends through their talk, often with a great deal of subtlety. The social nature of talk is not an undesirable side effect, but rather the heart of it: personal motivations fuel conversation and provide the energy for the considerable intellectual work it takes, whether the conversation in question is banter over morning coffee or discussing the composition of a journal paper.

In addition, conversation within the digital medium, has a property of great importance for our purposes: it can *persist*. Instantiated as text, whether typed in or spoken and recognized, persistence expands conversation beyond those within earshot, rendering it accessible to those in other places and at later times. Thus, digital conversation may be synchronous or asynchronous, and its audience intimate or vast. Its persistence opens the door to a variety of new uses and practices: persistent conversations may be searched, browsed, replayed, annotated, visualized, restructured, and recontextualized, with what are likely to be profound impacts on personal, social, and institutional practices.

Summary. This concludes our review of an array of research findings and practical techniques that have to do with the cognitive and social factors that come into play in the creation and communication of knowledge. We hasten to note that we do not claim this review is, in any sense, comprehensive. We have focused on pieces of the knowledge management puzzle that, while missing from many accounts, have informed our own explorations. We have no reason to believe that we have somehow, through either good luck or cleverness, uncovered all the missing pieces. In our view, even as it condenses into a coherent discipline, it is important to recognize that our understanding of the critical factors in managing knowledge is in its infancy.

New pictures: Socially informed knowledge management systems

So what are we to make of all this? We have outlined several new pieces to the knowledge management puzzle—but how do these new pieces fit with the old ones and provide a new and seamless picture of knowledge management? As we noted in the introduction, we think such a goal is premature, and that the field would be better served by a multiplicity of partial models and a commitment to explore multiple perspectives. In line with this view, rather than trying to present a unified framework, we offer two examples from our own work that suggest how all the pieces might start to come together in a socially informed approach to knowledge management systems.

The first example, a system called "Babble," comes from an area known as social computing.[56] Social computing has to do with digital systems that draw upon social information and context to enhance the activity and performance of people, organizations, and systems. Instances of social computing systems include recommender programs (e.g., for movies or music), "wearware" (i.e., showing signs of "wear" or history in a digital system, such as how heavily traversed a Web link is), social navigation,[60] and social awareness indicators (e.g., visualizations of people and their behavior, buddy lists). Babble is an on-line multiuser environment that is intended to support the creation, explication, and sharing of knowledge through text-based conversation. The basic rationale underlying Babble is described first, and the system and usage experience with Babble follow.

The second example, knowledge socialization, describes a constellation of projects around the use of stories and storytelling in business settings. The rationale for knowledge socialization is described first, showing how stories can facilitate knowledge creation, sharing, and reuse. Following this, we describe an integrated suite of story-related tools that supports these ends.

The rationale for Babble: Visibility yields awareness yields accountability. Imagine a knowledge management system that was designed from a social perspective, a system predicated on the assumption that knowledge is rooted in a social context. Such a system would assume that knowledge is produced within, and dispersed among, a network of people; that only a small proportion of knowledge is captured in concrete form; that knowledge sharing involves social factors like relationships, trust, obligation, and reputation. One way in which a system might instantiate such assumptions is not just by providing access to data and documents, but also by interconnecting the social network of people who produced the knowledge.

Imagine further that we include not just the people who produce the knowledge, but those who use it as well. Suppose that—just as we look for crowded restaurants, eye fellow shoppers, or look for engaging conversations—we could see similar traces of those making use of information in a knowledge management system. After all, some of the knowledge users might have to invest considerable effort in order to apply the knowledge to their own ends, developing an understanding of its shortcomings and particularities, as well as building on it. If we could capture traces of this knowledge work, others with similar needs might find as much value in talking with users of this knowledge as with the original authors. Such a system would not be just a database from which workers retrieved knowledge, it would be a

focus for a knowledge community, a place within which people could discover, use, and manipulate knowledge, and encounter and interact with others who are doing likewise.[61]

How might we do this? One way in which this might be achieved is in the context of an on-line computer-supported communication environment. Thus we might imagine a computer-based system that not only allows records and documents to be stored, but allows people to converse with one another and to have some visible presence. Such a system would not just make people and knowledge visible, but it would make interactions among them visible. That is, it should be possible to see people interacting with explicitly expressed knowledge (e.g., reading), and it should be possible to see people conversing with one another (both as a means of explicating tacit knowledge and as a means of building and maintaining the social factors such as trust and relationships that are important in knowledge management).

For the last four years our research has been focused on ways of making such interactions visible in on-line environments. We have developed the notion of socially translucent systems[46] to guide us in designing such environments. By social translucence we mean systems that provide perceptually based information about the presence and activity of users, thus creating social resources that the group as well as individuals can use to structure and enhance their on-line interactions. When such information is made visible to all participants, people become aware of one another's presence and activity, allowing social conventions and other social dynamics to come into play. With mutual awareness comes accountability for one's actions (e.g., if "I know that you know that I know" of your presence and activity, my activity will be interpreted with respect to that knowledge; that is, I will be held accountable for my actions). By invoking social translucence as a framework, we are attempting to make people and their behavior more prominent, enabling the creation, exercise, and mutual observation of social behavior. By so doing, we aim to create a basis for more coherent, productive, and fluid interactions on line. Socially translucent systems make it easier for users to interact in purposeful, coherent ways; to observe and imitate others' actions; to engage in peer pressure; to create, notice, and conform to social conventions. We see social translucence as a fundamental requirement for supporting most types of communication and collaboration in digital spaces.

Babble: An infrastructure for a knowledge community. Babble is an on-line, digital space in which knowledge can be created, discovered, shared, and reused.[56] One of the principal means for fulfilling these core knowledge management processes in Babble is its support for blended expressive and instrumental communication through informal conversation. Babble is therefore particularly concerned with supporting long-running, contextual interactions (as opposed to short-term, task-focused activities). It very deliberately blends work and social talk, synchronous and asynchronous interactions, and private and public discourse. The aim is to provide a digital substrate upon which knowledge communities can grow, and where "discourse bases," rather than databases, can provide a medium for people to develop, share, and reuse experiences and knowledge, and watch others do the same.

The Babble system. Babble resembles a multichannel, text-based chat system to which many users can connect, and either select from a list of conversations to participate in, or create their own. Babble differs, however, from conventional chat in two ways, both of which stem from our goal of creating a socially translucent system that supports knowledge management. First, the textual conversation that occurs in Babble is persistent: that is, unlike conventional chat where newly arriving users see only what has transpired since they have joined a channel, Babble users can see everything entered 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 through a variety of means, but principally through what we call a social proxy.

Figure 1 shows the Babble user interface. The upper left pane contains a list of the names of users currently connected to Babble. The middle upper pane contains the social proxy, which we describe shortly. The upper right pane displays a hierarchical list of the conversation topics, grouped in categories and subcategories. 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 the date and time of its creation. Babble conversations need not be synchronous; indeed, some are asynchronous, with hours, days, or weeks separating comments. A variety of other functions may be invoked via the menu bar, or through context-sensitive menus

Figure 1 The Babble interface

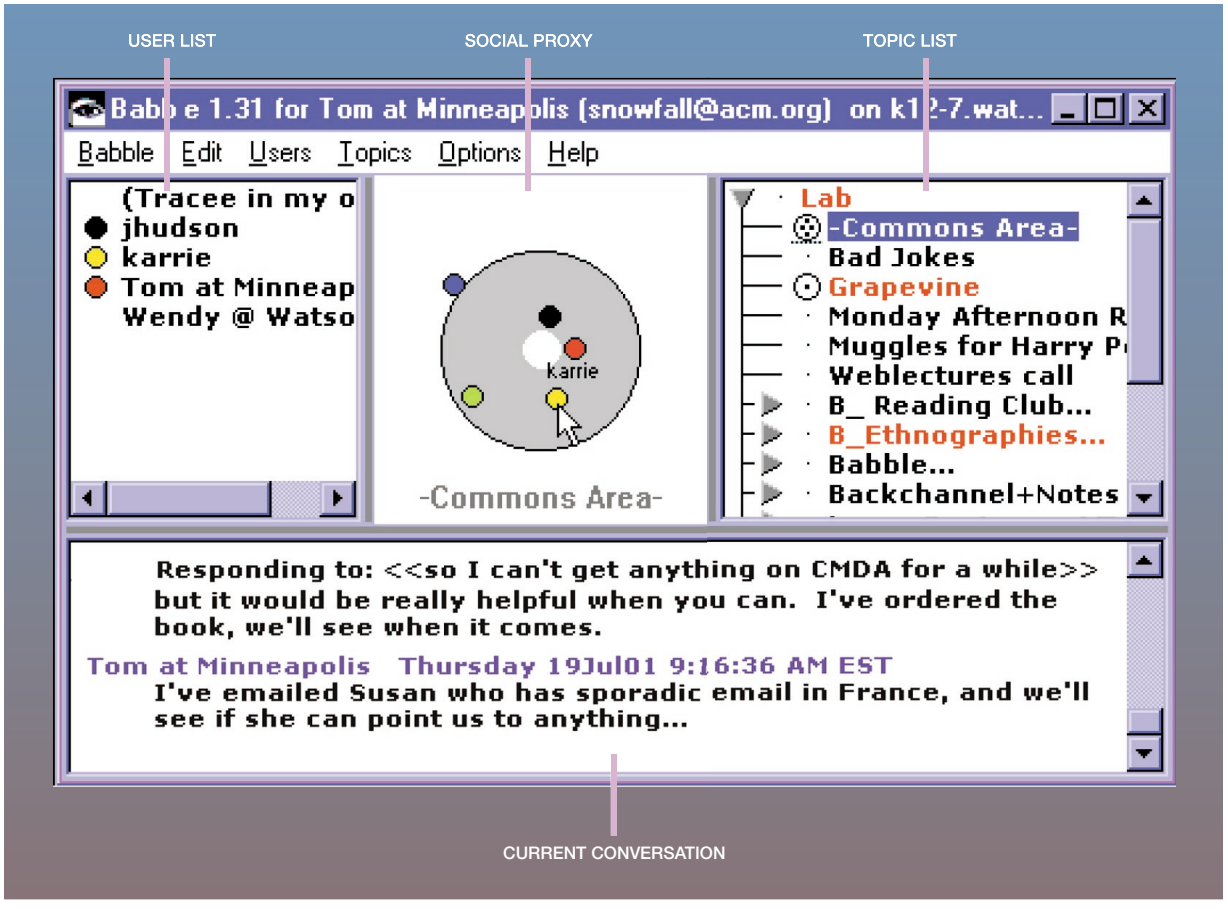

accessed via right clicks, and keyboard shortcuts. These include functions for creating messages, creating, changing, and deleting topics and categories, conducting private, ephemeral chats, and so forth.

The social proxy, also known as the cookie, in the upper middle pane, represents the current conversation as a large circle, and the participants as colored dots, also known 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 to the inner periphery of the circle over the course of about 20 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 three people have been recently active (i.e., speaking or listening) in the current conversation, and one other has been idle for a while (and a fifth person is off in the "Grapevine" topic).

When people leave the current conversation their marbles move to the outside of the circle's periphery (as with the marble at "eleven o'clock"); when they enter the conversation, their marbles move inside the circle. When a person logs onto the system, a virtual wedge is created for the person's marble,

Figure 2 The timeline

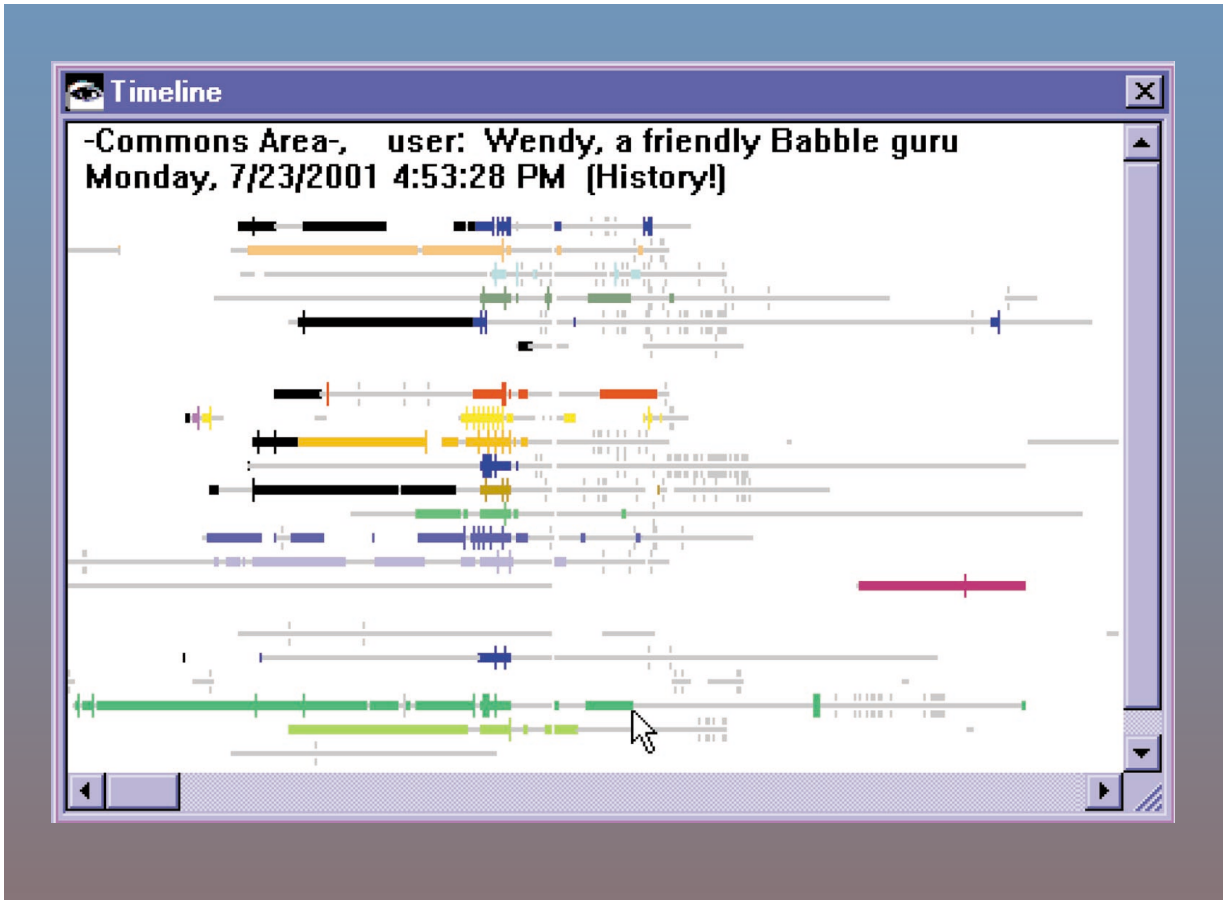

adjusting the position of all the marbles in the social proxy; when the person departs, the wedge is destroyed, and the remaining marbles adjust to uniformly occupy the space. All marble movements are shown with animation, thus making arrivals, movements, and departures visually salient. Although simple in concept, this social proxy gives a sense of the size of the audience and the degree to which the audience is actively listening or contributing, indicates whether people are gathering or dispersing, and who is coming or going.

In addition to the social proxy, Babble uses additional mechanisms to reveal the presence and activity of users. In the topic list, to the left of the topic names, are "minicookies," thumbnails of the social proxy for each topic with at least one participant in it. So, in Figure 1, we can see that there is a single person in the third topic, "Grapevine." Babble also highlights information that the user has not yet seen: the names of topics and categories with new material in them are shown in red (e.g., Grapevine and B_Ethnographies), and within the conversation pane, comments that have been added since the user last "touched" Babble have their authors' names in red.

The cookie shows only synchronous interactions—that is, it shows only the presence and activities of people who are currently logged on to Babble. This may be a drawback because often 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, the timeline,[62] illustrated in Figure 2.

The basic goal of the timeline is to provide a way for a "speaker" to see that people were "listening" (or not), even when the listening was offset in time. Each user logged on to Babble is represented by a row. When the user "speaks," a vertical mark or blip appears on the line. If the line/blip is in color, it means the user was active in the conversation currently being viewed by the user of the timeline; otherwise the line/blip is shown in gray (and the line is thinner). As the user moves the mouse over the timeline, the name of the topic, the user, and the time being examined are shown in the upper left corner of the window; the user can scroll back through as much as one week of activity. Other functions of the timeline may be invoked by right-clicking on another user's row (e.g., private chats).

The timeline in Figure 2 covers about half a day's worth of activity. We can see that over the course of the afternoon about 20 people have logged onto Babble (shown by the number of rows), most of them have spent some time in the current conversation (shown by the color/increased thickness of the lines), and many, but not all, have "spoken" (shown by the blips). Gaps in the line indicate intervals when the person logged off. In the center of the timeline, a flurry of concentrated activity can be seen. This represents an on-line brainstorming session that took place in midafternoon, involving a majority of the people who logged onto Babble that day. Since this view of the timeline is from the Commons Area, we can see that the brainstorming session started out with people arriving and "hanging out" (i.e., not necessarily saying very much) in the Commons Area (as shown by the multiple colored lines), followed by a lot of interaction in topics that the group created for more focused brainstorming by subgroups (as shown by the colored lines changing to gray as people switch to the focused topics). Note that this synchronous use of Babble occurred in conjunction with a conference call (not shown in the timeline visualization), so the lack of activity in the Commons Area may have been due to the simultaneous conference call, followed by a flurry of brainstorming activity in specific topics just after the phone call ended. After the synchronous interaction, the timeline shows "listeners" entering the various topics and spending time there. We can infer that these visitors are reading through conversations they may have missed.

The timeline thus can reveal that others in the knowledge community have been paying attention to conversations (i.e., listening), even if they do not post a comment (i.e., speak). Although user interviews we have conducted suggest that the timeline is not routinely used by most Babble users, it has been used by more sophisticated users, for example, the hosts of on-line group interaction, to get a sense of how well the group is functioning and who is participating over time.[62] This portrayal of on-line participation that includes both speakers and listeners is meant to increase the amount of social feedback available in the environment.

Usage experiences with Babble. Having described the basic Babble functionality, we now turn to describing some of our experiences in using Babble ourselves, and in watching others use Babble. While one must be wary about drawing conclusions concerning the usability of software when its developers use it, our aim here is to simply provide a sense for how Babble is actually used by a group and to give some examples of how Babble functions as a knowledge community.

Our group has used Babble for the almost four years of its existence. The group consists of a software development group 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. The variance is due in part to the ebb and flow of people characteristic of groups in large organizations, and in part to current members inviting "associates"—colleagues with whom they had strong social or professional ties—to join.

Over the last several years we have deployed Babble to a number of other work groups: about 15 groups within IBM and one group formed by a university class outside IBM. We have studied the deployments in a variety of ways, ranging from detailed ethnographic studies—see Reference 63 for a study of six IBM Babble groups—to studies based on surveys and analysis of log data and conversation archives.[64]

We have had mixed experiences with the adoption of Babble. If we consider a Babble deployment successful when it is used more or less daily, by several people, for more than six weeks, then we can say that about half of our Babble deployments have met with success. Currently we have eight Babble groups running, not including our own. Five of these have passed the six-week mark (some are far past it), with four showing continued robust daily activity, and one in decline. In one of the recent deployments we have seen Babble support a time-limited, specific group

exercise—a month-long, global brainstorming exercise (the source of the data in Figure 2).

It is evident that when Babble catches on in a group, it supports a variety of communication purposes and practices, often similar to those we have observed in our own usage. Here we describe three social phenomena that are most relevant to the knowledge community vision.

One social practice we have observed is "waylaying," 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 by greeting the person in a public conversation, via private chat, telephone, or other means. Because the movement of the marble occurs when the user has just begun interacting with the system, it indicates an opportune moment for contact (since the user's attention has just shifted to communication with the group). Waylaying is used for purposes ranging from asking questions to initiating casual social chat.

Babble also supports group awareness through the persistence of its conversation. For example, when members of a Babble group 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 is 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 another participant is beginning to write a paper.

In addition to the persistence of conversation, 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 had not posted recently had been on line; and using the timeline to get a sense for the activity of the community as a whole. One user wrote: "It's a little like reading an electrocardiogram, the heartbeat of the community. I noticed that I missed [Susan] by an hour on Monday morning . . . . [Daphney] comes in every so often as a blip. [Frasier] jumps from space to space . . . ."

Another phenomenon that can be observed in Babble 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 on-line nick-

Babble provides an environment where social capital can be built.

name (e.g., in some Babble groups, users append "@mylocation" after their name), the types of on-line conversations created (e.g., some Babble groups have categories for "personal places" or "offices"), and naming conventions (e.g., one Babble group uses the term "chitchat" 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).

Waylay, group awareness, and the development of social norms are examples of the effect of social translucence: they are made possible by the mutual visibility and awareness of Babble participants and their activity. These social practices help to forge an identity for the group and allow individuals to become more fully dimensional as communicative partners: that is, they emerge not just as disembodied words on the screen, but as real people who might be liked or disliked, trusted or treated with caution, with reputations that can grow or be tarnished, and so forth. The fact that these effects emerge from long-running, day-to-day, work-related interactions in Babble is also important. As Cohen and Prusak[65] state, "Social capital is mainly created and strengthened (and sometimes damaged) in the context of real work. The conditions and durable connections that we experience day after day have vastly more influence on it than special events and team-building exercises."

In our experience, Babble provides an environment where social capital can be built. The following comments were drawn from a Babble group whose membership is composed of a worldwide cross-section of people in IBM and Lotus interested in on-line communities. In this group, participant comments indicate that the social interaction that occurs within the environment is valuable. Several of the participants

see the lightweight conversation that is possible in Babble as important in building social capital. One participant points to the role of ongoing, chat-like conversations in establishing what Clark and Brennan[66] have referred to as common ground:

> [Lightweight conversations] *can have more value than is immediately apparent. This (rather than with technical discussions) is often where personalities are exposed. That can make a big difference over time in feeling comfortable about asking for or offering opinions and help. Rapid exchanges often make all the difference in building mutual understanding.*

Another participant feels that informal interaction helps to build trust among remote collaborators:

> *In today's world . . . you'd want threaded discussions . . . and also have a chat space that would provide for real-time dialog, not necessarily staying on a particular topic, but a way to build trust, establish deeper relationships, a way to complement what you're trying to address over in the threaded dialog space. It's needed in the widely distributed, no-travel, matrix managed environment that we have today.*

Finally, another participant describes a deepening of relationships with colleagues through the daily interactions on Babble:

> *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.*

These remarks confirm that the informal interaction in Babble, and the blend of social and work talk, contribute to the formation and maintenance of a social fabric that underlies collaboration with distant colleagues. Through our work on Babble, we have begun to create an infrastructure that can support rich forms of social interaction. We have found that social proxies are a promising development, and we 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.

We now turn to our second example of a socially informed approach to knowledge management.

Knowledge socialization: Using stories to support knowledge creation, sharing, and reuse. We chose the term "knowledge socialization" to describe our work for several reasons. First, as discussed earlier, knowledge is heavily influenced by social and cultural factors: it is entwined, on the one hand, with human cognition, and, on the other, with the social context of teams, organizations, and communities. Second, the term "knowledge socialization" is meant to stand in contrast to the many approaches to knowledge management that take a particular technology or family of technologies as a starting point. Third, it connotes a holistic growth of knowledge through a complex, emergent system of richly interconnected processes (somewhat akin, metaphorically, to the growth of a snow crystal). In contrast, we believe the production line metaphor[67] of knowledge being created, then captured, then disseminated and then internalized can be quite misleading as an overall scheme for knowledge management. While some methods and technologies may legitimately focus on providing support for one of these processes, our work has focused on stories and storytelling as an exemplary holistic knowledge socialization process. We claim that storytelling is useful in creating, capturing, disseminating, *and* internalizing knowledge and that it accomplishes all of these simultaneously, not sequentially. Storytelling is also a representative knowledge socialization process in that it typically includes both instrumental and expressive aspects. In this section, we expand on the role of storytelling as a process that can be used throughout an organization and report on some preliminary tools to support storytelling. We end by claiming that an understanding of story as a knowledge socialization process is necessary for a deeper understanding of the social aspects of knowledge regardless of whether knowledge is explicitly presented as story.

There are many uses of story and storytelling in business. Stories can be useful ways for a business to find out about the needs of its customers in a deeper way. Stories can also help advertise a product or service; they help in showing the proper context for the use of a product or service and allow us to see the benefits. Stories can be used as educational materials within a company. Stories can be used as a tool in the design process.[68,69] Informally, colocated communities of practice may spontaneously share experiences in the form of oral stories.[26] Wider, more distributed communities of practice may share stories

in electronic forms; e.g., stories, as well as other kinds of knowledge, were shared successfully via the IBM VM (virtual machine) forums. Even more formally, stories may be collected and arranged into scenario-based learning systems.[18] Stories can be used to help establish or change corporate culture. Scenarios, a variant of stories, can be used to help organize the design process and keep it focused on real customer needs. Scenarios can be used for strategic planning. Scenarios can also be a useful way for team members from different functions to see how they can relate to solve a problem.

As instantiations of a type of knowledge that can be used in so many business processes, stories also have the advantage that they can help knowledge flow through the organization. Not only are stories capable of being used by many different business functions (marketing, design, management), they are also capable of being understood by various professions. Thus, stories can serve not only to support communities of practice with a common vocabulary; they can also serve an important coordinating role within a team whose members come from different communities of practice. A story might start with a customer expressing a need, be used as a scenario during design of a service to meet that need, and then be included as part of a marketing campaign to show how that need can be met. A different story might encapsulate the experiences of a consultant to the petrochemical industry. This story might seem to offer a lesson learned that is at odds with the experiences of another consultant. By comparing the stories and examining the apparent contradiction, the two consultants themselves (or even a third party) could find the differentiating factors between the two situations and, in effect, use the story combination to create new knowledge.

Although stories have the capability of serving as a kind of cognitive glue across the many functions and levels of a large organization, there is no guarantee that they will do so. Either the formal organization or the corporate culture may introduce roadblocks of various kinds to the use of stories. Perhaps the organization does not reward people for sharing their experiences. Or, even if the formal organization puts in explicit rewards for sharing experiences, the informal corporate culture may discourage people in various ways; stories may be seen as a kind of second-class knowledge compared to an algorithm or a formula; or stories from the marketing department may be seen as suspect by people in the engineering department. In the latter case, the potential flow of rich information about users and their context that could serve as a competitive differentiator for the company is blocked; instead people in the engineering department design products based on their own traditions or biases.

Even where the organizational and cultural factors are favorable to the use of stories, however, there may well be technological barriers. Stories are quite a natural way for small groups of trusted colleagues to exchange information. Scaling such a process to a large, global organization requires an integrated set of story tools such as we are developing at IBM Research. Dave Snowden, an IBM colleague working with stories, uses an apt analogy to explain why it may now be necessary to take storytelling to the next level. When the modern Olympics began in 1896, a natural athlete who trained hard had a good chance at winning a gold medal. At the end of this century, that is no longer true. Only a good athlete who trains hard and trains scientifically, with expert advice in nutrition, biomechanics, sports medicine, and other fields has a chance at a gold medal.

In the past, great leaders in business have instinctively told stories to help motivate people and to create an organizational reality. Workers have also shared knowledge by telling stories in small, face-to-face groups. But today, we live in a world at once faster paced, more competitive, and more global. Science and technology might now be used to make stories and storytelling more effective, more appropriate, more scalable to large organizations.

An integrated suite of story-related tools. In order to provide a common underpinning for the various story-related tools that we have developed, we have proposed a first pass at a "StoryML," that is, a markup language specifically geared toward stories. The initial representation is based on a distillation of many different approaches to story.[52,70–80] Our initial formulation has three different but related "views" of story: Story Form (what is in the story); Story Function (what are the purposes of the story); and Story Trace (what is the history of the story). In turn, the Story Form can be broken down into dimensions of environment, character, plot, and narrative. The idea of the StoryML is that it is expandable according to purpose. For some purposes, the user (e.g., a student studying mystery plots) may be satisfied with minimal detail concerning function and trace but may need to expand certain aspects of the Story Form in great detail. In another context, a different user (e.g., a historian comparing certain

themes across time and cultures) might have a very high-level view of Story Form and Story Function but want to see a detailed description of Story Trace. At this point, the meta-data in StoryML must be supplied by a knowledgeable human being. However, increasingly, it could become feasible to partially automate this process.

The following scenario illustrates how a StoryML might support reasoning about a business process.

> Jane is under the gun to cut costs in the fulfillment process without increasing delivery time or decreasing customer satisfaction. In fact, her boss, Betty, has strongly suggested that rethinking the fulfillment process *should* allow her to *decrease* delivery time and *increase* customer satisfaction. Jane's knowledge portal is already personalized to her general profile mostly via a dynamic background process that takes note of what Web sites she visits, what the topics of her e-mail are, and with whom she communicates. She can turn a software "dial" on any given knowledge scan that determines how much the scan will be influenced by her general profile and how much by the specific search terms she uses. In this case, Jane wants to see an overall story frequency map. Since stories generally arise when things do not go as planned, the story "hot spots" show her likely places where current fulfillment processes are probably inefficient or subject to breakdown. At this point, she is not very concerned with the structure of the story or even the function. She is mainly concerned with the Story Trace. Jane focuses on the two most likely trouble areas and sets up two separate story-exchange meetings of people expert in these two areas.
>
> The story exchange meetings only last an hour each and the stories exchanged are all digitally recorded with an associated (imperfect) transcript. Although the transcripts are imperfect, they serve adequately to allow her to zoom in accurately on audio versions of some very telling stories. These are referenced throughout the subsequent process re-engineering. Jane quickly assembles another team to explore possible ways to improve the fulfillment process. In this creative synectics[35] session, along with other techniques, she again accesses a corporate-wide story base, but this time, she is primarily concerned with Story Function. She is looking for stories that help people think about things in new ways and break old thought habits. From a host of potentially useful ideas, she and her team pick out a few high-leverage, quickly implementable, and practical ideas to pursue.
>
> In order to concretize these ideas for dissemination and also to double-check on their practicality, Jane develops some scenarios for how fulfillment will be done under the new process. Before committing further resources, she uses these scenarios to get feedback from a small but varied set of customers. These customers provide some additional insights and requirements.
>
> In parallel with the development of an improved process, training materials are produced to explain the new process as well as the design rationale behind it. In this case, a story creation tool (which incorporates examples and guidelines) focuses on Story Form. The materials make it quite clear how to use the new process and also explain why. In addition, related stories are created for marketing materials stressing to the customers how it is even more desirable now to do business with Jane's corporation.

In the scenario above, we showed how stories can be used in many ways. In each case, however, we were describing what might be termed "endogenous stories." That is, the "complete story" in some sense was captured and contained in some explicit record. But the potential use of stories extends considerably beyond "endogenous stories."

Suppose that Albert Einstein writes an equation such as "Energy = mass times speed-of-light squared." This is clearly not a story, per se. It is, after all, an equation. And, yet, in the larger sense, stories emanate in every direction from this equation. How did Einstein come up with this? How did it lead to the atom bomb? These might be termed "exogenous stories"—stories created around knowledge.

In order to take natural language processing to the next level, it will be necessary to understand such "exogenous stories." We will need to understand agents, goals, obstacles, communicative strategies, and intentions. Otherwise, it will not, in general, be possible to understand the import of even such a simple statement as "Alice left the party early." Who is making this statement, and to whom? What do they intend to communicate? Is it truthful or a lie? Only by understanding the larger "exogenous story" can we possibly know the function of this statement. Such functional variations in context can easily project into the semantic and even the syntactic domain as the

well-known example "Time flies like an arrow" illustrates. We cannot parse this sentence nor assign lexical items to all the surface tokens without understanding the exogenous story of which this sentence is a part.

Indeed, understanding the "exogenous story" of which individual statements partake, extends beyond what is typically termed "natural language processing." Current attempts to incorporate "intelligent agents" into various systems often lead to problems, even in something as simple as automatic spelling and capitalization correction. The reason is essentially that the system has no knowledge of the user's current context and intention. Therefore, a particular "correction" offered by an agent may or may not make any sense in the current context. As an example, in the immediately preceding sentence, Microsoft Word** caused a menu to pop up over both instances of the token "may," inviting me to substitute "May 6, 2001," which happens to be the current date. In order for computer systems to be more than passive conduits for human knowledge, we will need to develop knowledge representations that can account for and represent the essential elements of stories in terms of their form, their function, and their history. This is not to say that all knowledge is in the form of story; we claim only that once we understand and can represent story (StoryML is an initial proposal in this direction), we will have the concepts necessary to produce true knowledge-based systems. Until we build such representations, "intelligent agents" will as often constitute an amusing (or frustrating) distraction as a collaborative knowledge partner. Without such representations, so-called knowledge-based systems will not be capable "social actors," although it is possible that they will be temporarily perceived as such. Since the social aspects of knowledge management constitute an absolutely critical aspect of the general problem of knowledge management, until we can understand and represent story, we will not have the tools to build the underlying architecture for a knowledge-based system in which humans and computers can effectively collaborate.

Summary and conclusion

The simple picture of knowledge management as getting the right information to the right people at the right time is wrong. Knowledge management is not just a matter of managing information. It is, we have argued, deeply social in nature, and must be approached by taking human and social factors into account. We have provided some extra pieces for the knowledge management puzzle and demonstrated how we have used them in our own work to assemble some knowledge management systems that are strongly shaped by human and social factors. As the field of knowledge management develops, and more widespread and varied experience with different approaches to KM is gained, we believe not only that additional critical pieces of the KM puzzle will be revealed, but that it will become clearer how all the pieces fit together to create a rich picture of social and intellectual capital within organizations. Certainly, looking toward the future of work, as it becomes more centered in virtual relationships and spaces both within and across organizations, creating and maintaining knowledge and its social context will only become more vital.

We believe that one of the most important aspects of a knowledge management system is that it becomes what we have termed a *knowledge community*: a place within which people discover, use, and manipulate knowledge, and can encounter and interact with others who are doing likewise.[46,61] We have talked about two approaches for supporting knowledge communities, namely social computing and knowledge socialization. A fundamental characteristic of a knowledge community is that it includes conversation and other forms of narrative, for example stories, and/or unguarded discussion among people who know one another, who share professional interests, and who understand the contexts within which their remarks are being made. We have outlined a variety of specific techniques that can contribute to a realistic and effective approach to knowledge management, including supporting new forms of group interaction (e.g., Bohm Dialogue, stories), methods for enhancing creativity (e.g., the use of metaphor), and support for expressive communication. When such techniques are incorporated into knowledge communities, they result in organizational opportunities to build social capital, including trust and cooperation among colleagues. This notion of a knowledge management environment as a "trusted place" is an interesting and challenging one for system designers and for organizations. How—technically, socially, and organizationally—can we balance the need for a safe and trusting place, within which so much knowledge creation and social capital building takes place, with the organizational imperative to share information more broadly? We believe that a greater understanding of how to design socially translucent systems that permit social mechanisms to come into play will help developers of technological systems to negotiate such issues. Similarly,

we believe that understanding better how to socialize knowledge through techniques such as storytelling and scenarios will offer organizations greater mastery and scope in creating, sharing, and reusing the knowledge that is critical to survival in the twenty-first century.

Acknowledgments

This work is highly collaborative, and we owe a great debt to our colleagues, past and present. Thanks in particular to Andrew Gordon, Christine Halverson, Cynthia Kurtz, Mark Laff, Peter Malkin, Tracee Wolf, Mark Ackerman, John Richards, Rachel Bellamy, Cal Swart, David N. Smith, Erin Bradner, Jason Ellis, Beth Veinott, Jason Elliot, James Reed, James Hudson, Karrie Karahalios, Barbara Kelly, and Brent Hailpern. We also thank three anonymous reviewers for their comments on an earlier version of this paper.

Cited references and notes

1. F. Capra, *The Web of Life: A New Understanding of Living Systems*, Doubleday, New York (1997).
2. M. Wheatley, *Leadership and the New Science*, Berrett-Koehler, San Francisco, CA (1993).
3. G. Lakoff and M. Johnson, *Metaphors We Live By*, University of Chicago Press, Chicago, IL (1983).
4. R. Sternberg, *Defying the Crowd: Cultivating Creativity in a Culture of Conformity*, Free Press, New York (1995).
5. R. Sternberg, *Advances in the Study of Human Intelligence*, Erlbaum, New York (1986).
6. J. Guilford, "Intellectual Resources and Their Values as Seen by Scientists," *Scientific Creativity: Its Recognition and Development*, C. Taylor and F. Barron, Editors, John Wiley & Sons, Inc., New York (1963).
7. T. G. Evans, "A Heuristic Program to Solve Geometric Analogy Problems," *Semantic Information Processing*, M. Minsky, Editor, MIT Press, Cambridge, MA (1968). Based on Ph.D. dissertation, Department of Electrical Engineering, MIT, 1963.
8. J. Darley and B. Latane, "Bystander Intervention in Emergencies: Diffusion of Responsibility," *Journal of Personality and Social Psychology* **8**, 377–383 (1968).
9. T. DeMarco and T. Lister, *Peopleware: Productive Projects and Teams*, Dorset House, New York (1987).
10. J. C. Collins and J. I. Porras, *Built to Last: Successful Habits of Visionary Companies*, Harper, New York (1997).
11. D. Cohen and L. Prusak, *In Good Company: How Social Capital Makes Organizations Work*, Harvard Business School Press, Boston, MA (2001).
12. J. S. Brown and P. Duguid, *The Social Life of Documents. Release 1.0*, EDventure Holdings Inc., New York (October 11, 1995), pp. 1–18. Also available at www.parc.xerox.com/ops/members/brown/papers/sociallife.html.
13. J. S. Brown and P. Duguid, *The Social Life of Information*, Harvard Business School Press, Boston, MA (2000).
14. J. C. Thomas, "A Design-Interpretation Analysis of Natural English," *International Journal of Man-Machine Studies* **10**, 651–668 (1978).
15. J. Dewey, *How We Think: A Restatement of the Relation of Reflective Thinking to the Educational Process*, Heath, Boston, MA (1933).
16. L. S. Vygotsky, *Thought and Language*, MIT Press, Cambridge, MA (1962).
17. J. Piaget and B. Inhelder, *The Psychology of the Child*, Basic Books, New York (1969).
18. R. Schank, *Virtual Learning: A Revolutionary Approach to Building a Highly Skilled Workforce*, McGraw-Hill, New York (1997).
19. G. Miller, "The Magical Number Seven Plus or Minus Two: Some Limits on Our Capacity for Processing Information," *Psychological Review* **63**, 81–97 (1956).
20. R. Stuart, *The Design of Virtual Environments*, McGraw-Hill, New York (1996).
21. *Frontiers of Human-Centered Computing, Online Communities, and Virtual Environments*, R. Earnshaw, R. Guedj, A. van Dam, and J. Vince, Editors, Springer-Verlag, London (2001).
22. W. Gaver, "Synthesizing Auditory Icons," S. Ashlund, K. Mullet, A. Henderson, E. Hollnagel, and T. White, Editors, *Proceedings of INTERCHI'93*, ACM, New York (1993), pp. 228–235.
23. S. Ford, J. Forlizzi, and S. Ishizaki, "Kinetic Typography: Time-Based Issues in the Presentation of Text," *Proceedings of the Conference on Human Factors in Computing Systems* (CHI'97), *Extended Abstracts*, ACM, New York (1997). See also http://www.maedastudio.com.
24. J. Lave and E. Wenger, *Situated Learning: Legitimate Peripheral Participation*, Cambridge University Press, Cambridge, UK (1991).
25. E. Wenger, *Communities of Practice: Learning, Meaning, and Identity*, Cambridge University Press, Cambridge, UK (1998).
26. J. Orr, *Talking About Machines: An Ethnography of a Modern Job*, Cornell University Press, Ithaca, NY (1996).
27. G. Olson and J. Olson, "Distance Matters," *Human-Computer Interaction* **15**, Nos. 2–3, 107–137 (2000).
28. H. H. Clark, *Using Language*, Cambridge University Press, Cambridge, UK (1996).
29. W. Orlikowski, "Learning from Notes: Organizational Issues in Groupware Implementation," *Proceedings of the Conference on Computer Supported Cooperative Work*, ACM, New York (1992), pp. 362–369.
30. E. F. Churchill and S. Bly, "Virtual Environments at Work: Ongoing Use of MUDs in the Workplace," *Proceedings of International Joint Conference on Work Activities Coordination and Collaboration* (WACC'99), San Francisco, CA, ACM Press, New York (1999).
31. A. Huxor, "An Active Worlds Interface to BSCW to Enhance Chance Encounters," D. Snowdon and E. F. Churchill, Editors, *Proceedings of CVE'98* (Collaborative Virtual Environments), Manchester, UK (June 1998).
32. B. Hillier, *Space Is the Machine*, Cambridge University Press, Cambridge, UK (1996).
33. B. Shneiderman, "Creating Creativity: User Interfaces for Supporting Innovation," *ACM Transactions on Computer-Human Interaction* **7**, No. 1, 114–138 (2000).
34. D. Bohm, *On Dialogue*, David Bohm Seminars, Ojai, CA (1990).

35. W. Gordon, *Synectics*, Harper, New York (1961).
36. A. S. Gordon, "The Representational Requirements of Strategic Planning," Fifth Symposium on Logical Formalizations of Commonsense Reasoning, New York University (May 20–22, 2001), http://www.cs.nyu.edu/faculty/davise/commonsense01/papers.html.
37. J. C. Thomas, "Narrative Technologies for the New Millennium," *Knowledge Management* **2**, No. 9, 14–17 (1999).
38. W. Isaacs, *Dialogue and the Art of Thinking Together*, Doubleday, New York (1999).
39. E. Arias, H. Eden, G. Fischer, A. Gorfman, and E. Scharff, "Transcending the Individual Human Mind—Creating Shared Understanding Through Collaborative Design," *ACM Transactions on Computer-Human Interaction* **7**, No. 1, 84–113 (2000).
40. B. Mattimore, *99% Inspiration: Tips, Tales, and Techniques for Liberating Your Business Creativity*, Amacon, New York (1993).
41. J. C. Thomas, D. Lyon, and L. A. Miller, "Aids for Problem Solving," *Research Report RC-6468*, IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 (1977).
42. D. Snowden, "The Paradox of Story," *Journal of Scenario and Strategy Planning* **1**, No. 5 (Dec. 1999/Jan. 2000).
43. D. Snowden, "Storytelling: An Old Skill in a New Context," *Business Information Review* **16**, No. 1, 30–37 (1999).
44. A. S. Gordon, "Accessing Stories with Activities," E. Amitay, Editor, *Textual Display in Search*, MIT Press, Cambridge, MA (in press).
45. D. Lawrence and J. Thomas, "Social Dynamics of Storytelling: Implications for Story-Based Design," presented at AAAI workshop on Narrative Intelligence, N. Falmouth, MA (November, 1999).
46. T. Erickson and W. A. Kellogg, "Social Translucence: An Approach to Designing Systems That Support Social Processes," *ACM Transactions on Computer-Human Interaction* **7**, No. 1, 59–83 (2000).
47. J. Harris and A. Henderson, "A Better Mythology for System Design," *Proceedings of the Conference on Human Factors in Computing Systems* (CHI'99), ACM Press, New York (1999), pp. 88–95.
48. H. Beyer and K. Holtzblatt, *Contextual Design: Defining Customer-Centric Systems*, Morgan Kaufmann, San Francisco, CA (1998).
49. A. P. De Geus, *The Living Company*, Harvard Business School Press, Boston, MA (1997).
50. R. D. Putnam, R. Leonardi, and R. Nanetti, *Making Democracy Work: Civic Traditions in Modern Italy*, Princeton University Press, Princeton, NJ (1993).
51. Aristotle, *Poetics*, M. Heath, Translator, Penguin, New York (1997).
52. R. McKee, *Story*, Harper & Row, New York (1997).
53. J. C. Thomas, "Studies in Office Systems I: The Effect of Communication Medium on Person Perception," *Office Systems Research Journal* **1**, No. 2, 75–88 (1983).
54. D. Armstrong, *Managing by Storying Around*, Doubleday, New York (1992).
55. H. Gardner, *Leading Minds: An Anatomy of Leadership*, Basic Books, New York (1996).
56. T. Erickson, D. Smith, W. A. Kellogg, M. Laff, J. Richards, and E. Bradner, "Socially Translucent Systems: Social Proxies, Persistent Conversation and the Design of Babble," *Proceedings of the Conference on Human Factors in Computing Systems* (CHI'99), ACM Press, New York (1999), pp. 72–79.
57. E. Goffman, *Behavior in Public Places: Notes on the Social Organization of Gatherings*, Macmillan Publishing Co., New York (1963).
58. E. Goffman, *Interaction Ritual*, Anchor Books, New York (1967).
59. A. Kendon, *Conducting Interaction: Patterns of Behavior in Focused Encounters*, Cambridge University Press, Cambridge, UK (1990).
60. See *Social Navigation of Information Space (Computer Supported Cooperative Work)*, A. J. Munro, K. Hook, and D. Benyon, Editors, Springer-Verlag (1999). "Social navigation is a vibrant new field that examines how we navigate information spaces in 'real' and 'virtual' environments, how we orient and guide ourselves, and how we interact with and use others to find our way in information spaces."
61. T. Erickson and W. A. Kellogg, "Knowledge Communities: Online Environments for Supporting Knowledge Management in its Social Context," M. Ackerman, V. Pipek, and V. Wulf, Editors, *Beyond Knowledge Management: Sharing Expertise*, MIT Press, Cambridge, MA (in press).
62. T. Erickson and M. R. Laff, "The Design of the 'Babble' Timeline: A Social Proxy for Visualizing Group Activity over Time," *The Proceedings of CHI 2001, Extended Abstracts*, ACM Press, New York (2001).
63. E. Bradner, W. A. Kellogg, and T. Erickson, "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), Kluwer Academic Publishers, Norwell, MA (1999), pp. 139–158.
64. T. Erickson, "Making Sense of Computer-Mediated Communication (CMC): Conversations as Genres, CMC Systems as Genre Ecologies," J. F. Nunamaker, Jr., and R. H. Sprague, Jr., Editors, *Proceedings of the Thirty-Third Hawaii International Conference on Systems Science*, IEEE Press, New York (January 2000).
65. See Reference 11, p. 22.
66. H. H. Clark and S. E. Brennan, "Grounding in Communication," L. Resnick, J. M. Levine, and S. D. Teasley, Editors, *Perspectives on Socially Shared Cognition*, APA, Washington, DC (1991), pp. 127–149.
67. K. Nonaka and H. Takeuchi, *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation*, Oxford University Press, New York (1995).
68. T. Erickson, "Notes on Design Practice: Stories and Prototypes as Catalysts for Communication," J. Carroll, Editor, *Scenario-Based Design: Envisioning Work and Technology in System Development*, John Wiley & Sons, Inc., New York (1995).
69. T. Erickson, "Design as Storytelling," *interactions*, ACM Press, New York (July/August 1996).
70. M. Bakhtin, *The Dialogic Imagination*, C. Emerson and M. Holquist, Translators, University of Texas Press, Austin, TX (1981).
71. M. Bal, *Narratology: Introduction to the Theory of Narrative*, University of Toronto Press, Toronto (1992).
72. D. Boje, "The Storytelling Organization: A Study of Story Performance in an Office-Supply Firm," *Administrative Science Quarterly* **36**, 106–121 (1991).
73. E. Brannigan, *Narrative Comprehension and Film*, Routledge, New York (1992).
74. A. Dundes, *The Study of Folklore*, Prentice-Hall, Inc., Englewood Cliffs, NJ (1963).
75. W. Labov and J. Waletzsky, "Narrative Analysis: Oral Version of Personal Experience," J. Helm, Editor, *Essays on the Verbal and Visual Arts*, University of Washington Press, Seattle, WA (1967), pp. 12–44.

76. N. Norrick, "Twice-Told Tales: Collaborative Narration of Familiar Stories," *Language in Society* **26**, 199–220 (1997).
77. E. Ochs, C. Taylor, D. Rudolph, and R. Smith, "Storytelling as a Theory-Building Activity," *Discourse Processes* **15**, 32–72 (January–March 1992).
78. G. Polti, *The Thirty-Six Dramatic Situations*, The Writer, Inc., Boston, MA (1916).
79. R. Schank, *Tell Me a Story: Narrative and Intelligence*, Northwestern University Press, Evanston, IL (1990).
80. P. Underwood, *Three Strands in the Braid: A Guide for Learning Enablers*, A Tribe of Two Press, San Anselmo, CA (1999).

Accepted for publication August 6, 2001.

John C. Thomas *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 704, Yorktown Heights, New York 10598 (electronic mail: jcthomas@us.ibm.com).* Dr. Thomas (www.truthtable.com) is manager of knowledge socialization at the Thomas J. Watson Research Center. He received his Ph.D. degree in experimental psychology in 1971 from the University of Michigan. He has worked in a number of areas of human-computer interaction including speech systems, query systems, natural language systems, and design problem solving. In 1986, he founded the Artificial Intelligence Laboratory at NYNEX where work was carried out in expert systems, human computer interaction, machine vision, robotics, and computer-aided instruction. Dr. Thomas rejoined IBM Research in 1998 to work in the area of knowledge management (www.research.ibm.com/knowsoc/).

Wendy A. Kellogg *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 704, Yorktown Heights, New York 10598 (electronic mail: wkellogg@us.ibm.com).* Dr. Kellogg is manager of social computing at the Thomas J. Watson Research Center. Her current work involves designing and studying socially translucent systems for computer-mediated collaboration in groups and organizations. Dr. Kellogg's work in human-computer interaction over the last 15 years has spanned research areas including HCI theory, evaluation methods, design, and development. She is an author of numerous papers and was coeditor in 2000 of a millennial special issue of *Human-Computer Interaction* entitled "New Agendas for Human-Computer Interaction." Dr. Kellogg holds a Ph.D. degree in cognitive psychology from the University of Oregon.

Thomas Erickson *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 704, Yorktown Heights, New York 10598 (electronic mail: snowfall@us.ibm.com).* Mr. Erickson (http://www.pliant.org/personal/Tom_Erickson) is an interaction designer and researcher whose approach to systems design is shaped by work in sociology, rhetoric, architecture, and urban design. He has contributed to the design of a number of products and authored about 40 publications on topics ranging from personal electronic notebooks to pattern languages and virtual community. Originally trained as a cognitive psychologist at the University of California, San Diego, he spent five years at a startup company, nine years at Apple Research, and finally joined IBM in 1997 as a research staff member. His primary agenda is studying and designing systems that support network mediated group interaction.