

To appear in K. A. Renninger and W. Shumar (Eds.), Building virtual communities: Learning and change in cyberspace. NY: Cambridge University Press.

Finding the ties that bind: Tools in support of a knowledge-building community

Christopher M. Hoadley, Roy D. Pea

Center for Technology in Learning

SRI International

Abstract: Traditionally, collaborative technologies are intended to directly support joint, collaborative activity, taking their cues from communication and media. Here, empirical findings are presented about the types of information needs associated with the formation of a knowledge-building community among professional learning technology researchers. Several issues are outlined in designing, facilitating, supporting, and measuring knowledge-building activity in such a community of practice. It is argued that, rather than communication tools, a knowledge-building community is better served by knowledge-networking tools that support individual information needs (both social and topical) relevant to participating in the community.

Introduction

Finding a professional connection with a colleague seems like a simple task but can devour hours of time. An anecdote illustrates why this is hard. A researcher whom we'll call David got a call with a question about research on interactive toys. David had some experience in that area and immediately recalled several people who did similar work, but who didn't quite fit the bill of this request. He vaguely remembered someone he had heard about who did do that sort of work—the researcher was a Canadian woman who had recently won an award for women in computer science. He thought but wasn't sure that the woman was from Western Canada. With these recollections in mind, he set about trying to find her.

First, he tried searching based on the topic. He began with a Web search on the topic area but found far too many results. He tried narrowing his search but had no luck. He tried a number of refinements, including searching on words related to the award, and so on. After spending nearly half an hour, he decided to try a different strategy.

This time, David tried to find the researcher through his social network. He began by asking a coworker down the hall. A short conversation didn't yield any leads. Continuing down the hall, he asked another colleague. Again, the colleague didn't know the person he was seeking, but this person did suggest another related researcher who might know the mystery woman's identity. David knew that the related person (let's call her Renee) worked in Los Angeles and had written a book that he thought had cited the mystery person. David tried to find the book. When a quick glance through his own library didn't yield a copy, he tried to look the book up

on the Internet through searching, this time armed with an author, institution, and an approximate title. Ten to 20 minutes later, not having found the book, he moved to searching for the author's home page, hoping for a link to the publication, a phone number so he could call Renee, or a link to the home page of the Canadian researcher. Again, quick searches yielded no results. Not finding Renee's home page through a search engine, David tried a less direct approach. He started at the home page of the Los Angeles university in an attempt to drill down to Renee's home page directly. Lost in the vast Web site of the institution, he eventually aborted this attempt.

After pursuing a number of dead-end search strategies, he gave up on Renee entirely. He finally did discover the mystery researcher by a brute-force search, starting with the home pages of several universities in Western Canada and eventually stumbling on the right person by sifting through a number of computer science department Web pages. This search odyssey lasted hours before David finally reached his goal.

This example is important for two reasons. First, it demonstrates the high cost of finding and making connections to people. In this case, finding the collaborator took much longer than the collaboration, which consisted of a brief conversation and skimming one of the Canadian's articles. Second, it shows how social context is interwoven with finding information. David wanted some information, but he did not search for information in the traditional library sense. Rather, he searched for a person that he knew could help him. The specific information he sought was impossible to find directly, so he had to find its author, the mystery woman in Western Canada. As Harold "Doc" Edgerton, the inventor of the strobe light and one of the century's most prominent engineers, once explained, when he wanted to find something out, first he would ask around to see whether anybody knew the answer, then he would try it out in the lab himself, and only then would he try looking the information up in a book or library (Edgerton, personal communication, 1989). The social connection to knowledge is often the most expedient.

Even when finding information through a social network may be the best way, it is by no means an easy way. Finding this woman was difficult. The topic of interactive toys wasn't really helpful in locating her, but seemingly irrelevant contextual information was—her gender, geographic location, and an award she'd won. The Internet's vast information did contain exactly what David needed—contact information for the researcher, her profile, even some of her work. But traditional search engines did nothing to help connect her to David.

This chapter describes how we came to use technology in support of pre-collaboration activities like finding social and topical information, instead of the more traditional role of supporting communication during a collaboration. Our problem is an example of the more general problem of *knowledge networking*: how to get knowledge to those members of a community who need it. Often, when considering what collaborators need, we think of technology to directly support interaction, such as fancy telecommunications systems or "shared workspaces" in the computer. David would have been served much better by a way to find the researcher than by any traditional groupware to help him talk to her.

In the following sections, we will describe some of the general aspects of the problems associated with building collaborative technologies for knowledge networking. We discuss some findings from examining the knowledge-sharing practices of a group of scholars. Finally, we describe our experiences in implementing a knowledge-networking tool with a nascent, distributed community of educational technology users, researchers, and businesspeople called CILT. CILT, the Center for Innovative Learning Technologies, is funded by the U.S. National Science Foundation to foster a productive knowledge-building community among learning technology researchers and stakeholders. We explore the development of technologies for CILT as a case study of what is involved in creating technologies to support knowledge building.

Conceptualizing a learning community

Why collaborate? Humans need to coordinate in joint action to achieve tasks larger than any one person could accomplish. In addition, we communicate to express ourselves, to transmit information, and to learn. Through the processes of acculturation, knowledge and culture are perpetuated and transformed as we interact, define new problems, and take on new challenges. People generally highlight collaboration as good and are interested in creating tools to support it. But what is good collaboration? By examining some models of knowledge and organizations, we can get insight into what types of collaboration we might want to support with technology.

Models of collaboration

In the world of business studies of organizational behavior, the processes of collective action have been simplified into a number of models. The hierarchical model of Taylor was concerned primarily with a top-down control structure in which commands propagated downward from management to labor, purportedly dividing and delegating the tasks of the organization for efficient, coordinated action (Fischer, 1999). Knowledge in this model is generally simplified to the issue of information transmission—when somebody needs to know something, you tell it to him or her. This model dominated in the early 20th century, and the collaboration technologies we have inherited from it support the goal of information transmission: telephones, radio, loudspeakers, and the ever-present photocopier all support transmission of information (see especially Pea & Gomez, 1992).

Over time, this model proved ineffective. Flatter organizational structures, team-based work groups, and information management techniques began to emerge. Organizational knowledge was highlighted as an important type of institutional capital. In this more complex model, information transmission gave way to information management. Large organizations, such as companies, developed "management information systems" (MIS) departments whose job it was to collect, process, and route information to the right people. In this model, there were two ways to bring the right knowledge to bear on a problem: one was to move the people who knew the right things, assembling project teams with ready-made expertise; the other was to codify the information needed and use information technologies to help people find what they were looking for. This model yielded our standard view of corporate training and centralized information technologies for organizations. Technologies in this case were less communicative and more data oriented. The technologies used included relational databases, automatically generated statistics and reports, and codified sources of information such as manuals, corporate training documents, and the like. A few innovative applications attempted to do automatic knowledge management by means of techniques such as data mining or automatic information capture.

However, this model also has been proven ineffective. Corporate training cannot keep up with changing skill requirements, and MIS departments have a hard time ensuring access to the right kinds of information. Furthermore, there has been a growing awareness that information does not necessarily lead to knowledge. Until information has been comprehended and interpreted to the point that it can be applied to a situation, it is not knowledge. Decentralization became *de rigueur*, and the idea of the "learning organization" (Garratt, 1987) was born. In this model, individuals are constantly learning new skills and working to discover and propagate knowledge. We define a knowledge-building community as a community with a shared goal of individual learning and knowledge transfer within the group.

How does knowledge move around in organizations? A famous study of photocopier repair technicians (Orr, 1990) demonstrated that storytelling in social, water-cooler settings was the main way expertise was being

passed around. Not only did this social network help the technicians in a community of practice uncover and transmit information, but as the technicians applied the stories in their own repair work, the knowledge in the organization increased (Brown & Duguid, 1991). Individuals were constantly transforming information into knowledge and knowledge into outcomes (in this case, repaired copiers) through their social interactions.

The copier study provides an excellent example of the kinds of reasons why learning is an important function of collaboration. The copier repair people did not really need to collaborate to coordinate their actions; a dispatch system could easily have been rigged up that did not require the repair people to talk to each other. Nor was collecting and routing information a primary benefit; the individual copier repair people primarily invented repair techniques on their service calls, and a system could have been put in place to impersonally capture that information. (Indeed, the U. S. military attempts to capture all processes and make them explicit in manuals, although these explicit process instructions rarely capture what really occurs. This approach is fundamentally flawed, in part because it ignores the constructive nature of understanding and learning. See Hutchins, 1995). No, the primary benefit of this community was how the technicians could learn from one another, increasing their knowledge, thereby enabling the company to solve more copier problems for more people in less time. Getting others to know what one person had figured out augmented the overall knowledge in the community and improved every repairperson's ability to fix copiers. Figuring out a tough copier problem was of some benefit, but sharing that knowledge with others was the real success. One term for this type of community is a "knowledge-building community" (Scardamalia & Bereiter, 1994), where individuals are committed to sharing information for the purpose of building understanding (knowledge) in all the participants. This knowledge-building activity benefits not only individuals but also groups (Pea, 1992). One example of a type of knowledge-building community that has existed for many centuries is the scholarly community, where sharing information (via publishing) and boosting overall knowledge (through teaching) are as important as the creation of new information for oneself (through research).

A fascinating project at the University of Waterloo's Electronic Library has been devoted to documenting the history of scholarly societies. These historians highlight how the sharing and publication of scholarly knowledge emerged from the scholarly societies formed starting in the 14th century. Their chronology documents how until the 19th century, scholarly societies were generally of broad scope (e.g., all sciences, all arts, or both) and geographically based. Afterwards, increasingly specialized scholarly societies came into being. They highlight the seminal publication in 1938 by Martha Ornstein of her book on the role of scientific societies in the 17th century (Ornstein & Cohn, 1938). In this work, Ornstein documents that the goal of these early societies (such as Accademia del Cimento of Florence or the Royal Society) was generally to promote research by providing a place for researchers to meet one another and discuss or even carry out research. This same knowledge-building function, and its affiliated objectives of finding people and collaborating with them, follows to the present day.

Today we find tools that provide not only access to information but access to people. Access to people includes referral technologies for help-on-demand; "customer relationship management software" that aims to help phone operators interact with customers in a consistent way; and participatory news services like the now-famous "slashdot.org" Web site, with news articles and discussion intertwined throughout the site. These technologies help us find not only information but also knowledge by connecting us to people and not only to facts.

How can technology help a learning community?

Given a model of a learning community as a community that builds knowledge in all its participants through collaboration, how do we support such a community? Certainly, many collaborations are dysfunctional and actually prevent learning through encouraging "groupthink" or by disadvantaging some participants (Linn & Burbules, 1993). Can we help make effective collaboration easier through technology?

It is often observed that a community of practice is embedded in and overlaps with other communities of practice, or CoPs (Lave & Wenger, 1991; Wenger, 1999). CoPs are diverse in nature, and, like organisms in ecological niches, they originate, evolve, and may become extinct. Individuals play membership roles in many different communities. Different kinds of communities (e.g., a theatre group, K-12 mathematics education researchers) provide different identifiable roles, thus providing diverse routes into becoming a member of a given CoP.

Members often share work, lifestyles, activities, and identity badges such as ways of speaking and clothing, and these members are interconnected in that they contribute to co-constructing what aspects of activity and choice define a sense of membership. It is an inherent part of communities of practice that members carry out what might be called tacit or indigenous assessments—gauging one another according to the perceived appropriateness of talk, activity, lifestyle, competence, commitment, and other realms of behavior or being.

Part of the process of learning within communities of practice is described by Lave and Wenger (1991) as *legitimate peripheral participation* (LPP), a relationship that individual learners have to the activities of established communities of practice when they act with the goal of increasing their sense of membership in and acceptance by these communities. LPP is "a way of gaining access to sources for understanding through growing involvement." In their development of this view, Lave and Wenger do much to explicate "the relations between newcomers and old-timers, and about activities, identities, artifacts and communities of knowledge and practice."

Lave and Wenger's analysis of CoPs indicates the importance of learning by membership. Learning energy is devoted to becoming a member of a community of practice, and what is learned is how to be a member. Consider what this might imply for knowledge-building communities: since the community is oriented toward the production and dissemination of knowledge, the process of joining the community involves learning how to become committed to these goals in a way that the community values. In short, learning how to learn is the price of entry into a knowledge-building community.

The realization that learning is an inherent property of an effective, knowledgeable organization suggests that technologies for learning and technologies for collaboration may be one and the same. The realization that learning results when people participate in a community of practice has already been documented in social science research studies of apprenticeship systems (Lave & Wenger, 1991), and collaborative learning has been widely proposed as an important pedagogical technique (Cohen, 1994; Webb, 1995).

When attempting to design technology, it is important to remember the triad of components of activity (e.g., Kuutti & Bannon, 1993). This triad is based on Russian activity theory (Bedny'i & Meister, 1997). In every situation, there are tools, activities, and people. These three elements are interdependent. A change to one element affects the others. When a new tool is introduced, people and their activities change to accommodate it. For instance, a piece of bookkeeping software might be introduced into a company. Initially, people will try to use the system to replicate their prior bookkeeping practices (new forms for old functions). Differences in how the software does things and the prior system will most likely chafe the users. Over time, people begin to change, learning the new possibilities of the software and adapting their practices (activities) to take advantage

of its benefits and work around its shortcomings.

Technologies can thus change the practices of the people in an organization profoundly. For instance, studies of the introduction of email into companies revealed that underlying power structures in the organization were changed—in some cases, drastically (Francik, Rudman, Cooper, & Levine, 1991). These technologies have an impact by changing not only what is possible in the organization but also what is easy (and hard). In the email study, the power structures changed because it became easier for people to communicate with others outside their work group (including those in upper levels of the employee hierarchy). It had previously also been possible to communicate across departmental lines, but email made it vastly easier and thus encouraged people to do so.

One example of a technology that supports but does not supplant student communication is CSILE, Computer-Supported Intentional Learning Environments (Cohen, 1995; Scardamalia & Bereiter, 1992; Scardamalia, Bereiter, McLean, Swallow, & Woodruff, 1989), in which students collaborate to co-construct a shared database of knowledge. The system allows students to flag ideas in ways that invite social interaction, such as "My theory for now is..." or "What I need to know now is..." This technology was successfully used to change classroom culture with elementary school students, yielding an atmosphere in which students took more responsibility for forming and answering their own questions, and in addition learned at least as well as with traditional didactic methods (Scardamalia & Bereiter, 1991). These University of Toronto based researchers continue these activities today with the Knowledge Forum, a commercially available knowledge-building software system for communities (Hewitt & Scardamalia, 1999).

Studies on another collaboration tool called SpeakEasy revealed that the interactive and social nature of using the tool was far more important for learning than the information that was exchanged (Hoadley, 1999; Hoadley & Linn, in press). In this tool, students were able to learn from peers online through a structured discussion tool without having access to any expert information. Their learning was related most closely to the interactivity of the online medium and relatively unrelated to the information they encountered in the discussion. This is a surprising finding, which emphasizes the importance of establishing a social context oriented toward learning. It also suggests that technologies should be designed with at least as much attention to social context as to the information presented within. For instance, when using SpeakEasy in a middle school science class, the inclusion of features such as an anonymity option *erased* the typical significant gender differences in student participation and learning (Hoadley, 1999; Hsi & Hoadley, 1997). Such dramatic effects from technology indicate that we can indeed build tools that help form and sustain more effective learning communities.

Knowledge networking for learning communities

We have described different models of organizational collaboration and learning and suggested that an effective learning community is a knowledge-building community of practice, one in which members of the community interact to collaboratively help other individuals and the group to increase their knowledge. This interaction is in contrast to mere information management, which ignores the role of social interaction in helping individuals find and come to understand information, thereby transforming it into knowledge. We have characterized in a general sense how technological tools can help support knowledge building by influencing people and their activities. In contrast to knowledge management tools or information management tools, where the focus is on helping to route information, knowledge-networking tools help foster all the constituent activities that increase knowledge building. These activities include not only information capture and transmission, but also the establishment of social relationships in which people can collaboratively construct understanding. In the next section, we describe how some of these steps have been carried out in helping to foster a new learning community called CILT.

The CILT Community

The Center for Innovative Learning Technologies, or CILT (pronounced "silt"), is an attempt to engineer a learning community among people who work with learning and educational technologies. Although this is a burgeoning area of work, with billions of dollars being spent annually on research, development, and deployment of technology in education, there are few effective mechanisms for getting information about what types of research and development have been done in the area. Several situations contribute to this problem. First, there is and has traditionally been a divide between (usually academic) researchers and practitioners and industry (Kozma, 1996; Office of Technology Assessment, 1988). This has been an ongoing problem that has been recognized by the U.S. federal government at the highest levels, as in the PCAST report (President's Committee of Advisors on Science and Technology, 1997). Teachers are generally in contact with academic experts only during their preservice training. In-service professional development attempts to update the teachers' skills and knowledge with the latest research, but this is typically limited to a few days per year. A second related issue is that technology in general, and educational technology research and development in particular, is changing rapidly. With the typical shelf life of an educational technology hovering around 3 to 5 years, an incredible amount of information must be read simply to keep up with the changes in the field. A third difficulty is the multidisciplinary nature of research and development in this area. Even among academics, researchers might be housed in departments as diverse as psychology, computer science, education, sociology, communications, and media. Indeed, many educational technologists are housed in the department of the discipline they are teaching (math, science, English, foreign language, etc.) and have no connection to a general educational technology community. These diverse researchers tend to frequent different conferences, read and publish in different journals or trade publications, and have no way of collaborating with each other.

The results are disappointing. Although much research on learning and technology has been carried out for more than 20 years, it is nearly impossible to answer the simple question "What do we know about what technologies work for learning?" (President's Committee of Advisors on Science and Technology, 1997). The field of learning technology has been accused of a lack of cumulativeness, an irrelevance to everyday practitioners (technology developers and teachers), obsolescence compared with the rapid advances in technology, and a disconnectedness that prevents anybody from finding useful information even if it does already exist.

CILT has been designed to address these challenges as a distributed center for fostering collaboration, research, and dissemination in learning technologies (Pea et al., 1999). The CILT organization was founded by four nonprofit and educational institutions (SRI International, the University of California at Berkeley, Vanderbilt University, and the Concord Consortium) in cooperation with industry and school partners. CILT's slogan, "Uniting people, technology, and powerful ideas for learning," reflects one of its main goals: helping to ensure collaboration and effective transfer of knowledge among members of a community of practice devoted to improving scientific understanding of and ongoing practices with learning technologies. In short, CILT is a knowledge-networking organization that is attempting to form a learning community of researchers, developers, and practitioners from academia, government, industry, and education.

CILT was formed in late 1997 to foster a learning community and has identified and tested a number of strategies since then. These strategies have included hosting workshops and conferences, awarding small grants to help new collaborators seek funding for joint research and development, sponsoring a postdoctoral program for training new professionals, and providing technologies to support collaboration. This last goal—of providing technologies to support the community—is the focus of the remainder of this chapter. The efforts to address this goal led us to the surprising conclusion that the best collaborative software is not any sort of traditional groupware, but rather a ubiquitous collaborative infrastructure.

Initial failures

Initially, CILT attempted to support collaboration through traditional community tools on its Web site. These included a number of leading commercial and research products for collaboration. Several Web-based discussion systems were tried (such as Allaire Corporation's Forums and Berkeley's COOL system), as were more unusual technologies, such as Digital Knowledge Assets' intelligent-agent-based collaborative workspaces, SRI's URLex URL exchange program, Vanderbilt University's Webliographer URL bulletin board, and Vanderbilt's LTSeek daily news publishing system. With the exception of LTSeek, each of these technologies failed to attract a significant user base. When motivated groups tried to use them for directed collaborative activities, they quickly reverted to their prior collaboration technologies, including telephone conferences and email mailing lists. Although there are many possible explanations, the most likely is that these tools did not support the users' workflow or collaborative needs. These "bolt-on" technologies (Gifford, 1996) were tried out for the problems at issue without attending to the range of issues presented below. It was at this time, in 1998, that we began the development of CILTKN (CILT Knowledge Network), a knowledge-networking technology, for this community. We now discuss our experiences in this design process.

Below, we describe the process of designing the CILTKN in terms of eight areas of activity. Although these areas of work can be seen as stages to be completed in order, in our case they were loosely overlapping. Six areas have been at least initially addressed in our design and implementation phase, while two more areas are under investigation in our evaluation and redesign phase.

Designing and implementing the CILTKN

Defining the learning community

In our case, much of the work of defining the learning community and its goals had been already accomplished through the writing of the initial grant proposal for CILT and the following discussions about how CILT would operate. The perceived challenges of the field—lack of cumulativity, lack of connection between research and practice, obsolescence, and disciplinary isolation—drove the goals and activities of CILT. The CILT leadership team set the following as CILT's goals: identifying areas of high potential for research and development, supporting rapid innovation, stimulating collaborative development in the selected areas, fostering interdisciplinary research and dissemination, and helping train new professionals in the field of learning technology research (Pea et al., 1999). CILT was envisioned as a learning community in which researchers, teachers, developers, and policy-makers would collaborate to share and build knowledge about learning, education, and technology.

Examining existing practices

Our initial failures were a strong motivation to examine existing practices in the audience we were trying to reach. Certainly, participating in online CILT discussions was not part of examining existing practice, so we went back to the drawing board and tried to enlist friends and colleagues to tell us what they really did need. We realized that the audiences of teachers, researchers, and businesspeople were probably different; given this fact, we decided initially to focus on researchers, both because they were the bulk of CILT's membership and because we felt that they would form a good base on which to build.

Information needs

Our first step in trying to uncover latent information needs was to ask people at CILT's workshops, engaging hundreds of researchers and other participants, what they hoped CILT would provide. It became clear that an important part of cumulativity was simply information about the learning technology community, for the community. We began holding brainstorming sessions with researchers at the four CILT institutions on what types of information might be useful. We, as the designers, narrowed their suggestions into a smaller list. In addition to wanting information about the CILT organization, people wanted very basic information, such as:

- Who else is involved in learning technology research?
- How do I contact them?
- What do they work on?
- What institutions are doing work in this area?
- What are the important research results?
- How can I get up to speed in this area?
- How can I share or post information on a job listing I have?
- How can I share or post information on an upcoming conference?
- and so on.

The respondents seemed entirely capable of establishing collaborations and carrying them out, if only they had a better handle on whom to collaborate with. As the example with David points out, social collaboration often begins with information finding. It seemed that people needed a little more information than what was readily available to begin collaborating. We therefore shifted our focus from supporting online collaboration directly to people's pre-collaboration information needs.

We translated these information needs into information types that addressed the needs. The kinds of information people seemed to need included:

People

Names, interests, and contact information for people in the field.

Projects

Descriptions and pointers to more information on research projects, implementation projects, or other work being done in the field.

Places

University departments or labs, K-12 schools, or other places where substantial activity (research, development, implementation related to educational technology) is taking place.

Papers

Research results or bibliographic pointers to research results in learning technology.

Syllabi

Course syllabi or reading lists to get people up to speed on learning technology. This also might help people decide what research papers were foundational or important in the field.

Collaboration notices

"Classified ads" or similar types of notices for items like jobs, conferences, etc., typically distributed through email lists.

Once we had identified these types of information, we began to look at how people currently arrived at the information.

Sharing practices

"I can't imagine losing my date planner; I keep my whole life in there!" Following this idea that people tended to keep important information in a single place, we began to survey and interview professional researchers about the information types listed above. We asked where they acquired the information, where they kept it, how they used it, and how they shared it (if they shared it at all). Two surveys of an approximately 50-person research department touched on researchers, clerical staff, students, and teachers. Participants were rewarded for returning surveys with candy bars, yielding a very high participation rate (over 75%).

The surveys revealed several important facts. First, there was no predominant system for keeping most types of information. Second, most people had organizational systems for their personal and professional information that they felt were woefully inadequate. Yet, people did generally manage to function perfectly well with their current systems, even if using them did take a substantial amount of time. (For instance, the time it was reported to take to format a bibliography to a paper was anywhere from a few hours to several days. Though people were largely citing papers they had read and could remember readily, the details of citations were difficult to find.) Third, people rarely shared these types of information; when they did, they would either type them into an email message or photocopy them. A few exceptions cropped up where a coherent system for sharing existed: many people made use of their email program's ability to store email addresses, and people generally shared their own contact information in the form of business cards or electronic signature files appended to their outgoing email. But when it came time to find information, people were back in the quandary David encountered in the opening section. Basic information about people in this "community" was nearly impossible to find. This fact was serving as a substantial damper on activities that would help establish it as a knowledge-building community. An excerpt from one year's findings in Table 1 shows primary means of storing information types. People could list more than one primary means (for instance, if they used both a personal digital assistant, or PDA, and a desktop program in tandem). Note the wide variety of systems.

Table 1: Primary locations of information types

		Phone numbers	Postal addresses	Email addresses	Events	Papers
Text files	Barely digital	1	1			2
	Barely digital	1	1	0	0	2
Explorer	Browser					
Netscape	Browser	1	1	2	2	2
	Browser	1	1	2	2	2
Contact	Database	2	2	1	1	
EndNote	Database					2
FileMaker	Database	1	2	4		
Touchbase	Database	1	1			
	Database	4	5	5	1	2
E-mailer	Mail program			3		
Eudora	Mail program	1	1	8		
Mulberry	Mail program			1		
Outlook	Mail program	1	1	2	1	
Pine	Mail program			1		
QuickMail	Mail program			1		
	Mail program	2	2	16	1	0
Address book	Paper	4	4	2	1	
Anywhere	Paper					1
Date planner	Paper	2	2	1	4	
Desk calendar	Paper				4	
Filing cabinet	Paper					2
Notebooks	Paper		1			1
Paper	Paper	1	1		2	
Piles	Paper					1
Post-Its	Paper	1	1	1	3	1
Rolodex	Paper	2	4	2		
Wall calendar	Paper				1	
Wallet	Paper	2	2	1		
Whiteboard	Paper				1	
	Paper	12	15	7	16	6
HP95	PDA	1			1	
Newton	PDA	1	1		1	
Pilot	PDA	6	6	3	6	
Watch	PDA				1	
	PDA	8	7	3	9	0
Business cards	Rummage	2	2	2		
Email	Rummage	2	2	4	2	1
Old papers	Rummage					6
	Rummage	4	4	6	2	7
TAPPED IN	Net	1	1	1	1	1
Web servers	Net	1	1			2
	Net	2	2	1	1	3
	Total	34	37	40	32	22

Identifying potential improvements

By interviewing our audience, we were able to identify a number of issues that seemed addressable with technology. Some were areas that required no technical innovation, only a good implementation. Others were (and are) ongoing technical research areas in computer science and human-computer interaction.

Heterogeneity of formats

The most obvious problem was that there was no simple way to exchange information with others because the information was rarely in a format that could be used directly by another person. One obvious distinction was between people who kept their information in digital form vs. people who kept their information offline, in a paper-based format. Over and over again, ease of use was the determining factor for each individual. Several people would keep telephone contact information on a well-worn piece of paper, folded to the size of a business card and kept in a wallet. Others, especially those for whom searching was important, would keep contact information online in some sort of personal information management software. Even those who did, however, could not readily exchange information because of the wide variety of incompatible file formats. In the first survey, there were nine formats for digital contact information in one department. Only two people used the corporate standard software that had been site licensed (Netscape Communicator). Although it might have been possible to export and import data in text formats, this practice was nearly unheard of, and the general perception was that such actions required technical gurus and arcane knowledge to make the process work. Even within paper-based formats, there were no standards. For instance, only two people used a physical Rolodex system for phone numbers, even though this was a common paper-based standard at one time.

Lack of structure in the data

Much of the information people stored was not well structured. For instance, when storing bibliographic information online, people typically would glean needed references from the ends of word processing files scattered around their hard drives. Since full-text search of files on desktop computers is only now beginning to be standard, many people would need to manually open many documents to search for a particular reference. Once they found it, it often needed to be reformatted—for instance, from an American Psychological Association style format into an Association for Computing Machinery format. In this case, there was no substitute for human intervention—the reference would have to be retyped. Likewise, many people knew that they could find contact information in signatures at the end of email messages, but a lack of structure in this data prevented importing it into their own contact manager; again, retyping was necessary.

Different solutions to the rummaging problem

Another barrier to sharing information that suggested technological intervention was the wide variety of organizational styles people employed in keeping professional information. It was apparent from the interviews that people engaged, to varying degrees, in what we term "rummaging." Rummaging is searching through loosely organized information when it is needed. Think of it this way: a person may be very organized, cataloging and filing every piece of information in a comprehensive organizational scheme. This up-front effort yields very short search times when the information needs to be retrieved. Libraries use this strategy, for instance. On the other hand, this effort is wasted if the time saved retrieving the information doesn't balance out the costs of creating and maintaining the scheme. Many respondents reported using lightweight organizational schemes such as chronological filing or piling of documents, "clumping" by topic, and the like. These schemes yield longer search times, but if the person refers back to the material infrequently, the person has a net saving

of effort.

The difficulty in sharing arises when people fall on different ends of the organize-now/rummage-later continuum. People who rummage may not feel comfortable letting others do the work of retrieval by looking through their materials, and someone who organizes ahead of time may not understand why rummaging is necessary in the first place. The work done by the organized person doesn't really help when he or she sends information along to a rummager, since the rummager's system doesn't have a way to preserve the work done by the first person. Thus, these differences can serve as a barrier to information sharing.

The Diner's Dilemma (incentives)

A fourth difficulty in any collaborative situation is the problem of incentives. As noted by Glance and Huberman in their paper on the Diner's Dilemma (Glance & Huberman, 1994), individual needs often compete directly with group needs. This competition can yield a worse outcome for everybody when people are not willing to give up a little for the greater good. The Diner's Dilemma situation is easy to grasp. Imagine going out to dinner with a group of people; nobody has discussed in advance whether the bill will be divided equally or calculated exactly. The diner must decide whether to order the hot dog, lowering his or her bill, or gamble on an even split and order lobster, with the cost being borne by his or her fellow diners.

Is the effort of putting information out for the community worth it to me? In the case of sharing information with a knowledge-building community, that is the fundamental question. If everyone participates, the community benefits (as do all the individuals in it). But if some people contribute while others merely consume their efforts, the costs of sharing information are unfairly carried. This situation can lead people to act protectively, expending as little energy as possible. In this case, everyone loses. We realized that whatever system we set up not only had to take into account the group's well-being, but had to be enticing enough to individuals to nudge them into sharing their data.

Social metacognition (know-who, not know-how)

The final difficulty we noted has already been brought up: how do you know whom to talk to? In the case of a learning community, knowing people in the social network is at least as important as having a lot of information at your fingertips. We realized that "know-who" was just as important as "know-how" or "know-what" (Kahn, 1999). Research on how novices comprehend a discipline has shown that social cues can facilitate understanding of the discipline and that sustained social interactions over time likewise facilitate learning (Hoadley, 1999; Hoadley & Enyedy, 1999; Hoadley & Linn, in press).

Finding ways technology can help

Working from the list of areas ripe for improvement, we expected the following technologies to help establish collaboration.

Standard online database technologies make quick work of storing information in an easily searchable format and making it available over networks. The fact that use of the World Wide Web is nearly ubiquitous among CILT's audience suggested using this technology for dissemination.

Although no unique standards existed for the types of information we were interested in, several technologies did exist to provide information in a variety of formats, and some formats were more easily exchanged than others. A careful study of each information type helped us uncover the best existing formats (for instance, vCards or LDAP servers for contact information) or technologies to support multiple formats (such as the

ReferenceWebPublisher software, which allows Web download of bibliographic references in the three most common formats (ProCite, Reference Manager, and EndNote). Technologies (such as Corex's CardScan software and hardware) that allow users to readily take unstructured or differently structured information and convert it into a common, structured data format also seemed promising.

To help ensure that the data would be easily shared, we began work on the development of VLTIs, or Very Low Threshold Interfaces. The idea was that if information could be accessed in a very quick manner without disrupting workflow, then users would have fewer disincentives to contribute and would be more likely to make use of the information. As the databases were enriched with more and better information, individuals would have more and more incentives to participate in maintaining and using them. We identified several desktop technologies that seemed promising for quickly finding information, such as Apple Computer's Apple Data Detectors and Sherlock technologies, which allow selected text in any application to be parsed and fed to search engines without launching an Internet browser.

Finally, we realized that tools for "know-who" would be important in our system. We envisioned that the use of recommender engines (Greer, McCalla, Kumar, Collins, & Meagher, 1997) and innovative visualizations of social information (Kautz, Selman, & Shah, 1997) would help individuals find one another and view information about learning technology research in the social context of the community.

Designing and building the tool

We began designing the CILTKN tool to help people connect and share information. Since our budget did not permit development of all the features we had designed, we started small, with most of the data types we had identified but few of the advanced features, such as online synchronization with desktop databases or recommender systems.

Currently, the CILTKN software (see figures 1-3) is up and running at <http://kn.cilt.org/> and has several hundred active users. Information available in the network includes People (contact information for researchers, teachers, and businesspeople), Pedagogy (course syllabi for undergraduate and graduate courses in learning technology from premier institutions), Papers (bibliographic information for important papers in the field of learning technology), Personals (requests for collaboration), and Places (labs or organizations that study learning and technology). Two kinds of information can be downloaded directly into people's desktop software: contact information, through the vCard format, and bibliographic information, through ReferenceWebPublisher. A partnership with AT&T Research has allowed us to use ReferralWeb (a dynamic visualization tool) to show connections between researchers in learning technology, as evidenced by coauthorship of papers. A demo is available at <http://www.research.att.com/~kautz/referralweb/> (note that the demo requires using an IBM-compatible computer). Each type of information can be searched, browsed, and contributed by members of CILT.

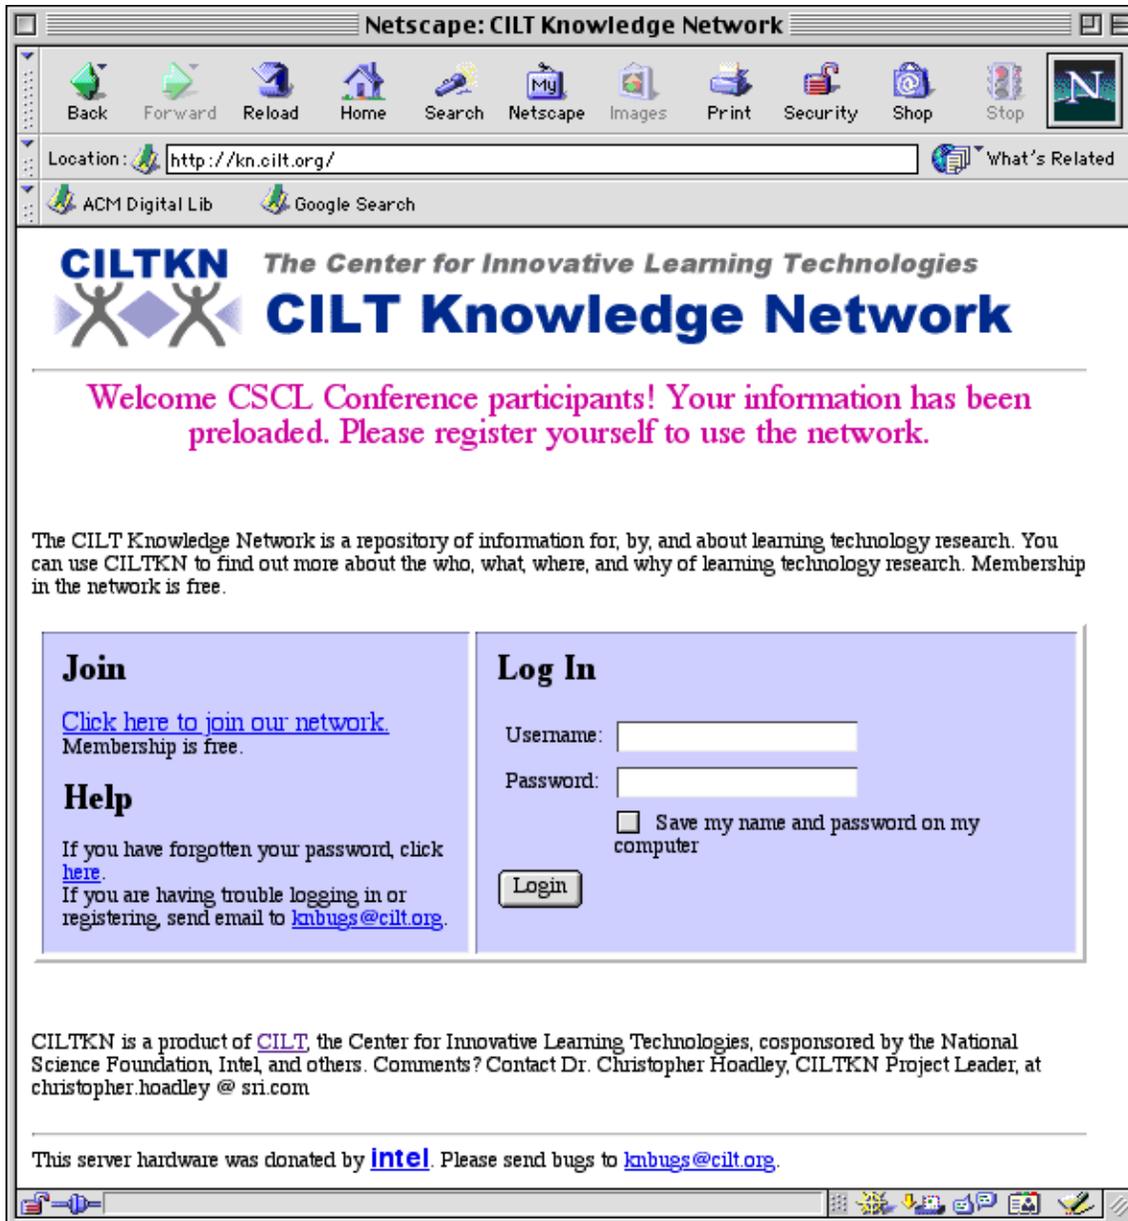


Figure 1: CILTKN

Opening Screen

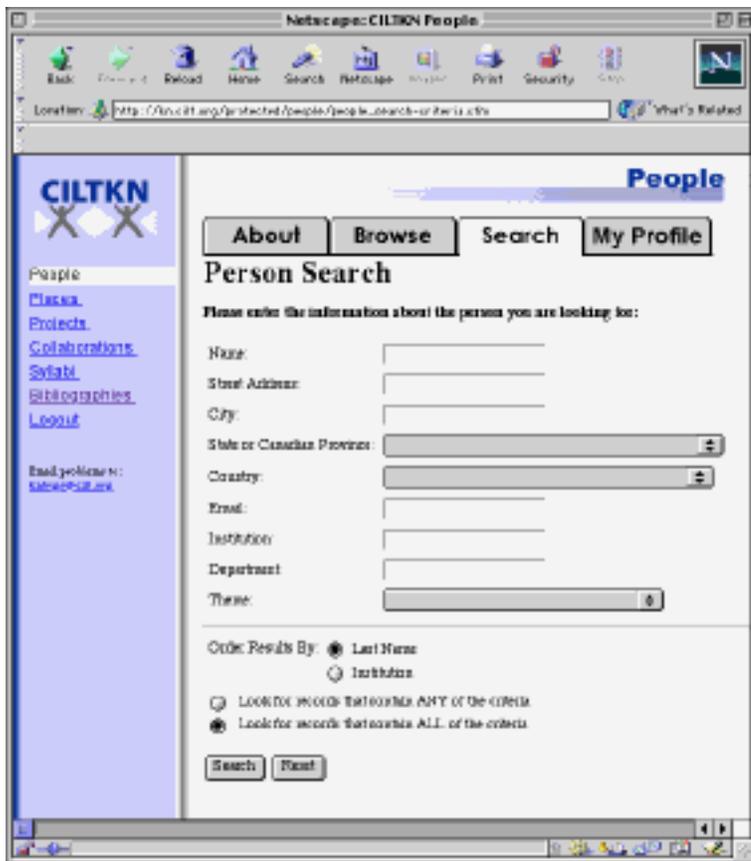


Figure 2: Searching the CILTKN

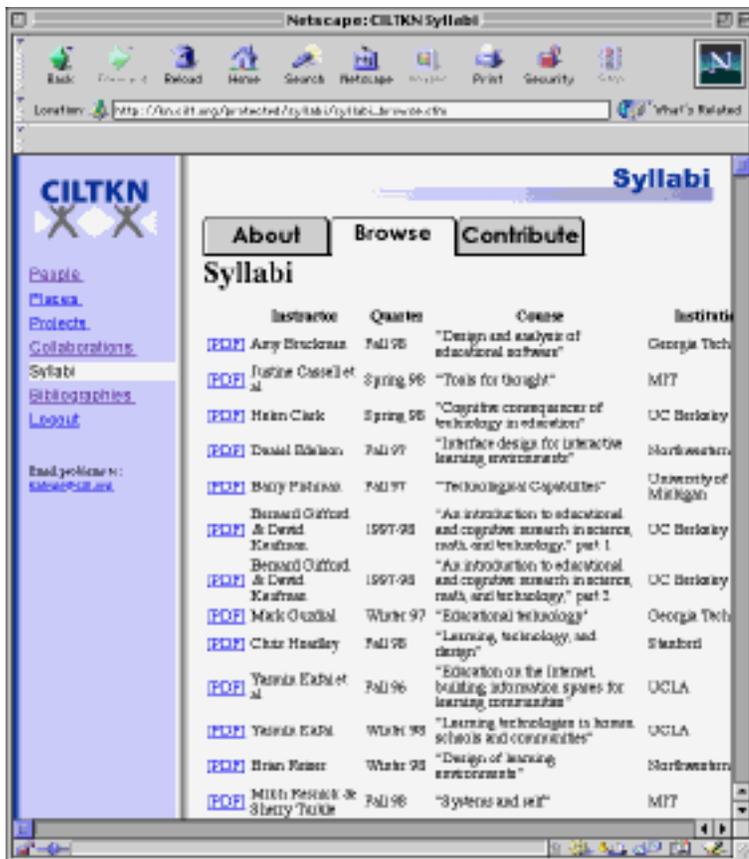


Figure 3: CILTKN Search Results

Cultivating a community of use

Fostering participation in CILTKN has taken several forms. First, we began laying groundwork by soliciting material. Some was collected from traditional sources, such as library or Web searching, but most was collected by personal appeal to members of the authors' social networks. Syllabi, in particular, had to be solicited from individual instructors since often they were not publicly available. By "passing the plate" for references and syllabi, searching for projects and places, and pre-entering hundreds of CILT members' contact information, we built a solid start to having databases that could describe the community.

The system was opened for public use at the CILT99 conference (April - May 1999). All attendees were encouraged to register themselves, and a subset of the databases (People, Syllabi, and Papers) were available for use and testing. Over the following months, additional data types were added, and the system was advertised through conference presentations and mailing lists.

One of the most powerful techniques we used to encourage appropriation of the tool was to employ CILTKN at the source of some of the knowledge-building activities already taking place within the community. For instance, CILTKN was used to collect submissions for the Computer Supported Collaborative Learning 1999 conference, and also for a contest sponsored by CILT in late 1999 for educational applications of handheld computers. Since people could also join CILT by simply checking a box, this encouraged more people to sign up. The use of CILTKN for conference submission allowed the capture of bibliographic information on papers as they were published. It even helped with maintenance of the databases, as current users were asked to confirm and update their contact information. We plan further integration of CILTKN with the learning

technology research community by using CILTKN to support registration in two of the field's professional organizations, the American Educational Research Association's Special Interest Groups in Advanced Technologies for Learning (ATL) and Education in Science and Technology (EST).

Last, but not least, personal reminders and social interactions outside the tool remain one of the most effective means to encourage participation. Invitations to participate in CILTKN always go out under the project leader's name (Hoadley) and often lead to brief conversations that serve to remind potential users of CILTKN that they are joining not just a mailing list but a community.

Future plans: Assessing our success

As mentioned earlier, CILTKN is already in use. Over 500 CILT members and over a thousand others use the system now. Most users return more than once, indicating that the tool is perceived to be useful. Our plan is to complete the design cycle by examining tool use and assessing its strengths and weaknesses.

Examining tool use

We have only a murky picture of how CILTKN is being used. Only recently did we begin to track individual users over time. We do know that several hundred unique users visit the site each month, and that these span many countries and include not only university researchers but also people from the education, government, nonprofit, and for-profit sectors. In fact, the most frequent users of the system (apart from CILT's leadership team) are non-researchers. We would like to conduct user interviews and possibly field observations to judge the impact CILTKN is having on daily workflow and to document the ways in which CILTKN is being used.

Evaluation

Although we don't have detailed analyses or surveys, anecdotal information suggests that the tool is succeeding at some of its goals. At least three people have reported that CILT is their first place to search for contact information after their personal address book, outranking even the search engines. This suggests that people are easier to find on CILTKN than anywhere else (and, hopefully, this will remove a barrier that previously existed for finding collaborators in this field). Several university instructors have used CILTKN in their undergraduate and graduate courses, pointing students to it for more information or even structuring student projects around the tool. We take this use as evidence of the kind of training CILT hopes to foster. At least two groups applying for one of CILT's mini-grants used CILTKN to do a background literature review before submitting their proposal to CILT. And we have a report that one officer of an international professional society in computer science used the syllabi in CILTKN to begin learning about educational technology. These incidents support the idea that CILT is fostering the kind of cumulativity and dissemination of results we had wished for.

We are beginning to operationalize measures of the learning community we hope to achieve. By defining our goals precisely enough to measure them, we hope to demonstrate real benefits from CILTKN and help guide further development by better characterizing how the tool is shaping the people and activities around it. Although CILTKN may not cure all the ills of learning technology researchers, we feel we have successfully demonstrated that, with care and attention, a learning community might be engineered where there really wasn't one before. By heeding all eight facets of creating a learning community, from definition of a learning community to evaluation, we came up with an innovative type of collaborative software that wasn't about supporting communication but about supporting a community and its need for information in a social context.

Pieces of the Puzzle

If our end goal is solving the puzzle of how to support learning communities, a number of questions must be considered. We reflect on the eight areas of inquiry we encountered in this project that may help achieve our goals: defining learning communities, examining existing practice, identifying potential changes to improve practice, finding ways that technology might effect these changes, designing and building the technology, advocating the technology and cultivating a community of use, understanding the consequences of the technology, and, finally, evaluating the community with respect to the original goal. If one were to attempt to change a particular community, one might view these as eight stages that occur more or less linearly (or cyclically). Although researchers are pursuing these are eight areas of inquiry in a number of settings, finding a case where all eight are present is quite unusual. Each area is essential to fully understanding how knowledge-networking technologies might help build learning communities, and each draws on a different research paradigm. Each of the eight types of inquiry is a type of research. As we step through these areas, we call attention to existing research paradigms that address each type of inquiry.

Defining learning communities

The notion of a learning community is not clearly understood. Indeed, this volume is a testament to the complexity of the question "What defines a learning community?" Even seemingly simple terms such as "collaborative learning," "shared goals," and "joint action" are hotly debated. The choice of definition is vital. Almost any group of individuals who interact might be called a community, and certainly people change and learn in some fashion as a result of every life experience, as we have indicated in our earlier discussion of the community-of-practice concept. Yet we need to be selective about what we hold up as exemplars of learning communities and how we recognize a community as a learning or knowledge-building community.

This volume contains a number of important efforts to define learning communities. In addition, others have discussed different definitions and indicators of learning communities. Organizational behaviorists identify the learning organization as important (Garratt, 1987) but offer few concrete measures of learning or of an organization as community. Woodruff (1999) describes some features that distinguish learning communities in terms of cohesion. Hsi (1997), following Pea (1993), defines learning communities as ones in which participants construct productive discussions (with productive discussions defined in terms of inclusiveness, knowledge integration processes, etc.). Scardamalia and Bereiter (1991) propose individual agency as an important feature of a learning community. Research is still needed to examine on a range of scales the different types of communities that exist and to characterize which ones may fairly be called learning communities. Continued philosophical and empirical inquiry is needed to define the nature of a learning community.

Examining existing practices

Before attempting to intervene with respect to a system to improve it, one generally characterizes its current state. Learning communities are no exception. Fieldwork could help pin down the existing state of affairs. Anthropologists, sociologists, and other social scientists study current work, home, and school environments for some of the characteristics that concern us—learning, both individual and group; collaboration, competition, and other forms of interaction; and the use of tools and their impact on the overall culture. This is often done by using ethnographic techniques, such as with our copier repair example (Orr, 1990), and is advocated for informing system design (Kling, 1991). This type of descriptive research is required to set the stage for

principled interventions.

Identifying potential improvements

Once a group has been characterized, and in some cases before, one can begin to identify areas that might be improved with respect to collective intelligence. This type of study is often the realm of industrial or process engineers or of management consultants. A careful examination of the groups and comparison with other collaborative groups often yields suggestions for how collaboration or knowledge sharing and knowledge building could be improved, for example, by "increasing communication between division X and site Y." Although it might be tempting to presume that these suggestions could simply be signed into marching orders, leaving the problem solved, identifying areas for improvement is not the same as discovering how to initiate reforms. Management experts frequently grapple with how to create a more learning-oriented organization (Cashman & Stroll, 1989; Davenport & Prusak, 1998). Although drawing implications from existing practices is far from an exact science, it is empirically informed by work on best practices drawn from studying many institutions.

Finding ways technology can help

Technology is too often thrown at problems with an attitude that it can solve any problem. This view, of course, is naïve. Much of the field of human-computer interaction is concerned with measuring how technologies change people and their behaviors, and theorizing how this interaction might be generalized or predicted. Because the technology affects the group only through its impact on individual people, supporting a community often means encouraging individuals to behave in a more group-oriented fashion. However, a user is unlikely to adopt tools that do not support his or her goals at least as well as other alternatives. So "win-win" situations must be ferreted out in which technology can enhance the community while minimizing costs to the individual user. Ehn's and Bodker's (Bodker, 1991; Ehn, 1989) work on participatory design illustrates research strategies for this goal. Like identifying improvements, this area of inquiry can benefit from best-practices research. It can also benefit from theories of human-computer interaction (cognitive, sociocultural, or otherwise), which predict the impact of technology on human systems.

Designing and building technologies

This aspect of changing a community through technology is perhaps the most visible—the actual design and creation of the technology tools. Design involves the balancing of the many constraints and multiple goals of the situation with the technological techniques available. Designers frequently have experience with what types of tools "work" in particular kinds of settings and must use their intuition, experience, and information that can be gathered (from user testing, for instance) to evolve a software or hardware design to fit the situation. Building the technology is another task, one that may be more difficult, given the designer's need to test and iterate the design. Typically, the design process is intimately tied to the advocacy of the intervention and cultivation of the community of users (Kling, 1991; Kyng, 1991). This phase can be driven by empirical research on design and engineering methodologies, and indeed many design methodologies have research methodologies (such as laboratory-based user testing) embedded within them.

Cultivating a community of use

A great deal of energy is needed to take a tool, introduce it to a community, and nurture it through adoption or, as we prefer to designate it, "appropriation" (Newman, Griffin, & Cole, 1989; Pea, 1992). Tool users come to appropriate a tool by establishing its fit with their work practices or changing their work practices to accommodate special properties of the tool as they come to perceive them. Community-oriented tools, in particular, need nurturing for such appropriation to take place, as do the communities they are intended to help (see, for instance, the description in this volume of the Math Forum). The proponents of the technology must help users overcome initial hurdles to appropriation. They then must help the community and the tool reach a productive equilibrium (which may include the development of very new practices or ways of working). Creating this culture of use is an important person-to-person task that goes beyond simply taking a technology and "throwing it over the wall" to the intended user community. It follows the aphorism that "Use is design"—that design does not end with what the technical designers have created but continues in what the user community makes of the tool in context. There is no one label for this class of activity, but it is practiced by technology coordinators, community facilitators, reformers, and community "champions" who help advocate use of the tool and participation in the community. It is a form of "reciprocal evolution" of technology, work practice, and basic research (Allen, 1991). We term it "cultivating a community of use." It is especially helpful if this participatory design process includes individuals who already have authority or power in the community, such as school administrators in the case of schools or, in business, managers and executives or, in some cases, unions. Although facilitating use of a tool may not initially seem like research, in fact research on collaborative tools cannot easily be separated from "community support." By definition, a research intervention requires the researchers to intervene in some way, and in this field the researchers are thus either directly or indirectly responsible for bringing the tool into the community. In developmental psychology, this type of activity has been practiced by "participant observers" (Becker & Geer, 1969a; Becker & Geer, 1969b; Trow, 1969); in anthropology, it derives from the ways observers participate in the cultures they study (Burgess, 1984; Charmaz, 1983). In tool design, it derives from the ways the tools are brought into the communities of study by the researchers or their agents. This type of action-oriented research is an essential component of studying technologies to support learning communities, and is perhaps the least well understood of the areas of inquiry.

Understanding technology's consequences

How is the technology used, and what effects is it having on the community it is being used in? These questions are often best answered by those in the thick of the matter, the users and participants. Again, anthropologists, ethnographers, and, to some extent, advocates study this question, as do media researchers. Many studies on email, for instance, study the outcomes as the tool has become more and more a part of organizational culture, even if the researchers themselves were not involved in the development of the software or the decision to use it in an organization (Perin, 1991; Reil & Levin, 1990). Participating in the community support (discussed above) almost always yields information on adoption and institutional change, although these may be studied separately (Orlikowski, 1992).

Evaluation

The last piece of the puzzle is formal documentation of what has happened and whether or not the technology, the community, and the individuals are successful. Obviously, success varies depending on the goals against which one wishes to measure it. In the case of learning communities, individuals might be assessed for learning, or groups of students might be assessed on their group skills for problem solving in the learning domain. Entire

communities might be evaluated on their size and the amount of participation, the degree to which members of the community help other members, or the net quality of the community's output (such as advances in a field made by a research community). A tool's success could be gauged by changes in these individual or group measures, or by looking at the tool's use directly: by investigating whether the tool is appropriated, by asking users how they use the tool and whether they find it helpful, by cataloging anecdotes of how the tool changes the community and individuals (Gay & Bennington, 1999).

Summary

The development of knowledge-building or learning communities is a complex, multifaceted task. By examining users like David, we came to understand that our goals for a learning community would not be addressed by any "magic bullet" technology solution. Instead, we undertook a lengthy design process that started with self-examination and self-definition and still continues today with community support, assessment, and evaluation. The challenges we faced are similar to those in other community-building efforts, and we have attempted to extract the intrinsic types of work required to engineer technologies to support an online community. Many of these areas of inquiry would exist even if we were not using technology to support our users, but they are all the more important when we consider designing software to support their needs. By now, the reader has probably noticed the wide variety of skills to be brought to the problem, from computer science and design to management and grassroots community building to social science research. To be successful at supporting learning communities, we need to address all the questions here in a multidisciplinary way that not only involves research on existing practices and definition of the goals for the community, but also supports design and implementation with community support, technologies that map to the users' needs, and reflection on community and individual outcomes.

Bibliography

- Allen, C. (1991). The reciprocal evolution of technology, work practice and basic research. In D. Shuler (Ed.), *Participatory design*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Becker, H. S., & Geer, B. (1969a). 'Participant observation and interviewing': A rejoinder. In G. J. McCall & J. L. Simmons (Eds.), *Issues in participant observation* (pp. 338-341). Reading, MA: Addison-Wesley.
- Becker, H. S., & Geer, B. (1969b). Participant Observation and Interviewing: A comparison. In G. J. McCall & J. L. Simmons (Eds.), *Issues in participant observation* (pp. 322-331). Reading, MA: Addison-Wesley.
- Bedny'i, G. Z., & Meister, D. (1997). *The Russian theory of activity: Current applications to design and learning*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Bodker, S. (1991). *Through the interface: A human activity approach to user interface design*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Brown, J. S., & Duguid, P. (1991). Organizational learning and communities-of-practice: Toward a unified view of working, learning, and innovation. *Organization Science*, 2(1), 40-57.
- Burgess, R. G. (1984). Multiple strategies in field research. In R. G. Burgess (Ed.), *In the field* (pp. 143-165). Winchester, MA: Allen & Unwin, Inc.

Cashman, P. M., & Stroll, D. (1989). Developing the management systems of the 1990s: The role of collaborative work. In M. H. Olson (Ed.), *Technological support for work group collaboration* (pp. 129-146). Hillsdale, NJ: Lawrence Erlbaum Associates.

Charmaz, K. (1983). The grounded theory method: An explication and interpretation. In R. M. Emerson (Ed.), *Contemporary field research: A collection of readings* (pp. 109-126). Boston, MA: Little, Brown, and Co.

Cohen, A. (1995, April). *Mediated collaborative learning — How CSILE supports a shift from knowledge in the head to knowledge in the world*. Paper presented at the Annual Meeting of the American Educational Research Association, San Francisco, CA.

Cohen, E. G. (1994). Restructuring the classroom: Conditions for productive small groups. *Review of Educational Research*, 64(1), 1-35.

Davenport, T. H., & Prusak, L. (1998). *Working knowledge: How organizations manage what they know*. Boston: Harvard Business School Press.

Ehn, P. (1989). *Work-oriented design of computer artifacts*. Stockholm: Arbetslivscentrum.

Fischer, G. (1999). Lifelong learning—More than training. In R. Mizoguchi & P. A. M. Kommers (Eds.), *International Journal of Continuing Engineering Education and Life-Long Learning* (Special issue on "Intelligent Systems/Tools in Training and Life-Long Learning").

Francik, E., Rudman, S. E., Cooper, D., & Levine, S. (1991). Putting innovation to work—Adoption strategies for multimedia communication systems. *Communications of the ACM*, 34(12), 52-63.

Garratt, B. (1987). *The learning organization: And the need for directors who think*. Aldershot, Hampshire, England: Ashgate.

Gay, G., & Bennington, T. L. (Eds.). (1999). *Information technologies in evaluation: Social, moral, epistemological, and practical implications* [*New Directions for Evaluation* #84]. San Francisco: Jossey-Bass Publishers.

Gifford, B. R. (1996). Bernie Gifford on changing the educational technical landscape. *EDUCOM Review*, 14-19.

Glance, N. S., & Huberman, B. A. (1994). The dynamics of social dilemmas. *Scientific American*, 270(3), 76-81.

Greer, J., McCalla, G., Kumar, V., Collins, J., & Meagher, P. (1997, December). *Facilitating collaborative learning in distributed organizations*. Paper presented at the Computer Support for Collaborative Learning 97 Conference, Toronto.

Hewitt, J., & Scardamalia, M. (1999, April). *Knowledge Forum: A next-generation discourse environment*. Paper presented at the Annual Meeting of the American Educational Research Association, Montreal, Canada.

Hoadley, C. M. (1999). *Scaffolding scientific discussion using socially relevant representations in networked multimedia*. Unpublished Ph.D. Dissertation, University of California, Berkeley, CA.

Hoadley, C. M., & Enyedy, N. (1999). *Between information and collaboration: Middle spaces in computer*

- media for learning. In C. M. Hoadley & J. Roschelle (Eds.), *CSCL '99: Proceedings of Computer Supported Collaborative Learning 1999*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Hoadley, C. M., & Linn, M. C. (in press). Teaching science through on-line, peer discussions: SpeakEasy in the Knowledge Integration Environment. *International Journal of Science Education*.
- Hsi, S., & Hoadley, C. M. (1997). Productive discussion in science: Gender equity through electronic discourse. *Journal of Science Education and Technology*, 6(1), 23-26.
- Hsi, S. H. (1997). *Facilitating knowledge integration in science through electronic discussion: The Multimedia Forum Kiosk*. Unpublished Ph.D. dissertation, University of California, Berkeley, CA.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.
- Kahn, T. M. (1999). Social "know-who" for building virtual communities: It's both what you know and who you know that count [World Wide Web page]. EOE (Educational Object Economy) Feature of the Week. Available: <http://www.designworlds.com/articles/KnowWho.html> [1999, November 1].
- Kautz, H., Selman, B., & Shah, M. (1997). The hidden web. *AI Magazine*, 18(2), 27-36.
- Kling, R. (1991). Cooperation, coordination, and control in computer-supported work. *Communications of the ACM*, 34(12), 83-88.
- Kozma, R. (1996). *Increasing the impact of advanced technology on the education market (Technical Report)*. Menlo Park, CA: Center for Technology in Learning, SRI International.
- Kuutti, K., & Bannon, L. J. (1993). *Searching for unity among diversity: Exploring the "interface" concept*. Paper presented at the INTERCHI '93 Conference, Amsterdam, the Netherlands.
- Kyng, M. (1991). Designing for cooperation: Cooperating in design. *Communications of the ACM*, 34(12), 64-73.
- Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. New York: Cambridge University Press.
- Linn, M. C., & Burbules, N. C. (1993). Construction of knowledge and group learning. In K. G. Tobin (Ed.), *The practice of constructivism in science education* (pp. 91-119). Washington, DC: American Association for the Advancement of Science (AAAS) Press.
- Newman, D., Griffin, P., & Cole, M. (1989). *The construction zone: Working for cognitive change in school*. New York: Cambridge University Press.
- Office of Technology Assessment. (1988). *Power On! New tools for teaching and learning* (GPO Stock No. 052-003-01125-5, ERIC number ED295677). Washington, DC: U. S. Government Printing Office.
- Orlikowski, W. J. (1992, November). *Learning from Notes: Organizational issues in groupware implementation*. Paper presented at the Computer-Supported Cooperative Work Conference, Toronto, Canada.
- Ornstein, M., & Cohn, A. E. (1938). *The role of scientific societies in the seventeenth century*. Chicago, IL: The University of Chicago Press.

- Orr, J. E. (1990). Sharing knowledge, celebrating identity: Community memory in a service culture. In D. Middleton & D. Edwards (Eds.), *Collective remembering* (pp. 169-189). Newbury Park, CA: Sage Publications.
- Pea, R. (1992). Augmenting the discourse of learning with computer-based learning environments. In E. De Corte, M. Linn, H. Mandl & L. Verschaffel (Eds.), *Computer-based learning environments and problem solving*. New York: Springer-Verlag.
- Pea, R., & Gomez, L. (1992). Distributed multimedia learning environments: Why and how? *Interactive Learning Environments*, 2, 73-109.
- Pea, R. D. (1993). Seeing what we build together: Distributed multimedia learning environments for transformative communications. *Journal of the Learning Sciences*, 3(3), 285-299.
- Pea, R. D., Tinker, R., Linn, M. C., Means, B., Bransford, J., Roschelle, J., Hsi, S., Brophy, S., & Songer, N. (1999). Toward a learning technologies knowledge network. *Educational Technology Research and Development*, 47(2), 19-38.
- Perin, C. (1991). Electronic social fields in bureaucracies. *Communications of the ACM*, 34(12), 75-82.
- President's Committee of Advisors on Science and Technology (PCAST). (1997). *Report to the President on the use of technology to strengthen K-12 education in the United States*. Washington, DC: Executive Office of the President of the United States.
- Reil, M., & Levin, J. (1990). Building electronic communities: Success and failure in computer networking. *Instructional Science*, 19, 145-169.
- Scardamalia, M., & Bereiter, C. (1991). Higher levels of agency for children in knowledge building: a challenge for the design of new knowledge media. *Journal of the Learning Sciences*, 1(1), 37-68.
- Scardamalia, M., & Bereiter, C. (1992). An architecture for collaborative knowledge-building. In E. De Corte, M. C. Linn, H. Mandl, & L. Verschaffel (Eds.), *Computer-based learning environments and problem solving*. Berlin: Springer-Verlag.
- Scardamalia, M., & Bereiter, C. (1994). Computer support for knowledge-building communities. *Journal of the Learning Sciences*, 3(3), 265-283.
- Scardamalia, M., Bereiter, C., McLean, R. S., Swallow, J., & Woodruff, E. (1989). Computer-supported intentional learning environments. *Journal of Educational Computing Research*, 6(1), 55-68.
- Trow, M. (1969). Comment on 'Participant observation and interviewing: A comparison'. In G. J. McCall & J. L. Simmons (Eds.), *Issues in participant observation* (pp. 331-338). Reading, MA: Addison-Wesley.
- Webb, N. (1995). Constructive activity and learning in collaborative small groups. *Journal of Educational Psychology*, 87(3), 406-423.
- Wenger, E. (1999). *Communities of practice*. New York: Cambridge University Press.
- Woodruff, E. E. (1999, December). *Concerning the cohesive nature of CSCL communities*. Paper presented at the Computer Supported Collaborative Learning '99 Conference, Palo Alto, CA.