

COMPUTER SUPPORT FOR KNOWLEDGE-BUILDING COMMUNITIES

Marlene Scardamalia
Carl Bereiter

Ontario Institute for Studies in Education
of the University of Toronto

Nobody wants to use technology to recreate education as it is, yet there is not much to distinguish what goes on in most computer-supported classrooms versus traditional classrooms. Kay (1991) has suggested that the phenomenon of reframing innovations to recreate the familiar is itself commonplace. Thus, one sees all manner of powerful technology (Hypercard, CD-ROM, Lego Logo, and so forth) used to conduct shopworn school activities: copying material from one resource into another (e.g., using Hypercard to assemble sound and visual bites produced by others), and following step-by-step procedures (e.g., creating Lego Logo machines by following steps in a manual). With new technologies, student-generated collages and reproductions appear more inventive and sophisticated—with impressive displays of sound, video, and typography—but from a cognitive perspective, it is not clear what, if any, knowledge content has been processed by the students.

In this chapter we offer a suggestion for how to escape the pattern of reinventing the familiar with educational technology. Knowledge-building discourse is at the heart of the superior education that we have in mind. We argue that the classroom needs to foster transformational thought, both on the part of students and teachers, and that the best way to do this is to replace classroom-bred discourse patterns with those having more immediate and natural extensions to the real world; patterns whereby ideas are conceived, responded to, reframed, and set in historical context. Our goal is to create communication systems where the relations between what is

said and what is written, between immediate and broader audiences, and between what is created in the here-and-now and archived are intimately related and are natural extensions of school-based activities, much as these processes are intertwined and are natural extensions of activities conducted in scholarly disciplines. Our efforts to create an enabling technology have led to the CSILE project—Computer Supported Intentional Learning Environments (Scardamalia & Bereiter, 1991a; Scardamalia et al., 1992). In this chapter, we focus on the educational ideas for knowledge-building discourse—with some discussion, toward the end of this chapter, on the technology.

The ideas represented in CSILE come from three lines of research and thought:

1. *Intentional learning.* Although a great deal of learning is unintentional, there are important kinds of school learning that appear to not take place unless the student is actively trying to achieve a cognitive objective—as distinct from simply trying to do well on school tasks or activities (Bereiter & Scardamalia, 1989; Chan, Burtis, Scardamalia, & Bereiter, 1992; Ng & Bereiter, 1991).
2. *The process of expertise.* Although expertise is usually gauged by performance, there is a process aspect to expertise, which we hypothesize to consist of reinvestment of mental resources that become available as a result of pattern learning and automaticity, and more particularly their reinvestment in progressive problem solving—addressing the problems of one's domain at increasing levels of complexity (Bereiter & Scardamalia, 1993; Scardamalia & Bereiter, 1991b). Progressive problem solving not only characterizes people on their way to becoming experts, but it also characterizes experts when they are working at the edges of their competence. Among students, the process of expertise manifests itself as intentional learning.

3. *Restructuring schools as knowledge-building communities.* The process of expertise is effortful and typically requires social support. By implication, the same is true of intentional learning. Most social environments do not provide such support. They are what we call first-order environments. Adaptation to the environment involves learning, but the learning is asymptotic. One becomes an old timer, comfortably integrated into a relatively stable system of routines (Lave & Wenger, 1991). As we explain further in later sections, there is good reason to characterize schools of both didactic and child-centered orientations as first-order environments. In second-order environments, learning is not asymptotic, because what one person does in adapting changes the environment so that others must readapt. Competitive sports and businesses are examples of second-order environments, where the accomplishments of participants keep raising the standard that the others strive for. More relevant examples as far as education is concerned are the sciences and other learned disciplines, where adaptation involves mak-

ing contributions to collective knowledge. Because this very activity increases the collective knowledge, continued adaptation requires contributions beyond what is already known, thus producing nonasymptotic learning. The idea of schools as knowledge-building communities is the idea of making them into second-order environments on this model.

In this chapter we focus on the third point, restructuring schools, but in a way that incorporates the other two points. Thus, the focus is on restructuring schools so that they become the kinds of environments that support the process of expertise, in particular, progressive problem solving as it applies to competence and understanding.

HOW SCHOOLS INHIBIT KNOWLEDGE BUILDING

Contemporary criticism of schools in the United States and Canada tends to be dominated by acute problems on one hand (drop-outs, drugs, violence; etc.) and on the other by comparisons with schools in other countries that score better on achievement tests. These criticisms, in turn, lead to reform proposals that address the acute problems or that advance means of bringing achievement up to European and Japanese standards. It cannot be said that school reform is being approached with much optimism, except in speeches by politicians—and with good reason. On the basis of demographic projections, the acute problems can be expected to get worse; as for achievement, there is little prospect of duplicating either the teaching force or the family support system that seems responsible for the high achievement of other societies. Furthermore, there is no reason to suppose that other nations will stand pat, waiting for us to catch up.

It has seemed to us that a more promising approach to school restructuring would start by examining how schools limit knowledge-building potential—including the high-achieving ones. By addressing fundamental shortcomings, it may be possible to do more than struggle to catch up.

The conception of expertise as a process affords a viewpoint on schooling that reveals certain drawbacks of a fundamental nature. Although schools are devoted to teaching useful cognitive skills and formal knowledge, they are not designed to foster the progressive problem solving that generates the vast informal knowledge that has been found to characterize expert competence. Instead, the following seem to be true of schools in general:

1. Schooling focuses on the individual student's abilities, dispositions, and prospects. Educators have failed to grasp the social structures and dynamics that are required for progressive, communal knowledge building.
2. Schooling deals with only the visible parts of knowledge: formal knowledge and demonstrable skills. Informal or tacit knowledge—both the kind

that students bring in with them and the kind that they will need in order to function expertly—is generally ignored in school curricula. The result, frequently, is inert knowledge, unconnected to the knowledge that actually informs thought and behavior.

3. The knowledge objectives that are pursued, limited as they may be, tend to be made invisible to the students. The objectives are translated into tasks and activities. The students' attention, and often that of the teachers as well, is concentrated on the activities and not on the objectives that gave rise to them.
4. Scope for the exercise of expertise—for progressive problem solving, in other words—is generally available only to the teacher, and schooling provides no mechanisms (such as those that exist in trade apprenticeships) for the teacher's expertise to be passed on to the students.

These defects are especially relevant to the development of experts and expertlike learners. Schools have never been designed with a conception of expertise as a process that can be fostered at all levels of development. They have all been built on a primitive conception of knowledge that leaves out most of what is required to become an expert.

KNOWLEDGE BUILDING: A THIRD WAY

For the most part, educational technology has accommodated itself to the conventional schizophrenia, in which didactic instruction and child-centered methods compete for control of the educational mind. Thus, we have drill-and-practice, tutoring, and instructional management programs on the one hand, and we have a variety of exploratory and activity-centered programs on the other. The arguments for and against didactic approaches and child-centered ones are so familiar that there is no reason to review or criticize them here. Suffice it to say that any hope for technology to have a role in restructuring education must take the form of searching for a third way—something that is neither didactic nor activity-centered nor a mere compromise between the two (which is what already exists in most schools).

In searching for a third viable form of schooling, educational thinkers have looked outside the school for models. Thus, traditional apprenticeship has been examined as a possible model, one that provides for a natural but highly goal-oriented kind of learning (Collins, Brown, & Newman, 1989). The learned disciplines themselves show promise as models for the redesign of schools. This notion makes most sense when considered in light of the ideas we have been trying to advance about expertise—conceiving of it as a process of progressive problem solving and advancement beyond present limits of

competence. In the sciences, problem redefinition at increasingly high levels is the goal, based on a fundamentally social process. Researchers benefit from the advances of others, with continual interplay of findings, not just among scientists working concurrently, but from generation to generation.

There have been previous efforts to capture the character and spirit of scientific inquiry in the classroom. Several elementary school science curricula and social studies curricula developed during the 1950s and early 1960s were of this kind (see Bruner, 1964). However, the emphasis was on students as individuals engaged in the processes of scientific inquiry, rather than on the class as a collective engaged in the processes of a scientific community.

In recent times, however, people have begun to attend more to the social processes of research teams and laboratories, which have a character and a power quite different from a mere aggregation of individual researchers. Whitehead (1925/1948) recognized this decades ago, when he credited the German universities of the 19th century with having discovered how to produce disciplined progress instead of having to wait for "the occasional genius, or the occasional lucky thought" (p. 99). So successful have research centers been that they have begun to be used as models for many other kinds of enterprises—for management teams, sales teams, even secretarial staffs (Peters, 1987). The restructuring of manufacturing processes around quality circles also owes something to the research team as a model. Why, then, should the research center not also inform school restructuring?

As we suggested earlier, by focusing on the individual student's abilities and dispositions, educators have failed to grasp the social structures and dynamics that are required for progressive knowledge-building of the kind to which Whitehead referred. In effect, they have remained fixed on a pre-19th-century model of science, dependent on "the occasional genius, or the occasional lucky thought." Their focal question has been, "To what extent can a child be expected to act like a physicist, biologist, historian, literary scholar, anthropologist, or whatever?" The answer to this question will necessarily be equivocal. Of course, children are curious about the world, and they can in some fashion collect and evaluate evidence, venture explanations, test conjectures, and so on. Thus they can be said to act like researchers. But it is doubtful how far these talents can take them, and so there are perennial questions about how much discovery methods can be relied on to develop students' knowledge. Furthermore, fixing on the individual talents, needs, and learning outcomes only suggests to didactic educators that research skills and laboratory activities should be incorporated into the curriculum, and confirms for child-centered educators the claim they have been making all along: that children's curiosity should be allowed to guide their activities. It does not suggest any new structure for schooling.

More significant implications follow if the question is reformulated at the level of the group rather than the individual: Can a classroom function as a knowledge-building community, similar to the knowledge-building communities that set the pace for their fields? In an earlier era, it would have been possible to dismiss this idea as romantic. Researchers are discovering or creating new knowledge; students are only learning what is already known. By now, however, it is generally recognized that students construct their knowledge. This is as true if they are learning from books and lectures as it is if they are acquiring knowledge through inquiry. A further implication is that creating new knowledge and learning existing knowledge are not very different, as far as psychological processes are concerned. Thus, there is no patent reason why schooling cannot have the dynamic character of scientific knowledge building. If there are insurmountable obstacles, they are more likely to be of a social or attitudinal kind than of a cognitive kind.

The idea of restructuring schools as intellectual communities of some sort is very much in the wind these days. Brown and Campione (1990) proposed communities of learners and thinkers; Lipman (1988) proposed "community of inquiry" (p. 67). We strongly prefer our own term, *knowledge-building community*. It suggests continuity with the other knowledge-building communities that exist beyond the schools, and the term *building* implies that the classroom community works to produce knowledge—a collective product and not merely a summary report of what is in individual minds or a collection of outputs from group work.

The idea of knowledge as a product, enjoying an existence independent of individual knowers, presents epistemological difficulties with which educators are not accustomed to contending.¹ More familiarly, the problems of objectified knowledge are being wrestled with in such contexts as technology transfer, institutional memory, and intellectual property law. In science, it is clear that when we talk about Newton's theory, we are not merely talking about something once encoded in Newton's brain but about something that even today is discussed, tested, taught, applied, evaluated, and credited with causal force. When we speak of schools as knowledge-building communities, we mean schools where people are engaged in producing knowledge objects that, although much more modest than Newton's theory, also lend themselves to being discussed, tested, and so forth without particular reference to the mental states of those involved, and where the students see their main job as producing and improving such objects. Restructuring schools as knowledge-building communities means, to our minds, getting the community's efforts directed toward social processes aimed at improving these objects, with technology providing a particularly facilitative infrastructure.

¹Popper (1972) addressed these problems with his work on objective knowledge and epistemology without a knowing subject.

WHAT MAKES KNOWLEDGE-BUILDING COMMUNITIES WORK?

In trying to develop ideas of how to achieve knowledge-building communities in schools, we first considered knowledge-building communities we are already familiar with, namely those that exist in research-oriented universities and in research centers. These have also been the focus of much recent research by sociologists of science.

According to Latour (1987) who, along with a number of other contemporary sociologists, has studied the workings of scientific laboratories at first hand, the selfless pursuit of knowledge is a story that is fabricated after some claim has achieved factual status and is no longer controversial. Before that point, scientific practice is more like politics, an effort to marshal support for one's position. We should not expect school students to act a great deal differently, and it seems likely that past efforts to bring scientific inquiry into schools have suffered from promoting an idealistic model that is at odds with reality. Protocol studies of students carrying on scientific discussions indeed show frequent evidence that discussion is treated like a contest (Eichinger, Anderson, Palincsar, & David, 1991). What the sociologists fail to explain is why science works as well as it does, given the unseemly characteristics they have observed.

The problem of accounting for the success of knowledge-building communities is like that of accounting for the performance of an old Swiss watch. On microscopic inspection, the watch will be found to contain so many irregularities and imperfections that it will seem unlikely that its readings could have much validity at all, and yet it keeps nearly perfect time. And so with science, as with watches, the major challenge is to explain how they work so well, given their imperfections. If schools are to be transformed into effective knowledge-building communities, we need that kind of information.

Our own analysis is necessarily limited and impressionistic. We started by considering the role of journals in the progress of learned disciplines. As it happens, Latour (1987) devotes a significant part of his analysis to journals as well. The focus is not on the journals themselves and their content, but on the whole journal publication process, with its editors, editorial boards, reviewers, and contributors.

The imperfections of the journal process are well known and again lead to the conclusion that such a flawed process could not possibly work to advance knowledge. Unreliability of judgment, bias, political maneuvering, conservatism, failure to detect gross errors—all are familiar (see Peters & Ceci, 1982, and the whole journal issue devoted to discussion of their experiment, in which previously published articles were slightly disguised and resubmitted to the same journals). Nevertheless, discipline-based journals manage to harness an enormous amount of energy and get it working toward

collective advance in knowledge, and so they surely hold a key to what makes knowledge-building communities work.

The fundamental point that distinguishes scholarly journals from other periodicals is the requirement that the articles be contributions to knowledge—that is, that they represent some advance over what is already known. Peer review, usually pointed to as the essential characteristic of scholarly journals, is subordinate to this criterion, a way of ensuring that it is met. The knowledge-advance criterion, universal in scholarly journals, is foreign to the writing students do in schools, even on into graduate school. "How could it be otherwise?" one might ask, given the unlikelihood of a novice's finding out something that would advance a discipline. But it should be recognized that the knowledge-advance criterion is always to some extent local. In psychology, for instance, there are occasional articles suggesting the relevance to psychology of methods or concepts that are already well known in some other field, such as economics or information science. During the whole Cold War period, there were articles informing American psychologists of the work of Soviet psychologists. Operationally speaking, an article represents an advance in knowledge if it is so experienced by the peer reviewers. By extension, then, if the reviewers were other students, a student contribution would meet the knowledge-advance criterion if the student reviewers found that it advanced their own knowledge. Thus there is no intrinsic reason why the knowledge-advance criterion cannot be applied to student efforts. But to structure classroom activity so that a peer review system could be fully functional would be a radical restructuring.

Creating the structures that make peer review of knowledge advances possible would not be sufficient to make a viable knowledge-building community, however. There must also be motivation to do the work that goes into the construction of collective knowledge. Here, again, we may look to the journal process in scholarly disciplines for pointers. There are strong material rewards motivating young academics to publish, but these do not explain the sustained publication effort of established academics, nor the work that goes into reviewing manuscripts, which is often considerable and (being usually anonymous) earns no rewards this side of heaven.

Some other motives that appear to figure in academic publishing are the desire for recognition and respect from the people one regards as peers, the desire to have impact (on conclusions being reached, on the development of the discipline, etc.), and the desire to participate in significant discourse.

These motives have recognizable counterparts in school students. The problem is to get them attached to knowledge-building activity. Recognition and respect from peers can come from many sources, and contribution to the group's collective knowledge is not usually prominent among them. The same goes for having impact. And what students find to be significant

discourse—the kind they will get truly involved in, struggling for a turn to speak, actually listening to and responding to what others say—will often deal with issues closer to their personal lives than the issues arising from scholarly inquiry.

Our focusing on journal publication may seem like a case of mistaking the wrapper for the candy bar. What about research? What about curiosity? We do not mean to slight either of these. Surely, scholarly disciplines would not exist without them. However, these have received ample consideration in previous thinking about school learning. What has been neglected until recently is the role of discourse and the role of motives other than purely epistemic ones. Decades ago, Popper (1962) recognized argument and criticism as the driving forces in the advancement of scientific knowledge, with research having its impact through these discourse processes. But only in the last few years has talking science (Lemke, 1990) begun to be recognized as a necessary adjunct to hands-on investigation in school science. The use of inquiry methods in schools has been based on a frequently disappointed confidence in the power of children's natural curiosity. The study of scholarly discourse, as embodied in the journal process, shows us how a wide range of human motives (including curiosity, of course) are marshaled in the actual progress of knowledge construction in the disciplines.

SPECIFICATIONS FOR KNOWLEDGE-BUILDING DISCOURSE

How does one characterize knowledge-building discourse and then recreate classroom activity to support it? We could imitate at the surface level—for instance, by having classes produce scholarly journals with peer review. In fact, the CSILE implementation we describe later in this chapter has provisions for doing that. But it is not likely that imitation of surface forms can produce the restructuring necessary to turn schools into real knowledge-building communities. The whole journal process could easily be degraded into just another form of schoolwork. That would happen if the essential point were lost, that publications should embody contributions to collective knowledge.

There is plenty of discourse in schools, but it has little resemblance to the kind that goes on in knowledge-building communities. Most of the oral discourse can be characterized as recitation (Doyle, 1986). When there are discussions that could be construed as building knowledge, they are generally led by the teacher. Socratic dialogue is the model, which means that the teacher, playing Socrates, gives the discussion such direction as it has and is therefore likely to be the only one whose goals have substantive influence on the outcome. The students' own goals may influence how

successful the discussion is, but mainly through influencing the extent of their cooperation. Transcripts of classroom discussion indicate that it typically consists of a string of three-step units, each unit being made up of the following conversational moves: teacher initiates, student responds, teacher evaluates (Heap, 1985). Whatever this formula may represent, it surely does not represent the pattern of discourse in a knowledge-building community.

In order to begin defining characteristics of such discourse, we have drawn on analogies with groups working at the forefront of their fields, and considered how new knowledge media might not only support, but enhance, their work. At the same time, we have kept the constraint in mind of defining characteristics applicable across the span from kindergartens to advanced research institutes. The result, presented next, is what we hope is the beginning of specifications for knowledge-building discourse to be enabled by new knowledge media.

KNOWLEDGE-BUILDING DISCOURSE: THE CLASSROOM AND BEYOND

We have roughly divided characteristics for knowledge-building discourse into three categories: (a) focus on problems and depth of understanding; (b) decentralized, open knowledge environments for collective understanding; and (c) productive interaction within broadly conceived knowledge-building communities.

Focus on Problems and Depth of Understanding

In knowledge-building contexts, the focus is on problems rather than on categories of knowledge or on topics. Explaining is the major challenge, with encouragement to produce and advance theories through using them to explain increasingly diverse and seemingly contrary ideas. Engagement is at the level of how things work, underlying causes and principles, and interrelatedness of ideas explored over lengthy periods of time, and returned to in new contexts.

Decentralized, Open Knowledge Building, With a Focus on Collective Knowledge

From the perspective of social interactions, there is an expectation of constructive response to one another's work. Inquiry on all sides is driven by questions and desire for understanding. Negotiating the terrain around ideas is marked by complex interactions with others, using purposeful and constructive ways to engage busy people, to distribute work among members, to sustain increasingly advanced inquiry, to monitor advances of distant groups working in related areas, and to ensure the local group is indeed working at

the forefront of their collective understanding. There is also a great deal of opportunistic work, oftentimes in small groups (as opposed to legislated schoolwork of the conventional kind, where students are working individually but all doing the same thing or are subdivided in some arbitrary fashion).

In knowledge-building discourse, more knowledgeable others do not stand outside the learning process (as teachers often do), but rather participate actively. Further, the knowledge of the most advanced participant does not circumscribe what is to be learned or investigated. There are other sources of information, and participants aim to point the way to other groups and resources that might prove helpful.

Less knowledgeable participants in the discourse play an important role, pointing out what is difficult to understand and, in turn, inadequacies in explanations. To the extent that novices can be engaged in pushing the discourse toward definition and clarification, their role is as important as that of those more knowledgeable. In all, knowledge-building begets knowledge-building: Important factors include the creation of a climate and desire to advance understanding rather than to display individual brilliance (although individual brilliance can certainly help in the collective effort), and opportunities more plentiful than restricted communities allow.

The Broader Knowledge Community

Peer review for scientific publication exemplifies working with ideas in contexts broader than one's immediate working community. We have re-written this chapter in response to reviewers who raised issues that had not been raised in more local review processes. Additionally, the different reviewers brought different perspectives depending on their areas of expertise. All of this has proved quite helpful in allowing us to address a broader audience and to advance our own understanding in the process.

Earlier we made a distinction between first- and second-order environments. In first-order environments, learning is asymptotic—one can become comfortably integrated into a relatively stable system of routines. In second-order environments, learning is not asymptotic because what one person does in adapting changes the environment so that others must readapt. Adaptation itself involves contributions to collective knowledge. Because this very activity increases the collective knowledge, continued adaptation requires contributions beyond what is already known, thus producing non-asymptotic learning. Working within the broader knowledge-building community places one in a second-order environment, and accustoms participants to viewing ideas from the perspective of multiple expertises and issues. (Such anticipation and writing to broader audiences could not be more different from the normal pattern of school writing.)

We have barely begun the process of extending CSILE into a wide-area configuration, and in turn dealing with the educational issues that will come

about in the process of having student discourses more broadly available. We see potential for new educational models of openness and decentralization powered by a communal database of the sort that underlies CSILE (see next section). It is a logical extension of this communal database to have all participants at all levels (including, but not limited to, student, teacher, administrators, researchers, curriculum designers, and assessors) entering ideas into the same database. Thus, for example, if teachers are discussing students' problems in understanding a concept, students might be engaged along with them in the discussion. Although openness is an important principle, it must also be recognized that knowledge building requires private and directed discussions at times, so one of the many challenges in coping with educational uses of a communal database is to interweave open and private discourses, and to provide conditions for freedom from irrelevant, boring, or otherwise unhelpful information.

With the advent of wide-area networks for schools, students will have access to all manner of databases, CD-ROMs, video, microworlds, and so forth, as well as links to live experts and more advanced students. The challenge we see for educational technology is to preserve a central role for the students themselves, lest they be reduced to passivity by the overwhelming amounts of authoritative external information available. The surest way to keep the students in the central role, it would seem, is to ensure that contacts with outside sources grow out of the local knowledge-building discourse and that the obtained information is brought back into that discourse in ways consistent with the goals and plans of the local group.

At this point, it is fanciful (but nonetheless exciting) to contemplate advantages of having communal structures that span the whole of the school years, and that also profitably engage those in research institutes and other knowledge-creation enterprises. The fancifulness is not with the technology—recent developments make that by far the easy part. The problems to be solved are educational. As the preceding discussion indicates, it is the nature of the classroom discourse that determines whether the classroom functions as a knowledge-building community rather than, say, a classroom focused on pursuit of individual interests or on teacher-organized activities. In the next section, we turn to the issue of CSILE as an enabling technology for knowledge-building discourse.

HOW TECHNOLOGY CAN HELP REFRAME CLASSROOM DISCOURSE TO SUPPORT KNOWLEDGE BUILDING

In this section we suggest means for reframing classroom discourse to support knowledge building in ways extensible to out-of-school knowledge-advancing enterprises and indicate how we are attempting to realize these through CSILE.

A Community Database at the Center of Classroom Discourse

CSILE's community database is created by students. Users produce public access material, not simply material to be turned in for grading, and do so in a context that engages others on their behalf. Although students can choose to keep material private, the default option is public. Using networked microcomputers, a number of users (located within or outside the school walls) can be simultaneously creating text or graphical notes to add to the database, searching existing notes, commenting on other students' notes, or organizing notes into more complex informational structures. The community database serves as an objectification of a group's advancing knowledge, much like the accumulating issues of a scholarly journal, but with additional facilities for reframing ideas and placing them in new contexts. In local-area configurations, students' writings are available to classmates, not just to the teacher, and that gives them a feel for speaking and being responsible to a broader audience. In wide-area configurations, the audience is expanded, and with that comes an increased need to address problems and represent knowledge in ways that are comprehensible to people outside the immediate context. CSILE is designed to frame students' ideas in ways extensible to the broader knowledge-building community, and concomitantly, to resist discourse frameworks only workable in schools. Commitment to the notion that students can serve as legitimate partners in knowledge building is reflected in the fact that they are placed centerfront in the knowledge-creation process as authors of databases, not simply reviewers of databases created by others.

The database, which is wholly created by students, consists of text and graphical notes. Graphical notes can be used to create organizing frameworks. Anyone can add a comment to a note or attach a graphic note subordinate to another graphic note, but only authors can edit or delete notes. Authors are notified when a comment has been made on one of their notes, and the notes of all participants are accessible through database search procedures. This basic set of features represents the core functionality of a system in which the construction of knowledge is a social activity. For an account of other features that are available and envisioned, see Scardamalia and Bereiter (1992, 1993), Scardamalia et al. (1992), and Scardamalia, Bereiter, McLean, Swallow, and Woodruff (1989).

Focus on Problems and Depth of Understanding

Specially Designed Discourse Environments

We are creating note-writing environments so that surrounds convey and support knowledge-building operations of the sort otherwise absent from student interchanges. For example, a discussion note encourages students

to frame their inquiry in light of a problem rather than a topic, and their interactions in light of statements of theory and information needed to advance that theory. The note type also encourages commentary (Hewitt, 1996).

Emphasis on Intentionality

Studies suggest that the hallmark of the intentional learner is the ability to diagnose one's own learning needs and to identify next steps. Accordingly, the CSILE approach is to have students write statements of what they need to understand in order to make conceptual advances, with others engaged in helpful support activity (offering references, suggesting alternatives, and so forth). Additionally, CSILE places intentional overhead on activities. For example, students do not simply link notes, they write justifications for links they create. CSILE's low-tech approach to diagnosis (students diagnose their own needs and write an "INTU"—I need to understand—note) contrasts sharply with that of intelligent tutoring systems. With intelligent tutoring systems, the intentionality resides in the system's own diagnostic and decision processes. The contrasting view, which we have embodied in CSILE, is that an important part of education is for students themselves to learn to carry out those diagnostic and decision processes.

Decentralized, Open Knowledge Building, With a Focus on Collective Knowledge

Reversing the "Teacher Initiates, Student Responds, Teacher Evaluates" Pattern for Oral and Written Discourse

In recent years, educational computing has shifted strongly toward what is called a distributed model. The idea seems to have two components. One is that information should flow freely among participants, without having to pass through a central authority. The other is that knowledge should be distributed across students, rather than each student being expected to know the same things, thus making for more productive exchanges between students (Brown et al., 1993; Pea, 1993; Resnick, 1996). CSILE is designed to support a distributed model in both these senses, through the following features:

Elimination of Turn-Taking Problems. Classroom discussions with 20 or 30 participants typically feature the teacher as leader, if only to manage the turn-taking. With asynchronous discussion over a computer network, any participant can take a turn at any time.

Peer Commentary and Notification. CSILE has facilities that encourage users to comment on each other's notes, and provides automatic notification to authors of the availability of comments.

Entry Points for All Ages and Ability Levels. When networks extend across classroom boundaries, younger students question and challenge older ones. Less language-proficient students can represent ideas graphically or copy and edit text from other notes to express their own ideas. Less knowledgeable students can contribute through their questions and their supportive comments. Although no medium is culturally neutral, open systems like CSILE offer opportunity for culturally different students to appropriate ideas in their own ways and for their own uses.

Maximizing the Interplay and Value of Different Communication Modes

CSILE-supported classrooms have as much opportunity for oral inter-change as students in other classrooms. Accordingly, CSILE-supported classrooms allow for the immediacy, spontaneity, and ease of conversation, as well as the more reflective and long-term benefits of written discourse. Additionally, different communication modes are supported within CSILE. Students choose the mode appropriate to their talents, goals, and problem at hand. As suggested earlier, CSILE's goal is to increase the range of expressive languages to include video, audio, and animation, as well as specially designed contexts for mathematical, historical, geographical expression. This framework has allowed us to maximize advantages of particular discourse modes, as well as to encourage the following kinds of contributions unique to the written word:

Reflection. Students using CSILE have frequently commented on the blessing of having time to think rather than needing to respond under the pressures of oral discourse.

Publication/Review Process. The system supports a publication process similar to that of scholarly journals. Students produce notes of various kinds and frequently revise them. When they think they have a note that makes a solid contribution to the knowledge base in some area, they can mark it as a candidate for publication. They then must complete a form that indicates, among other things, what they believe the distinctive contribution of their note is. After a review process (typically by other students, with final clearance by the teacher), the note becomes identified as published. Users searching the database may, if they wish, restrict their search to published notes on the topic they designate. At the end of the school year, a class can

decide on a selection of notes to remain in the database for the benefit of classes that come after them. Thus, as in the real world, each generation does not have to rediscover everything that the previous generation found out, but can instead attempt to go beyond it.

Cumulative, Progressive Results. Even when oral discourse proceeds optimally, it is difficult for it to achieve cumulative, progressive results because of its transitory nature. Hence the advantage for written discourse.

Independent Thought. Conversation tends to favor the ideas of the most vocal, and to limit independent processing of material for all but the responder and most intentional students. In CSILE, each student is responsible for contributing to the discourse.

Diverse Arrangements for Supporting Small-Group Interchanges

CSILE allows for small-group discussion, and additionally provides records that bring those discussions to a broader audience (Scardamalia, Bereiter, & Lamon, 1994).

Increased and Diversified Response to Ideas

Under classroom conditions, written communication tends to be centered on the teacher because of the practical difficulties in giving every student access and opportunities to respond to what other students have written. Having documents available through a computer network can solve these problems.

Next we provide two examples of CSILE use, both involving fifth- and sixth-grade students, to give an idea of CSILE's knowledge-building facilities. The first example is not notable as an advance in subject-matter knowledge, but as an advance in methodology, achieved by the students themselves and enabled by the technology. The class was studying medieval history, one of the topics being castle defenses. In addition to compiling text notes recording their findings and speculations on this topic, many students availed themselves of CSILE's graphics facilities to produce graphical notes depicting their understanding of castle defenses. Two students, generally regarded as below-average achievers, took a different tack. As they explained in a later interview, they had examined the graphical notes of their classmates and were dissatisfied with them. As one of them explained, with graphics you can show anything and you don't know if it would really work or not. Earlier in the year, they had used Interactive Physics™ in conjunction with CSILE in work on lever problems in elementary physics. Interactive Physics™ permits simulation of physical systems by assigning physical prop-

erties to simple geometric figures. The two students decided to use Interactive Physics™ to represent walls, drawbridges, portcullises, and attacking forces in ways that could actually be run as simulations to see how well they would work. Their CSILE notes referred to these simulations, which other students could access. Soon other students took up the simulation challenge, shifting the method of inquiry from graphically represented speculation to simulation constrained by laws of physics.

Obviously the availability of simulation software was essential for this methodological shift, but according to the students' own report, another critical element was dissatisfaction with the approach other students were taking. That dissatisfaction would not likely have occurred in a classroom where students had no opportunity to peruse one another's work fairly extensively. Also, the innovation would not have caught on, or would have done so only as a result of teacher endorsement, whereas in this case, the students themselves took up the new approach, some of them extending it beyond what its originators had done.

The second example illustrates more clearly the progressive character of knowledge-building that CSILE is designed to support. This incident occurred spontaneously and was not even known to the teacher until a researcher came upon it while exploring the student-produced database in CSILE. In the course of work on a biology unit, one student had entered a note reporting that sponges have three ways of reproducing. This fact caught the fancy of other students who came upon the note through database searches, and there followed a series of 12 notes and comments dealing with why nature would have contrived to provide sponges with such an array of options. Plausible conjectures were offered about the value of backup systems and the survival of a species unable to defend itself. One student, however, kept raising the question in comments to others: If three ways of reproducing are better than one, why don't other animals have them too? This is an illustration of progressive problem solving in the construction of knowledge. The solution to the first problem—why three ways?—gives rise to a higher-level problem that raises deeper issues about evolution. The answer that was finally proposed to the second question drew upon an idea that has figured prominently in evolutionary theory of recent decades: structural constraints on evolutionary possibilities. By going deeper into the study of reproduction, a student came to the insight that it is because they are structurally so simple that sponges are able to reproduce by budding and regeneration, in addition to sexual reproduction. Higher animals are too complex for this. As the student put it, "A stomach, lungs, a brain, and a heart, etc., could not grow on your finger if it was cut off."

Evaluations of CSILE to date indicate that CSILE students surpass students in ordinary classrooms on measures of depth of learning and reflection, awareness of what they have learned or need to learn, and in under-

standing of learning itself. Moreover, individual achievement, as conventionally measured, does not suffer. In fact, students do better on standardized tests in reading, language, and vocabulary (Scardamalia et al., 1992). What most impresses teachers and observers alike, however, is what the students are able to do collectively.

We do not want to suggest that the technology by itself can bring about the transformation of a school into a knowledge-building community. We already have evidence that teacher strategies can make a major difference in the extent to which students engage in collaborative knowledge-building (Bereiter & Scardamalia, 1992). Neither do we want to claim that a knowledge-building community, meeting the specifications set out previously, has actually been realized. Those specifications are ideals to work toward. The most that can be claimed is that, in the progress made to date in working toward those ideals, CSILE appears to provide a vital kind of support.

The computer technology that enables students to share knowledge with one another, as in CSILE, is rapidly being extended to give students access to the great bodies of information now being made available on the World Wide Web and on compact discs, and also access to live experts. In principle, this greatly expanded access to knowledge resources should be all to the good, but unless schools can be restructured into communities that actually work to build their own knowledge from those resources, the technology may be largely wasted.

ACKNOWLEDGMENTS

This chapter is a revised version of an article previously published in the *Journal of the Learning Sciences* by Lawrence Erlbaum Associates. The authors acknowledge the generous support of Apple Computer, Inc., External Research Division, and the James S. McDonnell Foundation. We are indebted to the students and teachers at Huron Public School who contributed their time and talents to this project, and to the entire CSILE team, without whose contributions the work reported here would not have been possible. We are particularly grateful to reviewers Bill Clancey, Paul Feltovich, and Roy Pea, whose insights and comments on previous drafts have helped us rethink, and, we hope, improve the ideas presented in this chapter.

REFERENCES

- Bereiter, C., & Scardamalia, M. (1989). Intentional learning as a goal of instruction. In L. B. Resnick (Ed.), *Knowing, learning, and instruction: Essays in honor of Robert Glaser* (pp. 361-392). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Bereiter, C., & Scardamalia, M. (1992). Two models of classroom learning using a communal database. In S. Dijkstra (Ed.), *Instructional models in computer-based learning environments* (NATO-ASI Series F: Computer and systems sciences; pp. 229-241). Berlin: Springer-Verlag.
- Bereiter, C., & Scardamalia, M. (1993). *Surpassing ourselves: An inquiry into the nature and implications of expertise*. Chicago: Open Court.
- Brown, A. L., Ashi, D., Rutherford, M., Nakagawa, K., Gordon, A., & Campione, J. C. (1993). Distributed expertise in the classroom. In G. Salomon (Ed.), *Distributed cognitions: Psychological and educational considerations* (pp. 188-228). Cambridge, UK: Cambridge University Press.
- Brown, A. L., & Campione, J. C. (1990). Communities of learning and thinking, or a context by any other name. *Contributions to Human Development*, 21, 108-126.
- Bruner, J. (1964). *The process of education*. Cambridge, MA: Harvard University Press.
- Chan, C. K. K., Burtis, P. J., Scardamalia, M., & Bereiter, C. (1992). Constructive activity in learning from text. *American Educational Research Journal*, 29(1), 97-118.
- Collins, A., Brown, J. S., & Newman, S. E. (1989). Cognitive apprenticeship: Teaching the crafts of reading, writing, and mathematics. In L. B. Resnick (Ed.), *Knowing, learning, and instruction: Essays in honor of Robert Glaser* (pp. 453-494). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Doyle, W. (1986). Classroom organization and management. In M. C. Wittrock (Ed.), *Handbook of research on teaching* (pp. 392-431). New York: Macmillan.
- Eichinger, D., Anderson, C., Palincsar, A. S., & David, Y. (1991, April). *An illustration of the roles of content knowledge, scientific argument, and social norms in collaborative problem solving*. Paper presented at the meeting of the American Educational Research Association, Chicago.
- Heap, J. L. (1985). Discourse in the production of classroom knowledge: Reading lessons. *Curriculum Inquiry*, 15(3), 245-280.
- Hewitt, J. (1996). *Progress toward a knowledge-building community*. Unpublished doctoral thesis, University of Toronto.
- Kay, A. C. (1991, September). Computers, networks and education. *Scientific American* [Special Issue], 138-148.
- Latour, B. (1987). *Science in action*. Milton Keynes, UK: Open University Press.
- Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge, England: Cambridge University Press.
- Lemke, J. (1990). *Talking science: Language, learning and values*. Norwood, NJ: Ablex.
- Lipman, M. (1988). *Philosophy goes to school*. Philadelphia: Temple University Press.
- Ng, E., & Bereiter, C. (1991). Three levels of goal orientation in learning. *Journal of the Learning Sciences*, 1(3 & 4), 243-271.
- Pea, R. D. (1993). Distributed intelligence and designs for education. In G. Salomon (Ed.), *Distributed cognitions: Psychological and educational considerations* (pp. 47-87). Cambridge, UK: Cambridge University Press.
- Peters, D. P., & Ceci, S. J. (1982). Peer-review practices of psychological journals: The fate of published articles, submitted again. *Behavioral and Brain Sciences*, 5(2), 187-255.
- Peters, T. (1987). *Thriving on chaos: Handbook for a management revolution*. New York: Knopf.
- Popper, K. R. (1962). *Conjectures and refutations: The growth of scientific knowledge*. New York: Basic Books.
- Popper, K. R. (1972). *Objective knowledge: An evolutionary approach*. Oxford, England: Clarendon.
- Resnick, M. (1996). Beyond the centralized mind-set. *Journal of the Learning Sciences*, 5, 1-22.
- Scardamalia, M., & Bereiter, C. (1991a). 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. (1991b). Literate expertise. In K. A. Ericsson & J. Smith (Eds.), *Toward a general theory of expertise: Prospects and limits* (pp. 172-194). Cambridge, UK: Cambridge University Press.
- Scardamalia, M., & Bereiter, C. (1992). An architecture for collaborative knowledge-building. In E. De Corte, M. Linn, H. Mandl, & L. Verschaffel (Eds.), *Computer-based learning environments and problem solving* (NATO-ASI Series F: Computer and Systems Science, pp. 41-46). Berlin: Springer-Verlag.

10. KNOWLEDGE-BUILDING COMMUNITIES

- Scardamalia, M., & Bereiter, C. (1993). Technologies for knowledge-building discourse. *Communications of the ACM*, 36(5), 37-41.
- Scardamalia, M., Bereiter, C., Brett, C., Burtis, P. J., Calhoun, C., & Smith Lea, N. (1992). Educational applications of a networked communal database. *Interactive Learning Environments*, 2(1), 45-71.
- Scardamalia, M., Bereiter, C., & Lamon, M. (1994). The CSILE project: Trying to bring the classroom into world 3. In K. McGilly (Ed.), *Classroom lessons: Integrating cognitive theory & classroom practice* (pp. 201-228). Cambridge, MA: MIT Press.
- Scardamalia, M., Bereiter, C., McLean, R. S., Swallow, J., & Woodruff, E. (1989). Computer supported intentional learning environments. *Journal of Educational Computing Research*, 5, 51-68.
- Whitehead, A. N. (1925/1948). *Science and the modern world*. New York: New American Library.