Knowledge Construction and Sharing in Quorum[1]

Alberto J. Cañas,* Kenneth M. Ford,* Patrick J. Hayes,†
John Brennan,* and Thomas Reichherzer*

* Institute for Human and Machine Cognition
University of West Florida, Pensacola FL 32514

† Beckman Institute
University of Illinois, Urbana IL 61801

May 1995

[1] Proceedings of the 1995 AI in Education Conference, August 1995, pp. 218-225.

Knowledge Construction and Sharing in Quorum

Alberto J. Cañas[1], Kenneth M. Ford[1], Patrick J. Hayes[2], John Brennan[1], Thomas Reichherzer[1]

[(1)]Institute for Human and Machine Cognition
University of West Florida, Pensacola FL 32514

[(2)]Beckman Institute
University of Illinois, Urbana IL 61801

E-mail: acanas@ai.uwf.edu

Abstract

The purpose of this paper is to report on a continuing research effort aimed at the design and development of large-scale international computer network connecting schools in at least ten countries throughout Latin America. Toward this end, the University of West Florida and IBM Latin America have initiated a joint partnership called "Quorum: Collaboration without Boundaries." The project is large and multifaceted. In addition to providing physical connectivity, it includes teacher training and development of the curriculum material and building advanced software tools necessary to support student collaboration across classrooms and countries. This paper is focused on Project Quorum's support for collaboration at the knowledge level across classrooms and continents.

Introduction

As part of Project Quorum, we are developing a system (provisionally called *Pangea*) to empower students to collaborate in the construction and sharing of models of their beliefs about specific domains. Project Quorum is a three-year cooperative partnership between the University of West Florida and IBM Latin America to implement a telecommunications network connecting Latin American schools participating in Project Genesis (Cañas & Schaffer, 1990; Escorcia et al., 1991), an IBM initiative that has grown to over a thousand schools spread throughout more than ten Latin American countries. Genesis distinguishes itself by a fundamental emphasis on training, preparation, and support of teachers. Building on this strong foundation, Quorum involves the creation of an infrastructure for collaboration, including the implementation of an international computer network, development of the curriculum material necessary to support collaborative work among schools, teacher-training workshops throughout the countries, coordination in the design and development of projects, access to experts in the various areas of the curriculum, the development of advanced software tools to support more extensive collaboration as new technology becomes available in the schools, and most importantly, a momentum of interaction between schools that will result in a self-nourishing project.

Collaborative learning (Ausubel, Novak & Hanesian, 1978) is an enterprise in which the learners, and perhaps their teachers, cooperatively build an explicit knowledge model which gives coherent expression to their understanding. From a constructivist perspective, the most important outcome of the modeling process is less the model itself than the insight we gain as we struggle to articulate, organize, critically evaluate and assent to it (Cañas & Ford, 1992). Likewise, the collaborative process we envision will derive its value chiefly from our success in framing the activity as a self-correcting enterprise in which learners can subject any part of the model — including their own background assumptions — to critical scrutiny. From this standpoint, the crucial question for us is "how useful is the modeling process as a means of facilitating the learner's understanding?" rather than "is the model correct?" Our research agenda, therefore, is oriented towards the development of tools and methods to aid learners and teachers to express, elaborate, share, improve, and understand their constructions of their world.

In this paper we first provide a summary of the current status of Quorum and then describe the software tools being developed in *Pangea* to support collaborative knowledge building.

Quorum Overview: Current Status

Quorum is a means for enabling students to collaborate, to help them understand their cultural and social differences, as well as their commonalities. Students collaborate on issues that are global in nature, involving not only the problems of their communities, but also those of neighboring communities and the world as a whole. We anticipate that by their participation in the Quorum Project, students are gaining an enhanced appreciation of their social responsibilities and a deeper appreciation of the world at large.

Telecommunications initiatives between schools often limit the interaction to trivial topics that have little to do with the student's learning. In Quorum, care has been taken in developing collaborative projects that take full advantage of the participation of students from different countries and that could not be carried out without telecommunications.

We want to encourage students to fully embrace telecommunications — to make it seem natural to them. We believe that to achieve this end, every student and teacher in the school must have his or her own unique e-mail address and be able to exchange messages with any other student or teacher in the network. Not all projects require that students participate individually, but they should know that they have the capability of doing so. We also want the mail tool program to be accessible to students of all ages. LogoWriter is the main software tool used throughout Genesis, together with its successor MicroMundos (MicroWorlds). In Quorum, LogoWriter and MicroMundos are also both the mail tool and the transmission medium; in this way, electronic communication becomes a part of the student's work very naturally, without the need to learn a separate, dedicated software package. Although simple to use, each tool is a very powerful mail medium. Quorum students and teachers are able to send not only text, but graphics, animation and sound, and more importantly, executable programs, without fussing with file access protocols. Most e-mail users in the world are limited to text messages. Very few, if any, e-mail systems allow users to execute the mail message they receive.

The Quorum network is divided into regions. Each region consists of the group of schools located physically close to a Quorum Server, which is usually installed at an IBM regional office. All IBM offices are linked through reliable communications using its proprietary network VNet. Schools need only to call the closest IBM office to connect to the network. A gateway has been implemented to exchange documents between Quorum and Internet users. Quorum became operational in May of 1994 and currently includes schools from Rio de Janeiro, Brazil; Caracas, Venezuela; Aguascalientes, Mexico; and Sao Paulo, Brazil. During the next two years, Project Quorum will include most of the Latin American countries. A more extensive description of the implementation can be found in Cañas (1993).

Concept Maps

Concept maps, developed in an educational setting by Novak (1977), are used as the primary language for description and communication of concepts within assimilation theory, a cognitive learning theory that has had extensive application to education (Ausubel, Novak & Hanesian, 1978). It is based on a constructivist model of human cognitive processes. In particular, assimilation theory focuses on describing how concepts are acquired and organized into a learner's cognitive structure. Ausubel argues that learning is synonymous with a change in the meaning of experience. His fundamental premise is deceptively simple:

> Meaningful learning results when new information is acquired by deliberate effort on the part of the learner to link the new information with relevant, preexisting concepts or propositions in the learner's own cognitive structure. (Ausubel et al., 1978, p. 159)

Assimilation theory stresses that meaningful learning requires that the learner's cognitive structure contain anchoring concepts to which new material can be related or linked. For this reason, Ausubel argues that "the most important single factor influencing learning is what the learner already knows. Ascertain this and teach him accordingly."

A concept map is a graphic display of concept names connected by directed arcs encoding propositions in the form of simplified sentences. The simplest concept map would be two nodes connected by an arc representing a simple sentence such as 'grass is green,' but they can also become quite intricate. Figure 1 shows a concept map about concept maps (from Novak & Gowin, 1984. By convention, links run top-down unless annotated with an arrowhead.) When concepts and linking words are carefully chosen, these maps can be useful classroom tools for observing nuances of meaning, helping students organize their thinking, and summarizing subjects of study. Mapping techniques are employed to help students "learn how to learn" by bringing to the surface cognitive structures and self-constructed knowledge (Novak & Gowin, 1984).

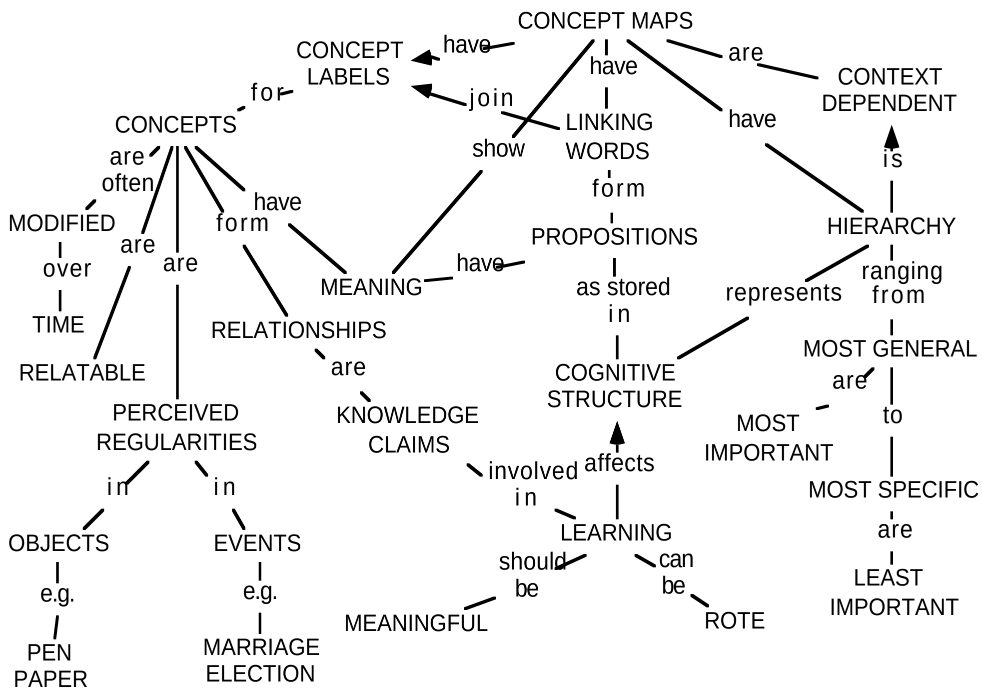

Figure 1: A concept map about concept maps

From an AI perspective, concept maps seem similar to such familiar graphical knowledge representations as semantic networks and conceptual graphs. This similarity is useful and suggestive, and we employ it in our work, but it is important to realize that concept maps are not “knowledge representations” in a computational or logical sense. They are much more loosely defined, with no firm syntactic rules and no formal rules of interpretation or semantics. Nodes in a concept map are often labelled with nouns interpretable as object names and arcs labelled with verbs interpretable as relation names, but this is by no means required. Their function is pedagogical, to help students clarify and organize their own conceptual framework, rather than formal.

There are several rules and conventions for a well-designed concept map. More general, inclusive concepts should be found at the highest levels, with progressively more specific, less inclusive concepts arranged below them. Beginners sometimes produce very “linear” maps — essentially merely decorations of long sentences — which are not well constructed. However, these conventions are not strict conditions of grammatical well-formedness, and even a poorly constructed concept map is, indeed, a concept map.

A concept map is never finished. New propositions can always be added to it. Two concept maps may also be merged to form a larger, more complex and enriched view of a subject. This process of building and merging is at the heart of the use of concept maps in Quorum.

Pangea: Collaboration at the knowledge level

Quorum affords several levels and kinds of collaboration. The simpler forms of support for collaboration include text-based mail messages and newsgroup-type discussions. Richer levels of participation include students exchanging LogoWriter pages and executable programs. For example, they can construct a simulation of the project they are studying and share that simulation with other students.

We are currently developing a collection of tools intended to take the collaboration one step further and provide a means for the students to collaborate at the “knowledge level.” (Newell 1982) This system, called *Pangea,* allows students to build concept maps and share the knowledge expressed in their concept maps with other students.

A soup with many cooks

The system has many provisions to facilitate the process of map construction. The software allows students to build concept maps individually, under the guidance of the teacher, or as part of a group map-building project. Moreover, concept maps can be sent by electronic mail. This part of the system is fully implemented, and its use the schools will begin shortly.

Pangea will also support ways in which the knowledge encoded in concept maps can be selectively distributed among a community of pupils. This requires a rather different representation of the information in the map. As noted, a concept map can be regarded as an organized collection of propositions relating together a collection of topics. Each proposition is expressed by a simplified sentence which can be extracted from the map by following a series of arcs beginning and ending at nodes. For example, the map in Figure 1 contains the sentences 'CONCEPTS have MEANING' and 'RELATIONSHIPS are KNOWLEDGE CLAIMS involved in LEARNING.' A student creates a sentence by highlighting part of the concept map, and can edit a sentence by altering the map. The resulting sentences appear in a separate window (see Figure 2). This allows two very different (although "logically" equivalent) representations of a student's ideas; one embedded in the concept map's graphical structure and the other more textual in nature.

A student may *publish* a sentence, which makes it potentially visible to other students. We call this process *making a claim.* These published sentences — claims — become part of the class "knowledge soup," which consists of a highly organized "database" of simple sentences representing the growing knowledge of the group. It is through these knowledge soups that collaboration and sharing take place. Knowledge soups have many interpretations and can be displayed in several ways. They can be thought of as a body of text, an encoding of a large concept map, or an annotated collection of discussions between students.

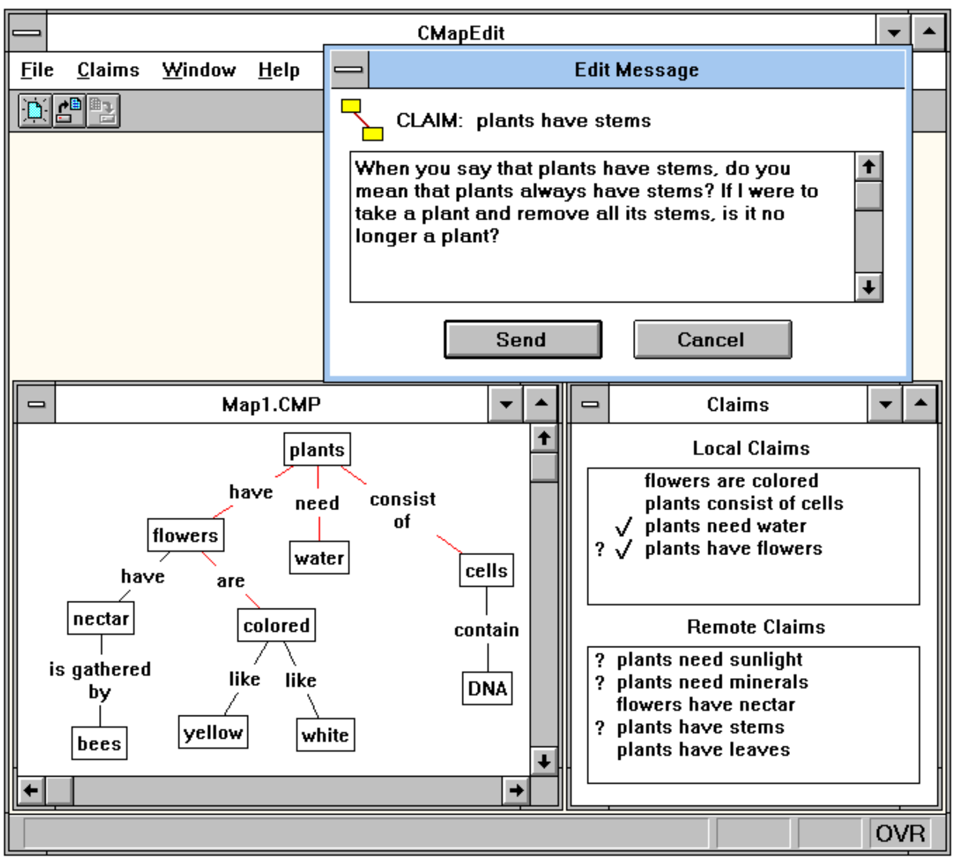

Figure 2: Concept map, claims, and threads

Challenges and dialogs — peer review, not simple authority

Published claims can be seen by other students and can be utilized in their own map-building process, but a student can't see *all* claims published by the other students, as this would often be cognitively unmanageable. The system has heuristics about the relatedness of knowledge claims. The only claims that a student sees are those directly related to the ones that he contributed to the soup. As a student publishes more, a wider range of other related claims becomes visible. This strategy is intended to encourage and reward students for participation. The source of a claim remains anonymous.

A student can query a claim submitted by another student, if he disagrees with it or finds it puzzling, and the originator of the claim can respond. Querying a claim causes it to be displayed with a question mark after it to indicate to the author and any third parties that it is under dispute or discussion. The querier types a message which becomes invisibly attached to the query mark. Anyone — including, of course, the originator of the claim — can read this message and respond to it with a further comment, or an explanation or defense of the original claim. In this way, a published claim may become the locus of an extended discussion on some topic. The student's own claims are likewise subject to peer review.

The system keeps track of the authorship and time of all this publication, but it is not made public to the students. We judged that the advantages of anonymity in encouraging less socially confident students to risk publication and discouraging "flaming" outweighed the advantages of identifiability in encouraging students' pride in their work. (This strategy could be easily modified to allow signatures, identity-masking pseudonyms, etc., or even allow the teacher to decide which mode was best.)

The way in which electronic communication is a "leveler" of social hierarchies is well known, and we expect to be able to take powerful advantage of this phenomenon. This mode of electronic communication and debate, leading to "threads" of related comments from several individuals, is an evolution of the widespread practice in CompuServe, Usenet, etc. This sharing of claims is similar to the interaction that takes place under CSILE (Scardamalia et al., 1992; Scardamalia & Bereiter, 1993). However, a key difference is that the discussion threads are attached to specific knowledge claims; thus, each student will see only dialogs that are related to claims they themselves have published — and hence that they can contribute to.

Another consequence of anonymity is that this is not a communication channel. Students cannot send personal messages to one another. The content, rather than the source, determines which parts of the soup are displayed to each student. In this way, we hope that Pangea will establish a new kind of group interaction amounting to a collaboration between students entirely in terms of what they know, a collaboration at the knowledge level.

Soups from distant lands

In addition to supporting knowledge level collaboration within a single classroom, knowledge soups can be made public for access by classrooms in other countries or continents. A teacher can import a knowledge soup from another classroom (or school) for his students to interact with. The same rich level of interaction supported for local soups is also provided for remote soups. A teacher in Caracas and another in Brazil, for example, could swap knowledge soups about the rain forest (Figure 3). Likewise, teachers in several geographically dispersed schools could agree to have their students collaborate while developing concept maps on the same specific topic. The overall result will be a "kettle" of knowledge soups, collaboratively developed by students, and available to other students for their use in their knowledge construction.

An "artificial idiot"

We plan to make available an optional third mode of interaction, in which the system takes on the role of a friendly, helpful agent with whom the child may converse in a very simple way (see Figure 4). This agent will have access to all the information in the soup and is equipped with some heuristics for drawing plausible conclusions and asking questions.

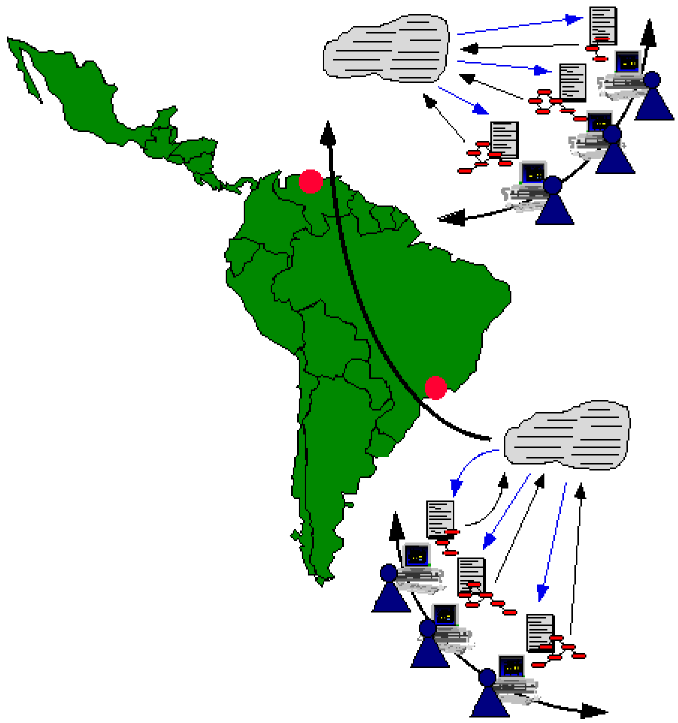

Figure 3: A Knowledge Soup being shared by schools in different countries

The system will use information about a restricted sub-vocabulary of arc labels and some easily obtained grammatical information to draw tentative conclusions about relationships between the meanings of some of the assertions made in the student's growing cognitive map. These conclusions might use assertions taken from many maps. If Juan has previously claimed that 'fish have wings,' and Marta claims that 'fish can-be edible,' then the system might notice the connection, tentatively draw a (false) conclusion, and ask her if anything with wings can be edible. We hope that the friendly idiot's questions might often be stimulating to the learner's imagination, while acting in effect as an additional channel of communication between learners' understandings of the subject matter. Notice how different this system is from the overtly "debating" flavor of the assert-query-respond game in the threads associated with knowledge claims. We believe that it represents a new kind of interaction, in which a learner's ideas are contrasted or reinforced by those of other students without anything like communication, even anonymous communication, occurring between them.

We have chosen the personality of the agent with care, based in part on the recommendations of Novak and Gowin, (1984) for how teachers should interview students about their concept maps. It will be a kind of idiot savant who is nonjudgmental, nonintrusive, friendly, and hopefully amusing. (For example, we are developing a simple graphical interface using a cartoon face to convey puzzlement, interest, happiness, etc.) It uses the student's own language as much as possible and always tries to be helpful. It always leaves the student alone if told to. It knows a lot but sometimes reveals surprising ignorance. It never takes offense at being told it is a fool. In fact, it may ask the student for help or clarification when puzzled and always enjoys being taught new things by the student. The student may mislead it, of course, but to do so deliberately would be irresponsible, and such behavior is easily detectable since all transactions with the agent are internally recorded.

This agent gives a persona to the growing body of information concerning the domain of study. Just as this body grows as the class project advances, the knowledgeable fool can be expected to become more knowledgeable and his opinion more reliable as time advances. However, he should never be considered to be an authority figure, as he will almost certainly make silly mistakes from time to time no matter how much he comes to know. His role is not that of a Socratic source of wisdom, but that of a friendly colleague designed to encourage the students to reflect upon their knowledge claims.

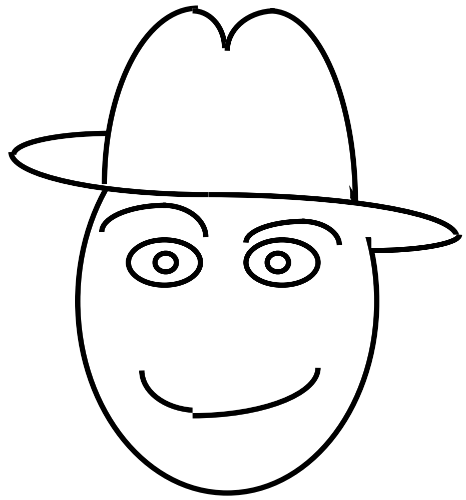

Figure 4: The Artificial Idiot's image can be tailored by the student

The system has only a very weak grasp of the what the various claims mean and does not have most of the everyday knowledge which humans use to help them interpret simple sentences. We cannot afford to put a large-scale knowledge base of commonsense information into the system, and in any case available evidence suggests strongly that this would not work. The conclusions it could support would almost certainly be too fragile for the often rather imprecisely expressed claims, and too idiosyncratic for the wide range of topics with which the system must be able to function. Rather, we plan to look for rather general patterns in the groups' own knowledge (as represented in the growing knowledge soup) which suggest gaps, misunderstandings, or inconsistencies in an individual student's claims.

The system will make inferences from the answers it receives, and it can sometimes check its tentative conclusions by simply asking the learner. This is the sense in which we say that the system is always willing to learn from the student. (The answers can of course be confirmed by going directly to a more authoritative source such as the teacher.) The agent can also sometimes learn a great deal by being told that its question was silly. For example, suppose Marta replies that it is silly — as opposed to simply wrong or right — to ask whether all things with wings are edible, when prompted by the claim that fish are edible. The system can infer that something is suspicious about Juan's assertion concerning winged fish, information which might be otherwise difficult to obtain.

We are investigating several ways in which the system can put together information from several claims to draw tentative conclusions about the coherence or completeness of one learner's partial concept map. These ways involve the ideas of similarity, possibility, causal relations between types of events, and 'using' (for example, 'Plants use Sunlight for Photosynthesis'). We find that many of the arc labels used in classroom concept maps fall into one of these categories. Another key observation is that in a well-constructed concept map, the nodes are usually labelled with nouns, and the arcs with verbs or supporting vocabulary such as prepositions, and we expect to be able to recognize this distinction automatically. (The system currently uses it to initially translate maps to sentences.)

The system understands the basic notion of a concept hierarchy and recognizes some words which are commonly used as arc labels to express hierarchical relationships. These include 'is,' 'are,' 'is-usually,' 'may-be' and 'e.g.' It also understands the notion of a classification, in which one concept is divided into subclasses according to some property. The arc labels 'are-usually,' 'can-be,' and 'some-are' often indicate such a classification, so that we might find the claims 'Leaves can-be Deciduous,' 'Leaves can-be Evergreen,' etc.

One common pattern is a series of claims of the form 'A can-be B1' with an associated hierarchical series 'B1 is-a C,' which names the classification. (For example, 'Cows can-be White,' 'Cows can-be Black,' etc.; 'White is-a Color,' 'Black is-a Color,' etc. Notice however, 'Cows can-be Thin' makes sense, but 'Thin is-a Color' is a conceptual error.) The system can draw several useful tentative conclusions even from only a fragment of such information. For example, a single 'A can-be B' claim indicates fairly clearly that some A's are not B. This obsrevation might prompt a question "Which A's cannot be B?" for example, and a positive answer (or a suitable adjustment to the concept map) of the form ‘A can-be D’ would support the tentative conclusion that B and D are of the same classificatory type. If learners have made several such claims, a question like “What are B, D, ... all called?” asking directly for the name of the type, can provide useful information. Once the type name is known, the system can ask directly whether a new concept appearing in a claim like ‘A can-be X’ is one of that type. This could have led to the question ‘Is Thin a Color?’ in the example above, and again, a negative answer allows the system to draw useful conclusions.

For example, a concept map containing any of ‘White can-be Color,’ ‘White is Color,’ or ‘White are Cows’ can now be immediately recognized as faulty; and more interestingly, a repair can be suggested (“Do you mean that white is a color, and some cows are white?” — [yes] — “You could say: ‘White is-a Color’ and ‘Cows can-be White’.” More subtly, a concept map which contains the arcs ‘Cows can-be White’ and ‘Cows can-be Thin’, but no other arcs of the form ‘Cows can-be...’, can now also be recognized as possibly indicating a conceptual confusion in the mind of the learner, and suitable questions could be designed to probe it, or the situation could simply be recorded for the teacher’s benefit.

We should emphasize that much of the information in the soup will be meaningless to the agent. Failures which might be unacceptable for a natural-language understanding system, such as being unable to resolve ambiguities or to notice alternative lexical usages, are not a disaster here. On the other hand, new heuristic rules can be smoothly incorporated into the agent’s repertoire of inferential abilities without affecting its overall classroom situation. We expect to treat this part of the system as an on going experiment.

This part of the system is still under development and has not yet been implemented, but we are excited by its potential. For example, while the “fool” is in many ways very limited in its abilities, it might access information in distant knowledge soups to bring ideas from other cultures, saying things like: “In Peru I heard that people eat guinea pigs.”

Summary

Quorum is not a project that will vanish after three years. Our objective is to establish an infrastructure that includes hardware and software connectivity, technical and pedagogical support, and a momentum of interaction between schools that results in a self-perpetuating project which continues to grow. The continuing fundamental theme is to foster communication and collaboration at many levels, from simple exchange of electronic messages to collaboration across classrooms and continents at the knowledge level. We anticipate something which has never happened before. Whole generations of people in more than ten countries will be able to understand and empathize with one another in a way hitherto impossible, because they will all have gone to school together.

References

Ausubel, D. P., Novak, J. D. & H. Hanesian, H. (1978). *Educational Psychology: A Cognitive View* (2nd ed.). New York: Holt, Rinehart& Winston. Reprinted, 1986. New York: Warbel & Peck.

Cañas, A. J. (1993). *Niños Colaborando a Travis de Fronteras: Logo y Telecomunicaciones*, Proceedings of the VI International Logo Congress, Caracas, Venezuela.

Cañas, A. J., Schaffer, M. (1990). *Project Genesis: Technology and Innovation in Education*, First International Seminar: the Implementation of Computers in Education, International Society for Technology in Education (ISTE), Guatemala City, Guatemala.

Cañas, A. J. & Ford, K. (1992). An Environment for Collaborative Knowledge Building, presented at the *Workshop on the Technology and Pedagogy for Collaborative Problem Solving as a Context for Learning*, Toronto.

Escorcia, G., Cañas, A. J., & Schaffer, M. (1991). *The Genesis Concept*, Second International Seminar: the Implementation of Computers in Education, International Society for Technology in Education (ISTE), Mexico City, Mexico.

Newell, A. (1992) The Knowledge Level. *Artificial Intelligence* **18**, 87-127

Novak, J. D. (1977). A Theory of Education. Ithaca, NY: Cornell University Press.

Novak, J. D. & Gowin, D. B. (1984). Learning How to Learn. New York: Cambridge University Press.

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: 45-71.

Scardamalia, M. & Bereiter, C. (1993). Technologies for Knowledge-Building Discourse, *CACM* 36: 37-41.