PART II. ANALYSIS OF COLLABORATIVE KNOWLEDGE BUILDING

Introduction to Part II: Studies of Interaction Analysis

Retro-perspective on the Studies in Part I

Looking back on the preceding historical documents—little modified from several years ago when the reported investigations were just winding down—from the perspective of 2004 as I compile this book, the contributions of the individual studies to the overall inquiry becomes more clear. I will reflect from this hindsight briefly here as a way of motivating part II.

Study 1, TCA. This effort always seemed ahead of its time. It failed to attract continued funding because while the funding source approved of its goals, the reviewers were worried how the quality of the curriculum content and its indexing could be guaranteed. Similarly, although our contacts at Apple Corporation liked the concept, they could not see how such software would help their quarterly financial bottom line in the near future. In the intervening decade, the NSF has spent millions on a major national science digital library initiative (NSDL) and Apple tried to build its own educational object economy (EOE) online repository of curricular artifacts. Today, the Math Forum, one of NSDL's most successful digital libraries, consists primarily of content submitted by users, demonstrating the power of community knowledge construction. Large international efforts have gone into trying to define and negotiate standard metadata ontologies (e.g., Dublin Core); they remain incomplete, inadequate and superficial because systems of interpretive categories cannot be legislated, once and for all, from on high; they must evolve with usage and understanding within specific communities. Ten years of technological advance, the pervasive growth of the Web and the establishment of digital libraries have changed the way TCA would work today. But they have not eliminated the basic needs that TCA was designed to address. Publishers and teams with NSF funding have meanwhile produced integrated curricula along constructivist principles, but still-isolated teachers have benefited little from the collaboration potential of the Internet. Some of the ideas of TCA have yet to be tried by groups of teachers and to be evolved in response to their usage. In general, the idea of virtual communities of teachers sharing best practices has not really taken off yet, despite various efforts like Tapped-In, Merlot or ENC.org. Perhaps it is no coincidence that I am now working with the Math Forum and my TCA collaborator, Tamara Sumner, is working with DLESE, another major educational digital library project, trying to increase their utility for constructivist learning.

Study 2, Essence. The use of LSA (latent semantic analysis)—explored in this study—is still being pushed as both a panacea for automating student evaluation and as a model of human semantic understanding. This book problematizes the assumptions underlying LSA by discussing many means by which people understand language in non-algorithmic ways: through their life experience situated in specific activities within a meaningful physical, cultural and historical world; by means of the collaborative negotiation of shared meaning in small groups and communities; through personal, professional and cultural perspectives; and thanks to subtle processes of contextualized interpretation. The success of LSA in `State the Essence` is due to the fine tuning that took place as the algorithms, coefficients and mechanisms were co-adapted with the research effort, the teacher presentations and the student expectations or behaviors. This suggests that it is possible to take advantage of this technology, but only in certain carefully designed applications and through extensive trial and adjustment. The `Essence` software was originally conceived as a way of evaluating the product of individual work, as LSA is usually applied. However, the study suggests also using such software to stimulate collaborative interactions and to provide feedback and motivation to small groups. This is an approach that deserves further thoughtful exploration in the context of CSCL.

Study 3, `CREW`. The problem of group formation is one of the first issues I face each term when starting a course organized around small-group projects. Supporting the self-organization of students into effective work groups is a function that is particularly needed in virtual communities, but has not been extensively researched. This has been identified as an important area for exploration in the Virtual Math Teams (VMT) project that is now underway at Drexel, but is mostly beyond the scope of the present book (see chapters 17 and 21). The approach reported in the `CREW` study is quite different. It attempts to aid administrators in the formation of astronaut crew groupings by giving them feedback on the probability that given individuals will perform well psychologically under given mission conditions. It develops a temporal model of individual factors based on case study data under analogous conditions, using a combination of adapted AI and statistics methods. Pushing this approach to its limits, it shows the enormous requirements such a system has for high quality data across the whole range of interest. Given that little relevant data currently exists and the difficulty of evaluating its status, as well as the practical and political barriers to collecting much more data, it seems doubtful that this sort of approach can succeed in many realistic situations. Without adequate data, such AI methods are empty promises: garbage in / garbage out. In addition, the expectation that relevant data can easily be collected in an explicit and context-free manner was shattered with the failure of the expert system craze.

Study 4, `Hermes`. This system suffered the problem of all DODEs. The effort required to configure a system for a particular application domain is enormous. One must be an expert in that domain, understand the detailed work-flow, spot the functions that can usefully be supported and seed the system with vast amounts of domain knowledge. All this is necessary before anyone would even consider trying out the system. In order to support the work of a group of domain experts, the system must be complex and sophisticated, combining advanced features as well as all the basics. To develop such a system requires the combined talents of software developers and domain experts, along with a budget on the order of a million dollars—just to produce a system that can be tested by a small community of friendly users. There are engineering and medical applications where such an effort might be financially practical, but it does not seem to be a workable approach in the grossly under-funded education arena.

Study 5, `CIE`. As the studies of part I progress up through this one, there is a growing awareness that the important information is not just domain-specific and gradually evolving, but is specific to each community of practice and is constantly changing and being re-negotiated. The studies from this point on try to respond to this finding. The idea of supporting communities of practice with specially tailored computational media seems plausible. The issue then is one of attracting a whole community to a particular platform and getting them in the habit of using that system for their regular community participation.

Study 6, WebGuide. This study confronts the crisis of adoption. Designers can go to great lengths to design systems to provide wonderful tools, but people—even students studying the design and use of such tools—resist using them. WebGuide provided a sophisticated system of inter-connected perspectival views on an asynchronous discourse. The intention was to support on-going knowledge building. But it was under-used and at best served to exchange personal opinions. It failed to merge ideas of different people together into effective group cognition.

Study 7, Synergeia. This system added many features, based on a review of what typical CSCL and CSCW systems offered. In particular, it supported group negotiation as well as perspectives and the definition of groups, providing a structure for the interaction of ideas at the individual, small-group and community levels. As necessary or useful as such features are or could be, the proliferation of features is not sufficient to overcome the barriers to adoption confronted by all collaboration systems. To design more effective media (the goal of part I), we will need better models of computer-mediated collaboration, clearer conceptions of group negotiation, and detailed studies of small-group interactions (as is illustratively undertaken in part II).

Study 8, BSCL. The BSCL study carefully conceptualized negotiation based on current understandings of collaboration. The failure of the BSCL negotiation mechanism to be used as intended showed the need for more detailed analysis of how people actually collaborate and negotiate in normal life. Parts II and III will therefore empirically and theoretically investigate how knowledge is actually constructed and negotiated in small-group interactions.

Assessing the Studies in Part I

How should one evaluate the success of these studies? Each provided a valuable learning experience in the design of groupware. Some never got much past the conceptual design phase—perhaps producing a detailed scenario, a set of interface designs or a limited working prototype—while others have survived in one form or another.

TCA ended before ever being tried by teachers, but its designers are now deeply involved with major digital library projects that carry on much of that vision. Several versions of Essence were tested in classrooms for two years, eventually demonstrating statistically significant improvement in learning outcomes in controlled experiments reported elsewhere. The technology refined in Essence is now used in scoring various national tests. Essence itself is now being used in dozens of schools in the state of Colorado. The CREW software was turned over to NASA and its fate is not publicly known. The Hermes software was further developed and used to deploy NASA's outer-space design rationale manuals as an online hypertext system.

Versions of WebGuide were used in classes as reported in the study, although the final implemented version was never actually deployed. Its perspectives concept reappeared in Synergeia and BSCL, simplified and integrated with negotiation support. Another version of personal, group and class perspectives on threaded discussion appeared in the Polaris system from the University of Maastricht. I have tried to put all these perspective systems to good effect in my classes, with little success. I now use simple html websites for student, group and class repositories, with no computational support. We simply do not know how to design more sophisticated systems that people will really use to support group cognition.

As other researchers have also discovered, threaded discussion and chat systems, as their names suggest, are generally used for the relatively superficial exchange of opinions rather than deep, interactive knowledge building. To the extent that the systems presented here were designed to support group cognition (knowledge building, situated interpretation, intertwining of perspectives, knowledge negotiation, etc.), this shows how far we still have to go. The fostering of group cognition is a socio-technical problem, which is not automatically solved by offering certain functionality in a technical system (Kling & Courtright, 2004). It will require designing whole activity systems or shared worlds around such systems, based on a detailed analysis and understanding of collaboration.

Two releases of Synergeia were fielded in European schools as reported. Extensive surveying of teachers and students who used it showed that they liked it and had no fundamental criticisms of it. However, this form of evaluation provided little guidance for further software development. In particular, it is not clear that the negotiation mechanism was even used in the schools. I used BSCL in two of my Drexel courses on human-computer interaction and had my students design extensions to the negotiation mechanism. In this process, I experienced again how hard it is to adopt the use of that mechanism in an effective manner. The BSCL version of the software has now been integrated as an option of the popular BSCW collaboration system and is being used in many European classrooms.

Groupware is hard to assess. To see how it really supports groups, one must have groups use it under relatively naturalistic conditions and for a long-enough time to become comfortable with it. But this requires not only building a sufficiently complete and robust prototype for group usage, but also finding an appropriate group of users, training them in its use and involving them in a meaningful application of the software that properly exercises the functionality of interest.

Getting a group to use software is not easy, even once all the preconditions have been established. I found this repeatedly when using groupware prototypes in my classes. Users resist. Information science majors who are interested in the design of innovative software, students whose grades depend on entering comments by using the software, and participants in courses whose activities have been designed around the use of the software, are all reluctant to use the software, and they constantly look for more familiar alternatives: meetings, conference calls, email, instant messaging, etc. Adoption becomes *the* issue. It dominates over all the technical issues of groupware design. As one experiences the studies of this part, it becomes clearer and clearer that the problem is to design socio-technical systems, where the technological product is simply an artifact to mediate the important, complex and poorly understood processes of group collaboration. That brings us to the need to increase our understanding of the social-systems aspects of groupware design through analysis and theory of small-group interaction.

The preceding eight groupware design studies thus supply a sense of the potentials, the issues and the challenges inherent in the design of collaboration technology. They provide a shared experience to motivate and lead into the parts that follow. Like all case studies, the experiences they offer are limited by their specifics; they should not be looked to for conclusions concerning the effectiveness of their innovations. They may, however, be legitimate and worthwhile explorations of what is possible through investigation of what actually happened under unique and irreproducible conditions. Taken in this sense, the experiences of part I furnish useful occasions for the situated interpretation of what it means today to support collaborative knowledge building and group cognition.

Theoretical Background to Part II

In developing the studies of part II, two analytic perspectives played a major role: socio-cultural psychology and communication analysis. I actively pursued an understanding of them in order to resolve some of the mysteries that arose in my earlier software studies.

Socio-Cultural Psychology. Vygotsky's thinking had an immediate catalytic effect on me when I first read his *Mind in Society* (1930/1978). I was excited by his deep and original appropriation of Hegel and Marx, and by his materialist theory of mind. I was intrigued not so much by what he actually explored in his experiments and what is generally interpreted as a psychology still centered on the individual mind, but by the vision he sketched, often between the lines, of a truly socially-constructed mind, whose consciousness is derivative of the culture in which it was constructed. While my reading of Vygotsky is explicated more in part III, his emphasis on the role of artifacts in mediated cognition is already central to part II. In particular, these studies pursue the question of how people come to understand the meaning or affordances of artifacts and what implications this has for the design of groupware conceptualized as a mediating artifact.

Communication Analysis. In my search to understand perspectives and negotiation, I turned to communication analysis. This choice was obviously also compatible with Vygotsky's emphasis on language and interaction. Colleagues, methods and ideas from the discipline of communication made possible the analyses of this part, particularly chapter 12. The most relevant work for me was that of ethnomethodology (Garfinkel, 1967) and conversation analysis (Sacks, 1992). In general, I think that interactionist theories of communication have led the way in understanding the philosophical and methodological issues that are essential for developing a theoretical framework, empirical analysis and software support design practice for collaboration.

The Studies in Part II

These five essays try to analyze the nature of small-group interaction as it actually occurs. They:

- propose a model of small-group knowledge building. (chapter 9)
- critique prevalent CSCL methodologies for systematically ignoring the group interactions. (chapter 10)

- suggest new approaches that focus on the group discourse. (chapter 11)
- conduct micro-analyses of small-group interaction, detailing a group decision-making negotiation and looking at its cognitive ramifications for group understanding and activity. (chapters 12 and 13)

Chapter 9, a model. This study started as a tentative working paper when it first occurred to me that we needed to have some kind of graphical representation of the important process called *knowledge building* (Scardamalia & Bereiter, 1996) that I had started to refer to constantly, but that seemed only vaguely defined. Gradually, as I circulated the paper for comment, shortcomings of the model became apparent, both in terms of its representation and its content. Nevertheless, the diagram has endured with only minor modifications and continues to prove useful. This paper—with its graphical model—has always been one of my most popular and suggestive writings because it starts to articulate what goes into collaborative knowledge building.

Chapter 10, rediscovery. When asked to look at a couple of representative CSCL papers, I began to question the adequacy of available analytic methods. In particular, I bumped into the old lamppost problem: people tend to search where it is easiest to see and measure things, even if the important things lie elsewhere. It struck me that some of the most essential phenomena of computer-mediated collaboration were being systematically eliminated by the very methodological procedures that were recommended for rigorously analyzing them.

Chapter 11, contributions. At the CSCL '02 conference, which I viewed as an occasion for injecting a more theoretical perspective into the field, I proposed a set of four notions that could contribute to a deeper understanding of collaboration: knowledge building, group and personal perspectives, mediation by artifacts, and conversation analysis. By pulling together these four themes, the paper effected a transition from the design issues of knowledge-building support, perspectives mechanisms and software artifacts to the micro-analysis of collaborative interactions.

Chapter 12, a moment. This study looks closely at a transcript from an intense half-minute interaction among five students involved in an activity with `SimRocket`, a computer simulation of model rockets. A thick description of this hard-to-interpret discussion shows how the small group constructed group knowledge, which each participant came to share. The phenomenon of group cognition appeared here, where the indexical, elliptical and projective character of the utterances showed that their meaning only existed at the small-group unit of analysis, not as something attributable to individual cognition. This study provides a pivotal point for the book. Its transcript is repeatedly referred to in the subsequent chapters. In fact, much of the theory presented in part III is derived from this 30 second episode, illustrating how much can be learned from detailed reflection on a brief case study.

Chapter 13, references. The preceding study's analysis is expanded here to dissect the nature of the group cognition that took place around the `SimRocket` artifact. It is argued that, before the collaborative moment, the group could not see the structure of the `SimRocket` list of rocket characteristics, but that through their interaction they learned to see the new kind of structure and taught each other to see it. This group conceptual change allowed the group to repair the breakdown in relational references of their utterances to the artifact. This incident provides a key case study for the theoretical reflections on group cognition in part III.