

Information Ecology of Collaborations in Educational Settings: Influence of Tool

Mark Guzdial

EduTech Institute and GVU Center, College of Computing, Georgia Institute of Technology

Abstract

The information ecology perspective [2] helps to understand information spaces in terms of the creation, searching, and use (consumption) of information. An information ecology perspective of Computer-Supported Collaborative Learning (CSCL) environments describes the flow of information into the ecology (who writes notes? how many and when?) and the use of that information (who reads? how many and when?). From that perspective, use of two CSCL tools is compared to note similarities (e.g., notes written per student) and dissimilarities (e.g., thread length).

I. An Information Ecology Perspective of CSCL

A useful perspective on studying computer-supported collaborative learning is analysis at a high level of aggregation: Multiple group or whole class discussion forums, such as studying an entire CSILE knowledgebase [1] or the newsgroup of an entire class. (I refer to the entire discussion space for a class generically as a *forum*..) The questions at this level are about the behavior of all the participants in the forum. When do students read notes? When do they write notes? What is the level of participation in the class? Does the kind of computer technology and how it is used impact reading and writing behaviors of students?

At such a high level of aggregation it is difficult to make statements about content of notes. We really cannot even determine much about what students are learning and whether they are learning. However, there are benefits to analysis of aggregate behaviors in CSCL forums:

- **Mediating behaviors.** While we cannot determine if individuals are learning or not at the aggregate level, we can determine whether some mediating conditions for learning in a

collaborative setting are being met. For example, if a small percentage of students is writing all the notes in the forum, we may suspect that not all students are using the opportunity to articulate their positions and have them reviewed by others. Or, if the forum can be characterized as simply question-and-answer exchanges, in-depth analysis and discussion may not be taking place. By looking at how a whole class reads and writes notes, we can learn about the kind of aggregate behaviors that suggest a successful collaborative learning situation or suggest that there are problems with the learning setting.

- **Informing designers.** Designers of tools for computer-supported collaborative learning can use aggregate behaviors to inform designs and for high-level checks on their designs. Through looking at aggregate behaviors in a variety of settings, we can inform designers about the kinds of aggregate behaviors that can be affected by tool design and which are affected more by kind of use, age of students, or other variables. Further, it is simpler to conduct an aggregate analysis than an in-depth content analysis of a discussion, so an aggregate analysis can provide early, rough estimates on how well use of a new tool is proceeding. For example, if students in a class of 100 only write about two notes each in the first two weeks of a class using a new tool, it would be useful for the designer to know if this is within the range of normal behavior (based on the analysis in this paper, we suspect that it is) or if there is a severe problem that needs immediate attention.
- **Measuring actual practice.** Looking at patterns of use across multiple classes can inform us about what actual classes do, as opposed to carefully controlled experimental groups or

classes. Taking stock of current practice helps us understand existing problems and how much of good practice has infiltrated the average class.

Stuart Card and his colleagues at Xerox PARC refer to aggregate behavior within an information space such as the World Wide Web as an “information ecology” [2]. We participants in an information ecology (referred to as “informavores” in Card’s paper) are producers, gatherers, and consumers of information. Research in information ecology are developing models of the WWW, for example, that describe when pages are created or deleted, and when they are accessed [3]. By studying the rules of behavior and the relationships between variables in the information ecology, we can learn better how to maximize the ecology (i.e., achieve more information at lower cost).

Collaboration forums are a kind of information space. An understanding of student reading and writing behaviors from an information ecologies perspective may help us better understand and better design CSCL environments. Papers that present new tools for CSCL often do present information ecology statistics in describing use of the tools (e.g., MFK/Speakeasy [4], CoNote [5], CoVis [6]). In this paper, I describe the information ecology for forums used in educational settings in two different collaboration tools, and then contrast these results with those for other systems. The result provides a picture of the information ecology of CSCL forums, and in particular, which features seem consistent across tools and which are tool dependent.

The results in this paper are based on analysis of 35 collaborative forums, which totals over 7000 notes in the collaboration spaces written by 1300 students, teachers, and teaching assistants. These forums are split almost equally between two different kinds of CSCL tools: CaMILE [7, 8] and newsgroups [9]. I also present reading behavior analysis of one forum, and then contrast the findings with those for other information ecologies.

II. Data and Methods

The CSCL forums used in these analyses come from two different sources: class newsgroups used in the College of Computing at Georgia Tech and CaMILE class discussions from a variety of different academic units at Georgia Tech. 17 CaMILE discussions and 18 class newsgroups were analyzed. I first describe the two different kinds of collaboration tools (summarized

in Table 1), and then describe how the data sets were selected and analyzed.

Newsgroups: Newsgroups are an old form of asynchronous collaboration support on the Internet. The newsgroup is distributed across multiple machines, which means that access is improved but is difficult to track. Notes are *threaded*—the newsgroup protocol tracks which notes were composed in response to other notes.

Users read newsgroup messages using one of many available newsgroup readers. The interface and even the modality of messages depends on the newsgroup reader used by an individual. Most newsgroup readers, by default, show a note only once—unless the participant makes an explicit effort, a viewed note will not be shown ever again. The lack of persistence may make a difference in sustaining discussion—if a note is not commented upon immediately, it may be difficult to retrieve for later comment or review. At Georgia Tech, all College of Computing classes have an associated newsgroup for use by students to discuss the class, ask and answer questions, and perhaps interact with the class teacher or teaching assistants.

CaMILE: CaMILE (Collaborative and Multimedia Interactive Learning Environment) is an asynchronous collaboration support designed by me and my colleagues in the EduTech Institute at Georgia Tech. CaMILE is a Web-based application, where all access is through a Web browser accessing a single server. The interface is forms-based. CaMILE discussions are also threaded, as in a newsgroup. Unlike a newsgroup, CaMILE threads are persistent—they are always available to users and do not disappear after viewing. CaMILE notes can contain anything that a Web page can contain. In one forum, approximately 30% of all notes contained some kind of HTML tag (e.g., links out from the note, embedded images, etc.) [10].

An important distinction between newsgroups and CaMILE is that CaMILE supports *anchored collaboration*. [10, 11]. Each individual note can be referenced uniquely through a Web browser. Direct addressing of notes allows for the creation of Web pages that can contain single-click hyperlinks (say, a report to discuss) to a thread of discussion (a collaboration space) related to the given Web page. Anchors serve as indices (e.g., all the notes related to a given assignment are in the thread of notes accessed from the assignment Web page) and as reminders of what students are to talk about in a given thread. Typically, teachers create the anchors.

	Newsgroups	CaMILE
General structure	Threaded notes in an asynchronous forum	Threaded notes in an asynchronous forum
Searching	Newsreader dependent	None
Indexing	None	Index through anchors
Persistence of Notes	Newsreader dependent, but default is not persistent	Persistent
Use of multiple media	Newsreader dependent, but not typical	In anchors and notes
Location of notes	Distributed	Centralized

Table 1: Describing and Contrasting Newsgroup and CaMILE Collaboration Tools

Selection of DataSets: Data sets were selected to emphasize larger classes (where more forum activity may occur) and a predominantly undergraduate population. More undergraduate than graduate forums were available, and I predicted (but did not test) that use would differ between undergraduates and graduates. Overall, there were 7262 notes analyzed, with 1300 authors. There were 3007 CaMILE notes by 526 authors, and 4255 Newsgroup notes by 774 authors.

The class forums are described in summary in Table 2. I used 17 CaMILE undergraduate class discussions from over the two years of use (eliminating four graduate classes). The units represented are Computer Science (CS), Chemical Engineering (CHE), English (ENGL), History

(HIST), and Literature, Culture, and Communication (LCC). Some of the CaMILE use was very sparse (e.g., two to four notes in the entire quarter in two classes) and some very narrow distribution of authors (e.g., one author out of a class of 31). These forums were still included in the analysis, as part of the broad range of use which might be expected with a new tool. Since not all academic units at Georgia Tech provide course newsgroups to every course, I used 18 Computer Science undergraduate course newsgroups at Sophomore-level or above, to be sure that the audience was familiar with newsgroups (from first year CS course newsgroup use) and were at the same academic level (if not same unit) as the CaMILE users.

CaMILE Classes	# Notes	# Authors	# In Class	Newsgroup Classes	# Notes	# Authors	# In Class
CS2390 f96	409	61	81	2360 Sp	446	59	81
CS2390 sp96	464	65	79	2360 Wi	1110	103	75
CS2390 w96	487	57	79	2430 Sp	587	83	92
CS2390 w97	503	60	80	2430 Wi	536	98	89
CS2390 sp97	452	109	92	2760 Sp	159	45	61
CS4345 w97	35	15	30	2760 Wi	108	54	57
CS6397 w97	141	23	32	3156 Sp	40	20	51
CS6398 sp96	15	7	16	3156 Wi	159	54	49
CHE2208 sp97	13	1	31	3158 Sp	62	16	50
CHE2210 sum96	71	16	40	3158 Wi	26	9	44
CHE2210 win97	103	18	66	3302 Sp	14	6	50
CHE4803 win96	42	9	20	3302 Wi	88	27	47
ENGL1002e sp97	75	29	35	3361 Sp	186	37	49
ENGL1002l sp97	76	28	37	3361 Wi	233	45	47
HIST3043 sp97	4	3	40	3411 Sp	214	43	49
LCC4875 f96	115	23	24	3411 Wi	204	44	50
LCC6607 f96	2	2	4	3431 Sp	79	28	60
				3431 Wi	4	3	45

Table 2: Summary statistics for CaMILE and Newsgroup-using dataset classes

How Used: In general, I can make few assumptions about how the forum was used in the class. Use was not required in any of these classes¹. The main purpose was question asking and answering. Assignments were presented to students either on paper or on the Web, but not in the collaborative forums themselves. We might also assume that CaMILE-using teachers, since they sought out use of a new tool, were more interested in collaboration in the classes and may have encouraged its use more (perhaps subtly or implicitly).

Analysis Methods: Analysis focused on writing (information-producing) behavior and reading (information-consuming) behavior. Writing behavior analysis looked at the entire dataset. Reading behavior, however, only looked at the CaMILE CS2390 Spring '97 data (452 notes with 109 authors). Since use of newsgroups is distributed, it is very difficult to get reading behavior data in that tool. CaMILE is centralized, so access data are possible to collect. The Spring '97 quarter was the first forum in which usage data has been collected and analyzed. Three questions about writing (information production) behavior were addressed, with the goal of understanding how much information was available in the ecology and that information was structured:

- **How much do individual students write over time?**
- **How broad is participation (operationally defined as writing, not simply reading) in the forum?**
- **How many of the notes are in response to others' notes (i.e., threaded)?** A sustained discussion is probably necessary for a successful computer-supported collaborative learning forum. A simple question-and-answer pair of postings (thread length of 2) is probably not a broad group discussion, a gathering of different kinds of evidence, a comparison of alternatives, or an exploration of issues. Threading, that is, the number of notes in response to a given note, is a measure of sustained discussion. I computed the length of a thread from each top-level note (that has no parent note). So if notes A and B were both in

¹ However, students may *perceive* use as being required. In some of our CaMILE-using classes, students have told us that they use CaMILE because it was required, even if the teacher did not say so explicitly.

response to note C which was a response to note D, note D would have a thread length of four in my calculation, and only one thread (top-level) would be counted.

Three questions about reading (information consumption) behavior were addressed, with a goal of understanding the rate of flow and desirability of the information in the space:

- **How much reading do students do?**
- **How much reading does each note receive?** As reported by others, individuals' reading behavior can vary dramatically within a forum [4, 5]. In WWW access, looking at reading (information-consuming) behavior from a piece of information perspective has been more consistent [3], so I apply that perspective to CSCL forums. I computed, for each note, how many times the note is read over the course of the ten week quarter.
- **When does a note get read?** Research on the WWW as an information ecology has suggested that recency is perhaps the most critical factor driving the desirability of a piece of information [12]. To explore when notes are accessed, I computed the *lifetime* of a note, where birth is defined as the first day that the note is ever accessed and death is defined as the last day that the note is ever accessed.

III. Writing (Information-Producing) Behavior

How much do individual students write over time? On average, a student using either tool wrote 4.8 notes (standard deviation of 9.8). Newsgroup authors wrote slightly less (4.4, SD 10.1) and CaMILE authors wrote slightly more (5.2, SD 9.3). The difference is not reliable ($p=0.72$ on a two-tailed t-test). Overall, this is about 0.4 notes per student per week of the course.

Figure 1 depicts the distribution of authors and the number of notes that they wrote. 87% of all authors wrote between 1 to 10 notes in the ten week quarter. 92% wrote between 1 and 20 notes. Only 5% of authors wrote more than 50 notes, that is, more than five notes per week.

The authors that write relatively little produce the majority of the notes in a forum. Authors writing 1-10 notes produce 44% of all notes in a forum, authors writing 1-20 produce 60% of the notes. The high-end authors (writing 50 or more notes in a quarter)

account for 16% of all the notes in a forum. These findings suggest that forums are not typically dominated by a small number of authors.

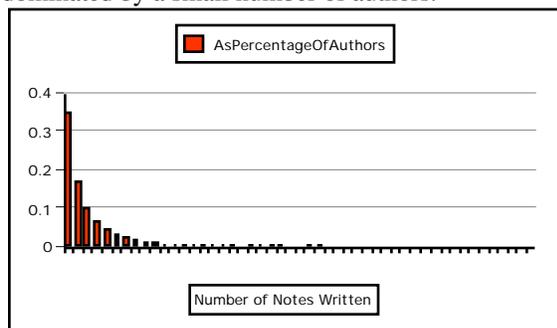


Figure 1: Percentage of Authors by Number of Notes Written

However, it is not the case that a handful of students are dominating the conversation. Rather, it is the teacher or TA who is most commonly the prolific author. Table 3 lists the most prolific two authors for a sample of the forums. Almost always, it is a teacher or TA who is the biggest contributor, and the second most prolific author is a much smaller contributor. The overall picture is that involvement of students in most class forums is fairly egalitarian, without a handful of students dominating the conversation. In general, the teacher is more often a prolific author in CaMILE than in newsgroups. This may be one of the subtle ways that CaMILE-using teachers convey their enthusiasm and encourage participation.

was slightly lower (60%, SD 30%), and Newsgroup participation ratio was slightly higher (70%, SD 37%). The difference is not reliable ($p=0.40$).

While disappointing, this difference is not too unexpected. Notice that some of the newsgroup classes have a participation ratio of *over* 100% – past students, TAs, and faculty do occasionally participate in newsgroups. Access to CaMILE is more difficult than newsgroups. There are many newsgroup readers, any student account on any computer on-campus can access a newsgroup, and newsgroup access is distributed. CaMILE requires a special username and password, and must be accessed via the WWW. That latter point, access via the WWW, is a particular *disincentive* for some Computer Science students. Some CS students complain in surveys and via email that they do not like using graphical WWW browsers and do not want to be forced to use the WWW to engage in class discussions. While a minority, there is a hacker culture that insists on use of text-only tools (including WWW browsers) and revolts against graphical user interfaces as used in CaMILE.

How many of the notes are in response to others' notes (i.e., threaded)? Overall, 55% notes posted in a forum are in response to other notes. In CaMILE, it's higher at 60%, and in newsgroups, it's lower at 50%. The average length of a thread in across all forums is 2.8 notes (SD 6.5), which suggests that most notes get a response and many get a third note in the thread. Newsgroup threads are shorter: 2.2 notes (SD 2.1). This implies that most threads in a newsgroup are simply a note

Class	Top Author	N of Notes	% of Notes	Second Author	N	%
<i>CaMILE</i>						
CS2390 f96	Teacher	108	26%	TA	17	4%
CS2390 w96	Teacher	94	19%	Student	28	6%
CHE2210 sum96	Teacher	42	25%	Student	4	6%
CHE2210 w97	Teacher	38	37%	Student	26	25%
CHE4803 w96	Teacher	23	55%	Student	4	10%
<i>Newsgroups</i>						
CS2360 Sp	TA	83	19%	TA	53	12%
CS2360 Wi	TA	296	27%	TA	69	6%
CS2430 Sp	TA	152	26%	Student	30	5%
CS2430 Wi	Student	104	19%	Student	20	4%
CS2760 Wi	Student	17	11%	Student	17	11%

Table 3: Top two most prolific authors in selected forums in the dataset

How broad is participation in the forum?

Overall, the ratio of the number of authors to the number of students registered for the course is 64% (SD 33%). CaMILE participation ratio in each class

(perhaps a question) and a response (perhaps an answer). In CaMILE, the average thread length is significantly higher ($p<.001$, two-tailed t-test): 4.2 notes (SD 10.9). The maximum thread length in any newsgroup was 56 notes, while the maximum in a CaMILE forum was 176 notes.

IV. Reading (Information-Consuming) Behavior

How much reading do students do? There were 452 notes in the Spring '97 CS2390 CaMILE forum. The average number of notes read per student was 163 (36%), with the maximum being 543 (multiple reads were counted). The standard deviation on reads per student is very large, at 158. On average, students in this forum wrote 3.8 notes each (SD 12.4), giving a read/write ratio of 42.04. (The high number of notes written was 117, by me.)

How much reading does each note receive? Figure 2 presents the distribution of the number of reads per note (aggregated across the entire course and all students). The log graph makes the observation more obvious that there are a bunch of notes that get a reasonable amount of reads (between 10 and 100 references over the course of the quarter by 92 students), but there are a few that get almost no attention and another few that are markedly popular with many reads. The maximum number of reads was 229 (for a single note).

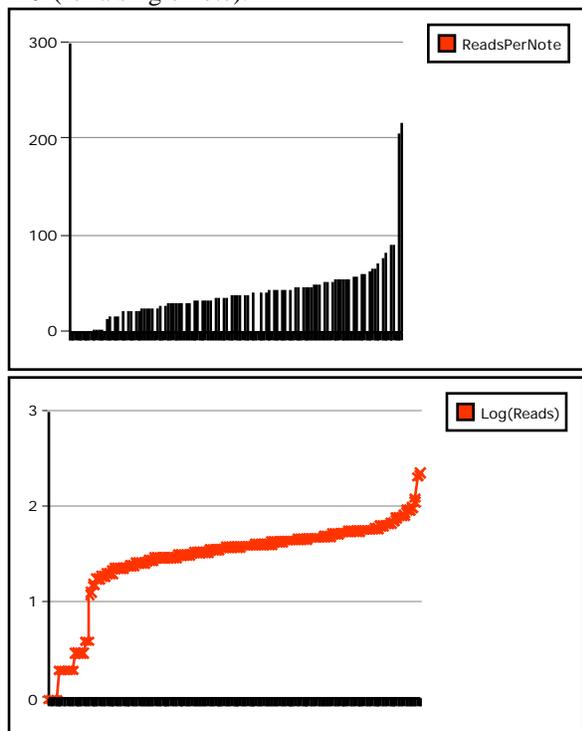


Figure 2: Number of Reads per Note (above) and the Log Graph of the Number of Reads (below)

When does a note get read? However, the number of reads tell us little about when reading occurs. Figure 3 shows the percentage of notes across the length of the note's lifetime (difference between first time read and last time read) in days. 38% of all

notes were dead in a week or less—they were never accessed again after a week of writing. 81% of all notes were dead in a month or less. Two notes (out of 452) had a lifetime of 67 days, out of the 73 days (from the start of the quarter to final exam) in the forum.

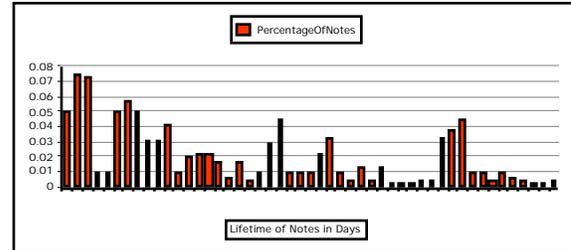


Figure 3: Percentage of Notes across the Lifetime of a Note in Days

V. Discussion and Contrast with Literature

Writing Behavior: A positive finding in these results is that participation in these forums is rather broad-based. Most students participate, and the average author (in terms of number of notes written) creates the majority of the notes.

What is more disturbing is that most authors are writing very little. Four or five notes (the average writing by an author in either forum) over the course of *ten* weeks is not what one might call a broad-based dialogue where individuals are presenting their views and responding to others. However, in contrast with literature on similar forums, these are not surprising findings.

- Results with CoNote, an asynchronous collaboration tool through which students annotated WWW pages as part of a studying activity [5], report that 65 students wrote 428 notes over 8 weeks, for an average of 0.80 note per student per week.
- Results with AnswerGarden, an email-based system for tracking questions-and-answers [5], reported that 59 students and researchers wrote 121 messages over 20 weeks, for an average of 0.10 notes per student per week.

It may be that low rates of student participation are normal in an asynchronous forum, independent of tool, where use is driven by student interest. There are contrasting examples of collaboration where the dialogue has been much more active, to where one might believe that a sharing of views and perspectives.

- In almost any kind of face-to-face dialogue, one would imagine more than half-a-comment per week. However, there are factors such as gender which can reduce classroom face-to-face dialogue [4].
- In a synchronous computer-based forum, more than one note a week is much more common than in an asynchronous forum. In MOOs, for example, interactions are much more frequent [13], though may not be as thoughtful (per interaction) as in an asynchronous forum. In Belvedere, students wrote relatively few notes, but these were augmented with a synchronous conversation, so that the asynchronous forum became more of an artifact of the collaboration than the medium for the conversation [14].
- Creating a more specific purpose for students' participation can dramatically increase the amount of interaction. The CoVis project had students create collaborative notes to fill a shared research notebook, where students created (on average) 9 notes in a two week period [6]. Researchers developing CLARE asked students to discuss a specific paper in a week [15], which resulted in student writing rates of 20 to 30 notes in that week. Taylor (1996) required use of newsgroup for the course grade and was able to get very long thread lengths (e.g., hundreds of notes).

While the newsgroup average thread length of 2.2 notes is another indication that students are not conducting much of a dialogue in these forums, the average thread length of 4.2 (with a large standard deviation of 10.9 notes) in CaMILE suggests that thread length is a variable that a tool can influence. However, we have argued elsewhere that thread lengths are longer in CaMILE due to anchored collaboration [10]. In datasets where we could carefully track use of anchors, we have found that CaMILE threads that were anchored by an external Web page of interest to students had longer threads than those unanchored CaMILE threads or threads in a newsgroup. There are other factors, besides anchoring, that might be influencing the longer threads in broader use of CaMILE. The courses using CaMILE were more varied and were potentially less familiar with technology (e.g. English and History classes) than those using newsgroups (all Computer Science students). Further, the persistence of notes in CaMILE may be enabling students to revisit and extend discussions, and the multimedia in notes or

anchors may be holding students' attention and may be encouraging revisiting of notes. For a designer, the good news here is that design of a tool can facilitate what is probably a desirable characteristic, a mediating factor of a successful CSCL forum.

Reading Behavior: Hsi and Hoadley pointed out that reading behavior and the reading-to-writing ratio varied dramatically among students using MFK/Speakeasy [4]. The results in this paper are showing a similar large variance in reading behavior using CaMILE.

Focusing on notes access may provide better insight. Models that describe desirability (in terms of the amount of attention or usage some information receives) have had some success explaining page usage on the WWW. Pitkow in his dissertation [12] found that recency drives access—if a page has been accessed recently (i.e., was found desirable by somebody), it would likely be accessed again soon. But as soon as recency dropped, access dropped very quickly. What drove the desirability (which led to the peak of accesses over time for a page) is varied. In the undergraduate settings for CaMILE and for newsgroups, assignment release and due dates probably drive much of the desirability. The key point is that access on the WWW can be modeled to a high degree of reliability by simply looking at use over time.

The CaMILE results look similar to Pitkow's findings for the WWW overall. Most pages are accessed *only* soon after they are written, and their desirability drops quickly over time. Only 20% of the pages have an information lifespan of longer than a month (that is, that they get accessed at all within a month of being written). There are several possible explanations for this result. Perhaps only 20% of the content was worth revisiting. Better indexing or searching mechanisms might have driven up revisiting. In any case, this result tells us that CaMILE notes are being accessed more like a newspaper than as a database of useful information.

More generally, all asynchronous CSCL forums may be subject to the same patterns of access as other information ecologies, such as the WWW. CaMILE usage data may be driven by recency, as on the WWW in general. Results of use on CoNote [5], for example, are consistent with CaMILE results and with Pitkow's results. Davis and Huttenlocher found that access to CoNote had enormous spikes, where access would increase dramatically (by almost a magnitude) in a short period, and then drop down quickly. They found that these usage spikes correlated very strongly with the dates that problem sets are due. Though they did not specify which annotations were

accessed during these spikes, one might imagine that different annotations would be read for different problem sets, which would lead to similar short lifetimes and recency-driven access as in CaMILE. Thus, the CoNote usage data is consistent with the same recency effect seen in WWW usage data and hypothesized in the CaMILE data.

VI. Conclusions

The results presented in this paper begin to paint a picture of the information ecology of CSCL forums such as CaMILE and newsgroups.

- Access is spread broadly across a class, but not deeply. Few students contribute many notes to the conversations, and many students contribute few notes to the conversation. The many create the majority of notes in the conversation.
- The only tool-dependent characteristic is the length of a thread, which suggests that how sustained the conversation is in the forum is amenable to influence by the tool designer.
- Not all notes are read equally. A few notes get next to no attention, while a few notes get a lot of attention.
- Students tend to read the most recently written notes. Access in CaMILE, as in other information ecologies, tends to be spikey.

These lessons can inform designers of new CSCL tools:

- If more participation is desired, it must be explicitly encouraged—either through classroom practice (e.g., requiring use, providing a context for use) or through features such as synchronous collaboration.
- If sustained conversation is desired, anchoring collaboration to objects of interest, use of multimedia in notes and anchors, and providing persistent notes and threads may lead to longer threads of discussion.
- Ways of accessing notes of interest or notes that you want students to read are useful. Students inclination seems to be for the most recent notes. Indexes to recent notes are useful. Mechanisms to highlight or index useful notes may encourage broader use of notes.

As networking technologies continue to improve and large information spaces such as the WWW are created and utilized, theory of information ecologies can be expected to develop. CSCL forums are also information ecologies, in some ways unique from general access on the WWW but in other ways quite similar. As we better understand the information ecologies of CSCL forums, we can better design and use these facilities in order to better facilitate learning.

Acknowledgments

This research was funded by the National Science Foundation through grants RED-9550458 and CDA-9414227.

References

- [1] M. Scardamalia, C. Bereiter, R. McLean, J. Swallow, and E. Woodruff, "Computer-supported intentional learning environments," *Journal of Educational Computing Research*, vol. 5, pp. 51-68, 1989.
- [2] S. K. Card, G. G. Robertson, and W. York, "The WebBook and the Web Forager: An information workspace for the World-Wide Web," in *CHI96 Conference Proceedings*, M. J. Tauber, Ed. Vancouver, BC: ACM, 1996, pp. 111-117.
- [3] J. Pitkow and P. Pirolli, "Life, Death, and Lawfulness on the Electronic Frontier," in *CHI97 Conference Proceedings*, S. Pemberton, Ed. Atlanta, GA: ACM, 1997, pp. 383-390.
- [4] S. Hsi and C. M. Hoadley, "Productive discussion in science: Gender equity through electronic discourse," *Journal of Science Education and Technology*, vol. 6, pp. 23-36, 1997.
- [5] J. R. Davis and D. P. Huttenlocher, "Shared Annotation for Cooperative Learning," in *CSCL'95 Proceedings*, J. L. Schnase and E. L. Cunnius, Eds. Bloomington, IN: Lawrence Erlbaum and Associates, 1995, pp. 84-88.
- [6] D. K. O'Neill, D. C. Edelson, L. M. Gomez, and L. D'Amico, "Learning to Weave Collaborative Hypermedia into Classroom Practice," in *CSCL'96 Proceedings*, J. L. Schanse and E. L. Cunnius, Eds. Bloomington, IN: Lawrence Erlbaum and Associates, 1995, pp. 255-258.
- [7] M. Guzdial, J. Turns, N. Rappin, and D. Carlson, "Collaborative support for learning in complex domains," in *Computer Support for Collaborative Learning (CSCL '95)*, J. L. Schnase and E. L. Cunnius, Eds. Bloomington, IN: Lawrence Erlbaum Associates, 1995, pp. 157-160.

- [8] M. Guzdial, J. L. Kolodner, C. Hmelo, H. Narayanan, D. Carlson, N. Rappin, R. Hübscher, J. Turns, and W. Newstetter, "Computer support for learning through complex problem-solving," *Communications of the ACM*, vol. 39, pp. 43-45, 1996.
- [9] D. Taylor, "Process metrics for asynchronous concurrent engineering: Communication in an Internet newsgroup," in *Proceedings of the 1996 ASME Design Engineering Technical Conferences and Computers in Engineering Conference*. Irvine, California: ASME, 1996.
- [10] M. Guzdial and J. Turns, "Technological Support for Anchored Collaboration," : Draft, 1997.
- [11] M. Guzdial, D. Carlson, and J. Turns, "Facilitating learning design with software-realized scaffolding for collaboration," in *Proceedings of the Frontiers in Education Conference: American Society for Engineering Education*, 1995.
- [12] J. Pitkow, "Characterizing WWW Ecologies," : Georgia Institute of Technology, 1997.
- [13] A. Bruckman, "Programming for Fun: MUDs as a Context for Collaborative Learning," in *Proceedings of the National Educational Computing Conference (NECC'94)*. Eugene, OR: International Society for Technology in Education (ISTE), 1994.
- [14] D. Suthers and A. Weiner, "Groupware for Developing Critical Discussion Skills," in *CSCCL'95 Proceedings*, J. L. Schnase and E. L. Cunniss, Eds. Bloomington, IN: Lawrence Erlbaum and Associates, 1995, pp. 341-348.
- [15] D. Wan and P. M. Johnson, "Computer Supported Collaborative Learning using CLARE: The Approach and Experimental Findings," in *Proceedings of CSCW'94*, R. Furuta and C. Neuwirth, Eds. Chapel Hill, NC: ACM, 1994, pp. 187-198.