Learning Via Distributed Dialogue: Livenotes and Handheld Wireless Technology

DRAFT Final Report for the Center for Innovative Learning Technologies (CILT)

Alastair Iles, Daniel Glaser, Matthew Kam, John Canny, Ellen Do

1. Introduction

The rationale for doing this research is that learning is more likely to occur if people can collaborate productively in small groups. Prior research by Eric Mazur and the Stanford TVI project have each demonstrated that small groups lead to more active questioning, better results (as measured by test scores), and greater participation by group members in the classroom. Our research, however, concentrates on what occurs inside the groups, rather than on what comes out of the groups. We believe that whether and how dialogue is *distributed* among group members is key to learning. People tend to learn more in small groups and in collaborative dialogue because they have experiences, skills, and perspectives that no-one else will have in the same way. Even a single lecture, presentation, or design will be evaluated differently by individuals. Thus, being able to distribute dialogue can lead to learning. Our earlier research supported this proposition.[1]

Therefore, for the CILT project, we investigated distributed dialogue via handheld wireless tablets as one (not the only) modality. Writing and drawing on shared whiteboards allows people to interact with each other in different ways compared to verbal dialogue. People are able to distil points, engage in textual or drawn conversations alongside the verbal dialogue, and assemble arguments interactively and in real time (i.e., writing points in response to others and then having the others add further elaboration immediately). They can draw objects or ideas and attach their own evaluations without needing to put these into a verbal form. They can refer back to earlier ideas or drawings in later whiteboard pages, which allows take-up of ideas. Alternatively, people can observe the dialogue and decide to contribute when they feel that they have something to add, which may not be supported by a verbal dialogue. The common thread is that people can do so by putting together their responses in a shared format that permits a memory of what is being articulated and that gives them degrees of freedom to collaborate.

To some extent, these activities can be captured by educational tests such as quizzes set to measure the knowledge of people acquired in response to material provided via distributed dialogue. Our interest, however, is in measuring the variables of distributed dialogue that seem to lead to the greater ability of people to absorb, evaluate, and generate material as a result of their collaborative work. Much of the knowledge and skills acquired through distributed dialogue may not be easily measured by educational tests. For example, design classes deal extensively with developing design solutions and the process and value of the solutions are what matters. The extent to which students take account of design constraints is both a process and part of the solution, and is less open to testing. Rather, it is evaluated by other practitioners based on their experience and knowledge. The generation of new ideas or insights is also difficult to measure through educational testing.

In this (interim) report, we focus on the initial development of learning metrics based on analysis of two classes (a graduate discussion group on technology for sustainable technology at the

[1] Alastair Iles, Daniel Glaser, Matthew Kam, and John Canny. **Learning via Distributed Dialogue: Livenotes and Handheld Wireless Technology.** In *Proceedings of Conference on Computer Support for Collaborative Learning*, Boulder, Colorado, January 2002.

University of California at Berkeley during late fall 2002 and a graduate design studio at the University of Washington in Seattle during mid-spring 2003). These deployments have contributed markedly to our evolving conceptions of what learning via distributed dialogue is and how to measure it. The report provides initial details of the studies, their methodologies, and key findings, with much more detail being given to the Seattle findings because these are highly intriguing. We plan to add further data analysis based on further review of videotape and quantitative coding analysis. We also plan to complete developing visual conversation maps as a powerful metric of learning that reveals how people distribute insights and are inspired in their dialogue by what others contribute.

As expected, the Seattle study yielded many new insights into the nature of distributed dialogue. We had suspected that drawing could be an important means of learning, but prior deployments had not made this a central feature of the dialogue. In Seattle, drawing emerged as a major mechanism of learning. Students did not necessarily engage in drawing as a dominant part of their dialogue, but integrated drawings as part of their dialogue. As a result, the metrics first developed for the Berkeley study have been changed markedly to reflect what happened in the Seattle study. For example, learning is signaled by not simply the expansion of a point in an ongoing conversation thread based on text, but also by adding to a drawing with annotations and questions and answers. There are endless ways in which expansion or embellishment could occur. An underlying learning metric, however, can capture this behavior by targeting whether or not an input into the dialogue is amplified or developed.

Significant cross-campus collaboration took place during this research. Without Professor Ellen Do's insight into analyzing learning and design studio dynamics, we would not have been able to understand more about what was happening in the studio classes. Without access to the Design Machine Group at the University of Washington, we would not have been able to study the learning of design students. Likewise, the Seattle participants would not have been able to analyze learning without the skills and input of the Berkeley team members (Alastair Iles on learning metrics and transcript analysis, Daniel Glaser on methodology and videotape analysis, Matthew Kam on programming and Livenotes system support, and John Canny on learning metrics and system design).

The learning metrics proposed in this report as a result of the empirical data will help in the design of future deployments of handheld wireless tablets, and in evaluating many varieties of distributed dialogue in the educational field. Simply put, these metrics deal with the longevity and character of conversation threads; the extent to which participants ask prompting questions or make "provocative" statements that inspire further dialogue; the degree to which participants make comments reinforcing, contradicting, or feeding back on dialogue; the extent to which participants create memories of prior dialogue on a given transcript page; (in the case of drawings) the degree to which a drawing is annotated or added to by group members; and the charting of dialogue structure and dynamics through visual conversation maps to see how ideas are generated. We have found that learning – as defined in terms of outcomes such as generating new ideas or changing the state of knowledge – appears to be much more likely to occur

2. Definition and Measurement of Learning

For our study, we define learning through distributed dialogue as:

- changes in understanding related to the subject matter (as marked by the take-up of ideas or design concepts, the explanation of points, the specificity of commentary on the subject matter);

- additions of new knowledge (such as discussion of a previously unknown U.S. government program aimed at sustainable cities);
- acquired ability to ask questions or to contribute to the larger group discussion (for example, group members may be better placed to ask a really good question if they have already discussed the issue in Livenotes)…

that result from the collaborative interactions of group members when using the shared whiteboard.

At a finer scale, to capture the learning that may occur through dialogue, we need to attend to:

1. **The context in which a specific input occurs** (measured by reviewing the transcripts: does it happen in response to a lecture slide or prior commentary, does it cross over from a previous page to a new page…?)

2. **The aim of the input** (gleaned from content, context, and interviews with participants: is the input note-taking, clarification of the lecture, discussion of something not directly connected to the class, organizing the class, a humorous interlude…?)

3. **The impact of the input on the subsequent dialogue** (measured by reviewing the transcripts and interviews with participants: do people respond, do people clarify what they said, do people ask questions, do they keep discussing the point, do they refer to or take up a concept in later pages, do they ignore or drop the input, do participants react to other participants…?)

4. **The change in thinking or knowledge resulting from the input** (measured by tracking the development of thoughts across the dialogue: can we see different participants adding their ideas or insights to a thread, can we see evidence of new awareness of what a concept means (not just on the tablet screen but in verbal dialogue), do participants correct or "provoke" each other…?)

Each input needs to be evaluated in terms of these features by looking at the Livenotes transcripts, videotape analysis of images and audio (which can tell us, for example, how far the transcripts correlate with the verbal dialogue), and feedback from users following sessions. We are still performing this analysis for the Seattle study and revisiting the Berkeley study as a result of what we have learned in Seattle.

At a broader scale, learning is heavily context-dependent. That is, what counts as learning will depend on the format, goals, membership, and process of a particular classroom. Design classes generally do not provide opportunities to students to ask questions or to contribute directly to some joint product (a discussion or a design outcome). In contrast, the TASED group was by definition a group meant to ask questions about the readings. Thus, it would be expected that in the design classes, there would be fewer questions posed to the class as a whole and more attention to analyzing the design (which in fact transpired). We therefore provide detailed elaboration later in this report of learning in terms of each study and in each session. An advantage of conducting the two different classes is that we can investigate learning according to each class and thereby determine whether or not there are similar, cross-cutting features of learning. If these features exist, they may be generalized across the educational field.

The definition above differs from two kinds of learning that we have observed in previous deployments of Livenotes, namely learning how to engage in distributed dialogue and learning how to use Livenotes. In the Seattle study, for example, students were initially confused as to who were using what ink colors, and did not know how to use the new awareness widget to

identify which pages users were on. To some degree, they learned how to resolve these issues. Instead, our focus is on how people interacted cognitively with the subject material of the classes.

3. Methodologies

In all sessions in Berkeley and Seattle, we used Clios, notebooks, or tablet PCs. We used Clio handheld tablets for the first Berkeley session, but they did not work with the new user interface, which demands both significant computing power and a finely touch-sensitive screen. We therefore temporarily switched to using notebooks on an ad hoc basis to keep the study going. Migrating to laptops meant that keyboard support and lecture slide features were added to Livenotes, thus facilitating distributed dialogue in ways previously not feasible, but also inadvertently changing the experimental design part-way through the first study. Still, the new user interface is far superior to the original version used in earlier studies, leading to much more and better user input. Users appreciated being able to type into Livenotes as well as being handwriting into it. In all sessions at Berkeley and at Seattle, we found the use of text as input to be an important part of the conversation. In turn, a new lecture slide feature allows us to superimpose Powerpoint slides on the tablet screen, and users can write on top of the slides, enabling direct interaction with lecture content. The TASED group meetings used this superimposition feature extensively, but the Seattle meetings did not use the feature because it would have been incompatible with the nature of the design classes. As a result, a new learning metric was made possible, namely measuring the extent to which users annotated or commented on the slides.

For the Seattle deployment, we were finally able to use 10 Tablet PCs given to Professor John Canny through a delayed Microsoft Learning Sciences grant.[2] These computers allowed for input via sophisticated, more easily held and manipulated pens (as well as keyboards), had much superior computing power compared to the Clios, and had far more sensitive screens. These features significantly enhanced the amount and quality of the graphical input of the users. This was especially timely since the Seattle deployment involved architectural design students. We tested two types of Tablet PCs: the Toshiba Portege 3500 and the Compaq TC1000. Both machines had similar functionality for running Livenotes, but had some important differences from the perspective of users. The Toshibas were computationally faster machines, but were less ergonomic to use since their stylus was rectangular in shape, whereas the Compaq stylus had smooth contouring that was far easier to grip.

In both the Seattle and Berkeley deployments, we used peer to peer 802.11b networking. The internal network protocol of Livenotes uses a client/server model which was invisible to users except during the cumbersome startup routine. This allowed us to run Livenotes sessions among laptops or Tablet PCs in any building regardless of the networking infrastructure. In the Berkeley study, the sessions met in one of the computer science school buildings, and therefore could piggyback on the building network and gain access to the internet. In the Seattle study, the sessions met in the design studio building which is not wirelessly networked, but used an "ad hoc network" in which one of the Tablet PCs acted as the server and supported the network connection. As a result, the Tablet PCs and Livenotes can be used in a broad range of settings including the outdoors (which may be relevant to, say, environmental science classes).

[2] We have continued to use the Tablet PCs in the ongoing TASED group meetings and in the CS160 course that Professor Canny is teaching in Spring 2003.

We have accumulated a vast amount of data from the two studies in the form of Livenotes transcript pages, videotape and audio recordings, questionnaires distributed to class members, and feedback discussion after the last two sessions in Seattle.

The Livenotes program is our most important data source by definition. It captures all the inputs that users make on the shared whiteboard and therefore serves as a permanent record that both users and researchers can refer to. Our basic unit of analysis is one transcript page. We analyze what happens on each page, both individually and in the context of the other pages. It is possible to store Livenotes transcripts at the end of each session on the server Tablet PC or notebook, for later retrieval and uploading to the project website in html format. This makes it easy to print and view the transcripts. The website is: http://kettle.cs.berkeley.edu/livenotes (user name and password needed; these can be obtained on request to the researchers).

In Berkeley, we filmed the last two TASED sessions with a digital camera on a tripod at one corner of the room overlooking the table. The audio turned out to be reasonably good but spotty in places. In Seattle, we chose to use two different, better cameras, one mounted on a tripod and set up to cover the whole room, the other carried by Daniel Glaser around the room to zoom in on individual users and Tablet PCs.[3] The handheld camera turned out to have excellent audio whereas the stationary camera was too far from the discussion, but the audio coverage is adequate for our analysis. We have learned much about the logistics of setting up the equipment for our future research. We have stored our videotape and audio data on DVDs and are analyzing it with special video editing software (Adobe).

We collected a grand total of 134 Livenotes transcript pages (Berkeley – 51 pages and Seattle – 83 pages; see below for specific break-downs). We also made videotapes of two sessions in the Berkeley class and of two sessions in the Seattle class, amounting to a total of about 4.5 hours, mostly accompanied by useable audio recordings of what was being uttered verbally.

Our analysis is exemplified by looking at Figure 1, which shows how we are cross-correlating (especially for the Seattle deployment) the different sources of data to reconstruct the Livenotes dialogue, its broader classroom setting, how the dialogue relates to the physical design object and the verbal classroom dialogue, and the time and space over which these interactions occur (both within the Tablet PC group and the classroom as a whole). It can be seen that this is a complicated process. But we are realizing, as a result, that the Livenotes dialogue is only a partial representation of the classroom dialogue and is simultaneously its own conversation. It becomes very important to determine what exactly is transcribed onto the shared whiteboard to see what and how participants are selectively choosing to emphasize or discuss. This is also true even when the Powerpoint slides are superimposed onto the whiteboard, since participants will only discuss some of the points raised, and may not even track the slides at all, preferring to raise relevant but extraneous matters.

This makes Livenotes a sensitive means of following learning, since the participants typically transcribe what they are attending to and thinking about or processing. They are not just passively following the classroom dialogue or the Livenotes conversation, but are actively constructing their own inputs.

[3] This was the first time that we had used a mobile camera as well as a stationary camera, and we learned that the mobile perspective can capture images that are very useful. It was only though the mobile camera that we learned that some of the Toshiba group data may have been lost as a result of the temporary group network disconnection.

Figure 1. An image of the analysis workbench constructed by (clockwise from top left) a hand camera with audio, a classroom camera on tripod, the LiveNotes transcript, and speech in class. We found cross correlating the data to be helpful in understanding the many conversations in the classroom.

4. Deployments and Initial Analysis

In the following discussion, we outline what each study did, the kinds of meetings that took place, and the participants. We also define what learning means in the context of each study and analyze some of the most important findings.

4.1 The Berkeley Class

We decided to set up a seminar and reading group on Technology and Sustainable Economic Development (TaSED) that began meeting weekly on November 16, 2002. Co-convened by Matthew Kam and Alastair Iles as part of their own research interests, the group has a website: http://www.cs.berkeley.edu/~mattkam/tased/index.html. The group included at least 10 U.C. Berkeley graduate students from engineering, computer science, and the Energy and Resources Group (environmental policy and science). Each meeting lasted for over one hour, began with a lecture presentation (in the form of 10-20 overheads) by a volunteer lecturer, and discussion either tracked the presentation or occurred for 30-40 minutes after the lecture. The topics examined different aspects of sustainable development and technology.

The TaSED group met four times, with three meetings providing a total of 51 pages of dialogue transcripts (Session I: 24 pages; Session II: 10 pages; Session III: 17 pages). Attendance has varied from 4 to 6 people, reflecting the late semester start. Livenotes formed an integral part of

the discussion in that participants were asked to conduct as much discussion on the tablet screen as they could. Lecture slides were exclusively provided on the tablet screens, encouraging dialogue connected to the slides. (TaSED has continued to meet during the Spring 2003 semester, but not for CILT purposes.)

Figure 2. Pictures of the Third TASED session to illustrate the meeting space and the use of the notebooks.

We are revisiting the Berkeley data on the basis of our insights into the Seattle data. We have accumulated quantitative data to measure inputs and marks by all participants, but need to go back and review the transcripts to see how ideas are developed across the pages. A full set of transcript pages is posted on the project website. The important analytical results are explained in section 5 below. For now, we present one example from the meetings to illustrate what we are looking at in terms of learning metrics. Figure 3 shows several conversation threads taking place on one transcript page, along with inputs that did not spark further dialogue.

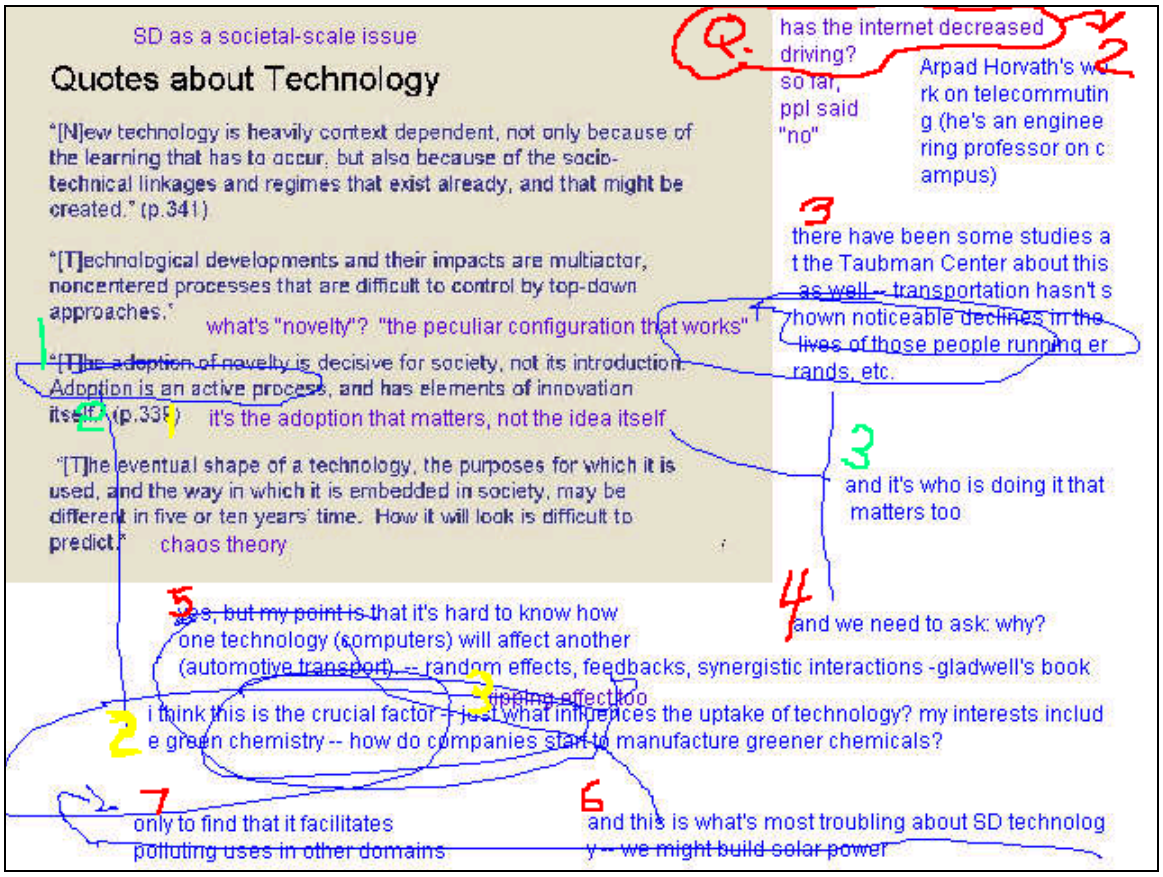

Figure 3. A transcript showing how the dialogue starts with a question "has the internet decreased driving?" and then threads of 7, 3, and 3 step lengths each. Also on the page are remarks that were not extended into threads.

4.2 The Seattle Deployment

For our second CILT deployment, we chose to perform a study of a design studio class that Professor Ellen Do teaches at the School of Architecture at the University of Washington. Professor Do is one of the co-PIs of the study and directs the Design Machine Group where the study took place. The class is Architecture 401/504, "Inventing Futures: Space, Computing & Toys", a winter/spring 2003 course intended to teach visual thinking, spatial reasoning and innovation through exposing students to art, architecture, computer science, and cognitive science.[4] The class comprises 12 graduate students who are engaged in developing projects inspired by play with toys. It usually meets three times a week. Therefore, we considered that the studio would provide an ideal setting for exploring learning via distributed dialogue in a relatively short visit to Seattle. The course website is: http://courses.washington.edu/studio03/gallery.html.

In summary, we conducted three sessions under different conditions:

- Session 1 was a lecture in the regular weekly Tuesday afternoon Design Machine Group seminar on Tuesday, March 4, 2003.

[4] According to the *Seattle Post Intelligencer* (a Seattle newspaper), "Do's objectives for the class include teaching students to work in teams and to participate in class discussions. The topics and disciplines covered range from art to computer science to architecture. However, the core of the class is thinking: how people think when they solve problems, from mundane interior design, say, to helping a blind person create animation." http://seattlepi.nwsource.com/local/86491_visual11.shtml

- Session 2 was the full studio meeting of the design class meeting on Wednesday, March 5, 2003.
- Session 3 was a smaller subset of the larger design studio meeting on Friday, March 7, 2003 (conducted during the regular studio hours).

Figure 4: The design classroom where the Seattle study took place. This shows a scene from Session 2, the design studio class on Wednesday, March 5, 2003.

Session 1

Session 1 was designed to allow interested design studio members and faculty at the Design Machine Group to become more familiar with Livenotes before actually engaging in the major tasks that we proposed to set. Daniel Glaser gave a lecture between 4 and 5.30 pm on Tuesday, March 4, 2003, explaining how Livenotes worked and outlining the potential uses of handheld wireless tablets in design. The audience comprised 10 students and faculty (plus the Livenotes team). They were meant to take notes of the lecture. Thus, we measured learning in terms of how far they had reported the lecture, and how far they had expanded on or interrogated points that were made during the lecture. Four Tablet PCs were connected in a network, and at least 6 individuals used these during the lecture, trying out the capabilities of the whiteboard. This alerted us to the possibility that people would swap Tablets and make it difficult to correlate users with inputs.

As expected, the first and final transcript pages (and for all Seattle sessions) were filled by humorous drawings or writing, suggesting that this is a widespread way of getting ready, determining if people are on the same page, or practicing the stylus writing action before beginning a session. Also, if the drawings take place towards the end of the session, this signals that people are growing tired and are ready to terminate dialogue. We will not analyze the session for this interim report, except to note that only 2 pages out of a total of 6 were directly concerned

with the lecture, and that one person took the notes and helped prompt others to comment on the lecture.

Session 2
Session 2 was designed to probe how group members learned through discussing the design models that each classmate presented for review. In this session, which ran for 3 hours between 1.30 pm and 4.30 pm on Wednesday, March 5, 2003, 12 models were presented for group critique. These models were:

1. Tommy - transforming sail with servo motors [DISCUSSED IN LIVENOTES]
2. Koo-young - musical wire [DISCUSSED]
3. Sorin - music interface for DJ [DISCUSSED]
4. Babak - light bulbs for energy saving [DISCUSSED]
5. Cornel - the wooden box/public furniture in plaza [DISCUSSED]
6. Patrick - puzzle space and furniture (all with puzzle pieces)
7. Heidi - luggage to go
8. TJ - rock-paper-scissors [DISCUSSED]
9. Junta - architecture blocks with database scoring
10. Golnaz - transparency of spaces, game, collage
11. Yeonjoo - software ripple, painting in transparency
12. Markus - flexible cube/pin pressure

Not all presentations were discussed in Livenotes because of a combination of technical system failures (temporary network disconnection), fading group concentration, and the nature of the last presentations which necessitated the group moving across the room to look at computer screens.

The features of the design studio class need to be carefully understood because these distinguish it from most other class formats. Designers do not usually work collaboratively in the design studio (during their education) but individually. In the design critique process, however, everyone has the opportunity to briefly comment on the designs. In this case, each presenter spoke for 5 minutes about his or her design model, either using a physical model to demonstrate the design, or projecting a short computer-generated show onto a screen at one end of the room. The Tablet PCs were added to this environment. There are few precedents for shared note-taking in design critique since the critique takes place through people commenting directly on projects. Therefore, by using the Tablet PCs, the studio members were being asked to adopt a new practice, which may have led to initial uncertainty over how to use Livenotes.[5]

Class members needed to focus on multiple targets of concentration: the physical model or the visual projection, the verbal presentation and ensuing questions and answers, and the dialogue taking place in Livenotes itself. Because the students did not have any specific task set in advance, they were free to comment on whatever parts or presentations they found most interesting. This unstructured dialogue also meant that students could pick up and leave Tablets on chairs if they wanted to focus on the verbal dialogue, potentially resulting in temporal gaps in the Livenotes dialogue.

The Toshiba group consisted of 4-6 users working with 4 Toshiba Tablet PCs, while the Compaq group had 4-6 members using a Compaq Tablet PC. Overall, both groups exhibited an impoverished discourse (compared to Session 3), highlighting how the setting and the group

[5] However, the design students and faculty responded very positively in the Ftiday session, suggesting that experience with Livenotes can lead to acceptance.

dynamics mattered for learning. Importantly, both groups were relatively unstable in terms of their membership, with several members giving their Tablets to others to use, meaning that user identity was not always obvious, especially for the Compaq group.[6] At one point, the professor, Ellen Do, was given one of the Tablets for about 10 minutes and she contributed an important part of the discussion on Babak's presentation. Nonetheless, the groups constituted a valuable comparison of how some variables may influence the nature of the dialogue that takes place.

For each group, we measured whether or not the students had learned through distributed dialogue by looking at:

- the extent to which students referred to, analyzed, criticized, and expanded on, the presentations, either contemporaneously or later in their dialogue (in text or drawing) [Category A];
- the ways in which the students using Livenotes contributed to the verbal dialogue (such as by asking questions) [Category B];
- the degree to which students generated their own comments and insights as a result of evaluating the presentations[Category C];
- the degree to which students gave feedback on each other's dialogue through clarifying, correcting, or reinforcing statements [Category D]; and
- the numbers of dialogue not related to the design presentations [Category E].

Compared to the Berkeley study, it is much harder to trace conversation threads because the dialogue is more widely distributed across pages, involves drawings, and is shorter in attention. Even so, the concept of a thread remains a worthwhile analytical tool to use.

The data from the Compaq group is very rich, compared to the Toshiba group. The groups differed in that the Toshiba group experienced technical problems, namely network disconnection for about 15 minutes early in the studio meeting; had weak preexisting social relationships (the Toshiba group were not friends, whereas the Compaq group were all friends); and had a different spatial and social lay-out in the room (the Toshiba group were more widely spaced out, some sat on chairs adjacent to or behind the sofa, meaning that members could not see each other as well, or observe what others were doing on their screens; in contrast, the Compaq group were densely bunched together on the sofa and could watch each other closely).

[6] In Livenotes, it is possible to identify an user in the new user interface by tapping on the user icon. However, this does not allow for the possibility that users will switch their identities, leading to confusion over who is engaged in the dialogue.

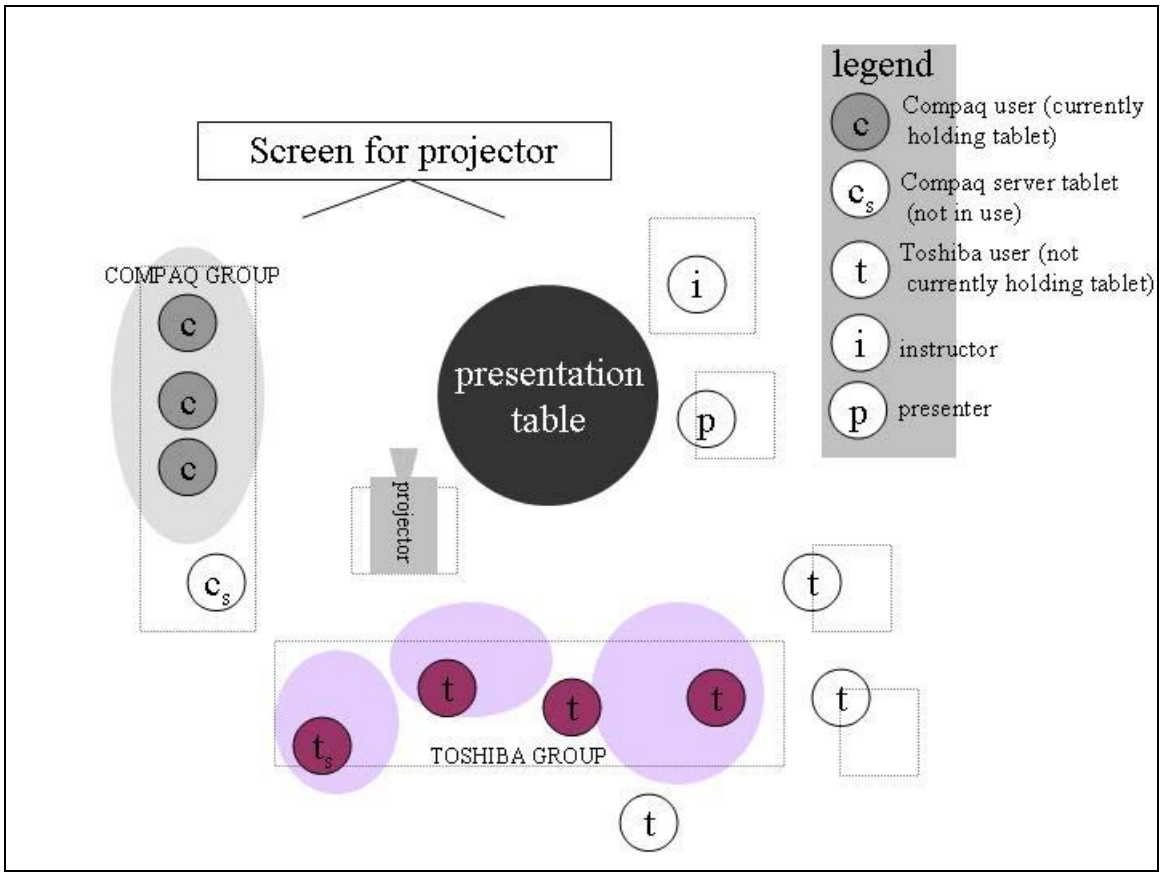

Figure 5. Classroom layout during the second session

We present two examples (for now) to illustrate how learning occurred in the groups.

A. Critiquing Tommy's transforming sail project.

The project was a sail sculpture (in the form of a curving bird wing) that would rotate in a landscape with the aid of motors mounted on a platform at the top of the sail's stand. A remote control would drive the motors to work at different rates, causing the platform to rotate with limited degrees of freedom, and therefore deforming the sail into different planar appearances. The Compaq group's analysis of this design was brief (only half of one page) but was relatively sophisticated in its understanding of the design concept and problems. With a few marks, the group summed up the presentation intelligently. However, as the videotape transcript and close-up images of Tablet PC screens reveal, the presenter was inarticulate in explaining his design, thus losing the concentration of two group members who proceeded to draw his face on the page towards the end of his presentation.

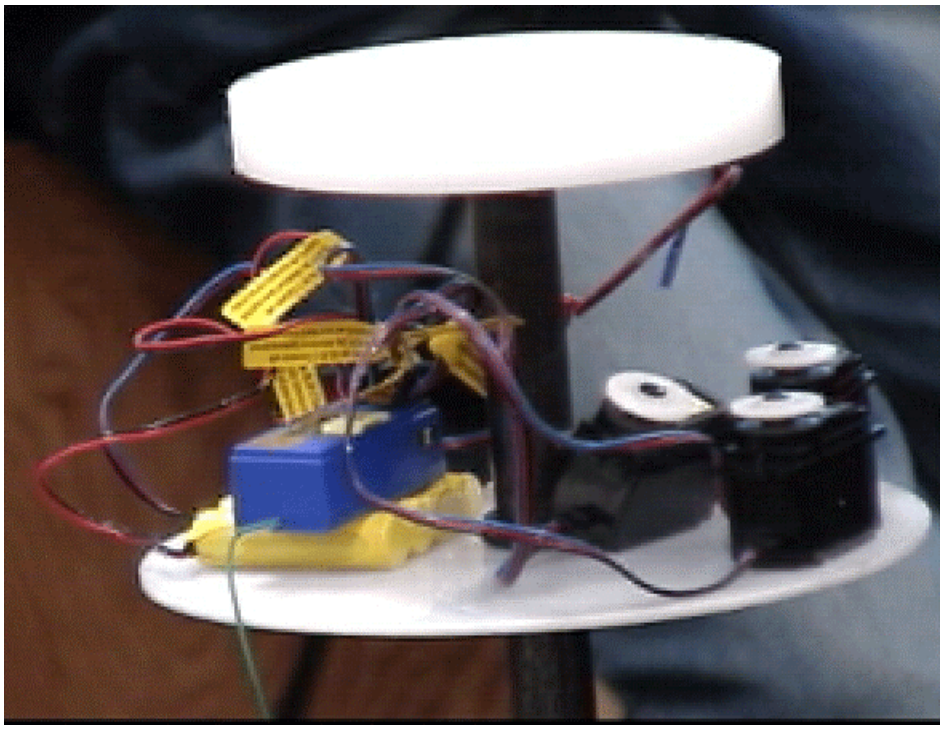

Figure 6. A closeup of the servos on Tommy's project.

As Tommy introduced the sculpture, “Blue” drew the stand. In the presentation, Tommy explained that the stand would rotate though servos moving sideways instead of vertically. “Blue” drew some machinery on the round platform above the base and an arrow indicating the direction in which the stand would rotate. “Cyan Blue” added his own arrow to indicate that he understood the design concept. He also wrote “Servos and stuff” at the top of the page to allude to the machinery, and therefore how the stand would move.

This reflected the observation that although Tommy talked about rotation and movement with the servos, he did not articulate well how the jumble of electronics (as seen in the image) was linked to the movement. The dialogue on Livenotes was inhibited by the verbal dialogue: there were few points at which group members could discuss the proposed design. “Servos and stuff” is a vague statement that does not explain how the motion will actually work, or what its source will be. “Brown” repeated “servos and stuff” underneath to signal that he was paying attention to the presentation but was possibly confused. These marks showed that the group was closely following the presentation and extracting several fundamental points (namely, rotation and mechanism of movement). However, on analyzing the video and the design itself, it becomes clear that the group selected only some details to question, and did not depict the upper part of the stand in accurate relief.

In response to Tommy’s saying to the class that he did not quite know yet how to make it hold together, or whether the design was strong enough, “Cyan Blue” asked a question about how the stand moved (“counterweight?”), and “Blue” then wrote “Heavy base” and linked this to the base with a long looping arrow. This was an attempt to critique a crucial feature of the design – and to underscore how Tommy had already addressed the issue.

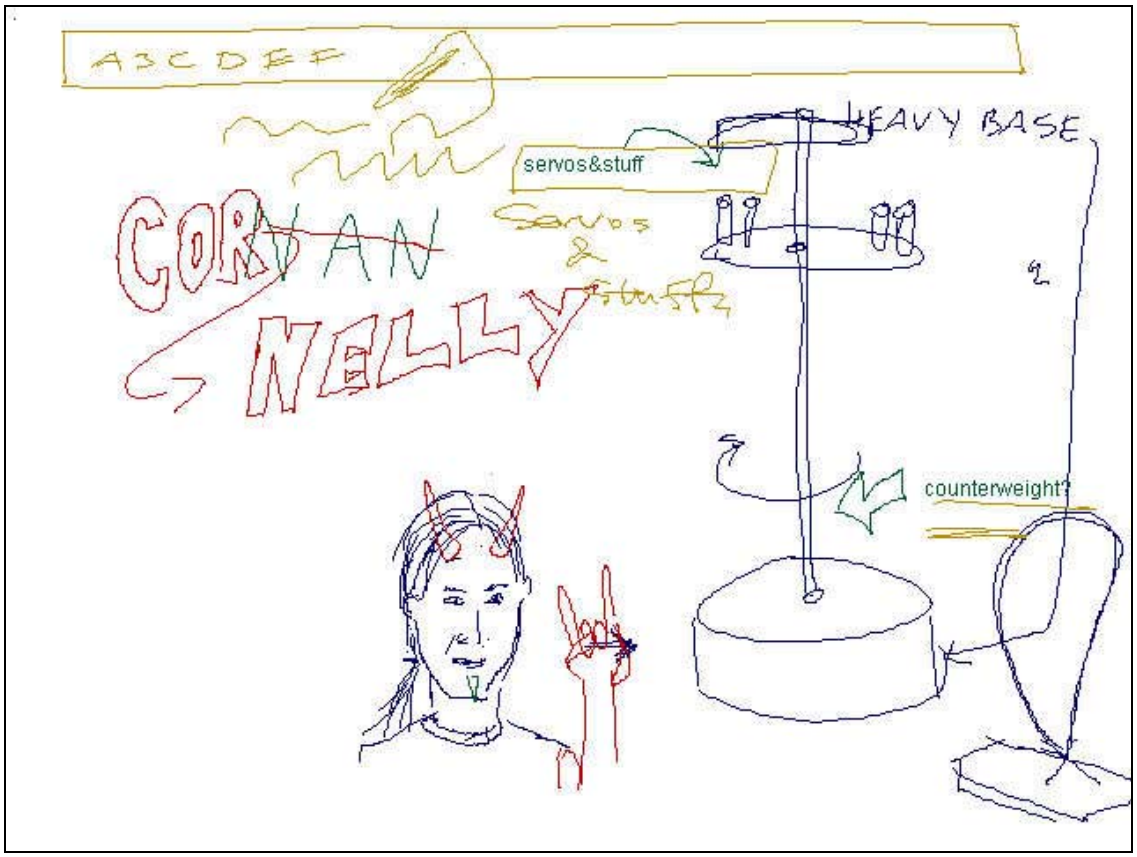

Figure 7. Screen capture of the Compaq's group discussion on Tomas' work

The videotape analysis shows that the Compaq group asked a crucial question to Tommy about how he would address the design problem of the third axis of movement. This question was inspired directly by the prior Livenotes dialogue; the other participants did not ask such a penetrating question.

However, the page also demonstrates that play occurred alongside the note-taking, both when the presentation began, and when members became confused over the design due to the inaccessibility of the verbal dialogue and the non-participation of the expert questioner ("Cyan Blue") who had previously sparked discussion of the role of the base. "Blue" drew the presenter's face, with contributions by "Red" who had not participated in the note-taking. Then "Blue" drew the deformed sail in one corner of the page to show what would be placed on top of the stand. Yet no-one made a comment in Livenotes regarding how these would be connected in practice. "Blue" did not try to explain how the sail would be deformed, for example.

In this case, learning occurred through the group not just summarizing the presentation, but deconstructing the design to identify important unresolved design issues. The dialogue threads were short, and combined text and drawing, but were still very expressive. The above sequence, however, could not be understood without matching the transcript to the verbal presentation and questions and answers. This suggests that the group members did not use Livenotes as their primary mode of learning for this particular presentation, but instead used it to distil a few key points.

In contrast, the Toshiba Group yielded this transcript page of the same presentation.

Figure 8. Toshiba "group" transcript. Note a single author who asks for "comments", then "Any comments?" This transcript markedly differs from the Compaq group shown in Figure 7.

The page shows a much less sophisticated understanding of the presentation. Only one person participated, namely "Red". "Red" draws the sail and the stand as they would look in the landscape, as a conceptual drawing, and mentions wind and light as features that will distinguish the sculpture. Significantly, "Red" asks repeatedly for feedback, "Any comments?", but no-one else responds to this prompting question. Other group members do not add to the drawing. This suggests that either the other members were not viewing the page since it was far ahead of the page that they were using, or were not paying attention to it.

B. Critiquing Sorin's musician aid project.

Sorin presented a project, "Music Box", which sought to assist disc jockeys trying to create and play digital music, such as drums, without a physical sense of the music's tempo and rhythm (which would ordinarily come through actions such as tapping). The Music Box assembles several pieces of electronics and sound equipment to ease the work of the disc jockey. Again, the presenter was relatively inarticulate in his explanation, providing a dense mass of material for the class to absorb.

"Red" provides a highly technical, pedantic summary of the design on one page. His sequence of drawing mirrors the presentation itself, by drawing the physical model on show, and then drawing the sequence of overheads that Sorin projects but selecting important features to highlight. "Red" notes some key features of the design: "Similar to an Eq", (comparing to an existing music, the equalizer, that the presenter refers to), "16 faders" (devices that control the sounds being sent to speakers for broadcast), "BPM controller" (beat per minute regulator), and "snare kick" (a method of generating a beat rhythm within the system that does not need the disc jockey to tap physically; and which appears to be the key innovation in this design). At the end, "Red" draws a speaker emitting music, illustrating the process in which the beat is generated for the speaker. Notably, a single person is doing all the commentary on this page, without any participation by anyone (except "Green" who seems to be absent-mindedly doodling). This page demonstrates significant learning through transcribing and extracting the technical features of the design.

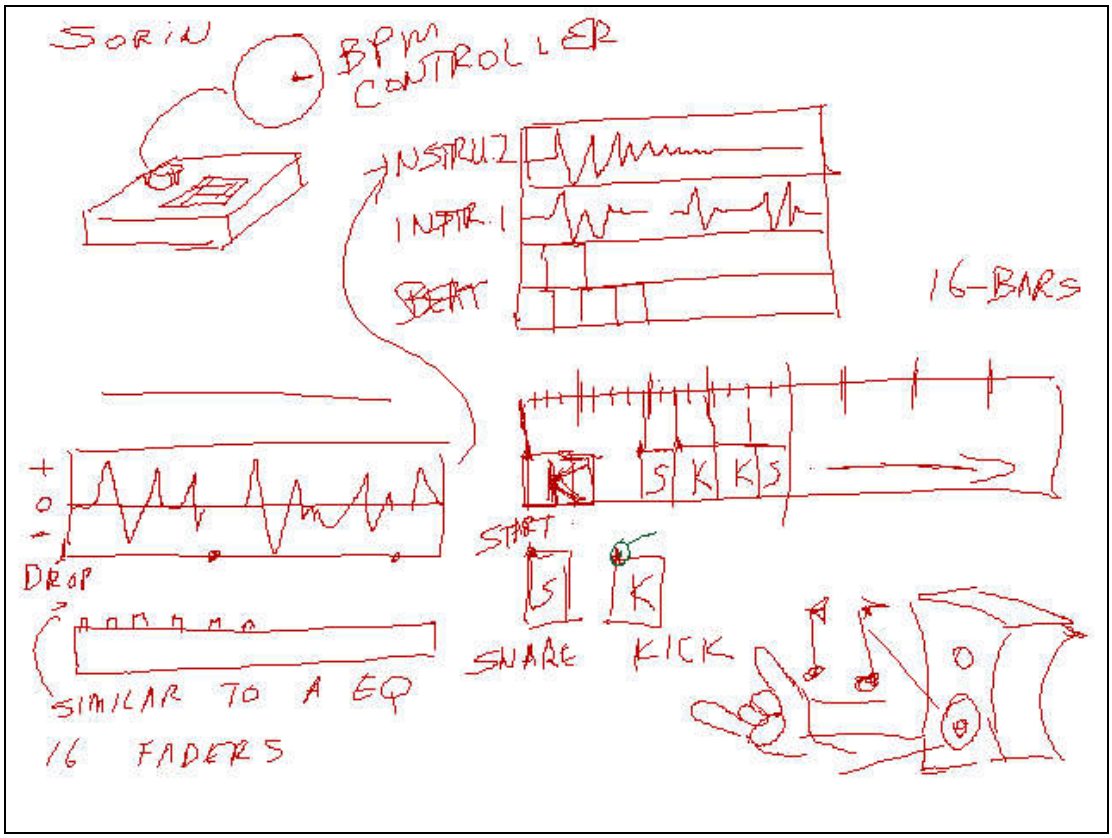

Figure 9. The first transcript of Sorin's presentation.

However, a completely different perspective on the design is offered in parallel. Several pages later, "Blue" gives a creative analysis of the design that states the basic area of work, the design problem and who needs it, the proposed solution, how the solution will work, the name of the concept, and then ends with a joke. He does so by writing a conceptually inventive ballad without drawing the project:

> Music [AREA OF WORK]
> Rhythm music composer [DESIGN USER]
> Musician – if you are
> If you are missing a drummer [DESIGN PROBLEM]
> Device [SOLUTION]
> Make beats without tapping [HOW SOLUTION WORKS]
> 8x10…4x4 beat…
> kick snare, kick snare [NAME/METHOD OF SOLUTION]
> snare kick, snare
> bump kick bum bum kick

By conducting video analysis, we see that the lines are inspired by, yet go beyond, what the presenter says.

Then "Blue" explains critically that "I'm lost in all the technical stuff". He does not understand the intricate details that Sorin has projected onto the screen and sees this as undermining the transparency of the design. "Brown" agrees and says "Good!" "Green" draws a voice box

surrounding the ballad emerging from a figure that "Brown" draws following the ballad. In response, "Red" sarcastically writes "We know!" and demands "Who's this?" when he sees the drawing of the presenter. It appears that "Red" has a different learning approach based on careful recording of the technical presentation, whereas "Blue" focuses on the underlying principles of the design and tries to figure out what the function of the design is. It is striking that much more interactive dialogue and drawing takes place on this page, and that the design is not just analyzed, but also criticized, whereas "Red" merely summed up the technical features of the design.

Even the dialogue of "Red" with regard to Sorin's project far exceeds that of the Toshiba group. The Toshiba group seems to have been producing some discussion based on this project (as captured in videotape close-ups of Tablet screens), but this discussion was not captured in the final transcript because of a network disconnection.

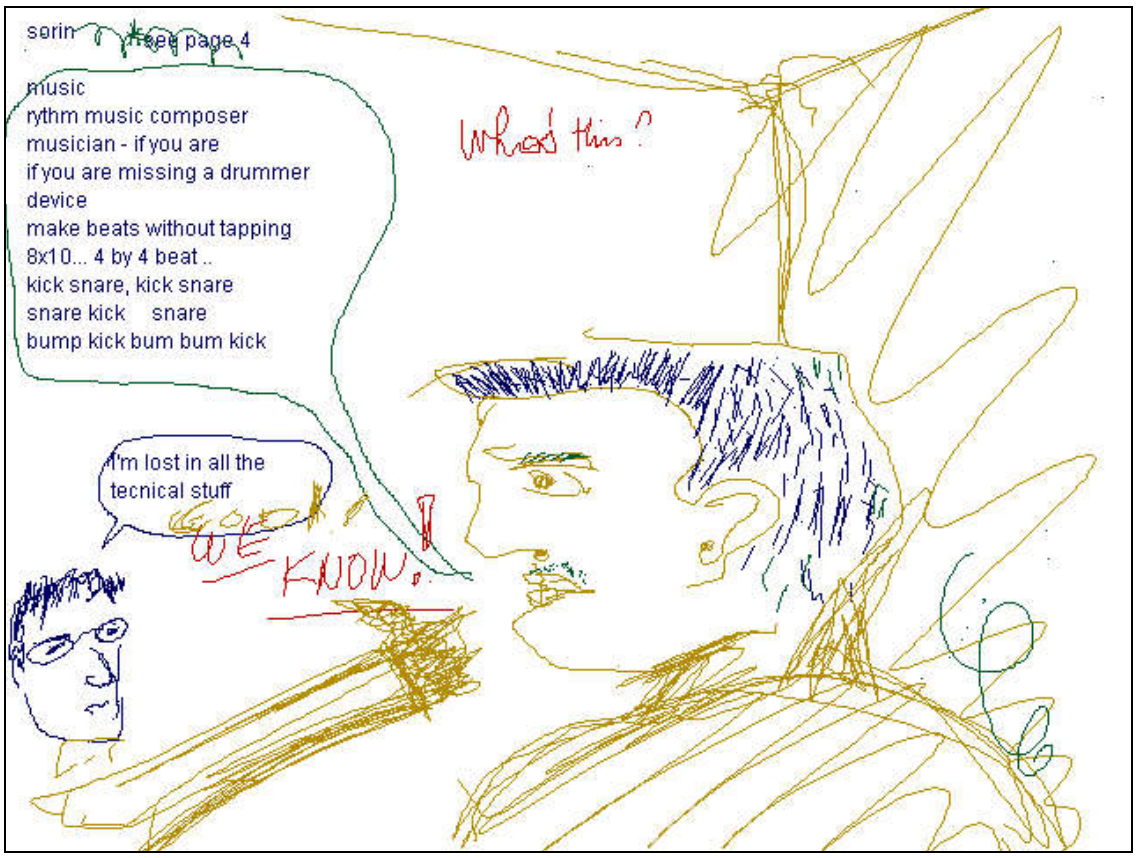

Figure 10. The second transcript of Sorin's presentation.

Therefore, the Toshiba group serves as a good control case to see why the Compaq group was moderately better in terms of learning than the Toshiba group. Differences in learning are marked by three sets of observations:

1. *Qualitative analysis.* Individuals in the Toshiba group appear not to have discussed, or absorbed, the presentations in detail; not to have taken up important ideas or drawings (such as the concept of "public furniture") in later portions of dialogue; and not to have given feedback to each other on the ideas being expressed. No-one in the Toshiba group was willing to take a leadership role and scribe notes, draw, or ask prompting questions, therefore inhibiting dialogue.

Yet members of the Compaq group were much more willing to scribe details of the presentations. Threads of dialogue in the Toshiba group were almost non-existent: they petered out and seldom were taken up again in later pages, whereas several discussions in the Compaq group had longevity. The Toshiba group drawings were relatively less elaborate and layered, and had almost no text or explanation attached to them. They also tended not to be collaborative or interactive, but to stand on their own,[7] while the Compaq group drawings were more collaborative.

2. *Quantitative analysis.* Coding the transcript inputs in terms of the categories set out above shows quantitatively that:

	A	B	C	D	E
Compaq Group					
Toshiba Group					

[We will elaborate on this analysis.]

3. *Surveys and questionnaires*. In the feedback discussion held immediately after the session, Toshiba group members reported that little learning had occurred. They complained that Livenotes was a distraction from the verbal class dialogue. One student argued that Livenotes was a "pain in the ass". Compaq group members seemed more positive about their learning experience, with one member stating that Livenotes was an useful means of sharing comments.[8]

As well as the network connection failure already mentioned, these differences may be attributed to the different degrees of stability in identity and membership between groups. The Toshiba group swapped Tablets often, leading to questions about who was writing what, and perhaps a much more anonymous tenor of dialogue. In contrast, the Compaq group remained highly stable in identity and noticeably there were no questions about who was writing what. The Toshiba group was also much more concerned about criticizing each other or making inappropriate statements (such as those pertaining to "Little Buddha", a nickname assigned to a presenter who encountered problems in setting up his presentation).

Session 3

Session 3, held between 3.30 and 5 pm on Friday, March 7, 2003, was designed to pose a structured design problem to the participants, instead of having an unstructured discussion that depended on whether or not participants took an interest in the designs being presented. Participants were recruited from the design studio on the basis of their interest in the design problem and in testing Livenotes in greater depth. The problem was "Design an alarm clock for Alastair", a non-hearing member of the CILT research team, using the constraints that he could not hear an alarm, many light clocks did not succeed reliably in waking him up, and vibrating devices were likely to fall off the bed. (Previous semesters of the design studio had involved design projects aimed at addressing disabled people needs, thus motivating the choice of this particular problem.) Alastair Iles wrote problem specifications and design constraints in Windows Journal, which were projected onto a screen that the groups could see. Daniel Glaser then gave a short lecture outlining these to the two groups that had volunteered to take part.

[7] An example is the public furniture presentation on page 3 which does have a single good drawing by one individual of a feature but does not attract overt written comment from other members and has no explanation.

[8] Interestingly, two members of the Toshiba group participated in the Friday design session and found it more satisfying as an experience. This implies that different settings and learning dynamics may affect the accessibility of distributed dialogue to participation.

Group 1 comprised 2 experienced designers and 1 student; while Group 2 had 3 students from the design studio. Crucially, they all knew each other as friends or colleagues or both, as they later stressed in the feedback session. Thus, as long as they knew the identity of the inks, they did not worry about whether or not to censor themselves – whereas in Session 2, there were many instances of people worrying if they should say something or not. This meant that the dialogue was much more fluid and focused. One significant difference may have been that Group 1 had a better sense of eye contact because they were bunched closely together at one table end whereas Group 2 were spread out on both sides towards the other end. Thus, Group 1 may have found it easier to monitor each other and to resolve usage issues by looking or gesturing at each other.

For this session, we measured learning through indicators highly specific to the problem-solving setting:

- the extent to which members discussed the design problem and design constraints, working these into their deliberations on proposed ideas;
- the numbers and nature of design ideas that were proposed;
- the insight that participants gained into the most salient principles underlying the design problem (that is, what would wake Alastair up reliably?); and
- the ability of the group to arrive at an agreed design solution through assembling the dialogue.

[We are still working on coding the input points which is a very time-consuming activity; our final report will present quantitative analysis of how the two groups differed.]

We present three examples of different stages in the dialogue in Group 1 to illustrate what it included. Importantly, Session 3 was not affected by technical equipment failures, and ran perfectly on this point. The impact of the technical failure on the Toshiba group in Session 2 may therefore have been very important in not only leading to some data loss, but also to a loss of group engagement. Our analysis at this point focuses on transcript captures since we have not yet considered the videotape and audio material pertaining to this session.[9]

Diagram x (below) is the second page of the Group 1 dialogue and represents the group's efforts to comprehend the design problem and constraints as outlined by Alastair Iles and Daniel Glaser. "Blue" takes excellent notes, aided here and there by "Green", to summarize the short lecture, extracting what the group needs to keep in mind as it discusses potential design solutions. "Blue" repeatedly underlines the key design constraints (without sound, light, robust, portable, vibration, easy to build, and so on) and begins to imagine design solutions ("cushions"). This page does not reveal distributed dialogue in the way that the TASED group sessions at Berkeley did. In these sessions, people shared the note-taking work. Here, a single person shouldered almost all of the work (as well as providing most of the prompting questions).

Nonetheless, this page is significant because there is no equivalent in Group 2. Group 2 did not take notes and instead continued to engage in humorous drawings and "getting on the same page". Consequently, as the transcripts for Group 2 demonstrate, Group 2 members repeatedly proposed design ideas that would not satisfy the constraints and even suggested that third party actors (namely, pets or housemates) should intervene to wake the sleeper. They did not appear to have learned about the constraints.

[9] However, we can say that – in strong contrast to the previous two sessions – the groups did not engage in much verbal dialogue and instead focused on Livenotes as the medium of interaction. Thus, the transcripts represent the most reliable, informative source of data that we have for Session 3.

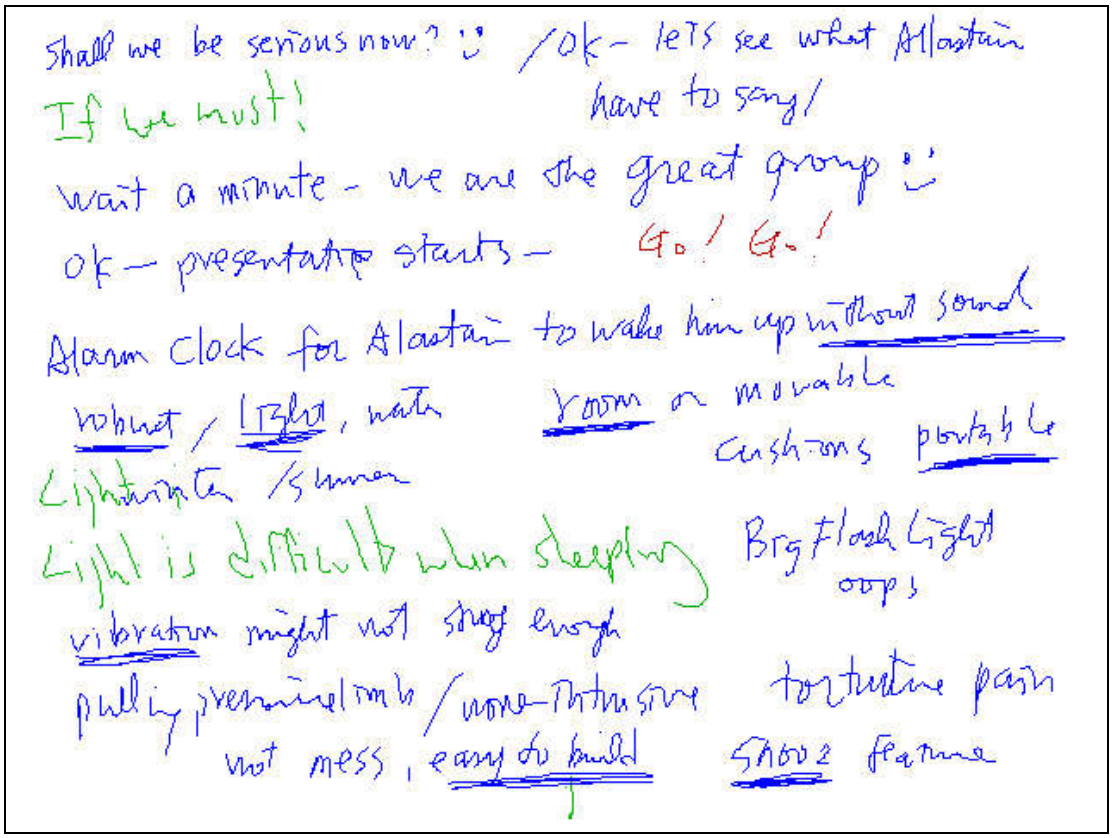

Figure 11. The screen capture of the note-taking of the brief lecture in which the design problem and the design constraints were presented.

The transcript in Figure x (below) shows a very rich set of conversation threads that generates a major new design principle initially leading to creative combinations of design ideas on this page, and later to the invention of a very different type of clock (the stink clock) based on the design principle. Without distributed dialogue in Livenotes, this learning would have been much less likely to occur.

First, "Blue" acts as a memory maintainer, ferrying over a design idea from the previous page that she thinks summarizes the current state of play in the design process. She then asks "What else?" therefore prompting the group to continue its dialogue (whereas the other group had no such prompting presence or memory ferry, leading to an impoverished, slow-moving dialogue). "Red" suggests massage as another approach, following on from the vibrating sensory experience previously analyzed, and "Blue" draws a person lying on a bed with an undulating arrow under the bed before noting "Some kind of motion head to toe". In response, "Red" draws another bed that repeats the concept introduced in the upper right corner but noting that the desired massaging motion can be "a moving ball along bone" with "Blue" adding "in circular motion" and a small drawing of a rotating ball.

Meanwhile, "Green" suggests "electric shock" and "Blue" instantly has a negative feedback, "!That's scary" (invoking the design constraint that the solution should be non-intrusive and painless), leading "Green" to suggest "fire/heat" instead. "Blue" writes "I like the hot tub one", draws an arrow, and then a tiny picture of a bed tilting to throw a sleeping figure into a hot tub. But "Red" underlines "shock" and says, "Good!", suggesting that she sees the possibility that

what will most effectively awaken the sleeper is a kind of shock. She writes "Shock", but "Blue" argues "forget that" and draws an arrow pointing to "shock". "Red" then expands and says "and then water – good sequence!" "Blue" says "Ah sauna – first massage then hot tub" which is not what "Red" had in mind but is a plausible design outcome. This is an important example of how group members monitor each other's contributions to the dialogue, seeing whether or not these are compatible with the design constraints.

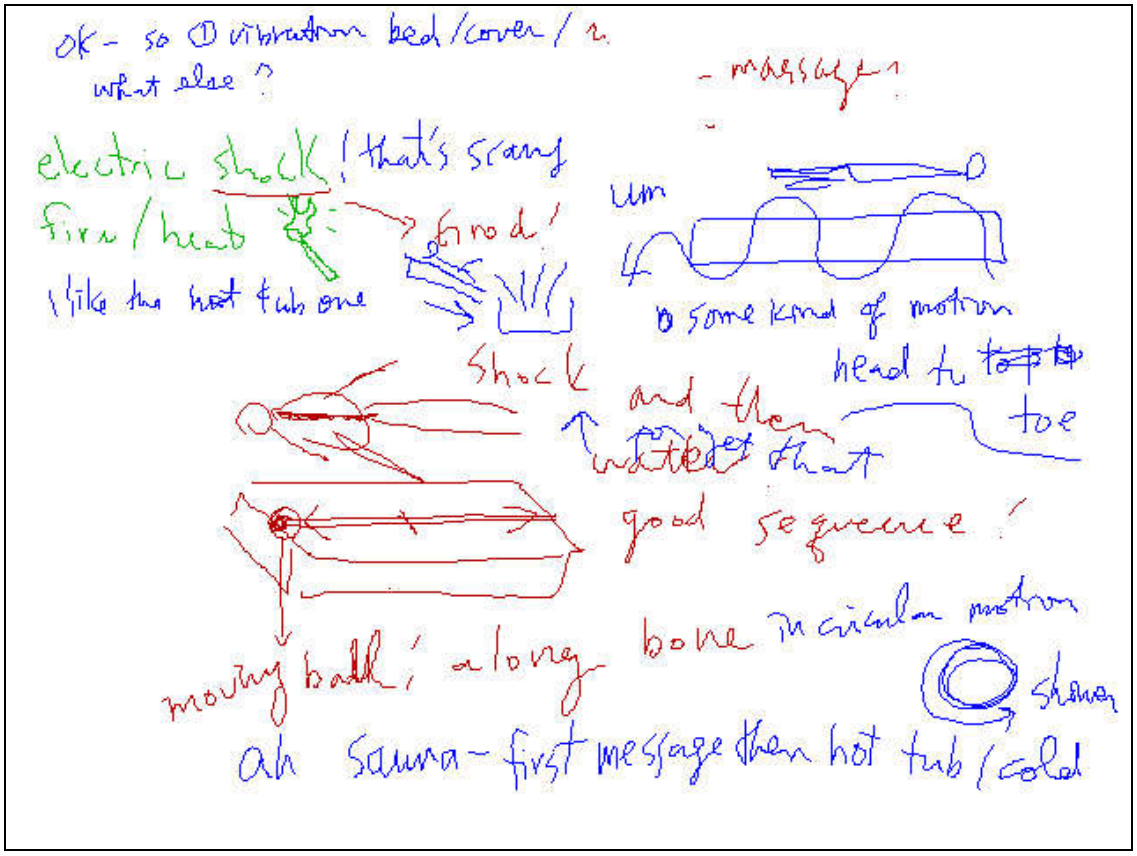

Figure 12. Screen capture during the "alarm clock" design problem at DMG.

In other words, the spatial layout of the whiteboard is not a trustworthy guide to the actual sequence and structure of the dialogue. Simply reading the transcript, moreover, does not tell you whether or not group members are learning. The different parts of the dialogue are located all over the page and refer back to earlier pages. More important, the salient learning that occurred here is fully assimilated only on later pages. The two design ideas (electric shock and fire/heat) and associated reinforcements or negative feedbacks are apparently disconnected. But a new idea, "shock", is developed as the underlying design principle. "Green" had noted that electric shock might work, and "Red" realized that there was an underlying design idea – that of shocking the sleeping person awake. "Green" notes 3 pages later, "It would have to be an extreme change to wake a person up", which evidences learning about how the design constraints can be resolved.

As seen in Figure x, this leads to the abrupt realization by "Green" that a stink clock– based on awakening the sleeping person with the shock of pungent odors – might be an effective design solution. "Blue" – who had previously rejected the idea that shock might be important – begrudgingly allowed "sure, at least that's harmless" and underlined "stink". She then adds smell

to her evolving list of environmental conditions that can be adjusted, signaling that she is receptive to the idea. This is a wholly new, unexpected idea that appeared only because the group members had already developed the design principle of shock. The rest of the page is taken up with a freewheeling debate over the different variants of odors that could be experimented with. The group members also begin to design the procedure by which the stink clock could work.

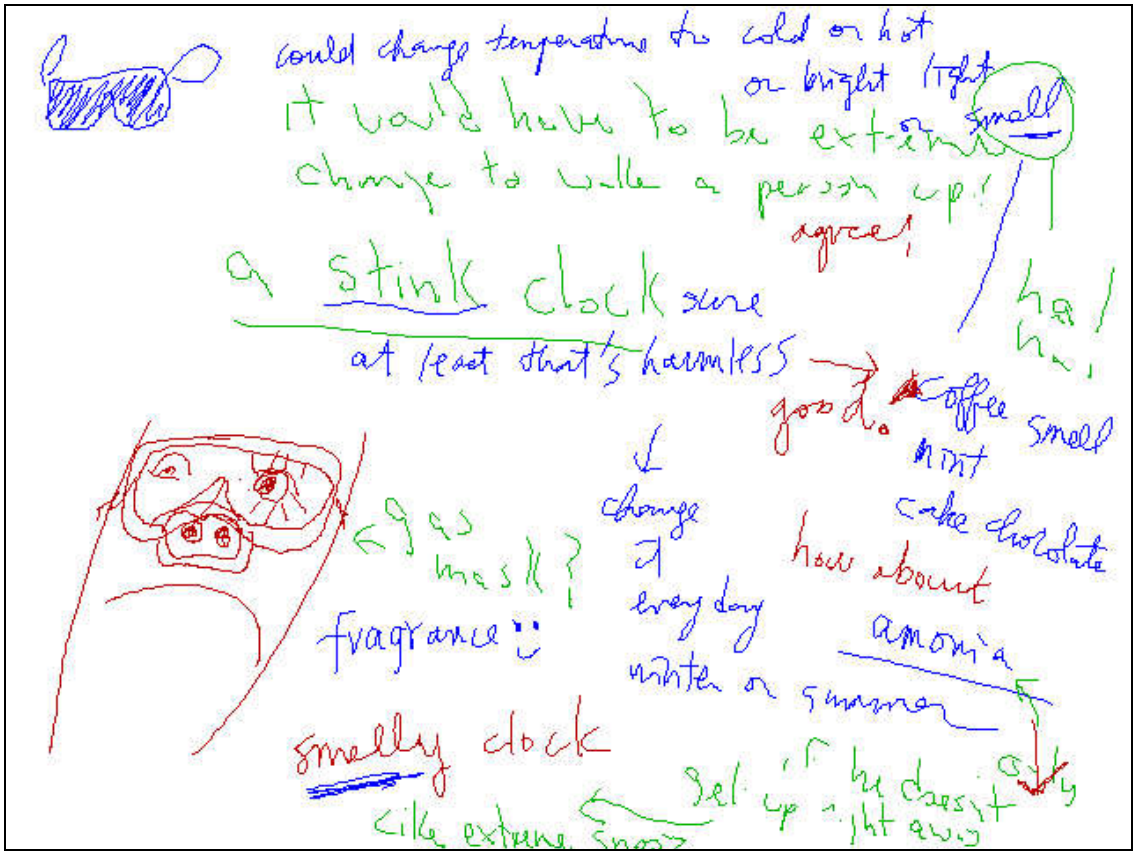

Figure 13. **The screen capture that shows the appearance of the stink clock concept.**

Another key feature to note in Figure x is that "Blue" again acts to ferry a design idea from the previous page to spark the dialogue, but this time by drawing a pair of eye covers made of a material that can adjust its temperature from cold to hot and therefore shocking the person into wakefulness. Otherwise "Green" would not have made the observation that gave rise to his own insight regarding the stink clock. Equally important, the page is peppered with feedbacks and reinforcements such as "Good", "Ha, ha!", "agree", or "how about" contributed by all participants. They help the members distribute the dialogue, provide feedback on whether or not the design idea might meet the constraints or address the client's needs, and build a group camaraderie conducive to learning. Learning, then, occurs through members generating and commenting on design ideas.

In contrast, these feedbacks and reinforcements are almost completely missing from the transcripts of the other small group engaged in the design session. It can be seen that most of their discussion covers things that are not directly germane to the design problem or that transgress the design constraints. Indeed, there is great confusion over who the client is, and there are frequent side-conversations.

In terms of session 3, we conclude that learning (as defined above) can be observed through the following empirical data:

- the addition of new ideas such as "shock" or "stink clock" to the dialogue (new design concepts) [this can be measured in terms of major/minor ideas]
- adding different layers to a design drawing on the page (if people add features and write comments or questions alongside, it becomes more complex and meaningful, and is more likely to be picked up and used in later dialogue)
- the development of a thread (building an idea up, leading to a "workable" design, as we see with the hot tub and massage thread in Group 1)
- the creation of a permutation of ideas (mainly in Group 2: e.g. alarm clock wakes cat that wakes A; cat food is dropped on A, cat licks A; drops cat on A) to see what might work
- questioning whether or not a proposed idea or design is compatible with the design constraints set out at the beginning of the session and also, realizing that a design feature would violate the design constraints
- references to external events or materials that other group members do not know about and that become incorporated into the dialogue to enrich it (e.g., citing films as the source of inspiration for design ideas; saying in Group 2 that a housemate or a pet might serve as a clock, which does not come from the design specification)
- explanations being offered about how a proposed design might work (this occurs much more in Group 2 which may be why they are so slow in their dialogue whereas Group 1 seems more focused on generating design ideas).

We are working on building a set of data like the above. Generally, Group 2 seems much less active and more impoverished in their design dialogue compared to Group 1. There were only 2 major ideas put forward in Group 2 whereas 7 were proposed in Group 1. Group 2 only has 4 pages of design dialogue whereas Group 1 has ~9 pages, indicating a great disparity of dialogue development. Group 2 drew a lot more than Group 1, but never really referred to the design constraints. They were much more uncertain when deciding to move between pages. Group 1 was much more consistent and smooth than Group 2 in their dialogue.

5. Discussion

[We plan to expand on these points and show how metrics can be designed.]

Learning (as measured in the distributed dialogue) seems more likely to occur if:

- If someone acts as a facilitator, asks questions, or makes a "provocative" statement (such as "I really don't think sustainable development is workable") when the page is first accessed. These interventions help stimulate people's interest or attention (just as a question can be important in catalyzing verbal dialogue). Many more conversation threads typically occur when such inputs are found on the page, for both the Berkeley and the Seattle studies.
- If someone draws or writes on the page from a prior page to maintain the memory of what the dialogue has been covering. This aids group members to continue the dialogue and to build constructively on what has already been said.
- There is someone taking careful notes of dialogue in the broader setting, or of the requirements and tasks that the group is being asked to perform. This helps the group to frame the conversation and to monitor what is being said in terms of whether or not it is germane to the group. For example, Group 1 in the third Berkeley session had detailed notes of the design constraints, but not Group 2. As a result, we see much less reference in Group 2 to the design constraints. Someone says "door open/shut loudly" without anyone pointing out that Alastair will not hear this.

- If there are lecture notes, Powerpoint slides, or images provided on the Livenotes screen, people appear much more likely to comment on these rather than to engage in their own side conversations. Providing some structured input as part of the Livenotes environment helps stimulate learning related to the input.
- There is reinforcement, feedback, or commentary on people's points. If people are continually commenting on or giving feedback on dialogue (even if only in monosyllables such as "good!" or "scary"), they are showing that they are engaged in the dialogue and absorbing the material. They are motivating the author to continue his or her input and creating a setting in which people are distributing the work of the dialogue. But if there are no annotations attached to the input, this often means that no-one comments on the input: it is free-standing and the dialogue dies. Feedback seems to lead to further development. Strikingly, when there was a prompting question or reinforcement in Group 2 of the third Seattle session, the dialogue became richer and more focused, and Group 2 instantly changed its favored design away from a design idea that would not have met the design needs.
- A drawing or textual input has multiple contributors who each add something. If people do not draw or write a comment interactively or collaboratively (i.e., they draw a picture or comment that stands by itself and that is not added to by other people), then they are much less likely to build on the drawing in subsequent conversation. If they do, then the input is much more likely to be an entrenched part of the dialogue and to be connected to later pages (i.e., development over time).

Each observation helps predict the extent to which learning may occur. In short, the degree to which dialogue is distributed, and how it is distributed, marks the likelihood of learning. Learning metrics can be designed to capture these features and therefore assist teachers at all levels to develop, monitor, and evaluate distributed dialogue.

In some cases, however, dialogue dies out quickly and is not continued beyond 1-2 thread components. Counter-intuitively, lecture slides with predominantly visual content do not seem to elicit significant dialogue among participants. In contrast, lecture slides with exclusively textual content seem to stimulate tablet screen dialogue. Similarly, we discovered in the Seattle study that if a speaker is inarticulate and inaccessible, the group dialogue atrophies and people lose attention or become less able to specify on the whiteboard what they want to discuss.

These observations imply that learning may not occur under some conditions, or may be shifted to a different modality (such as verbal dialogue). It is crucial to understand these conditions. For instance, the placement of the input on the tablet screen may be important. If input is hard to see on the page, it may be ignored. If it runs up against tablet screen space constraints (i.e., the page filling up), it may not be transferred to the next page unless participants are motivated to do so. Participants may also not find the input salient, preferring to focus on some other input. That is, the ways in which the dialogue is being *distributed* on the whiteboard screen itself may inhibit some learning (as well as stimulate this learning). Alternatively, what makes input salient to participants may reveal much about how they interpret distributed dialogue. Metrics of "missing learning" may be just as important as learning metrics, in indicating when distributed dialogue is not likely to support learning constructively.

We recognized that our first iteration of metrics analysis (as performed for the CSCL paper) was inadequate to deal with the far richer level of dialogue on Livenotes. This analysis simply classified ink strokes into a set of categories based on tablet screen management and content. These categories could tell us much about the distributed character of the dialogue (e.g., how many inputs were made by participants, individually and collectively; how much of the dialogue

is related to the lecture itself; when input occurs throughout the session) in quantitative terms. But they could not tell us about the *learning* that happens, if any. As a result, we developed an initial set of learning metrics that are now updated to reflect the Seattle insights:

1. Learning occurs through developing "threads" of dialogue that build on each other, usually with 5-7 items contributed by 2-4 people. The threads can have multiple content: e.g., lecture slide-> question -> expansion of point -> expansion of point -> reference to bibliography. Threads can go up to 18 steps over three transcript pages in the longest thread observed to date. Sometimes, the threads occur concurrently, with participants engaging in separate conversations on the tablet screen. These threads clearly indicate that distribution of thinking about the lecture is happening. Signals of learning may include: the sudden appearance of a comment that takes the dialogue in a new direction while building on the lecture; listings of book references to help some participants; commentary on the lecture slides that go beyond their content. (These are examples from the sessions.) **We think that major insights into learning will occur through constructing metrics and maps of the thread progression that record: (1) the "steps" in the dialogue (2) the sequence of specific steps in both content and timing; and (3) the spatial movement of steps across the tablet screen page.** These maps will count not only the quantitative details, but the transition in thought, content, and framing throughout the sequence. Right now, our draft maps are intricate and hard to follow on the transcripts. Our goal, therefore, is to devise means of making the maps more transparent to researchers.

2. Extended dialogue is much more likely to occur when prompted by a question or a "provocative comment" (that is, a statement that challenges the lecture slide content or what another participant has said, or that poses a problem not otherwise addressed in the lecture), or if someone ferries a memory of the dialogue (or a fragment thereof) from previous pages. The question or provocative comment seems to stimulate learning through creating a shared group consciousness. **A powerful metric is the number of questions and provocative comments, along with where in the dialogue they occur. Prompting and memory ferrying is likely to be associated with learning.**

3. Strikingly, multiple threads on the one page often somehow link together eventually, such as two threads trying to define "sustainable technology" from different angles that link together by referring to some shared concept or word. That is, learning may occur through people realizing that threads point towards a common conclusion. Alternatively, this may be evidence that some participants are focused on a specific point, and repeatedly try to bring the dialogue back to the point. **Another metric is to measure the extent to which threads become joined on the tablet screen.**

4. Learning also occurs through complex, dynamic, unfolding processes of distribution on the page, as is seen in the "shock" example in the third Seattle session. These are processes by which group members inspire each other, give each other information and ideas that others can use and combine in novel ways, and add ideas to create an overall idea or argument. This is a very difficult dimension of the dialogue to capture, but we believe that it underlies session 3 of the Seattle study and the last two Berkeley sessions in which members were able to generate important insights by pooling the comments of disparate individuals. **This can be turned into a visual conversation map to highlight the learning that occurs, the ways in which distribution facilitates it, and the specific dialogue components involved**. This map consists of a series of fields and tags superimposed onto the page. We envisage that it will be possible to build into the conversation map:
   - links to video and audio clips of actions such as the group asking a question

- a play-back function in which the various components appear in sequence on the page (using Snag-It software to capture the images and create time stamps, which we did for one of the Seattle sessions)
- symbols denoting the context, impact, and intention of each specific input to enable teachers to measure

This map will not be applicable to every page, because it demands a lot of resources and time to construct it. However, some features can be automated. The map will reveal the key underlying sequence and structure of the dialogue. Much dialogue appears to be "all over the place" and not necessarily connected directly to the classroom activity. Yet, this dialogue may be joined in very creative, surprising ways. Different members may contribute ideas, and then another member may see how to connect them. Below is an example of the visual conversation map that we are designing at present.

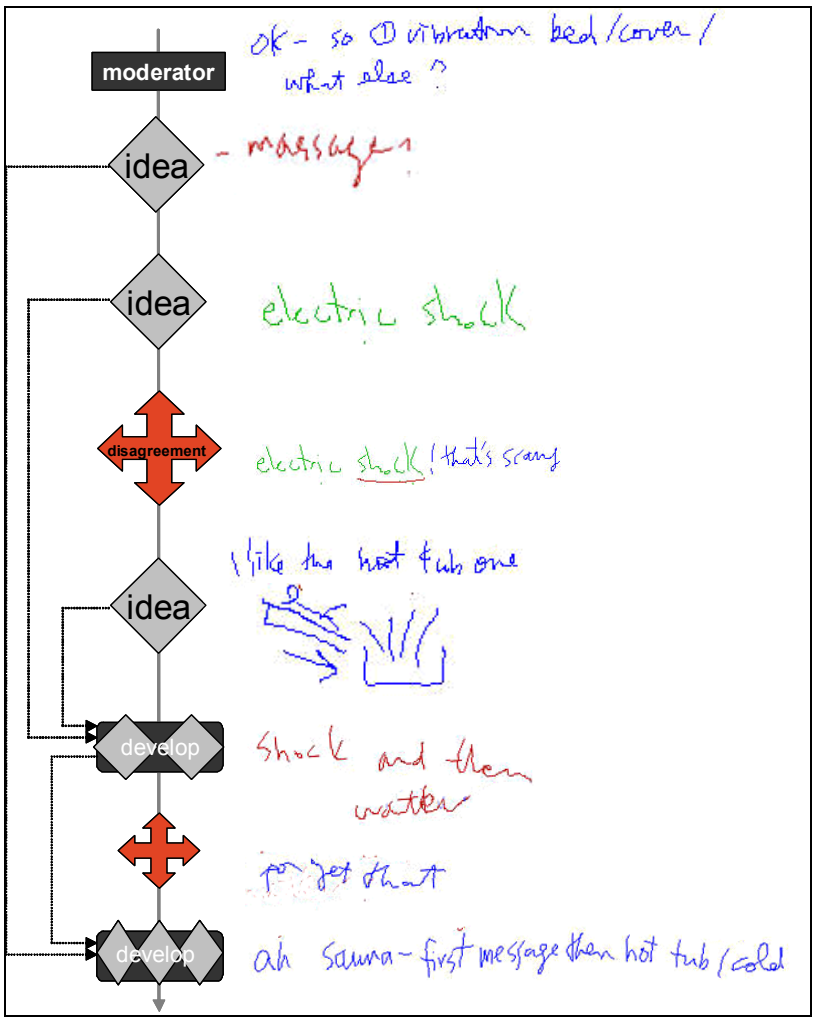

Figure 14. A partial conversation map of the notes in FIgure X

Based on these metrics, we will develop software to automate the extraction of indicia of learning. This will include grouping ink strokes and generating conversation threads to estimate how much input occurred and their diversity; looking for questions or prompting statements that

spark further dialogue; and searching for the uptake of key ideas (such as by doing keyword or drawing structure searches) throughout the transcript domain. The Livenotes research team is now expanding with new programmers. We are assigning some of these programmers, semester projects to design and implement individual metrics, especially as part of CS160, a course that Professor John Canny teaches.