

Multi-level Coordination in a Computer-Mediated Conversation

Pamela W. Jordan and Megan Moser
DRAFT

Abstract

Conversational participants use their evolving common ground to interpret and generate utterances. Participants assume their common ground is coordinated, a correct assumption in most cases. To study this coordination of common ground, we collected computer-mediated conversations between pairs of subjects collaboratively solving a problem. During the conversation, each subject maintained an independent record of the current partial solution, which functions as a direct representation of an important part of the common ground. Our presentation describes the immediate coordination that happens with each turn and the resolution of mis coordinations as they are detected during another part of the task.

1 Introduction

Conversational participants use their evolving common ground to interpret and generate utterances [Sta78, Tho90, CS87]. Participants assume their common ground is coordinated, a correct assumption in most cases. The questions we wish to explore with respect to common ground address the coordination issue. We wish to characterize the processes conversational participants use to update the common ground and the processes they use to stay coordinated on what is in common ground. To characterize these processes, we must understand what is in common ground, when updates occur and by what means the conversational participants know that the intended meaning is recognized. In particular, the hearer must know that he has understood the intended meaning and the speaker must know that the hearer has understood the intended meaning.

2 Related Research

Work in cognitive science, psycholinguistics, conversation analysis and natural language processing are relevant to questions about the coordination of common ground.

First, in the field of cognitive science and psycholinguistics, Clark & Schaefer [CS87] have characterized coordination in terms of presentation and acceptance pairs. They showed how tightly coordinated speakers are in normal speech-to-speech conversations but their analysis focused more on accepting the utterance into the conversation than on whether the intended meaning was added to the hearer's common ground. Consider the states of understanding described by Clark & Schaefer [CS87]:

State 0. B didn't notice that A uttered any u'

State 1. B noticed that A uttered some u' (but wasn't in State 2)

State 2. B correctly heard u' (but wasn't in State 3)

State 3. B understood what A meant by u'

State 3 seems to be at the wrong level of detail compared to states 0-2. There are three states that involve hearing an utterance and only one for understanding it. One would think that understanding what is meant would be more complex than this (e.g. recognized the literal meaning, recognized the inferred meaning, etc.). For example, how would the speaker know that the hearer added to his common ground the possibly intended ironic meaning for “This is such a beautiful day!” when it has been raining all day, instead of the literal meaning?

Clark & Schaefer showed that every presentation must have overt evidence of acceptance. Since every acceptance is itself a presentation then there must be levels of acceptance where a strong form of acceptance (e.g. a paraphrase) will require a weaker form of acceptance (e.g. acknowledgement) until the weakest form of acceptance (e.g. continued attention) is displayed. This reflects the strength of evidence principle proposed by Clark & Schaefer [CS87]:

“Strength of evidence principle. The participants expect that, if evidence e_0 is needed for accepting presentation u_0 , and e_1 for accepting the presentation of e_0 , then e_1 will be weaker than e_0 .”

Given this lack of attention to understanding what is meant by a contribution, we suggest that the “moving on to the next contribution” form of acceptance needs more study. What are the next contributions that can serve as an acknowledgement while signalling that the intended meaning was understood and when would a next contribution signal that a coordination error has occurred in the common ground as a result of some intended meaning not having been recognized?

In the field of conversation analysis, work by Schegloff et.al. [SJS77, Sch87] on repairing errors is also relevant to the issues we are addressing. While Schegloff et.al. addressed issues of error repairs, they have not looked at error recognition as it relates to common ground. Finally in the field of natural language processing, work by Green & Carberry [GC94] and by Grosz & Kraus [GK93] is something that we can build upon by adding processing mechanisms that utilize common ground. Green & Carberry looked at indirect answers and finding the most plausible chain of hypotheses to bridge question-answer pairs, however they did not incorporate the notion of common ground. Grosz & Kraus looked at shared-plan construction which may be one type of information found in common ground.

3 An Environment for Studying Coordination

To study the coordination of common ground, we collected computer-mediated conversations between pairs of subjects who were collaboratively solving a problem. By using a computer-mediated conversation, we were able to collect conversational data that did not depend on an analysis of prosodic and facial cues. We chose a collaborative design task so that the common ground would be more narrowly defined. Our task was for the subjects to buy furniture that fulfilled a minimum functional requirement for the living room and dining room of a house

(based on tasks used in [Wal93, WGR93]). Each subject was given a separate budget and inventory of furniture that listed the quantities, colors, and prices for each available item. Although each subject's budget and inventory are private information, the subjects can share this information during their conversation. By sharing this information, the subjects can combine their budgets and can select furniture from each other's inventories. While it is conceivable that subjects might just list all of their inventory for the other subject at the beginning of the design task, the inventories are generally large enough that this was never done in practice.

3.1 Defining Common Ground

While the subjects' main goal is to negotiate the design of a two room house, the subjects also have specific task subgoals to achieve which better define their common ground. Subjects are instructed to try to meet as many of these subgoals as possible. The subgoals are:

- Match colors within a room.
- Plan to buy as much furniture as you can.
- Plan to spend all your money.
- The items of highest priority are basic furniture for each room:
 - a sofa for the living room and
 - a high table and 4 straight chairs for the dining room.

With a well-defined task, mutually-known goals define the common ground for both the subjects and the analyst and the task also allows the analyst to enumerate the possible inferences. For example:

- If no location is given for the furniture item and the item is a priority item for a particular room then infer that room as the location (e.g. sofas are in living rooms and high tables are in dining rooms).
- If no color is given for an item then infer the color matches any items already in or proposed for the room.
- If there are two subclasses for an item (e.g. high and low table), and the subclass is not given then infer it is the higher priority subclass.
- If a furniture item has already been mentioned and quantity requirements are fulfilled then infer this item is a replacement for the previously mentioned item.
- If a proposed item matches the colors in the room but is a different type of furniture then infer it is an addition to the room and not a replacement for existing items in the room.
- If a proposed item is cheaper than another item of the same type that was previously mentioned then infer it is a replacement for the previously mentioned item.

To show an example of how these inferences are used, we can construct a minimal conversation that leaves out any information that can be inferred. While subjects do leave out some information during their conversations they do not leave out everything that can be inferred. What they leave out is some varying subset of what it is possible to infer. This extreme example is just to illustrate the usage of the possible inferences:

A: green sofa 400

B: red 300

A: blue table 450

B: chairs 150 each

A: 100 each

The resulting design for this minimal conversation is that the living room contains 1 red sofa costing \$300 and the dining room contains 1 high blue table costing \$450 and 4 straight blue chairs costing \$400 total for the 4 chairs. An expanded version of this conversation might be as such:

A: I have a green sofa that costs \$400. Let's put it in the living room.

B: Well I have a red sofa that costs \$300. Since it costs less than your green one, let's put the red sofa in the living room instead.

A: Okay, we'll put your red sofa in the living room. I have a blue table that costs \$450. Let's put it in the dining room.

B: Okay, it sounds like a good idea to use your blue table. I have 4 straight blue chairs that cost \$150 each. Since they match the table, let's put these 4 chairs in the dining room.

A: I forgot to mention that I have 4 straight blue chairs that cost \$100 each. Since my chairs are cheaper than yours, let's use my 4 blue chairs.

During the conversation, each subject maintains an independent record of the current partial solution in the form of a graphical display of a floor plan. This floor plan functions as a direct representation of an important part of the common ground. By collecting both the conversations between the two subjects and their private copies of what they think the current agreed upon floor plan is, we have two sources of information on what is happening to the participants' individual representations of the common ground. The analyst can tell from the utterances what the subjects intended to do and can confirm what was meant by the utterances by examining the updates made to the individual floor plans.

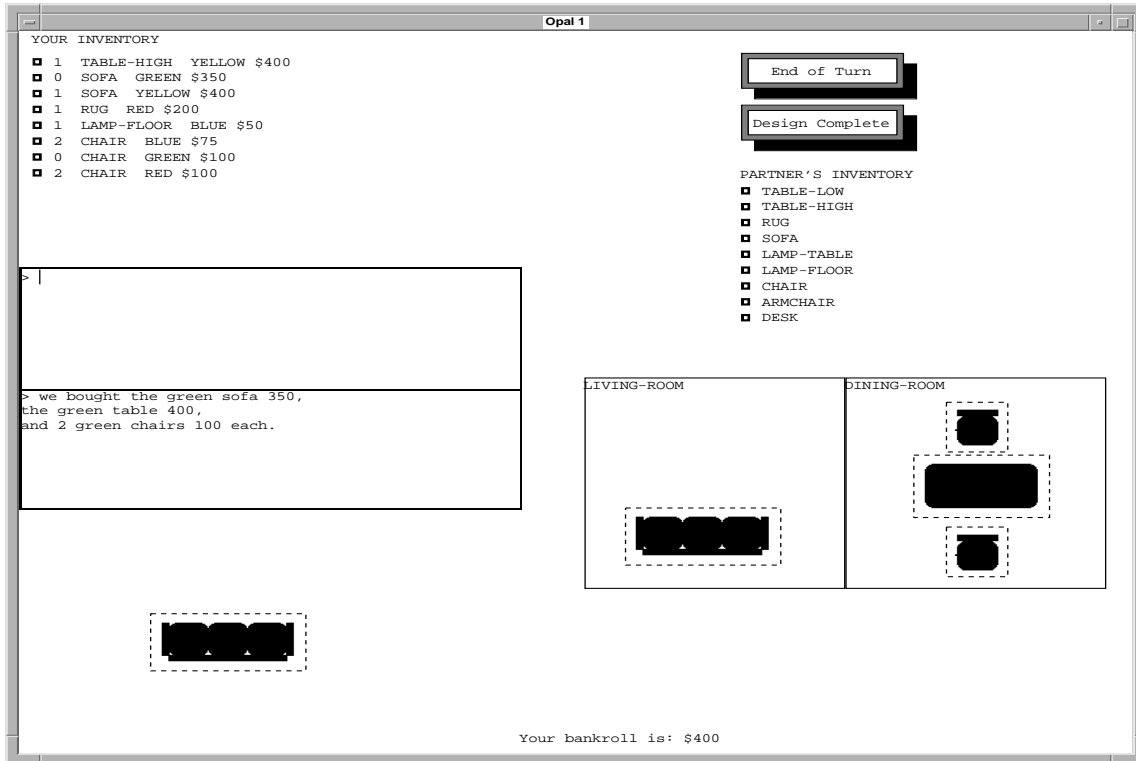


Figure 1: Matt's View Just After Receiving a Message from Dolph

3.2 The Interface

The subjects are in separate rooms and can communicate via the computer interface only. They share dialogue windows but the inventories, budgets and updated floor plans are private and show up only on the owner's display. Figure 1 shows the interface as it looks in the middle of a design session. First, we will give an overall description of the interface screen as a particular subject sees it and then we will show how the conversation progresses and how the graphics representation of the floor plan can reveal when a miscoordination arises. Looking at figure 1, in the upper left-hand corner is the private inventory that was allocated to Matt. The inventory lists the quantity of the item described by the line that is available (e.g. Matt has 1 high yellow table costing \$400, 2 blue straight chairs costing \$75 each, etc.). Matt can use items in the inventory by clicking the mouse on the square at the beginning of the line. Each time Matt clicks the mouse button on a square, the item described by the line is allocated and the number in the inventory is decreased by one. If the quantity is 0, no more items of this description are available and nothing happens as a result of clicking the mouse. When an item is allocated from the inventory, a graphics icon that represents the furniture item appears in the lower left-hand area of the screen. This icon can be dragged via the mouse into the living-room and dining-room spaces shown in the lower right-hand area of the screen. Furniture icons can be returned to inventory via mouse clicks as well.

The middle, left-hand side of the screen contains two dialogue boxes. The top dialogue box is cleared automatically, anticipating that the subject will type in a new message to send to

his partner. The lower dialogue box shows the last message that the subject received from his partner. In this case it shows the message that Matt (the subject) received from his partner Dolph.

Near the top right-hand side of the display is a menu for creating icons to represent the furniture the subject's partner is contributing to the design. This menu behaves similarly to the subject's inventory area except that there is no quantity indication and the subject must provide color and price information according to what he gathers and infers from the conversation with his partner.

The buttons above the partner's inventory menu, "End of Turn" and "Design Complete", enforce turn-taking and initiate the incremental recording of the conversation and the graphics updates. The subject is unaware of the recording effects and uses the "End of Turn" button merely as a means of transmitting his message in the top dialogue box to his partner. Once the subject hits the "End of Turn" button the subject's screen is frozen until the partner's reply is received.

Finally, the subject's budget allocation is shown at the bottom of the display. This number reflects the initial allocation of money to the subject and does not update during the session.

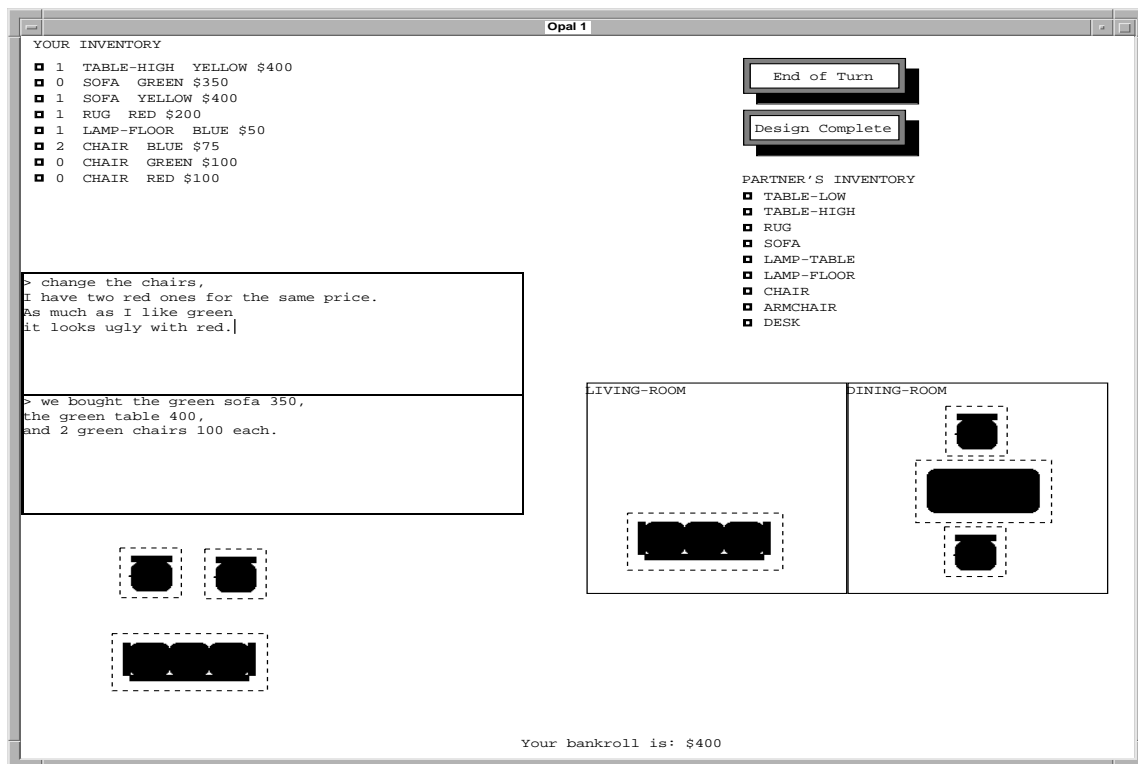


Figure 2: Matt's View After Responding to a Message from Dolph

Now that we have described the layout of the screen, we can look at what happens during the interaction. As we said earlier, the subjects Matt and Dolph are in the middle of a session. We see that Matt has allocated some of his inventory and that he has a green sofa in the living room and two green chairs in the dining room along with a red high table. He has just received

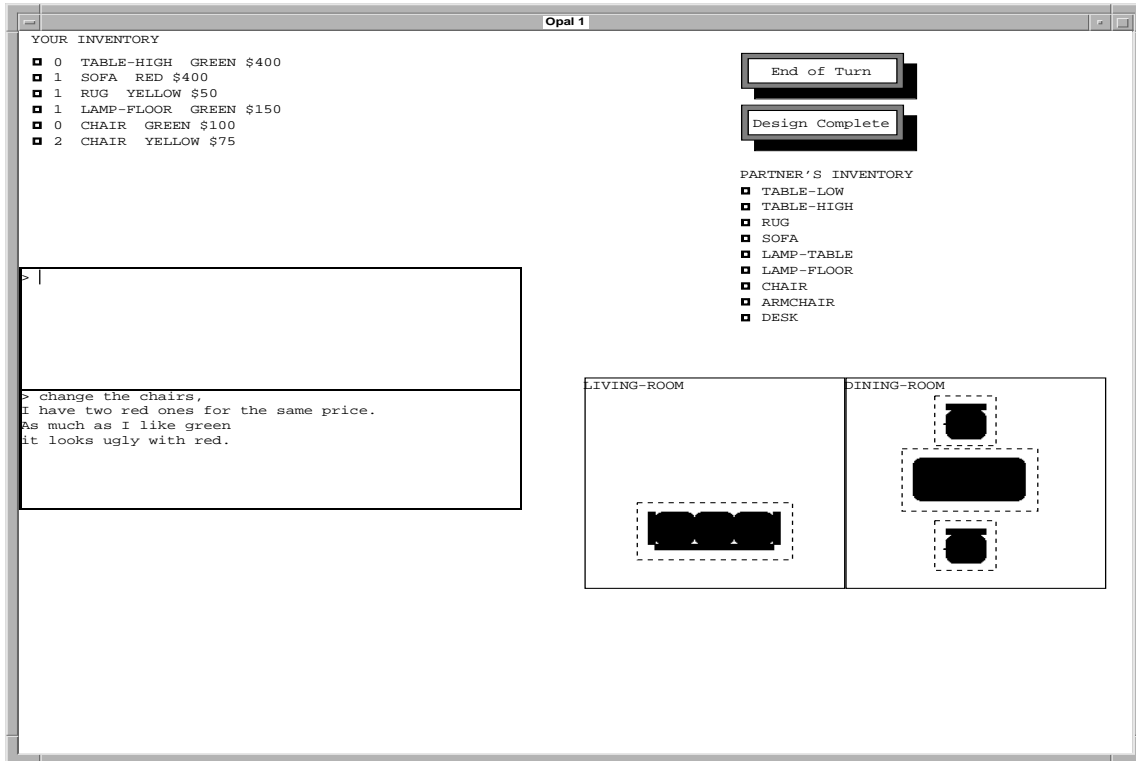


Figure 3: Dolph's View Just After Receiving a Message from Matt

a message from Dolph that summarizes the decisions they have made so far. Note that Dolph's description does not match what Matt has represented in his floor plan. Dolph has a green table while Matt has a red table. Matt now makes some updates and types a message to send to Dolph. Matt's interface just before he sends a message to Dolph is shown in figure 2.

The floor plan shows that Matt allocated two red chairs from his inventory and has used them to replace the green ones that were in the dining room. Instead of returning the green chairs to the inventory, he moves the two icons back to the lower left-hand corner as a reminder that these are items that have been mentioned before. He then types in a description of the changes he wishes to make "change the chairs, I have two red ones for the same price...". Now when Matt clicks on the "End of Turn" button his display freezes and Dolph's display is unfrozen so that Dolph may take his turn.

Matt's message is displayed in Dolph's lower dialogue box as shown in figure 3. The previous message that Matt had sent was cleared out of the lower dialogue box before the current message was displayed and Dolph's prior message to Matt is cleared from the top dialogue box in readiness for Dolph's next message.

A collection environment such as this is flexible enough to give us the opportunity to study many questions. We can modify the design task to study the effect of giving the subjects different information distributions as well as the effect of giving the players different goals. In our case, we have set up the design task to focus on the issue of coordination; the subjects have similar information distributions and the same goals.

DOLPH> "Gotta love that inflation, they give a little more money and the prices go up a lot. I have a red sofa 400 and a table high 400 with 550 to spend let's start here."
DOLPH> ADD allocated:
((TABLE-HIGH GREEN 400 DOLPH INTRODUCED))
(SOFA RED 400 DOLPH INTRODUCED))

MATT> "that party must have been rough. I am still hung over and we destroyed the furniture. a good thing we charged money. I have 400. that is it. everything is expensive. lets just get the basics. i have a sofa, green for 350."
MATT> ADD living-room contents:
((SOFA GREEN 350 MATT INTRODUCED))
MATT> ADD dining-room contents:
((TABLE-HIGH RED 400 DOLPH INTRODUCED))
MATT> ADD allocated:
((SOFA RED 400 DOLPH INTRODUCED))

DOLPH> "I don't remember much either, so it must have been a good one as I have a headache myself. the sofa sounds good, i have 2 green chairs we could put with the diningroom table 100 each. and that would spend all of our money. Next party we should either charge more or invite more people so we have more money!!"
DOLPH> ADD living-room contents:
((SOFA GREEN 350 MATT INTRODUCED))
DOLPH> ADD dining-room contents:
((TABLE-HIGH GREEN 400 DOLPH INTRODUCED))
(CHAIR GREEN 100 DOLPH INTRODUCED)
(CHAIR GREEN 100 DOLPH INTRODUCED))
DOLPH> DEL allocated:
((TABLE-HIGH GREEN 400 DOLPH INTRODUCED))
(SOFA RED 400 DOLPH INTRODUCED))

Figure 4: Examples of Local Coordination and An Enduring Miscoordination along with the Subjects' Updates to the Floor Plan as Captured by the Collection Interface

4 The Preliminary Analysis

We have collected data during a number of preliminary trials while we were testing the interface, the design problem, and the instructions to the subjects, and then during 18 design sessions using 6 subject pairs from a pool of undergraduate students taking an introductory psychology course. The undergraduate students did not know one another and did not meet until after we had finished our data collection trials with them. In addition to asking for students who did not know one another, we also asked for students who were acquainted with the Macintosh interface so that they would be able to quickly learn to use our click and drag interface.

A preliminary analysis of this data led us to hypothesize that there are at least two levels of coordination opportunities. These coordination opportunities can be characterized as immediate and delayed coordination. Immediate coordination occurs at each turn. The hearer

.

.

MATT> "change the chairs. I have two red ones for the same price. as much as i like green it looks ugly with red."

DOLPH> "what is red?"

MATT> "the chairs. I forgot that we had them. The alcohol is still in my system"

DOLPH> "You must have had a red table that you did not tell me about because I had a green one for400 that is why i put the green chairs"

MATT> "wait a minute. I have on my screen, a green sofa, red table and 2 red chairs. I have the sofa and the chairs, you gave me the red table. So, it is a green table, not a red one?"

Figure 5: An Example of Delayed Coordination (Recorded Conversation Minus Any Updates to the Floor Plan)

will either make inferences or ask questions in order to fill in any missing information. We hypothesize that a hearer will ask questions when he is unsure of the inferences he has made or when he is unable to make the necessary inferences. Delayed coordination, on the other hand, arises from enduring miscoordinations. These enduring miscoordinations may be caused by "speech" errors (in our case, it is a cross between a typo and a speech error) or by erroneous bridging inferences. The delayed coordination appears later in the plan when a conflict occurs at more abstract levels of the plan. For example, the more abstract level of the plan may be an evaluation of the plan to see if the colors in a room match. When a conflict is recognized, this initiates a repair to the plan. The repair is accomplished either by summarizing the design as the subject understands it to be, questioning the conflict, or both.

Figures 4 and 5 are excerpts from the data that we collected during a session with some undergraduate subjects that supports this preliminary analysis. The excerpts are part of the session that led to the interface screens we showed earlier. In figure 4, note that Dolph mentions a red sofa, a high table with no color indicated, and tells his partner what his budget allotment is. Matt, in the next turn, reciprocates with his budget allotment, an overall evaluation of the situation and mentions a green sofa that he has. Looking at the graphics update we see that Matt intended his mention of the green sofa to be a proposal to put the sofa in the living room and we also see that he accepts Dolph's mention of the high table as a proposal for the dining room. Matt has inferred that Dolph intended the red sofa as a proposal for the living room and the table as a proposal for the dining room. However, Matt intends the green sofa as a replacement for the red sofa since it is cheaper. In addition Matt has inferred that the table is red, since Dolph did not tell him the color. However, we can see from Dolph's earlier graphics update that he meant the color of the table to be green but he neglected to mention the color (i.e. a speech error/typo). In this case Matt has made an erroneous bridging inference in order to supply the color value. As we will see, this will be the source of a delayed coordination later in the conversation.

In Dolph’s next turn, he mentions that he has two green chairs to go with the dining room table. This is an opportunity for Matt to perform a delayed coordination but he misses the opportunity and it is not until later in the conversation (as shown in figure 5) that the delayed coordination opportunity is recognized by Dolph.

In figure 5, Matt indicates that he would prefer to use the red chairs instead of the green ones since green and red do not match. Dolph then wonders what furniture item in the dining room is red since everything in his floor plan is green. Matt still does not realize there is a miscoordination and just repeats the information that he has red chairs. He goes on to justify why he did not mention them earlier by saying that he forgot to consider them. Dolph then makes an inference about why Matt wants red chairs and what the clashing red item was for the green chairs Matt wants to replace. Dolph tells Matt what he thinks the miscoordination is “you must have had a red table...”. Matt finally recognizes that there is a miscoordination and the repair is initiated. He summarizes what he has in his floor plan and asks for confirmation of the repair.

5 Summary and Future Directions

In this paper we have described a collection environment that allows us to influence the resulting conversation so that we may focus on particular research problems. By choosing a collaborative design task where the conversational participants have similar knowledge distributions and the same goals, we were able to focus on coordination issues. We could define the common ground from these goals and were able to enumerate the possible bridging inferences that could be expected and these bridging inferences were confirmed by the data that we collected. We could hypothesize with greater certainty the changes to the common ground and the source of these changes since we had the utterances as well as a graphic representations of what was intended and what was understood. We were able to identify two levels of coordination opportunities, immediate and delayed. The immediate coordination resulted in updates to common ground and the delayed coordination opportunities resulted in repairs to common ground.

As our next step, we will do a more detailed study of the instances of underdetermined proposals and their contexts. This study will enable us to discriminate between those situations where a conversational participant is confident about a bridging inference and those in which the participant needs to ask a question. It will also characterize the contexts in which a participant is more likely to detect an enduring miscoordination. And finally, the study will help us characterize the contexts in which a participant leaves out information. The omission may be unintentional; the participant forgets to provide the information, or it may be intentional; the participant expects his partner to be able to easily infer what it is. We need a characterization of the situations in which both intentional and unintentional omissions arise.

References

- [CS87] Herbert H. Clark and Edward F. Schaefer. Collaborating on contributions to conversations. *Language and Cognitive Processes*, 2:19–41, 1987.
- [GC94] Nancy Green and Sandra Carberry. A hybrid reasoning model for indirect answers. In *Proceedings of ACL*, 1994.

- [GK93] Barbara Grosz and Sarit Kraus. Collaborative plans for group activities. In *IJCAI-93*, 1993.
- [Sch87] Emanuel A. Schegloff. Recycled turn beginnings: A precise repair mechanism in conversation's turn-taking organisation. In Graham Button and John R. E. Lee, editors, *Talk and Social Organisation*. Multilingual Matters Ltd, 1987.
- [SJS77] E.A. Schegloff, G. Jefferson, and H. Sacks. The preference for self-correction in the organisation of repair in conversation. *Language*, 53:361–382, 1977.
- [Sta78] Robert C. Stalnaker. Assertion. In Peter Cole, editor, *Syntax and Semantics, Volume 9: Pragmatics*, pages 315–332. Academic Press, 1978.
- [Tho90] Richmond Thomason. Propagating epistemic coordination through mutual defaults I. In R. Parikh, editor, *Proceedings of the Third Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 29–39. Morgan Kaufmann, 1990.
- [Wal93] Marilyn A. Walker. *Informational Redundancy and Resource Bounds in Dialogue*. PhD thesis, University of Pennsylvania, 1993.
- [WGR93] Steve Whittaker, Erik Geelhoed, and Elizabeth Robinson. Shared workspaces: How do they work and when are they useful? *IJMMS*, 39:813–842, 1993.