Coordination of Talk; Coordination of Action

Richard Alterman, Alex Feinman, Seth Landsman, Josh Introne
Department of Computer Science
Brandeis University
Waltham, MA 02454 USA

Technical Report: CS-01-217
Revised October, 2001

ABSTRACT

In this paper, the analysis of "talking about things" is framed in terms of the interdisciplinary literature on coordination in "face-to-face" interaction. The working assumption is that talking about things is an interaction that is embedded in a larger (social and historic) interaction that has a purpose. The focus is on the coordination work that goes on during a same time/different place computer-mediated cooperative activity in support of cooperative work that involves shared domain objects. The main work of this paper is to explore the ramifications of (and methodology for) introducing secondary structures that are realized in specific data structures (called coordinating representations) to help organize both conversations about things and the cooperative behaviors involving them.

INTRODUCTION

Within any community of actors who have a common set of overlapping goals and tasks, and who must coordinate their efforts in accomplishing these tasks, there is an ongoing design process at work that creates secondary structures for organizing behavior. In future cycles of behavior, as the designs become debugged and deployed, they simplify the reasoning of the participants and improve their performance. For non face-to-face activities these designs can become realized in a physical structure that helps solve a problem of coordination: these will be referred to as *coordinating representations*.

This perspective frames our ideas and analyses of how participants in a computer-mediated collaboration talk about shared domain objects. From this vantage point, the coordination of activity involving “things” requires the coordination of talk about them. For regularly occurring patterns of online interaction, secondary structures that organize talk naturally emerge and serve to improve performance in several ways. Coordinating representations can be specifically designed for a recurrent online coordination problem.

The first part of the paper develops our theoretical framework. Some of the ideas and terms that are defined and discussed include the problem of coordination, the role of context and structure in organizing behavior, and the emergence of convention. Two features of coordinating representations – as external representation and as mediating artifact – are used to situate the idea of a coordinating representation into the interdisciplinary literature on distributed cognition. The notion of a coordinating representation is contrasted to other methods that have been developed in CSCW that add structure to improve coordination in online communication. The model of communication that is assumed features the interactive and contextual elements of synchronous communication, but within a framework that focuses on talk that develops to support conventional activities. A crucial issue is to develop a methodology that will reliably produce effective coordinating representations for particular groupware systems.

The discussion of these issues will be developed in the context of a same time/different place groupware system (VesselWorld), where the root form of communication is electronic chatting. In VesselWorld, three participants engage in a computer-mediated problem solving session. To achieve a set of tasks in a simulated environment the participants must communicate and jointly problem-solve. The only way they can communicate with one another is via the computer. Access to the environment, and objects in the environment, is also mediated through representations provided via the computer. A task cannot be broken down into pieces solvable by an individual actor working alone. The problem solving sessions require joint sensemaking, cooperation, and coordination. Most of the participants’ task (and communication about their tasks) is concerned with objects existing in the simulated environment.

In a base version of the VesselWorld system, the participants can only exchange information using a text-based chat relay. To create useful structures for representing information about objects that supports communication among the participants requires an analysis of the participants’ discourse while solving problems using the base system. There are three indicators for the development of coordinating representations that are examined; each of them reflects the theoretical framework that is developed. What are the recurrent areas of behavior where coordination requires increased effort? What sort of errors do the participants make? What secondary structures for organizing conversation develop among the participants? This method of analysis was applied to the pilot study data we collected of subjects using the base system, and then three coordinating representations were developed.

The latter parts of the paper present an empirical evaluation of the performance differences between groups of subjects that had access to, and did not have access to, coordinating representations. The analysis of the data is framed by the early discussion of the contextual, emergent, and conventional aspects of the structure of behavior. One important theme of this analysis is to explore how the participants exchange information about, and maintain consistent representations of, shared domain objects with, and without, coordinating representations. A concluding discussion section re-visits the topic of methodology.

STRUCTURE FOR BEHAVIOR

Secondary Structures for Recurrent Joint Behaviors

The view that the primary site of everyday activity is a face-to-face interaction developed within Sociology (e.g., Goffman, 1983; Garfinkel, 1967; Schutz, 1967). A significant area of analysis has been to deconstruct conversation as "talk in action" (e.g., Sacks, Schegloff, and Jefferson, 1974; Schegloff & Sacks, 1973). The work of Suchman (1986; see also Agre, 1997) developed critiques of Artificial Intelligence from this latter perspective. Within Cognitive Science and Artificial Intelligence, a number of general models of cooperative activity have emerged (Clark, 1996; Bratman, 1999; Grosz & Sidner, 1990). One topic that has gotten considerable attention in the interdisciplinary literature is how participants in a cooperative activity manage to coordinate their behavior.

Greeting someone, planning a potluck dinner party, moving through a doorway, forming a queue at the coffee shop, turn taking in conversation, and the exchange of information within an organization are all occasions of cooperative activity – and in each case problems of coordination among the participants are entailed. Some examples of the types coordination problems are the assignment of tasks (you bring the desert, I will bring an appetizer) and roles (you lead the discussion), the establishment of location (the end of a queue), sequencing (you go first), timing (let's lift together), manner (proceed slowly so as to avoid accidents), and co-reference (pointing to a part of a figure on the whiteboard). Joint decisions on matters of coordination are ubiquitous, endemic, inherent, and unavoidable. Even the participants of a heated political argument are likely to coordinate on the ground rules and boundaries of a debate.

The participants exploit the physical, social and cultural features of their "context" in order to cooperatively reach decisions about their shared behaviors (e.g., Hutchins, 1995a; Greeno, 1998). The participants share a common ground (Clark & Brennan, 1991): the presupposed expectations, facts, and referents that frame their cooperative behavior. The fact that they are co-present allows them to monitor the progress of their activity. Throughout their activity they can speak to one another; their comments to one another are exchanged without delay, in the course of their joint behavior. Because they can see one another, they can also use body position and language, as well as gaze, to communicate information. The actions that form their conversation and activity occur sequentially.

The term *structure for behavior* is used here to refer to the kinds of organizational information that support coordination and shared among participants. Examples of structures established as part of the common ground in order to organize behavior are agreements about the assignment of roles, the path, the manner of an activity, a schedule that orders a set of joint actions, et cetera. Turn taking is an example of a structure that emerges from the interaction among the participants, but does not exist prior to conversational interaction (Sacks, Schegloff, and Jefferson, 1974). For the difficult portions of their task, the participants may explicitly create a shared plan (Grosz & Sidner, 1990), an agreed to structure – you do this and I'll do that – for the behavior. Not all the information exchanged serves the purpose of organizing the current behavior. Nor are all structures for joint behavior exchanged at runtime: participants who share common ground are likely to have prior experience with activities that recur.

Cognitively, the joint behavior of two actors engaged in a cooperative activity can be modeled as having both individual and social parts (Alterman & Garland, 1998). Individually, each actor reasons from prior episodes of joint behavior, using them as the basis from which to construct the organization of the current behavior. At runtime these expectations can be confirmed or adjusted by means of a social interaction. Using both the social exchanges of information about structure and the recollection of prior related experiences, the participants must jointly reason out and construct a behavior which achieves their shared goals. For recurrent problems of coordination within a community of actors, a design process is at work. The goal of this design process is to produce secondary structures that will better organize behavior, in the future, for similar sorts of activities. Over time, conventions for coordinating behavior during the course of a recurrent cooperative activity within a community of actors begin to emerge. These conventions are designs for the structure of recurrent cooperative activities. Prior knowledge of a convention creates expectations that must be either confirmed or revised as the current activity progresses.

Coordinating Representation

Some designed structures for behavior become realized as materials objects, available at the scene of a non face-to-face cooperative activity, to signify an organization for a particular sort of behavior – we will refer to these as *coordinating representations*. The stop sign is an everyday example of a coordinating representation. It is a material object that signifies an organization for an expected behavior. The stop sign is a representation shared among the participants at a traffic setting. The stop sign presents a structure for organizing the collective behavior of drivers, pedestrians, and cyclists at a busy intersection. The interpretation of the structure imposed by the stop sign is negotiated during the activity. Things may run smoothly at the intersection – but there will also be interruptions. An impatient driver piggybacks on the driver in front of him. A pedestrian decides to ignore the stop sign altogether.

A coordinating representation is an external representation shared among participants in a joint activity. It is designed for the activity at hand and reduces the complexity of the coordination task. It mediates and structures the activity. It has the designated purpose of helping participants to achieve coordination in non face-to-face cooperative activities. Its meaning is based on conventional interpretation. It signals to the participants – without dictating action – that a conventional organization of behavior is in place.

An important topic in Cognitive Science has been the role of external representations in cognition. The general scheme is to expand the unit of analysis for the study of cognition to include a person, plus the context of their activity (*distributed cognition:* Hutchins, 1995; *person plus:* Perkins, 1993; *epistemic acts:* Kirsh, 1993; *context:* Lave, 1988). Context, knowledge, memory, reasoning, the structure of behavior, and "seeing" depend on interaction with external representations that are available in the task environment (e.g., Lave, 1988; Hutchins, 1995ab; Goodwin & Goodwin, 1996). By virtue of the fact that a coordinating representation is external (and material), it functions to distribute some of the coordination work of the participants in a cooperative activity into the design of the task environment.

As a mediating artifact, a coordinating representation has both a sign and a tool function (Vygotsky, 1968; Cole & Engestrom, 1993). As a sign it modifies how the individual thinks and as tool it changes the ways in which the collaborating actors proceed. Mediating artifacts are critical in the accumulation of modifications to practice over generations of actors within a community (see also, Norman, 1988; Alterman, Zito-Wolf, and Carpenter, 1998). By virtue of its role in the accumulation of improvements and changes in behavior, a coordinating representation codifies a "solution" to a recurrent problem of coordination.

In the everyday world, there are many examples of coordinating representations that serve the function of organizing the exchange of information. An appointment slip helps a patient to return to the dentist's office on the right day at the right time. A mail order catalogue helps the customer and the sales office reach agreement on purchase items, sizes, and prices. Jargon is created to encode information about common situations. In each of these cases, an exchange of information of a particular type is organized by a structure that was designed to coordinate that sort of exchange of information. The remainder of the paper will explore the consequences of, and a methodology for, introducing coordinating representations into a computer-mediated activity so as to improve the participants' exchange of information about shared domain objects for recurrent activities.

Computer-Mediation

Within the literature on CSCW, computer-mediated communication has engendered a fair amount of research. For synchronous communication, the canonical example is to convert an everyday task of several actors engaged in planning out some kind of activity in front of a whiteboard into a task that could be computer-mediated. Given the shared workspace, two issues of interest are how the participants in such an activity organize their talk, and how they organize their task. For asynchronous communication, general-purpose structures that effectively coordinate the exchange of information have been produced. The idea is that these "languages" are general enough to be re-usable for more than one application.

WYSIWIS (What You See Is What I See) systems (such as Cognoter: Stefik et al., 1987) provided a shared virtual "whiteboard" for exchanging ideas, drawings, and texts during a same time different place computer-mediated activity. Experimental studies examined and evaluated the function and utility of whiteboards both outside and inside the context of computer-mediation; these studies produced insights and constraints that inform the design and development of groupware applications (e.g., Greenberg, Roseman, and Webster, 1992). The drawing space itself is an important resource for the group in mediating their collaboration (Tang, 1991). Permanent media such as the whiteboard enable users to construct shared data structures around which to organize their activity (Whittaker, Brennan, and Clark, 1991). Both speech and shared data structures were used to represent some of the content of the communication among the participants, but the coordination of talk (e.g., clarifying references) which required more interaction were primarily accomplished using speech. Collaborators used hand gestures to communicate significant information; process of creating and using drawings conveys much information not contained in the resulting drawings (Tang, 1991; Bly 1988). Communication that is supported by external representation media (i.e., the whiteboard) can be modeled using an interactive model of communication (Tatar, Foster, and Bobrow, 1991). Making the creation of a contribution available to all participants, and not just its final form, improved communication among the participants. The utility of adding a shared workspace to audio depends on the complexity of the task (Whittaker, Geelhoed, and Robinson, 1993).

Where a whiteboard has general properties of interest that characterize all kinds of social interaction supported by external media, a coordinating representation is one kind of content realized in an external media. It is specifically designed for a particular context and it addresses a problem of coordination that emerges in the performance of a recurring cooperative activity across sessions of cooperation. For some applications, general mechanisms that support communication and coordination, like the whiteboard, will suffice. But for tasks that occur over extended periods of time, the introduction of a coordinating representation will potentially improve the interaction among the participants during the recurrent problematic areas of joint behavior.

Other investigations in CSCW focused on specific kinds of structures that could potentially better organize online asynchronous communication for all kinds of computer-mediated applications. One example is The COORDINATOR (Flores et al., 1988), which structured, in such a way as to make explicit, the negotiation of commitments among the participants of a collaborative activity; The COORDINATOR assumed a language-action framework for the analysis of communication. Another example was the Information Lens system (Malone et al., 1989); it added more structure to the exchange of email. The claim of the Information Lens system was that the additional structure would help users to sort, filter, and prioritize their messages, and find messages that they might not otherwise receive. NoteCards (Trigg, Suchman, and Halasz, 1986) added structure to the online sharing of information on an ongoing project. The NoteCard system is a hypertext environment, where the users can construct notecards that contain text, graphics, or images. The notecards can be arranged into hierarchies or arbitrarily linked together forming networks of ideas (as they represented on the individual notecards).

These latter sorts of mechanisms are general prescriptions that apply to any number of potential groupware applications. Collectively these methods of adding structure to communication provide groupware system designers and builders with a toolbox of representational devices for engineering online coordination in the asynchronous exchange of information. Because they are general, the designer of a groupware system is confronted with a methodological gap between a set of representational tools, on the one hand, and a reliable analysis of the specific coordination problems that exist for a given application on the other. The coordinating representation tackles the problem of adding structure to online communication from the other direction. A coordinating representation is a fix for a coordination problem that emerges for a recurrent activity within a community of actors; it is not a general solution to the problem of coordination but rather it depends on the specific context of activity. It can be represented in any number of shared data structures. But the key problem, the crux of the matter, is how to identify the critical problems of coordination and then how to address them.

Summary

In everyday activities, the participants of a cooperative activity must continuously solve problems of coordination; a problem of coordination occurs when increased activity is needed for the participants to continue the cooperative behavior, or where errors in performance frequently occur. Structures that support, simplify, and organize behavior can be expected, planned, and/or remembered during the course of an activity. Using both the social exchanges of information about structure and the recollection of prior related experiences, the participants must jointly reason out and construct a behavior which achieves their shared goals. Over time, for joint activities that recur, secondary structures for those behaviors – which function to reduce the complexity of the behavior – begin to emerge. For non face-to-face activities, designed structures for joint behavior become realized as material objects that signify at the scene of activity an organization for a certain kind of behavior – these representations will be referred to as *coordinating representations*.

The research within CSCW on the role of external media in face-to-face activity has developed a set of design criteria for re-creating face-to-face interactions online. The work (within CSCW) on general approaches to structuring communication provides designers with a toolbox of representation devices that can be further developed to meet the demands of any number of groupware applications. Both of these are general prescriptions that apply to any number of potential groupware applications. A coordinating representation is a fix to a recurrent problem of coordination in a non face-to-face cooperation that improves performance beyond what can be achieved using general mechanisms. The research we describe in this paper introduced coordinating representations so as to organize talk, so as to coordinate action. From a methodological perspective, the crux of the matter was how to identify the critical coordination problems in the rapid and effective exchange of information during an online collaboration and then implement a set of coordinating representations to simplify those interactions.

INTRODUCING COORDINATING REPRESENTATIONS INTO A COMPUTER-MEDIATED ACTIVITY

The first part of this section presents an overview of VesselWorld (VW), the testbed we used for exploring issues in developing groupware systems that include coordinating representations. VesselWorld is specifically designed such that the participants must engage in a collaborative online problem-solving session. All exchanges between the participants are mediated by the system. Access to the environment, and objects in the environment, is also mediated by system. In a base version of the system, the participants can only collaborate by electronic chatting.

The latter parts of this section discuss how we developed three coordinating representations to add to the base version of the system. Many of the themes introduced in the preceding section on coordination, secondary structures and convention, and shared data structures will be elaborated upon. A critical feature of the methodology that was used to develop the coordinating representations was that it included both an analysis of social interaction among the participants and an analysis of the individual interacting and reasoning vis-à-vis a coordinating representation. This feature of the methodology reflects our commitment at the outset of the research that developing technology to support groupware systems requires an analysis of the individual's reasoning and decision-making in a context of the larger social interaction. Ideas for potential coordinating representations were generated from an analysis of the pilot study data[1], interviews with users, and our own experiences with the system; in each of these cases our analysis focussed on problem areas of coordination and the development of secondary structures by participants. Another important source of intuition was prior theoretical work on cooperative joint activities. In the discussion section at the end of the paper, we will revisit the question of methodology and propose some refinements, extensions, and tools that would improve upon the methodology we used in this initial study.

[1] The data for this pilot study was collected by S. Kirschenbaum at NUWC.

VesselWorld

For the last several years we have been building a same time/different place groupware system (VesselWorld) as an experimental platform for analyzing real time computer-mediated collaborations. A demo of the system was run at CSCW 2000 (Landsman et al., 2000). We are also extracting from VesselWorld a framework (called ADAPTIVE) for the rapid development of distributed, component-oriented groupware systems.

In VesselWorld, three users, situated at three physically separate computers, engage in a set of cooperative tasks that require the coordination of behavior in a simulated environment. In the simulated world, each participant is a captain of a ship, and their joint task is to find and remove barrels of toxic waste from a harbor and load them onto a large barge. Two of the users operate cranes that can be used to lift toxic waste from the floor of the harbor. The third user is captain of a tugboat that can be used to drag small barges from one place to another. The cranes are able to individually lift and carry small or medium toxic waste barrels, jointly lift large barrels, and jointly lift (but not carry) extra large barrels. The tugboat cannot lift barrels, but can attach to, and move, small barges. Small barges can hold multiple barrels, but have a limited carrying capacity. Each captain has a small radius of perception. Many barrels require the use of additional equipment (either a dredge or a net); one of the cranes has the dredge and the other has the net. The tugboat captain is the only one who can examine barrels to determine equipment needs. Barrels can be leaking – and will begin to leak if the users make a mistake – in which case the leak must be contained by the tug.

There are three participants (as opposed to two) in a VesselWorld problem solving session mainly because it increases the complexity of the collaboration and increases the cognitive load of the participants. It was also believed that this feature would make our results more generalizable than a system that had subjects work in dyads; for example, this arrangement allows participants to form subgroups.

The participants communicated in the base system via electronic chatting, instead of audio, because chatting leaves a reviewable record of the information being exchanged. This reduces the load on working memory. It also means the recipient of a message can delay accessing the information it contains until attention is freed up from some other task; this is especially important given the multi-tasking nature of the VesselWorld task. It is also significant that electronic chatting and messaging is pervasive for synchronous activities. Finally this simplified data collection and analysis.

In VesselWorld segments of activity are divided into *rounds.* During a round of activity, any number of *events* may occur: participants can chat, examine any or all panels of information, and plan out their actions. Once a participant has submitted to the system her next action, she can no longer change it. When all three participants have submitted actions, the round ends, the system updates the state of world, and the next round begins.

The WorldView (the large window in Figure 1) is a segment of common ground that graphically represents several kinds of information about the location and status of objects from the perspective of an individual participant. The WorldView depicts the harbor from the point of view of a participant, who can only see a limited region at one time. A second window of information is used for planning. A third window allows a user to access more detailed information about visible objects. A chat window allows participants to communicate with one another using an electronic chat.

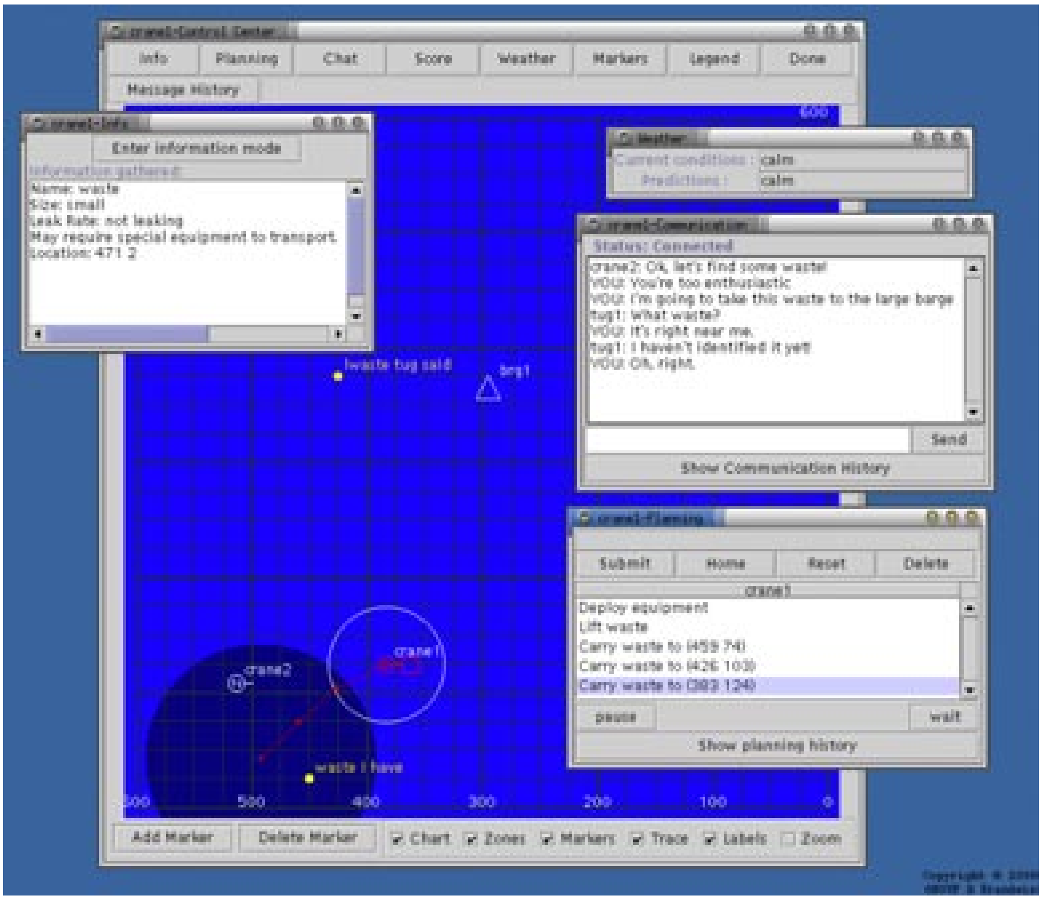

Figure 1: The interface for the basic system.

During a VesselWorld problem solving session most of the participant dialogue is centered on the barrels, and how effort can be coordinated in removing the barrels from the harbor and transporting them to a large barge. Figure 2 summarizes the information that must be shared about the domain objects.

The participants must keep track of what areas of the harbor have (or have not) been searched. The participants must discover and then keep track of the location of wastes. Initially this is the location of a waste in the harbor; later this includes whether a waste has been moved and to where, and if it is on a small barge, in what order it was stacked. References to the wastes must be shared; these references can change depending on the circumstances. The equipment needed to handle a waste must be reported by the tug and is relevant to the actions of both cranes. Information about whether a waste is leaking must also be distributed. The location of the small barges is relevant to the actions of all the participants. Knowing what wastes are currently on a small barge and their cumulative weight determines whether or not additional wastes can be added to the barge or if the barge first needs to be unloaded. For a session to be completed successfully all the wastes must be removed from the harbor and transported to the large barge, so the location of the large barge must be known by all participants. The participants need to keep in mind which crane has which equipment and whether the equipment has been deployed for a given waste. An effective organization of their joint activity requires that the participants keep track of each other's locations. Coordination of participants in the joint removal of wastes also requires that the participants know, at the relevant time, the intended next action or the recently completed action of the other actors.

1. Harbor
   a. Areas that have been searched
2. Wastes
   a. The location of wastes
   b. References to the wastes
   c. The equipment needed for handling the wastes.
   d. The size of the waste (thus determining if it takes two cranes to lift)
   e. Whether or not it is leaking
3. Small barge
   a. Location
   b. What is on it
4. Large barge
   a. Location
   b. What is on it
5. Equipment
   a. Who has which equipment (this is permanent)
   b. Whether the equipment is currently deployed
6. Other Actors
   a. Their location
   b. Their intended future action for (or recently completed action with) some shared domain object

Figure 2: Information about domain objects that is shared amongst the participants.

There are several important characteristics of the cooperative activity of participants in a VesselWorld problem solving session as they jointly manage the removal of wastes from the harbor. Participants have different roles (both predefined and emergent). Cooperation and collaboration are needed to succeed. Participants must develop a shared understanding of an unfolding situation to improve their performance. Uncertainty at the outset makes pre-planning inefficient in many circumstances. There are numerous problems of coordination.

Methodological Considerations

From the point of view of the designer, the primary task is to situate the technology under development into the larger context of a specific task and community. The methodological process for developing coordinating representations begins after an initial analysis of the work environment has been completed. For some collaborative computer-mediated tasks general-purpose mechanisms like whiteboards or electronic chat rooms will suffice. For other applications, these general-purpose mechanisms will not provide enough structure for the task. In these cases, the designer’s goal is to develop an analysis of the interaction among the participants that highlights the critical features of a particular cooperative activity. The question then is – how to proceed?

The development of a coordinating representation requires a stage of analysis where the social interaction among the participants is investigated, and then a stage of analysis where the initial approximations for coordinating representations are converted into external representations that are shared among the users.

Of the two stages of analysis, the second one seems a bit more straightforward. It has been widely argued that external representations, and their design, play a significant role in both human computer interaction and, more generally, cognition (e.g., Vygotsky, 1978; Norman, 1991; Perkins, 1993; Hutchins, 1995b). The extent to which these shared data structures are external representations is the extent to which the design task at this point becomes the province of traditional HCI. In other words, during the second phase of development, coordinating representations are structured in a manner that reduces the physical work and the cognitive load of the individual user in creating and accessing the coordination information that is shared among the participants.

The first stage of analysis – the social part – is more problematic, as there is no existing formal methodology that fits the task. The coordinating representations that are needed will depend on the character of the interaction that emerges from the participants, once there is a tool in place and they begin to collaborate. This necessitates that a basic version of the system should be deployed first, one that only includes a general-purpose coordination mechanism (e.g., whiteboard or electronic chat). After the basic system is deployed and pilot study data is collected, a discourse analysis can proceed which will identify the critical problem areas of coordination for the participants.

Discourse Analysis: The social part of the analysis

The discourse analysis of participant dialogue taken from the pilot study of VesselWorld focussed on three indications that the introduction of a coordinating representation may be advisable.

1. Recurrent patterns of coordination.
2. Repeating errors.
3. The development of secondary structures to organize talk.

Both recurrent patterns of coordination and repeating errors are the kinds of situations in which the participants potentially would want to introduce secondary structures to better organize activity so as to improve performance. The last indicator goes a step further; in this case, the participants have both determined a potential area of improvement and have devised a structure to improve the situation.

Of the three indicators the third indicator is the surest bet for the analyst; this is because the participants have added corroborating evidence for a particular analysis of the situation. There are problems, however, with an analysis strategy that exclusively relies on the third indicator. First of all, the pilot study data would have to be more extensive so the participants would have time to generate all the most useful secondary structures, and there is no telling just how extensive it would need to be. Another problem is that the participants may identify a coordination problem they would like to fix, but do not have the means necessary to fix the problem; there is some evidence for this in our data. A final problem is that introducing a coordinating representation where a secondary structure already exists may have no additional benefit.

One final cautionary note is that there is no guarantee that either the secondary structures that the users add or the coordinating representations that the designer adds will improve the situation at all; it is entirely possible that some problems of coordination are best dealt with using a free form of communication like electronic chatting.

Adding Three Coordinating Representations to VesselWorld

The electronic chatting amongst participants during a VesselWorld problem solving session was used as a basis for developing some coordinating representations. As the discourse was reviewed, the analyst looked closely at using coordinating representations to simplify the most common interactions, fix repeated errors in coordination, and replace conventions developed by users during the course of a problem solving session. The goal was not to entirely replace other forms of communication with coordinating representations.

Rather the analyst wanted to use coordinating representations to simplify the interaction among the participants at the critical points in the ongoing cooperation among participants.

The analysis was framed by cognitive work on the problem of coordination that was presented earlier. Figure 3 shows a list of the kinds of methods that were used by participants to coordinate their joint activities. The participants did some planning by assigning roles or agreeing to sets of actions. During the activity, a fair amount of chatting was used to initiate joint actions that were tightly coupled; for example, to lift an extra large waste, the cranes have to begin lifting during the same round. Also found in the discourse were examples of the participants creating conventions to simplify the exchange of information for recurrent problems of coordination. Chatting was continuously used throughout each session to establish references and exchange information about shared domain objects.

- Plan (provide orientation: delimit tasks)
  - Plan to do; Role assignments
- During Activity (Entry & Exit into Phases)
  - Synchronization; sequencing; step; turn-taking; Action taken; See; Initiating Statement
- Develop conventions
  - For both action and talk
- Co-Referencing and the exchange of information
  - Refer to status, location, feature, identity of object

Figure 3: Taxonomy of coordination methods

Figure 4 shows a sample dialogue of the kind of close coordination users needed to do in order to time closely coupled activities. At 1 and 2, after jointly lifting a large barrel, Crane1 and Crane2 agree to do a joint carry followed by a joint load onto a barge. It will take three moves to reach their destination. In lines 3, 4, and 5, they tell each other they submitted their first move. At 8 the tug suggests a convention to simplify coordination. At 9 and 10, Crane1 and Crane2 tell each other they are ready to do the second part of the move. At 14, Crane1 states she is doing the third move. At 15-18 they plan, and then they submit actions, to do the joint load. At 19 and 20, they celebrate.

1. Crane1: now a joint carry, clicked at 375,140 got 3 carrys
2. Crane2: i will do same
3. Crane2: move to first location
4. Crane1: submitted first
5. Crane2: ditto
6. Crane1: again?
7. Crane2: yes
8. Tug1: do you want to just type something in after submitting each turn
9. Crane1: submitted second
10. Crane2: ditto
11. Tug1: just some shorthand or something, for everyone so we know whats going on
12. Crane1: submitted third
13. Tug1: submitted
14. Crane2: submitted third
15. Crane2: Crane1: load, and then i'll to the same
16. Crane1: submitted load
17. Crane2: ditto
18. Tug1: submitted move
19. Crane2: hey, i think that worked!

20. Crane1: looks like it's Miller time. I think we did it.

Figure 4: A sample dialogue.

The analysis of the pilot study discourse identified three recurrent areas of coordination activity:

1. Timing of closely coupled cooperative activities involving the toxic wastes
2. Establishing references for, and exchanging information about, shared domain objects and their status.
3. Higher-level planning to manage multiple cooperative activities in searching the harbor and organizing the removal of all the wastes

Each of these has been suggested by prior theoretical analysis. But there are also other potential problem areas. The problem was to determine which things were most problematic for the task-at-hand.   Associated with all three of these recurrent areas of coordination were participant errors and increased effort. For the first of these, the participants in the pilot study generated secondary structures that simplified talk about these coordination issues.   In the case of high-level planning, there was a fair amount of talk on this topic throughout the sessions, but we were unable to detect any secondary structures that were created in order to simplify communication.

In a second version of the VesselWorld system, coordinating representations were introduced that structure and simplify the exchange of information in the problem areas of coordination.

All the coordinating representations were designed to allow the participants to better manage the mutual exploration of (and joint operations upon) the shared domain and the objects it contains. The coordinating representation embeds the representation of shared domain objects in a structure that simplifies coordination at two levels: 1) in the conversational exchange of information about the shared objects; and 2) by organizing some of the joint actions that are taken with regard to those objects.

Some sketches of three coordinating representations were developed and later refined through an interview with one of the test groups in the pilot study. After the interview, the iterative design process continued by a cycle of (re)design, implementation, and evaluation. The periodic evaluation came in several forms, including expert reviews, in-group experimentation, and study groups paid for at Brandeis University. What resulted from this process were three coordinating representations that were designed both to simplify the interaction among participants (the social part) and structure it so as to reduce the cognitive load of each user (the individual part) in her use of the mediating representation.

The coordinating representation showed in Figure 5 (the planning window) allows a user to compare his projected actions to those of the other participants. The next few planned steps for each actor are displayed in a labeled column for each participant. The actions are listed in order from top to bottom. (So, the next projected step of Crane1 is to do deploy equipment and then he will lift some waste.)   Each user has control of only one column, his/her own. This representation improves timing on exit and entry of phases for tightly coordinated phases of activity by allowing participants to compare each other's next few projected actions.

Figure 5: Timing of joint actions.

A list of objects (with relevant properties) allows users to more systematically keep track of the object in the domain (see Figure 6)**.** Columns provided information about the name, object type, location, and equipment needed for a given object. The rows can be sorted by any column; thus, each participant can maintain a different view on the information shared in the object list. An important feature of the object list was that entries in the object list could be displayed on the WorldView map as markers; each user had a toggle switch that would allow her to display or not display the markers. Each marker was named automatically using the name entered in the name field of the object list.

This representation allows participants to exchange information about shared objects. The organization of this information reduces the cognitive load for the individual, by organizing information relevant for decision making into predetermined representational structure.

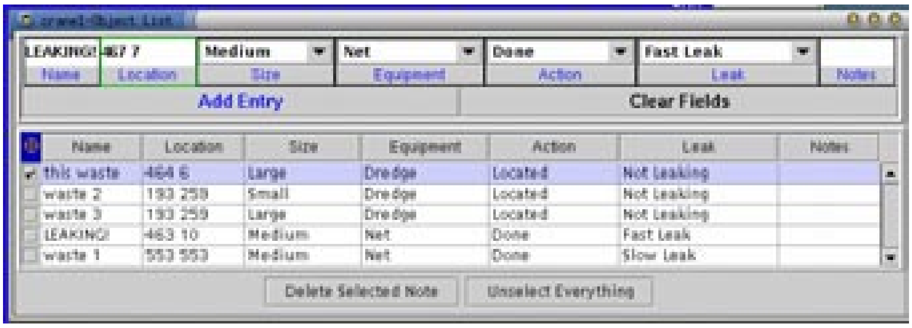

Figure 6: The object list.

A third coordinating representation was designed to allow the users to do high-level planning. The idea was to create a structured space where the participants could rapidly sketch a high-level plan that would help them to manage multiple open tasks. The high-level plan depicted in Figure 7 shows that the three actors are in the midst of an organized search of the harbor. After this, they are committed to a plan to move, in order, wastes 1, 2, and 3 onto the small barge brg1. The palette at the top allows users to rapidly build a description of a joint action sequence. Actions are one of a small set of action primitives, e.g., MOVE, SEARCH, and CONTAIN. Once formed these descriptions added to lines in the plan shared among the participants (*shared plans*: Grosz & Sidner, 1990). Color-coding of entries in the high-level plans allows participants to indicate both accomplished tasks and future commitments.

Figure 7: High-level planning.

In the empirical evaluation described below, the first two coordinating representations are shown to improve the performance of participants in managing shared domain objects. The third coordinating representation, high-level planning, was not used at all.

AN OVERVIEW OF THE EMPIRICAL EVALUATION

An empirical evaluation conducted at Brandeis compared the performance of teams of participants with and without the coordinating representation. Three groups could only electronically chat during problem solving sessions (non-CR groups), and three groups could chat but also had access to coordinating representations (CR groups). Each team was trained and then played for about 10 hours over several sessions of problem solving.

Some Performance Comparisons

Given the size of the experiment one has to be cautious about extracting too many conclusions from a statistical analysis. It would be nice to conduct a much more extensive experiment, but there are practical reasons why this is not easily done. The most important of these has to do with the problem of coordination. Each problem-solving session required three subjects to coordinate their schedules with the experimenter. Because of the complexity of the task, it was necessary for subjects to use the system for extended periods of time over multiple problem-solving sessions. This issue tended to compound the scheduling problem. It is for this very reason that most studies of synchronous cooperative behavior have tended to be more qualitative that quantitative. That is not to say there have not been more robust quantitative studies of collaboration, but these have frequently focussed on less complicated reasoning tasks. It is for these reasons that within the study of groupware systems there have been very few formal studies of the performance.

Despite the limitations of our data there were some quantitative comparisons that we were able to make that were significant. Unless otherwise noted the measures reported have a 95% confidence level and are normalized for the complexity of the problem. Problem complexity is a weighted sum over all wastes taking into account size, equipment needed, and distance from large barge for each waste.

One measure of general performance is the amount of clock time it took the participants to solve a problems: there was a 49% improvement in clock time to complete task for those groups using coordinating representations. Another measure of user work was a 38% reduction in the number of events generated while completing tasks of comparable difficulty. Because the coordinating representations pre-structure certain exchanges of information we expected to see a reduction in the quantity of electronic chats: there was a 57% reduction in the amount of electronic chatting. Because one of the coordinating representations dealt with commitment (high-level planning), another with timing, and a third with the exchange of information about equipment requirements for lifting barrels, we expected to see a reduction in domain errors: total errors were down 61% (this number has only a 80% confidence level). However, a closer analysis of the data reveals that the high-level planning coordinating representation was used hardly at all.

Where the above measure of performance compares CR and non-CR groups during their last problem-solving session, we also wanted to compare differences in performance from the early session to the later sessions for each group with and without CR's. All groups (regardless of whether they had access to CR's) improved their average performance when comparing their first 5 hours of problem solving (after training) to their second 5 hours of problem solving. In the case of communication and clock time (for problems of equivalent complexity) the CR groups outperformed the non-CR groups both during the first and second five hours of problem solving. In addition to the improved performance of the CR groups during the first five hours of play, the CR groups also generated fewer utterances for problems of equivalent difficulty. The average reduction in communication costs (as measured by the number of utterances) was similar (51% versus 56%) for CR and non-CR groups – but the variance is very high. For clock time the average improvement for CR groups was significantly better than non-CR groups (80% versus 66%) – with a 90% confidence interval.

THE ANALYSIS TOOL

All events that occur during a VesselWorld problem solving session are recorded in a log file by the system. A VCR-like program was built that was used to review the decision making of each group and examine how the participants in a VesselWorld session coordinate the exchange of information so they can maintain a common sense of domain objects. This tool was used to perform a detailed analysis of the data. Figure 8 shows the analysis tool. Starting in the upper left-hand corner of the figure and moving counter clockwise, the following are some of the windows of information that are provided for the analyst:

- The history of chatting
- The current plans of each of the participants.
- The layout of the harbor and the location of all objects.
- An annotation window that can be used by the analyst to comment at a particular point of the problem-solving session.
- The controller of the tool.
- The object list.

The analyst can also open other windows of information. The controller allows the analyst to step through the data using any number of metrics, e.g., it can move forward to the next communication, round, or bookmark. The analyst can also fast forward through the data. The controller also displays information about the current round being represented, the current time, and the number of rounds in the session. Further details on the analysis tool and its implementation can be found in Landsman & Alterman (2001). Currently under development is a second version of the analysis tool that allows the analyst to segment and sort the conversation into threads and semi-automatically profile and model interactions among participants.

The first part of qualitative analysis discusses the work of participants to maintain consistent private representations of shared domain objects; review the information contained in Figure 2 for a list of the information the participants must keep track of. The analysis also details the design work of participants to create conversation structures that simplify the exchange of information. The second part of the qualitative analysis examines the functioning of all three coordinating representations.

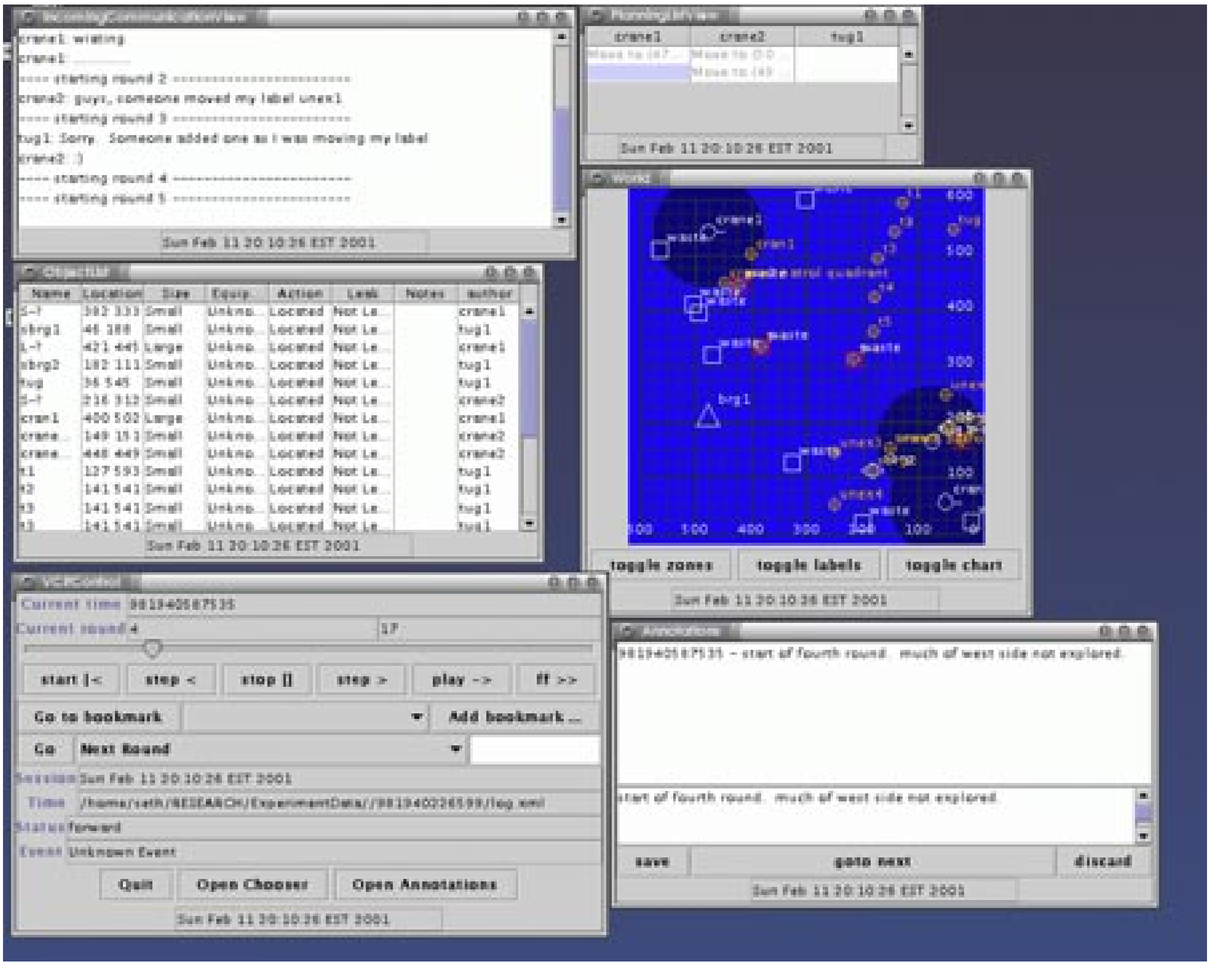

Figure 8: The analysis tool.

WITHOUT COORDINATING REPRESENTATIONS

The participants must coordinate their efforts in the exchange of information about shared domain objects. Similar conceptions of the status of shared objects and actions to be taken upon them must be maintained in order to perform effectively at their common tasks. For the users of the base system, electronic chatting is the primary method available for exchanging information. When participants are in the same area of the harbor, and can see each other, they can also access information about objects within their shared purview by using either the information window or the worldview window. This latter source of sharing information can lead to some confusion, as the visible areas for two users may overlap but are never identical. Sometimes a situation occurs such that two participants are referring to what they think is the same object, but are in fact two objects close together, with one visible to one participant and the other visible to other participant.

For the recurrent problems of coordination in the exchange of information and the maintenance of a common viewpoint, the participants engage in a design process that results in a set of common expectations, designs, for the structure of their behavior in maintaining a common viewpoint of shared objects. Although conventions begin to emerge, the actors must continue to communicate at the critical junctures (the coordination points) of their joint activity. The participants need to exchange information about newly discovered information, confirm that expectations are being met, or adjust to variance in the flow of activity.

Consistency

Whatever inconsistencies develop among the participants' private assessments of shared domain objects (or about the state of the world) must eventually be repaired. The ongoing effort that is needed to maintain a mutually consistent understanding of a shared situation is the basis of intersubjectivity (Garfinkel, 1967; Schegloff, 1992). What is "common" or "shared" is a result of continuous coordinated activity.

There were four methods that were used for maintaining consistent understandings of the status of shared objects:

1. Report
2. Review
3. Repair
4. Confirm

Report

The predominant method for maintaining consistency was for the participants to report on their current activity (via electronic chatting). Typical statements of this sort were:

- Waste at 554,41 is small, dredge
- New XL! At 200 431
- I guess I'll sweep the bottom, west to east
- [all] remove marker at 100,425
- 8 clicks for me to hit BB
- I'll grab sX at 500 275
- The first on the barge is sX
- w8ting
- killed sX at 500,275
- I still and 2 or 3 moves till I get there
- I'm waiting at SW corner
- I have a leaker on my hands

In many cases the report of information was in response to a request from another actor, e.g.,

- [Tug] check 0,575
- load then get XLX?
- LL XLX to east?
- How fast can Crane 2 get here?
- Did Crane 2 deploy?
- Which one is leeking?

The flow of information is continuous. It is the responsibility of each actor to add that information to his or her private representation (either by taking notes, marking the map, or remembering), or be prepared to examine the history of chatting at some appropriate future time. Any information that is lost, misunderstood, never recorded, or never transmitted in the first place, can lead to discrepancies between the participants' individual assessments of the status of shared objects in the domain.

Review

One strategy for avoiding differences in assessment is to engage in a conversation to review the status of one or another of the shared domain objects. In Figure 9 we see the “fleet” engaging in a review discussion, confirming that a portion of the harbor has been searched and is in fact clean.

1. Crane1: center is clean – nothing in the center 4 quads
2. Crane2: ok, we forgot the west, I didn’t get down that far
3. Tug1: 600-400 square is clean, I don’t have to worry about tit?
4. Crane2: correct, that’s where I started
5. Crane2 about 100 N of sb
6. Tug1: okay cool, grabbing the sb

Figure 9: Reviewing the status of the harbor.

The review of what part of the harbor had already been searched was a common topic of review for all the groups. Another topic of review was the size and weights of toxic wastes at any given point and time on the small barge. A third topic of review was for the participants to remind one another of what toxic wastes were at what location; a fourth was to remind each other of the procedure that was to be followed or the equipment that was needed for a particular waste.

Repair

Whenever discrepancies in the assessment of a situation unexpectedly develop, the participants engage in repair work to re-mediate between alternate representations of “reality”. An example of where repair work is needed is where one of cranes intends to pick up a waste only to discover it is not there. He may immediately comment upon this to the other participants, whereupon one of the other participants responds that that waste has already been removed.

Figure 10 shows another example of the kind of repair work that the participants must do to maintain consistency. In this example, the chief problem is that the participants do not have the option of pointing to the thing they are referring to; they must produce a description that uniquely identifies the referent in the light of the limited “perceptual” capabilities of each of the actors. At line 1, Crane2 broadcasts that he has discovered two new wastes. At line 3, Crane2 refers to one of the wastes by giving its location. At line 4, the Tug reports that the waste at that location is a mN (i.e., medium waste that requires a net to be removed from the harbor). At line 5, he states that both the wastes are mNs, that he “id’d“ them before. But Crane2 disagrees by stating there is a new waste at 175,225 (line 6) and that it is a small one (line 7). It takes several lines for the Tug and Crane2 to reach agreement, which is confirmed when the Tug (at line 15) identifies it as a sX (a small waste requiring no equipment), not a mN, that is at location 175, 225.

1. Crane2: see tug and 2 new wastes
2. Crane1: let tug ID them and announce them
3. Crane2: tug, wast at 50,225
4. Tug1: mN at 50 225
5. Tug1: they’re both mNs, ID’d the other one before
6. Crane2: New waste at 175, 225
7. Crane2: It’s a small
8. Tug1: are you sure about that c2? I was just there
9. Crane2: moving NW
10. Tug1: oh wait, my bad

11. Crane2: yes, see mN and small
12. Tug1: ok
13. Tug1: going to check it out
14. Crane2: see it tug
15. Tug1: sX at 175 225

Figure 10: Repair work.

Figure 11 gives an example of a dialogue where repair work is triggered by two of the participants "simultaneously" finding the same waste. At line 1, the Tug says he found a medium waste at 400 125 that requires no equipment to remove. At line 2, Crane1 states he found a medium waste at 392 127 (only the Tug can determine what equipment is needed to remove a waste from the harbor), and then at line 3 realizes that it is the same waste as the one the Tug mentioned. At line 4, the Tug confirms that they are referring to the same waste, and at line 6 repeats that it is a mX. From this dialogue, Crane2 (who can see neither Crane1 nor the Tug) infers that his co-participants are in the south portion of the harbor, which is confirmed by the Tug on line 7. At line 9, Crane1 states that he is going to move the mX to the BB (big barge) that is (visibly) nearby.

1. TUG1: mX at 400 125
2. CRANE1: medium at 392 127
3. CRANE1: that's got to be the same one
4. TUG1: yep
5. TUG1: that's an mX
6. CRANE2: so you guys are in S?
7. TUG1: yes
8. …
9. CRANE1: I'm going to trash that mX on the BB

Figure 11: Establishing a referent.

In many cases a description like the one produced by Tug1 in line 1 of Figure 11 would be sufficient to avoid ambiguity. But in this case, because Crane1 has "simultaneously" offered another description, the actors must determine if there is actually only one waste or if there are two wastes near one another. In this case there is only one waste, and it requires neither equipment nor the help of the other crane to remove. At line 9, because it has been established there is only one waste at the "location" of the Tug and Crane1, Crane1 can broadcast to the other participants that he is removing "that mX" and the other actors will know which waste he is referring to. In the event the waste could not have been handled immediately, without correction, the waste could later by referred to by one of at least two names, either "the mX at 400 125" or "the mX at 392 127", which would provide another opportunity for the participants' individual assessments of the situation to become misaligned.

Confirm

When information is reported it is not always "received" by the other participants; e.g., it may be read but not remembered or recorded. This can lead to divergence in the actors' individual representations of the status of shared objects. A fourth method used by the participants to maintain consistent representations of shared domain objects is by confirming that somebody else's report or repair was received. A simple example is when the Tug reports "72,368 extra lge, net" and one of the Cranes confirms that report has been received by stating "[tug] copy".

As shown in the dialogue segments depicted in Figure 10 (line 12) and Figure 11 (lines 4, 5, and 7), confirmation can be used to reach agreement in establishing whether a domain object exists or not. Figure 12 provides another example of the use of confirmation. At line 5, the Tug proposes a vague description of a number of wastes around the location 100 380. Although this reference is vague, Crane1 is able to confirm that he knows which ones the Tug is referring to buy adding detail to the description at line 6, i.e., “I marked both”. If Crane2 had chosen to help remove the wastes around 100 380, and she had not marked these wastes on her map, she would have to use her “perception” to discover the two wastes in that general location or remember their locations, but that adds to the short term memory load.

1. Crane1: done! Now we go to the small wastes
2. Crane2: split and go find other waste now huh
3. Crane1: yep
4. Crane2: here we go
5. Tug1: Cool. I left a center marker on my screen at 100h/380v The wastes are around that marker all of them re small.
6. Crane1: i marked both, ill head to the one i need to get
7. Tug1: And I’ll hang out and you can load on me.

Figure 12: Confirming a reference.

Secondary Structures Created to Simplify Talk

When dialogue to maintain a consistent interpretation of the status of shared domain objects recur, the participants can develop designs (conversation structures and methods) that simplify coordination in the exchange of relevant information. Here we discuss three examples of conventions that emerged during the problem solving sessions of the groups that used the basic system.

Each group developed shorthand notational conventions for describing various features of the wastes. For example, one group developed a set of abbreviations to reduce effort in describing the size, equipment needed, and location of a barrel of waste. The sizes of the wastes are one of xl (extra large), l (large), m (medium), or s (small). The equipment needed is indicated by a single letter: X (no equipment), D (dredge), and N (net). The location is indicated by x and y coordinates; for example, 400 125. So when the Tug says “mX at 400 125” or “mN at 150 200” or “XL D at 200 425”, she means, respectively, there is a medium waste at 400 125 that requires no equipment to move, a medium waste at 150 200 that needs a net to be removed, and there is an extra large waste at 200 425 that will need a dredge to be removed from the harbor. These conventions of naming allowed the participants to rapidly describe, in a few keystrokes, the relevant information about a particular waste. Thereafter the participants had a useful handle that reminded them of many of the relevant properties of the waste.

The problems inherent in the cranes jointly lifting or moving a large or extra large toxic waste was also the site of recurrent dialogue. To remove a large or extra large waste required two levels of coordination.

1. Coordination of talk
    a. Adjacency pairs to propose and confirm next action
    b. Expectations that adjacency pairs will occur for each of the actions in an extended sequence of tightly coupled cooperative behaviors
2. Coordination of action

Lifting a large toxic waste that requires a dredge and loading it onto a small barge requires several coordinated actions: 1) the correct crane must deploy the dredge; 2) jointly, during the same round of activity, the cranes must lift the waste; 3) jointly, during the same round(s) of activity, the cranes must carry the large waste to the barge; and 4) jointly, during the same round of activity, the cranes must load the large waste onto the small barge. Each of these actions must occur in that exact order and must be synchronized. Any errors in coordination can result in failure and leakage of toxic waste.

Each of the groups that did not have access to a coordinating representation used "adjacency pair" (p. 295: Schegloff & Sacks, 1973) as the basic unit for coordinating joint operations on large and extra large wastes. The first part of the adjacency pair was for one actor to propose to take a given joint action on the next round. The second part of the adjacency pair was for the other actor to confirm that he would take the corresponding action. So, if Crane1 proposes to do a joint load, Crane2 can confirm. Another aspect of the design of the talk that occurs for closely coupled joint activities is the expectation that over the course of removing a large or extra waste from the harbor to a barge, adjacency pairs will occur for each of the coordinated actions. In addition to the conversational part of their joint action, there is also a domain action part, where the cranes submit their actions to the system. In other words, first the talk is coordinated, and then the actions are coordinated.

Figure 13 shows a sample of dialogue that occurs over several rounds of activity and involves the cranes unloading wastes from the small barge onto the large barge. At line 1, Crane1 states he is submitting a joint lift to the system. Simultaneously, Crane2 had stated he was submitting a "lift large together" (line 2). At line 3, Crane2 confirms Crane1's proposal in line1. Next, Crane2 states he is submitting a loading, which Crane1 confirms at line 6. Because the "conversation" of the users is mediated through an electronic chatting frame, adjacency pairs do not strictly speaking occur one after the other. Other kinds of comments can be interposed along the way, e.g., the Tug's comment that the next waste will be an extra large that "needs nothing".

1. Crane1: sub lift
2. Crane2: LL
3. Crane2: k
4. Crane2: sub load
5. Tug1: the next XL needs nothing
6. Crane1: k
7. Crane2: ok, then XLD right?
8. Crane2: sub Lift
9. Tug1: yep
10. Crane1: k
11. Crane2: sub load
12. Crane1: k
13. Crane2: sub sep
14. Crane1: sep

Figure 13: Adjacency pairs in unloading the small barge.

A third example of a conventional structure for conveying information emerged to support a review of the wastes that had so far been discovered. In order to insure that each participant was maintaining the same list of toxic wastes, one of the groups using the basic system would periodically do a "marker check" (see Figure 14). At line 1, Crane1 proposes to do a marker check. There is some intervening dialogue and then at line 3, Crane1 produces a legend, at line 4 establishes an ordering, and at line 5 proceeds to list all the wastes he thinks they have found "south of the equator". At lines 6-10 the Tug and Crane2 initiate some repair work. At line 11, Crane1 continues listing all the toxic wastes, their locations, and the equipment needed. At lines 15-22, the repair work initiated earlier in the dialogue continues until all the differences in the participants' individual assessments of the known wastes are resolved.

---

1. Crane1: [ALL] ok I will dump all the markers ok with every1?
2. …
3. Crane1: Legend: [Sm|L|XL] – [Ni [no id'd] Net | Dr]
4. Crane1: from south east clockwise
5. Crane1: [Sm-Ni 50,0][Sm-Net 150,25][Sm-Ni 350,150][Sm-Ni 550,50] [Sm-Ni 600,100] That's all south of equator. NORTH coming up.
6. Tug1: 97441 and 72,368 already ID'd
7. Crane2: 350,150 is barge, isn't it?
8. Crane2: that's the problem
9. Crane2: stop dump I was there to… that's the Sbarge at 350,150
10. Crane2: confirm with TUG
11. Crane1: [xL-Ni 475,425][Sm-Ni 450,450][Sm-Ni 525,500][Sm-Ni 250,500][XL-Ni 200,475] [Sm-Net 100,425][Xi-net 75,375][Sm-Ni 25,575]
12. Crane1:
13. Crane1:
14. Crane1: –- END – Tug, confirm you have all those
15. Tug1: large barge at 400,325
16. Crane2: repeat: I say the sbarge ar 350,150, not a small waste
17. Crane2: [TUG] where are the small barges?
18. Crane1: [c2] – you sure? If you saw that, ok ill cancel as waste
19. Crane1: of yeah tug can tell us
20. Tug1: 350,150 is small barge
21. Crane2: ok, 12 wastes and no problems, let's get back to work
22. Crane1: okdoke. Still sweeping west

Figure 14: A marker check.

WITH COORDINATING REPRESENTATIONS

The general advantage of the coordinating representation is that it simplifies the problem of establishing a consistent representation of the situation among the participants, i.e., all of the coordinating representations are by definition shared representations. It also provides a shortcut for coordinating the exchange of information about shared domain objects. A problem is that the participants do not always access the information that is available in the coordinating representation at the appropriate time.

Close Coordination

The idea behind the shared planning window was that it would reduce errors and the work involved in the close coordination needed to handle the large and extra large wastes. Earlier we examined (see Figure 13) the structured talk that the participants engaged in during these kinds of cooperative activities. Figure 15 shows the cranes' entire conversation while jointly lifting a barrel of waste and loading it on a small barge. During round 7, the cranes and the tug all agree to remove wastes G and B from the harbor. During round 8, the Tug makes a joke, which Crane1 responds to during round 10. During round 11, the Tug announces that G1 is loaded.

1. CRANE1: G, then B
2. CRANE2: Okay.
3. Tug1: sounds like a plan
4. TUG1:
5. —— End of round 7 ——
6. TUG1: a man, a plan, a canal, panama.
7. —— End of round 8 ——
8. CRANE1:: that was *last* time : -)
9. —— End of round 10 ——
10. TUG1: G1, loaded.
11. —— End of round 11 ——

Figure 15: Conversation while lifting and loading an extra large waste.

What happened to the extensive dialogue that the groups using the basic system engaged in to coordinate their operations for the removal of large and extra large wastes? As can be seen in Figure 16, almost all of the structure that the non-CR groups produced in chat to organize their activity is being created by the CR groups in the planning window, without any extra effort. By the end of round 7, all three participants have "planned out" their next few actions and these are visible to all three participants via the shared planning window. Crane1 will wait while Crane2 deploys equipment. In the succeeding rounds of activity, the cranes will "join" to one another, jointly lift the waste, and then load it on the barge. Through all of these actions, the Tug holds steady waiting for the small barge to be loaded.

The main advantage that accrues to the users who have access to the shared planning window is that it requires no extra work on the part of the participants to build: i.e., report, review, repair, and confirm require less effort. In order to submit an action to the system the users need to add it to their "plan" anyway. So, from the point of the view of the users who have access to the shared planning window, having to talk about their cooperative activity is just extra work. Another advantage is that one actor now has the opportunity to spot potential problems in another actor's plan. So if one actor is about to lift a waste but forgets to deploy his crane first, another actor can remind him to deploy his crane. Without the coordinating representation this sort of reviewing of one another's plans during the course of action could only occur when plans are being reported in the chat window. On more than one occasion it was observed that one crane would adjust his plan to match the plan of the other crane within the same turn, without any discussion. At other times it was observed that when the cranes failed to match their plans, one of the participants would bring the disagreement up for discussion in the chat window.

Figure 16: The planning window at the end of round 7.

Exchanging information using the Object List

The object list is a shared representation structure that allows the participants to coordinate their efforts in sharing some of the information about shared domain objects. Figure 17 shows the opening dialogue in a VesselWorld problem solving session where users had access to coordinating representations. This dialogue ensues before all of the participants have submitted their actions to the system for the first round of action. At line 1, Crane1 ecstatically declares that he can see an extra large waste. At line 2, the Tug expresses his "envy". At line 3, Crane2 expresses his excitement that he can see both an extra large and large waste. The rest of the opening dialogue is mostly concerned with planning.

1. CRANE1: I got an XL!
2. TUG1: I got nothing, you luck basrstartd.
3. CRANE2: I got an Xl and an L, mommy! ;)
4. TUG1: Merry christmas kids….
5. CRANE1: I'll map North third?
6. TUG1:
7. TUG1:
8. TUG1:
9. TUG1: I'll take middle 3rd.
10. CRANE2: I'm at south-central. Tug, where are you?
11. TUG1: I'm jus nw of the barge, let me put that on the map…
12. TUG1: actually more w than n.
13. CRANE2: With the LB in the corner, maybe our best bet is moving the SB NW and loading it with all the NW corner's goodies, which CRANE1: can map
14. CRANE1: not a bad plan…
15. TUG1: Ok, I'll make a bit of a sweep around here while CRANE1: looks around.
16. CRANE1: Tug, can you pick up the SB at your earlier opp?
17. TUG1: CRANE2: can map up on the way?

Figure 17: Opening dialogue.

Figure 18 shows the object list that is constructed by the time all the participants have submitted their first action. Only three of the entries into the object list were explicitly mentioned in the opening dialogue, and none of these were explicitly named. Much of the dialogue that accompanied the discovery of a new waste in the groups using the basic system is now occurring "offline" and is mediated by a coordinating representation. Identifiers are attached to each of the "objects" that are found; this will simplify future conversational exchanges about the wastes. Because pointing and clicking on wastes can add entries to the object list, precise locations for each of wastes that are found can be stored. This will simplify future interactions for disambiguating references and referents as now the actors have the precise location descriptions to match against the information that they get about wastes from the information window. An important feature of the object list is that all the information is shared. Much of the consistency checking that the users of the base system had to engage in is no longer necessary. Rather than having three private representations that must periodically be reconciled by electronic chatting, the users can share a single representation. This scheme reduces the number of conflicts between different conceptions of the shared workspace, but it also eliminates the work involved in re-mediating discrepancies between alternate views of the shared domain objects. Another behavior that was observed was that actors would not always access the relevant information from the object at the outset of a course of action: in one case it was observed that one of the cranes, on more that one occasion, failed to deploy his equipment before lifting a medium or small waste that he was responsible for removing from the harbor – even though he had that information available.

Name	Location	Size	Equipm...	Action	Leak	Notes
G1	556 465	XLarge	Unknown	Located	Not Lea...	
A1	186 107	XLarge	Unknown	Located	Not Lea...	
m	550 447	Small	None	Located	Not Lea...	
A2	249 21	Large	Unknown	Located	Not Lea...	
m	250 149	Small	None	Located	Not Lea...	
m	449 349	Small	None	Located	Not Lea...	
SB	305 310	XLarge	None	Located	Not Lea...	

Figure 18: The object list

The object list shown in Figure 18 has a few other interesting features. This group explicitly developed naming conventions for labeling entries into the object list. Labels for entries began with A (as in Alpha), B (as in Beta), and G (as in Gamma), depending on which participant originated an entry. (An alternate scheme that was used by the development team was to label entries by their size characteristics, and handle duplicate name problems as they arose.) Another interesting feature of this object list is that this group used the object list to keep track of what parts of the harbor had already been explored. This group entered an "m" to indicate that an area of the harbor had been visited; a slightly different scheme of this sort was used by another group. Because all entries in the object list can be displayed on the WorldView map, these groups had an external representation of the parts of the harbor that were (or were not) explored. Both the labeling convention and the marking convention are indications that the groups with access to coordinating representations also created secondary structures to support their interaction.

We did some analysis of the usage of each column in the object list. The hypothesis (after the fact) was that not all columns of the object list were used equally. The name, location, size, and equipment fields were used most of the time by each of the CR groups. The action field was used infrequently; the group that used it most stopped using it all together after the first 3 hours of play. About 30% of the leaks that occurred were recorded on the object list. Only 3% of the entries created had notes associated with them. Our belief is that the differences of usage are a function of the duration of relevance of each kind of information. Those information fields that are relevant for extended periods of time (like the equipment needed to lift a waste) were recorded in the object list, while those that were relevant for only a short period of time, and could be easily retained in working memory, or information was easily accessible by other means, were less likely to be recorded.

High-Level Plans

The high-level planning window was not used by any of the groups. The surveys we collected from the subjects show that the chief problem with the high-level planning windows was that, given the rewards it provided, it required too much work to complete. Further analysis shows that the problems that the high-level planning window was designed to fix continued to occur. Further analysis of the lack of utility of the high-level planning window is included in the discussion section.

DISCUSSION

Recall that there were three indicators for the development of coordinating representations that were used in the analysis of the pilot study data.

- Recurrent patterns of coordination
- Repeating errors
- The development of secondary structures to organize talk

The first indicator requires that the analyst be able to sort the conversation among the participants so as to identify the coordination problems the participants spend the most time talking about. The second indicator requires the analyst to determine the situations in which the participants were most likely to make errors. The third indicator is the most participatory in nature, and it depends on the analyst being able to recognize where the participants have developed secondary structures to coordinate their talk by fixing repeated errors or organizing complicated recurrent interactions.

The analysis of the pilot study data identified the object list as a possible coordinating representation to introduce because the exchange of information about objects was a regularly occurring activity among the participants (indicator 1). It is noteworthy that in the more extensive data we collected for the empirical study, one of the non-CR groups created a secondary structure (indicator 3) that did some of the work of the object list: the marker check procedure that was performed by one of the groups was intended to ensure that each participant had the same list of toxic wastes. The empirical study showed that all the CR groups effectively used the object list coordinating representation.

The analysis of the pilot study data predicted the utility of the shared planning window for two reasons: because closely coupled actions involving jointly lifting and loading wastes was a source of participant errors (indicator 2), and also because the participants developed secondary structure that organized their talk about these kinds of actions (indicator 3). There are several pieces of evidence that the shared planning window was an effective coordinating representation. We observed several occasions where the crane operators in the CR groups used the shared planning window to synchronize their joint actions without having to resort to talk. The quantified comparisons showed that on average the groups with access to the shared planning window spent less time, and generated fewer utterances and words, when coordinating their activities during rounds in which joint actions occurred. With regard to the effectiveness of our implementation of this coordinating representation into a specific shared data structure, the implementation required no extra work for the users to post the information; the work required to access the information was not evaluated.

Perhaps the most interesting case to look at is the failed coordinating representation, the shared representation for high-level planning. The indicator that pointed to the development of this coordinating representation was that there were numerous occasions where the participants in the pilot study engaged in high-level planning (indicator 1). It is perhaps significant that even for the more extensive data we collected in the empirical study, none of the non-CR groups developed secondary structures to support high-level planning – this is in contrast to the object list. Perhaps it is the case that those places where the participants develop secondary structures to coordinate conversation are a more reliable indicator of places where coordinating representations are appropriate. Alternately, it could be the case that recurrent problem areas of coordination are sufficient to predict where a coordinating representation could be developed to introduce a secondary structure, but that our implementation was inadequate.

If the problem is in the implementation, then the design methodology for constructing coordinating representations needs to be improved. Perhaps too much functionality was bundled into the high-level planning window, thus necessitating a cumbersome interface that was too much like end-user programming. Separating out the task of marking which parts of the harbor had been searched from the tasks involved in performing a series of related operations might have produced a more useful coordinating representation. Another potential source of difficulty was that much of the data entered into the high-level planning window had a relatively short lifetime of relevance. With the object list, those columns where information was relevant for an extended period of time was more likely to be used than those with a short lifetime of use. Similar constraints were undoubtedly at work in the high-level planning window. Given the relatively short period of relevance of much of the high-level planning information, it may have been easier for the participants to chat and then retain that information in short term memory than to do the work required to maintain that information in an external representation. An alternate approach to remedying the situation is to introduce some AI techniques that would allow the system to fill out portions of the high-level planning window semi-automatically (Introne & Alterman, 2000).

CONCLUDING REMARKS

The participants in a joint activity must work hard to maintain coordination. They may have expectations about how the activity will end up being organized. These expectations can be shared and a part of the common ground, but because of variance in circumstance, it will be necessary to interact in order to communicate with one another to organize joint behavior or confirm expectations about the course of the cooperative activity. For more complicated activities, involving large amounts of coordination, proportionally more talk is needed in order to stay on course. For novel activities, extra conversational work is needed to work out the details. Over time, for recurrent cooperative behaviors, the participants will organize their talk as a means of organizing their actions.

The data collected from the empirical study exhibited these features of the theoretical framework. There were several examples of the subjects (with or without coordinating representations) working hard to stay coordinated. The performance of the groups improved as they introduced secondary structures that organized their talk. As their talk became better organized so did their joint actions. For the non-CR groups, the introduction of "marker checks" and the development of an adjacency pair structure to orchestrate tightly coupled actions are two examples of this sort of phenomena at work. For the CR groups, the introduction of the coordinating representations provided the subjects with pre-packaged structure that served a similar function.

For some collaborative computer-mediated tasks general-purpose mechanisms like an electronic chat room will suffice to effectively organize a community of users to accomplish some joint task. For other applications, these general-purpose mechanisms will not provide enough structure for the task. In these cases, developing a set of coordinating representations for the recurrent problems of coordination that occur both in the level of conversation and in action will result in significant improvements in the performance of the participants. Because coordinating representations must be tailor-made for the coordination problems that emerge for recurrent activities, it is necessary to develop a methodology for producing effective coordinating representations. The task for the developer is to add structure to the problem areas of interaction. In this paper, three indicators that a problem area of coordination have been examined: What are the recurrent areas of behavior where coordination requires increased effort? What sort of errors do the participants make? What secondary structures for organizing conversation developed among the participants?

ACKNOWLEDGMENTS

This work was supported by ONR Contracts N00014-96-1-0440 and N66001-00-1-8965.

REFERENCES

Agre, P. (1997). *Computation and Human Experience*. Cambridge University Press.

Alterman, R., Zito-Wolf, R., and Carpenter, T. (1998) *Cognitive Science,* 22(1), 53-105.

Alterman, R. and Garland, A. (1998). Convention in Joint Activity. *Cognitive Science*, 25(4), 611-657.

Bly, S. (1988) A use of drawing surfaces in different collaborative settings. In *Proceedings of the Conference on Computer-Supported Cooperative Work*, 25-256.

Bratman, M. (1999) Shared cooperative activity. In *Faces of Intention,* pages 93-108, Cambridge University Press, Cambridge, U.K.

Clark, H. (1996). *Using Language*. Cambridge University Press.

Clark. H. and Brennan, S. (1991). Grounding in Communication. In Resnick, Levine, and Teasly editors, *Perspectives on Socially Shared Cognition*, 127-149.

Cole, M., & Engeström, Y. (1993) A cultural histgoric approach to distributed cognition. In G. Salomon (Ed.), *Distributed Cognitions,* 1-46. Cambridge, MA: Cambridge University Press.

Flores, F., Graves, Hartfield, B., and Winograd, T. (1988) Computer systems and the design of organizational interaction. *ACM Transactions on Office Information Systems,* (6) 153-172.

Feinman, A., and Alterman, R. (2001) Modeling Communicative Behavior in a Groupware System. TR CS-00-211 Computer Science Department, Brandeis University. To appear in HCI 2001.

Garfinkel, H. (1967). *Studies in Ethnomethodology*. Polity.

Goffman, E. (1967). *Interaction Ritual: Essays in face-to-face behavior.* Chicago: Aldine Publishing Co.

Goffman, E. (1983) The interaction order. *American Sociological Review,*48:1-17.

Goldman, S. (1996) Mediating microworlds: Collaboration on high school science activities. In Koschman, R. (Ed.) *CSCL: Theory and practice of an emerging paradigm.* Lawrence Erlbaum Associates. Mahway, New Jersey.

Goodwin, C. and Goodwin, M.H. (1996). Seeing as situated activity: Formulating planes. In Engeström and Middleton (Eds), *Cognition an communication in work.* Cambridge University Press, Cambridge, U.K.

Greenberg, S., Roseman, M., and Webster, D. (1992) Issues and experiences designing and implementing two group drawing tools. In *Proceedings of the 25th Annual Hawaii International Conference on the System Sciences,* (4):139-150. Reprinted in *Groupware and Computer-Supported Cooperative Work*, 1993.

Greeno, J. (1998). The situativity of knowing, learning, and research. *American Psychologist*, 53:5-26.

Grosz, B. and Kraus, S. (1996). Collaborative plans for complex group action. *Artificial Intelligence*, 86:269-357.

Grosz, B. and Sidner, C. (1990). Plans for discourse. In Cohen, P. R., Morgan, J., and Pollack, M. E., editors, *Intentions in Communication*, pages 417-444. Bradford Books, Cambridge, MA..

Hutchins, E. (1995a) *Cognition in the Wild.* MIT Press. Cambridge, MA.

Hutchins, E. (1995b). How a cockpit remembers its speed. *Cognitive Science*, 19:265-288.

Introne, J. and Alterman, R. (2001) Segmenting Usage Data in Collaborative Systems. TR CS-01-215. Appeared in workshop on "Dealing with Community Data" at CSCW 2000.

Kirsh, D. & Maglio, P. (1994) On distinguishing epistemic from pragmatic action. Cognitive Science, 18(4), 513-549.

Landsman, S., Alterman, R., Feinman, A, Introne, J., VesselWorld and ADAPTIVE, TR CS-01-213, Brandeis University. Demo given at CSCW 2000.

Lave, J. (1988). *Cognition and Practice.* Cambridge University Press, Cambridge.

Malone, T., Grant, K., Lai, K-Y.,Rao, Rao, R., and Rosenblitt, D. (1989) The information lens: An intelligent system for information sharing and coordination. In *Technological Support for Work Group Collaboration,* edited by Olson, M.H., (65-88).

Mitchell, T. Keller, R. and Kedar-Cabelli, S. (1986) Explanation-based generalization: A unifying view. *Machine Learning,* 1:47-80.

Newell, A. and Simon, H (1972). *Human Problem Solving.* Prentice-Hall, Englewood Cliffs, N.J.

Norman, D.A. (1988) *The Psychology of Everyday Things.* Basic Books.

Norman, D.A. (1991) Cognitive Artifacts. In J.M. Carroll (Ed.) *Designing Interactions: Psychology at the Human-Computer Interface,* 17-38. Cambridge. Cambridge University Press.

Perkins, D.N. (1993) Person-Plus: A distributed view of thinking and learning. In G. Saloman (Ed.), *Distributed Cognitions,* 88-110. Cambridge, MA. Cambridge University Press.

Sacks, H., Schegloff, E., and Jefferson, G. (1974). A simplest systematics for the organization of turn-taking in conversation. *Language*, 50:696-735.

Schutz, A. *The Phenomenology of the Social World.* Translated by Walsh, G. and Lehnert, F. Northwestern University Press, 1967.

Shegloff, E. and Sacks, H. (1973) Opening up closings. *Semiotica*, 8:289-327.

Shegloff, E. (1992) Repair after next turn: The last structurally provided defense of intersubjectivity in conversation. *American Journal of Sociology,* 97:1295-1345.

Stefik, M., Bobrow, D.G., Foster, G., Lanning, S., Tatar, D. WYSIWIS Revised: Early experiences with Multiuser Interfaces. *ACM Transactions on Office Information Systems,* 5(2), 147-167.

Suchman, L. A. (1987). *Plans and Situated Actions*. Cambridge University Press, Cambridge.

Tang, J. (1991) Findings from observational studeis of collaborative work. *International Journal of Man-Machine Studies,* 34, 143-160.

Tatar, D., Foster, G., and Bobrow, D. (1991) Design for conversation: lessons from Cognoter. *International Journal of Man-Machine Studies,* 34:185-209.

Trigg, R.H., Suchman, L.A., and Halsz, F.G. (1986) Supporting collaboration in NoteCards, In *Procceedings of Conference on Computer-Supported Cooperative Work* , 153-162.

Vygotsky, L. S. (1978). *Mind in Society*. Harvard University Press, Cambridge, MA.

Whittaker, S., Brennan, S., and Clark, H. (1991) Co-ordinating activity: An analysis of interaction in Computer-Supported Co-operative Work. In *Proceeding of the Conference on Computer Human Interaction,* 360-367.

Whittaker, S., Geelhoed, E., and Robinson, E. (1993) Shared Workspaces: How do they work and when are they useful? *International Journal of Man-Machine Studies,* 39:813-842.