

Coordination of Talk: Coordination of Action

Richard Alterman, Alex Feinman, Seth Landsman, Josh Introne

Computer Science Department

Center for Complex Systems

Brandeis University

RUNNING HEAD: TALK AND ACTION

Corresponding Author's Contact Information:

Richard Alterman
Computer Science Department
Center for Complex Systems
Brandeis University
Waltham, MA 02454
alterman@cs.brandeis.edu

Brief Authors' Biographies:

Richard Alterman is a Cognitive Scientist with interests in cognitive modeling, activity, collaboration, and discourse; he is a professor with a joint appointment in Computer Science and the Center for Complex Systems at Brandeis University. **Alex Feinman** is a computer scientist with an interest in human computer interaction; he is a graduate student in Computer Science at Brandeis University. **Seth Landsman** is a computer scientist with an interest in software development and groupware systems; he is a graduate student in Computer Science at Brandeis University. **Josh Introne** is a computer scientist with an interest in artificial intelligence; he is a graduate student in Computer Science at Brandeis University.

ABSTRACT

The participants in a joint activity must work hard to maintain coordination. For complicated and/or novel activities, even more talk is needed to proceed. Over time, for recurrent cooperative behaviors, the participants will organize their talk as a means of organizing their actions. For recurrent activities, a sign may be introduced at the scene to fix a recurrent problem of coordination by providing some organizational structure (e.g., a stoplight). We will refer to permanent structure designed and implemented prior to a cooperative activity by a non-participant that mediates and organizes the activity as a *coordinating representation*. The main work of this paper is to explore the ramifications of, and methodology for, introducing coordinating representations into same-time / different-place computer-mediated cooperative activities.

The discussion of these issues will be developed in the context of a same-time / different-place groupware system called VesselWorld, where the root form of communication is textual chat. In VesselWorld, three participants engage in a computer-mediated problem solving session. To accomplish a set of cooperative 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 by the system. Most of the participants' task (and communication about their task) is concerned with objects existing in the simulated environment.

A pilot study was performed that collected 30 hours of VesselWorld problem-solving data. A discourse analysis identified recurrent areas of coordination for subjects of the pilot study. Three coordinating representations were implemented to facilitate interaction during these critical junctures of joint behavior. A formal experiment was then run that consisted of two populations, one in which groups made use of the coordinating representations (the CR groups), and another in which they performed the task without them (the non-CR groups). The results show that the CR groups out-performed the non-CR groups right from the start. CR groups talked less than non-CR groups and did less work at maintaining consistent representations of shared domain objects. Our study also shows that the planning work of the CR groups was improved over that of the non-CR groups.

CONTENTS

Abstract	2
Introduction	4
Background.....	5
Coordinating Representations.....	6
Computer-Mediation of Activity	7
Method: The Vesselworld simulation environment.....	8
VesselWorld: Task and Interface.....	9
Pilot Study	10
Discourse Analysis.....	11
An Example of Discourse Analysis	11
Design of the Coordinating Representations	13
Evaluation.....	15
The Experiment.....	15
Findings	16
Improvement.....	17
Discussion.....	18
The Shared Planning Coordinating Representation	19
The Object List	20
The High-Level Planning Coordinating Representation.....	23
Conclusions	24
Comments on the Design of Coordinating Representations.....	25
Future Work.....	26
Acknowledgements.....	27
References	27

INTRODUCTION

People work together to achieve coordination in their cooperative efforts. The organization of a behavior is determined during the interaction among the participants. For recurrent activities there is an accumulation of knowledge about the effective organization for a specific context-dependent behavior. Even with the accumulation of knowledge, the participants of a repeating cooperative activity must continue to interact to adjust their behavior to the dynamics and details of the immediate context. Trouble spots in activity (and novel activity) require relatively more interaction to organize behavior. For recurrent activities, signs may be introduced at the scene that mediate the interaction among the participants and provide some organization (e.g., a stoplight).

Some interesting features of this arrangement are:

1. Interaction is external (and therefore accessible to analysis).
2. Cooperative behavior necessarily entails interaction.
3. Trouble spots in the organization of activity require increased amounts of interaction.
4. Structure is created to organize recurrent context-dependent interactions.
5. Coordination of activities can be organized by mediating artifacts.

For the development of computer-mediated cooperative activities, these features are especially significant in that the interaction is mediated by the computer and therefore recordable and more readily available for analysis. If the analyst can identify the most costly repeating coordination tasks and any secondary structure that the users develop to organize interaction, she can potentially introduce an external representation that will mediate that interaction in a way that improves performance. We will refer to a permanent structure designed by a non-participant that mediates an interaction of a particular sort as a *coordinating representation*.

In this paper, we are interested in the development of coordinating representations that support the participants' interaction with regards to shared domain objects. Two issues of relevance are:

1. Can a designer construct a coordinating representation that improves performance for the participants in a computer-mediated task beyond what the participants achieve on their own from their repeated interactions? If so, in what manner does it improve performance?
2. Can a methodology be developed for the analyst that capitalizes on the work participants do to organize their interaction? How can the analyst reliably produce a coordinating representation that will effectively mediate and organize the interaction among the participants?

These issues will be explored in the context of VesselWorld, a research platform for studying distributed collaboration within small groups. VesselWorld is a simulated world consisting of a harbor with sunken barrels of toxic waste that the participants must discover and then clean up using a tugboat and two cranes (one participant controlling

each). The tools through which participants can achieve coordination in VesselWorld consist of textual chat, a map that shows the layout of the harbor, and non-identical views of each participant's near surroundings.

This paper describes and discusses a pilot study and a subsequent experiment. In the pilot study, 30 hours of VesselWorld data were collected; the main purpose was to create a corpus of analyzable data. A discourse analysis identified recurrent patterns of coordination, repeated errors, and the development of discourse structures for organizing talk and behavior. The analysis provided a basis for designing and implementing three coordinating representations for VesselWorld. The experiment consisted of two populations, one in which groups made use of the coordinating representations (the CR groups), and another in which they performed the task without them (the non-CR groups). Each population consisted of three groups of three subjects, and each group used the system for about 10 hours over several sessions. Our analysis looked at performance (e.g., time taken to complete tasks, number of events produced, and number of errors), and closely examined the interaction among participants in each group.

Our results show that all groups in both populations improved their performance over time. The CR groups out-performed the non-CR groups right from the start. CR groups talked less than non-CR groups and did less work at maintaining consistent representations of shared domain objects. Our study also shows that the planning work of the CR groups was improved over that of the non-CR groups. Clock time, interface work, and confusion among the users about the details of each other's plans were all reduced. The detection of potential errors at plan time was simplified, thus reducing errors in action.

BACKGROUND

Sociology developed the view that the primary site of everyday activity is face-to-face interaction (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; Suchman, 1987; Clark, 1996). One topic that has gotten considerable attention in the interdisciplinary literature is how participants in a cooperative activity manage to coordinate their behavior. Coordinating communication has been a significant topic in the literature on computer-mediated activities (e.g., Malone & Crowston, 1990).

In a face-to-face activity, 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 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. Because they can see one another, they can also use body position, language, and gaze to communicate information. The actions that form their conversation and activity occur sequentially.

The term *structure for behavior* will be used to refer to the kinds of organizational information shared among participants in order to support coordination. Examples of structures established as part of the common ground are agreements about the assignment of roles, the path, the manner of an activity, and a schedule that orders a set of joint actions. For the difficult portions of their task, the participants may explicitly create a shared plan (Grosz & Krauz, 1996), an agreed-to structure – you do this and I'll do that – or design a structure on the spot (e.g., Whitaker, Brennan, and Clark, 1991) to support the coordination of 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 and expectations about the organization of those behaviors (*expectation about organization of behavior*: Alterman & Garland, 2001).

Cognitively, the joint behavior of two actors engaged in a cooperative activity can be modeled as having both individual and social parts (Alterman & Garland, *ibid*). 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 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 to produce secondary structures that will better organize behavior, in the future, for similar sorts of activities. Prior knowledge of conventional organization for recurrent activities creates expectations that must be either confirmed or revised as the current activity progresses. Nevertheless, expectations reduce communication and planning costs, while improving overall performance at cooperative activities. The necessity of communication in order to act and the ongoing process of design among the participants to organize their behavior provides the basis for a participatory approach to tailor-making groupware systems.

Coordinating Representations

Representational structures designed to support individual and joint behavior can become realized as external representations that distribute cognitive activity into the environment (e.g., Goodwin & Goodwin, 1996; Perkins 1993; Norman, 1993). Hutchins provides several illustrative examples of how cognition is distributed. One example is how ‘speed bugs’ on the airspeed indicator in the cockpit of an airplane transform a short term memory task for the pilot into a spatial reasoning task that is easier to perform during the landing of an airplane (1995b). Another example is the Mercator projection chart, which is used by Navy personnel to do much of the computation in “fixing a position” for a vessel underway (1995a); the Mercator projection chart is a specially constructed artifact that supports this computation.

Those external representational structures that are available at the scene of a recurrent cooperative activity and are explicitly (primarily) designed (by a non-participant) to improve coordination among the participants will be referred to 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 an external 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. However, 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.

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 into the design of the task environment. As a mediating artifact, a coordinating representation has both a sign and a tool function (Vygotsky, 1978). As a sign it modifies how the individual thinks; it creates expectations among participants and thereby reduces the effort required to achieve coordination. As a tool it changes the ways in which the collaborating actors proceed: it organizes their activity. Mediating artifacts are critical in the accumulation of modifications to practice over generations of actors within a community (e.g., Cole & Engeström, 1993). 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.

Computer-Mediation of Activity

Within the literature on CSCW, the coordination of computer-mediated communication has engendered significant 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 frameworks that can be implemented to 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 (Stefik et al., 1987) provided a shared virtual “whiteboard” for exchanging ideas, drawings, and texts during a same-time / different-place computer-mediated activity (Ellis et al, 1991). Experimental studies examined and evaluated the function and utility of whiteboards (Tang, 1991; Bly 1988). These studies produced insights and constraints that inform the design and development of groupware applications (e.g., Greenberg, Roseman, and Webster, 1992). External media such as the whiteboard enable users to construct shared data structures that help to organize their activity (Whittaker, Brennan, and Clark, 1991). Communication that is supported by external representational media (i.e., the whiteboard) can be modeled using an interactive model of communication (Tatar, Foster, and Bobrow, 1991).

A whiteboard enables participants to create shared representations as they go along. For continued activities, some of these representations may be convertible into coordinating representations. For some applications a whiteboard will suffice; for other applications the introduction of a coordinating representation may be preferable.

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. The COORDINATOR adds structures to the negotiation of commitments among the participants of a collaborative activity (Flores et al., 1988). The Information Lens system adds formatting structure to email messages based on the type and function of the message being exchanged (Malone et al., 1989). NoteCards is a hypertext environment that organizes the online sharing of information among the participants of an ongoing project (Trigg, Suchman, and Halasz, 1986). The Ariadne language provides a set of criteria and a toolbox for developing computational artifacts that support articulation of collaborative work (Schmidt & Simone, 1996).

These latter sorts of efforts provide representational languages that can be used to structure communication for a wide variety of tasks and domains of activity. A complimentary problem is to determine how to use these representational languages to facilitate a specific set of communications. For example, an analysis of a corpus of roughly 2000 electronic messages exchanged among the 117 developers of Common Lisp identified several kinds of typical communicative acts (i.e., *genres*) with a socially defined and recognized communicative purpose (Orlikowski & Yates, 1994). Three examples of genres that they discovered were the memo, the proposal, and the ballot. Given this sort of analysis, an asynchronous communication system like Information Lens could be tailored to provide organizational structure to email messages for each expected genre of communication within the community.

The coordinating representation also begins with an analysis of a particular computer mediated practice. It is a fix to a particular coordination problem that emerges within a community of users who engage in a set of synchronous activities mediated by a particular groupware system. A coordinating representation can be constructed using any number of representation languages.

METHOD: THE VESSEL WORLD SIMULATION ENVIRONMENT

The coordinating representations that are needed for a given application will depend on the character of the interaction that emerges from an activity within a context, once there is a tool in place and the users begin to collaborate. This necessitates that a basic version of the system should be deployed first, one that only includes coordination mechanisms flexible enough that the users can come up with their own secondary structures in response to the problems that arise. 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.

Figure 1 shows the basic methodology we propose for adding coordinating representations to a groupware system. A groupware system is developed using context-free coordination tools (e.g., whiteboard or textual chat) to support the users' cooperative activity. In some cases there are difficult problems in coordination that confront the users during the normal course of their mediated behaviors that are not easily or efficiently resolved using context-free tools. In these cases, data generated in the pilot study is analyzed to find the recurrent areas of coordination and to determine what sorts of

organizational structure emerges in conversation to handle them. Based on this analysis, a second version of the system is developed that includes coordinating representations. If necessary steps 3 and 4 are repeatable.

1. Build a base system that includes general-purpose coordination methods only (e.g., whiteboard, textual chat)
Sometimes this is enough
2. Perform pilot study with base system
3. Analyze data to discover recurrent problems of coordination and secondary structures users devised to organize those behaviors.
4. Rebuild system using coordinating representations suggested by analysis.

Figure 1: Basic Methodology

To some this methodology may be surprising. A rule of thumb for the development of software is that design done up front significantly reduces development cost (e.g., Brooks, 1995). There are advantages to a participatory development process that includes user feedback at each stage of design up to the implementation of a prototype (Scaife, Rogers, Aldrich and Davies, 1997), but the process of user-informed development does not have to stop there. The introduction of technology changes the collaboration within a community, and consequently further adjustments to the system may prove necessary. A technique of analysis that examines the interaction among participants as they collaborate online is potentially a significant method of accessing input from the user. The remainder of this paper will explore the consequences of, and the methodology we propose for, introducing coordinating representations into a computer-mediated activity.

VesselWorld: Task and Interface

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 the captain of a ship. 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. There are many complications in clearing the harbor. An efficient solution requires planning. Some barrels are large and require the two cranes to join together and lift them simultaneously. There is a lot of information (e.g., what equipment is needed to retrieve a particular barrel) that needs to be discovered and shared in order to complete a session. Additional complications arise because the participants have limited (and non-identical) areas of perception, and the harbor must be searched to discover the toxic waste. VesselWorld was demonstrated at CSCW 2000 (Landsman et al, 2001).

In VesselWorld, segments of activity are divided into *rounds*. During a round of activity, participants plan out their future actions explicitly, and then submit them to the system. Once a participant has submitted 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 the world, and the next round begins.

The WorldView (the large window in Figure 2) 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. It depicts the harbor from the participant's point of view; only a limited region of the whole harbor is visible at any one time. A second window of information is used for editing and displaying the user's current plan. A third window allows the user to access more detailed information about visible objects. There is also a textual chat window that enables the participants to communicate with one another.

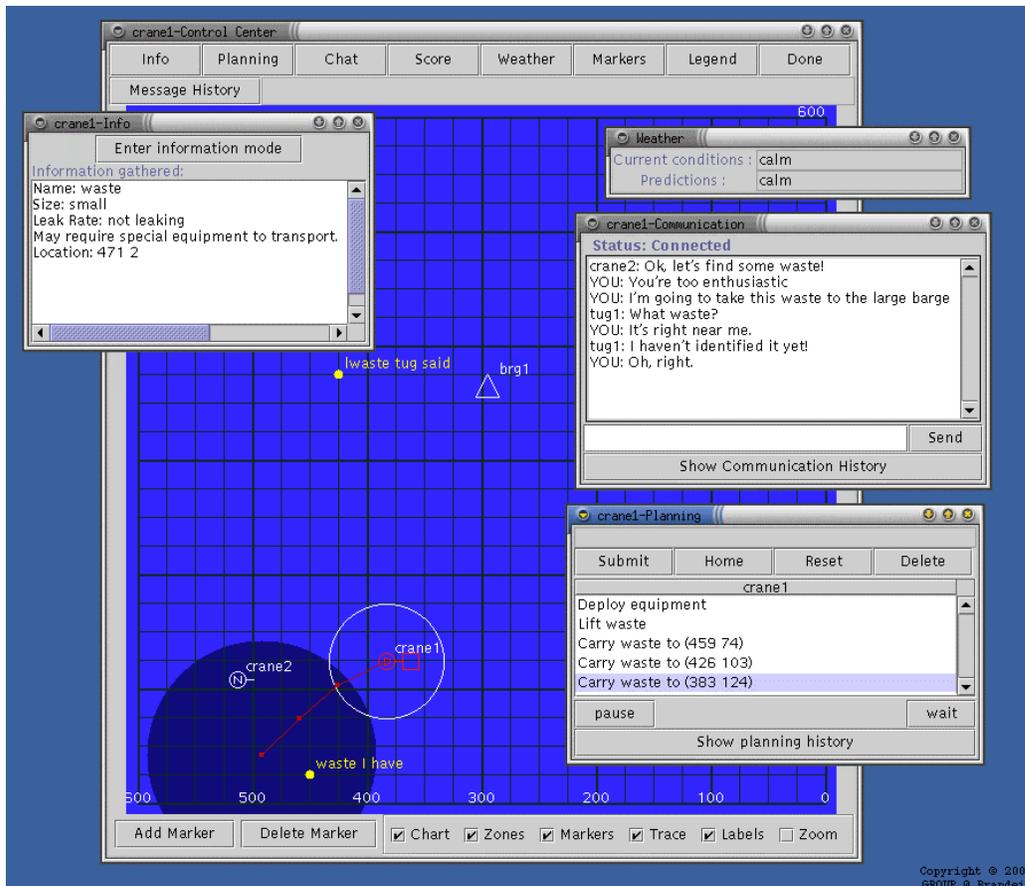


Figure 2: The interface for the basic system.

PILOT STUDY

A group of three subjects used an early version of the VesselWorld for more than 30 hours, split over many sessions of problem solving; this data was collected by Susan Kirschenbaum at NUWC. The subjects were Navy professionals, and were not paid. The

group worked in a partitioned space where they could not see each other's screens. They were not permitted to speak to one another. This set-up ensured that all communication was performed via the system, where it was logged for subsequent analysis.

After this initial data was gathered, a discourse analysis was performed. The analysis of the logged data was supported by interviews with the subjects, who shared their observations about the system as it supported coordination. From these sources came both redesign ideas for VesselWorld, and a clarification of the structure of our methodology.

Discourse Analysis

The discourse analysis of participant dialogue taken from the pilot study of VesselWorld focused on three indicators that the introduction of a coordinating representation might be advisable:

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

Both recurrent patterns of coordination and repeating errors are the kinds of situations in which the participants potentially would introduce secondary structure 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 structure to improve the situation.

Of the three indicators, the third seems to be 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 relies exclusively on the third indicator. 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.

There is no guarantee that either the organizational structure 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 context-free form of communication like textual chatting.

An Example of Discourse Analysis

The problems inherent in the cranes jointly lifting or moving a large or extra large toxic waste made for a recurring source of difficulty for the users. To remove a large waste requires two levels of coordination:

1. Coordination of action
 - a. The correct crane must deploy the equipment necessary to lift the waste.
 - b. The cranes must join together
 - c. During the same round of activity, the cranes must jointly lift the waste
 - d. During the same round(s) of activity, the cranes must jointly carry the large waste to the barge, if necessary
 - e. During the same round of activity, the cranes must jointly load the large waste onto the small barge.
2. 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

The first level of coordination is dictated by the task domain. These actions must occur in order and must be synchronized; errors in coordination or ordering can result in failure and leakage of toxic waste. The second level of coordination emerges in conversation structure as the participants attempt to coordinate these kinds of activities. In other words, the talk is coordinated so that the actions may be coordinated.

The pilot study data showed that participants used adjacency pairs (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 that adjacency pairs will occur for each of the coordinated actions.

Figure 3 shows a sample of the dialogue where the participants used adjacency pairs to coordinate the handling of a large barrel of toxic waste. 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. Because the conversation of the users is mediated through textual chat, adjacency pairs do not strictly speaking occur one after the other; their positioning depends more on the typing speed of the users. Other kinds of comments may end up interposed along the way, e.g., the Tug's comment that the next waste will be an extra large that "needs nothing".

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 3: A sample dialogue.

The secondary structure created by the participants of the pilot experiment is an indicator that a coordinating representation could potentially improve the coordination of talk on these occasions. A coordinating representation could thereby improve the participants' performance at these kinds of cooperative activities.

Design of the Coordinating Representations

In total, the analysis of the pilot study discourse for VesselWorld identified three recurrent areas of extensive coordination work:

1. **Timing** of closely coupled cooperative activities involving the domain objects.
2. **Exchanging information about**, and establishing references for, **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 general areas has been suggested by prior theoretical analysis. There are, however, many other potential problem areas. The designer's task was to determine which things were most problematic for the task at hand. Participant errors and increased effort were strongly associated with these three recurrent areas of coordination. The participants in the pilot study generated secondary structure that simplified talk about the timing of joint actions. They also generated secondary structure (notational conventions) that was used to describe properties of the wastes (e.g., size and location). 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 well-defined secondary structure that was created in order to simplify communication.

After these problem areas were identified, coordinating representations were developed. During this phase of development, coordinating representations were structured in a manner that reduced the physical work and the cognitive load of the individual user in creating and accessing the shared coordination information.

The coordinating representation shown in Figure 4 (the shared planning CR) 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 few planned steps of Crane1 are to deploy equipment and then lift some waste.) Each user has control of only one column, her own. This representation improves timing on entry and exit of phases within a tightly coordinated activity by allowing participants to compare each other's next few projected actions.



Submit	Home	Reset	Delete
crane1	crane2	tuq1	
Deploy equipment	Move to (424 146)	Move to (119 92)	
Lift waste	Move to (409 153)	Move to (200 145)	
Carry waste to (453 67)	Move to (373 243)	Move to (363 251)	
Carry waste to (414 93)	Move to (338 329)	Move to (511 338)	
Carry waste to (390 109)		Attach to sbrg1	

pause wait

Show planning history

Figure 4: Timing of joint actions.

The second coordinating representation mediates communication of information about domain objects; it does not organize domain actions. A list of objects (with relevant properties) allows users to more systematically keep track of objects in the domain (see Figure 5). This information is visible to all users and can be edited by any user. Columns provide information about the name, object type, location, and equipment needed for a given object. Each participant can maintain a different view of the information shared in the object list. An important feature of the object list was that entries could be displayed on the WorldView as markers; each user could choose when to display the markers on their own map. Each marker was named automatically using the name entered in the name field of the object list. The organization of this information reduces the cognitive load for the individual, by organizing information relevant for decision making into a predetermined representational structure.



Figure 5: 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 6 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 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., GO, SEARCH, and CONTAIN. Color-coding of entries in the high-level plans allows participants to indicate both accomplished tasks and future commitments.

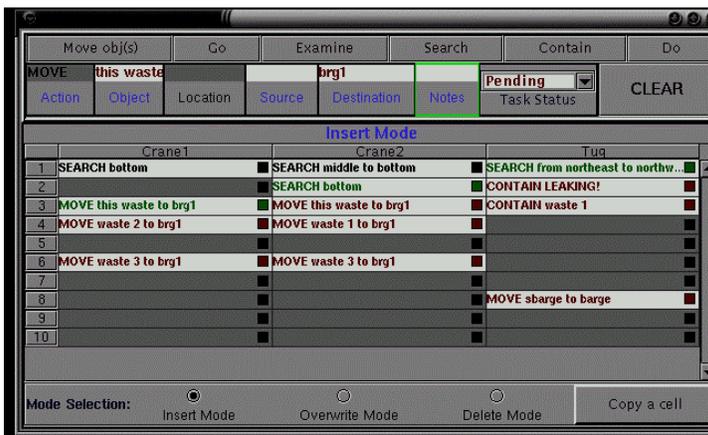


Figure 6: High-level planning.

EVALUATION

The Experiment

After a redesign of the VesselWorld system, we conducted an experiment to compare group performance with and without coordinating representations. We performed a single-variable experiment to assess the impact of coordinating representations on the performance of groups of subjects using VesselWorld. One set of groups, the control groups (which we will call the non-CR groups), used an improved, more stable version of VesselWorld otherwise similar to the one in the pilot study. The other set of groups (the

CR groups) used a version of VesselWorld with the coordinating representations enabled. Each set consisted of three groups of three subjects. The groups consisted of a mix of area professionals, mostly in computer-related industries, and undergraduate students; all were paid a flat fee for the experiment.

Each group was trained together for two hours in use of their system, and then solved VesselWorld problems for approximately ten hours. To alleviate fatigue concerns, the experiment was split into three four-hour sessions. Subjects were asked to fill out entrance surveys to obtain population data, and exit surveys where they could give feedback about their experience with the system and the coordination issues arising in their group.

Findings

The experiment produced number of major results. In general, the performance of groups with coordinating representations improved in many measures: clock time, number of system events generated (an indicator of interface work), and number of errors committed. Performance in the trouble areas we had previously identified (close coordination, domain object reference) was notably improved, with errors due to miscommunication of object information significantly reduced.

However, the data is not as clear as one would like. Because of the relatively small sample population (due primarily to the arduous nature of producing data), variability of group performance due to individual differences was high. For example, personal strife between subjects in one group led to severely reduced performance in early sessions, skewing the data slightly. Likewise, one subject's comparatively low computer proficiency introduced a small bias in that group's clock time. These sorts of issues seem unavoidable when dealing with small, self-selected user populations.

Another issue was quantifying the difficulty of the varying problems presented to subjects. A set of random problems was produced, and subjects were given a succession of problems drawn from this set. However, groups did not necessarily see the same problems, nor in the same order, and because of differences in performance, did not complete the same number of problems over their ten hours of problem solving. To account for this, a general measure of the complexity of a particular problem was devised, taking into account the quantity and type of the wastes in the harbor, their distance from the large barge, and the number of small barges available to the subjects. This metric was used to normalize results.

General quantitative results are presented in Figure 7. These results represent comparison of the final five hours of play for each group, by which point the performance of the groups had stabilized. As discussed above, these results are normalized over the computed complexity of the problems being solved. The most significant effect, though not the one of greatest magnitude, is the 57% reduction in communication generated. Also highly significant is the 49% reduction in clock time. Only slightly less significant is the reduction in system events (mouse clicks, etc.), down 38%. These reductions were all expected. Also as expected, overall domain errors (errors in performing domain

actions which led to a toxic spill) were reduced by 61%. The variance of this measure was quite high due to the low frequency of errors; this reduced its confidence below statistical significance ($p < 0.2$). One measure that we expected to fall significantly was rounds of activity. However, as can be seen below, the small reduction in this measure was not at all significant ($p < 0.35$).

	Improvement
Communication	57% ($p < 0.01$)
Rounds of Activity	22% ($p < 0.35$)
Domain Errors	61% ($p < 0.2$)
System Events	38% ($p < 0.06$)
Clock time	49% ($p < 0.01$)

Figure 7: Improvement of CR groups over non-CR groups; final 5 hours of play

Where the above measures of performance compare CR and non-CR groups during the last 5 hours of problem solving, we also wanted to compare differences in performance from the early sessions to the later sessions for each group with and without CRs (see Figure 8). All groups improved their average performance, as seen by 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. The average reduction in communication costs (as measured by the number of utterances) was similar and significant (51% versus 56%, $p < 0.05$) for CR and non-CR groups. For clock time the average improvement for CR groups was significantly better than non-CR groups (80% versus 66%, $p < 0.1$).

	Non-CR groups	CR groups	
Communication	51%	56%	$p < 0.05$
Clock Time	66%	80%	$p < 0.1$

Figure 8: Within-group Improvement over time

We also asked subjects to fill out an exit survey to gauge their general reactions to the system and to working with each other (see Figure 9). We asked seven targeted questions, each answerable on a scale of 1 to 7, and four open-ended questions asking for reaction to the system and possible improvements. The results of the survey show that, on average,

subjects in the CR groups reported that their system was harder to use: with 1 being “Easy”, non-CR subjects rated their system 2.4, and CR subjects rated their system 4 ($p<0.01$). However, they did not find the game itself harder (2.8 vs. 3.0, $p<0.35$), nor did they find it harder or easier to stay coordinated with the other captains (3.2 vs. 3.0, $p<0.7$). Both sets rated their reaction to the system as average (4.3 vs. 4.2, neither “A breeze” nor “Frustrating”, $p<0.88$).

	Non-CR groups	CR groups	
Difficulty using system	2.4	4	$p<0.01$
Difficulty of game	2.8	3.0	$p<0.35$
Difficulty of staying coordinated	3.2	3.0	$p<0.7$
Reaction to system	4.3	4.2	$p<0.88$

Figure 9: Results from exit survey

There were some unexpected results as well. One of the coordinating representations, the high-level strategy planning CR, was not used once training in it was complete. Certain columns of the object list went unused. Subjects in the CR groups gave some insight into why: “...the Object list had too many options. Many weren't used because we were in constant chat contact.” Another wrote: “We never used the Strategy Window because we could see what we were doing in the planning window.” One user’s analysis of the fundamental problems with the high-level planning CR was especially interesting: “...since all plans must de facto be agreed upon by all (relevant) players, negotiation via the Chat window is required. Since the plans are discussed in detail there, putting those plans in the Strategy window would be redundant.”

While the non-CR groups were less verbose, one comment was especially telling: “A list of objects that everyone can see would help a lot with the coordination.” This request is, of course, fulfilled in the creation of the object list.

Discussion

All events that occur during a VesselWorld problem solving session are recorded in a log file by the system. A VCR-like program (called SAGE) was built to review the decision making of each group and examine how the participants in a VesselWorld session coordinate their activities and the exchange of information (Landsman & Alterman, 2001). SAGE enables an analyst to replay a session of problem solving from an omniscient perspective; it consolidates the individual views into a useful format for the analyst. Because the data saved has an inherent structure, the analyst can search through the data using any number of criteria; e.g., he can move forward to the next

communication, round, plan action, or other such action within the system, allowing for easier review of the bulky data logs.

Close coordination during some phases of activity requires special attention be paid by the subjects to the timing of actions. Discovering new barrels of toxic waste, establishing if, when, and what kind of equipment is needed are tasks that develop throughout a problem session, and progress is made intermittently. Because of uncertainty it is not possible to reason out at the beginning of the session a complete plan for the activity. To a certain extent, the participants must improvise as they go along. Nevertheless, it is important for the participants to plan out some of their activity to generate an efficient solution.

Below we characterize differences between how the CR and non-CR groups organized their activities. The comparisons are presented so as to highlight the effects of each coordinating representation.

The Shared Planning Coordinating Representation

• Reporting on plans to one another is extra work. The CR groups share this information via the shared planning coordinating representation and do not need to represent this information in the textual chat.

All three of the non-CR groups in our experiment invented secondary structures involving adjacency pairs to stay coordinated during the closely coupled joint activities that were necessary for handling large and extra large wastes. Figure 10 shows, for one of the CR groups, the 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. — End of round 9 —
9. CRANE1:: that was *last* time : -)
10. — End of round 10 —
11. TUG1: G1, loaded.
12. — End of round 11 —

Figure 10: Conversation while lifting and loading an extra large waste.

All of the structure that the non-CR groups produced in chat to organize their activity (e.g., see Figure 3) was being created by the CR groups in the shared planning coordinating representation, without any extra effort. In order to submit an action to the

system the users need to add it to their “plan” anyway. 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 CR. 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 removal of a need for redundant descriptions of the users’ plans reduced clock time, interface work, and confusion among the users about the details of each other’s plans.*

The non-CR groups resorted to adjacency pair structure to organize tightly coupled actions because there were confusions about timing for actions like joint lifting, which require precision timing. These confusions occurred because the participant’s access to each other’s plans was indirect, depending on a redundant description produced in the textual chat window. The creation and maintenance of the textual descriptions of plans increased clock time and interface work for the non-CR groups.

- *The shared planning coordinating representation affords the opportunity for participants to check each other’s plans during closely coupled actions, and thereby reduces errors.*

With access to the shared planning CR 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 equipment, another actor can remind him. 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 before they acted (and before errors could be produced). Without the coordinating representation this sort of reviewing of one another’s plans during the course of action could only occur when plans were being reported in the chat window, and even then detection of mismatched activities depended on the precision of the description.

The Object List

- *Participants who have access to the object list do less chatting than non-CR groups because they can share information about shared domain objects via the coordinating representation.*

The object list is an external representation that allows the participants to coordinate their efforts in exchanging information about domain objects over an extended period of time. Figure 11 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 basstartd.
3. CRANE2: I got an XI 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 11: Opening dialogue.

Figure 12 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 non-CR groups is now replaced by exchanging information using the object list coordinating representation. Identifiers are attached to each of the objects that are found; this simplifies future conversational exchanges about the wastes. Because pointing and clicking on wastes can add entries to the object list, precise locations for each of the 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.

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 12: The object list

- *The non-CR groups must do more work (elapsed clock time and number of system events) because they must keep redundant separate private representations of the same information. This is in contrast to the CR groups, which keep a single shared representation.*

For the non-CR groups, textual chatting is the only method available for exchanging information about wastes and the location of the small barges. It is the responsibility of each actor to add that information to his or her private representation (either by marking the map or remembering), or be prepared to examine the history of chatting at some appropriate future time.

- *Extra work was needed by the non-CR groups to compare and align their individual representations of shared domain objects.*

An important feature of the object list is that all the information is shared. Much of the consistency checking that the non-CR groups 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 repairing these discrepancies. The non-CR groups had to talk extensively in order to reconcile their individual representations of the shared domain objects. One example of a strategy used by the non-CR groups for avoiding differences in assessment is to engage in a conversation to review the status of one or another of the shared domain objects. Whenever discrepancies in the assessment of a situation unexpectedly develop, the participants engage in repair work to reconcile alternate representations of “reality”.

- *Evidence that the work required to review and maintain a consistent representation of shared domain objects was problematic for the non-CR groups is provided by one of those groups, which invented a secondary structure to organize this activity.*

In order to ensure that each participant was maintaining the same list of toxic wastes, one of the non-CR groups would periodically do a “marker check” (see Figure 13). 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 13: A marker check

- *An analysis of the usage of the object list showed that it was used mostly for maintaining information about shared domain objects that was likely to have a long lifetime of relevance.*

The object list was mainly used to maintain information that was intermittently relevant over an extended period of time. 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. Only 30% of the leaks that occurred were recorded in the leaking field of the object list. Only 3% of the entries created used the notes field. Our belief is that the differences of usage are a function of the duration of relevance of each kind of information. Information that was relevant for extended periods of time (like the equipment needed to lift a waste) was recorded in the object list. Information that was relevant for only a short period of time and could be easily retained in working memory, or was easily accessible by other means, was less likely to be recorded.

The High-Level Planning Coordinating Representation.

The high-level planning coordinating representation was not used by any of the CR groups. The subjects indicated in the exit surveys that they thought the chief problem with the high-level planning CR was that it required too much work to use given the

rewards it provided. Further analysis shows that the problems that the high-level planning CR was designed to fix continued to occur. The lack of utility of the high-level planning CR is included in the discussion section below on the design process for coordinating representations.

CONCLUSIONS

In order to act cooperatively the actors need to communicate. For the portions of the activity where the coordination is complicated, or for novel activities, it is especially important for the participants to talk. By introducing structure into their conversation, they organize their activities; i.e., the organization of activity is preceded by the organization of talk about the activity. However, even for activities that are well organized and have structure available as a part of the common ground, communication will still be necessary. The participants will need to communicate, for example, to time the actions they expect to occur. The amount of conversation depends inversely on the similarity of unstated expectations and the amount of secondary structure available as a part of common ground.

Many of these features of the connection between talk and action were exhibited in the data we collected on the non-CR groups in our experiment. The participants tended to talk less for problems of equivalent difficulty as time went on. The participants introduced secondary structure to organize both their talk (e.g., notational conventions) and their action (adjacency pair structures for tightly coupled actions). As the amount of secondary structure increased the amount of conversation decreased.

Methodologically, what is significant about the connection between talk and action is that for computer-mediated tasks involving whiteboards or textual chatting, all coordination work that results from the misalignment of expectations is externally represented and is in a form that is analyzable. The main issue from an engineering point of view is to determine if it is possible to capitalize on this arrangement. One issue is whether a method of discourse analysis can be established that can reliably identify areas of complicated and/or inefficient cooperative behavior. Once a method of analysis is established, is it possible for the system builder to introduce coordinating representations that improve user performance beyond the performance they achieve from textual chatting or whiteboard interactions?

Our study provides evidence that the introduction of coordinating representations improves performance over that achieved by textual chatting. It also showed that the CR groups talked less than non-CR groups. By sharing a representation of domain objects it was no longer necessary to keep individual representations; without the object list the participants were keeping redundant representations. Removing the need for redundant representations reduced clock time and interface work. It also reduced work at repair and reconciliation in maintaining a common viewpoint. The fact that one of the non-CR groups invented the marker check to address this problem is an indication of the complexity of the task.

Our study also showed that the participants' planning activity -- their ongoing work to organize their emergent activity -- was improved. The shared planning CR removed the need for a redundant description (communicated as textual chat) of each actor's plan for tightly coupled actions; this reduced clock time, interface work, and confusion among the users about the details of each other's plans. The shared planning CR also simplified the correction of errors at plan time, thus reducing errors in action. Finally, because the participants could use CRs to organize parts of their interaction, the textual chatting could be used to emphasize long range planning issues.

Comments on the Design of Coordinating Representations

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 in coordination
- The development of secondary structures to organize talk to support coordination.

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 made the most errors. The third indicator depends on the analyst being able to recognize where the participants have developed secondary structures to coordinate their talk to fix repeated errors or organize 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 experiment, all of the non-CR groups created shorthand notational conventions (indicator 3) to describe domain objects. One of the non-CR groups also created a secondary structure that did some of the work of the object list: the marker check procedure that was performed by this group was intended to ensure that each participant had the same list of toxic wastes. The experiment showed that all the CR groups more effectively achieved this functionality via the object list coordinating representation.

The analysis of the pilot study data predicted the utility of the shared planning CR 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 CR was an effective coordinating representation. We observed several occasions where the crane operators in the CR groups used it to synchronize their joint actions without having to resort to talk. The quantified comparisons showed that on average the CR groups spent less time, and

generated fewer utterances, 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 experiment, none of the groups developed secondary structures to support high-level planning to manage multiple open plans and commitments. 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 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 CR, 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 other sorts of high-level planning might have produced a more useful coordinating representation. Another potential source of difficulty was that much of the data the high-level planning CR was designed to contain had a relatively short lifetime of relevance. With the object list, those columns where information was relevant for an extended period of time were more likely to be used than those with a short lifetime of use. Similar constraints were undoubtedly at work in the high-level planning CR. 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 CR semi-automatically (Introne & Alterman, 2000).

Future Work

We are currently developing (using the methodology described in this paper) a groupware system that supports the collaborative development of websites; this groupware system will be used in an Internet Studies course at Brandeis that is taken by students from a mix of disciplines. We also intend to refine our methodology and will revisit the design problem we encountered with the high level planning CR.

THYME is a framework for building component-oriented groupware that is being developed at Brandeis; it provides a set of basic groupware constructs (e.g., a chat room) from which groupware systems can be assembled. A RAD tool for THYME is being built that will be used by teams of undergraduates to develop groupware systems for a term project in an HCI class taught at Brandeis. The VCR-like program we used for the analysis of our experimental data (SAGE) is a prototype of an analyst's tool that will

automatically be supplied for a THYME-built groupware system. THYME and SAGE provide the engineering foundation to efficiently and effectively build tailor-made groupware systems that include coordinating representations. We are also working on devising AI-based adaptive technology that exploits the existence of coordinating representations; the key idea is that the coordinating representations provide relevant information in a structured format.

Also under development is a second version of the analysis tool that allows the designer to partition and sort the discourse into segments and semi-automatically profile and model interactions among participants. By comparing discourses about similar coordination problems, the analyst can more readily extract patterns of talk and action (indicator 1), recurring errors in coordination (indicator 2), and note secondary structures created by participants (indicator 3).

ACKNOWLEDGMENTS

The authors are appreciative of the helpful comments we received on earlier drafts of this paper. We would like to thank Elizabeth Churchill and Tom Erickson for their editorial input.

This work was supported by ONR Contracts N00014-96-1-0440 and N66001-00-1-8965. Additional funding came from NSF grant EIA-0082393.

REFERENCES

- Agre, P. (1997) *Computation and Human Experience*. Cambridge University Press.
- Alterman, R. and Garland, A. (2001). 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.
- Brooks, F. (1995) *The Mythical Man-Month*. Addison-Wesley.
- Clark, H. (1996). *Using Language*. Cambridge University Press.
- Clark, H. and Brennan, S. (1991). Grounding in Communication. In Resnick, Levine, and Teasley editors, *Perspectives on Socially Shared Cognition*, 127-149.
- Cole, M., & Engeström, Y. (1993) A cultural historic approach to distributed cognition. In G. Salomon (Ed.), *Distributed Cognitions*, 1-46. Cambridge, MA: Cambridge University Press.

- Ellis, C., Gibbs, S., and Rein, G. (1991). Groupware: some issues and experiences. *Communications of the ACM*, 34,1, 9-28.
- 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.
- Garfinkel, H. (1967). *Studies in Ethnomethodology*. Polity Press.
- Goffman, E. (1983) The interaction order. *American Sociological Review*, 48:1-17.
- 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.
- 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.
- Landsman, S. and Alterman, R., (2001) Analyzing Groupware Usage In Context. Technical Report: CS-01-218, Brandeis University.
- Landsman, S., Alterman, R., Feinman, A, Introne, J. (2001) VesselWorld and ADAPTIVE. Technical Report: CS-01-213, Brandeis University. Demo given at CSCW 2000.
- Lave, J. (1988). *Cognition and Practice*. Cambridge University Press.
- Malone, T., Grant, K., Lai, K-Y., 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).

- Malone, T. and K. Crowston (1990) What is coordination theory and how can it help design cooperative work systems?, in CSCW'90, Proceedings of the Conference on Computer-Supported Cooperative Work, ACM, New York.
- Norman, D.A. (1993) *Things That Make Us Smart*. Perseus Books, Cambridge, MA.
- Orlikowski, W., and Yates, J. (1994) Genre Repertoire: Norms and Forms for Work and Interaction. MIT CCS Technical Report 166.
- 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.
- Scaife, M., Rogers, Y., Aldrich, F., and Davies, M. (1997) Designing For Or Designing With? Informant Design for Interactive Learning Environments. CHI 97, Atlanta, GA. p. 343-350.
- Schmidt, K. and Simone, C. (1996): Coordination mechanisms: Towards a Conceptual Foundation of CSCW Systems Design, Computer Supported Cooperative Work 5, p. 155-200.
- Schutz, A. (1967) *The Phenomenology of the Social World*. Translated by Walsh, G. and Lehnert, F. Northwestern University Press.
- Schegloff, E. and Sacks, H. (1973) Opening up closings. *Semiotica*, 8:289-327.
- Stefik, M., Bobrow, D.G., Foster, G., Lanning, S., Tatar, D.(1987) 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 studies 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 *Proceedings 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.