Toward Identifying Process Models in Ad Hoc and Distributed Teams

Joseph A. Giampapa

The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213-3890
(U.S.A.)

garof+@cs.cmu.edu

Katia P. Sycara

The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213-3890
(U.S.A.)

katia+@cs.cmu.edu

Gita Sukthankar

School of EECS
University of Central Florida
Orlando, FL 32816-2362
(U.S.A.)

gitars@eecs.ucf.edu

ABSTRACT

This article reports work on first steps toward characterizing a negotiation process model for ad hoc and distributed groups or teams, so that automation can more accurately track the states of a negotiation from human discourse. We devised three experimental scenarios and ran human subject experiments that involved group decision-making and consensus building. Our experiments showed that the communication patterns of successful distributed ad hoc teams differed in two significantly different conditions. We describe our motivations, experimental design and results.

Keywords

Communication process model, consensus building, ad hoc distributed teams, negotiation, teamwork

INTRODUCTION

Negotiation is a process by which two or more people reach a common understanding, and ultimately, a decision. It consists of complex, interdependent activities that involve gathering and transmitting information from and to the other parties in the negotiation, e.g., acquiring / expressing individual perspectives. A frequent negotiation activity is determining where there are gaps and mismatches in the participants' knowledge. Such gaps or mismatches can be in terms of: understanding, term usage, cultural and social protocol, or simply a difference in valuation that the negotiators apply to the issues they are negotiating. A negotiation process may include an attempt at clarification, often effected through generate-and-test protocols, to rectify the gaps and mismatches of understanding and expectations. The use of diverse rhetorical styles so as to influence and persuade others to a conclusion; and the assessment of a "converted" partner's ability to derive implications from their new position in a way that is consistent with a negotiator's expectations, may also be part of the process. A negotiation may have a process to decide the consequences or the outcome of the negotiation itself.

One form of such a decision could be a consensus opinion, in which the majority of the group members are in agreement. A final step in the negotiation process can be compliance monitoring. While compliance monitoring is not usually considered part of the negotiation process, negotiations that do not conclude with a simple decision or action often conclude with a series of steps that must be followed as a verification of good-faith commitment to the negotiation process. Any deviation or sign of non-compliance can provide justification for the other parties to revoke their intentions to concede, as well.

There are many ways in which negotiation support systems (NSS) assist people in specific aspects of the negotiation process, but participants of negotiations execute a mix of the above process steps in a variety of orders, often omitting steps as well. Negotiators may execute the whole sequence as a very quick and tight-looped spiral, each time concentrating on either one specific sub-issue of the topic of negotiation, or perhaps iteratively deepening it. As the participants progress in the process, they may discover that they need to clarify certain terms, resolve some misunderstandings, and perhaps even follow a nested negotiation process in order to do so. As a long-term research goal, we are interested in developing algorithms to automatically recognize the following features in human-human negotiations: (a) when a negotiation process step is initiated, abandoned, or resolved; (b) when there is a transition from one process step to another; (c) when a new negotiation process is initialized, be it the top-level problem, or for a nested sub-problem of the larger issue; and (d) when a negotiation process is concluded, the algorithms should be able to identify the outcomes and commitments of the interested parties from the negotiation dialogue. With such capabilities, we believe that negotiation support systems will have access to a broader range of problem domains.

In this paper we narrow our attention to the type of negotiation process that takes place in small, ad hoc and distributed teams that are engaged in teamwork. For convenience we refer to this negotiation process as

consensus building for ad hoc and distributed team decision making, since most of the negotiations of the participants will be for reaching a common understanding of their status and progress as a team and for jointly deciding what their course of action, as a team, should be. An ad hoc team is a group that is brought together for the duration of a task and in which the members lack prior experience training together as a team. Hence, one key characteristics of this type of consensus building that we address with this research is that it involves ad hoc teams of three persons. Other characteristics are: (2) there are no pre-assigned roles in the team; (3) there is no hierarchical organizational or authority structure among the team members; (4) there is likely very little shared pre-existing knowledge of the performance capabilities of the teammates; and (5) the members of the team are geographically distributed. Work in the team literature (Fiore et al., 2003) has found that establishing effective communication patterns are the key to creating effective ad hoc human teams; this is especially true for distributed ad hoc teams in which communication is cited as a key problem area (Pascual et al., 1999). By limiting our attention to groups that have: (a) no prior experience together; (b) no pre-specified organizational structure; (c) specific objectives to attain; and a need for high-bandwidth communications, we hope to observe many of the characteristics of the above negotiation process through human verbal and textual communications.

To test our hypothesis, we designed a human subject experiment that situated triadic teams in three randomized test conditions of varying interdependence and complexity. We also devised an utterance classification scheme of 34 codes that correspond to ad hoc team behaviors. (See Figure 1.) These utterance categories were determined in a bottom-up way, based on our observations of the transcribed text of the conversations rather than being based on a theory or preconceptions. The utterances could then be segmented into conversational moves (Hirokawa, 1983), which could then be correlated with a team performance metric in order to determine the degree to which they contribute to the team success. A conversational move unit is an uninterrupted utterance of a team member that has a discrete problem solving function.

Our research is an initial step towards the problem of identifying communication patterns of ad hoc and distributed teams in time critical situations so that suitable NSS aiding strategies can be developed. The identification of patterns is through communication logs collected from human teams. In contrast to prior team communication research, in which the results are finalized to develop guidelines for team training, we are interested in using team communication results for monitoring team performance by automation, such as negotiation support systems, that would be used for team aiding. The focus of our initial human team experimentation is to (a) establish a baseline of human-only teamwork for a given task domain and (b) ascertain the relative importance of different information flows

for the team task in order to derive "insertion points" for assistance in the process. These insertion points are not merely limited to coordination and information flows, but include self-organization, maintenance and task completion. We feel that these results are relevant to human factors in negotiation research because the focus is on identifying steps and transitions within a communicative process.

RELATED WORK

There is much research on group decision processes, consensus building in groups, and techniques for identifying elements of participant dialogues that indicate specific negotiation processes, as exemplified through (Morimoto et al., 2006) and (Weingart et al., 2007). While the labels are applied by human classifiers on an utterance-by-utterance level of granularity, the process is purely for explanatory purposes, and not for "live transition detection" of negotiation state. In the Computational Linguistics community, the focus of dialogue systems is primarily for the purposes of Intelligent Tutoring System (ITS) interfaces, in which case the interface has a very active and specifically didactic role as tutor. The research that best approximated the problem of trying to linguistically identify transitions from the top-level discourse to a sub-topic, followed by a return of focus, was presented by (Carberry & Lambert, 1999). These three references are representative of the current status of research that identifies process state based on textual information in group decision-making dialogues: it is partially addressed from multiple angles, not in one, concerted effort.

(Carberry & Lambert, 1999) consider negotiation in function of dyadic dialogues in which an expert is consulted. The implemented system combines three types of knowledge -- *linguistic*, *contextual*, and *world* knowledge -- in a unified framework that is represented in data structures that encapsulate expressions of first order logic. The system recognizes complex discourse acts and reasons about expressions of belief and certainty when one agent communicates a proposition to another. Should the receiving agent doubt the veracity of the proposition, it may need to temporarily enter into a subdialogue to clarifying the doubts. The system is able to recognize when the subproblem has been solved and the conversation returns to the top-level topic.

(Morimoto et al., 2006) report work in a consensus-based decision-making task by "naïve negotiation" among triads of the same gender, in Japan. The authors coined the term, "naïve negotiation," *to refer to a style of negotiation in which participants share neither particular knowledge on the topic nor the specific skills to reach a consensus, unlike the established style used in formal meetings* (Morimoto et al., 2006). The results of the study reached two conclusions: (1) that the participants often reach a consensus without a thorough examination of proposed opinions; and (2) that the naive negotiation process is characterized by four features that demonstrate the need for a negotiation

support system for novices. These features are: (1) topics were chosen at random; (2) discussants would not discuss topics for substantial amounts of time; (3) participants would abruptly shift to the consensus phase; and (4) participants would present claims without any evidence or reasoning.

(Weingart et al., 2007) report a study that used multi-level modeling and sequential analysis to test hypotheses regarding how negotiators' social motives and the composition of the negotiating group influence group members' negotiation strategies. The authors labeled transcriptions of the negotiation sessions with 32 individual behavioral codes. The labeled utterances were used to detect transformational sequences in the negotiation strategies of the participants, and to reliably predict the strategic behaviors of the subjects. While the research reports on a different domain and uses a different utterance coding scheme, it identified significant strategies, behaviors and transitions in negotiations that could be used to extend this work.

Inspired by such results, our hypothesis was that we could discover and observe team processes, such as decision-making by consensus, though similar methods. Unlike the tasks of the related work, however, there would be a component of plan execution as well as planning. Based on self-assessment of performance in the execution environment, the human team would probably need to revise and re-form their consensus on the course of action to follow.

THE EXPERIMENT

To create a baseline of decentralized ad hoc team performance, we monitored teams of human subjects performing a collaborative search task, a scavenger hunt, in simulation. To implement the task, we reconfigured a scenario in the multi-player game and battlefield simulator, Operation Flashpoint (OFP version 1.96), by customizing the virtual characters, objects on the map, pre-game briefing, map, object triggers, and the scoring mechanism.

In the team scavenger hunt, human subjects must read a map, navigate a 3-D simulated environment and recover a collection of objects (bottles) within a bounded amount of time. The task is designed to evaluate the team's ability to develop and execute a search plan under time-stress. The scenario also creates natural and spontaneous opportunities for group consensus building and decision-making. The team performance metric was simply the number of objects that could be collected in twenty minutes.

Seventeen teams of three paid subjects, each, were recruited to participate in the pilot study. Human subjects self-assessed and reported their abilities to play first-person video games in terms of the following classification: novice, medium expertise, or expert. Combined expertise of the teams varied from "two novices and a medium expert" to a team of "three experts." (See Figure 2.) Please note that expertise in playing video games does not imply expertise in the scavenger hunt task.

Each team member played the game through an assigned and dedicated laptop. The seating of the three team members was such that they could not look at each other's screen. The human subjects were forbidden from sharing computer screens, note sheets or other such aids. They could only describe their locations, intentions and actions in the game by using verbal communications and the 2-D OFP map. All verbal communications were recorded using TeamSpeak, for later transcription.

Time was taken during a practice session to instruct the players on the key and mouse commands for the game. Players were instructed on how to move their characters, find and crush bottles, query bottle counts, and how to use additional aids that were available to their avatars. After sighting a bottle, a player must move their avatar to within a couple of meters of it in order to crush it and get credit for the crush. Audible and visible feedback was presented to the player in multiple ways.

The virtual OFP environment offered five ways in which a bottle could be detected. These sensing techniques varied in effectiveness according to the terrain, the virtual character, and the scenario. This was one of the randomizing effects of the OFP environment that the human subjects needed to discover, detect and adapt to.

Three experimental conditions were evaluated:

Number of Bottles is Known, in which the subjects knew how many total bottles they were trying to recover; the performance measure was a team score, where each found bottle counted for one point.

Number of Bottles is Unknown, in which the subjects did not know how many bottles were hidden in the search area; the performance measure was a team score, where each found bottle counted for one point.

Bottle Portfolio Condition, where subjects were given a bonus of 93 points if they collected a "portfolio" of 7 bottles, each of a different kind. For this condition, there were 28 hidden bottles, thus enabling the formation of 4 portfolios in the ideal case.

The three experimental conditions were counterbalanced, namely the conditions were executed in varying order among the teams to mitigate order effects. The subjects had 20 minutes to perform the task in each experimental condition.

We created these different experimental conditions in order to observe the communication patterns, their similarities and differences in the different conditions, in particular in the non-portfolio vs. the portfolio conditions. While the focus of our analysis was on teamwork and team processes, we feel that the contrast in utterance frequencies between the portfolio and non-portfolio conditions illustrates the prominence of communications for consensus-building and group decision-making, and their situation within the overall team process.

	Utterances
ID	**Description**
1	Requesting or Communicating Team Members' Location
2	Any Reference to Terrain or Map Features
3	Question / Indication of Bottle Locations
4	Personal Tips
5	Communication for Role Allocation Before Execution
6	Communication for Division of Execution Space Before Execution
7	Communication for Role Allocation During Execution
8	Communication for Division of Execution Space During Execution
9	Bottle Count
10	Object Count
11	Coverage Progress
12	Hypothesis about Bottle Locations
13	System Problems: Keyboard Lockups, Game Window Loses Focus, etc.
14	Personal Discussion
15	Communicating Intent with Respect to Action
16	Communicating Intent with Respect to Changing Location
17	Acknowledgment / Agreement / Answers to Questions
18	Plan or Activity Critiquing / Plan Suggestions
19	Social Encouragement
20	Situation Assessment of Equipment
21	Commands
22	Justifying Behavior, Action, or Plan
23	Considering / Evaluating Alternatives
24	Suggesting Individual or Team Strategies
25	Clarification Question
26	Question about Means of Doing Something
27	Disagreement
28	Permission / Request for Agreement to Do Something
29	Reporting Status / Activity
30	Offer of Help
31	Request for Status Check
32	Discussion of Capability
33	Monitoring Remaining Time in Session
34	Request for Help

Figure 1: The coding scheme that was used to label team communications

Team		# Bottles Recovered / Total # Bottles in Scenario			Portfolio Score	
Number	**Subject Expertise**	**# Bottles Known**	**# Bottles Unknown**	**Portfolio Condition**	**Raw**	**%age**
1	novice, medium, expert	84.62%	78.57%	64.29%	204 / 307 =	66.45%
2	novice, expert, expert	73.33%	54.76%	28.57%	8 / 307 =	2.61%
3	novice, novice, expert	60.53%	66.67%	42.86%	12 / 307 =	3.91%
4	novice, novice, medium	77.27%	35.71%	39.29%	11 / 307 =	3.58%
5	medium, expert, expert	86.36%	73.81%	78.57%	208 / 307 =	67.75%
6	novice, medium, medium	59.52%	52.38%	46.43%	12 / 307 =	3.91%
7	expert, expert, expert	81.82%	76.19%	60.71%	17 / 307 =	5.54%
8	medium, expert, expert	94.00%	80.95%	57.14%	109 / 307 =	35.50%
9	medium, medium, expert	76.60%	80.95%	71.43%	113 / 307 =	36.81%
10	expert, expert, expert	97.78%	64.29%	50.00%	107 / 307 =	34.85%
11	medium, expert, expert	80.85%	90.48%	42.86%	105 / 307 =	34.20%
12	novice, novice, expert	82.93%	83.33%	57.14%	109 / 307 =	35.50%
13	medium, expert, expert	88.24%	69.05%	53.57%	201 / 307 =	65.47%
14	medium, expert, expert	67.44%	73.81%	42.86%	12 / 307 =	3.91%
15	expert, expert, expert	95.35%	90.48%	71.43%	113 / 307 =	36.81%
16	medium, medium, medium	84.62%	76.19%	28.57%	8 / 307 =	2.61%
17	novice, medium, medium	74.42%	71.43%	42.86%	12 / 307 =	3.91%
mean	—	80.33 ± 11.07%	71.71 ± 14.03%	51.68±14.51%	26.08±24.32%	

Figure 2: The performance of the teams in all three test conditions

.

Communication Code	$F_{2,14}$	Sig	Unknown vs. Portfolio	Known vs. Portfolio
(3) reference bottle locations	19.258	p=.0001	p=.001	p=.0001
(8) division of search space	6.437	p=.004	p=.007	N.S.
(12) hypothesis, bottle locations	4.109	p=.026	N.S.	p=.05
(15) action intent	15.021	p=.0001	p=.003	p=.0001
(21) commands	7.979	p=.002	p=.043	p=.004
(22) justification for action/plan	3.516	p=.042	N.S.	N.S.
(25) clarification of question	12.988	p=.001	p=.001	p=.014
(29) reporting status	6.202	p=.005	p=.01	N.S.
(30) offer of help	10.464	p=.0001	p=.001	p=.041
(32) discuss capability	3.630	p=.038	p=.05	N.S.
(34) request help	33.549	p=.0001	p=.0001	p=.0001

Figure 3: Significant Differences in Communications across Conditions
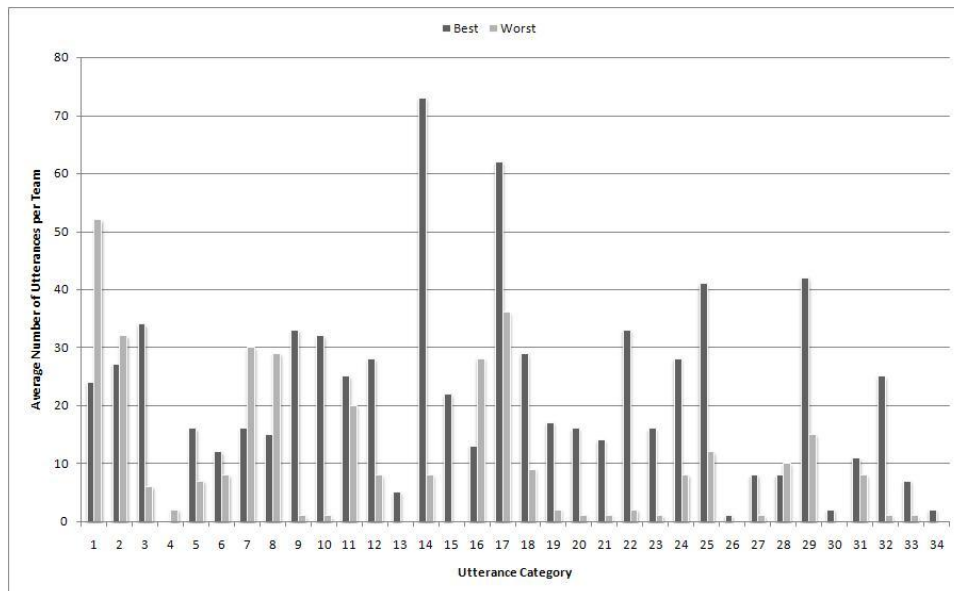

Figure 4: Comparison between the Best and Worst Performing Teams in the **# Bottles Known** Condition

To analyze the coordination demands of the collaborative search task, we recorded all audio communications between team members. The audio files were manually transcribed to produce the logs of utterances that were segmented into conversational moves (Hirokawa, 1983). According to this scheme, a conversational move unit is an uninterrupted utterance of a team member that has a discrete problem solving function. Figure 1 shows the main categories of the codes we used plus their descriptions. The codes represent task-related problem solving and coordination categories that are consistent with current teamwork literature (Fischer et al., 2007). In addition to the communication categories used by other researchers, we included categories of utterances relevant to team self-organization. These aspects were not needed in prior work where, unlike for ad hoc teams, team structure and role allocation were already in place.

RESULTS
Data were collected in a three condition repeated measures experimental design and were analyzed using the SPSS software.

Team Performance
Figure 2 reports the performance of all the teams in our initial set of experiments, measured by percentage of bottles crushed by each team. The subjects participated in an initial practice session during which they learned the user interface (results not shown).

A repeated measures ANOVA shows that team performance in the **# Bottles Unknown** condition was poorer than in the **# Bottles Known** condition ($F1, 15$

= 6.659, p = .02) indicating that knowing the goal (total number of bottles to be found) improved performance. In the **Portfolio** condition, there were greater coordination demands on the team. In this condition, the reward associated with collecting individual bottles is dwarfed by the potential portfolio construction reward. Assembling a portfolio, however, may require players to pass up duplicate bottles, locate and report bottles needed for another's portfolio, and other more complex interdependent behavior. We should therefore expect to see substantial differences in communication patterns between the **Portfolio** and the other two conditions. Because collecting bottles is secondary to assembling portfolios for this group we should also expect lower bottle scores. Figure 2 shows the mean percent bottles found for each of these groups. Differences in performance were found between the three conditions ($F2$, 14 = 36.617, p < .0001) with paired comparisons showing teams in the portfolio condition to find fewer (p < .0001) bottles than in either of the other conditions.

Differences in communications patterns across the three conditions were found for eleven of the thirty-four coded categories using a repeated measures ANOVA with paired comparisons.

Figure 3 shows the significant differences in communications across conditions. (*N.S.* means *non-significant*.) Significance was found between the Portfolio and each of the two bottle search conditions; no significance was found between the two non-portfolio conditions. The differences involving communications such as bottle locations and hypotheses about locations, intended actions and their justifications, reaching consensus on the correct strategy and course of action, and communications characterizing teamwork such as offering and requesting help are of the sort that might be expected to distinguish the relatively independent bottle collection tasks from the high coordination demand required for assembling portfolios.

Team Communication
To assess the communication demands of the collaborative search task, we compiled frequency counts of the different types of team communication.

Five of the 34 utterance categories: (2) reference to terrain or map features, (12) hypotheses about bottle locations, (14) unrelated personal discussion, (15) communication of intended action, and (23) evaluating alternatives, were found to be related to team performance in the non portfolio conditions through a series of step-wise regressions. Teams in the **# Bottles Unknown** condition that exerted the least coordination demand produced the simplest model containing only two independent variables, (14) and (15). The frequency of unrelated discussions, β = .433, t_{14} = 2.205, p = .045, and communicated intent, β = -.433, t_{14} = -2.305, p = .037, significantly predicted the bottles found with a regression explaining about half of the variance, R^2 = .504, $F_{2.14}$ = 7.116, p = .007, in the number of bottles. Our finding that extraneous communication helped — and conveying intent, hurt — in this condition, reinforces our contention that the **# Bottles Unknown** task requires independent search by teammates. Communications such as (15) that encourage unnecessary coordination hurt performance while those that supplant potentially disrupting alternative communications (14) actually helped.

The **# Bottles Known** condition, which stands to benefit from coordination, yields more complex communication patterns. Unlike the **# Bottles Unknown** regression, (15) communication of intent, now contributes positively, β = .323, t_{12} = 2.68, p = .02, to predicting bottles found. Related communications (23) involving sharing hypotheses also contributes positively, β = .502, t_{12} = 4.162, p = .001, as does (14) unrelated discussion, β = .323, t_{12} = 2.68, p = .02, that we have argued may benefit performance by supplanting inappropriate task related communications. Communication code (2) reference to terrain or maps enters the model negatively, β = -.392, t_{12} = -3.32, p = .006. We speculate that this may be due to the fact that players that talked a lot about terrain features may have been disoriented and lost; such players retrieved few bottles. This four variable regression significantly predicted the bottles found and explained almost all of the variance, R^2 = .842, $F_{4,12}$ = 16.003, p = .0001, in the number of bottles.

Teams in the **Portfolio** condition had a substantially different task since their scores were dominated by the portfolio bonus. Nevertheless, we examined the relation between their communication patterns and *number of bottles found* before looking at the more salient relation between communications and scores, where the score includes the bonus for constructing portfolios. Our **Portfolio** data were fit by a five variable regression on *bottles found*. (29) reports of status, β = .687, t_{11} = 6.17, p = .0001, (17) acknowledgment/agreement, β = .504, t_{11} = 4.777, p = .001, and (34) request for help, β = .265, t_{11} = 2.522, p = .028, were all positively related to *bottles found*. (1) Communicating location, β = -.518, t_{11} = -4.851, p = .001, and (31) request for status, β = -.353, t_{11} = -3.353, p = .006, were negatively related. As for the earlier conditions we presume that communications (1) and (31) place cooperative demands on team members that detract from the independent task of accumulating bottles. This five variable regression significantly predicted the *bottles found* and explained almost all of the variance, R^2 = .916, $F_{5,11}$ = 24.029, p = .0001, in the number of bottles.

Regression on *score*, the more salient performance measure in the **Portfolio** condition, revealed a different set of predictors. Data were fit by a four variable regression on team score with all variables providing positive contributions. Communication types: (10) object count, β = .85, t_{12} = 6.827, p = .0001; (19) social encouragement, β = .450, t_{12} = 3.615, p = .003; (11)

reported coverage, $\beta = .315$, $t_{12} = 2.708$, $p = .018$; and (15) communication of intent, $\beta = .278$, $t_{12} = 2.293$, $p = .041$ contribute positively to the model. This four variable regression significantly predicted team scores and explained a substantial amount of the variance, $R^2 = .906$, $F_{4,12} = 29.075$, $p = .0001$, in scores. These differences in predictors for bottle count and scores suggest that different team processes may be involved. To explore the possibility that teams were using different processes to attain these goals we examined the correlation between *number of bottles found* and *team scores*. The positive correlation between bottles found and score, $r_{15} = .73$, $p < .001$, suggests that teams with good searchers were better at both finding more bottles and assembling portfolios. This debunks the notion that the negative relation observed between (1) communicating location, and (31) request for status, in the regression on *bottles found* resulted from allowing team members to search independently, thus accumulating more bottles by freeing them from coordination demands. A more likely explanation is that subjects who talk a lot about position may have been lost and thus do not retrieve many bottles.

To further explore communication patterns in teams we report the results of the highest and lowest performing teams in the non-portfolio condition, **# Bottles Known**. Moreover, we report here the highest and lowest performing team so that the range of their differences may give us some clue as to the types of assistance that would be beneficial.

From the graph of Figure 4, we make a variety of observations. First, we find that the lower-performing team communicates less in all but four communication categories: utterances 1, 2, 7, and 8. This is consistent with social theories of teamwork, where communication is considered one of the most important teamwork dimensions. In addition we see that, consistent with theories of teamwork (Salas & Fiore, 2004) the communication frequency peaks of high performing teams in both conditions were for teamwork behaviors such as inventory monitoring (code 9), object counts (code 10), offering suggestions (code 18), describing intentions (code 15), acknowledgments (code 17), and monitoring the passing of time (code 33). We also note some differences. In the **non-portfolio** condition, offering help (code 30) is not utilized much due possibly to the more loosely coupled nature of the task, whereas in the **portfolio** condition, there is a frequency peak there.

DISCUSSION AND FUTURE WORK
Bottom-up utterance classification is a useful first step for any automation to understand negotiation processes in human teams. It serves as a "semantic grammar" that is quicker to generate and apply to the textual representations of the utterances, than to perform robust syntactic parsing and processing of the same.

The classification process, itself, is labor-intensive however, and is a significant hindrance to experimenting with and evolving classification schemes. Future work will certainly involve seeking methods by which codes can be applied more effortlessly.

While statistical analysis of semantic utterance codes is an important tool for identifying and correlating groups of utterances that impact human team performance, it cannot be used to infer direct explanation for why, nor situate the correlated utterances in an overall process model. This analysis technique can also miss significant transitions that occur in one turn, as was indirectly reported in (Weingart et al., 2007). Future work would seek to develop techniques that not only attempt to remedy this specific shortcoming, but also attempt to extend it to robustly recognize other process model transitions, as well.

The utterance classifications were also useful in the way that they revealed a weakness of the simulation environment. Namely, the environment did not constrain the human subjects to carefully plan their team activity. Indeed, this is evidenced by the need to identify only two utterance codes, 5 and 6, as being applicable before the start of plan execution. The other codes described utterances during execution.

A second weakness of the experiment was that it did not record any means by which team processes could be associated with a temporal sequence. This prevented us from being able to derive a dynamic process model of the consensus building activities. Future work might revisit the encoded transcripts and reconstruct the activity sequences.

CONCLUSIONS
We presented: (a) a characterization of a negotiation process model; (b) a taxonomy of team-oriented classification codes which can be used to identify such process models in transcriptions of ad hoc distributed team activities; (c) an experimental design for correlating the utterance codes with team performance metrics so as to identify team processes and areas where automation could possibly improve human performance, and (d) a statistical analysis of the results from the experiments. Our experiences with this technique showed that while some groups of utterance categories are correlated to the performance of the team in statistically significant ways, the utterances are nonetheless too coarse-grained to index changes in team state, which sometimes occur in one utterance or one utterance turn (e.g., statement and reply).

REFERENCES

Carberry, S. & Lambert, L. (1999). A process model for recognizing communicative acts and modeling negotiation subdialogues. *Computational Linguistics, 25* (1), 1-53.

Fiore, S., Salas, E., Cuevas, H., & Bowers, C. (2003). Distributed coordination space: towards a theory of distributed team performance. *Theoretical Issues in Ergonomic Science, 4*.

Fischer, U., McDonnel, L., & Orasanu, J. (2007). Linguistic correlates of team performance: Toward a tool for monitoring team functioning during space missions. *Aviation, Space, and Environmental Medicine, 78* (5), B87-B95. Section II.

Hirokawa, R. Y. (1983). Group communication and problem-solving effectiveness: an investigation of group phases. *Human Communication Research, 9*, 291-305.

Morimoto, I., Saijo, M., Nohara, K., Takagi, K., Otsuka, H., Suzuki, K., & Okumura, M. (2006). How Do Ordinary Japanese Reach Consensus in Group Decision Making?: Identifying and Analyzing "Naïve Negotiation." *Group Decision and Negotiation, 15* (2), 157-169.

Pascual, R. G., Mills, M. C., & Blendell, C. (1999). Supporting distributed and ad-hoc team interaction. In *People in Control: An International Conference On Human Interfaces in Control Rooms, Cockpits and Command Centres*, number 463 in Conference Publication, 64-71. IEE.

Salas, E. & Fiore, S. (Eds.). (2004). *Team Cognition: Understanding the Factors that Drive Process and Performance*. American Psychological Association.

Weingart, L. R., Brett, J. M., Olekalns, M., & Smith, P. L. (2007). Conflicting social motives in Negotiating groups. *Journal of Personality and Social Psychology, 93*, 994-1010. Draft.