Communication for Conflict Resolution in Multi-Agent Collaborative Planning*

Jennifer Chu-Carroll and Sandra Carberry
Department of Computer and Information Sciences
University of Delaware
Newark, DE 19716, USA
E-mail: {jchu,carberry}@cis.udel.edu

Abstract

Conflict management, communication, and negotiation are important components of collaborative multi-agent activity. Thus, a collaborative agent must be able to handle situations in which conflicts arise and must be capable of negotiating with other agents to reach an agreement. This paper presents a model which 1) captures multi-agent collaboration in a *Propose-Evaluate-Modify* cycle of actions, 2) initiates negotiation with the executing agent to resolve detected conflicts regarding proposed actions and proposed beliefs, 3) selects the focus of the modification process when multiple conflicts arise, and 4) handles the negotiation of proposed domain actions, proposed problem-solving actions, and proposed beliefs in a unified manner.

Introduction

Conflict management, communication, and negotiation are important components of multi-agent activity. Conflict resolution involves communication among agents – communication for the purpose of *squaring away* (Joshi 1982) discrepancies among agents' beliefs. Successful communication requires strategic planning to say the right thing at the right time. In the context of conflict resolution, this means determining whether 1) to select and provide convincing evidence so that another agent will change his beliefs, or 2) to gather information from others so that the conflicting beliefs can be re-evaluated and potentially modified.

In this paper, we present a plan-based model that captures multi-agent collaborative activities in a sequence of *Propose*, *Evaluate*, and *Modify* actions. We describe how our model evaluates an agent's proposals, including both proposals of actions to be performed and proposals of beliefs to be held jointly among the participants, as well as how it attempts to resolve the conflicts detected during the evaluation process. We discuss how our model selects the focus of modification when multiple conflicts arise and how it captures situations involving embedded negotiation during conflict resolution.

*This material is based upon work supported by the National Science Foundation under Grant No. IRI-9122026.

Communication and Collaboration

Sidner formulated an artificial language for modeling collaborative discourse using proposal/acceptance and proposal/rejection sequences (Sidner 1994). While Sidner's work is descriptive, our research is prescriptive in that we have identified appropriate response generation strategies for agents involved in collaborative interactions. Our research focuses on a particular kind of collaborative activity in which two participants,[1] an executing agent (EA) and a consulting agent (CA), work together to construct a plan for achieving EA's domain goal. The agents involved in such collaboration bring to the plan construction task different knowledge about the domain and the desirable characteristics of the domain plan. In particular, EA has knowledge about his particular circumstances and preferences that are potential influencers (Bratman 1990) of the domain plan being developed. Thus CA,[2] as a collaborator on EA's domain plan, must evaluate proposals with respect to her beliefs about EA's preferences; consequently, CA will generally attempt to modify unaccepted proposals to a form acceptable to both agents, rather than rejecting EA's proposals entirely. Thus we capture collaborative activities in a *Propose-Evaluate-Modify* cycle of actions (Chu-Carroll & Carberry 1994). This model is a recursive one: the *Modify* action in itself contains a full collaborative cycle – an agent's proposal of a modification, the other agent's evaluation of the proposal, and potential modification of the proposed modification!

In this *Propose-Evaluate-Modify* cycle of collaboration, communication may occur at all three stages. At the *proposal* stage, communication is necessary to con-

[1] Although our work has focused on conflict management, communication, and negotiation between two agents, we believe that the overall mechanism we have identified can be extended to model activities among multiple agents.

[2] Since our system plays the role of a consulting agent, we will use the terms CA and system interchangeably in the rest of this paper.

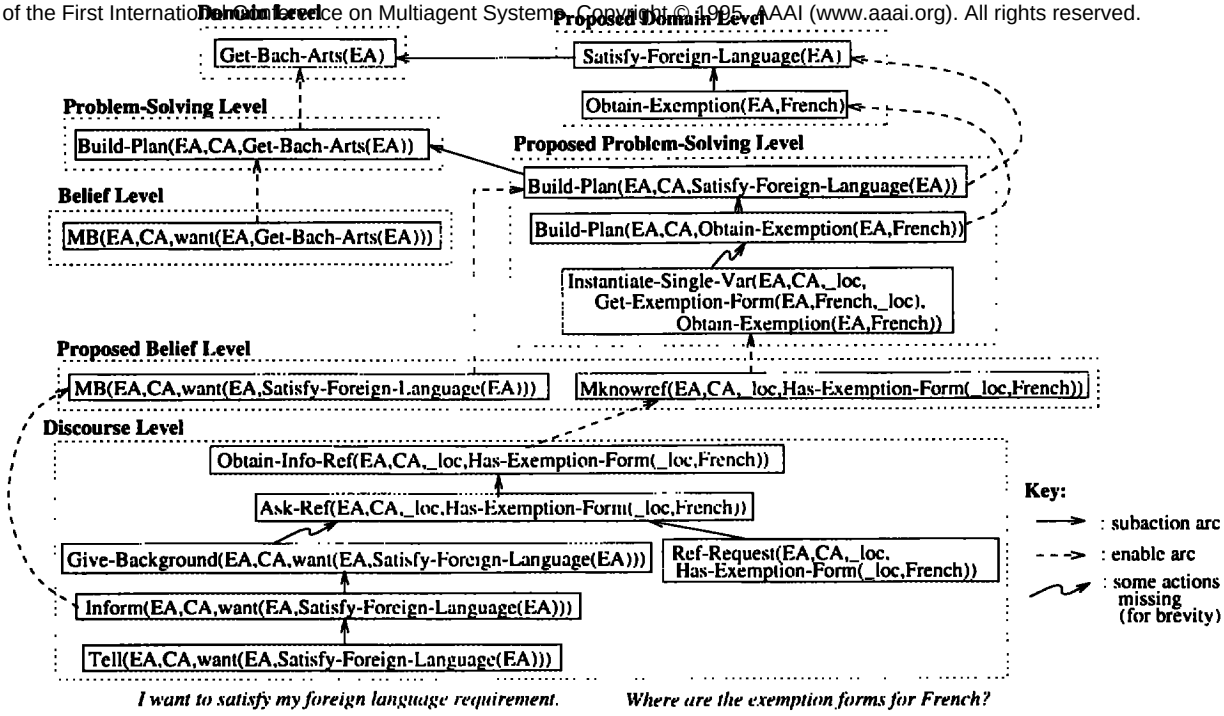

Figure 1: Dialogue Model for Utterances (1) and (2)

vey proposed additions to the *shared plan* (Grosz & Sidner 1990; Allen 1991) being constructed. At the *evaluation* stage, communication is necessary when an agent cannot determine whether or not to accept a proposal and attempts to share information with the other agent so that each agent can knowledgably re-evaluate the proposal. The generation of such *information-sharing* subdialogues is discussed in (Chu-Carroll & Carberry 1995a). At the *modification* stage, communication is an integral part of the conflict resolution process since prior to modifying a proposal, a collaborative agent must first obtain the other agent's consent regarding the modification. If the other agent does not immediately agree, an extended negotiation subdialogue can ensue to resolve the detected conflict.

In the following sections, we present our plan-based model for communication in collaborative task-oriented dialogues. We focus on situations in which conflicts arise among agents, and we identify appropriate communication, negotiation, and conflict resolution strategies. The examples in this paper are taken from a university course advisement domain, although the model can easily be applied to other domains.

Evaluating Proposals

Agents collaborate on more than just which domain actions to include in a plan. For example, the agents might negotiate the strategies they will use to construct the domain plan. Agents also collaborate on their beliefs, forming a set of mutual beliefs that are relevant to the task at hand. Thus we use an enhanced version of the dialogue model presented in (Lambert & Carberry 1991) to capture the current intentions of the dialogue participants. The dialogue model has four levels: the *domain* level which consists of the domain plan being constructed for EA's later execution, the *problem-solving* level which contains the actions being performed to construct the domain plan, the *belief* level which consists of the mutual beliefs pursued during the planning process in order to further the problem-solving intentions, and the *discourse* level which contains the communicative actions initiated to achieve the mutual beliefs. Actions at the discourse level can contribute to other discourse actions and also enable mutual beliefs. Mutual beliefs can support other beliefs and also enable problem-solving actions. Problem-solving actions can be part of other problem-solving actions and also enable domain actions.

In order to properly model collaboration, proposals must be distinguished from the shared beliefs and plans that have been agreed upon by the participants. Therefore, we separate the dialogue model into an *existing model* and a set of *proposed additions*, following Allen who differentiated among private, proposed, and shared modules of beliefs (Allen 1991). For example, suppose that earlier dialogue suggests that EA has the goal of getting a Bachelor of Arts degree (*Get-Bach-Arts(EA)*). Figure 1 shows the dialogue model that would be built after the following utterances:

(1) EA: I want to satisfy my foreign language requirement.

(2) Where are the exemption forms for French?

As the dialogue model indicates, utterance (2) proposes, at the belief level, that both agents come to know the referent of the location of the exemption forms for French. This belief is a prerequisite for the problem-solving action of instantiating the variable _loc, which is part of building a plan for obtaining an exemption for French. Having such a plan is again a prerequisite for executing the domain action of obtaining an exemption for French as part of satisfying the foreign language requirement. Notice that these newly inferred actions and suggested mutual beliefs are treated as a set of proposed additions to the existing model consisting of the shared plan and shared beliefs already established between the agents.

Instead of slavishly[3] responding to questions, a collaborative agent presented with a proposal (such as the one implicitly conveyed by utterances (1) and (2) above) needs to first decide whether she can accept the proposal as a valid and reasonably efficient contribution to achieving the high-level goal. Our *evaluator* evaluates the proposed domain actions, proposed problem-solving actions, and proposed mutual beliefs in that order. For the domain and problem-solving levels, the evaluator performs a top-down analysis to detect invalid as well as suboptimal plans. A plan is considered invalid if either an action in the plan cannot be performed (an action is *infeasible*) or a child action does not contribute to its parent action as intended (the plan is *ill-formed* (Pollack 1986)). A plan is considered suboptimal if there exists a better way to perform the desired action. We evaluate the optimality of a plan with respect to EA's preferences (Elzer, Chu-Carroll, & Carberry 1994). The evaluation of actions is a top-down process which terminates as soon as a conflict regarding an action or parent-child relationship is detected, since conflicts about child actions are irrelevant if the parent action is rejected.

The belief level of our dialogue model consists of one or more belief trees where the belief represented by a child node is intended to support that represented by its parent. The evaluation of such belief trees is a bottom-up process because acceptance of a proposed belief, _bel, may be influenced by acceptance of the beliefs represented by its children in the belief tree. A simplified version of Galliers' belief revision mechanism (Galliers 1992; Logan *et al.* 1994) is used to determine whether or not a proposed belief or evidential relationship should be accepted. For each such belief _bel, the belief revision mechanism constructs an evidence set which contains evidence from CA's existing beliefs that supports or attacks _bel and the evidence accepted by CA that was proposed by EA as support for _bel. It then determines whether or not _bel warrants CA's acceptance based on this evidence set. The process continues until the beliefs at the root nodes of the proposed belief trees (henceforth referred to as *top-level proposed beliefs*) are evaluated. Conflict resolution and negotiation are necessary only if the top-level proposed beliefs are not accepted since if the agents agree on a particular belief relevant to the domain plan being constructed, it is irrelevant whether they both agree on all the evidence for that belief (Young, Moore, & Pollack 1994).

[3]Grosz and Sidner (1990) argued against a master-slave relationship among collaborative dialogue participants.

Examples of Conflict Detection

To illustrate the evaluation process, we return to the example depicted in the dialogue model in Figure 1. The system evaluates the proposal beginning with the proposed domain actions. Since the system believes that *Satisfy-Foreign-Language(EA)* contributes to its parent action *Get-Bach-Arts(EA)*, that *Satisfy-Foreign-Language(EA)* is feasible, and that there is no better alternative to *Satisfy-Foreign-Language(EA)*, the system evaluates its child action *Obtain-Exemption(EA,French)*. Again, the system believes that *Obtain-Exemption(EA,French)* contributes to *Satisfy-Foreign-Language(EA)*. However, the system's recipe library[4] indicates that an applicability condition[5] of *Obtain-Exemption* is that the agent is not a native North American, but the system believes that EA *is* a native North American. Thus, EA's proposal is not accepted because the system believes that *Obtain-Exemption(EA,French)* is infeasible.

In order to demonstrate the differences between the process of evaluating actions and that of evaluating beliefs, consider the following exchange:

(3) CA: Dr. Smith is going on sabbatical next semester.

(4) EA: Dr. Smith is not going on sabbatical next semester.

(5) He is teaching AI next semester.

The dialogue model for utterances (4) and (5), part of which is shown in Figure 2, suggests that EA is proposing three mutual beliefs: 1) Dr. Smith is not going on sabbatical next semester, 2) Dr. Smith is teaching AI next semester, and 3) Dr. Smith teaching AI provides support for the belief that he is not going on sabbatical.

The process for evaluating proposed beliefs starts at the leaf node of the proposed belief tree, *Teaches(Smith,AI,next semester)*. The system

[4]A recipe (Pollack 1986) is a template for performing actions. It contains the preconditions of an action, the subactions comprising the body of an action, etc.

[5]Applicability conditions are conditions that must already be satisfied in order for an action to be reasonable to pursue, whereas an agent can try to achieve unsatisfied preconditions.

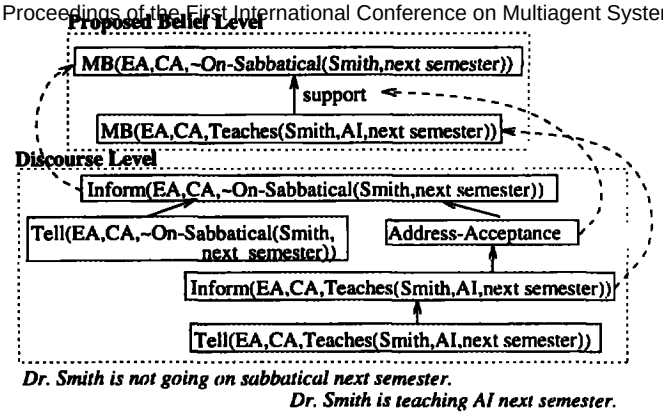

Figure 2: Belief and Discourse Levels for (4) and (5)

Action:	Correct-Node(_s1, _s2, _proposed)
Type:	Decomposition
Appl Cond:	believe(_s1, ¬acceptable(_node))
	believe(_s2, acceptable(_node))
Const:	error-in-plan(_node,_proposed)
Body:	Modify-Node(_s1,_s2,_proposed,_node)
	Insert-Correction(_s1, _s2, _proposed)
Goal:	acceptable(_proposed)
Action:	Modify-Node(_s1,_s2,_proposed,_node)
Type:	Specialization
Appl Cond:	believe(_s1, ¬acceptable(_node))
Precond:	believe(_s2, ¬acceptable(_node))
Body:	Remove-Node(_s1, _s2, _proposed, _node)
	Alter-Node(_s1, _s2, _proposed, _node)
Goal:	modified(_proposed)

Figure 3: The *Correct-Node* and *Modify-Node* Recipes

searches for its evidence pertaining to the proposed belief and invokes the belief revision mechanism to determine the acceptance of the belief. Since the system strongly[6] believes that Dr. Brown is teaching AI next semester, and that only one person teaches a course, the proposed belief that Dr. Smith is teaching AI next semester is rejected. However, the process for belief evaluation does not terminate when a proposed belief is rejected, because the acceptance of a parent belief is not necessarily contingent on the acceptance of its children. The system believes that teaching a course implies a faculty member is not on sabbatical; therefore it accepts the proposed evidential relationship. However, the system does not accept the top-level proposed belief *¬On-Sabbatical(Smith, next semester)*, because the only evidence provided by EA was an implication whose antecedent was not accepted by the system. Since the top-level proposed belief is rejected by the system, EA's proposal of mutual beliefs is rejected.

Resolving Conflicts – A Negotiated Process

Once CA detects a relevant conflict, she must notify EA of the conflict and attempt to resolve it – to do otherwise is to fail in her responsibilities as a collaborative participant (Walker 1992). This results in *collaborative negotiation*. Such negotiation differs from other kinds of negotiation, such as labor negotiation (Sycara 1989), in that the participants are not trying to enforce their views on one another or to maximize their own benefits, but rather are trying to share their individual knowledge and beliefs in order to determine what *really* is best (Chu-Carroll & Carberry 1995b).

[6]The strength of an agent's beliefs is modeled with *endorsements* (Cawsey *et al.* 1993), which are explicit records of factors that affect one's certainty in a hypothesis (Cohen 1985). Our endorsements include the semantics of the utterance used to convey a belief, the level of expertise of the agent conveying the belief, stereotypical knowledge, etc.

The *Modify-Proposal* problem-solving action is invoked when an agent does not accept a proposal and attempts to modify it. Its goal is to modify the proposal to a form that will potentially be accepted by both agents. *Modify-Proposal* has four specializations: 1) *Correct-Node*, for when a proposed action is infeasible or when a proposed belief is not accepted; 2) *Correct-Relation*, for when the proposal is ill-formed or when the evidential relationship between two proposed beliefs is not accepted; 3) *Improve-Action*, for when a better generic action is found; and 4) *Improve-Parameter*, for when a better instantiation of a parameter is found. Each specialization eventually decomposes into some primitive action which modifies the proposal. However, an agent will be considered uncooperative if he modifies an agent's proposal without the agent's consent; therefore, all four specializations share a common precondition – both agents must come to an agreement that the original proposal is faulty before any modification can take place. It is the attempt to satisfy this precondition that leads to the generation of natural language utterances to square away the conflict in the agents' beliefs.

Figure 3 shows problem-solving recipes for *Correct-Node* and one of its subactions, *Modify-Node*. The applicability conditions of *Correct-Node* specify that the action can only be performed when _s1 believes that _node is not acceptable while _s2 believes that it is (when _s1 and _s2 disagree about the feasibility of _node when it is instantiated as an action, or about the truth of _node when it is instantiated as a belief). The applicability condition and precondition of *Modify-Node*, however, show that the action can only be performed if both _s1 and _s2 believe that _node is not acceptable — that is, the conflict between _s1 and _s2 must have been resolved. The attempt to satisfy this precondition causes CA (the system) to post as a mutual belief to be achieved the belief that _node is not acceptable, leading CA to invoke discourse actions

to modify EA's beliefs. If EA accepts CA's beliefs, thus satisfying the precondition of *Modify-Node*, the original dialogue model can be modified; however, if EA does not accept CA's beliefs, he may try to modify CA's suggested modification of his original proposal, resulting in a recursive process.

Selecting the Focus of Modification

When the *Modify-Proposal* action invokes one of its subactions, it must determine the focus of modification – i.e., the aspect of the proposal that an agent will address in her pursuit of conflict resolution. For instance, in the case of invoking *Correct-Node*, it must determine how the parameter _node in *Correct-Node* (Figure 3) should be instantiated. If the reason for proposal rejection occurs at the domain or problem-solving level, the focus of modification is the action or parent-child relationship about which the agents disagree. However, if the reason for rejection occurs at the belief level, the process for selecting the focus of modification is more complex, since there may be multiple proposed beliefs or evidential relationships that were rejected during the evaluation process.

A proposal at the belief level is rejected if at least one of the top-level proposed beliefs is rejected. If CA's evidence against a top-level belief is so strong that the belief is rejected even though all of its descendents are accepted, the focus of modification should be the rejected belief. On the other hand, a top-level proposed belief may fail to be accepted because evidence provided by EA was not accepted. In these cases, CA has to select its focus of modification from among the top-level belief and its supporting evidence (its descendents in the dialogue model).

Collaborative participants are expected to engage in effective and efficient dialogues, and not to argue for the sake of arguing. Thus CA should address the rejected belief that it predicts will most quickly resolve the conflict regarding the top-level proposed belief. Therefore, for each rejected top-level belief, our process for selecting the focus of modification involves two steps: identifying a candidate foci tree from the proposed belief tree, and selecting a focus from the candidate foci tree using the heuristic rule "attack the belief(s) which will most likely resolve the conflict about the top-level belief." The candidate foci tree contains the pieces of evidence in the proposed belief tree which, if disbelieved by EA, might change EA's view of the rejected top-level proposed belief. It is identified by performing a depth-first search on the proposed belief tree. When a belief node is visited, both the belief and the evidential relationship between the belief and its parent are examined. If both the belief and relationship were accepted by CA during the evaluation stage, the search on the current branch terminates, since once CA accepts a belief, it is irrelevant whether it accepts EA's evidence for that belief. Otherwise, this piece of evidence is included in the candidate foci tree and

Select-Focus-Modification(*_bel*)

1. If *_bel* is a leaf node in the candidate foci tree,
   - 1.1 If **Predict**(_bel, system's evidence against _bel) = ¬_bel,
     Return _bel
   - 1.2 Else return nil
2. Select focus for each of *_bel*'s children, $\_bel_1, \ldots, \_bel_n$, in the candidate foci tree:
   - 2.1 If supports($\_bel_i$, *_bel*) is accepted but $\_bel_i$ is not, **Select-Focus-Modification**($\_bel_i$).
   - 2.2 Else if $\_bel_i$ is accepted but supports($\_bel_i$, *_bel*) is not, **Select-Focus-Modification**(supports($\_bel_i$, *_bel*)).
   - 2.3 Else **Select-Focus-Modification**($\_bel_i$) ∪ **Select-Focus-Modification**(supports($\_bel_i$, *_bel*)).
3. Choose between attacking the proposed evidence for *_bel* and attacking *_bel* itself:
   - 3.1 If **Predict**(*_bel*, {evidence against children of *_bel*})=¬*_bel* (i.e., evidence against children of *_bel* causes EA to reject *_bel*), return the list of children along with the evidence against them.
   - 3.2 Else if **Predict**(*_bel*, {evidence against *_bel*})=¬*_bel* (i.e., evidence against *_bel* itself causes EA to reject *_bel*), return *_bel* and all evidence against it.
   - 3.3 Else if **Predict**(*_bel*, {evidence attacking *_bel* and its children})=¬*_bel*, return *_bel*, its children, and all evidence against them.
   - 3.4 Else return nil.

Figure 4: Selecting the Focus of Modification

CA continues to search through the beliefs proposed to support the rejected belief and/or relationship[7]

Once CA identified a candidate foci tree, she should select the focus of modification based on the likelihood of each choice changing EA's beliefs. An outline of our algorithm for selecting the focus of modification, **Select-Focus-Modification**, is shown in Figure 4.[8] Given a rejected belief (*_bel*), **Select-Focus-Modification** determines whether to attack its supporting evidence (its children) separately (step 3.1), thereby eliminating EA's reasons for believing *_bel*, or to attack *_bel* itself (step 3.2). However, in evaluating the effectiveness of attacking each of *_bel*'s support, CA must determine whether it is more effective to attack the child itself or its support, thus resulting in a recursive process (step 2). Notice that steps 1 and 3 of the algorithm invoke a function, **Predict**, that predicts whether or not EA's belief in a proposition will change after a set of evidence is presented to him. **Predict** makes use of the belief revision mechanism (Galliers 1992) discussed earlier to predict EA's view on *_bel*

[7] It is possible to provide support/attack for an evidential relationship.

[8] For a full account of the algorithm, see (Chu-Carroll & Carberry 1995b).

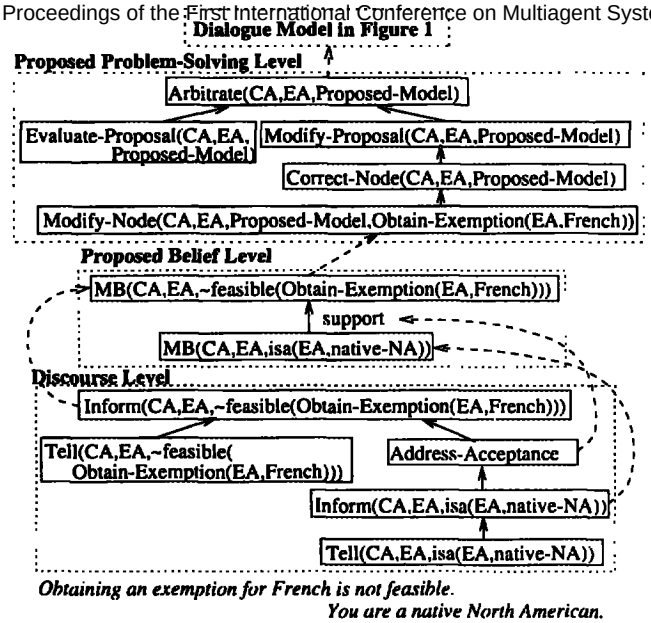

Figure 5: Responding to Implicitly-Conveyed Conflicts

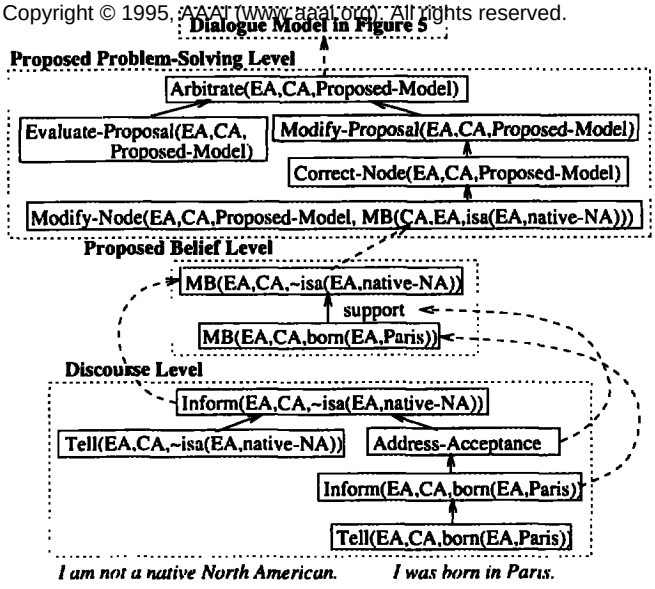

Figure 6: Dialogue model for Utterances (8) and (9)

based on the system's knowledge of EA's existing beliefs and the evidence to be presented to him (Logan *et al.* 1994).

An Example of Conflict Negotiation

In the example given in utterances (1) and (2), whose dialogue model was shown in Figure 1, the evaluator rejected the proposed domain action *Obtain-Exemption(EA, French)* because it believes that the action is infeasible. Thus the *Modify-Proposal* action is invoked to resolve the conflict. Since an infeasible action is detected, *Correct-Node* (Figure 3) is selected as the specialization of *Modify-Proposal*. Figure 5 illustrates the dialogue model for the modification process. In order to satisfy the precondition of *Modify-Node*, the system attempts to establish the mutual belief that the *Obtain-Exemption* action is infeasible and to support it with the mutual belief that EA is a North American. The system invokes *Inform* discourse actions to convey these beliefs, thus generating the following two utterances:

(6) CA: Obtaining an exemption for French is not feasible.

(7) You are a native North American.

Notice that in Figure 5, the actions for modifying EA's proposal operate on the entire dialogue model in Figure 1, and therefore are represented as meta-level problem-solving actions. The original dialogue model will be returned to when the meta-level actions have been completed or abandoned. If EA accepts the system's proposed beliefs, thus satisfying the precondition that the conflict be resolved, *Modify-Node* can be performed and changes made to the original proposal. In this case, the *Obtain-Exemption* action along with all actions that contribute to it (including the proposal that the agents mutually know where exemption forms are located) will be removed, and the remaining actions in the proposal will be incorporated into the shared plan. Thus our model is capable of capturing a common feature of naturally occurring collaborative planning dialogues — questions may never be answered because they become superfluous to a correct means of achieving an agent's goals.

On the other hand, consider the case in which EA responds to the system's utterances as follows:

(8) EA: I am not a native North American.

(9) I was born in Paris.

Utterances (8) and (9) would be interpreted as an indication that EA does not accept the previously proposed belief *MB(CA,EA,isa(EA, native-NA))*. Figure 6 shows the dialogue model that would be constructed after these utterances.

Suppose the system believes that EA has been living in North America since age five and that for the purpose of foreign language exemption, one is considered a native North American if he has been in North America since age six. We demonstrate how our system will evaluate the proposed beliefs. This evaluation will begin at the leaf node of the belief tree. The system will accept the belief that EA was born in Paris; however, it will not accept the proposed evidential relationship "*since EA was born in Paris, he is not a native North American.*" Consequently, the system will reject EA's proposed evidence for $\neg$*isa(EA, native-NA)*, thus retaining its prior belief that EA is a native North American. Since EA's top-level proposed belief is rejected,

the modification process will be invoked.

In modifying the proposal, the focus of modification must be identified in order to determine how the subactions of *Modify-Proposal* should be selected and instantiated. The candidate foci tree is identical to the proposed belief tree represented at the belief level in Figure 6, since both the top-level proposed belief and the evidence proposed to support it were rejected. This indicates that the focus of modification could be either *¬isa(EA, native-NA)* or the *supports* relationship (since *born(EA, Paris)* was accepted during the evaluation process). The algorithm for selecting the focus of modification (Figure 4) will then be invoked with *_bel* bound to *¬isa(EA,native-NA)*. Evidence against the *supports* relationship will be gathered (step 2.2). If this evidence is predicted to be strong enough to convince EA that the relationship is faulty, then the *supports* relationship will be selected as the focus of modification (step 3.1), and *Correct-Relation* will be chosen as the specialization of *Modify-Proposal*. Again, in order to satisfy the precondition that conflicts between agents be resolved before any modification can take place, *¬supports(born(EA,Paris),¬isa(EA,native-NA))* will be posted as a mutual belief to be achieved. Discourse actions will then be adopted to convey the belief as well as its supporting evidence to EA,[9] leading to the generation of the following utterances:

(10) CA: Having been born in Paris does not support you not being considered a native North American.

(11) For the purpose of foreign language exemption, if you have been in North America since age six, you are considered a native North American whether or not you were born outside of North America.

Now if EA expresses acceptance of the proposed mutual beliefs, therefore satisfying the precondition of *Correct-Relation*, the proposed evidential relationship at the belief level in Figure 6 will be removed. This leaves EA with no more reason to believe that he is not a North American;[10] thus the *Modify* actions in Figure 6 are abandoned and the dialogue model in Figure 5 is returned to. EA accepts the proposed mutual belief that he is a native North American and thus the mutual belief that he is not eligible for a foreign language exemption, the latter of which satisfies the precondition of *Modify-Node*. The system will therefore execute *Modify-Node* and remove the *Obtain-Exemption* action and all actions that contribute to the proposal of *Obtain-Exemption* in Figure 1.

[9] In cases where multiple pieces of evidence are available to justify a belief, our model is capable of selecting and presenting a subset of the evidence which it predicts will successfully convince EA to accept the belief. This selection process is described in (Chu-Carroll & Carberry 1995b).

[10] If EA had other reasons to believe that he is a North American, he, as a collaborative agent, would provide them as further support for this belief, instead of merely accepting (10) and (11).

Related Work

Researchers have developed conflict resolution strategies for various types of conflicts, including resolving conflicting goals between non-fully cooperative agents (Sycara 1989), resolving conflicts in resource allocation among cooperative agents (Conry, Meyer, & Lesser 1988), resolving conflicts between coordinating subproblems in distributed problem-solving (Klein 1991; Lander & Lesser 1992), maintaining consistency among beliefs of computational agents (Huhns & Bridgeland 1991), etc. In addition, Rosenschein and Zlotkin (1994) proposed a general theory characterizing the relationship between domains and appropriate negotiation mechanisms. These research efforts have focused on different aspects of conflict resolution from ours. Grosz, Sidner and Kraus developed a Shared-Plan model for collaborative activities which captures the agents' intentions in developing a plan to be carried out jointly by the agents (Grosz & Sidner 1990; Grosz & Kraus 1993). However, in their model the agents will avoid adopting conflicting intentions, instead of trying to resolve them.

Cawsey et al., in developing their automated librarian (Cawsey *et al.* 1993; Logan *et al.* 1994), introduced the idea of utilizing a belief revision mechanism (Galliers 1992) to predict whether a given set of evidence is sufficient to change a user's existing belief. They argued that in the information retrieval dialogues they analyzed, "in no cases does negotiation extend beyond the initial belief conflict and its immediate resolution." (Logan *et al.* 1994, page 141); thus they do not provide a mechanism for collaborative negotiation. On the other hand, our analysis of naturally-occuring consultation dialogues shows that conflict resolution does extend beyond a single exchange of conflicting beliefs; therefore our model employs a recursive model for multi-agent collaboration which captures extended negotiation. Furthermore, their system deals with one conflict at a time, while our model is capable of selecting a focus in its pursuit of conflict resolution when multiple conflicts arise.

Conclusion

In order for heterogeneous agents to successfully collaborate in achieving a shared goal, they must be capable of dealing with situations in which conflicts occur. Conflict resolution involves communication among agents for the purpose of squaring away discrepancies among their beliefs. This paper has presented a plan-based model that specifies how the system (as consulting agent) should detect and attempt to resolve conflicts with the executing agent. The model provides an overall mechanism for modeling such activity by capturing multi-agent collaboration in a cycle of *Propose-Evaluate-Modify* actions. It specifies how

the system should engage in negotiation with the executing agent regarding the validity and optimality of the proposed actions and the truth of the proposed beliefs. Furthermore, our mechanism for conflict resolution identifies, in cases where multiple conflicting beliefs exist, the focus of modification based on its predicted success in resolving the conflict about the top-level belief, thus leading the system to engage in effective and efficient dialogues. Finally, in cases where the system's first attempt to resolve a conflict fails, our model applies the same conflict detection and resolution method to the executing agent's new proposal, resulting in an embedded negotiation dialogue captured by recursive meta-plans.

Acknowledgments

The authors would like to thank Candy Sidner, Stephanie Elzer, and Kathy McCoy for their helpful comments and discussions.

References

Allen, J. 1991. Discourse structure in the TRAINS project. In *Darpa Speech and Natural Language Workshop.*

Bratman, M. 1990. What is intention? In Cohen; Morgan; and Pollack., eds., *Intentions in Communication.* MIT Press. chapter 2, 15–31.

Cawsey, A.; Galliers, J.; Logan, B.; Reece, S.; and Sparck Jones, K. 1993. Revising beliefs and intentions: A unified framework for agent interaction. In *The Ninth Biennial Conference of the Society for the Study of Artificial Intelligence and Simulation of Behaviour*, 130–139.

Chu-Carroll, J., and Carberry, S. 1994. A plan-based model for response generation in collaborative task-oriented dialogues. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, 799 805.

Chu-Carroll, J., and Carberry. S. 1995a. Generating information-sharing subdialogues in expert-user consultation. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence.* To appear.

Chu-Carroll, J., and Carberry, S. 1995b. Response generation in collaborative negotiation. In *Proceedings of the 33rd Annual Meeting of the Association for Computational Linguistics.* To appear.

Cohen, P. R. 1985. *Heuristic Reasoning about Uncertainty: An Artificial Intelligence Approach.* Pitman Publishing Company.

Conry, S. E.; Meyer, R. A.; and Lesser, V. R. 1988. Multistage negotiation in distributed planning. In Bond, A. H., and Gasser, L., eds., *Readings in Distributed Artificial Intelligence.* Morgan Kaufmann Publishers, Inc. 367–383.

Elzer, S.; Chu-Carroll, J.; and Carberry, S. 1994. Recognizing and utilizing user preferences in collaborative consultation dialogues. In *Proceedings of the Fourth International Conference on User Modeling*, 19- 24.

Galliers, J. R. 1992. Autonomous belief revision and communication. In Gardenfors., ed., *Belief Revision.* Cambridge University Press.

Grosz, B.. and Kraus, S. 1993. Collaborative plans for group activities. In *Proceedings of the 13th International Joint Conference on Artificial Intelligence.*

Grosz, B. J., and Sidner, C. L. 1990. Plans for discourse. In Cohen; Morgan; and Pollack., eds., *Intentions in Communication.* MIT Press. 417–444.

Huhns, M. N., and Bridgeland, D. M. 1991. Multiagent truth maintenance. *IEEE Transactions on Systems, Man, and Cybernetics* 21(6):1437 1445.

Joshi, A. K. 1982. Mutual beliefs in question-answer systems. In Smith, N., ed., *Mutual Knowledge.* Academic Press. chapter 4. 181 197.

Klein, M. 1991. Supporting conflict resolution in cooperative design systems. *IEEE Transactions on Systems, Man, and Cybernetics* 21(6):1379–1389.

Lambert, L., and Carberry, S. 1991. A tripartite plan-based model of dialogue. In *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*, 47–54.

Lander, S. E., and Lesser, V. R. 1992. Negotiated search: Cooperative search among heterogeneous expert agents. In *AAAI-92 Workshop: Cooperation Among Heterogeneous Intelligent Systems*, 74–83.

Logan, B.; Reece, S.; Cawsey, A.; Galliers, J.; and Sparck Jones, K. 1994. Belief revision and dialogue management in information retrieval. Technical Report 339, Univ. of Cambridge, Computer Laboratory.

Pollack, M. E. 1986. A model of plan inference that distinguishes between the beliefs of actors and observers. In *Proceedings of the 24th Annual Meeting of the Association for Computational Linguistics*, 207-214.

Rosenschein, J. S., and Zlotkin, G. 1994. *Rules of Encounter Designing Conventions for Automated Negotiation among Computers.* MIT Press.

Sidner, C. L. 1994. An artificial discourse language for collaborative negotiation. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, 814–819.

Sycara, K. 1989. Argumentation: Planning other agents' plans. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, 517 523.

Walker, M. A. 1992. Redundancy in collaborative dialogue. In *Proceedings of the 15th International Conference on Computational Linguistics*, 345–351.

Young, R. M.; Moore, J. D.; and Pollack, M. E. 1994. Towards a principled representation of discourse plans. In *Proceedings of the Sixteenth Annual Meeting of the Cognitive Science Society*, 946- 951.