

Progressive Negotiation for Resolving Conflicts among Distributed Heterogeneous Cooperating Agents

Taha Khedro
Department of Civil Engineering
Stanford University
Stanford, CA 94305-4020
khedro@cive.stanford.edu

Michael R. Genesereth
Department of Computer Science
Stanford University
Stanford, CA 94305-2140
genesereth@cs.stanford.edu

Abstract

Progressive negotiation is a strategy for resolving conflicts among distributed heterogeneous cooperating agents. This strategy aims at minimizing backtracking to previous solutions and provably ensures consistency of agents' distributed solutions and convergence on a globally-satisfiable solution. The progressive negotiation strategy is enforced by a task-independent agent called Facilitator, which coordinates and controls the interaction of cooperating agents. The interaction of cooperating agents includes the communication of messages, the identification of conflicts, and the negotiation of conflicts as a way to resolve them. In this paper, we formally present our conceptualization of cooperating agents and their interaction via the facilitator. We next discuss the conflict types identified by agents and then present the progressive negotiation strategy for resolving conflicts. We then present two theorems that discuss the consistency and convergence of distributed solutions ensured by the strategy. Finally, we conclude with a summary of this paper and remarks about the strategy.

Introduction

The topic of negotiation has been a subject of central interest in Distributed Artificial Intelligence (DAI) (Zlotkin & Rosenschein 1993). The word has been used in a variety of ways although it generally refers to communication mechanisms that improve coordination (Kuwabara & Lesser 1989; Conry, Meyer, & Lesser 1988). Negotiation procedures have included the exchange of partial global plans (Durfee 1988), the communication of information intended to alter other agents' goals (Sycara 1989), and the use of incremental suggestions leading to joint plans of action (Kraus & Wilkenfield 1991). In this paper, we propose a negotiation strategy called Progressive Negotiation for resolving conflicts among distributed heterogeneous agents. The strategy aims at minimizing backtracking to previous solutions while agents are cooperating to reach a globally-consistent satisfiable solution. We assume that cooperating agents have disparate knowledge and interact via a task-independent agent called Facilitator by sending and receiving messages, which are assertions and

retractions of predicate logic sentences (Genesereth 1992). Each agent has a theory, which involves a vocabulary of predicate symbols, function symbols, and constant symbols, a set of predicate-logic axioms expressing the agent's task-specific knowledge, and another set of predicate-logic axioms expressing the agent's criteria constraints, which can be relaxed. Because of the nature of the tasks performed by cooperating agents, their theories overlap and subsets of the vocabularies are shared among them. Also, cooperating agents are allowed to have part of their vocabulary not shared with other agents. In addition, agents are allowed to have vocabularies that are related by a set of predicate logic axioms, provided in the facilitator.

In this paper, we first formally describe cooperating agents and their interaction via the facilitator. We then formalize conflict types and introduce the strategy of progressive negotiation for resolving conflicts. Then, we present theorems that discuss the consistency of distributed solutions and the solution convergence of the progressive negotiation strategy. Finally, the paper concludes with a summary and a few remarks about the strategy.

Cooperating Agents

We consider that a cooperating agent α has knowledge K^α as a set of predicate logic axioms, a set of criteria constraints C^α as predicate logic axioms, and a database D^α as a set of ground predicate logic atoms, all expressed over a vocabulary consisting a set of predicate, function, and constant symbols, X^α . A cooperating agent α has authority to make a final decision over a vocabulary $Y^\alpha \subseteq X^\alpha$. Final decisions are those decisions that conclude a disagreement between agents. The goal of every cooperating agent α is to find a complete local solution G^α that is consistent with both its knowledge K^α and constraints C^α . Formally, this can be expressed as follows:

$$G^\alpha \cup K^\alpha \cup C^\alpha \text{ is consistent.} \quad (1)$$

At time t , agent α maintains a partial local solution, which consists of a set of predicate-logic atoms D_t^α . In the process of finding a solution, agent α generates a set of assertions and retractions of atoms V_t^α that is consistent with its knowledge K_t^α and constraints C_t^α , and updates its partial local solution to D_t^α . Agent α also updates its solution when it receives a set of messages reflecting assertions and retractions of predicate logic sentences. Formally the solution update step can be expressed as follows:

$$\begin{aligned} \forall \text{ assertion}(v) \in V_t^\alpha \mid D_t^\alpha = D_t^\alpha \cup v, \text{ and} \\ \forall \text{ retraction}(v) \in V_t^\alpha \mid D_t^\alpha = D_t^\alpha - v. \end{aligned} \quad (2)$$

Agent Interaction

Cooperating agents interact via a task-independent agent called facilitator, which coordinates and controls the exchange of messages. The facilitator captures the interests of agents and performs various functions aimed at facilitating the exchange of assertions and retractions of predicate-logic sentences. In the event of receiving messages, the facilitator determines the appropriate recipient agents of the messages and forwards them accordingly. In addition, it translates between vocabularies used by different agents in their exchange of sentences. The translation is achieved through a set of predicate-logic axioms R^ϕ defined over a subset of all agents' vocabularies $X^\phi \subseteq X$.

Consider the group of agents $\Gamma = \{\alpha, \beta, \dots, \zeta\}$ that are cooperating on solving a problem defined by agents' knowledge $K^\alpha, K^\beta, \dots, K^\zeta$ over a vocabulary X . The group of agents interact via the facilitator ϕ , which captures the agents' interests expressed as sets of predicate logic axioms $I^\alpha, I^\beta, \dots, I^\zeta$ and a set of translation axioms R^ϕ over X^ϕ . This can formally be expressed as follows:

$$\begin{aligned} X^\phi \subseteq X = X^\alpha \cup X^\beta \cup \dots \cup X^\zeta, \\ K = K^\alpha \cup K^\beta \cup \dots \cup K^\zeta \cup R^\phi, \text{ and} \\ I = I^\alpha \cup I^\beta \cup \dots \cup I^\zeta. \end{aligned} \quad (3)$$

At time t , when an agent $\xi \in \Gamma$ generates messages V_t^ξ , it updates its current local solution to D_t^ξ and then communicates V_t^ξ to the facilitator. When the facilitator receives the set of messages V_t^ξ , it first deduces additional sentences based on the axioms R^ϕ . The result of this translation step is the set of sentences U_t^ϕ whose number is typically greater than that in V_t^ξ . This translation step can formally be expressed as follows:

$$U_t^\phi = \text{sentences}(V_t^\xi) \cup R^\phi \text{ is consistent.} \quad (4)$$

Then the facilitator checks for the agents that are interested in the communicated sentences, U_t^ϕ . For every $\delta \in \Gamma - \{\xi\}$, if $I^\delta \cup U_t^\phi$ is consistent, agent δ is interested and is added to the set of interested agents Δ . For every interested agent $\delta \in \Delta$, the facilitator forwards an appropriate set of messages W_t^δ . When agent δ receives the set of messages W_t^δ , it updates its current local solution to D_t^δ .

Conflict Types

Agent δ checks the consistency of the updated local solution D_t^δ with respect to its knowledge K_t^δ and constraints C_t^δ . If

$$D_t^\delta \cup K_t^\delta \cup C_t^\delta \text{ is consistent,} \quad (5)$$

there is no conflict and agent δ accepts the messages. Conversely, if

$$D_t^\delta \cup K_t^\delta \cup C_t^\delta \text{ is inconsistent,} \quad (6)$$

there is a conflict and agent δ identifies the conflict as one of three types: critical conflict, non-critical conflict with authority, and non-critical conflict without authority. In this section, the three conflict types are formally discussed.

- *Critical Conflict*: a critical conflict is a conflict in which the updated solution based on the received messages is inconsistent with the agent's knowledge K_t^δ . Formally, a critical conflict can be expressed as follows:

$$D_t^\delta \cup K_t^\delta \text{ is inconsistent.} \quad (7)$$

- *Non-Critical Conflict with Authority*: a non-critical conflict with authority is a conflict in which the updated solution is consistent with the agent's knowledge K_t^δ but inconsistent with the agent's constraints C_t^δ , and the vocabularies of each sentence in W_t^δ belong to X^δ over which agent δ has authority. Formally, a non-critical conflict with agent δ having authority over the vocabulary can be expressed as follows:

$$\begin{aligned} D_t^\delta \cup K_t^\delta \text{ is consistent,} \\ D_t^\delta \cup C_t^\delta \text{ is inconsistent, and} \\ \forall \text{ message}(w) \in W_t^\delta \mid \text{vocabulary}(w) \in X^\delta. \end{aligned} \quad (8)$$

- *Non-Critical Conflict without Authority*: a non-critical conflict without authority is a conflict in

which the updated solution is consistent with the agent's knowledge K_t^δ but inconsistent with the agent's constraints C_t^δ , and the vocabularies of each message in W_t^δ do not belong to X^δ over which agent δ has authority. Formally, this can be expressed as follows:

$$\begin{aligned} D_t^\delta \cup K_t^\delta &\text{ is consistent,} \\ D_t^\delta \cup C_t^\delta &\text{ is inconsistent, and} \\ \forall \text{ message}(w) \in W_t^\delta \mid \text{vocabulary}(w) &\notin X^\delta. \end{aligned} \quad (9)$$

Progressive Negotiation: A Conflict Resolution Strategy

Conflict resolution is an essential requirement for cooperation for autonomous, intelligent, interacting agents (Adler et al. 1989). In conflict resolution, the role of negotiation has been emphasized as the focal point for conflict resolution in distributed problem solving for different domains (Durfee & Lesser 1987; Laasri, Laasri, & Lesser 1990; Lander & Lesser 1989). Our research has focused on developing a strategy called progressive negotiation for resolving conflicts that aims at minimizing backtracking to previous solutions. In this strategy, conflict resolution is carried out by agents. Depending on the type of conflict, negotiation takes place in an attempt to resolve the conflict. In this strategy, critical conflicts are always resolved because they result from the agent's knowledge, which must be satisfied in order to have a satisfiable solution. Non-critical conflicts are resolved by getting one of the agents to relax some of its violated criteria constraints in order for an agreement to be reached. The following is a formal treatment of how conflicts are resolved for the three types outlined in the previous section.

Critical Conflicts

When a critical conflict is identified, agent δ determines a set of axioms $Q_t^\delta \subseteq K_t^\delta$ that caused the conflict and sends it to the facilitator. The facilitator in turn forwards appropriate axioms to the sending agent ξ and other agents in Δ interested in Q_t^δ . Formally, the violated axioms can be expressed as follows:

$$Q_t^\delta = \{q \mid \forall q \in K_t^\delta \text{ such that } D_t^\delta \cup q \text{ is inconsistent}\}. \quad (10)$$

Once the sending agent ξ receives the set of axioms Q_t^ξ , at time t' , forwarded by the facilitator, it checks the consistency of the axioms Q_t^ξ with its knowledge K_t^ξ . If $Q_t^\xi \cup K_t^\xi$ is inconsistent, there is no solution that is consistent with the knowledge of both agents.

If $Q_t^\xi \cup K_t^\xi$ is consistent, however, then there could be a solution that is consistent with the knowledge of both agents. In this case, the agent's constraints are updated to C_t^ξ ensuring that $C_t^\xi \cup Q_t^\xi$ is consistent. This update may involve relaxation of some constraints in C_t^ξ . After updating its constraints, agent ξ generates new messages V_t^ξ and updates its local solution to D_t^ξ , in a way that ensures the consistency of the updated solution with its knowledge and updated constraints. Formally, we can write:

$$D_t^\xi \cup K_t^\xi \cup C_t^\xi \text{ is consistent.} \quad (11)$$

The messages V_t^ξ are then sent to the facilitator, which forwards appropriate messages to all interested both agents. At time t'' , agent δ receives the new messages $W_{t''}^\delta$ and updates its local solution to $D_{t''}^\delta$. The new local solution, at time t'' , for agent δ will probably be consistent with the agent's knowledge since $K_{t''}^\delta = K_{t''}^\delta$. Formally, we can write:

$$D_{t''}^\delta \cup K_{t''}^\delta \text{ is consistent.} \quad (12)$$

Non-Critical Conflicts with Authority

When a critical conflict is identified, agent δ determines a set of axioms $P_t^\delta \subseteq K_t^\delta$ that caused the conflict and sends it to the facilitator. The facilitator in turn forwards appropriate axioms to the sending agent ξ and other agents in Δ interested in P_t^δ . Formally, the violated axioms can be expressed as follows:

$$P_t^\delta = \{p \mid \forall p \in C_t^\delta \text{ such that } D_t^\delta \cup p \text{ is inconsistent}\}. \quad (13)$$

Once the sending agent ξ receives the set of axioms P_t^ξ , at time t' forwarded by the facilitator, it checks the consistency of the axioms P_t^ξ with its knowledge K_t^ξ knowing that $K_t^\xi = K_t^\xi$. If $P_t^\xi \cup K_t^\xi$ is inconsistent (i.e., the sets of violated axioms and the agent's knowledge do not lead to a solution), agent ξ rejects the received set of axioms P_t^ξ , and they are sent back to agent δ , which relaxes its set of constraints to $C_{t'}^\delta$ such that $D_{t'}^\delta \cup C_{t'}^\delta$ is consistent since $D_{t'}^\delta = D_{t'}^\delta$.

If $P_t^\xi \cup K_t^\xi$ is consistent, however, then agent ξ relaxes some of its constraints to C_t^ξ , if necessary, such that $P_t^\xi \cup C_t^\xi$ is consistent. Then agent ξ generates new messages V_t^ξ and updates its local solution to D_t^ξ in a way that ensures the consistency of the updated solution with the agent's knowledge and constraints,

$$D_i^\xi \cup K_i^\xi \cup C_i^\xi \text{ is consistent.} \quad (14)$$

The messages V_i^ξ are then sent to the facilitator, which forwards it to all interested agents. Agent δ receives the new messages W_i^δ at time t'' , and updates its local solution to D_i^δ . The new local partial solution for agent δ will provably be consistent with the agent's knowledge and constraints since $K_i^\delta = K_i^\xi$ and $C_i^\delta = C_i^\xi$,

$$D_i^\delta \cup K_i^\delta \cup C_i^\delta \text{ is consistent.} \quad (15)$$

Non-Critical Conflicts without Authority

This case is similar to the case of non-critical conflicts with authority. The only difference is that when agent ξ receives the constraints P_i^ξ forwarded by the facilitator, it checks $P_i^\xi \cup K_i^\xi \cup C_i^\xi$ knowing that $K_i^\xi = K_i^\delta$ and $C_i^\xi = C_i^\delta$. If $P_i^\xi \cup K_i^\xi \cup C_i^\xi$ is inconsistent, agent ξ rejects the set of constraints, and they are sent back to agent δ , which relaxes its set of constraints to C_i^δ such that $D_i^\delta \cup C_i^\delta$ is consistent since $D_i^\delta = D_i^\xi$.

If $P_i^\xi \cup K_i^\xi \cup C_i^\xi$ is consistent, however, then agent ξ generates new messages V_i^ξ and updates its local solution to D_i^ξ in a way that ensures the consistency of its knowledge and constraints,

$$D_i^\xi \cup P_i^\xi \cup K_i^\xi \cup C_i^\xi \text{ is consistent.} \quad (16)$$

Again the messages V_i^ξ are then sent to the facilitator, which forwards it to all interested agents. For agent δ , the new messages W_i^δ are received at time t'' and the local solution is updated to D_i^δ . The new local solution for agent δ will provably satisfy the agent's knowledge since $K_i^\delta = K_i^\xi$ and $C_i^\delta = C_i^\xi$,

$$D_i^\delta \cup K_i^\delta \cup C_i^\delta \text{ is consistent.} \quad (17)$$

Solution Consistency and Convergence

In this section, we present two theorems for proving the consistency of distributed local solutions and the convergence on a common globally satisfiable solution under certain conditions for the progressive negotiation strategy. Before presenting these two theorems, we give a number of definitions that are used in the proofs.

Definition 1: Let ξ denote an agent from the group of agents, $\xi \in \Gamma$. The interests of agent ξ are said to be complete if and only if:

$$\forall x \in X^\xi, \exists s \in I^\xi \text{ expressing interest in } x. \quad (18)$$

Definition 2: Let $K^\alpha, K^\beta, \dots, K^\zeta$, denote the knowledge for the group of agents $\Gamma = \{\alpha, \beta, \dots, \zeta\}$ and R^ϕ denote the translation axioms captured in the facilitator ϕ . The initial sets of axioms expressing agents' knowledge are said to be consistent if and only if:

$$K = K^\alpha \cup K^\beta \cup \dots \cup K^\zeta \cup R^\phi \text{ is consistent.} \quad (19)$$

Definition 3: Let $C^\alpha, C^\beta, \dots, C^\zeta$ denote the initial sets of constraints for the group of agents $\Gamma = \{\alpha, \beta, \dots, \zeta\}$. The initial sets of constraints are said to be consistent if and only if:

$$C = C^\alpha \cup C^\beta \cup \dots \cup C^\zeta \text{ is consistent.} \quad (20)$$

Definition 4: Let $Y^\alpha, Y^\beta, \dots, Y^\zeta$ respectively denote the authorities of the group of agents $\Gamma = \{\alpha, \beta, \dots, \zeta\}$. The authorities of the group of agents Γ are said to be exhaustive if and only if:

$$Y^\alpha \cup Y^\beta \cup \dots \cup Y^\zeta = X. \quad (21)$$

Definition 5: Let $Y^\alpha, Y^\beta, \dots, Y^\zeta$ respectively denote the authorities of the group of agents $\Gamma = \{\alpha, \beta, \dots, \zeta\}$. The authorities of the group of agents Γ are said to be disjoint (i.e., authorities over sets of interrelated vocabularies do not overlap) if and only if:

$$\begin{aligned} \forall \xi, \psi \in \Gamma \mid Y^\xi \subseteq X^\xi \wedge Y^\psi \subseteq X^\psi \\ \Rightarrow Y^\xi \cap Y^\psi = \emptyset, \text{ and} \\ \forall r \in R^\phi \text{ over } Z^\phi \subseteq X^\phi \wedge \forall \xi, \psi \in \Gamma \mid \\ Y^\xi \cap Z^\phi \neq \emptyset \wedge Y^\psi \cap Z^\phi = \emptyset. \end{aligned} \quad (22)$$

Consistency of Distributed Solutions

In this section, we present a theorem that states the consistency of distributed local solutions for the progressive negotiation strategy.

Theorem 1: For the group of agents Γ , if the interests of the agents are complete (which is usually the case), the progressive negotiation strategy guarantees the consistency of the distributed local solutions generated after every exchange among agents. Formally, the distributed local solutions after an exchange at time tn can be expressed as follows:

$$D_m = D_m^\alpha \cup D_m^\beta \cup \dots \cup D_m^\zeta \text{ is consistent.} \quad (23)$$

Proof: Suppose that $D_m = D_m^\alpha \cup D_m^\beta \cup \dots \cup D_m^\zeta$ is inconsistent. Then there must exist at least two local solutions D_m^α and D_m^β such that

$$D_m^\alpha \cup D_m^\beta \text{ is inconsistent.} \quad (24)$$

This means that there are at least two ground atoms, d^α and d^β , in the solutions D_m^α and D_m^β respectively such that $d^\alpha = \neg d^\beta$, which in turn, means that one of the agents α or β did not receive a message about it or update its solution. This is not possible because it contradicts the solution update step of agent interaction (2). Therefore, there cannot be any two partial solutions such that $D_m^\alpha \cup D_m^\beta$ is inconsistent. Thus, $D_m = D_m^\alpha \cup D_m^\beta \cup \dots \cup D_m^\zeta$ is always consistent according to the progressive negotiation strategy.

Convergence on a Common Solution

In this section, we present a theorem that states the convergence conditions of the progressive negotiation strategy.

Theorem 2: For the group of agents Γ , if the interests are complete (which is usually the case), and the sets of axioms expressing agents' knowledge are consistent, whereas the set of initial axioms expressing agents' constraints is not necessarily consistent, and the authorities are exhaustive and disjoint, then the progressive negotiation strategy guarantees convergence on a common solution that satisfies all agents' knowledge, and a relaxed version of the initial constraints. Formally, the following must be proven:

$$D_m \cup K_m^\alpha \cup \dots \cup K_m^\zeta \cup R^\phi \text{ is consistent.} \quad (25)$$

Proof: Suppose that $D_m \cup K_m^\alpha \cup \dots \cup K_m^\zeta \cup R^\phi$ is inconsistent. Thus, one can write:

$$D_m^\alpha \cup \dots \cup D_m^\zeta \cup K_m^\alpha \cup \dots \cup K_m^\zeta \cup R^\phi \text{ is inconsistent.} \quad (26)$$

Given that $K_m^\alpha \cup \dots \cup K_m^\zeta \cup R^\phi$ is consistent as stated in the theorem conditions, and having proved that $D_m^\alpha \cup \dots \cup D_m^\zeta$ is always consistent in Theorem 1, one can conclude that equation (26) can hold if and only if either of two possibilities holds:

- There exists at least one agent $\xi \in \Gamma$ such that its local solution is inconsistent with the facilitator's translation axioms: $D_m^\xi \cup R^\phi$, or
- There exist at least two agents $\xi, \psi \in \Gamma$ such that their local solutions are inconsistent with their knowledge: $D_m^\xi \cup D_m^\psi \cup K_m^\xi \cup K_m^\psi$.

Case 1: In order for $D_m^\xi \cup R^\phi$ to be inconsistent, there must be sent or received messages (respectively V_m^ξ or W_m^ξ) that are inconsistent with R^ϕ since the local solution atoms are typically updates of atoms in V_m^ξ or W_m^ξ . Formally, one of the following must hold:

$$\begin{aligned} V_m^\xi \cup R^\phi \text{ is inconsistent, or} \\ W_m^\xi \cup R^\phi \text{ is inconsistent.} \end{aligned} \quad (27)$$

This is not possible because of the facilitator's translation step expressed in equation (4) and the completeness of agents' interests. Therefore, $D_m^\xi \cup R^\phi$ is consistent and there cannot be any agent whose solution does not satisfy the facilitator's translation axioms.

Case 2: $D_m^\xi \cup D_m^\psi \cup K_m^\xi \cup K_m^\psi$ can never be inconsistent because $D_m^\xi \cup D_m^\psi$ is consistent from Theorem 1, $K_m^\xi \cup K_m^\psi$ is consistent from the theorem conditions, and $D_m^\xi \cup K_m^\xi$ is consistent and $D_m^\psi \cup K_m^\psi$ is consistent from equations (11, 12, 14, 15, 16, and 17) and other equations of the progressive negotiation strategy. Thus, there cannot be any agent whose local solution is inconsistent with other agents' knowledge. Therefore, $D_m \cup K_m^\alpha \cup \dots \cup K_m^\zeta \cup R^\phi$ is always consistent and the solution convergence of the progressive negotiation strategy is guaranteed.

Summary and Concluding Remarks

In this paper, we formally presented a conflict resolution strategy called Progressive Negotiation that guarantees the consistency of distributed solutions and convergence on a globally-satisfiable solution. The strategy involves communication of predicate-logic axioms to alter distributed agents' local solutions incrementally to reach a globally-consistent satisfiable solution. It aims at minimizing backtracking to previous solutions by getting agents to communicate their violated axioms and thus to inform all agents involved in a conflict about those axioms, which ensures they will not commit that violation again. The strategy assumes that interacting agents exchange their messages via a task-independent agent called facilitator, which controls the exchange of messages in a way that ensures the satisfaction of agents' knowledge and constraints with respect to their current local solutions. The theorems presented for proving convergence of the progressive negotiation strategy show the conditions under which a group of distributed cooperating heterogeneous agents are guaranteed to reach an agreement in the course of negotiation. It is possible that an agreement

among agents can be reached in some cases even if they do not comply with all the convergence conditions. However, if these conditions are not followed, convergence is not guaranteed in every exchange.

References

- Adler, M. R., Davis, A. B., Weihmayer, R., and Worrest, R. W. 1989. Conflict-Resolution Strategies for Non-hierarchical Distributed Agents. *Research Notes in Artificial Intelligence, Distributed Artificial Intelligence*, Morgan Kaufmann Publishers, Inc., pages 139-161.
- Conry S., Meyer, R., and Lesser, V. 1988. Multiagent Negotiation in Distributed Planning. *Readings in Distributed Artificial Intelligence*, A. Bond and L. Gasser, editors, pages 367-384. Morgan-Kaufmann Publishers, Inc.
- Durfee, E. H. and Lesser, V. R. 1987. Using Partial Global Plans to Coordinate Distributed Problem Solvers. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, pages 857-883.
- Durfee, E. H. 1988. *Coordination of Distributed Problem Solvers*. Kluwer Academic Publishers.
- Genesereth, M. R. 1992. An Agent-Based Framework for Software Interoperability. In *Proceedings of DARPA Software Technology Conference*, pages 359-366.
- Kraus, S. and Wilkenfeld, J. 1991. Negotiations over Time in a Multi Agent Environment: Preliminary Report. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, pages 56-61.
- Kuwabara, K. and Lesser, V. R. 1989. Extended Protocol for Multistage Negotiation. In *Proceedings of the Ninth Workshop on Distributed Artificial Intelligence*, pages 129-161.
- Laasri, B. Laasri, H. and Lesser, V. R. 1990. Negotiation and Its Role in Cooperative Distributed Problem Solving. In *Proceedings of the Tenth AAAI International Workshop on Distributed Artificial Intelligence*.
- Lander, S. and Lesser, V. R. 1989. A Framework for the Integration of Cooperative Knowledge-Based Systems. In *Proceedings of the IEEE International Symposium on Intelligent Control*, Albany, NY.
- Sycara, K. P. 1989. Argumentation: Planning other Agents' Plans. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 517-523.
- Zlotkin, G. and Rosenschein, J. S. 1993. A Domain Theory for Task Oriented Negotiation. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, pages 416-422.