The University of Kansas

Information and Telecommunication Technology Center

Technical Report

Literature Review on Agent-Based Negotiations

Leen-Kiat Soh

ITTC-FY2001-TR-19860-02

October 2000

Project Sponsor:
U.S. Air Force
Award Number: F30602-99-2-0502

Literature Review on Agent-Based Negotiations

Leen-Kiat Soh

Information and Telecommunication Technology Center (ITTC)
2291 Irving Hill Road
Lawrence, KS 66044

1. Introduction

Basically, agents negotiate to exchange information and knowledge. Each agent maintains its own knowledge base or information in a distributed network of agents. When an agent needs to collaborate with another agent to perform a task, the agent negotiates with that another agent. The negotiation usually involves exchanges of information or knowledge until a deal is agreed or rejected. There are several approaches to agent-based negotiations that we will discuss in this review. Multi-agent negotiations can be used to perform task allocations (collaborative agents), resource allocations (cooperative agents), knowledge distribution (learning agents), etc.

2. Bazaar and Bayesian Negotiation

The Bayesian negotiation in Bazaar (Zeng and Sycara 1996; 1997; 1998) is powered by a Bayesian belief network. Given each response from the counterpart, an agent updates its beliefs (through the belief network) regarding the counterpart's reservation price. From each response, an agent updates its beliefs (through the belief network) regarding its counterpart's reservation price. A supplier's reservation price, for example, is the price below which the supplier agent will not accept an offer. Then, the agent makes another offer based on the newly-updated reservation price. The unique characteristic of this work is its use of Bayesian belief network, which allows modeling of constraints and hypotheses formally. However, under close examination, the Bazaar's reliance on solely the offer or counter-offer to drive the updates in the belief network, in my opinion, constrains the approach. If additional information is exchanged between agents, then the belief network can provide formality and consistency. One also has to come up with the initial conditional probabilities to build the belief network.

Sycara's early work on Persuader (Sycara 1989a, 1989b, 1990a, 1990b, 1992) is closely related.

3. Negotiation Decision Models

Faratin *et al.* (1998) presented a formal model of negotiation between autonomous agents. It describes multi-party, multi-issue, and single-encounter negotiations among competitive and cooperative agents. The discussion was comprehensive and insightful. They have a family of negotiation tactics (which are similar to our strategies): (1) time-dependent, (2) resource-dependent, and (3) imitative. In my opinion, the time-dependent and resource-dependent tactics can be grouped together. It involves basically changing the function to vary based on two different issues. The imitative tactics are similar to our 'Bayesian' strategies (see Section 3.2 of this paper). In (Faratin *et al.* 1998), a strategy is defined as a combination of tactics. The authors also used criteria. The tactics (such as polynomial, exponential, Boulware, and Conceder functions) presented by Faratin *et al.* (1998) were reflective but prescribed. For example, the tactics are temporal aware because the functions used are time-dependent; the tactics are situational aware because of functions used can also be resource-dependent. In the imitative tactics, the negotiation depends on the counterpart. In my opinion, however, since agents are operating in a dynamic environment, not only the agents have prescribed reflectivity, but also event-driven reflectivity. In the absence of events, the negotiation process will depend on the agent model (current status), the target model, and the neighbor model. But if there is an event, then an agent should be ready to evaluate the new situation and react. In addition, an agent should also be historically reflective—it remembers what it did before through its retrieval and adaptation of old cases. Faratin's models did not address these two issues. In short, most negotiation models discussed only a prescribed reflectivity—following a set function based on some varying entities.

The authors defined *negotiation issues* as price, volumes, duration, quality, etc. Thus, if one agent is trying to sell an item to another agent and both are negotiating about the price, then the price is the negotiation issue. The authors defined *negotiation tactics* as the set of functions that determine how to compute the value of an issue by considering a single criteria (time, resources, etc.). The criteria the they described included time, resources, previous offers and counter offers. The tactics are:

(1) Time-dependent tactics – polynomial, exponential, Boulware tactics, and conceder.

(2) Resource-dependent tactics – (a) dynamic-deadline tactics in which the greater the number of a gents who are negotiating with agent A for a particular service, the lower the pressure on agent A to reach an agreement with any specific individual; but the longer the negotiation, the greater the pressure on A to come to an agreement, (b) resource-estimation tactics in which diminishing resource (based on time) is factored into the Boulware tactics.

(3) Behavior-dependent tactics – (a) relative tit-for-tat in which an agent reproduces, in percentage terms, the behavior that its opponent performed several steps ago, (b) random absolute tit-for-tat which is very similar to (a) but uses absolute increase or decrease, and (c) averaged tit-for-tat in which an agent computes the average of percentages of changes in a window of a certain size of its opponent's history when determining its new offer.

The authors also defined a *negotiation strategy* as the following. First, for an issue, there is a weighted counter proposal, made up of a linear combination of the tactics that generates the value at the current time step of a negotiation, and the set of tactics is finite. Given a set of tactics, different types of negotiation behavior can be obtained by weighting the tactics in a different way and the weights are specified in a matrix. A negotiation strategy is then any function that is based on the matrix of the weights and the mental state of the agent. In my opinion, this approach provides a strong and formal foundation to multi-tactic strategy generation and real-time negotiation strategy modifications during negotiation. Matos *et al.* (1998) described an evolutionary approach to determine successful negotiation strategies based on the above model.

4. Argumentative Negotiation Models

The argumentation-based negotiation model described in (Parsons and Jennings 1996, Sierra *et al.* 1997a, 1997b, Doran *et al.* 1997, Jennings *et al.* 1998, Parsons *et al.* 1998) was comprehensive. The authors used bridge rules with their logics and theories (Noriega and Sierra 1996). A bridge rule is a rule of inference with premises and conclusions in different units.

They also incorporated explicitly and rigorously multi-context Belief-Desire-Intention (BDI) agents (Rao and Georgeff 1991, 1995, Jennings 1995) and the formalism described in (Kraus *et al.* 1993) to construct arguments to evaluate proposals and counterproposals in negotiation.

Parsons and Jennings make a strong case for sending over rules of an agent to another agent. This allows the other agent to know the actual reasoning process behind its counterpart's claims. In an agency of diverse and competitive agents, this makes a lot of sense. When each agent has its own brain and own agendas, it is logical to have different reasoning processes. And so naturally, to convince another agent, an agent lets this another agent know about its reasoning steps. This approach allows diversification and efficient *knowledge* (and not just information!) distribution within an agency. Of course, in a system of homogeneous agent in which all agents share the exactly same reasoning, then the argumentation-based negotiation model of Parsons and Jennings might not be necessary. In my opinion, I see these models as true argumentative negotiations: not only agents argues about what evidence there is, they also argue about how one should interpret the evidence.

According to (Parsons *et al.*1998), there are several scenarios when an agent B dose not agree with the arguments of an agent A:

(1) arguments directly conflicts with the objective of B.
(2) arguments partially conflicts with the objective of B.
(3) Arguments are acceptable but B does not have the resources to accept the deal.

By the same token, there are different levels of acceptance, based on notions of argument defeat (Elvang-Goransson *et al.* 1993; Fox *et al.* 1992; Krause *et al.* 1995; Loui 1987). For example, when an agent evaluates a proposal and it falls into class A4 or A5, then the agent accepts the proposal; if the proposal falls into class A2 or A3, then the agent accepts it partially; and so on. This is amiable.

5. Game Theory

Kraus (1997a) applied game-theoretic techniques to multi-agent environments. If agents are self-motivated and try to maximize their own benefits, the designers of the agents may agree in advance on regulations and protocols for the agents' interaction. In each interaction, the number of agents is usually small; there are only automated agents which can communicate and have computational capabilities. In such situations, the authors recommended the application of game-theoretic techniques.

Game theory studies mathematical models of conflict and cooperation between people. The models of game theory are highly abstract representations of classes of real life situations that involve individuals who have different goals of preferences. The active entity in all game-theoretic models is a player. Game-theoretic models are divided into two main types: (1) non-cooperative models in which the sets of possible actions of individual players are primitives, and (2) cooperative models in which the sets of possible joint actions of groups of players are primitive (Osborn and Rubinstein 1994).

The authors found that formalizing and implementing informal models of behavioral and social sciences can be beneficial. These negotiation guides do not accept the strong restrictions and assumptions presented in the game-theoretic models such as Diplomat (Kraus and Lehmann 1995).

Diplomat (Kraus and Lehmann 1995) uses a 1-to-N negotiation approach. The agent uses a notice board that keeps track of the common expected profits of each negotiation so that the agent can make a selection among the N concurrent negotiated deals. The notice board also records the contradictions among the deals. Based on the profits and contradictions, each agent computes the *intention* of honoring a deal. Each agent also keeps track of a *loyalty* measure of its neighbors.

See (Shehory and Kraus 1995a, 1995b, 1996a, 1996b) on coalition formation among autonomous agents.

See also Section 6 of this review.

6. Model of Alternative Offers

Kraus and Wilkenfeld (1993) proposed a negotiation model on hostage crisis. The model is a modification of Rubinstein's model (Rubinstein and Osborne 1994) of alternative offers which focuses on the passage of time and the preferences of the players for different agreements as well as for opting out of the negotiations (Rubinstein 1982). The authors assumed that agents cared only about the nature of the agreement or opting out, and the time at which the outcome was reached, and not about the sequence of offers and counter-offers that led to the agreement. In particular, no agent regrets either making an offer that was rejected or rejecting an offer. They also assumed disagreement is the worst outcome and the initiator gained over time and the receiver lost over time, since this was applied to hostage crisis.

This paper talks about a strategic model for negotiation of alternative offers which takes into account the effect of time on the negotiation process. In the model, both players can opt out, and one of the players loses over time and the other gains over time. The paper shows that if there is an agreement zone, an agreement will be reached in the first or the second step of the negotiation as the players either have the complete knowledge of each other to begin with or have the complete knowledge of each other after the first offer. In my opinion, this model does not in general apply to multi-agent negotiations. It is not competitive in the way that an agent wants to gain as much as possible. It is the initiating agent's drive to keep negotiating before the counterpart opts out.

Kraus and Wilkenfeld (1991a, 1991b), Kraus *et al.* (1995) and Kraus (1997b) investigated further the negotiation model under time constraints. The authors examined the following situations and their corresponding negotiation strategies:

(1) *Situation*: Two agents that lose over time need to share a common resource and each agent knows all relevant information about the other agent. They have no alternative but to negotiate until an agreement is reached.
Strategy: The agents will reach the perfect equilibrium (a deal) in one handshake (or after the first offer).

(2) *Situation*: Two agents that lose over time need to cooperate to satisfy a common goal and each agent knows all relevant information about the other agent. Each agent can unilaterally leave the negotiations.
Strategy: The agents will reach the perfect equilibrium (a deal) in one handshake (or after the first offer).
(3) *Situation*: Two agents need to share a resource. One of the agents already has access to the resource and is using it during the negotiation process. It is thus gaining over time. The other agent is waiting to use the resource and loses over time. Both agents have full information and can unilaterally leave the negotiations.
Strategy: The agreement is guaranteed at the latest in the second step since even though the agent that has access to the resource prefers to continue the negotiation indefinitely, it is afraid that the counterpart might threaten to opt out at any given time. In my opinion, this strategy is simplistic. If all agents strive to achieve an agreement, then all agents can actually threaten to opt out of a negotiation!
(4) *Situation*: Similar to situation (3) but the agents do not have complete information about each other.
Strategy: The agents conduct a *sequential equilibrium* (Kreps and Wilson 1982) negotiation. The agents will reach the perfect equilibrium (a deal) in at most two steps.
(5) *Situation*: Several agents need to cooperate to satisfy a common goal. All of them are losing over time, have full information about each other, and can unilaterally leave the negotiations.
Strategy: At each time period, one of the agents makes an offer to all the other agents. Each of the agents either accepts the offer, or rejects it, or opts out of the negotiation. If the offer is accepted by all the agents, then the negotiation ends and the agreement is implemented. Also, opting out by one of the agents ends the negotiation. After a rejection by at least one agent, another agent must make a counter offer, and so on. This mechanism provides each of the agents with a veto power on the agreements that will be reached. In my opinion, this is a powerful and useful negotiation strategy.

The sequential equilibrium (Kreps and Wilson 1982) strategy involves a sequence of strategies and a system of belief with the following properties: each agent has a belief about its counterpart's type. At each negotiation step, the strategy for an agent A is optimal given its current belief and its counterparts possible strategies in the sequential equilibrium. Ad each negotiation step, each agent's belief is consistent with the history of the negotiation. The history is a sequence of previous proposals and responses. Therefore, this approach requires the agents to know the profile of the strategies and the beliefs of the agents according to the following three conditions: (1) *sequential rationality* in which an agent tries to maximize its expected utility with regard to the strategies of its opponents and its beliefs about the probabilities of its opponent's type according to the given history, (2) *consistency* in which an agent's subsequent beliefs should be consistent with its initial belief and with the possible strategies of its opponent, and (3) *never dissuaded once convinced*.

There is a philosophical issue regarding this approach. The agents are non-cooperative. However, in my opinion, by basing the model on Rubinstein's alternative offers, the agents are *implicitly* cooperative. In a complete knowledge model, an agent knows exactly what the other agent wants and thus makes an offer accordingly. In a partially complete information model, an agent first makes an offer based on partial information. The other agent learns full information from the offer and thus makes an offer accordingly. Even though the agents are non-cooperative, they do share pre-negotiation information and thus are, in my opinion, implicitly cooperative.

Moreover, I think the approach is particular useful when inter-agent communication is minimal, individual agent knowledge is small, and less *intelligent* as a result. I also think that this approach can be very useful in investigating emergent behaviors of agents.

See also a closely related paper (Kraus *et al.* 1991).

7. Unified Negotiation Protocol

Zlotkin and Rosenschein (1991) and Rosenschein and Zlotkin (1994a) outlined a theoretical negotiation model for rational agents in general non-cooperative domains. In general non-cooperative domains, the negotiation set is sometimes empty. They presented a unified negotiation protocol (UNP) for conflict resolutions and agents taking partial cooperative steps. The protocol can be applied to two general agent scenarios. First, agents have been centrally designed to coexist in a single system and are predisposed toward cooperative activity—or in the least, there is some notion of global utility that the system is trying to maximize. Second, agents are self-serving, have their own utility functions and no global notion of utility. These agents have disparate goals and are individually motivated. Negotiation can be used to share the work associated with carrying out a joint plan (for the agents' mutual benefits), or to resolve outright conflict arising from limited resources.

The authors presented several assumptions of an agent (Zlotkin and Rosenschein 1989):

(1) Utility maximizer – Each agent wants to maximize its expected utility.
(2) Complete knowledge – Each agent knows all relevant information.
(3) No history – There is no consideration given by the agents to the past or future; each negotiation stands alone.
(4) Fixed goals – Though the agents negotiate with one another over operations, their goals remain fixed, i.e., agents do not change their pre-defined goals.
(5) Bilateral negotiation – In a multiagent encounter, negotiation is done between a pair of agents at a time, i.e., no 1-to-N negotiations.
(6) Symmetric abilities – all agents are able to perform the same set of operations in the world. In my opinion, this assumption is necessary since their agents have to operate in a non-cooperative, non-collaborative environment. Also, since the agents during negotiation can swap plan segments, it is important to have symmetric abilities in this unified negotiation protocol. In my opinion, in general, agents can be heterogeneous.
(7) Deterministic world – The world is deterministic and altered only by agents' actions. In my opinion, this assumption is generally not true.

The utility of an agent from a deal is simply the difference between the cost of achieving its goal alone and its expected part of the deal. And (1) a deal is *individual rational* if the utility of that deal to each partner of the deal is not negative, (2) a deal is *pareto-optimal* if there does not exist another deal that dominates it—there does not exist another deal that is better for one of the agents and not worse of the other, and (3) the *negotiation set* is the set of all the deals that are both individual rational and pareto-optimal. These defintitions are from game theory and bargaining theory (Harsanyi 1977, Luce and Raiffa 1957, Nash 1950).

A necessary condition for the negotiation set to be non-empty is that there is no contradiction between the two agents' goals. However, this condition is not sufficient since there may still be a conflict between the agents. A *conflict* is where any joint plan that satisfies the union of goals will cost one agent (or both) more than it would have spent achieving his own goal in isolation—i.e., no deal is individual rational. In addition, a joint plan is said to satisfy the *sum condition* if the sum of all costs for all agents to perform the task individually is greater than or equal to the sum of all costs for all agents to perform the parts of the joint plan. Moreover, a joint plan is said to satisfy the *min condition* if the minimum cost for an agent to perform the task individually is greater than or equal to the minimum cost for an agent to perform the joint plan. When these two conditions are true, then the agents are in *cooperative* situations. In my opinion, this defines cooperative agents generally and implies selfishness and collaborativeness.

Now, in non-conflict situations, if neither the min nor the sum conditions are true, then in order for the agents to cooperate, then at least one of the agents will have to do more than if it were alone in the world and achieved only its own goals. In order to be motivated to do so, an agent has to redefine its utility or the "worth of a goal". The authors defined the *worth* of a goal as the maximum *expected* cost that an agent is willing to pay in order to achieve its goal. And thus the utility of an agent from a deal can be modified as the difference between the worth of a goal and the agent's cost of its expected part of the deal. Given this redefinition of utility, the agents may now be involved in three possible interactions:

(1) A *cooperative* situation in which there exists a deal in the negotiation set that is preferred by an agent over achieving his goal alone. Here, an agent welcomes the existence of the other agents. However, in my opinion, in order for the other agents to accept the plan, then this particular agent has to convey voluntarily its preference to the other agents.

(2) A *compromise* situation in which there are individual rational deals for an agent. However, an agent would prefer to be alone in the world, and to accomplish its goal alone. All of the deals in the negotiation set are better for the agent than leaving the world in its initial state.

(3) A *conflict* situation in which the negotiation set is empty.

When a conflict arises, the authors proposed a conflict resolution via coin toss based on weights. So, one agent sacrifices to help another agent achieve that agent's goals. When this occurs, we have a *semicooperative deal*, a tuple of (t, J, q) where t is a desired world state, J is a mixed joint plan that moves the world from the initial state to t, and q is the weighting of the coin toss—the probability that an agent will achieve its goal.

In cooperative and compromise situations, the agents negotiate on deals that are mixed joint plans (or cooperative deals) while in conflict situations, on semicooperative deals. In addition, the authors also defined a multiplan deal concept in which a deal has two mixed joint plans (A and B) that each is able to move the world to a state that satisfies an agent's goal and a probability that the agent will perform one of the plans. This allows a negotiation over two joint plans simultaneously. The unified negotiation protocol (UNP) that the authors proposed was to be used in any of the above situations, without having to compute, prior to the negotiation process, the type of situation that the agents are in.

The authors also define a hierarchy of deal types. First, a pure deal is a joint plan, and a mixed deal is a mixed joint plan. A joint plan is a mixed joint plan, a mixed joint plan is a semiccooperative deal, and a semicooperative deal is a multiplan deal.

In short, the various definitions of deals together with the weighting of the coin toss or the probability of choice between plans allow the utility of a plan to be computed for each agent. The utilities enable the agents to make decisions. Also, the agent-based negotiations proposed by the authors do not involve the local elements of the negotiations as discussed in other sections of this review. This unique approach to negotiations is refreshing.

The authors further extended their theoretical negotiation model to task-oriented, state-oriented, and worth-oriented functions (Rosenschein and Zlotkin 1994b, Zlotkin and Rosenschein 1996a, 1996b, 1996c). In the task-oriented domain—the lowest and simplest level, agents have nonconflicting jobs to do and these jobs or tasks can be redistributed among the agents. Thus, the agents receive some list of jobs that they have to accomplish and the object of negotiation in this environment is to redistributed tasks among the agents to everyone's mutual benefit if possible. In the state-oriented domain—the second level, and a superset of the task-oriented domains, actions in these domains can have *side effects*, where an agent performing one task might hinder or help another agent. The object of negotiation is to develop joint plans or schedules for the agents (as previously discussed above). In the worth-oriented domain—the highest level, and a superset of the state-oriented domains, agents strive for better states through a decision-theoretic formulation. The object of negotiation is joint plans or schedules and goal relaxation.

Rosenschein and Zlotkin (1994b) identified three types of task-oriented domains: (1) subadditive where tasks may be inter-related and by combining them may cost less, (2) concave (a subset of subadditive cases) where an additional task adds less cost to a task X than to a task Y, and (3) modular where tasks are not related and cannot be less costly done if combined. The authors also outlined some interesting agent behaviors for a deceitful agent: "letter-hiding" to withhold information, "phantom letters" and "decoy task" to mislead opponents.

Zlotkin and Rosenschein (1996c) detailed the goal relaxation in a multiplan deal. They described a process where the states with the highest value of worth might be thought of as those that satisfy the full goal, while others, with lower worth values, only partially satisfy the goal. They also focused on product maximizing mechanisms for summing up worth values. The authors identified different worth functions: (1) subgoal set, (2) distance between states, and (3) probabilistic distance to handle uncertainty.

Zlotkin and Rosenschein (1996a) further defined their UNP model and design mechanisms. They see the protocol as the rules of negotiation, but the exact deals that an agent proposes is a result of the strategy that the agent's designer has put into it. They also identified several attributes of standards for agents:

(1) Efficiency – the agents should not squander resources when they come to an agreement.

(2) Stability – No designer should have an incentive to deviate from agreed-upon strategies. This is assuming heterogeneous agents.
(3) Simplicity – It will be desirable for the overall interaction environment to make low computational demands on the agents and to require little communication overhead.
(4) Distribution – Preferably, the interaction rules will not require a central decision maker.
(5) Symmetry – No designer wants the negotiation process to be arbitrarily biased against its agent.

Zlotkin and Rosenschein (1996b) also defined altruistic agents, tidy agents, strict and tolerant product maximizing mechanisms for dealing with a conflict when both agents declare worths lower than their stand-alone costs.

8. Partial Global Planning, TEAM, and TAEMS

Research by Lesser and Durfee and colleagues (Decker and Lesser 1993, Durfee and Lesser 1991) conducted research in which agent-based negotiations were essentially for the collection and combination of various information pieces to derive meaningful and correct observations of the world.

In Durfee and Lesser (1991), the authors presented a partial global planning overview. This planning gives an agent the ability to (1) represent its own expected interpretation activities (of sensory data), (2) communicate about these expectations with others, (3) model the collective activities of multiple systems, (4) propose changes to one of or more systems' interpretation activities to improve group performance, and (5) modify its planned local activities in accordance with the more coordinated proposal. Thus, the object of negotiation here is to combine information from several agents in order to reach an outcome for the entire group of agents. Agents are also organized hierarchically: (1) subordinate agents – an agent reviews the plan information from these subordinates, makes modifications, and sends the modified plan information back, (2) supervisor agents – an agent sends its plan information to the these agents, receives the modified plan information and can adopt the modified plans, and (3) peer agents – an agent sends information to another agent and receives modified plan information back and can carry out its tasks locally. During the information sending process, the goals, long-term strategy and the rating of a plan are exchanged.

In Lander and Lesser (1992), the authors introduced a framework called TEAM that implements the cooperative search and conflict resolution among heterogeneous, reusable, expert agents, based on the partial global planning approach. They identified two types of heterogeneity: (1) logical where agents may have different long-term knowledge (expertise), goals, views or perpectives, constraints or preferences, or criteria for evaluating solutions, and (2) implementational where agents may have different knowledge representations, languages, architectures, inference engines, etc. However, in my opinion, the communication language has to be one that is known to all agents. The authors also defined a set of requirements for the agents:
(1) Agents sets are dynamically formed by grouping agents with the specific expertise required for the problem.
(2) Agents do not have prior knowledge of what other agents will be included in the set and what their capabilities will be.

(3) The agent sets are cooperative—agents are not hostile and will not intentionally mislead or otherwise try to sabotage another agent's reasoning.

(4) Agents are willing to contribute both knowledge and solutions to other agents as appropriate and to accept solutions that are not locally optimal in order to find a mutually-acceptable solution.

In my opinion, the above definitions are one of the most general discussed in this review. The second definition is particularly powerful as it implies that agents must learn from each other to form a group among which they can exchange information (or negotiate). The authors also characterized two basic approaches to negotiated search. First, when an agent recognizes a conflict with another agent in an existing solution, it extends its local search until a solution is found that does not conflict and the general search solutions are heuristic search, case-based search, and searching for alternate goal expansions. Second, an agent relaxes some requirement on a solution, thereby expanding its local search space and the general methods are relaxing or relinquishing constraints, relaxing or relinquishing goals, manipulating constraints, or manipulating evaluation criteria

The authors investigated further on the negotiated-search strategies: (1) local search, (2) integration of local search, and (3) general negotiated search. A local search is performed individually by an agent within its current view of the shared solution space. This search includes an initiation, a critique, and a solution relaxation. In an integration of local search, solutions are critiqued by other agents. From the critiques, conflicts are determined locally by each agent. The agent then performs relaxation (which lowers the utility of the solution to the agent) and computes the acceptability of the solution. The acceptability is based on the number of agents considering the solution to be acceptable. Finally, a general negotiated search is an opportunistic search augmented by the communication and assimilation of conflict information. It behaves in the following manner:

(1) One or more agents produce base proposals,

(2) Other agents critique the partial solutions created from those proposals.

(3) If a conflict is detected, any constraining information is communicated to other agents.

(4) Agents that receive conflict information attempt to assimilate that informatio.

(5) If an agent has successfully assimilated conflict information from another agent and later attempts to generate a proposal, the new proposal will avoid that conflict.

(6) If the required number of solutions is not found within a specified number (the *relaxation threshold*) of agent cycles, each agent will apply the negotiated-search operator *relax-solution-requirement* to expand the solution space.

(7) Agents continue in a cycle of search and relaxation until an acceptable solution is found, or until further relaxation is impossible (at which point a failure is declared).

The TEAM framework itself consists of distinct agents that communicate through a shared memory (a blackboard) and the framework is controlled by a framework controller. During run-time, there are two distinct phases: (1) an agent cycle, and (2) a framework cycle. During the agent cycle, each agent is invoked sequentially. The agent uses information in shared memory to choose applicable negotiated-search operators and add them to its agenda. It then invokes its highest-priority operators and returns the result. After all agents have executed, the framework controller is invoked to update the shared memory based on messages from agents and to propagate the effect of changes to shared memory objects. Thus, in my opinion, the TEAM framework is synchronous.

Each agent maintains a local knowledge base. There is no consistency enforcement across agents. When there are explicit inconsistencies in the local and the received knowledge, an agent must choose what information to use in making decisions, whether to override local knowledge and use the conflicting external information, or ignore conflicting external information and use the local knowledge, or relax local solution requirements to reduce the inconsistency and reexamine the situation within the new solution boundaries. The authors proposed a flexibility value that measures the degree for each piece of information to which an agent is willing to relax that information. This is a good idea.

Lander and Lesser (1993) elaborated further on the negotiated search operators of the TEAM framework. Decker and Lesser (1994) outlines an extendable family of coordination mechanisms, called generalized partial global planning, and implemented the TAEMS framework (Task Analysis, Environment Modeling, and Simulation) to represent coordination problems in a formal, domain-independent way. The TAEMS framework also attaches a negotiability index to every commitment (high, medium or low) that indicates the difficulty in rescheduling a taskgroup (a group of tasks) if the commitment is broken.

See Durfee *et al.* (1989) also for a discussion on various trends in cooperative distributed problem solving

In short, there is no explicit negotiation involved in the above approach. Agents exchange information through blackboards, and re-plan based on the critiques and conflicts. Then they re-post their solutions to the blackboard and the framework controller sends the information to every agent in the group. The solutions are refined until acceptable by all agents or otherwise. However, the negotiated search in addressing inconsistencies such as relaxation is very important in agent-based negotiations, and the taskgroup and commitment sharing is critical in negotiations as well.

9. A Logical Model -- ANA

In (Kraus *et al.* 1998), the authors described a logical model for multi-agent argumentative negotiations and an implementation. The paper described a logical model of the mental states of the agents based on a representation of their beliefs, desires, intentions, and goals, and defined argumentation as an iterative process emerging from exchanges among agents to persuade each other and bring a bout a change in intentions. The implemented agent is called the Automated Negotiation Agent or ANA. The authors extended the Belief, Desire, and Intention model to include also Goal.

The paper defined arguments as utterances whose aim is to change the intentions (and consequently the actions) of the listener. In order to negotiate effectively, an agent needs the ability to (1) represent and maintain a model of its own beliefs, desires, goals, and intentions, (2) reason with other agents' beliefs, desires, goals, and intentions, and (3) influence other agents' beliefs, intentions, and behavior. The mental state of an agent is thus characterized by using the notions of beliefs, goals, desires, intentions, and local *preferences*—that determines the degrees of desires, intentions, and goals. The paper also described several agent characteristics:

(1) Bounded – a bounded agent does not believe in falsehood.
(2) Omniscient – an omniscient agent has beliefs that are closed under inferences, i.e., its beliefs are consistent.
(3) Knowledgeable – an agent is knowledgeable if its beliefs are correct.
(4) An agent will keep believing in something until it observes something otherwise; or an agent is memoryless; or an agent does not forget anything.
(5) Confident – an agent is confident if it believes that it will succeed in carrying out its intended actions.

The paper identified six argument types:

(1) Threats to produce goal adoption or goal abandonment on the part of the responder.
(2) Enticing the responder with a promise of a future reward
(3) Appeal to past promise
(4) Appeal to precedents as counterexamples to convey to the responder a contradiction between what he/she says and the past actions
(5) Appeal to prevailing practice to convey to the responder that the proposed action will further his/her goals since it has furthered others' goals in the past
(6) Appeal to self-interest to convince a responder that taking this action will enable achievement of a high-importance goal

In my opinion, to conduct in the above argument types, an agent has to know another agent is doing. As claimed in the paper, an agent A can threaten B or entice B to not perform a task. That means, A has to have the knowledge that B is trying to perform a specific task. Without explicit knowledge exchange, how can A learn such information? If B knows that A is afraid of B performing a task T, then can B keep threatening A that it is going to perform T unless A promises a future reward? In another example regarding the appeal to past promise, a similar assumption has to exist. The scenario described by the authors was the following: Suppose A promises B that it will do T1 at time t if B does T2. So when time t arrives, B asks A to perform task T1. If A refuses, B reminds A of the past promise. In my opinion, this arrangement is not logical since by the same token, A has to know whether B has performed task T2.

The paper also described the structure of an agent, i.e., the Automated Negotiation Agent (ANA). It consists of a mental state (beliefs, desires, goals, and intentions), characteristics (agent type, capabilities, belief verification capabilities), and inference rules (mental state update, argument generation, argument selection, request evaluation). Thus, at the designer of a specific agent can influence its general behavior by providing the mental state update rules, argument generation and selection mechanism, request evaluation rules, and initial beliefs and desires at agent creation time.

The paper also outlines some negotiation meta-rules in the areas of argument generation, argument selection, and request evaluation. Since these meta-rules are critical to the negotiation process, we will discuss them in details here.

First, the argument generation rules are used when an agent identifies an intention that it cannot execute. Here are some rules:

(1) If the opponent is a memoryless agent, then do not choose "Appeal to Past Promise".

(2) If the agent received a request from the opponent in the past, which included a future reward argument, and if that reward was the intended action right now, then choose "Appeal to Past Promise".

(3) Use "Appeal to Self Interest" when the agent believes the opponent is not aware of the implications—therefore this should not be used with a knowledgeable or reasonable agent, since the inferences of such agents are closed under consequences.

(4) Use "Appeal to Prevailing Practice" if you can find a third agent that has done this before. In my opinion, this requires a record of past events.

(5) Use "Counterexample" if the agent knows about the past history of the responder.

(6) Use "Promise of a Future Reward" if the agent can find something the opponent likes. This is an interesting approach. Usually, I expect a negotiation to go this way: A intends that B do X. B counter-offers, "If I do X, then you must do Y for me." And A responds accordingly. However, with use of "Promise of a Future Reward", the negotiation goes this way: A intends that B do X. B says no. Then A finds something that B might be interested and then makes another offer, "Would you please do X for me if I promise to do Y for you?" Note that since A wants X performed, so naturally A should work harder to achieve that goal—that is, A should come up with some enticements to convince B to perform X. But this is not practical since it requires A to know what B wants or might be interested, a feature that may not be affordable in many agent cases.

(7) Use "Threat" if the agent can find something that the opponent hates. First, obtain the list of desires of the opponent. Consider first desires with the highest preference value for the opponent which are not included in the agent's own desires set. Also, try to find a desire which involves action that the agent can perform while its opponent cannot. Then find a contradicting action to that desire. If the agent cannot, then choose "Appeal to Self Interest". Otherwise, threatens.

Second, an agent uses argument selection rules to choose from a set of potential arguments the best argument to communicate to its opponent. The authors implemented only one selection rule by severity ordering. The agent will first try to use the weakest argument and if it does not succeed, it will follow with stronger arguments. The order is set as follows: (1) an Appeal to Prevailing Practice, (2) a Counter-example, (3) an Appeal to Past Promise, (4) an Appeal to Self Interest, (5) a Promise of a Future Reward, and (6) a Threat.

Third, an agent uses request evaluation rules to evaluate a request (or a counter-request) to decide whether to accept or reject it. An agent computes three parameters regarding a request: (1) collision_flag to indicate whether the results of the requested action conflict with the agent's current goals, (2) convincing_factor to indicate how convincing the argument is given the requested action, and (3) acceptance_value to indicate the overall preference of the results of the requested action as opposed to all the other desires of the agent. In my opinion, the convincing factor should be continuous instead of a 0-or-1 binary variable. The agent also considers the agent's reliability, the other agent's reliability for keeping promises, and the other agent's percentage of threat executing when computing for the acceptance_value.

10. Others

Smith and Davis (Smith 1979, 1980, and Smith and Davis 1988, 1983) proposed a system of exchanging commitments and goals and task sharing and allocation called the contract-net. This works spawns the later work in negotiations, described in this review. Matwin *et al.* (1989) proposed a decision support system called Negoplan for conducting negotiations. Conry *et al.* (1988, 1991) termed *multistage negotiation* in distributed planning and satisfaction in search.

Lâasri *et al.* (1992) described a recursive negotiation model that defines where and how negotiation can be applied during problem solving based on structuring problem solving into four stages: (1) problem formulation, (2) focus-of-attention, (3) allocation of goals or tasks to agents, and (4) achievement of goals or tasks. The negotiation is based on formulation and exchange of information especially in terms of proposals, critiques, explanations, and meta-information. A proposal is a specific solution, a partial result, a cluster solution or a partial cluster. Each agent may generate proposals independently, or based on proposals generated by the other agents, or based on critiques made by other agents of previous proposals. A critique consists of a positive and a negative components. The positive component describes the features of the proposal the agent agrees with, whereas the negative those the agent disagrees with. An explanation is information about the knowledge and reasoning used to evaluate the proposal (such as inference rules). The meta-information guides negotiation search space and defines the general character of the agent's problem-solving capabilities. The authors also provided a detailed overview of some related work in negotiations and cooperative distributed problem solving.

Sandholm (1993) and implemented a contract net called TRACONET (Transportation Cooperation Net). Sandholm and Lesser (1995a) extended TRACONET with multiple levels of commitments for negotiation protocols. The use allows the agents to (1) move a low-commitment search focus around in the global task allocation space, covering more space to avoid risky commitments, (2) add flexibility to the agent's local deliberation control, (3) be able to make the same low-commitment offer to multiple agents (in case more than one accepts, the agent has, of course, to pay the penalty to all but one of them, but the speedup of being able to address multiple agents in committal mode may outweigh the risk), (4) have a lower decommitting penalty by trading off paying a higher price to its contractee if the agent is risk averse, and (5) condition the payments and commitment functions on future negotiation events or domain events. In contract-net, tasks were negotiated one at a time. TRACONET handles tasks interactions by having the announcer cluster tasks into sets to be negotiated atomically and equipping the bidders with the ability to re-cluster by counter-proposing. The authors also addressed issues of message congestion such as the tragedy of the commons where too much traffic decimates the utility. Sandholm and Lesser (1995b) also addressed coalition formation for agents.

Chavez and Maes (1996) described a simplistic negotiation approach in Kasbah. It is not adaptive and not reflective. It follows one function throughout the negotiation process and does not have situational awareness. Strictly speaking, it *does* have some sort of inflexible temporal awareness since the functions are time-dependent decays (e.g., linear, quadratic, or cubic). In this manner, an agent makes an offer based on the function it chose before the negotiation.

Walton and Krabbe (1995) listed several types of negotiation dialogs: persuasion, inquiry, negotiation, information seeking, and deliberation. A persuasion happens when there is a conflict of opinions and an agent has to persuade the other agent to resolve the conflict. As discussed above, our argumentative behavior includes this persuasion dialog when two agents are trying to agree on a set of consistent evidence or support. An inquiry happens when one agent needs to have proof for what the other agent describes. Again, this dialog is also part of the argumentative behavior as one agent asks the other agent to supply evidence, and subsequently analyzes the evidence. A negotiation occurs when there is conflict of interests. The information seeking dialog occurs when one agent lacks information and the other agent has to provide the information. This dialog mode is very similar to the inquiry mode. However, during the inquiry mode, the receiving agent will rigorously prove the validity of the evidence; whereas during the information seeking mode, the receiving agent will receive the information as facts. Finally, the deliberation occurs when an agent has to make a choice.

Liu and Sycara (1997) also used negotiation to describe searching for solutions in distributed resource allocation problems. Chavez *et al.* (1997) described a system called Challenger for distributed resource allocation. Dworman *et al.* (1996) and Prasad *et al.* (1997) integrated learning into negotiation. Pruitt (1981) discussed human negotiation from the viewpoint of social psychology.

13. References

Chavez, A. and P. Maes (1996). Kasbah: An Agent Marketplace for Buying and Selling Goods, in *Proceedings of the 1st International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology*, 75-90.

Chavez, A., A. Moukas and P. Maes (1997). Challenger: A Multi-agent System for Distributed Resource Allocation, in *Proceedings of the International Conference on Autonomous Agents*, Marina Del Ray, CA.

Conry, S. E., R. A. Meyer, and V. R. Lesser (1988). Multistage Negotiation in Distributed Planning, in A. H. Bond and L. Gaser (eds.), *Readings in Distributed Artificial Intelligence*, 367-384, San Mateo, CA: Morgan Kaufmann.

Conry, S. E., K. Kuwabara, V. R. Lesser, and R. A. Meyer (1991). Multistage Negotiation for Distributed Satisfaction, *IEEE Transactions on Systems, Man, and Cybernetics, Special Issue on Distributed Artificial Intelligence*, **21**(6):1463-1477.

Decker, K. and V. Lesser (1993). A One-Shot Dynamic Coordination Algorithm for Distributed Sensor Networks, in *Proceedings of AAAI-93*, pp. 210-216.

Decker, K. and V. Lesser (1994). Designing a Family of Coordination Algorithms, in *Proceedings of the 13th International Workshop on Distributed Artificial Intelligence*, July.

Doran, J. E., S. Franklin, N. R. Jennings, and T. J. Norman (1997). On Cooperation in Multi-Agent Systems, *The Knowledge Engineering Review*, **12**(3):309-314.

Durfee, E. H. and V. R. Lesser (1991). Partial Global Planning: A Coordination Framework for Distributed Hypothesis Formation, *IEEE Transactions on Systems, Man, and Cybernetics*, **21**(5):1167-1183.

Dworman, G., S. O. Kimbrough, and J. D. Laing (1996). Bargaining by Artificial Agents in Two Coalition Games: A Study in Genetic Programming for Electronic Commerce, in *Proceedings of AAAI Genetic Programming Conference.*

Elvang-Goransson, M., P. Krause, and J. Fox (1993). Dialectic Reasoning with Inconsistent Information, in *Proceedings of the 9th Conference on Uncertainty in Artificial Intelligence.*

Faratin, P., C. Sierra, and N. R. Jennings (1998). Negotiation Decision Functions for Autonomous Agents, *International Journal of Robotics and Autonomous Systems*, **24**(3-4):159-182.

Fox, J., P. Krause, and S. Ambler (1992). Arguments, Contradictions and Practical Reasoning, in *Proceedings of the 10th European Conference on Artificial Intelligence.*

Harsanyi, J. C. (1977). *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*, Cambridge: Cambridge University Press.

Jennings, N. R. (1995). Controlling Cooperative Problem Solving in Industrial Multi-Agent Systems Using Joint Intentions, *Artificial Intelligence*, **75**:195-240.

Jennings, N. R., P. Faratin, M. J. Johnson, T. J. Norman, P. O'Brien, and M. E. Wiegand (1996). Agent-Based Business Process Management, *International Journal of Cooperative Information Systems*, **5**(2-3):105-130.

Jennings, N. R., S. Parsons, P. Noriega, and C. Sierra (1998). On Argumentation-Based Negotiation, in Proceedings of International Workshop on Multi-Agent Systems, Boston, MA.

Kraus, S. (1997a). Negotiation and Cooperation in Multi-Agent Environments, *Artificial Intelligence*, **94**(1-2):79-97. Special Issue on economic Principles of Multi-Agent Systems.

Kraus, S. (1997b). Beliefs, Time, and Incomplete Information in Multiple Encounter Negotiations among Autonomous Agents, *Annals of Mathematics and Artificial Intelligence*, **20**(1-4):111-159.

Kraus, S., E. Ephrati, and D. Lehmann (1991). Negotiation in a Non-Cooperative Environment, *Journal of Experimental and Theoretical Artificial Intelligence*, **3**(4):255-282.

Kraus, S. and D. Lehmann (1995). Designing and Building a Negotiating Automated Agent, *Computational Intelligence*, **11**(1):132-171.

Kraus, S., M. Nirkhe, and K. Sycara (1993). Reaching Agreements through Argumentation: A Logic Model (Preliminary Report), in *Proceedings of the Workshop on Distributed Artificial Intelligence.*

Kraus, S., K. Sycara, and A. Evenchik (1998). Reaching Agreements through Argumentation: A Logical Model and Implementation, *Artificial Intelligence Journal*, **104**(1-2):1-69.

Kraus, S. and V. S. Subrahmanian (1995). Multiagent Reasoning with Probability, Time, and Beliefs, *International Journal of Intelligent Systems*, **10**(5):459-499.

Kraus, S. and J. Wilkenfeld (1993). A Strategic Negotiations Model with Applications to an International Crisis, *IEEE Transactions on Systems, Man, and Cybernetics*, **23**(1):313-323.

Kraus, S. and J. Wilkenfeld (1991a). Negotiations over Time in a Multiagent Environment: Preliminary Report, in *Proceedings of IJCAI-91*, Australia, 56-61.

Kraus, S. and J. Wilkenfeld (1991b). The Function of Time in Cooperative Negotiations, in *Proceedings of AAAI-91*, California, 179-184.

Kraus, S., J. Wilkenfeld, and G. Zlotkin (1995). Multiagent Negotiation under Time Constraints, *Artificial Intelligence*, **75**:297-345.

Krause, P., S. Ambler, M. Elvang-Gøransson, and J. Fox (1995). A Logic of Argumentation for Reasoning under Uncertainty, *Computational Intelligence*, **11**:113-131.

Kreps, D. and R. Wilson (1982). Sequential Equilibria, *Econometrica*, **50**:863-894.

Lâasri, B., H. Lâasri, S. Lander, and v. Lesser (1992). A Generic Model for Intelligent Negotiating Agents, *International Journal of Intelligent and Cooperative Information Systems*, **1**(2):291-317.

Lander, S. and V. Lesser (1992). Customizing Distributed Search Among Agents with Heterogeneous Knowledge, in *Proceedings of the 1st International Conference on Information and Knowledge Management*, November, Baltimore, MD, 335-344.

Lander, S. and V. Lesser (1993). Understanding the Role of Negotiation in Distributed Search among Heterogeneous Agents, in *Proceedings of the 12th International Workshop on DAI*, Hidden Valley, PA.

Liu, J. S. and K. Sycara (1997). Coordination of Multiple Agents for Production Management, *Annals of Operations Research*, **75**:235-289.

Luce, R. D. and H. Raiffa (1957). Games and Decisions, New York: Wiley.

Loui, R. (1987). Defeat among Arguments: A System of Defeasible Inference, *Computational Intelligence*, **3**:100-106.

Maes, P. (1996). Agents that Reduce Work and Information Overload, *Communications of the ACM*, **37**(7):31-40.

Matos, N., C. Sierra, and N. R. Jennings (1998). Determining Successful Negotiation Strategies: An Evolutionary Approach, in *Proceedings of the 3rd International Conference on Multi-Agent Systems (ICMAS-98)*, Paris, France, 182-189.

Matwin, S., S. Szpakowicz, Z. Koperczak, G. E. Kersten, and W. Michalowski (1989). Negoplan: An Expert System Shell for Negotiation Support, *IEEE Expert*, **4**(4):50-62.

Nash, J. F. (1950). The Bargaining Problem, *Econometrica*, **28**:155-162.

Noriega, P. and C. Sierra (1996). Towards Layered Dialogical Agents, in *Proceedings of the 3rd International Workshop on Agents Theories, Architectures and Languages*, 157-171.

Osborne, M. J. and A .Rubinstein (1994). *A Course in Game Theory*, Cambridge, MA: MIT Press.

Parsons, S. and N. R. Jennings (1996). Negotiation through argumentation—A Preliminary Report, in *Proceedings of the International Conference on Multi-Agent Systems*, Kyoto, Japan.

Parsons, S., C. Sierra, and N. R. Jennings (1998). Agents that Reason and Negotiate by Arguing, *Journal of Logic and Computation*, **8**(3):261-292.

Prasad, M. V. N., V. R. Lesser, and S. E. Lander (1997). Learning Organizational Roles for Negotiated Search in a Multi-Agent System, *International Journal of Human Computer Studies, Special Issue on Evolution and Learning in Multiagent Systems.*

Pruitt, D. G. (1981). *Negotiation Behavior*, Academic Press.

Rao, A. and M. Georgeff (1991). Modeling Rational Agents within a BDI-Architecture, in *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*, 473-484.

Rao, A. and M. Georgeff (1995). BDI Agents: From Theory to Practice, in *Proceedings of the 1st International Conference on Multi-Agent Systems*, 312-319.

Rosenschein, J. S. and G. Zlotkin (1994a). Rules of Encounter, Cambridge, MA: The MIT Press.

Rosenschein, J. S. and G. Zlotkin (1994b). Designing Conventions for Automated Negotiation, *AI Magazine*, **15**(3):29-46.

Rubinstein, A. (1982). Perfect Equilibrium in a Bargaining Model, *Econometrica*, **50**(1):97-109.

Sandholm, T. (1993). An Implementation of the Contract Net Protocol Based on Marginal Cost Calculations, in *Proceedings of the 11th National Conference on Artificial Intelligence*, Washington, D.C., 256-262.

Sandholm, T. and V. Lesser (1995a). Issues in Automated Negotiation and Electronic Commerce: Extending the Contract Net Framework, in *Proceedings of the 1st International Conference on Multiagent Systems (ICMAS-95)*, San Francisco, CA.

Sandholm, T. and V. Lesser (1995b). Coalition Formation among Bounded Rational Agents, in *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI-95)*, Montral, Canada.

Shehory, O. and S. Kraus (1995a). Coalition Formation Among Autonomous Agents: Strategies and Complexity, in *From Reactions to Cognition*, Berlin:Springer, 57-72.

Shehory, O. and S. Kraus (1995b). Task Allocation via Coalition Formation Among Autonomous Agents, in *Proceedings of IJCAI-95*, Montreal, Quebec, Canada, 655-661.

Shehory, O. and S. Kraus (1996a). A Kernel-Oriented Model for Coalition-Formation in General Environments: Implementation and Results, in *Proceedings of AAAI-96*, Portland, OR, 134-140.

Shehory, O. and S. Kraus (1996b). Emergent Cooperative Goal-Satisfaction in Large Scale Automated-Agent Systems, in *Proceedings of ECAI-96*, Budapest, Hungary, 544-548.

Sierra, C., P. Faratin, and N. R. Jennings (1997a). A Service-Oriented Negotiation Model Between Autonomous Agents, in *Proceedings of the 8th European Workshop on Modelling Autonomous Agents in a Multi-Agent World*, 17-35.

Sierra, C., N. R. Jennings, P. Noriega, and S. Parsons (1997b). A Framework for Argumentation-Based Negotiation, in *Proceedings of the 4th International Workshop on Agent Theories, Architectures and Languages*, 167-182.

Smith, R. G. (1979). A Framework for Distributed Problem Solving, in *Proceedings of the Sixth International Joint Conference on Artificial Intelligence*, August, Cambridge, MA, 863-841.

Smith, R. G. (1980). The Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver, *IEEE Transactions on Computer*, **C-29**(12):1104-1113.

Smith, R. G. and R. Davis (1981). Frameworks for Cooperation in Distributed problem Solving, *IEEE Transactions on Systems, Man, and Cybernetics*, **SMC-11**(1):61-70. (Also published in *Readings in Distributed Artificial Intelligence*, A. h. Bond and L. Gasser (eds.), Morgan Kaufmann, 333-356.)

Smith, R. G. and R. Davis (1983). Negotiation as a Metaphor for Distributed Problem Solving, *Artificial Intelligence*, **20**:63-109.

Sycara, K. (1989a). Argumentation: Planning Other Agents' Plans, in Proceedings of the 11th International Joint Conference on Artificial Intelligence, 517-523, Detroit, MI.

Sycara, K. (1989b). Multi-Agent Compromise via Negotiation, in *Distributed Artificial Intelligence*, Gasser, L. and M. Huhns (eds.), Los Altos, CA: Morgan Kaufman.

Sycara, K. (1990a). Persuasive Argumentation in Negotiation, *Theory and Decision*, **28**(3):203-242.

Sycara, K. (1990b). Negotiation Planning: An AI Approach, *European Journal of Operational Research*, **46**(2):216-234.

Sycara, K. (1992). The PERSUADER, in *The Encyclopedia of Artificial Intelligence*, Shapiro, D. (ed.), New York, NY: John Wiley and Sons.

Zeng, D. and K. Sycara (1996). Bayesian Learning in Negotiation, in *Working Notes of the AAAI Spring Symposium on Adaptation, Co-Evolution and Learning in Multiagent Systems*, Stanford, CA.

Zeng, D. and K. Sycara (1997). Benefits of Learning in Negotiation, in *Proceedings of AAAI-97.*

Zeng, D. and K. Sycara (1998). Bayesian Learning in Negotiation, *International Journal of Human-Computer Studies*, **48**:125-141.

Walton, D. N. and E. C. W. Krabbe (1995). *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*, Albany, NY: State University of New York Press.

Zlotkin, G. and J. S. Rosenschein (1989). Negotiation and Task Sharing Among Autonomous Agents in Cooperative Domains, in *Proceedings of the Eleventh International Joint Conference of Artificial Intelligenc (IJCAI)*, Detroit, MI, August, 912-917.

Zlotkin, G. and J. S. Rosenschein (1991). Cooperation and Conflict Resolution via negotiation among Autonomous Agents in Noncooperative Domains, *IEEE Transactions on Systems, Man, and Cybernetics, Special Issue on Distributed Artificial Intelligence*, **21**(6):1317-1324.

Zlotkin, G. and J. S. Rosenschein (1996a). Mechanism Design for Automated Negotiation, and Its Application to Task Oriented Domains, *Artificial Intelligence*, **86**(2):195-244.

Zlotkin, G. and J. S. Rosenschein (1996b). Mechanisms for Automated Negotiation in State Oriented Domains, *Journal of Artificial Intelligence Research*, **5**:163-238.

Zlotkin, G. and J. S. Rosenschein (1996c). Compromise in Negotiation: Exploiting Worth Functions Over States, *Artificial Intelligence*, **84**(1-2):151-176.