Coalition, Cryptography, and Stability: Mechanisms for Coalition Formation in Task Oriented Domains

Gilad Zlotkin
Center of Coordination Science
Sloan School of Management, MIT
1 Amherst Street, E40-179
Cambridge, MA 02139 USA
gilad@mit.edu

Jeffrey S. Rosenschein
Computer Science Department
Hebrew University
Givat Ram, Jerusalem
Israel
jeff@cs.huji.ac.il

Abstract

Negotiation among multiple agents remains an important topic of research in Distributed Artificial Intelligence (DAI). Most previous work on this subject, however, has focused on bilateral negotiation, deals that are reached between two agents. There has also been research on n-agent agreement which has considered "consensus mechanisms" (such as voting), that allow the full group to coordinate itself. These group decision-making techniques, however, assume that the entire group will (or has to) coordinate its actions. Sub-groups cannot make sub-agreements that exclude other members of the group.

In some domains, however, it may be possible for beneficial agreements to be reached among sub-groups of agents, who might be individually motivated to work together to the exclusion of others outside the group. This paper considers this more general case of n-agent coalition formation. We present a simple coalition formation mechanism that uses cryptographic techniques for subadditive Task Oriented Domains. The mechanism is efficient, symmetric, and individual rational. When the domain is also concave, the mechanism also satisfies coalition rationality.

Introduction

In multi-agent domains, agents can often benefit by coordinating their actions with one another; in some domains, this coordination is actually required. In two-agent encounters, the situation is relatively simple: either the agents reach an agreement (i.e., coordinate their actions), or they do not. With more than two agents, however, the situation becomes more complicated, since agreement may be reached by sub-groups.

The process of agent coordination, and of reaching agreement, has been the focus of much research in Distributed Artificial Intelligence (DAI). The general term used for this process is "negotiation" (usually in the 2-agent case) (Conry, Meyer, & Lesser 1988; Kraus & Wilkenfeld 1991; Kreifelts & von Martial 1990; Kuwabara & Lesser 1989; Sycara 1988; Zlotkin & Rosenschein 1993a; Rosenschein & Zlotkin 1994), and "reaching consensus" (in the n-agent case) (Ephrati & Rosenschein 1991; 1993). Both approaches, though dealing with different numbers of agents, share one underlying assumption: the agreement, if it is reached, will include all relevant members of the encounter. Thus, even in the n-agent case where a voting procedure might enable consensus to be reached, the entire group will be bound by the group decision. Sub-groups cannot make sub-agreements that exclude other members of the group. Interesting variations on these approaches, which nonetheless remain bilateral in essence, are the Contract Net (Smith 1978), which allows bilateral agreement in n-agent environments, and bilateral negotiation among two sub-groups discussed in (Kraus, Wilkenfeld, & Zlotkin 1995).

In some domains, however, it may be possible for beneficial agreements to be reached among sub-groups of agents, who might be individually motivated to work together to the exclusion of others outside the group. Voting procedures are not applicable here, because the full coalition may not be able to satisfy all its members, who are free to create more satisfying sub-coalitions. This paper considers this more general case of n-agent coalition formation (recent pieces of work on similar topics are (Ketchpel 1993; Shechory & Kraus 1993)). Building on our previous work (Zlotkin & Rosenschein 1993a), which dealt only with bilateral negotiation mechanisms, we here analyze the kinds of n-agent coordination mechanisms that can be used in specific classes of domains.

Coalitions

An Example—The Tileworld

Consider the following simple example in a multi-agent version of the Tileworld (Pollack & Ringuette 1990) (see Figure 1). A single hole in the grid is represented by a framed letter (such as [a]). Each agent's position is marked by its name (such as A_1). Tiles are represented by black squares (■) inside the grid squares.

Agents can move from one grid square to another horizontally or vertically (unless the square is occupied by a hole—multiple agents can be in the same grid square at the same time). When a tile is pushed into any grid square that is part of a hole, the square is filled and becomes navigable as if it were a regular

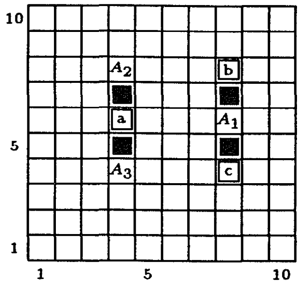

Figure 1: Three-Agent Encounter in the Tileworld

grid square. The domain is static except for changes brought about by the agents.

Agent 1's goal is to fill hole [a], while agents 2 and 3 need to fill holes [b] and [c] respectively. To fill its hole, each agent needs to do 7 steps. Agents can cooperate and help each other to reduce the cost of achieving their goals. There are several kinds of joint plans that the agents can execute that will reduce the cost of achieving their goals. Some of those joint plans are listed in the table on the right side of Figure 1.

The *coalition structure* {1,3},{2} means that there are two *coalitions*, one consisting of the agents 1 and 3, and the other consisting only of agent 2. When two agents form a coalition it means that they are coordinating their actions. The utility of an agent from a joint plan that achieves his goal is the difference between the cost of achieving his goal alone and the cost of his part of the joint plan (Zlotkin & Rosenschein 1991).

The coalition that gives the maximal total utility is the full coalition that involves all 3 agents, where they all coordinate their actions to mutual benefit (total utility is 17).[1] Although this full coalition is globally optimal, Agent 1's utility is only 4, and he would prefer to reach agreement with either agent 2 or agent 3 (with utility of 6), but not with both.

The agents in the above scenario are able to transfer utility to each other, but in a non-continuous way. Agent 1, for example, can "transfer" to agent 2 seven points of utility by achieving his goal. He cannot, however, transfer an arbitrary amount. Without this arbitrary, continuous utility transfer capability, agent 1 will prefer to form a coalition with either one of the other two agents, rather than with both.

Coalition Games

The definitions below are standard ones from coalition theory (Kahan & Rapoport 1984).

Definition 1 *A* coalition *game with transferable utility in normal characteristic form is* (N, v) *where:* $N = \{1, 2, \ldots, n\}$ *set of agents, and* $v\colon 2^N \rightarrow \mathbb{R}$. *For each coalition which is a subset of agents* $S \subseteq N$, $v(S)$ *is the* value *of the coalition* S, *which is the total utility that the members of* S *can achieve by coordinating and acting together.*

[1] The joint plan where agent 1 achieves both 2 and 3's goals (with cost of 3), while either agent 2 or 3 achieves 1's goal (each with expected cost of $\frac{1}{2}$).

Coalitions	u_1	u_2	u_3
{1},{2},{3}	0	0	0
{1,2},{3}	6	6	0
{1,3},{2}	6	0	6
{2,3},{1}	0	0	0
{1,2,3}	4	$6\frac{1}{2}$	$6\frac{1}{2}$

Table 1: Possible Coalitions in the Tileworld Example

The Tileworld example from Figure 1 can be described as a coalition game (N, v) such that: $v(\{1\}) = v(\{2\}) = v(\{3\}) = v(\{2,3\}) = 0$, $v(\{1,2\}) = v(\{1,3\}) = 12$, and $v(\{1,2,3\}) = 17$.

Note that the value derived by a coalition is independent of the coalition structure. A given coalition is guaranteed to get a certain utility, regardless of what coalitions are formed by the other agents. In the Tileworld domain this assumption is not necessarily true—though it is true in the example we gave above. We will see below that in Task Oriented Domains (Zlotkin & Rosenschein 1993a) this definition of the coalition value is directly applicable.

Task Oriented Domains

Definition 2 *A* Task Oriented Domain *(TOD) is a tuple* $< \mathcal{T}, \mathcal{A}, c >$ *where:* $\mathcal{T}$ *is the set of all possible tasks;* $\mathcal{A} = \{A_1, \ldots A_n\}$ *is an ordered list of agents;* c *is a* monotonic *function* $c\colon [2^{\mathcal{T}}] \rightarrow \mathbb{R}^+$. $[2^{\mathcal{T}}]$ *stands for all the finite subsets of* $\mathcal{T}$. *For each finite set of tasks* $X \subseteq \mathcal{T}$, $c(X)$ *is the cost of executing all the tasks in* X *by a* single *agent.* c *is monotonic, i.e., for any two finite subsets* $X \subseteq Y \subseteq \mathcal{T}, c(X) \leq c(Y); c(\emptyset) = 0$.

An encounter *within a TOD* $< \mathcal{T}, \mathcal{A}, c >$ *is an ordered list* $(T_1, \ldots, T_n)$ *such that for all* $k \in \{1 \ldots n\}, T_k$ *is a* finite *set of tasks from* $\mathcal{T}$ *that* A_k *needs to achieve.* T_k *will also be called* A_k*'s* goal.

The Postmen Domain (Zlotkin & Rosenschein 1989) is one classic example of a TOD. In this domain, each agent is given a set of letters to deliver to various nodes on a graph; starting and ending at the Post Office, the agents are to traverse the graph and make their deliveries. Agents can reach agreements to carry one another's letters, and save on their travel.

In multi-agent Task Oriented Domains, agents can reach agreements about the re-distribution of tasks among themselves. When there are more than two agents, the agents can also form coalitions such that tasks are re-distributed only among the members of the same coalition. When mixed deals are being used by agents (those are agreements where agents settle on a probabilistic distribution of tasks), it can be useful to conceive of the interaction as a coalition game with transferable utility. The use of probability smooths the discontinuous distribution of tasks, and therefore of utility. However, utility is still not money in a classic TOD; utility is the difference between the cost of achieving your goal alone, and the cost of your part of the deal. Therefore, there is an upper bound on the amount of utility that each agent can get—no agent can get more utility than his stand-alone cost. As we shall see below, however, our model never attempts to violate this upper bound on utility.

Subadditive Task Oriented Domains

In some domains, by combining sets of tasks we may reduce (and can never increase) the total cost, as compared with the sum of the costs of achieving the sets separately. The Postmen Domain, for example, has this property, which is called *subadditivity*. If X and Y are two sets of addresses, and we need to visit all of them ($X \cup Y$), then in the worst case we will be able to do the minimal cycle visiting the X addresses, then do the minimal cycle visiting the Y addresses. This might be our best plan if the addresses are disjoint and decoupled (the topology of the graph is against us). In that case, the cost of visiting all the addresses is equal to visiting one set plus the cost of visiting the other set. However, in some cases we may be able to do better, and visit some addresses on the way to others.

Definition 3 *TOD* $< \mathcal{T}, \mathcal{A}, c >$ *will be called* **subadditive** *if for all finite sets of tasks* $X, Y \subseteq \mathcal{T}$, *we have* $c(X \cup Y) \leq c(X) + c(Y)$.

Coalitions in Subadditive Task Oriented Domains

In a TOD, a group of agents (a coalition) can coordinate by redistributing their tasks among themselves. In a subadditive TOD, the way to minimize total cost is to aggregate as many tasks as possible into one execution batch (since the cost of the union of tasks is always less than the sum of the costs). Therefore, the maximum utility that a group can derive in a subadditive TOD is the difference between the sum of standalone costs and the cost of the overall union of tasks. This difference will be defined to be the value of the coalition.

Definition 4 *Given an encounter* $(T_1, \ldots, T_n)$ *in a subadditive TOD* $< \mathcal{T}, \mathcal{A}, c >$, *we will define the coalition game* induced *by this encounter to be* (N, v), *such that* $N = \{1, 2, \ldots, n\}$, *and* $\forall S \subseteq N$, $v(S) = \sum_{i \in S} c(T_i) - c(\bigcup_{i \in S} T_i)$.

Superadditive Coalition Games

It seems intuitively reasonable that agents in a coalition game should not suffer by coordinating their actions with a larger group. In other words, if you take two disjoint coalitions, the utility they can derive together should not be less than the sum of their separate utilities (at the worst, they could "coordinate" by ignoring each other). This property (which, however, is not always present) is called *superadditivity*.

Definition 5 *A* coalition *game with transferable utility in normal characteristic form* (N, v) *is* superadditive *if for any disjoint coalitions* $S, V \subset N, S \cap V = \emptyset$, *then* $v(S) + v(V) \leq v(S \cup V)$.

Theorem 1 *Any encounter* $(T_1, \ldots, T_n)$ *in a subadditive TOD induces a* superadditive *coalition game* (N, v).

Proof. Proofs can be found in (Zlotkin & Rosenschein 1993b). □

Mechanisms for Subadditive TODs

We would like to set up rules of interaction such that communities of self-interested agents will form beneficial coalitions. There are several attributes of the rules of interaction that might be important to the designers of these self-interested agents (as discussed further in (Rosenschein 1993)):

1. Efficiency: The agents should not squander resources when they come to an agreement; there should not be wasted utility when an agreement is reached. Since the coalition game is superadditive it means that the sum of utilities of the agents should be equal to $v(N)$.

2. Stability: Since the coalition game is superadditive, the full coalition can always satisfy the efficiency condition, and therefore we will assume that the full coalition will be formed. The stability condition then relates to the payoff vector $(u_1, \ldots, u_n)$ that assigns to each agent i a utility of u_i. There are three levels of stability (rationality) conditions: individual, group, and coalition rationality. *Individual Rationality* means that that no individual agent would like to opt out of the full coalition; i.e., $u_i \geq v(\{i\}) = 0$. *Group Rationality (Pareto Optimality)* means that the group as a whole would not prefer any other payoff vector over this vector; i.e., $\sum_{i=1}^{n} u_i = v(n)$. This condition is equivalent to the efficiency condition above. *Coalition Rationality* means that no group of agents should have an incentive to deviate from the full coalition and create a sub-coalition; i.e., for each subset of agents $S \subseteq N$, $\sum_{i \in S} u_i \geq v(S)$.

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, for all the obvious reasons. We do not want our distributed system to have a performance bottleneck, nor collapse due to the single failure of a special node.

5. Symmetry: Two symmetric agents should be assigned the same utility by the mechanism (two agents

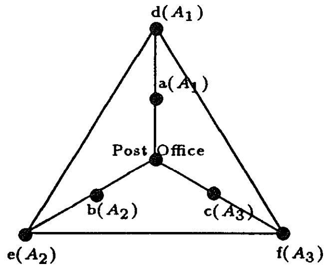

Figure 2: Example of an Unstable Encounter

are symmetric when they contribute exactly the same value to all possible coalitions).

We will develop a mechanism for subadditive TODs such that agents agree on the all-or-nothing deal, in which each agent has some probability of executing all the tasks. The question that we will try to answer now is "What should be the division of utilities among all agents in the full coalition?"

Coalition rationality is the strongest stability condition, and implies individual rationality and group rationality.[2] However, this condition is very strong, and cannot always be satisfied.

Consider the encounter from a three-agent Postmen Domain that can be seen in Figure 2.

The Post Office is in the center. The length of each arch is 1. The encounter is $(T_1 = \{a, d\}, T_2 = \{b, e\}, T_3 = \{c, f\})$.[3] Each agent can deliver his letters with a cost of 4. The cost of delivering the union of the letters of any two agents is 5. Therefore, the value of any two agents' coalition is $(2*4) - 5 = 3$. The cost of delivering all the letters is 8. Therefore, the value of the full coalition is $(3*4) - 8 = 4$. We would like to find a payoff vector (u_1, u_2, u_3) that satisfies the following conditions:

(1) $\forall i \in \{1,2,3\}\ u_i \geq v(\{i\}) = 0$;
(2) $\forall i \neq j \in \{1,2,3\}\ u_i + u_j \geq v(\{i,j\}) = 3$;
(3) $u_1 + u_2 + u_3 \geq v(\{1,2,3\}) = 4$.

Since the full coalition is also the maximal valued configuration, condition (3) is satisfied by equality (i.e., $u_1 + u_2 + u_3 = 4$). If we add up all the inequalities, we will have $u_1+u_2+u_3 >= 4\frac{1}{2}$, which cannot be satisfied. This means that in any division of the value of the full coalition among the agents there will be at least two agents that will prefer to opt out of the coalition and form a sub-coalition! For example, assume that the full coalition is formed with payoff vector $(1, 1, 2)$. Agents 1 and 2 can get more by forming a coalition (i.e., by excluding agent 3 from the coalition). The new payoff vector can then be $(1\frac{1}{2}, 1\frac{1}{2}, 0)$. This coalition and payoff vector is also not stable, since now agent 3 can tempt agent 2 (for example) to form a coalition with 3 by promising 2 more utility. The new payoff vector can then be $(0, 2, 1)$. However, now agent 1 can convince the two agents that they all can do better by forming the full coalition again. The new payoff vector can then be $(\frac{1}{3}, 2\frac{1}{3}, 1\frac{1}{3})$. This coalition is also not stable...

[2] All payoffs that satisfy the coalition rationality conditions are called the *core* of the game in the game theory literature. See, for example, (Kahan & Rapoport 1984).

[3] Agent 1 has to deliver letters to addresses a and d, agent 2 has to deliver letters to addresses b and e, and agent 3 has to deliver letters to addresses c and f.

Shapley Value

The Shapley Value (Shapley 1988; Young 1988) for agent i is a weighted average of all the utilities that i contributes to all possible coalitions. The weight of each coalition is the probability that this coalition will be formed in a random process that starts with the one-agent coalition, and in which this coalition grows by one agent at a time such that each agent that joins the coalition is credited with his contribution to the coalition. The Shapley Value is actually the expected utility that each agent will have from such a random process (assuming any coalition and permutation is equally likely).

Definition 6 *Given a superadditive coalition game with transferable utility in normal characteristic form (N, v), the Shapley Value is defined to be:* $u_i = \sum_{S \subset N, i \notin S} \frac{(n-|S|-1)!|S|!}{n!} v(S \cup \{i\}) - v(S)$.

The Shapley Value satisfies the efficiency, symmetry, and individual rationality conditions (Shapley 1988; Kahan & Rapoport 1984). However, it does not necessarily satisfy the coalition rationality condition.

Theorem 2 *The Shapley Value is also:* $u_i = c(T_i) - \sum_{S \subset N, i \notin S} \frac{(n-|S|-1)!|S|!}{n!} \Delta_c^i(S)$. $\Delta_c^i(S) \equiv c(S \cup \{i\}) - c(S)$, *i.e., the additional* cost *that agent i adds to a coalition S.*

Agent i's Shapley Value is the difference between the cost of its goal and its weighted average cost contribution to all possible coalitions. The cost that agent i can contribute to a coalition is bounded by $c(T_i)$. Therefore, the average contribution is also bounded by $c(T_i)$, which also means that the Shapley Value is positive (i.e., satisfies the individual rationality contribution) and bounded by $c(T_i)$ (which is also the maximal utility that an agent can get according to our model). Thus (as we promised above in Section), our model never attempts to transfer to an agent more utility than he can get by simply having his tasks performed by others.

Mechanisms for Subadditive TODs

We can define a Shapley Value-based mechanism for subadditive TODs that forms the full coalition and divides the value of the full coalition using the Shapley Value. The mechanism simply chooses the following (all-or-nothing) mixed deal, $(p_1, \ldots, p_n)$, such that $p_i = \frac{\sum_{S \subset N, i \notin S} \frac{(n-|S|-1)!|S|!}{n!} \Delta_c^i(S)}{c(N)}$.

Theorem 3 *The above all-or-nothing deal is well-defined, (i.e., $\forall i \in N: 0 \le p_i \le 1; \sum_{i=1}^{n} p_i = 1$) and gives each agent i an expected utility that is exactly the Shapley Value u_i.*

Evaluation of the Mechanism

The above mechanism gives each agent its Shapley Value. The mechanism is thus symmetric and efficient (i.e., satisfying group rationality), and also satisfies the criterion of individual rationality. However, as was seen in Example 2, no mechanism can guarantee coalition rationality. Besides failing to guarantee coalition rationality, the mechanism also does not satisfy the simplicity condition. It requires agents to calculate the Shapley Value, a computation that has exponential computational complexity.

The computational complexity of a mechanism should be measured relative to the complexity of the agent's standalone planning problem. This relative measurement would then signify the computational overhead of the mechanism. Each agent in a Task Oriented Domain needs to calculate the cost of his set of tasks, i.e., to find the best plan to achieve them. Calculation of the value of a coalition is linear in the number of agents in the coalition.[4] The calculation of the Shapley Value requires an evaluation of the value of all (2^n) possible coalitions. In Section below we will show that there exists another Shapley-based mechanism that has linear computational complexity.

Concave TODs

Definition 7 [Concavity]:[5] *TOD $< \mathcal{T}, \mathcal{A}, c >$ will be called* **concave** *if for all finite sets of tasks $X \subseteq Y$, $Z \subseteq \mathcal{T}$, we have $c(Y \cup Z) - c(Y) \le c(X \cup Z) - c(X)$.*

All concave TODs are also subadditive. It turns out that general subadditive Task Oriented Domains can be restricted, becoming concave Task Oriented Domains. For example, the Postmen Domain is subadditive, when the graphs over which agents travel can assume any topology. By restricting legal topologies to trees, the Postmen Domain becomes concave.

Definition 8 *A* coalition *game with transferable utility in normal characteristic form (N, v) is* convex *if for any coalitions S, V, $v(S) + v(V) \le v(S \cup V) + v(S \cap V)$.*

In convex coalition games, the incentive for an agent to join a coalition grows as the coalition grows.

Theorem 4 *Any encounter $(T_1, \ldots, T_n)$ in a concave TOD induces a* convex *coalition game (N, v).*

Theorem 5 [Shapley (1971)] (Shapley 1971): *In convex coalition games, the Shapley Value always satisfies the criterion of coalition rationality.*

[4]The cost of a set of tasks needs to be calculated only a linear number of times.

[5]The definition is from (Zlotkin & Rosenschein 1993a).

In concave TODs, the Shapley-based mechanism introduced above is fully stable, i.e., satisfies individual, group, and coalition rationality.

The Random Permutation Mechanism

The Shapley Value is equal to the expected contribution of an agent to the full coalition, assuming that all possible orders of agents joining and forming the full coalition are equally likely. This leads us to a much simpler mechanism called the *Random Permutation Mechanism*: agents choose a random permutation and form the full coalition, one agent after another, according to the chosen permutation. Each agent (i) gets utility (w_i) that is equal to its contribution to the coalition, at the time he joined it. This is done by agreeing on the all-or-nothing deal, $(p_1, \ldots, p_n)$, such that $p_i = \frac{c(T_i) - w_i}{c(N)}$.

Theorem 6 *If each permutation has an equal chance of being chosen, then the Random Permutation Mechanism gives each agent an expected utility that is equal to its Shapley Value.*

The Shapley-based Random Permutation Mechanism does not explicitly calculate the Shapley Value, but instead calculates the cost of only n sets of tasks. Therefore, it has linear computational complexity. The problem of coalition formation is reduced to the problem of reaching consensus on a random permutation.

Consensus on Permutation

No agent would like to be the first one that starts the formation of the full coalition (since this agent by definition gets zero utility). If the domain is concave (and therefore the coalition game is convex), each agent has an incentive to join the coalition as late as possible. To ensure stability, we need to find a consensus mechanism that is resistant to any coalition manipulation. No coalition should be able, by coordination, to influence the resulting permutation such that the members of the coalition will be the last ones to join the full coalition. For example, this means that no coalition of $n-1$ agents could force the single agent that is out of the coalition to go first.

We will use the simple cryptographic mechanism that allows an agent to encrypt a message using a private key, to send the encrypted message, and then to send the key such that the message can be unencrypted. Using these tools, each agent chooses a random permutation and a key, encrypts the permutation using the key, and broadcasts the encrypted message to all other agents. After he has received all encrypted messages, the agent broadcasts the key. Each agent unencrypts all messages using the associated keys. The consensus permutation is the combination of all permutations.

Each agent can make sure that each permutation has an equal chance of being chosen even if he assumes that the rest of the agents are all coordinating their permutations against him (i.e., trying to make him be the first). All he needs to do is to choose a *random* permutation. Since his permutation will also be combined into the final permutation, everything will be shuffled in a way that no one can predict.

Conclusions

We have considered the kinds of n-agent coordination mechanisms that can be used in Task Oriented Domains (TODs), when any sub-group of agents may engage in task exchange to the exclusion of others.

We presented a simple, efficient, symmetric, and individual rational Shapley Value-based coalition formation mechanism that uses cryptographic techniques for subadditive TODs. When the domain is also concave, the mechanism also satisfies coalition rationality.

Future research will consider non-subadditive TODs. It will also consider issues of incentive compatibility in multi-agent coalition formation, investigating mechanisms that can be employed when agents have partial information about the goals of other group members and can deceive one another about this private information.

References

Conry, S.; Meyer, R.; and Lesser, V. 1988. Multistage negotiation in distributed planning. In Bond, A., and Gasser, L., eds., *Readings in Distributed Artificial Intelligence*. San Mateo: Morgan Kaufmann Publishers, Inc. 367–384.

Ephrati, E., and Rosenschein, J. S. 1991. The Clarke Tax as a consensus mechanism among automated agents. In *Proceedings of the Ninth National Conference on Artificial Intelligence.*

Ephrati, E., and Rosenschein, J. S. 1993. Distributed consensus mechanisms for self-interested heterogeneous agents. In *First International Conference on Intelligent and Cooperative Information Systems*, 71–79.

Kahan, J. P., and Rapoport, A. 1984. *Theories of Coalition Formation.* London: Lawrence Erlbaum Associates.

Ketchpel, S. P. 1993. Coalition formation among autonomous agents. In *Pre-Proceedings of the Fifth European Workshop on Modeling Autonomous Agents in a Multi-Agent World.*

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*, 56–61.

Kraus, S.; Wilkenfeld, J.; and Zlotkin, G. 1995. Multiagent negotiation under time constraints. *Artificial Intelligence.* to appear. A preliminary version appeared in CS-TR-2975, University of Maryland.

Kreifelts, T., and von Martial, F. 1990. A negotiation framework for autonomous agents. In *Proceedings of the Second European Workshop on Modeling Autonomous Agents and Multi-Agent Worlds*, 169–182.

Kuwabara, K., and Lesser, V. R. 1989. Extended protocol for multistage negotiation. In *Proceedings of the Ninth Workshop on Distributed Artificial Intelligence*, 129–161.

Pollack, M. E., and Ringuette, M. 1990. Introducing the Tileworld: Experimentally evaluating agent architectures. In *Proceedings of the National Conference on Artificial Intelligence*, 183–189.

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

Rosenschein, J. S. 1993. Consenting agents: Negotiation mechanisms for multi-agent systems. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 792–799.

Shapley, L. S. 1971. Cores of convex games. *International Journal of Game Theory* 1:11–26.

Shapley, L. S. 1988. A value for n-Person games. In Roth, A. E., ed., *The Shapley Value.* Cambridge: Cambridge University Press. chapter 2, 31–40.

Shechory, O., and Kraus, S. 1993. Coalition formation among autonomous agents: Strategies and complexity. In *Pre-Proceedings of the Fifth European Workshop on Modeling Autonomous Agents in a Multi-Agent World.*

Smith, R. G. 1978. *A Framework for Problem Solving in a Distributed Processing Environment.* Ph.D. Dissertation, Stanford University.

Sycara, K. P. 1988. Resolving goal conflicts via negotiation. In *Proceedings of the Seventh National Conference on Artificial Intelligence*, 245–250.

Young, H. P. 1988. Individual contribution and just compensation. In Roth, A. E., ed., *The Shapley Value.* Cambridge: Cambridge University Press. chapter 17, 267–278.

Zlotkin, G., and Rosenschein, J. S. 1989. Negotiation and task sharing among autonomous agents in cooperative domains. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, 912–917.

Zlotkin, G., and Rosenschein, J. S. 1991. Cooperation and conflict resolution via negotiation among autonomous agents in noncooperative domains. *IEEE Transactions on Systems, Man, and Cybernetics* 21(6):1317–1324.

Zlotkin, G., and Rosenschein, J. S. 1993a. A domain theory for task oriented negotiation. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 416–422.

Zlotkin, G., and Rosenschein, J. S. 1993b. One, two, many: Coalitions in multi-agent systems. In *Pre-Proceedings of the Fifth European Workshop on Modeling Autonomous Agents in a Multi-Agent World.*