Designing Progressive MultiAgent Negotiations

Lyndon C. Lee
Intelligent Systems Research, ART,
BT Laboratories, Martlesham, UK.*
lyndon@computer.org

Abstract

The progressive negotiation model is one approach for organizing multiagent negotiations among autonomous utility-maximizing agents. The model specifies how agents are divided into a number of *subgroups*, and how *sub-coalitions* emerge after *sub-negotiations* take place among agents in each of these subgroups progressively. These sub-coalitions will then participate in the subsequent sub-negotiation until a *grand coalition* involving all agents is formed. While the effect of subgroup size on the solution quality (in terms of efficiency and Pareto optimality) has been investigated, little is known on how other design settings in progressive negotiations would affect the solution in terms of *time* and *individual utilities*. This paper reports on work we have done to analyze formally the relationships between the order of the sub-negotiations (i.e. *participation points*) and agents' self-utilities, and between the subgroup size and negotiation time.

Introduction

One major concern in distributed artificial intelligence (DAI) is how to build effective mechanisms to further coordination among autonomous utility-maximizing agents. The problem is non-trivial because it demands a practical scheme to exchange and to process information at the right abstraction level. Moreover, it must provide these inherently non-cooperative agents with the incentive of coordinating with one another.

We are concerned with the coordination of utility-maximizing agents that have conflicting goals, and we analyze how to use *negotiation* to further coordination. A model of progressive multiagent negotiation has been proposed [Lee, 1996a; 1996b], which allows agents to conduct negotiation by exchanging minimal information, and provides them with a risk function whereby agents could evaluate their relative losses of expected-utility due to concession with that as a result of conflict. Furthermore, the model has some proven desirable properties such as finiteness, and sufficient conditions for deadlock avoidance. Negotiations proceed by dividing agents into a number of *subgroups* which conduct *sub-negotiations* progressively. A *sub-coalition*, which emerges after each sub-negotiation, will then participate in the subsequent sub-negotiation. This process continues until a *grand coalition* is formed.

To design a progressive negotiation, the size of subgroups and the order (participation points) of the sub-negotiations have to be specified; but different parameters lead to different negotiation processes, resulting in different solutions in terms of utility and negotiation time. This paper investigates the impact of these parameters on solutions.

Related Work in DAI

Research into multiagent negotiation has many different strands in DAI. One approach explores how to model computationally human negotiation strategies [Sycara, 1988] or how to build sophisticated distributed search techniques for artificial systems inspired by human negotiations [Durfee and Montgomery, 1991; Lander and Lesser, 1993]. It aims to find ways to understand and improve (usually cooperative) agents' coordination abilities. Another approach attempts to predict the properties of negotiations under formal theoretical frameworks such as game-theoretical tools to study how agents should react in a given specific interaction [Zeuthen, 1930; Rosenschein and Zlotkin, 1994]. A third approach in negotiation is to model the negotiation processes computationally, and then to analyze the impact of the negotiation processes on the solution quality [Wooldridge *et al.*, 1996]. Another trend is to use evolutionary techniques [Matos *et al.*, 1998] or learning [Zeng and Sycara, 1997] to refine the negotiation processes, and to use various algorithms to form coalitions [Shehory and Kraus, 1998; Zlotkin and Rosenschein, 1994].

This paper is concerned with the coordination of utility-maximizing agents. Since agents are *self-interested*, any model using the social consensus metaphor like voting or social laws [Shoham and Tennenholtz, 1995] can be unstable[1]. Agents are also thought to be *rationally bounded*, thus the use of some prescriptive solutions (e.g. Nash solution) could be too expensive

[1] Nonetheless, see [Tennenholtz, 1998].

in the multiagent case [Binmore, 1992]. Therefore, the third approach is adopted. Unlike many coalition formation models where characteristic functions of all possible coalitions are assumed to be common knowledge and utility is divided among agents using Shapley value [Zlotkin and Rosenschein, 1993], our model is basically descriptive and analytical. Agents does not know about the value of any coalition, and a coalition is only formed when agents have agreed of how their tasks (hence utility) should be re-distributed (subsection).

Much as Shehory and Kraus distinguish a social level and a strategy level in their model [Shehory and Kraus, 1996] where in social level agent designers agree in advance the negotiation designs, we also focus on the strategy level since we aim to uncover the trends between certain negotiation designs and the solution quality. Since designers are also self-interested entities (otherwise their agents would not be), finding a consensus among them also involves negotiation. Our work therefore have implications to this meta-level negotiation (section).

To reduce complexity for rationally bounded agents, Sandholm and Lesser defines their characteristic function as a function of the computational time [Sandholm and Lesser, 1997]. Our model extends monotonic negotiations [Rosenschein and Zlotkin, 1994] and demands progressive negotiations: negotiations always converge provided that agents can compute at least one proposal at a time. Ketchpel also proposes a similar progressive coalition formation model where 2-agent coalitions are formed by mutual selection [Ketchpel, 1994]; whilst ours generalize the process to $n : n \geq 2$ agents. We also report the empirical results of different *orders* of coalition formations (section).

A Model of MultiAgent Negotiation

The model $\mathbb{PEA}$ is a utility-driven iterative process of proposal announcement, evaluation, and adjustment. When every agent is given a set of tasks to achieve, they negotiate at discrete times on how their tasks can be re-distributed so that everyone benefits. At each time instant, every agent chooses and announces a *proposal* of tasks re-distribution from their *negotiation sets* to other agents, where the proposal is a tentative proposition of how tasks are distributed among them. The *utility* of any proposal to an agent is defined by the difference in cost between achieving its original tasks and the tasks according to the proposed new task re-distribution.

When agents receive proposals announced by other agents, they are in a particular *negotiation state* defined by the set of proposals announced. An agent's negotiation set is defined as a set of all possible known proposals organized in *monotonic decreasing* order of the agent's self-utilities. If many proposals have the same self-utility to the agent, they are assumed to be organized *randomly* among themselves. During negotiation, proposals are chosen on the basis that they would provide agents with the greatest *expected-utility* at the time. If there is more than one such proposal, agents select one randomly.

It is worth noting that the model defines and assumes the existence of a negotiation set; nevertheless, when agents negotiate based on this model, they are *not* required to pre-enumerate a complete set. As it will become clear later, the convergence of the model using the monotonic negotiation protocol () only requires agents to be able to compute proposals in *monotonic decreasing* order of their self-utilities during the course of negotiation, even if agents can only produce one proposal at every time instant.

To encourage convergence and avoid *deadlock*[2], the monotonic *negotiation protocol* that specifies the set of proposals *eligible* for announcement at any time and state is used. If there are more than one eligible proposal, then agents use their *negotiation strategies* to determine the actual proposal for announcement. When a negotiation state occurs where all proposals are identical in terms of utility distribution, the negotiation ends with any proposal in it as the *solution*.

Monotonic Negotiation Protocol

It is found that if at any time at least one agent will make a *concession* by either *accepting* another agent's proposal or *conceding* a new (never announced) proposal demanding no more utility to itself than any previous proposal, deadlock will never occur [Lee, 1996b]. The *monotonic negotiation protocol* (MNP) is thus defined to specify the eligibility of proposals for announcement at any time and state. It also determines which agents should make the next concession by means of a multiagent risk function extended based on Zeuthen's risk function [Zeuthen, 1930]:

$$R_i = \prod_{\substack{j=1 \\ j\neq i}}^{n} r_{i,j};$$

$$r_{i,j} = \begin{cases} 1 & \text{if } U_i(P_{i,t}) = U_i(P_c) \\ 1 & \text{if } U_i(P_{i,t}) \leq U_i(P_{j,t}) \\ 1 & \text{if } V(P_{i,t}) \geq V(P_{k,t}) \\ & \text{where } 1 \leq k \leq n,\ k \neq i \\ \frac{U_i(P_{i,t})-U_i(P_{j,t})}{U_i(P_{i,t})-U_i(P_c)} & \text{if } P_{i,t} \in NS_{i,t} \text{ and} \\ & P_{j,t} \neq P_{i,t-1}.^3 \end{cases}$$

The semantics of the four conditional equations, respectively, corresponds to four axioms: (1) individual rationality (prefer any proposal to conflict); (2) utility maximizing; (3) local-joint efficiency assumption (always

[2]Deadlock occurs when two identical negotiation states appear. The negotiation will terminate and agents must then pursue their original tasks individually.

[3]$P_{i,t}$ is the proposal of agent i at any time t, while P_c is called the conflict proposal (i.e. agents have no deal and pursue their individual tasks.). Furthermore, $U_j(P_{i,t})$ refers to the *utility* of $P_{i,t}$, i.e. the utility of proposal $P_{i,t}$ to agent j where i and j are not necessarily distinct. Let n be the number of agents, $V(P_{i,t}) = \langle v_i, \ldots, v_n \rangle$ refers to the *utility vector* of $P_{i,t}$ where $U_j(P_{i,t}) = v_j$. For any two vectors V, V', $V \geq V'$ iff $v_i \in V, v'_i \in V', v_i \geq v'_i, \forall i \in (1, n)$. $NS_{i,t}$ is the negotiation set of agent i at time t.

prefers the most efficient proposal); (4) Zeuthen's risk formula. This risk function allows agents to determine, independently but coherently, their *risk limits* by comparing their overall relative losses of expected-utility due to concession with that resulting from conflict. The lower the limit, the greater loss an agent will suffer should a conflict occur. If agents are assumed to be *rational* in such a way that they will always act to maximize their expected-utility, those agents which have the minimum risk limit should make the next concession. On the other hand, if an agent does not have the minimum risk limit, it may either accept other's proposal (if the offered utility is at least as great as that of its own proposal) or insist on its own proposal.

Progressive Negotiation

Progressive PEA replaces the grand negotiation by a number of *sub-negotiations* in a number of *stages*. Agents are first divided into a number of *subgroups* where each agent will participate in only one subgroup. These subgroups will perform sub-negotiation progressively; on completion the *sub-coalition* formed will then participate in the sub-negotiation at the subsequent stage as if it were an individual agent called the *sc-agent*. The utility obtained by an sc-agent from a sub-negotiation will be *proportionally distributed* to every individual agent (and sc-agent) in the subgroup from which the sc-agent is formed, and the distribution is based on the utility distribution agreed during its formation [Lee, 1996b]. The progressive negotiation ends with a *grand solution* when a *grand coalition* containing all the agents is formed[4].

In this paper, for simplicity, we confine our analysis to a specific class of progressive negotiations where only *one* sub-negotiation occurs at any stage. Moreover, every sub-coalition formed is assumed to participate in the sub-negotiation at the *next* stage.

Let us explain how a progressive negotiation proceeds by a simple example. Suppose in a progressive negotiation agents i, j, and k are divided into two subgroups $A^1 = \{i, j\}$ and $A^2 = \{a^1, k\}$ where a^1 is the sc-agent formed when the sub-negotiation of A^1 is finished. This progressive negotiation has two stages and each subgroup has size 2. Let the solution of the two sub-negotiations be P_{A^1} and P_{A^2}, respectively. After the first sub-negotiation, the utility allocated to i will be $U_i(P_{A^1})$. After the second sub-negotiation, i's utility share from P_{A^2} will then be

$$U_i(P_{A^2}) = U_{a^1}(P_{A^2}) \times \frac{U_i(P_{A^1})}{\sum_{x \in A^1} U_x(P_{A^1})},$$

and i's aggregate utility is therefore $U_i(P_{A^1}) + U_i(P_{A^2})$.

Analysis

To design a progressive negotiation, three factors are required to determine: (1) the *subgroup size*, i.e. the number of agents in sub-negotiations; (2) how these subgroups are organized into successive stages; and (3) how those agents are allocated to these subgroups, i.e. *participation points*. This paper investigates the impact of different participation points and subgroup sizes.

The Effect of Participation Points

According to the model, the *aggregate utility* that any agent will obtain in a progressive negotiation is the sum of all the utility obtained in each sub-negotiation. If any agent i, instead of randomly allocated to sub-negotiation, is fixedly allocated to a sub-negotiation at an *early* stage (i.e. *lower* participation point), i will participate in *more* sub-negotiations than if i were allocated to a sub-negotiation at a *later* stage (i.e. *higher* participation point). As a result, the hypothesis is that

Hypothesis 1 *The average aggregate utility to an agent decreases with its participation point.*

The Effect of SubGroup Sizes

Before we can analyze the effect of subgroup size on negotiation cycle (i.e. the number of iterations required for a negotiation, which is used to estimate the actual *negotiation time*), we need first to establish the fact that the MNP negotiation will converge. Let us call negotiation state at any time the *pre-solution state* if a solution will be reached at the next time instant.

Theorem 1 *MNP negotiations will always terminate with a solution in finite time if* $\forall i, j \in A, NS_{i,1} = NS_{j,1}$[5].

Proof: (Sketch) Let

$$\mathfrak{S} = \{\mathrm{S} \mid \forall i \in A, P_i \in \mathrm{S}, \exists P_j \in \mathrm{S}, U_i(P_i) \le U_i(P_j)\}$$

where S is the negotiation state (i.e. set of announced proposals), A is the set of agents, P is a proposal and $U(P)$ is the utility of a proposal. First we show that $\mathfrak{S}$ is the *minimal* set of pre-solution negotiation states because in each member at least one proposal is efficient. It is also easy to see that $\mathfrak{S}$ is non-empty if $\forall i, j \in A$, $\mathrm{NS}_{i,1} = \mathrm{NS}_{j,1}$.

Since the number of agents is finite, and so are their negotiation sets, and MNP demands that at least one agent must announce a new proposal at every time instant, therefore it will only take finite time for a member in $\mathfrak{S}$ to be announced, after which the negotiation will terminate with a solution.

Theorem 2 *Provided that (i) the MNP protocol is the only constraint; (ii) only one proposal is different between any two consecutive negotiation states; (iii)* $\forall i, j \in A, NS_{i,1} = NS_{j,1}$ *and* $n > 2$, *then*

$$t_{\max} = n \mid NS_{i,1} \mid -2$$

where $t_{\max}$ *is the maximum negotiation time.*

[4]Formal specifications can be found in [Lee, 1996a].

[5]It means that all agents have identical negotiation sets. Detailed proofs of theorems in this paper can be found in [Lee, 1996a]

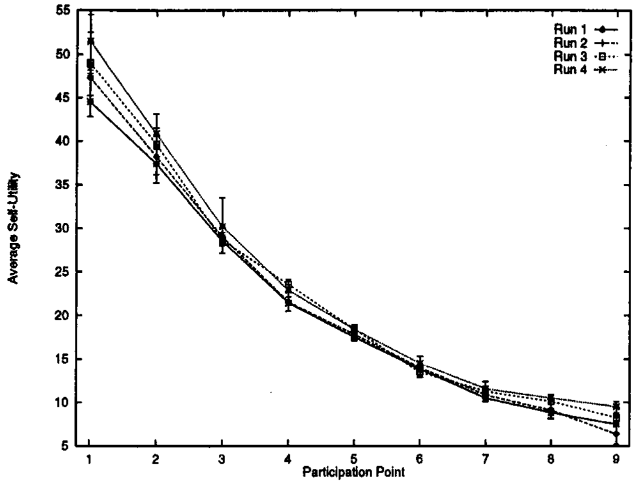

Figure 1: Self-utility as a function of participation point

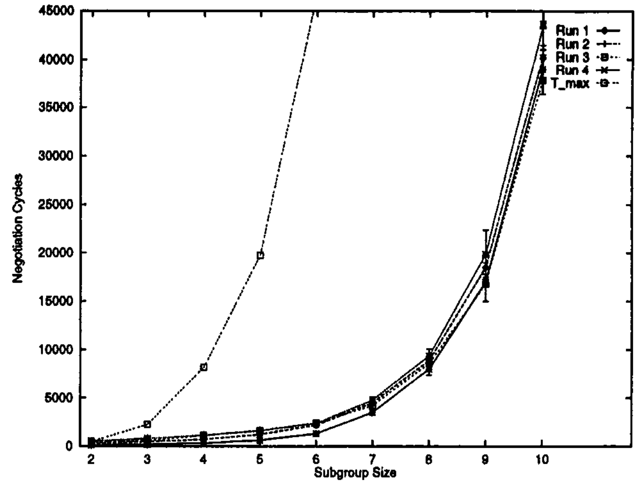

Figure 2: Negotiation cycle as a function of subgroup size

Proof: (Sketch) The result can be obtained from the following three cases: (1) the maximum number of time instants required for $(n-1)$ agents to announce proposals with the minimum utility – the utility they would have gained had they not engaged in any negotiation; (2) the maximum number of time instants that the last agent to produce a proposal with utility just above the minimum; and (3) the maximum number of time instants that during the process of (1) and (2) these agents could accept others' proposal but without resulting in a solution.

Since a pre-solution state cannot be reached unless agents' self-utility levels have been reduced through concession, which involves discarding a certain number of proposals, whilst the number of concessions required for reaching those self-utility levels depends on the *sizes* of the negotiation sets which in turn increases with the subgroup size [Lee, 1996a]. Thus,

Hypothesis 2 *The average negotiation time increases with subgroup size m.*

Empirical Results

Methodology

Four runs of simulation experiments were carried out based on the reformulated multiagent Tileworld domain [Pollack and Ringuette, 1990]. In each run, ten agents (initially situated at the four corners randomly) are given a set of 5-10 random holes to visit in a 2-dimensional 10 by 10 grid-cells. The only cost to agents is the traveling cost that is proportional with the number of steps agents move. Agents negotiate with one another to re-distribute their holes so as to reduce their traveling bills. The utility is defined as the the cost difference between the new set of holes (after re-distribution) and the original set of holes.

Participation Points

Here, we assume all sub-negotiations have size 2. In each run, 10000 simulations where participation points are randomly assigned to agents have been carried out[6]. Then we calculated the average and variance of the aggregate utility of each agent at a specific participation point. It is found that whichever agent is chosen as the subject of investigation, the result presented a consistent trend where average aggregate utility decreases with participation points, as Figure 1 shows. Furthermore, the variance of each aggregate utilities are also included in the graph, whose modest intervals suggest that the trends are likely to hold irrespective of other random factors such as how agents are allocated to subgroups. Hence hypothesis 1 is verified.

Subgroup Sizes

Here, a progressive negotiation of subgroup size m means as many sub-negotiations as possible has size m[7]. Each run has nine sets of 1000 simulations where each set corresponds to a particular subgroup size. In each set agents are allocated to subgroups randomly.

The average and variance of the number of negotiation cycles are then measured. Figure 2 shows that the average number of negotiation time (measured by the number of negotiation cycles) increases with subgroup sizes, which supports the hypothesis. Similarly, the relative small variance in number of cycles with respect to the mean indicates the trend is unlikely to be affected by other random factors. Note also that the actual number

[6]To perform an exhaustive participation point assignment is an extremely expensive process. For n agents and size m sub-groups, it involves C^n_m different allocations. In fact, we did perform some exhaustive simulations on runs with smaller number of agents, and compared them with results from random simulations. It is found that both of the results have consistent trends.

[7]For example, for a negotiation involving 10 agents has subgroup size 3, it means that the first four sub-negotiations all have 3 agents while the last sub-negotiation has two agents.

of negotiation cycles is well below the theoretical maximum bound as predicted in [Lee, 1996a].

An Opposite Trend

In all the above experiments, the *concession rate*, defined by the rate at which agents decrease their self-utility during concession, has so far been chosen to be the smallest. That is to say, agents will only choose a proposal with less self-utility than the existing one when there is no more proposal with equal self-utility (MNP prohibits agents to choose any proposal with greater utility than before). This also explain why the negotiation process is so lengthy.

When we repeated the experiment using the *fastest* concession rate, the trend for the negotiation time verse subgroup size was expected to remain unchanged although the overall scale of negotiation time should be lowered. However, we found that the negotiation time *decreases* with the subgroup size, as shown in figure 3.

Further investigation reveals that there are probably two factors that contribute to this surprising result. First, the size of the negotiation set at the *fast* concession rate becomes time-insignificant. Secondly, the *number of stages* in the progressive negotiation becomes the more time-dominating factor.

When agents use very fast concession rates, since only one (or a few) proposals from each *self-utility partition* of their negotiation sets[8] will be chosen for announcement, it will not take much time for a negotiation to reach a particular state where all agents have announced proposals with the minimum utility. With the fastest concession rate, it will only require L_i number of concessions for any agent i to reach minimum utility proposal, where L_i is the *number of utility levels* to i in its *NS*. A corollary from theorem 1 and 2 states that a solution must occur when every agent has reached its minimum utility proposals [Lee, 1996a], hence in this case $t_{\max} = \sum_{i=1}^{n} L_i$.

Although the size of *NS* increases with subgroup size, [Lee, 1996a] has found that that L_i does not vary very much as subgroup size rises. Hence, with fast concession rate, the negotiation time required, whatever the subgroup size m, is relatively of the same order of magnitude.

However, there are a fewer *number of stages* in any progressive negotiation with a larger m than that with a smaller m. For example, for five-agent progressive negotiations, if $m = 5$, there is only one stage; whereas there are two stages if $m = 4$, or $m = 3$, and four stages if $m = 2$. Since the time required for each sub-negotiation is relatively constant at fast concession rate, this implies the overall negotiation time now depends on the number of sub-negotiations to perform – the more number of sub-negotiations the longer the negotiation time. This explains why the negotiation time decreases with subgroup size at fast concession rate.

[8]Every self-utility partition of a negotiation set refers to the set of proposals that has the same self-utility. See Section.

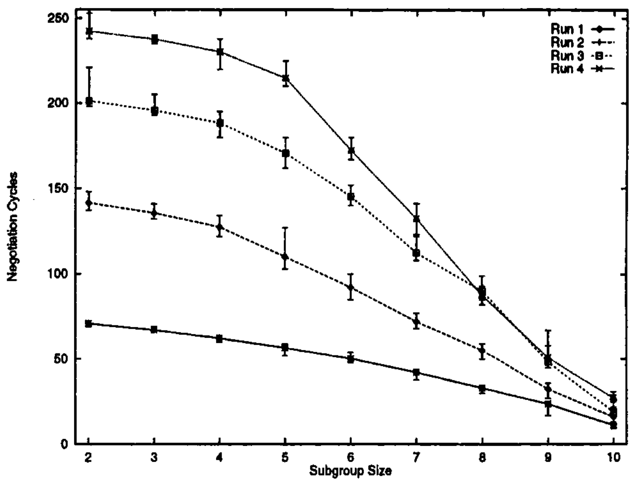

Figure 3: Negotiation cycle as a function of subgroup size

Conclusion

We have presented a model of progressive multiagent negotiation, which guarantees the convergence of the negotiation processes. It also provides a risk function that allows agents to determine *independently* but *coherently* which agents should make the next concession on the basis of relative loss of expected-utility due to concession with that as a result of conflict. This model does not assume unbounded rationality: provided an agent can compute new proposals in monotonic decreasing order of self-utility during the course of negotiation (even if only one proposal at every time instant), the model guarantees the negotiation to terminate with a solution. Furthermore, formal analysis and an empirical study suggest that there are consistent trends between various design choices in the negotiation settings and solution quality.

This paper has analyzed the effect of participation points on self-utility and of subgroup sizes on negotiation time. The relations between subgroup sizes, solution efficiency and Pareto optimality have already been reported: it is found that as the hierarchical levels of progressive negotiations (i.e. subgroup size) increases, the solution quality in terms of efficiency and Pareto optimality would be worse off and computational cost increases [Lee, 1996b]. Here we found that the negotiation time would also increase. Nevertheless, under certain conditions such as fast concession, the negotiation time decreases with subgroup size.

This paper also showed that, on the other hand, the point at which an agent engages in progressive negotiation tends to be proportional to its resulting aggregate self-utility, while at the expense of computational cost since it is required to perform more number of sub-negotiations. These results provide valuable information as to how to design a progressive multiagent negotiation so that certain long term expectations on solution quality such as negotiation time, Pareto optimality, as well as self-utility, and resources considerations like computational cost can be intelligently compromised.

Our next step is to look into the inter-relationships among negotiation time, subgroup size, and concession rate: can we design negotiations with a particular (average) negotiation time? We also want to relax our initial constraint on the progressive negotiations: i.e. only *one* sub-negotiation at each stage. Further investigation is also needed to find out the impact of different utility propagation functions and various risk functions based on different rationality.

Our long term goal is, by thoroughly investigating how different negotiation processes affect the outcome of negotiations in several classes of negotiation models (where ℙ𝔼𝔸 focuses on superadditive domain), we could gain a better understanding of how to build negotiation models with analyzable as well as predictable properties.

References

[Binmore, 1992] K. Binmore. *Fun and Games : A Text on Game Theory.* D.C. Heath, 1992.

[Durfee and Montgomery, 1991] E.H. Durfee and T.A. Montgomery. Coordination as distributed search in a hierarchical behavior space. *IEEE Trans. SMC*, 21(6):1363–1378, 1991.

[Ketchpel, 1994] S. Ketchpel. Forming coalitions in the face of uncertain rewards. In *Proceedings of Twelfth National conference on Artificial Intelligence*, Seattle, Washington, August 1994.

[Lander and Lesser, 1993] S.E. Lander and V.R. Lesser. Understanding the role of negotiation in distributed search among heterogeneous agents. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, Chambery, France, 1993.

[Lee, 1996a] L.C. Lee. *Negotiation Strategies and their Effect in a Model of Multiagent Negotiation.* PhD thesis, University of Essex, July 1996.

[Lee, 1996b] L.C. Lee. Progressive multi-agent negotiation. In *Proceedings of the Second International Conference on Multiagent Systems*, Kyoto, Japan, 9-13, December 1996.

[Matos *et al.*, 1998] N. Matos, C. Sierra, and N.R. Jennings. Determining successful negotiation strategies: An evolutionary approach. In *Proceedings of the Third International Conference on Multiagent Systems*, pages 182–189, 1998.

[Pollack and Ringuette, 1990] M.E. Pollack and M. Ringuette. Introducing the tileworld: Experimentally evaluating agent architectures. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, pages 183–189, July 1990.

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

[Sandholm and Lesser, 1997] T. Sandholm and V.R. Lesser. Coalitions among computationally bounded agents. *Artificial Intelligence*, 94(1-2):99–137, 1997.

[Shehory and Kraus, 1996] O. Shehory and S. Kraus. A kernel-oriented model for coalition-formation in general environments: Implementation and results. In *Proceedings of the Thorteenth National conference on Artificial Intelligence*, pages 134–140, 1996.

[Shehory and Kraus, 1998] O. Shehory and S. Kraus. Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 101(1-2):165–200, 1998.

[Shoham and Tennenholtz, 1995] Y. Shoham and M. Tennenholtz. On social laws for artificial agent societies: Off-line design. *Artificial Intelligence*, 73(1-2):231–252, February 1995.

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

[Tennenholtz, 1998] M. Tennenholtz. On stable social laws and qualitative equilibria. *Artificial Intelligence*, 102(1):1–20, 1998.

[Wooldridge *et al.*, 1996] M.J. Wooldridge, S. Bussmann, and M. Klosterberg. Production sequencing as negotiation. In *Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology*, pages 709–726, London, 22-24 April 1996.

[Zeng and Sycara, 1997] D. Zeng and K. Sycara. Benefits of learning in negotiation. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*, 36-41, 1997.

[Zeuthen, 1930] F. Zeuthen. *Problems of Monopoly and Economic Warfare.* G. Routledge & Sons, 1930.

[Zlotkin and Rosenschein, 1993] G. Zlotkin and J.S. Rosenschein. One, two, many: Coalitions in multiagent systems. In *Proceedings of the Fifth European Workshop on Modelling Autonomous Agents in Multi-Agent World*, Neuchatel, Switzerland, August 1993.

[Zlotkin and Rosenschein, 1994] G. Zlotkin and J.S. Rosenschein. Coalition, cryptography, and stability: Mechanisms for coalition formation in task oriented domains. In *Proceedings of Twelfth National conference on Artificial Intelligence*, Seattle, Washington, August 1994.