

Automated Negotiation and Decision Making in Multiagent Environments*

Sarit Kraus

¹ Dept. of Mathematics and Computer Science
Bar-Ilan University

Ramat-Gan, 52900 Israel

² Institute for Advanced Computer Studies
University of Maryland, College Park, MD 20742, USA

sarit@umiacs.umd.edu

<http://www.umiacs.umd.edu/users/sarit>

Abstract. This paper presents some of the key techniques for reaching agreements in multi-agent environments. It discusses game-theory and economics based techniques: strategic negotiation, auctions, coalition formation, market-oriented programming and contracting. It also presents logical based mechanisms for argumentations. The focus of the survey is on negotiation of self-interested agents, but several mechanisms for cooperative agents who need to resolve conflicts that arise from conflicting beliefs about different aspects of their environment are also mentioned. For space reasons, we couldn't cover all the relevant works, and the papers that are mentioned only demonstrate the possible approaches. We present some of the properties of the approaches using our own previous work.

1 Introduction

Negotiation has been a subject of central interest in multi-agent systems, as it has been in economics and political science. The word has been used in a variety of ways, though in general it refers to communication processes that further coordination and cooperation. Negotiations can be used to resolve conflicts in a wide variety of multi-agent domains [28]. Examples of such applications include conflicts over the usage of joint resources or task assignments, conflicts concerning document allocation in multi-server environments and conflicts between a buyer and a seller in electronic commerce.

When building an autonomous agent which is capable of flexible and sophisticated negotiation, the main questions that should be considered are: (i) what

* This work is supported in part by NSF under grant No. IIS-9820657. This paper is based on chapters in S. Kraus. *Strategic Negotiation in Multi-Agent Environments*, MIT Press, Cambridge, USA, 2001 and on S. Kraus, K. Sycara and A. Evenchik. *Reaching agreements through argumentation: a logical model and implementation*, *Artificial Intelligence Journal*, 104(1-2):1-69, 1998.

negotiation protocol will be used? (ii) what reasoning model, decision making procedures and strategies will the agents employ?

Several protocols for auctions, strategic negotiation and coalition formation are considered and we discuss their applicability in various multi-agent domains. We will present formal models for agent reasoning and we will discuss methods for identifying strategies for agents interacting using a specific protocol.

Evaluation of the results of multi-agent protocols is not an easy task. Since the agents are self-interested, when saying, for example, a “negotiation was successful” the question “successful for whom?” must be asked, since each agent is concerned only about its own benefits or losses from the resolution of the negotiation. Nevertheless, there are certain parameters that can be used to evaluate different protocols.

Negotiation Time: Negotiations which end without delay are preferred over negotiations which are time-consuming.

It is assumed that a delay in reaching an agreement causes an increase in the cost of communication and computation time spent on the negotiation.

We want to prevent the agents from spending too much time on negotiations resulting in deviation from their timetables for satisfying their goals.

Efficiency: An efficient outcome of the negotiations is preferred. In other words, an outcome that increases the number of agents which will be satisfied by the negotiation results and the agents’ satisfaction levels from the negotiation results.

Thus, it is preferred that the agents reach *Pareto optimal* agreements¹ In addition, if there is an agreement that is better for all the agents than opting out, then it is preferred that the negotiations will end with an agreement.

Simplicity: Negotiation processes that are simple and efficient are better than complex processes. Being a “simple strategy” means that it is feasible to build it into an automated agent. A “simple strategy” also presumes that an agent will be able to compute the strategy in a reasonable amount of time.

Stability: A set of negotiation strategies are stable if, given that all the other agents included in the set are following their strategies, it is beneficial to an agent to follow its strategy too. Negotiation protocols which have stable strategies are more useful in multiagent environments than protocols which are unstable. If there are stable strategies, we can recommend to **all** agent designers to build the relevant strategies into their agents. No designer will benefit by building agents that use any other strategy.

Money transfer: Money transfer may be used to resolve conflicts. For example, a server may “sell” a data item to another server when relocating this item. This can be done by providing the agents with a monetary system and with a mechanism for secure payments. Since maintaining such a monetary system requires resources and efforts, negotiation protocols that do not require money transfers are preferred.

¹ An agreement is Pareto optimal if there is no other agreement that dominates it, i.e., there is no other agreement that is better for some of the agents and not worse for the others.

The remainder of this paper is structured as follows. In the next section we will present a short survey of negotiation approaches in Distributed Artificial Intelligence (DAI) and in social sciences. Then we will discuss the strategic-negotiation model (section 2.3) which is based on the game-theory model of bargaining with alternating offers. Auctions are discussed in section 3, and market-oriented programming is briefly discussed in section 4. Agents can also cooperate by forming coalitions. This form of cooperation is presented in section 5. Contracting which is another form of reaching cooperation is surveyed in section 6. Finally, logical approaches to negotiation are presented in section 7.

2 Negotiation Models

We will first present a short survey of various negotiation approaches in Distributed Artificial Intelligence. Then, we will describe the two main approaches to negotiations in the social sciences and will demonstrate the application of one of the approaches to multiagent systems.

2.1 Negotiation Models in DAI

Negotiations were used in DAI both in Distributed Problem Solving (DPS) where the agents are cooperative and in Multiagent Systems (MA) where the agents are self-interested. Several works in DPS use negotiation for distributed planning and distributed search for possible solutions for hard problems. For example, Conry et al. [10] suggest multi-stage negotiation to solve distributed constraint satisfaction problems when no central planner exists. Moehlman and Lesser [52] use negotiation as a tool for distributed planning: each agent has certain important constraints, and it tries to find a feasible solution using a negotiation process. They applied this approach in the Phoenix fireman array. Lander and Lesser [44] use a negotiation search, which is a multi-stage negotiation as a means of cooperation while searching and solving conflicts among the agents.

For the MA environments, Rosenschein and Zlotkin [73] identified three distinct domains where negotiation is applicable and found a different strategy for each domain: (i) Task-Oriented Domain: Finding ways in which agents can negotiate to come to an agreement, and allocating their tasks in a way that is beneficial to everyone; (ii) State-Oriented Domain: Finding actions which change the state of the “world” and serve the agents’ goals; and (iii) Worth-Oriented Domain: Same as (ii) above, but, in this domain, the decision is taken according to the maximum utility the agents gain from the states.

Sycara [100,99] presented a model of negotiation that combines case-based reasoning and optimization of multi-attribute utilities. In her work agents try to influence the goals and intentions of their opponents. Kraus and Lehmann [39] developed an automated Diplomacy player that negotiates and plays well in actual games against human players. Sierra et al. [94] present a model of negotiation for autonomous agents to reach agreements about the provision of service by one agent to another. Their model defines a range of strategies and

tactics, distilled from intuition about good behavioral practice in human negotiation, that agents can employ to generate offers and evaluate proposals. Zeng and Sycara [116] consider negotiation in a marketing environment with a learning process in which the buyer and the seller update their beliefs about the opponent's reservation price² using the Bayesian rule. Sandholm and Lesser [85] discuss issues, such as levels of commitment, that arise in automated negotiation among self-interested agents whose rationality is bounded by computational complexity.

2.2 Negotiation Approaches in the Social Sciences

In social sciences there are two main approaches to the development of theories relating to negotiation. The first approach is the formal theory of bargaining e.g., [75,62], constituting a formal, game-theoretic approach that provides clear analyses of various situations and precise results concerning the strategy a negotiator should choose. However, this approach can only be applied to situations satisfying very restricted assumptions. In particular, this approach assumes that the agents are acting rationally, have large computation capabilities and follow strict negotiation protocols.

The second approach, which we refer to as the negotiation guides approach, comprises informal theories which attempt to identify possible general beneficial strategies for a negotiator. The works based on this approach advise a negotiator how to behave in order to reach beneficial results in a negotiation (see, for example, [68,13,11,32,29,22]). These negotiation guides do not presuppose the strong restrictions and assumptions presented in the game-theoretic models. Applying these methods to automated systems is more difficult than using the first approach, since there are neither formal theories nor strategies that can be used.³ In the next section we demonstrate the application of the formal approach to multiagent systems.

2.3 Strategic Negotiation

The strategic-negotiation model is based on Rubinstein's model of alternating offers [77]. In the strategic model there are N agents, **Agents** = $\{A_1, \dots, A_N\}$. The agents need to reach an agreement on a given issue. It is assumed that the agents can take actions in the negotiation only at certain times in the set $\mathcal{T} = \{0, 1, 2, \dots\}$ that are determined in advance and are known to the agents.

In each period $t \in \mathcal{T}$ of the negotiation, if the negotiation has not terminated earlier, an agent whose turn it is to make an offer at time t , will suggest a possible

² The reservation price of the seller is the price below which the seller refuses to sell. The reservation price of the buyer is the price above which the buyer refuses to buy.

³ These methods can be used in domains where people interact with each other and with automated systems, and situations where automated systems interact in environments without predefined regulations. These informal models can serve as guides for the development of negotiation heuristics [39] or as a basis for the development of a logical negotiation model [40].

agreement (with respect to the specific negotiation issue), and each of the other agents may either accept the offer (choose Yes), reject it (choose No), or opt out of the negotiation (choose Opt). If an offer is accepted by all the agents (i.e., all of them choose Yes), then the negotiation ends, and this offer is implemented. If at least one of the agents opts out of the negotiation, then the negotiation ends and a conflictual outcome results. If no agent has chosen “Opt,” but at least one of the agents has rejected the offer, the negotiation proceeds to period $t + 1$, and the next agent makes a counteroffer, the other agents respond, and so on. We assume that an agent responding to an offer is not informed of the other responses during the current negotiation period. We call this protocol a *simultaneous response* protocol.⁴ $j(t)$ will denote the agent that makes an offer at time period t . The following example demonstrate these notions.

Example 1 (Data Allocation in Large Databases). There are several information servers, in different geographical areas. Each server stores data, which has to be accessible by clients not only from its geographical area but also from other areas. The topics of interest of each client change dynamically over time, and the set of clients may also change over time. Periodically, new data arrive at the system, and have to be located at one of the servers in the distributed system.

Each server is independent and has its own commercial interests. The servers would like to cooperate with each other in order to make more information available to their clients. Since each server has its own preferences regarding possible data allocations, its interests may conflict with the interests of some of the other servers.

A specific example of a distributed information system is the Data and Information System component of the Earth Observing System (EOSDIS) of NASA [56]. It is a distributed system which supports archival data and distribution of data at multiple and independent data centers (called DAACs). The current policy for data allocation in NASA is static: each DAAC specializes in some topics. When new data arrive at a DAAC, the DAAC checks if the data is relevant to one of its topics, and, if so, it uses criteria, such as storage cost, to determine whether or not to accept the data and store them in its database. The DAAC communicates with other DAACs in those instances in which the data item encompasses the topics of multiple DAACs, or when a data item presented to one DAAC is clearly in the jurisdiction of another DAAC, and then a discussion takes place among the relevant DAAC managers. However, this approach does not take into consideration the location of the information clients, and this may cause delays and transmission costs if data items are stored far from their potential users. Moreover, this method can cause rejection of data items if they do not fall within the criteria of any DAAC, or if they fall under the criteria of a DAAC which cannot support this new product because of budgetary problems.

In this example the agents negotiate to reach an agreement that specifies the location of *all* the relevant data items. In the first time period, the first server

⁴ A sequential protocol is considered in [14]. In this protocol an agent responding to an offer is informed of the responses of the preceding agents (assuming that the agents are arranged in a specific order).

offers an allocation, and the other agents either accept the offer, reject it or opt out of the negotiation. If an offer is accepted by all the agents, then the negotiation ends and the proposed allocation is implemented. If at least one of the agents opts out of the negotiation, then a predefined *conflict allocation* is implemented, as described in [87]. If no agent has chosen “Opt,” but at least one of the agents has rejected the offer, the negotiation proceeds to the next time period and another agent proposes an allocation, the other agents respond, and so on.

In the strategic-negotiation model there are no rules which bind the agents to any specific strategy. We do not make any assumptions about the offers the agents make during the negotiation. In particular, the agents are not bound to any previous offers that have been made. After an offer is rejected, an agent whose turn it is to suggest a new offer can decide whether to make the same offer again, or to propose a new offer. The protocol only provides a framework for the negotiation process and specifies the termination condition, but there is no limit on the number of periods.

A fair and reasonable method for deciding on the order in which agents will make offers is to arrange them randomly in a specific order before the negotiation begins.⁵ That is, the agents will be denoted randomly A_1, \dots, A_N . At each time t , $j(t)$ will be A_i where i is equal to $(t \bmod N) + 1$.

The set of possible agreements is denoted \mathcal{S} . An outcome of the negotiation may be that an agreement $s \in \mathcal{S}$ will be reached at time $t \in \mathcal{T}$. This outcome is denoted by a pair (s, t) . When one of the agents opts out of the negotiations at time period $t \in \mathcal{T}$, the outcome is denoted (\mathbf{Opt}, t) . For example, in the data allocation scenario (example 1), an agreement is an allocation which assigns each data item to one of the servers. In this case \mathcal{S} is the set of all possible allocations. The symbol **Disagreement** indicates a perpetual disagreement, i.e., the negotiation continues forever without reaching an agreement and without any of the agents opting out.

The agents’ *time preferences* and the preferences between agreements and opting out are the driving force of the model. They will influence the outcome of the negotiation. In particular, agents will not reach an agreement which is not at least as good as opting out for all of them. Otherwise, the agent which prefers opting out over the agreement, will opt out.

Negotiation Strategies An agent’s negotiation strategy specifies for the agent what to do next, for each sequence of offers $s_0, s_2, s_3, \dots, s_t$. In other words, for the agent whose turn it is to make an offer, it specifies which offer to make next. That is, it indicates to the agent which offer to make at $t + 1$, if in periods 0 until t the offers s_0, \dots, s_t had been made and were rejected by at least one of the agents, but none of them has opted out. Similarly, in time periods when it is the agent’s turn to respond to an offer, the strategy specifies whether to accept

⁵ A distributed algorithm for randomly ordering the agents can be based on the methods of [6].

the offer, reject it or opt out of the negotiation. A strategy profile is a collection of strategies, one for each agent [63].

Subgame Perfect Equilibria The main question is how a rational agent will choose its negotiation strategy. A useful notion is the Nash Equilibrium [57,47] which is defined as follows:

Definition 1 (Nash Equilibrium). *A strategy profile $F = \{f_1, \dots, f_N\}$ is a Nash equilibrium of a model of alternating offers, if each agent A_i does not have a different strategy yielding an outcome that it prefers to that generated when it chooses f_i , given that every other agent A_j chooses f_j . Briefly, no agent can profitably deviate, given the actions of the other agents.*

This means, that if all the agents use the strategies specified for them in the strategy profile of the Nash equilibrium, then no agent has a motivation to deviate and use another strategy. However, the use of Nash equilibrium in a model of alternating-offers leads to an absurd Nash equilibria [102]: an agent may use a threat that would not be carried out if the agent were put in the position to do so, since the threat move would give the agent lower payoff than it would get by not taking the threatened action. This is because Nash equilibrium strategies may be in equilibrium only in the beginning of the negotiation, but may be unstable in intermediate stages. The concept of subgame perfect equilibrium (SPE) [63], which is a stronger concept, is defined in the following definition and will be used in order to analyze the negotiation.

Definition 2 (Subgame perfect equilibrium:). *A strategy profile is a subgame perfect equilibrium of a model of alternating offers if the strategy profile induced in every subgame is a Nash equilibrium of that subgame.*

This means that at any step of the negotiation process, no matter what the history is, no agent has a motivation to deviate and use any strategy other than that defined in the strategy profile.

In situations of incomplete information there is no proper subgame. The *sequential equilibrium* [42], which takes the beliefs of the agents into consideration, can be used in the incomplete information situations.

Example 2. The application of the strategic-negotiation model to the data allocation problem of example 1 was presented in [3]. Using this model, the servers have simple and stable negotiation strategies that result in efficient agreements without delays. It was shown that these methods yield better results than the static allocation policy currently used in EOSDIS (see Example 1).

In particular, it was shown that when servers negotiate to reach an agreement on the allocation of data items and they have complete information various agreements can be reached. It was proved that for any possible allocation of the data items that is not worse for any of the agents than opting out, there is a set of stable strategies (one for each server) which leads to this outcome. That is, suppose *alloc* is a specific allocation which all the servers prefer than opting

out of the negotiation. A strategy for each of the servers can be designed such that the strategy profile will be an equilibrium. If the servers use this strategy profile, the negotiations will end at the first time period of the negotiation with the agreement *alloc*.

The details of the allocations that are not worse for any of the agents over opting out depend on the specific settings of the environment in a given negotiation session. Thus, there is no way to identify these allocations in advance. In addition, there are usually several allocations which are not worse for any of the agents than opting out. Finding all these allocations is intractable. In addition, after identifying these allocations the servers should agree upon one of them as the basis for the negotiation.⁶ Of course, each of the servers may prefer a different allocation because it may yield a higher utility. A mechanism by which the servers can choose one of these profiles of stable strategies is presented. It leads to satisfactory results for all of the servers. In this mechanism each server proposes an allocation and the one which maximizes a social welfare criterion (e.g., the sum of the servers' utilities) is selected. Several heuristic search algorithms to be used by the servers to find such allocations were proposed.

There are situations where the servers have incomplete information about each other. For such situations a preliminary step was added to the strategic negotiation where the servers reveal some of their private information. When the servers use the revelation mechanism, it is beneficial for them to truthfully report their private information. After the preliminary step, the negotiation continues as in the complete information case and yields better results for all the servers than the static allocation policy currently used in EOSDIS. Thus, the overall process in this case is: First, each server broadcasts its private information. If a lie is detected, then the liar is punished by the group. In the next step each server searches for an allocation and then simultaneously each of them proposes one. The allocation which maximizes the pre-defined social-welfare criterion is selected. Then, the servers construct the equilibrium strategies based on the chosen allocation and they start the negotiation using the alternating offers protocol. In the first step of negotiations, the first agent proposes the selected allocation, and the others accept it.

In addition to the theoretical results, simulation results which demonstrate the effect of different parameters of the environment on the negotiation results are also presented. For example, when the servers are more willing to store data locally, better agreements can be reached. The reason for this is that in such situations there are less constraints on finding agreements which are better for all the servers than opting out, and it is easier to find a beneficial compromise. The servers are more willing to store data locally when the storage costs and the cost of delivery of documents stored locally to other servers is low.

In summary, the strategic negotiation model provides a unified solution to a wide range of problems. It is appropriate for dynamic real-world domains.

⁶ In game-theory terminology, the game has multiple equilibria and the problem of the players is to convert into one of them.

In addition to the application of the strategic-negotiation model to data allocation problems in information servers, it was applied to resource allocation and task distribution problems, and the pollution allocation problem [38]. In all these domains the strategic-negotiation model provides the negotiators with ways to reach mutually beneficial agreements without delay. The application of the strategic-negotiation model to human high pressure crisis negotiations was also studied [113,41].

In the next section we will discuss using auctions, another game-theory based technique, for reaching agreements in multiagent environments.

3 Auctions for Resolving Conflicts

In many domains agreements should be reached by the agents concerning the distribution of a set of items. For example, in the information server environment, the agents need to decide on the allocation of datasets, i.e., the items under consideration are datasets. In resolving conflicts on the scheduling of the usage of a resource, an agreement should be reached on the time slots to be assigned to each agent. When the agents need to decide on task assignments, then the items are the tasks and a decision should be made on which agent will carry out a given task. Most of these conflicts can be resolved efficiently by providing the agents with a monetary system, modeling them as buyers and sellers, and resolving the conflicts using a money transfer [71]. For example, a server may “sell” a dataset to another server when relocating this dataset; a subcontractor may be paid in order to carry out a task.

Auctions have become an area of increased interest since a huge volume of economic transactions is conducted through these public sales. The formation of virtual electronic auction houses on the Internet [21] such as eBay [12] has even increased the interest in auctions.

There are two patterns of interactions in auctions. The most common are one-to-many auction protocols [79,1,17] where one agent initiates an auction and a number of other agents can bid in the auction, or many-to-many auction protocols [115] where several agents initiate an auction and several other agents can bid in the auction. Given the pattern of interaction, the first issue to determine is the type of protocols to use in the auction [34]. Given the protocol, the agents need to decide on their bidding strategy.

There are several types of one-to-many auctions which are used, including the English auction, first-price sealed-bid auction, second-price sealed-bid (Vickery auction), and the Dutch auction. The English auction is an ascending auction in which the price is successively raised until only one bidder remains, and that bidder wins the item at the final price. In one variant of the English auction the auctioneer calls higher prices successively until only one willing bidder remains, and the number of active bidders is publicly known at all times. In other variants the bidders call out prices themselves, or have the bids submitted electronically and the best current bid is posted. The first-price sealed bid auction is a sealed-bid auction in which the buyer making the highest bid claims the object and

pays the amount he has bid. The second-price auction is a sealed-bid auction in which the buyer making the highest bid claims the object, but pays only the amount of the second highest bid. In the Dutch auction, the auctioneer begins by naming a very high price and then lowers it continuously until some bidder stops the auction and claims the object for that price. In real world situations, each auction has its advantages and drawbacks [34,53]. In order to test various auction protocols in and to compare bidding strategies, a series of open-invitation events are conducted. These events are featuring software agents from all over the world competing in a market game. The agents need to bid to obtain travel packages for their clients [111]. Sandholm [83] surveys the existing auction protocols, and discusses certain known and new limitations of the protocol for multiagent systems, such as the possibility of bidder collusion and a lying auctioneer.

The Vickrey auction [107] is widely used in DAI [72,26,88,105,104] and in research on electronic commerce [105,104] for the case of one-to-many auctions. Under various assumptions, this protocol is incentive compatible, which means that each bidder has incentives to bid truthfully.

We demonstrate the application of the Vickery auction in the data-allocation problem.

Example 3. An auction protocol can be applied to the data-allocation problem discussed in examples 1 and 2 when a server is concerned with the data stored locally, but does not have preferences concerning the exact storage location of data stored in remote servers. For example, when each server provides information directly to a client which requires it, and obtains payments directly from this client.⁷ According to this approach, the location of each data unit will be determined using an auction protocol, where the server bidding the highest price for obtaining the data will actually obtain it, but will pay the second-highest bid. This approach yields an efficient and fair solution [88], its implementation is simple, and the servers are motivated to offer prices which really reflect their utility.

There are situations in which the value of some items to a bidder depends upon which other items he or she wins. In such cases, bidders may want to submit bids for combinations of items. Such auctions are called *combinatorial auctions*. The main problem in combinatorial auctions is to determine the revenue maximizing set of non-conflicting bids. The general problem is NP-complete. Several researchers have been trying to develop polynomial algorithms, either for specific cases (e.g., [76]) or for finding sub-optimal solutions (e.g., [46,15,27].) Nisan [58] considers bidding languages and the allocation algorithm for combinatorial auctions.

⁷ Note that the auction mechanism is not applicable in the environments that are considered in examples 1 and 2 where each server is concerned with the exact location of each dataset. In the auction mechanism if a server would like to store a dataset it can make a high bid, however, there is no way for a server to influence the location of datasets which are not stored locally.

Double auction is the most known auction protocol for many-to-many auctions. In a double auction, buyers and sellers are treated symmetrically with buyers submitting bids and sellers submitting minimal prices [114]. There are several algorithms used for matching buyers and sellers and for determining the transaction price. Preferably, the protocol will be *incentive compatible*, *individual rational* and Pareto optimal [115]. As mentioned above, an auction is incentive compatible if the agents optimize their expected utilities by bidding their true valuations of the goods. An auction is individual rational if participating in an auction does not make an agent worse off than not participating.

In the next section we will discuss an economic-based mechanism for distributed allocation which consists of auctions.

4 Market-Oriented Programming

Market-oriented programming is an approach to distributed computation based on market price mechanisms [109,112,16].

The idea of market-oriented programming is to exploit the institution of markets and models of them, and to build computational economies to solve particular problems of distributed resource allocation. This is inspired in part by economists' metaphors of market systems "computing" the activities of the agents involved. The modules, or agents, interact in a very restricted manner—by offering to buy or sell quantities of commodities at fixed unit prices. When this system reaches equilibrium, the computational market has indeed computed the allocation of resources throughout the system, and dictates the activities and consumptions of the various modules

(<http://ai.eecs.umich.edu/people/wellman/MOP.html>). Note that this approach does not necessarily require money transfer and it is applicable when there is incomplete information. However, it is applicable only when there are several units of each kind of goods and when the number of agents is large. Otherwise, it is not rational for the agents to ignore the effect of their behavior on the prices when they actually have an influence. Another issue is that there are situations in which reaching an equilibrium may be time consuming, and the system may not even converge [112]. It also requires some mechanism to manage the auctions, (possibly, a distributed mechanism, one for each type of goods.) A survey and a general discussion on the market-programming approach can be found in [112,110]. <http://www2.elec.qmw.ac.uk/~mikeg/text.html> is a market based multi agent systems resource page.

5 Coalition Formation

Another important way for agents to cooperate is by creating coalitions [86,91,92]. The formation of coalitions for executing tasks is useful both in Multi-Agent Systems (MA) and Distributed Problem Solving (DPS) environments. However, in DPS, there is usually no need to motivate the individual agent to join a coalition. The agents can be built to try to maximize the overall performance of

the system. Thus, only the problem of which coalitions should be formed (i.e., the structure of the coalitions) for maximizing the overall expected utility of the agents should be considered. However, finding the coalition structure that maximizes the overall utility of the system is NP-complete.

In Multi-Agent Systems (MA) of self-interested agents, an agent will join a coalition only if it gains more if it joins the coalition than it could gain previously. Thus, in addition to the issue of the coalition structure, the problem of the division of the coalition's joint utility is very important. Game theory techniques for coalition formation can be applied for solving this problem. Work in game theory such as [69,93,108,117] describes which coalitions will form in N-person games under different settings and how the players will distribute the benefits of the cooperation among themselves. This is done by applying several related stability notions such as the core, Shapley value and the kernel [30]. Each of the stability notions is motivated by a different method of measuring the relative strengths of the participating agents. However, the game-theory solutions to the coalition formation problem do not take into consideration the constraints of a multiagent environment, such as communication costs and limited computation time, and do not present algorithms for coalition formation.

The coalition formation of self-interested agents in order to satisfy goals is considered in [92]. Both the coalition structure and the division of the utility problems are handled. An anytime algorithm for forming coalitions that satisfy a certain stability based on the kernel stability criteria is developed. The properties of this algorithm were examined via simulations which showed the model increases the benefits of the agents within a reasonable time period, and more coalition formations provide more benefits to the agents. These results were applied to the formation of coalitions among information agents [35].

Sandholm et al. [80] focused on establishing the worst case bound on the coalition structure quality while only searching a small fraction of the coalition structures. They show that there is a minimal number of structures that should be searched in order to establish a bound. They present an anytime algorithm that establishes a tight bound within this minimal amount of search. If the algorithm is allowed to search further, it can establish a lower bound.

Sandholm and Lesser [84] developed a coalition formation model for bounded rational agents and present a general classification of coalition games. They concentrate on the problem of computing the value of a coalition and in their model this value depends on the computation time available to the agents.

Zlotkin and Rosenschein [118] study the problem of the utility division in Subadditive Task Oriented Domains that is a subset of the Task-Oriented Domains (see section 2.1). They consider only the grand coalition structure where all the agents belong to the same coalition and provide a linear algorithm that guarantees each agent an expected utility that is equal to its Shapley value. Ketchpel [33] presents a utility distribution mechanism designed to perform in similar situations where there is uncertainty in the utility that a coalition obtains.

Coalition formation in DPS environments in order to perform tasks is considered in [91]. In this case, only the coalition structure problem is considered. Efficient distributed algorithms with low ratio bounds and with low computational complexities are presented. Both agent coalition formation where each agent must be a member of only one coalition and overlapping coalitions are considered.

6 Contracting

An agent may try to contract out some of the tasks that it cannot perform by itself, or that may be performed more efficiently by other agents. One self-interested agent may convince another self-interested agent to help it with its task, by promises of rewards.

The main question in such a setting is how one agent can convince another agent to do something for it when the agents do not share a global task and the agents are self-interested. Furthermore, if the contractor-agent can choose different levels of effort when carrying out the task, how can the manager-agent convince the contractor-agent to carry out the task with the level of effort that the manager prefers without the need of the manager's close observation.

The issue of incentive contracting has been investigated in economics and game theory during the last three decades (e.g., [2,74,70,18,25,43]). These works in economics and game theory consider different types of contracts for different applications. Examples of these are contracts between a firm and an employer or employers (e.g., [55,4,5,48]); a government and taxpayers (e.g., [9]); a landlord and a tenant (e.g., [2]); an insurance company and a policy holder (e.g., [78,24,98,45]); a buyer and a seller (e.g., [50,54]); a government and firms (e.g., [51]); stockholders and managements (e.g., [2]); a professional and a client [90], etc. In these situations two parties usually exist. The first party (called "the agent" in economics literature) must choose an action or a level of effort from a number of possibilities, thereby affecting the outcome of both parties. The second party (named "the principal") has the additional function of prescribing payoff rules. Before the first party (i.e., the agent) chooses the action, the principal determines a rule (i.e., a contract) that specifies the fee to be paid to the other party as a function of the principal's observations. Despite the similarity of the above applications, they differ in several aspects, such as the amount of information that is available to the parties, the observations that are made by the principal, and the number of agents. Several concepts and techniques are applied to the principal-agent paradigm in the relevant economics and game theory literature.

A well-known framework for automated contracting is the Contract Net protocol. It was developed for DPS environments where all the agents work on the same goal. In the Contract Net protocol a contract is an explicit agreement between an agent that generates a task (the manager) and an agent that is willing to execute the task (the contractor). The manager is responsible for monitoring the execution of a task and processing the results of its execution, whereas the

contractor is responsible for the actual execution of the task. The manager of a task announces the task's existence to other agents. Available agents (potential contractors) then evaluate the task announcements made by several managers and submit bids for the tasks they are suited to perform. Since all the agents have a common goal and are designed to help one another, there is no need to motivate an agent to bid for tasks or to do its best in executing them if its bid is chosen. The main problems addressed by [95,97,96] are task decomposition, sub-tasks distribution, and synthesis of the overall solution.

The Contract Net was used in various domains [65,60,49,89]. For example, a modified version of the Contract Net protocol for competitive agents in the transportation domain was presented in [79]. It provides a formalization of the bidding and the decision awarding processes, based on marginal cost calculations according to local agent criteria. More important, an agent will submit a bid for a set of delivery tasks only if the maximum price mentioned in the tasks' announcement is greater than what the deliveries will cost that agent. A simple motivation technique is presented to convince agents to make bids; the actual price of a contract is half way between the price mentioned in the task announcement and the bid price.

Contracting in various situations of automated agent environments is considered in [37]. These situations include certainty vs. uncertainty, full information vs. partial information, symmetric information vs. asymmetric information and bilateral situations vs. situations where there are more than two automated agents in the environment. For each of these situations appropriate economic mechanisms and techniques that can be used for contracting in automated agents environments are fitted from the game theory or economics literature. In all the situations that are considered, the agent that designs the contract is provided with techniques to maximize its personal expected utilities, given the constraints of the other agent(s).

Sandholm and his colleagues [81,82] developed a backtracking method called *leveled commitment contract* where each party can unilaterally decommit to a contract by paying a predetermined penalty. They show that such contracts improve expected social welfare even when the agents decommit strategically in Nash equilibrium.

7 Logical Approaches to Argumentation

Several researchers developed frameworks for negotiation through argumentation in which agents exchange proposals and counter-proposals backed by arguments that summarize the reasons why the proposal should be accepted. The argumentation is persuasive because the exchanges are able to alter the mental state of the agents involved.

Most of these framework are based on logical models of the mental states of the agents representing, for example, their beliefs, desires, intentions, and goals. The formal models are used in two manners. One use is as a specification for agent design [19]. In this role, the model constrains certain planning and

negotiation processes. It can also be used to check the agents' behavior. Another use of the model is by the agents themselves.

Parsons and Jennings [64] drew upon a logic of argumentation to devise a system of argumentation and use it to implement a form of dialectic negotiation. In their context, an argument is a sequence of logical steps indicating support or doubt of a proposition. They have a function of flattening, which can measure the set of arguments into some metric of how favored the proposition is, by determining which class of acceptability the arguments belong to.

Qiu and Tambe [67] focus on negotiations between team members to resolve conflicts that arise from conflicting beliefs about different aspects of their environment, about resource availability, and about their own or teammates' capabilities and performance. The basis of such negotiations is inter-agent argumentation where agents assert their beliefs to others, with supporting arguments. Their approach is implemented in a system called CONSA (COLlaborative Negotiation System based on Argumentation).

In [40] a logic is used in the above two ways. Using categories identified in human multi-agent negotiation, demonstrate how the logic can be used to specify argument formulation and evaluation. Furthermore, [40] presents a general Automated Negotiation Agent which was implemented, based on the logical model. Using this system, a user can analyze and explore different methods to negotiate and argue in a non-cooperative environment where no centralized mechanism for coordination exists. The development of negotiating agents in the framework of the Automated Negotiation Agent is illustrated with an example where the agents plan, act, and resolve conflicts via negotiation in a Blocks World environment.

The formal model of [40] consists of a set of agents, not necessarily cooperative, with the ability to exchange messages. Their mental states are characterized by using the notions of beliefs, goals, desires, intentions, and local preferences. Each agent has a set of desires. The agent's activities are motivated by the will to fulfill these desires. At any given time, an agent selects a consistent subset of its desires. This serves as its set of current goals. An agent ascribes different degrees of importance to different goals. It prefers to fulfill goals of higher importance. The set of goals motivate the agent's planning process.

The planning process may generate several intentions. Some of these are in what we would like to classify as the "intend-to-do" category and refer to actions that are within the direct control of the agent. Others are among the "intend-that" category [8,19,20,106]. These are propositions not directly within the agent's realm of control, that it must rely on other agents for satisfying.⁸ Often, there is room for argumentation when intend-that actions are part of a plan. Argumentation is the means by which an agent, the persuader, attempts to modify the intention structure of another agent, the persuadee, to include

⁸ The proposition may include a negation. When fulfillment of the proposition is beyond the control of the agent, it can be achieved by convincing another agent to abandon a relevant intention, or by convincing it to take an action that will make the proposition true.

the actions the persuader wants it to do. While an agent tries to influence the intentions of other agents, other agents may try to convince it as well. The role of persuader and persuadee is not fixed, but dynamically assumed during the agent interactions. Thus, during a negotiation process, each agent may update its intentions and goals after receiving a message from another agent. If the argumentation happens to fail, the agent which sent it must revise its arguments, its plans, and/or seek other sources of satisfying the portion of its plan in question.

An agent's belief set includes beliefs concerning the world and beliefs concerning mental states of other agents. An agent may be mistaken in both kinds of beliefs. It may update its beliefs by observing the world and after receiving messages from other agents. Each agent's actions are based upon its mental model of other agents.

Arguments serve either to add an intention to the persuadee's set or to retract an intention or to change the preferences of the persuadee. Below we present a list of several argument types. These argument types are not meant to constitute an exhaustive typology of arguments. Indeed, it has been pointed out [103] that it is not possible to present such an authoritative classification, since arguments must be interpreted and are effective within a particular context and domain. The six argument types that we present are ones that are commonly thought to have persuasive force in human negotiations [61,31,66]. Argumentations which were shown to be successful in human negotiation, may be also successful in automated agents' negotiations. Furthermore, we want our agents to be able to negotiate with humans, and therefore they need to be able to at least understand human argumentation. Moreover, the designers of the agents can follow the negotiation of the agents, if it is similar to human negotiation. The argument types we present are:

1. Threats to produce goal adoption or goal abandonment on the part of the persuadee.
2. Enticing the persuadee with a promise of a future reward.
3. Appeal to past reward.
4. Appeal to precedents as counterexamples to convey to the persuadee a contradiction between what she/he says and past actions.
5. Appeal to "prevailing practice" to convey to the persuadee 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 persuadee that taking this action will enable achievement of a high-importance goal.

Threats and promises are the most common arguments used in human negotiations [7]. An appeal to prevailing practice is the most common argument used in the legal system. Furthermore, it was found that presenting example instances (prevailing practice cases) is much more persuasive than presenting statistical summaries [36,101,59,23]. An "appeal to past promise" is supported by the cognitive dissonance theory [61] that assumes that a person seeks to maximize the internal psychological consistency of his/her cognition, and thus will be willing to keep his/her promises. This argument is also important in repeated interactions

since agents prefer to maintain their credibility. The other two arguments, “an appeal to self interest” and “a counter example” are examples of arguments useful to persuade bounded rational agents which have limited inferential resources. More discussion on these arguments can be found in [40].

8 Conclusions

Game-theory and economics techniques seem to be very useful in the development of self-interested automated agents that act in a well-defined environment. Logical models provide a framework for argumentations. In this paper we emphasized formal techniques. We believe that using them in multi-agent systems is beneficial because there is a need to provide the agents with well-designed algorithms.

The choice of the specific technique for a given domain depends on the specification of the domain. For example, whether the agents are self interested, the number of agents in the environment, the type of agreement that they need to reach, and the amount and the type of information the agents have about each other.

While negotiation has been studied in other disciplines for many years, the study of negotiations of multi-agent environment is relatively new. In particular, there are many open questions with respect to applying formal models to multi-agent environments. The main challenge is how to maintain the useful results of the formal models, while adjusting them to real-world applications.

References

1. M. R. Andersson and T. W. Sandholm. Contract Types for Optimal Task Allocation: II Experimental Results. In *AAAI 1998 Spring Symposium: Satisficing Models*, Stanford University, California, 1998.
2. K. J. Arrow. The economics of agency. In J. Pratt and R. Zeckhauser, editors, *Principals and Agents: The Structure of Business*, pages 37–51. Harvard Business School Press, Cambridge, MA, 1985.
3. R. Azoulay-Schwartz and S. Kraus. Negotiation on data allocation in multi-agent environments. *Autonomous Agents and Multi-Agent Systems journal*, 2001. To appear.
4. S. Baiman and J. Demski. Economically optimal performance evaluation and control systems. *Journal of Accounting Research*, 18:184–220, 1980.
5. A. Banerjee and A. Beggs. Efficiency in hierarchies: implementing the first-best solution by sequential actions. *The Rand Journal of Economics*, 20(4):637–645, 1989.
6. M. Ben-Or and N. Linial. Collective coin flipping, robust voting games and minima of banzhaf values. In *Proc. 26th IEEE Symposium on the Foundations of Computer Science*, pages 408–416, Portland, 1985.
7. F. J. Boster and P. Mongeau. Fear-arousing persuasive messages. In R. N. Bostrom, editor, *Communication yearbook 8*. SAGE Publications, 1984.
8. M. E. Bratman. Shared cooperative activity. *The Philosophical Review*, 101:327–341, 1992.

9. B. Caillaud, R. Guesnerie, P. Rey, and J. Tirole. Government intervention in production and incentives theory: a review of recent contributions. *Rand Journal of Economics*, 19(1):1–26, 1988.
10. S.E. Conry, K. Kuwabara, V.R. Lesser, and R.A. Meyer. Multistage negotiation for distributed satisfaction. *IEEE Transactions on Systems, Man, and Cybernetics, Special Issue on Distributed Artificial Intelligence*, 21(6):1462–1477, December 1991.
11. D. Druckman. *Negotiations*. Sage Publications, London, 1977.
12. eBay. eBay - Your Personal Trading Community. <http://www.ebay.com/aw/>, 2001.
13. R. Fisher and W. Ury. *Getting To Yes: Negotiating Agreement Without Giving In*. Houghton Mifflin, Boston, 1981.
14. E. Freitsis. Negotiations in the pollution sharing problem. Master's thesis, Bar-Ilan University, 2000.
15. Y. Fujishima, K. Leyton-Brown, and Y. Shoham. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In Dean Thomas, editor, *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99-Vol1)*, pages 548–553, S.F., July 31–August 6 1999. Morgan Kaufmann Publishers.
16. C. Gerber, C. Russ, and G. Vierke. On the suitability of market-based mechanisms for telematics applications. In *Proc. of Agents99*, pages 408–409, 1999.
17. E. Gimenez-Funes, L. Godo, and J. A. Rodriguez-Aguilar. Designing bidding strategies for trading agents in electronic commerce. In *ICMAS98*, pages 136–143, Paris, 1998.
18. S. Grossman and O. Hart. An analysis of the principal-agent problem. *Econometrica*, 51(1):7–45, 1983.
19. B. J. Grosz and S. Kraus. Collaborative plans for complex group activities. *Artificial Intelligence Journal*, 86(2):269–357, 1996.
20. B. J. Grosz and S. Kraus. The evolution of sharedplans. In A. Rao and M. Wooldridge, editors, *Foundations and Theories of Rational Agency*, pages 227–262. Kluwer Academic Publishers, 1999.
21. R. H. Guttman and P. Maes. Cooperative vs. competitive multi-agent negotiations in retail electronic commerce. In *The Second International Workshop on Cooperative Information Agents (CIA98)*, pages 135–147, Paris, 1998.
22. Lavinia Hall, editor. *Negotiation: Strategies for Mutual Gain*. Sage, Beverly Hills, 1993.
23. R. Hamill, T. D. Wilson, and R. E. Nisbett. Insensitivity to sample bias: Generalizing from atypical cases. *Journal of Personality and Social Psychology*, 39:579–589, 1980.
24. M. Harris and A. Raviv. Some results on incentive contracts with applications to education and employment, health insurance, and law enforcement. *The American Economic Review*, 68(1):20–30, 1978.
25. J. Hirshleifer and J. Riley. *The Analytics of Uncertainty and Information*. Cambridge University Press, Cambridge, 1992.
26. B. Huberman and S. H. Clearwater. A multi-agent system for controlling building environments. In *Proceedings of ICMAS-95*, pages 171–176, San Francisco, 1995.
27. L. Hunsberger and B. Grosz. A combinatorial auction for collaborative planning. In *ICMAS-2000*, Boston, USA, 2000.
28. N. R. Jennings, P. Faratin, A.R. Lomuscio, S. Parsons, C. Sierra, and M. Wooldridge. Automated haggling: Building artificial negotiators. In *Pacific Rim International Conference on Artificial Intelligence*, page 1, 2000.

29. R. Johnson. *Negotiation basics*. Sage, Newbury Park, 1993.
30. J. P. Kahan and A. Rapoport. *Theories of coalition formation*. Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1984.
31. M. Karllins and H. I. Abelson. *Persuasion: How opinions and attitudes are changed*. Springer Publishing Company, Inc., second edition, 1970.
32. C.L. Karrass. *The Negotiating Game: How to Get What You Want*. Thomas Crowell Company, NY, 1970.
33. S. P. Ketchpel. Forming coalitions in the face of uncertain rewards. In *Proc. of AAAI94*, pages 414–419, Seattle, Washington, 1994.
34. P. Klemperer. Auction theory: A guide to literature. *Journal of Economic Surveys*, 13(3), 1999.
35. M. Klusch and O. Shehory. A polynomial kernel-oriented coalition formation algorithm for rational information agents. In *Proc. of ICMAS-96*, pages 157–164, Kyoto, Japan, 1996.
36. R. R. Jr. Koballa. Persuading teachers to reexamine the innovative elementary science programs of yesterday: The effect of anecdotal versus data-summary communications. *Journal of Research in Science Teaching*, 23:437–449, 1986.
37. S. Kraus. An overview of incentive contracting. *Artificial Intelligence Journal*, 83(2):297–346, 1996.
38. S. Kraus. *Strategic Negotiation in Multiagent Environments*. MIT Press, Cambridge, USA, 2001.
39. S. Kraus and D. Lehmann. Designing and building a negotiating automated agent. *Computational Intelligence*, 11(1):132–171, 1995.
40. S. Kraus, K. Sycara, and A. Evenchik. Reaching agreements through argumentation: a logical model and implementation. *Artificial Intelligence*, 104(1-2):1–69, 1998.
41. S. Kraus and J. Wilkenfeld. A strategic negotiations model with applications to an international crisis. *IEEE Transaction on Systems Man and Cybernetics*, 23(1):313–323, 1993.
42. D. Kreps and R. Wilson. Sequential equilibria. *Econometrica*, 50:863–894, 1982.
43. J. Laffont and J. Tirole. *A Theory Of Incentives in Procurement and Regulation*. The MIT Press, Cambridge, Massachusetts, 1993.
44. Susan E. Lander and Victor R. Lesser. Customizing distributed search among agents with heterogeneous knowledge. In *Proc. first int. conf. on Information Knowledge Management*, pages 335–344, Baltimore, 1992.
45. M. Landsberger and I. Meilijson. Monopoly insurance under adverse selection when agents differ in risk aversion. *Journal of Economic Theory*, 63:392–407, 1994.
46. Daniel Lehmann, Liadan Ita O’Callaghan, and Yoav Shoham. Truth revelation in rapid, approximately efficient combinatorial auctions. In *Proceedings of the ACM Conference on Electronic Commerce (EC’99)*, pages 96–102, N.Y., November 3–5 1999. ACM Press.
47. R. D. Luce and H. Raiffa. *Games and Decisions*. John Wiley and Sons, New York, 1957.
48. I. Macho-Stadler and J. Pérez-Castrillo. Moral hazard and cooperation. *Economics Letters*, 35:17–20, 1991.
49. T. W. Malone, R. E. Fikes, K. R. Grant, and M. T. Howard. Enterprise: A market-like task schedule for distributed computing environments. In B. A. Huberman, editor, *The Ecology of Computation*, pages 177–205. North Holland, 1988.
50. S. Matthews. Selling to risk averse buyers with unobservable tastes. *Journal of Economy Theory*, 30:370–400, 1983.

51. R. P. McAfee and J. McMillan. Bidding for contracts: a principal-agent analysis. *The Rand Journal of Economics*, 17(3):326–338, 1986.
52. T. Moehlman, V. Lesser, and B. Buteau. Decentralized negotiation: An approach to the distributed planning problem. *Group Decision and Negotiation*, 2:161–191, 1992.
53. D. Monderer and M. Tennenholtz. Optimal auctions revisited. *Artificial Intelligence Journal*, 120:29–42, 2000.
54. R. Myerson. Mechanism design by an informed principal. *Econometrica*, 51:1767–1798, 1983.
55. B. Nalebuff and J. Stiglitz. Information, competition, and markets. *American Economic Review*, 73(2):278–283, 1983.
56. NASA. EOSDIS Home Page.
<http://www-v0ims.gsfc.nasa.gov/v0ims/index.html>, 1996.
57. J. F. Nash. Two-person cooperative games. *Econometrica*, 21:128–140, 1953.
58. Noam Nisan. Bidding and allocation in combinatorial auctions. In *ACM Conference on Electronic Commerce*, pages 1–12, 2000.
59. R. E. Nisbett, E. Borgida, R. Crandall, and H. Reed. Popular induction: information is not necessarily informative. In J. S. Carroll and J. W. Payne, editors, *Cognition and social behavior*. Lawrence Erlbaum, 1976.
60. T. Ohko, K. Hiraki, and Y. Anzai. Reducing communication load on contract net by case-based reasoning: extension with directed contract and forgetting. In *Proceedings of the First International Conference on Multi-Agent Systems*, Cambridge, MA, 1995. MIT Press.
61. Daniel J. O’Keefe. *Persuasion: Theory and Research*. SAGE Publications, 1990.
62. M. J. Osborne and A. Rubinstein. *Bargaining and Markets*. Academic Press Inc., San Diego, California, 1990.
63. M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, Cambridge, Massachusetts, 1994.
64. S. Parsons, C. Sierra, and N. R. Jennings. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3):261–292, 1998.
65. H. van Dyke Parunak. Manufacturing experience with the contract net. In *Proceedings of the 1995 Distributed Artificial Intelligence Workshop*, pages 67–91, December 1995.
66. D. G. Pruitt. *Negotiation Behavior*. Academic Press, New York, N.Y., 1981.
67. Z. Qiu and M. Tambe. Flexible negotiations in teamwork: Extended abstract. In *Proceedings of the AAAI Fall Symposium on Distributed Continual Planning*, 1998.
68. H. Raiffa. *The Art and Science of Negotiation*. Harvard University Press, Cambridge, MA, 1982.
69. A. Rapoport. *N-Person Game Theory*. University of Michigan, Michigan, 1970.
70. E. Rasmusen. *Games and Information*. Basil Blackwell Ltd., Cambridge, MA, 1989.
71. Amir Ronen. Algorithms for rational agents. In *Conference on Current Trends in Theory and Practice of Informatics*, pages 56–70, 2000.
72. J. S. Rosenschein. *Rational Interaction: Cooperation Among Intelligent Agents*. PhD thesis, Stanford University, 1986.
73. J. S. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation Among Computers*. MIT Press, Boston, 1994.
74. S. Ross. The economic theory of agency: The principal’s problem. *The American Economic Review*, 63(2):134–139, 1973.

75. A. E. Roth. *Axiomatic Models of Bargaining*. Lecture Notes in Economics and Mathematical Systems No. 170. Springer-Verlag, Berlin, 1979.
76. M. H. Rothkopf, A. Pekec, and R. M. Harstad. Computationally manageable combinatorial auctions. Technical Report 95-09, DIMACS, April 19 1995.
77. A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, 1982.
78. A. Rubinstein and M. Yaari. Repeated insurance contracts and moral hazard. *Journal of Economic Theory*, 30:74–97, 1983.
79. T. Sandholm. An implementation of the contract net protocol based on marginal cost calculations. In *Proc. of AAAI-93*, pages 256–262, Washington D.C., July 1993.
80. T. Sandholm, K. Larson, M. R. Andersson, O. Shehory, and F. Tohm. Coalition structure generation with worst case guarantees. *Artificial Intelligence Journal*, 111:209–238, 1999.
81. T. Sandholm, S. Sikka, and S. Norden. Algorithms for optimizing leveled commitment contracts. In *Proc. of IJCAI99*, pages 535–540, Stockholm, Sweden, 1999.
82. T. Sandholm and Y. Zhou. Surplus equivalence of leveled commitment contracts. In *Prof. of ICMAS-2000*, pages 247–254, Boston, MA, 2000. IEEE Computer Society.
83. T. W. Sandholm. Limitations of the vickrey auction in computational multiagent systems. In *International Conference on Multiagent Systems (ICMAS-96)*, pages 299–306, Kyoto, Japan, 1996.
84. T. W. Sandholm and V. R. Lesser. Coalition formation among bounded rational agents. In *Proc. of IJCAI-95*, pages 662–669, Montréal, 1995.
85. T. W. Sandholm and V. R. Lesser. Issues in automated negotiation and electronic commerce: Extending the contract net framework. In *First International Conference on Multiagent Systems (ICMAS-95)*, pages 328–335, San Francisco, 1995.
86. T. W. Sandholm and V. R. Lesser. Coalition formation among bounded rational agents. *Artificial Intelligence*, 94(1-2):99–137, 1997. Special issue on Principles of Multi-Agent Systems.
87. R. Schwartz and S. Kraus. Negotiation on data allocation in multi-agent environments. In *Proc. of AAAI-97*, pages 29–35, Providence, Rhode Island, 1997.
88. R. Schwartz and S. Kraus. Bidding mechanisms for data allocation in multi-agent environments. In Munindar P. Singh, Anand S. Rao, and Michael J. Wooldridge, editors, *Intelligent Agents IV: Agent Theories, Architectures, and Languages*, pages 61–75. Springer-Verlag, 1998.
89. Sandip Sen and Edmund Durfee. A contracting model for flexible distributed scheduling. *Annals of Operations Research*, 65:195–222, 1996.
90. S. Shavell. Risk sharing and incentives in the principal and agent relationship. *Bell Journal of Economics*, 10:55–79, 1979.
91. O. Shehory and S. Kraus. Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 15(3):218–251, 1998.
92. O. Shehory and S. Kraus. Feasible formation of stable coalitions among autonomous agents in non-super-additive environments. *Computational Intelligence*, 15(3):218–251, 1999.
93. K. I. Shimomura. The bargaining set and coalition formation. Technical Report 95-11, Brown University, Dept. of Economics, 1995.

94. C. Sierra, P. Faratin, and N. Jennings. A service-oriented negotiation model between autonomous agents. In *Proc. 8th European Workshop on Modeling Autonomous Agents in a Multi-Agent World (MAAMAW-97)*, pages 17–35, Ronneby, Sweden, 1997.
95. R. Smith and R. Davis. Framework for cooperation in distributed problem solvers. *IEEE Transactions on Systems, Man and Cybernetic*, C-29(12):61–70, 1981.
96. R.G. Smith. The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on Computers*, 29:1104–1113, 1980.
97. R.G. Smith and R. Davis. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, 20:63–109, 1983.
98. M. Spence and R. Zeckhauser. Insurance, information and individual action. *The American Economic Review*, 61(1):380–391, 1971.
99. K. P. Sycara. *Resolving Adversarial Conflicts: An Approach to Integrating Case-Based and Analytic Methods*. PhD thesis, School of Information and Computer Science, Georgia Institute of Technology, 1987.
100. K. P. Sycara. Persuasive argumentation in negotiation. *Theory and Decision*, 28:203–242, 1990.
101. S. E. Taylor and S. C. Thompsons. Staking the elusive vividness effect. *Psychological Review*, 89:155–181, 1982.
102. Jean Tirole. *The Theory of Industrial Organization*. The MIT Press, Cambridge, MA, 1988.
103. S. Toulmin, R. Rieke, and A. Janik. *An introduction to reasoning*. Macmillan Publishing Co., Inc., 1979.
104. M. Tsvetovaty, M. Gini, B. Mobasher, and Z. Wieckowski. Magma: An agent-based virtual market for electronic commerce. *Applied Artificial Intelligence, special issue on Intelligent Agents*, 6:501–523, 1997.
105. M. B. Tsvetovaty and M. Gini. Toward a virtual marketplace: Architectures and strategies. In *The first International Conference on the Practical Application of Intelligent Agents and Multi Agents Technology*, pages 597–614, London, 1996.
106. B. Vermazen. Objects of intention. *Philosophical Studies*, 71:223–265, 1993.
107. William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *J. of Finance*, 16:8–37, March 1961.
108. R. Vohra. Coalitional non-cooperative approaches to cooperation. Technical Report 95-6, Brown University, Dept. of Economics, 1995.
109. M. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, 1:1–23, 1993.
110. M. Wellman. Market-oriented programming: Some early lessons. In S. Clearwater, editor, *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific, Singapore, 1996.
111. M. Wellman and M. Wurman. TAC: Trading agent competition. <http://auction2.eecs.umich.edu/game.html>, 2001.
112. M. P. Wellman and P.R. Wurman. Market-aware agents for a multiagent world. *Robotics and Autonomous Systems*, 24:115–125, 1998.
113. J. Wilkenfeld, S. Kraus, K. Holley, and M. Harris. Genie: A decision support system for crisis negotiations. *Decision Support Systems*, 14:369–391, 1995.
114. R. Wilson. Incentive efficiency of double auctions. *Econometrica*, 53(5):1101–1116, 1985.

115. P. R. Wurman, W. E. Walsh, and M. P. Wellman. Flexible double auctions for electronic commerce: Theory and implementation. *Decision Support Systems*, 24:17–27, 1998.
116. Dajun Zeng and Katia Sycara. Bayesian learning in negotiation. *International Journal of Human-Computer Studies*, 48:125–141, 1998.
117. L. Zhou. A new bargaining set of an n-person game and endogenous coalition formation. *Games and Economic Behavior*, 6:512–526, 1994.
118. G. Zlotkin and J. S. Rosenschein. Coalition, cryptography, and stability: Mechanisms for coalition formation in task oriented domains. In *Proc. of AAAI94*, pages 432–437, Seattle, Washington, 1994.