Optimal sequencing of individually rational contracts

Partha Sarathi Dutta
Math & CS Department
University of Tulsa
Tulsa, OK, USA
partha@ens.utulsa.edu

Sandip Sen
Math & CS Department
University of Tulsa
Tulsa, OK, USA
sandip@kolkata.mcs.utulsa.edu

ABSTRACT

Multiagent researchers have worked on the problem of determining optimal contracts between self-interested agents. In particular, Sandholm *et al.* [1, 8] have both theoretically and experimentally studied the necessity and usefulness of different contract types under the assumption of individually myopically rational contracting. We study a variant of sequential contracting where the goal is to maximize social welfare through a fixed-length sequence of individually rational contracts. The space of possible contract sequences is exponential. We compare a greedy deterministic heuristic with a stochastic genetic algorithm based approach for this optimal sequential contract selection problem. We focus on sub-additive domains where individually rational contracts are feasible with side payments. We show that the GA-based approach consistently outperforms the deterministic heuristic by generating larger social welfare.

1. INTRODUCTION

Research areas such as contract protocols, coalition structures, argumentation-based negotiation and combinatorial auctions have generated wide spread interest among multiagent system researchers [6, 7, 10, 13]. An important objective of all these approaches is to allocate or re-allocate resources among agents or agent groups so that performance is optimized (in terms of minimizing cost of performing a set of tasks or enhancing the payoff received agents).

In this paper, we are interested in the identification of an optimal sequence of contracts by which a group of agents can decrease their cost of performing allocated tasks. We assume the cost of performing a task depends on other tasks in the allocation. Thus an agent can contract out a task such that the decrease in cost for the contractor agent is less than the increase in cost of the contractee agent. The contractor agent can then be better off even after compensating the contractee agent for the cost of the task that is being transferred. We assume that a contract is executed only if it benefits both the contractor and the contractee, i.e., agents are individually rational. Such mutually beneficial contracts situation is feasible in sub-additive domains.

AAMAS'02, July 15-19, 2002, Bologna, Italy.

The space of possible contracts is exponential in the number of agents and the total number of tasks. Sandholm *et al.* [8] have studied the necessity and usefulness of different contract types under the assumption of individually myopically rational contracting, i.e., a contract is accepted only if it is beneficial to both parties (this is myopic because an individual contract that produces a loss can, in the long term, lead to a contrat with a larger total benefit). In particular, they observe that if the allowed contract types do not include the fully general OCSM contracts, then individually rational contracts may not be found to reach globally optimal allocations. Sandholm's work, however, provides no guidelines for selecting a sequence of OCSM contracts to produce optimal allocations. Their experimental work on evaluating original (O), cluster (C), swap (S), and multiagent (M) contracts evaluate an exhaustively enumerative, rather than heuristic, scheme for sequencing contracts [1]. We believe that it is necessary to develop a better characterization of the usefulness of simple contract types under constraints on time and number of task exchanges feasible. The motivation behind such a study is to develop insight on the best possible sequence of contracts that can be achieved given the restrictions.

In particular, we are interested in a variation of the optimal contracting problem where the number of contracts is fixed. Each contract is a transfer of one task from a contractor to a contractee, which is an O contract. The *optimal n-sequence O* contracting problem involves the generation of the optimal sequence consisting of n individually rational O contracts that minimizes the total cost incurred by all agents to process their final allocations. Minimizing total cost is equivalent to maximizing the total welfare in the system. *Optimality* in this context refers to the best possible contracting sequence under the bounds specified.

We evaluate two heuristic algorithms, a deterministic and a stochastic version, for selecting the sequence of n individually rational O contracts that maximizes social welfare given arbitrary initial task allocations among agents. The deterministic algorithm is a greedy procedure for selecting the maximal cost reducing individually rational contract. The stochastic version has been implemented using an order-based genetic algorithm (GA) [2] where each GA structure selects a fixed set of contracts. Our results show that under various combinations of tasks and agents, the stochastic contracting protocol produces significantly larger social welfare gains. The metric we have used to quantify the performance of these protocols and to compare between them (the metric that we would like to maximize for both the protocols) is the welfare obtained after a sequence of contracts, where welfare is the net difference of costs between the final and the initial allocations. The GA searches a number of contracting sequences in parallel to find the optimal solution. Apparently, this implies that the GA-based approach is not comparable to the deterministic heuristic that searches in a limited sequence of O-contracts. It should be noted, however, that the deterministic heuristic uses knowledge about the costs of each agent's task allocation to eliminate a large number of possible contracts from the entire set of contracts possible. The GA, on the contrary, uses random sampling in the set of possible contracts without any knowledge about the contracting costs of the agents. Thus, though the GA searches in a larger set of possible solutions compared to the deterministic heuristic, the search is unbiased and hence, potentially runs over a non-trivial amount of redundant sample solutions.

Both of the algorithms studied here are centralized in nature in terms of the knowledge of individual task allocations and cost metrics of individual agents. While this is a limitation of the current work, our goal here is to identify whether even centralized heuristic algorithms exist to effectively search the exponential space of possible contract sequences. To accommodate the constraint that a centralized arbiter will not be able to enforce contracts on individually rational agents, we require that each O-contract in the n-sequence is individually rational. Therefore, for each of the contracts identified by the centralized algorithm, it is beneficial for the participant agents to accept the task exchanges with corresponding accompanying side payments. This means that only those contracts will be selected that are acceptable for rational agents. This is critical as a centralized arbiter will not be able to enforce contracts that are not acceptable to at least one of the participant agents. In spite of the limitations of the centralized framework, we believe that this work is important to create a baseline for measuring the performance of future distributed contract sequencing heuristics. The importance of this work can be further argued against the backdrop of almost a complete lack of effective heuristics for identifying mutually beneficial contracts.

2. HEURISTIC ALGORITHMS FOR SELECTING CONTRACTS

We assume an initial allocation of a set of tasks, T, to M agents in the form of a partition of the set of tasks: $T = \bigcup_{i \in M} T_i$ and $\forall i, j$, $T_i \cap T_j = \phi$, where T_i is the set of tasks assigned to agent i. The cost of an allocation to an agent i is given by a cost function, $g_i(T_i)$.

Sandholm and his students [1, 8] have studied several contract types for exchange of tasks between *myopically individually rational* agents to maximize social welfare. Of these contract types, we use in our work only the O-contract (for original contract) which involves one agent allocating one of its tasks to another agent. In Sandholm's work such a contract is myopically individually rational only if the cost of the task being transferred is less for the recipient than for the giver. The recipient is compensated by the giver in the form of a side payment to cover the cost incurred in processing the task. As a result both the giver (because the cost to it for processing the task is less) and the recipient (because it gets a payment for the task that is higher than the cost of processing the task) benefits, and global welfare is improved. The contracting domain they have studied consists of each agent solving a traveling salesman's problem (TSP) and an O-contract involves one agent transferring one city in its route to another agent. This exchange can reduce the size of the optimal tour of the former by more than the increase of the optimal tour of the latter.

Such exchange of tasks can prove to be mutually beneficial to the contracting parties in *sub-additive domains*. These are domains where the sum of the cost of doing two sets of tasks separately may be more than the cost of doing the union of the two task sets.

For sub-additive domains, we define individually rational contracts as those for which the cost of the resultant allocation is less than or equal to the cost of the initial allocation for each agent i participating in the contract:

$$g_i(T_i^{new}) < g_i(T_i^{old})$$

where T_x^{old} refers to the allocation to agent x previous to the contract and T_x^{new} refers to the new allocation to agent x. The objective of the heuristic algorithms is to maximize *total welfare increment*, ΔW, which is the net difference between the costs of the *final* and the *initial* task allocations.

$$\Delta W(S) = \sum_{i \in M} g_i(T_i^{final}) - \sum_{i \in M} g_i(T_i^{initial}),$$

where $T_i^{initial}$ and T_i^{final} are the initial and final allocations to agent i at the start and end of the exchanging tasks as specified by the contract sequence S.

We do not concern ourselves with specification of side payments for exchanges of tasks in individually rational contracts. This is left at the discretion of the participating agents in the contracts. It suffices to note that for individually rational contracts one can always find side payments which will make the exchange of task attractive to both parties.

We now discuss the algorithms used for selecting a sequence of n contracts for reallocating tasks between individually rational agents, given initial task allocations. We present both a deterministic greedy contract sequencing heuristic and a stochastic GA-based contract sequence selection approach.

2.1 Deterministic heuristic

Andersson and Sandholm [1] present an enumerative algorithm for selecting sequences of individually rational O-contracts. The sequencing of this contracting procedure starts from agent 1, that tries to pass all its tasks, one at a time, to agent 2. This procedure is repeated for each agent until no more contracts can be made. This protocol leads to maximum benefit in terms of total cost saved by all the agents, but requires an exorbitant number of trials or contracts to achieve that [1].

Trying out a disproportionately large number of contracts to arrive at the optimum allocation is not feasible in practice. As designers of contracting protocols, we will be interested in measuring the quality of a solution generated within a given amount of time. We modify the enumerative contracting sequencing algorithm to heuristically order the sequence of contracts attempted with the goal of selecting the maximal total reduction in cost (increase in welfare) under the constraint of a fixed number of contracts. This corresponds to real-life negotiating scenarios where we are interested in identifying the few exchanges that lead to maximal utility increment.

Limiting the number of contracts is reasonable in many practical situations as there may be sufficient time or resources to execute only a fixed number of contracts. We adapt a greedy strategy for sequencing O-contracts, which states that the agent trying to get a contract, uses its costliest task at that time and at any time the agent with the costliest task (according to its present allocation structure) gets to call for a contract, first. More formally, assuming that the tasks are arranged in decreasing order of their costs for each agent i as $< t_{i1}, t_{i2}, ..., t_{iN} >$, where t_{ij} is the cost of the j^{th} task of agent i, then agent i will be able to call for a contract for its task t_{i1} if $t_{i1} \geq t_{j1}, \forall j$ [1]. The other greedy choice in our procedure is the selection of the agent who can be awarded the contract. The agent for whom receiving the announced task will lead to the minimal increase in allocation cost is identified and if the corresponding increase is less than the cost decrease obtained by the contractor agent, an individually rational contract has been identified. If such a contract is possible, then the new task allocations of the two participating agents are re-ordered by the updated costs and the above procedure is repeated. If agent i, on the other hand, fails to contract its costliest task t_{i1}, then the next agent with the highest task cost tries, and so on. If no agent can contract away their costliest task, the agents get to announce contracts in order of their second costliest task, and so on.

The contracting sequence ends either if the specified finite number of contracts have been performed or there are no more contracts to be made. The latter happens, if all the agents, in turn, try their costliest task, next to costliest task, next to next to costliest task, and so on, until all the tasks of each have been tried and yet no contract resulted (i.e, none of those resulted in a valid contract).

In our experiments we use different random initial task allocations among the agents for every combination of task and agent numbers. The results are obtained by averaging the final welfare obtained for each of the random initial task allocations.

2.2 Non-deterministic contracting algorithm

We have used an order-based genetic algorithm (OBGA) as the stochastic alternative of the O-contract protocol. GAs are a class of stochastic algorithms that have been effectively used in combinatorial optimization problems. Searching for the optimal contract is a combinatorial optimization problem with an exponential search space. As long as there is some regularity in the search space, GAs have the potential of detecting the regularity and finding the contracts that would perform effectively. GAs, however, do not guarantee finding an optimal solution or bounding the quality of the solution within a specified number of iterations.

We now discuss the representation that we used for the OBGA for the optimal sequential contract problem. Let us assume that the number of contracts allowed is k. Then each member in the GA population is a string of length $3k$ and contains k triplets. Each triplet consists of the *contracting* agent, the agent to which the contract is allocated, i.e., the *contractor* agent, and the task that the contracting agent gives to the contractor. One such triplet defines a contract (giver, receiver and the task).

The values taken by the individual genes in the string depend on what that position represents, i.e. whether it corresponds to an agent or a task. More formally, there are *allele sets* that each gene can take values from. Since, in each triplet, the first two genes represent agents and the third represents task, the first two genes in the triplet take values from the allele set $S_1 = \{1, 2..., A\}$, where A is the number of agents, and the third gene in the triplet has the allele set $S_2 = \{t_{i1}, t_{i2}..., t_{i|T_i|}\}$, where t_{ik} is the id of the k^{th} task of the contracting agent i having a total of $|T_i|$ tasks. We have used the GALib package and employed the ArrayAlleleGenome class to enforce the above-mentioned allele set constraints on the values taken by individual genes. We also enforce the constraint that no agent can contract out the same task to two different agents. The objective function that is optimized (maximized) by our GA is the change in welfare function, ΔW, defined previously.

3. PROBLEM DOMAIN

For running experiments, we chose a problem domain from an interesting subclass of sub-additive domains, where the cost of performing a set of tasks can actually decrease by adding a new task to the task set! We call this subset of domains *ultra sub-additive domains*. For example, if tasks are nodes in a graph, and the job assignment is to find the minimal spanning tree in the associated graph, the solution cost can decrease if more nodes are added to the graph, i.e., if more tasks are added to the current set of tasks.

We now define the cost function used in our experiments, and which belongs to the class of ultra sub-additive domains. In this function, the cost of a given task j to agent i, c_{ij}, in an allocation is dependent on the other tasks in the allocation, T_i:

$$c_{ij} = \min_{k \in T_i, k \neq j} |f_i(j) - f_i(k)|.$$

The total cost for the allocation T_i to agent i is then $g_i(T_i) = \sum_{j \in T_i} c_{ij}$. The motivation behind the design of this function is that we should be able to compute the cost functions quickly. Moreover, if the f_i function maps the tasks to the number line, the optimal allocation would be to allocate those tasks to an agent for which the f_i mappings are densely clustered. For our experiments each task is an integer value in the range $1, \ldots, t$ and $f(x) = x$. Such a simple function suffices to demonstrate the efficacy of effective contracting to generate desired clusters of allocations. Our results, however, should hold for cost functions representing other sub-additive domains.

4. EXPERIMENTAL RESULTS

We now report the results of our experiments that were conducted to compare the performances of the two contract sequencing heuristics in finding the optimal contract within the limited number of contracts allowed. We allow a limit of ten contracts, i.e., k=10, and examine the welfare earned by the deterministic and stochastic contracting protocols at the end of the contracting process for different combinations of task and agent numbers. For each combination of tasks and agents, we use ten different randomly generated initial

[1]We assume that the centralized arbiter has perfect knowledge about the task allocations and cost functions of individual agents.

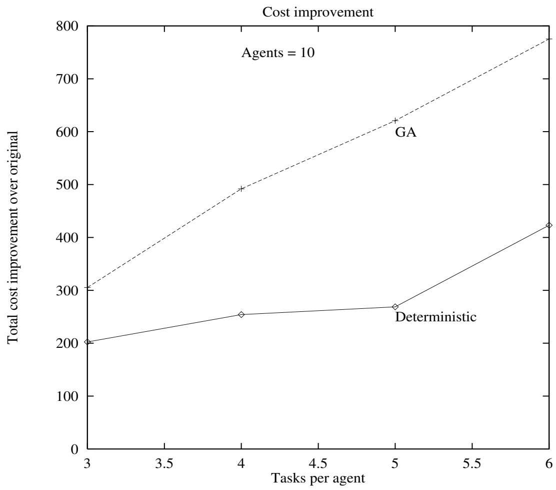

Figure 1: Welfare earned by the deterministic and stochastic contracting protocols (task per agent varying).

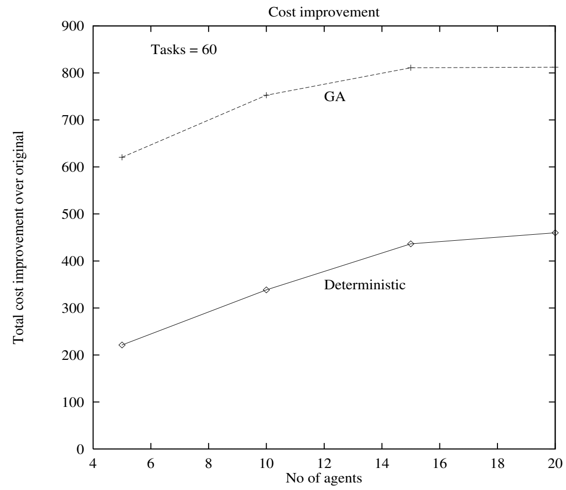

Figure 2: Welfare earned by the deterministic and stochastic contracting protocols (number of agents varying).

task allocations for each agent and calculate the welfare obtained at the end of the contracting process for each of the ten initial allocations. We report the welfare obtained by averaging the values obtained from the ten different initial allocations.

In the first set of experiments, the number of agents was fixed at 10 agents and the task per agent was varied from 3 to 6 in steps of 1. The GA parameters were selected as follows: population size = 100, probability of crossover = 0.9, probability of mutation = 0.005, crossover technique = "one point". GAs were ran for 1000 generations and averaged over 10 runs. The results from this experiment are summarized in Figure 1.

It is evident from the plot that with the increase in the per agent tasks, the welfare increases monotonically for both the deterministic and the stochastic protocols. This is due to the nature of the the cost function we have used to evaluate an allocation. Since the cost of an allocation is the sum of the minimum differences of each task cost from the rest, as number of tasks increases, so does the number of tasks initially allocated per agent and this leads to higher initial task set costs. The optimal costs, however, increase at a much lower rate, as optimal allocations involve consecutively numbered tasks being allocated to the same agent[2].

It is also evident from Figure 1 that the increase in global welfare produced by the GA for the same value of per agent tasks is substantially more than the corresponding values obtained using the deterministic contracting protocol. It is also observed from the graph that with increase in tasks per agent, the difference between the welfare obtained from the two protocols also increases. The rate of increment in welfare with per agent tasks using the OBGA is more than that obtained using the adapted original contracting protocol. As explained above, with increase in the per agent tasks, the diversity of the contracts increase. Since a GA searches in parallel several possible solutions (valid contracts), its performance is more effective in finding better contracts. It generates greater welfare than the deterministic protocol for the same values of agents and per agent tasks.

In the second set of experiments, the number of tasks was fixed at 60 tasks and the number of agents was varied from 5 to 20 in steps of 5. The GA parameters were selected as in the first set of experiments. The results of this experiment are summarized in Figure 2.

From the results we find an increase in welfare with an increase in the number of agents for both the deterministic and the stochastic protocols, with the welfare tending to level off as the number of agents continue to increase. It is also observed that the welfare values obtained using the OBGA is almost twice that obtained using the deterministic protocol for all values of agent numbers. OBGA outperforms the deterministic original contract protocol by consistently selecting more beneficial contract sequences.

It is also observed from the two figures that for the same values of task per agent, (e.g. task per agent = 3 in Figure 1 and number of agents = 20 in Figure 2), the welfare obtained are different. The welfare for the given values of task per agent and number of agents are more than twice in Figure 2 than that in Figure 1. We should note, however, that the total number of tasks and number of agents for the graph of Figure 2 are 60 and 20 while in Figure 1 they are 30 (3 tasks per agent * 10 agents) and 10 respectively. Hence the diversity of contracts possible in the former case is more than that in the latter. That is, with more number of agents and total tasks there is more diversity in initial allocation leading to higher initial costs. This creates the possibility for a contracting agent with the same number of tasks to be able to avail better contracting opportunities. This results in a bigger welfare earned for the same task by the contracting agent in the former case than the latter case.

5. RELATED WORK

The problem of distributed task allocation has been studied using the contract net framework [17]. In this framework,

[2]For the given evaluation function and task distribution used, there exists a number of allocations which optimize the social welfare metric. This includes an allocation where all tasks are assigned to any one agent. However, only a few of these solutions can be reached from the initial allocation state using only O-contracts.

cooperative and self-interested agents having different preferences contract tasks out to other agents. A contract, effectively, is equivalent to mutual selection of the agents (the contracting agent and the contractor) involved in it. The agents participate in a bidding and awarding sequence to obtain contracts, much like auctions. Agents having different local preferences, thus, have the potential of obtaining the tasks that match their preferences, leading to a more effective performance of the group without requiring a central mediator [7, 12].

Agents having different expertise levels and having a heterogeneous mixture of tasks may enter into cooperative coalitions to increase the performance level of the coalition as a whole. Coalition structure generation among cooperative agents is the strategy of finding the right group of agents that one should associate with to get the job done with minimal cost [13, 15]. Various models and algorithms have been studied by multi-agent systems researchers to obtain efficient and effective coalitions [11, 14].

A different approach of coordinating the task of solving problems by a group of agents is that of argumentation [4]. Negotiation through argumentation provides the agents with opportunities to resolve mutual conflicts. Self-interested rational agents come to know about the resources of other agents and hence can decide for/against an exchange of tasks among themselves. A rational agent can, using argumentative backup, win the cooperation of another non-cooperative agent [16].

Centralized and de-centralized allocations of multiple resource items among agents has been the focus of combinatorial auctions [3, 5, 9]. Extensive research in this area have been conducted resulting in valuable insight regarding possible algorithms that can be employed in different auction protocols to achieve effective allocations.

6. CONCLUSION

A number of agent and agent system design problems can be mapped into combinatorial optimization problems with huge state spaces. For a number of such problems, there are no known polynomial time optimal algorithms. Deterministic approximate algorithms often have the advantage of guaranteeing bounds about the quality of the solutions produced [11]. In some cases, such guarantees require the availability of significant domain knowledge which may not be readily available. In other cases, the bounds on optimality may not be tight. In either case, we believe there is a need for an algorithm that perform better given a bound on computational resources.

We argue for the need of more widespread application of stochastic optimization algorithms in agent and agent system design problem. We believe that a number of computationally challenging multi-agent systems problems can be fruitfully addressed by the use of such techniques.

In our previous work, we have shown that a GA-based algorithm can provide better results compared to deterministic algorithms with guarantees given bounds on the number of evaluations of potential coalition structures for the optimal coalition structure generation problem [13]. In this paper, we have shown that a GA implementation outperforms a reasonable greedy heuristic on the *optimal n-sequence contract* selection problem. In this work, we have used offline algorithms to heuristically generate desirable contract sequences. The GA and the deterministic algorithms are both heuristic approaches that have very distinctive bias of searching the space of contract sequences. We allow the deterministic heuristic to complete its biased search in the contract sequence space, but run the GA for a fixed number of generations. It would be instructive to compare the relative performance of these algorithms under bounds of computational resources.

We reiterate the fact that the centralized approaches presented here are designed to generate only individually rational contracts. This is necessary as the central arbiter can only suggest mutually beneficial contracts to participating agents, but cannot enforce any exchanges. Given that there has been little previous work on heuristics to find fixed length contract sequences that quickly improve both individual and social welfare, we believe that the current work provides a baseline to design and evaluate better contract sequencing algorithms. We do believe that the centralized approaches used in this paper should be followed up, or approximated, with distributed approaches that do not require a central arbiter with perfect information.

As a logical extension to the set of experiments we have already ran, We plan to run a set of experiments by varying the number of contracts allowed for a given number of agents and tasks. This would allow us to evaluate the relative progress in improving social welfare as more and more contracts can be executed.

Acknowledgments: This work has been supported in part by an NSF CAREER award IIS-9702672.

7. REFERENCES

[1] M. R. Andersson and T. W. Sandholm. Time-quality tradeoffs in reallocative negotiation with combinatorial contract types. In *Proceedings of the Sixteenth National Conference on Artificial Intelligence*, pages 3–10, Menlo Park, CA, 1999. AAAI Press/MIT Press.

[2] J. H. Holland. Genetic algorithms. *Scientific American*, 267:44–50, 1992.

[3] H. H. Hoos and C. Boutilier. Solving combinatorial auctions using stochastic local search. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence*, pages 22–29, Austin, TX, 2000.

[4] H. Jung, M. Tambe, and S. Kulkarni. Argumentation as distributed constraint satisfaction: Applications and results. In *Proceedings of the Fifth International Conference on Autonomous Agents*, pages 324–331, Montreal, Canada, 2001.

[5] K. Leyton-Brown, Y. Shoham, and M. Tennenholtz. An algorithm for multi-unit combinatorial auctions. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence*, pages 56–61, Austin, TX, 2000.

[6] S. D. Parsons and N. R. Jennings. Negotiation through argumentation-A preliminary report. In *Proceedings of the Second International Conference on Multi-Agent Systems*, pages 267–274, Kyoto, Japan, 1996.

[7] T. Sandholm. An implementation of the contract net protocol based on marginal cost calculations. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, pages 256–262, 1993.

[8] T. Sandholm. Contract types for satisficing task allocation: I theoretical results. In *Proceedings of the AAAI 1998 Spring Symposium: Satisficing Models.*, pages 68–75, 1998.

[9] T. Sandholm. An algorithm for optimal winner determination in combinatorial auctions. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 542–547, 1999.

[10] T. Sandholm. Approaches to winner determination in combinatorial auctions. *Decision Support Systems*, 28(1-2):165–176, 2000.

[11] T. W. Sandholm, K. Larson, M. Andersson, O. Shehory, and F. Tohme. Anytime coalition structure generation with worst case guarantees. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, pages 46–53, Menlo Park, CA, 1998. AAAI Press/MIT Press.

[12] S. Sen and E. H. Durfee. A contracting model for flexible distributed scheduling. *Annals of Operations Research*, 65:195–222, 1996.

[13] S. Sen and P. S. Dutta. Searching for optimal coalition structures. In *Proceedings of the The Fourth International Conference on Multi-Agent Systems*, pages 287–292, 2000.

[14] O. Shehory and S. Kraus. A kernel-oriented model for coalition-formation in general environments: Implementation and results. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, pages 133–140, August 1996.

[15] O. Shehory and S. Kraus. Feasible formation of coalitions among autonomous agents in non-super-additive environments. *Computational Intelligence*, 15(3):218–251, 1999.

[16] C. Sierra, N. R. Jennings, P. Noriega, and S. Parsons. A framework for argumentation-based negotiation. In *Proceedings of the Fourth International Workshop on Agent Theories, Architecture and Languages*, pages 167–182, Rhode Island, USA, 1997.

[17] R. Smith. The contract net protocol: High-level communication and control in distributed problem solver. *IEEE Transactions on Computers*, 29(12):104–1113, 1980.