

An Efficient Protocol for Negotiation over Multiple Indivisible Resources

Sabyasachi Saha and Sandip Sen

Department of Mathematical & Computer Sciences

The University of Tulsa

{sabyasachi-saha, sandip}@utulsa.edu

Abstract

We study the problem of autonomous agents negotiating the allocation of multiple indivisible resources. It is difficult to reach optimal outcomes in bilateral or multi-lateral negotiations over multiple resources when the agents' preferences for the resources are not common knowledge. Self-interested agents often end up negotiating inefficient agreements in such situations. We present a protocol for negotiation over multiple indivisible resources which can be used by rational agents to reach efficient outcomes. Our proposed protocol enables the negotiating agents to identify efficient solutions using systematic distributed search that visits only a subspace of the whole solution space.

1 Introduction

Negotiation is the most preferred approach for resolving conflicts in human and agent societies. Automated negotiation is being increasingly used in different multiagent domains including robotics, network bandwidth allocation, space applications, etc. where agents need to reach agreements on the allocation of one or more shared resources [Chevalyere *et al.*, 2006; Endriss *et al.*, 2006]. In multi-issue negotiation¹, agents with divergent preferences can cooperate to reach agreements beneficial for both agents. But when the preferences are not common knowledge, self-interested agents often fail to explore win-win possibilities using existing protocols and end up with inefficient agreements. Hence, there is a need for negotiation protocols which can lead rational agents to mutually preferred agreements. By efficient or optimal solution, we refer to a solution which is *Pareto optimal* [Brams and Taylor, 1996]. An outcome is Pareto-optimal or Pareto-efficient if there exists no other outcome which is at least as good as this outcome for both agents and strictly better for at least one agent. Another desirable property of an agreement is its fairness. An agreement can be considered to be fair if it is *envy-free*. An agreement is *envy-free* iff each agent's utility for its own share is at least as much as its utility for the share allocated to the other agent [Brams and Taylor, 1996]. Often

¹We identify negotiation over multiple resources as a special case of multi-issue negotiation.

no envy-free agreement exists when resources are indivisible. We, therefore, argue that the goal of the negotiation process should be to produce Pareto optimal agreements that maximizes *egalitarian social welfare*, the minimum utility of the negotiating agents.

When agents' preferences are common knowledge, negotiating agents can reach an efficient point of agreement using distributed protocols like *one-step monotonic concession protocol* or *monotonic concession protocol* [Rosenschein and Zlotkin, 1994] where each agent searches the entire space of possible agreements. Most negotiations in real life, however, take place under incomplete information where agents do not have complete knowledge about the preferences of the other agent. Negotiation protocols for such negotiation scenarios are proposed by Brams and Taylor [Brams and Taylor, 2000]. But their protocols do not guarantee optimality for allocation of indivisible resources. These protocols produce close to optimal agreements only for negotiations involving uncorrelated resources, *i.e.*, situations where the utility of possessing two resources is the sum of the utilities of possessing each individual resource. But real-life negotiations typically involve correlated resources. These protocols can produce very inefficient agreements for such negotiations.

Our objective is to design a protocol for negotiation over multiple indivisible resources, both correlated or uncorrelated, which can lead participating agents to efficient outcomes, without any prior information about opponent's preferences. This situation is common in E-commerce domains, where the number and diversity of agents is so large that an agent may not have any knowledge of the preferences of any particular agent it has to negotiate with. We present a three-phase protocol that guarantees Pareto optimal agreements between rational agents and ensures that the agreement reached will be at least as fair as the agreements reached by the protocols proposed by Brams and Taylor².

Our work has some similarity with the work on contract negotiation by Sandholm [Sandholm, 1998] and Endriss *et al* [Endriss *et al.*, 2003; 2006] in the context of task allocation, where agents negotiate to reallocate tasks. But these approaches are geared towards identifying, under complete information, necessary deals or exchanges required among

²We discuss the protocol for two agents and prescribe its extension for multiple agents in Section 8.

agents to reach efficient allocation from an arbitrary initial allocation. We, on the contrary, want to find mechanisms that lead agents, without any prior information about opponent's preferences, to efficient outcomes revealing as little information as possible.

Another advantage of our proposed protocol is that each agent is required to search only a subset of the whole set of possible allocations. Even under complete information, finding out necessary exchanges required to reach Pareto optimal solution requires $O(2^n)$ computations, where n is the number of issues [Chevalyre *et al.*, 2006]. Fortunately, in most negotiations, the number of issues is not very high (less than 20). Although the worst case complexity of our algorithm is $O(2^n)$, it requires significantly less search compared to exhaustive search in most negotiation instances.

2 Related research

We consider the problem of designing a negotiation protocol for two or more agents so that they can reach a mutually acceptable allocation of multiple indivisible resources or tasks without any external intervention. When agents' preferences are common knowledge, each agent can search the entire space of allocations and find out possible points of agreement. In case of multiple points there can be a problem of selecting the point of agreement. However, using distributed protocols like *one-step monotonic concession protocol* negotiating agents can reach an efficient point of agreement [Rosenschein and Zlotkin, 1994]. According to this protocol, both agents simultaneously make an offer. The offer that maximizes the product of the agents utilities will be chosen as the point of agreement. It can be shown that rational agents will reach Pareto optimal solution using this protocol. Here, each agent will choose the offer that has maximal product of utilities. So, each agent needs to know the other agent's complete preferences and search for all possible allocations. This protocol fails when agents do not have complete knowledge of other agent's preferences. Similar is the case with *monotonic concession protocol*. Most of the negotiations in real life, however, take place under incomplete information where agents do not have complete knowledge about the preferences of the other agent. In this paper, we focus on negotiations under incomplete information settings.

In similar research on multiagent resource allocation (MARA), researchers have investigated the efficiency of allocations of multiple tasks or resources among multiple agents [Endriss *et al.*, 2003; 2006]. Existing research have focused on identifying necessary exchanges, under complete information, to reach Pareto optimal solution from an initial allocation. In our research, agents do not have any information about the other agent's preferences and our objective is to design a protocol using which agents can reach an optimal agreement.

3 Negotiation framework

We formally describe a representative negotiation scenario for allocation of multiple indivisible resources as a 3-tuple $\langle \mathcal{A}, R, \mathcal{U} \rangle$, where $\mathcal{A} = \{1, 2\}$ is the set of agents, $R = \{r_1, r_2, \dots, r_H\}$, $H \geq 2$, is the set of H indivisible resources

whose allocation are being negotiated, and $\mathcal{U} = \{U_1, U_2\}$ is the set of utility functions, where U_i is the utility function of agent i . Each resource is considered as a negotiation issue. The negotiating agents must agree on the allocation of the resources. We assume a monotonic scenario, where an agent's utility increases if it possesses one more resource. Therefore, each agent wants to obtain each of the resources and for different bundles of resources, it has different utilities. The valuations for all possible bundles of resources is known as the agent's complete preference, which is not known to the other agent.

We now define some terms formally:

Outcome: An outcome O is an allocation of the resources to different agents. Formally, $O : R \mapsto R_{\{0,1\}}^{2 \times H}$, where $R_{\{0,1\}}^{2 \times H}$ is a $2 \times H$ matrix, whose (i, j) cell content, $x_{ij}^O \in \{0, 1\}$, presents agent i 's allocation of the j^{th} resource. $x_{ij}^O = 1$ if j^{th} resource is allocated to agent i , otherwise $x_{ij}^O = 0$. We can write $R_{\{0,1\}}^{2 \times H}$ as $(R_1^O, R_2^O)'$, where $R_i^O = (x_{i1}^O, \dots, x_{iH}^O)'$ is the column vector representing the allocation of resources to agent i under outcome O .

Utility: For agent i , utility of an outcome O is defined as: $U_i^O : R_{\{0,1\}}^{2 \times H} \mapsto \mathbb{R}^+$, where \mathbb{R}^+ is the set of non-negative real numbers.

Pareto optimality: An outcome O is Pareto optimal if there exists no other outcome O' such that all agent's utility in O' is at least as good as their utility in O and at least one agent's utility is more in O' . If outcome O is Pareto optimal, there cannot exist any outcome O' such that, $U_i^{O'} \geq U_i^O$ for both $i = 1, 2$ and $U_i^{O'} > U_i^O$ for at least one i .

Individually rational agent: An individually rational agent i will accept an agreement O' instead of O only if $U_i^{O'} \geq U_i^O$.

Egalitarian social welfare: For an outcome O , egalitarian social welfare, $ESW(O)$, is defined as the minimum utility among the agents under that outcome, *i.e.*, $ESW(O) = \min_{i \in \mathcal{A}} \{U_i^O\}$.

We now present an example of a bilateral negotiation scenario involving 4 indivisible resources: A, B, C , and D . The utilities of the two agents for different allocations are given in Table 1. Both agents want to maximize their individual utilities.

4 Existing negotiation protocols

In this section, we present existing negotiation protocols for allocation of indivisible resources among agents when agents do not have any knowledge of the preferences of the other agents:

Strict alteration: In this protocol, agents take alternate turns and in each turn an agent selects one resource from the set of resources not yet allocated. After an agent selects a resource, the resource will be removed from the set [Brams and Taylor, 2000]. The advantage of this protocol is its simplicity and the time required to reach an

Allocation	U_1	U_2
$(\{\},\{ABCD\})$	0	20
$(\{A\},\{BCD\})$	6	12
$(\{B\},\{ACD\})$	8	14
$(\{C\},\{ABD\})$	5	16
$(\{D\},\{ABC\})$	7	18
$(\{AB\},\{CD\})$	9	8
$(\{AC\},\{BD\})$	10	9
$(\{AD\},\{BC\})$	11	8
$(\{BC\},\{AD\})$	10	8
$(\{BD\},\{AC\})$	11	13
$(\{CD\},\{AB\})$	12	15
$(\{ABC\},\{D\})$	12	6
$(\{ABD\},\{C\})$	14	7
$(\{ACD\},\{B\})$	16	7
$(\{BCD\},\{A\})$	14	8
$(\{ABCD\},\{\})$	20	0

Table 1: Utilities of two agents for different allocations where (x,y) represents allocation of item sets x and y to agent 1 and 2 respectively.

agreement. The agreements reached, however, are often very inefficient.

Balanced alteration: This protocol is used to improve fairness. Here it is assumed that an agent who first chooses a resource has an advantage over the agent who chooses second. Hence, the second agent will have the opportunity to choose third, and so on [Brams and Taylor, 2000]. Therefore, one agent gets to choose in turns 1, 4, 5, 8, . . . , and the other agent in turns 2, 3, 6, 7, This protocol has similar advantages and disadvantages as the strict alternate protocol.

Exchange auctions: This protocol is an exchange-based extension of the Contract-Net protocol [Golfarelli *et al.*, 1997]. This protocol assumes an initial allocation of the resources to the agents. The protocol allows the agents to exchange resources to improve the allocation. An agent announces one or more resources for exchanges and other agents bid one or more resources that they are ready to exchange. But this protocol does not guarantee Pareto optimality unless repeated infinitely.

5 Our proposed negotiation protocol

In this section we present our proposed protocol: Protocol to reach Optimal agreement in Negotiation Over Multiple Indivisible Resources (PONOMIR). We first define some concepts required to describe our protocol.

Negotiation tree: We assume that the issues or resources are ordered in some way, *e.g.*, lexicographically. In this paper, we conceptualize the allocations of the resources as a tree, known as *negotiation tree*. For a negotiation over H resources, the maximum possible depth of the negotiation tree is H , the root being at level 0. The root node represents a null allocation to the agents and then each successive level represents allocation of the next

resource in the order chosen. For a bilateral negotiation, the negotiation tree is a binary tree³. The left and right branches at the l^{th} level imply that the l^{th} resource will be allocated to agent 1 and 2 respectively. Each leaf node at level H represents one possible allocation of the resources and the path to reach that leaf node from the root specifies the allocation of all the resources. Such a negotiation tree is shown in Figure 1 for the negotiation scenario presented in Table 1. A negotiation tree is created by the negotiating agents in a distributed manner. It starts with a root with $id = 0^4$. The tree is then created in a top-down process, where at any level agent 1 can only create the right child of a node in the previous level of the tree. Similarly, agent 2 can only create the left child nodes. Each agent, however, may choose not to create any node, which it can, at any level. If an agent does not create a node, that node will be marked as *black* node and no further sub-tree will be created from that *black* node. Therefore, a black node prunes that node and its descendants. The distributed formation of a negotiation tree can be implemented using message-passing. At any level, each agent knows the nodes created at the previous level and the nodes that either of the agents can create at this level. Hence both agents can send a message to the other agent stating the nodes it has created at this level before proceeding to the next level.

Best possible agreement (BPA): At each node of the negotiation tree, each agent has a best possible agreement which is the allocation where the resources until the current level are allocated according to the path from the tree root to this node and the remaining resources are allocated to this agent.

Now we present our three-phase protocol, PONOMIR. The first phase of PONOMIR consists of a primary allocation procedure using any one of *strict alteration* or *balanced alteration* protocol. We have used *strict alteration* protocol in this paper. The second phase consists of distributed formation of the negotiation tree by the negotiating agents. After the second phase, the agents will be left with few probable agreements. In the third phase, agents reach the final Pareto optimal solution by exchanging offers.

First phase of PONOMIR:

Step 1: A random device chooses one of the two agents and marks this agent as S . Denote the set of resources yet to be negotiated by G . Initially, $G = R$.

Step 2: Now, S will choose one of the remaining resources, $C \in G$. C is allocated to S .

Step 3: Mark the other agent as S and update G to $G - \{C\}$. If $|G| > 1$ return to Step 2, otherwise stop.

After the first phase, there is an initial allocation, L , of the resources as suggested by the *strict alteration protocol*.

³If n agents are negotiating, the tree will be n -ary.

⁴The node id of any other nodes in the tree will be $2 \times id(\text{parent}) + 1$ if it is a left node and $2 \times id(\text{parent}) + 2$ if it is a right node.

For this allocation L , agents have corresponding utilities $U_i^L, i = 1, 2$. If no mutual improvement is possible in the subsequent phases, agents will have to agree on this allocation⁵. This phase ensures that the agreement reached will be at least as good as L .

Second phase of PONOMIR: This phase involves the distributed generation of the negotiation tree by the negotiating agents.

Step 1: Let l denotes the level of the negotiation tree. Set $l = 0$. The root node is created with $id(root) = 0$.

Step 2: Agents 1 and 2, respectively, can create the right and left child nodes of each node in the level, l , of the tree and send a message to the other agent.

Step 3 Increase l by one. If $l < H$ and no node is created at level l , the negotiation terminates and the final allocation will be L . If $l < H$ and there is at least one node in this level of the tree, go to Step 2. If $l = H$, stop the second phase, collect all the allocations corresponding to the nodes at level H and proceed to the third phase. We refer to this set of leaf nodes (or allocations) as Q .

After this phase, the agents will be left with a small number of probable agreements in Q . The final agreement will be chosen in the third phase. Note that, at each level agent 1 can create only right child nodes and agent 2 can create only left child nodes. A right child implies that the resource representing that level is allocated to agent 2. Since we assume a monotonic scenario, the utility of the BPA of agent 2 will not decrease for allocating a resource to agent 2, but the utility of the BPA of agent 1 may decrease. So, it is sufficient if agent 1 only checks whether if it is interested in expanding this right subtree. Similarly agent 2 only needs to check the left child. If the utility of the BPA for an agent at any node is less than its utility under the initial allocation, L , it is sure that any allocation under the subtree starting at that tree node will always produce utility less than its utility under initial allocation L . An individually rational agent has no incentive to create that node. If a node is not created at any level by the corresponding agent, no further subtree will be generated from it. This implies that all the allocations under that subtree are discarded. If agents are individually rational, for each allocation in Q , the utilities of both agents are at least as much as their corresponding utilities in L . This is because all allocations which produce less utility to any of the agents will be discarded by that agent in the distributed tree generation process.

Third phase of PONOMIR: In this phase, agents will select the final agreement from allocations in the set Q . Let us define a set F , as the set of final agreements. Initially, it contains only the initial allocation L , i.e., $F = \{L\}$.

Step 1: One agent is chosen randomly to start. Mark this agent as S and the other agent as S' . Now, S needs to propose an allocation q from Q .

Step 2: S' can remove any other allocation O from the set Q and F if $U_{S'}^O \leq U_{S'}^q$. Update F to $F \cup \{q\}$.

Step 3: If Q is not empty, then swap the agents S and S' and go to Step 2. Otherwise, the set F contains the set of final agreements. If only one element remains in F , it will be selected as the final agreement. Otherwise, any one of them will be chosen randomly as the final agreement.

6 Properties of PONOMIR

The objective of an efficient protocol is to lead rational agents to efficient agreements. As we have discussed earlier, our goal is to obtain Pareto optimal agreements and to increase fairness as much as possible. As a measure of fairness, we use *egalitarian social welfare*.

Our proposed negotiation protocol, PONOMIR, is not strategy-proof and does not guarantee Pareto optimal agreements if agents are arbitrarily risk seeking. However, in this paper we consider a completely uncertain negotiation scenario where participating agents do not have any knowledge about the preferences of the other agents and agents do not want to reveal their utilities for different allocations. The rational behavior of the agents, under such complete uncertainty, depends on their risk attitudes. Bossert introduced and argued for uncertainty aversion for decision making under such complete uncertain scenarios [Bossert, 1997]. In this negotiation scenario, we assume that the rational agents will be *cooperative-individually rational*. An agent is *cooperative-individually rational* if it satisfies two properties: i) it does not take any risky action that can lead to an agreement which produces less utility than what it is already assured of, and, ii) if there exists two agreements which produces same utility to it but different utility to the opponent, then it will agree to accept any of the agreement proposed by the opponent. PONOMIR guarantees Pareto optimal agreements if the participating agents are *cooperative-individually rational*. The agreements reached also guarantees at least as much *egalitarian social welfare* as the agreements reached by the existing protocols. Since the agents are individually rational, both of them will discard all agreements that produce utilities less than the utility from the initial allocation, L , obtained in the first phase and hence the *egalitarian social welfare* will never decrease.

Proposition 1: *The agreement reached by cooperative-individually rational agents using the PONOMIR protocol is Pareto optimal.*

To prove this proposition, we first prove the following lemmas:

Lemma 1: *An allocation O will not be discarded in the second phase, only if $U_i^O \geq U_i^L, \forall i = 1, 2$, where, L is the initial allocation after first phase.*

Proof: There are two parts of this proof. In the first part we will prove that any allocation O' , for which there is at least one agent i such that $U_i^{O'} < U_i^L$, will be discarded in the second phase. In the second part we will show that the remaining allocations will not be discarded. For

⁵We assume that both agents prefer to agree on this allocation over disagreement.

an allocation like O' , such that $U_i^{O'} < U_i^L$, there will be at least one level in the distributed tree construction process when agent i will realize that its utility from the BPA at that node is less than U_i^L . Therefore, the agent has no incentive to create that node and hence the allocation will be discarded. The proof of the next part of the lemma is clear from the properties of the agents. Since they are *cooperative-individually rational*, they do not want to discard any possibility which can lead to an allocation which produces utility at least as much as the utility produced by the initial allocation.

Lemma 2: *While proposing an allocation in the third phase, an agent always propose the allocation that produces highest utility among the set of remaining allocations Q .*

Proof: In a complete uncertain environment, the likelihood of the opponent accepting any allocation is not known a priori. Hence, an allocation of higher utility to an agent should be proposed before an allocation with a lower utility. Therefore, a *cooperative-individually rational* agent proposes the allocation with the highest utility from the set of remaining allocations Q .

Lemma 3: *After an agent proposes an allocation in the third phase, the other agent will remove all allocations which produce less utility to it from the sets Q and F .*

Proof: An individually rational agent will remove all allocations that produce less utility to it compared to the offer proposed by the other agent as it is guaranteed to get the allocation proposed by the other agent. But if it does not remove any one of those it may end up with one of them which will not be individually rational. It will also remove the allocations which produce equal utility to it. According to the previous lemma a rational opponent proposes the offer which produces highest utility to the opponent. This implies that an allocation which is not offered by the opponent produces less than or equal utility to the opponent. Therefore even if it produces equal utility to this agent, it is either equivalent or Pareto dominated by the allocation proposed by the opponent. So, the agent should remove this from the set.

Also, the agent will not remove other allocations with higher utility as this may eliminate possible agreements which have better utility for it.

Proof of Proposition 1: From Lemmas 1, 2, and 3, it is evident that any allocation in the final set F is not Pareto dominated by any of the other possible allocations. Therefore, each of the allocations is Pareto optimal.

7 Experiments

In this section, we present experimental results to show the effectiveness of PONOMIR to reduce search requirement for the negotiating agents. Before that we present an illustration of our proposed protocol.

7.1 An illustration

In this subsection, we demonstrate the execution of PONOMIR on the example presented in Table 1. In the first




Figure 1: Negotiation tree formed for example in Table 1.

phase, agents take alternate turns to choose one resource. Let us assume that an agent, at first, chooses the resource which produces highest utility to it among the resources yet to be allocated. Then onwards, it will choose the resource, from the set of resources not yet allocated, which produce highest utility together with other resources already allocated to it. Agents with this strategy produce an initial allocation, L , where $L = (\{B, D\}, \{A, C\})$, i.e., resources B and D are allocated to agent 1 and resources A and C are allocated to agent 2. The corresponding utilities to the agents are 11 and 13 respectively.

Figure 1 shows the formation of the negotiation tree in the second phase. For example, consider the formation at level 1. While verifying the left child of the root node, which implies that the resource A is allocated to agent 1 (as shown in the corresponding edge of the graph), agent 2 observes that if resource A is allocated to agent 1, the best possible agreement for agent 2 will be $(\{A\}, \{B, C, D\})$ where all other resources will be allocated to it and the corresponding utility is 12, less than U_2^L , its utility from the initial allocation, L . So, it decides not to create the left child and therefore it is marked as a *black* node and no subtree will be generated from it. Hence, all allocations under this subtree are discarded from the set of possible agreements. Note that in the Figure 1, only one utility is given at each node which is the utility of the BPA of the agent who is deciding whether to create that node. The corresponding utility for the other agent is blank as it is not known to the deciding agent. At the end of the second phase only two allocations are chosen for the third phase: $O_1 = (\{B, D\}, \{A, C\})$ and $O_2 = (\{C, D\}, \{A, B\})$. Observe that both agents need to find the utility for only 6 nodes, instead of $2^4 = 16$ nodes. In the third phase one agent is chosen to propose one allocation. The chosen agent proposes the second allocation, O_2 , as that produces highest utility to it and the other agent will remove O_1 from the final set of chosen agreements, F . So, the final agreement will be $(\{C, D\}, \{A, B\})$, which is Pareto optimal and with an *egalitarian social welfare* 12, which is the highest possible *egalitarian social welfare* in the entire space of allocations.

7.2 Experimental results

We tested PONOMIR on a large number of scenarios varying agents' preferences and number of issues. We show the average reduction in the search effort required by each agent. We consider monotonic utility scenarios, which implies that for any agent the utility of a bundle of n resources is more than any of the n possible subsets of $(n - 1)$ resources. We vary the number of resources n . For each n , we generate 10,000 random examples of agent preferences and execute the negotiations using the PONOMIR protocol. We observed that they always reach Pareto optimal agreements and significant percentage of these produce maximum possible egalitarian social welfare. For the agreements where the fairness are not maximum, they are close to this maximum value. But, this result depends on the strategies used by the agents in the first phase of PONOMIR. It is, however, guaranteed that the agreement reached will be Pareto optimal. We have also calculated the average percentage of allocations that each agent searched during each negotiation process compared to all allocations that an agent needs to search to find a Pareto optimal outcome using protocols like *one-step monotonic concession protocol* under complete information. Table 2 shows significant reduction in average search by each individual agents using PONOMIR.

# of resources	% of allocations searched by each agent
4	53.8
8	50.4
12	47.4
14	45.9
16	44.1

Table 2: Reduction in search by agents

8 Conclusion and Future work

Existing research in multiagent systems for allocating multiple indivisible resources have focused on identifying necessary exchanges, under complete information, to reach Pareto optimal agreements. In a large number of real-world negotiations, however, agents are required to negotiate under a completely uncertain environment. In this paper, we propose an efficient three-phase negotiation protocol, PONOMIR, for negotiation over multiple indivisible resources when an agent does not know the preferences of the other agents. We have shown that PONOMIR leads rational agents to Pareto optimal agreements and the negotiating agents do not need to explicitly reveal their utilities for different allocations. The agreement reached depends on the initial allocation after phase 1, but will always be Pareto optimal if the negotiating agents are *cooperative-individually rational*.

The last two phases of PONOMIR can also be used separately for producing Pareto optimal outcomes given any initial allocation of the resources to the agents. This is a useful property of PONOMIR as in several multiagent resource allocation problems the negotiating agents may initially possess some resources. The second and third phase of PONOMIR guarantees Pareto optimal outcome in such situations.

We have also experimentally shown the effectiveness of the PONOMIR protocol for reducing the search costs of the participating agents. PONOMIR, however, cannot guarantee a timely agreement if the number of resources is very high, *e.g.*, 100. Fortunately most real-life negotiations do not involve very high number of resources and PONOMIR produces agreements quickly for 20 resources or less. We are currently working on improving the time complexity of our protocol to address the problem of negotiations over very large number of resources.

A major advantage of PONOMIR is its extensibility to multilateral negotiation. We have presented PONOMIR for bilateral negotiation, but this can be extended to multilateral negotiation with minor modifications.

We have experimentally observed that the fairness of the agreements reached using this protocol are high. But we can only guarantee that the fairness will not be reduced after the first phase. We are working to develop a novel first phase of our protocol to ensure increased guaranteed fairness of the negotiation outcome.

Acknowledgments: This work has been supported in part by an NSF award IIS-0209208 and an NASA EPSCoR RIG.

References

- [Bossert, 1997] Walter Bossert. Uncertainty aversion in non-probabilistic decision models. *Mathematical Social Sciences*, 34(3):191–203, 1997.
- [Brams and Taylor, 1996] Steven J. Brams and Alan D. Taylor. *Fair Division: From cake-cutting to dispute resolution*. Cambridge University Press, Cambridge: UK, 1996.
- [Brams and Taylor, 2000] Steven J. Brams and Alan D. Taylor. *The Win-Win Solution: Guaranteeing Fair Shares to Everybody*. W. W. Norton and Company, 2000.
- [Chevalleyre *et al.*, 2006] Yann Chevalleyre, Paul E. Dunne, Ulle Endriss, Jerome Lang, Michel Lemaître, Nicolas Maudet, Julian Padget, Steve Phelps, Juan A. Rodriguez-Aguilar, and Paulo Sousa. Issues in multiagent resource allocation. *Informatica*, 30:3–31, 2006.
- [Endriss *et al.*, 2003] Ulle Endriss, Nicolas Maudet, Fariba Sadri, and Francesca Toni. On optimal outcomes of negotiations over resources. In *AAMAS-03*, pages 177–184, New York, NY, USA, 2003. ACM Press.
- [Endriss *et al.*, 2006] Ulle Endriss, Nicolas Maudet, Fariba Sadri, and Francesca Toni. Negotiating socially optimal allocations of resources. *Journal of Artificial Intelligence Research*, 25:315–348, 2006.
- [Golfarelli *et al.*, 1997] Matteo Golfarelli, Dario Maio, and Stefano Rizzi. A task-swap negotiation protocol based on the contract net paradigm. Technical Report 005-97, CSITE, University of Bologna, 1997.
- [Rosenschein and Zlotkin, 1994] Jeffrey S. Rosenschein and Gilad Zlotkin. *Rules of Encounter*. MIT Press, 1994.
- [Sandholm, 1998] Tuomas W. Sandholm. Contract types for satisficing task allocation: I theoretical results. In *AAAI Spring Symposium on Satisficing Models*, 1998.