

Coalition Formation with Uncertain Heterogeneous Information*

Sarit Kraus

Dept. of Computer Science
Bar-Ilan Univ. Ramat-Gan, 52900
Israel and Univ. of Maryland College
Park, MD
sarit@cs.biu.ac.il

Onn Shehory

IBM Research Lab in Haifa
Haifa University
Mount Carmel, Haifa
31905 ISRAEL
onn@il.ibm.com

Gilad Taase

Dept. of CS
Bar-Ilan University
Ramat-Gan,
52900 Israel

ABSTRACT

Coalition formation methods allow agents to join together and are thus necessary in cases where tasks can only be performed cooperatively by groups. This is the case in the Request For Proposal (RFP) domain, where some requester business agent issues an RFP - a complex task comprised of sub-tasks - and several service provider agents need to join together to address this RFP. In such environments the value of the RFP may be common knowledge, however the costs that an agent incurs for performing a specific sub-task are unknown to other agents. Additionally, time for addressing RFPs is limited. These constraints make it hard to apply traditional coalition formation mechanisms, since those assume complete information, and time constraints are of lesser significance there.

To address this problem, we have developed a protocol that enables agents to negotiate and form coalitions, and provide them with simple heuristics for choosing coalition partners. The protocol and the heuristics allow the agents to form coalitions in the face of time constraints and incomplete information. The overall payoff of agents using our heuristics is very close to an experimentally measured optimal value, as our extensive experimental evaluation shows.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multi-agent Systems, Coherence and Coordination, Intelligent Agents.

General Terms

Algorithms, Design, Economics, Experimentation.

Keywords

Coalition formation, Incomplete information, RFP, task allocation, experimentation.

1. INTRODUCTION

Coalition formation is an important cooperation method in multi-agent systems. Within coalitions, agents may be able to jointly perform tasks that they would otherwise be unable to perform, or will perform poorly. To allow agents to form coalitions, one

should devise a coalition formation mechanism that includes a protocol as well as strategies to be implemented by the agents given the protocol. Coalition formation mechanisms proposed to date (e.g., [5][11][13]) commonly provide these, however they include several restrictive assumptions, which do not hold in real-world domains where coalitions are necessary. In this study we relax some of these assumptions thus arrive at an automated coalition formation mechanism better suited for real domains. In particular, we do not assume complete information. Rather, we assume, in similarity to practical economic situations, that agents value tasks differently—resulting in multiple valuations of a specific task—and that agents do not know the value of a task to other agents, although they may have a rough idea of the range of values the task may have. Additionally, we assume that the coalition formation process, as an economic process, is bounded in time, and that the value of a task is discounted as time elapses during the process. These assumptions, which are unique to our solution, are necessary for providing a solution applicable to a real-world coalition formation.

Under the assumptions of incomplete information, heterogeneous task valuations, and short time for completion of the coalition formation process, traditional coalition formation mechanisms are inadequate. Therefore we need to devise a new mechanism. Ideally, we would like a coalition formation mechanism to not only allow agents form coalitions for joint task execution, but also arrive at a coalition configuration which is optimal (in terms of utility maximization), stable, and fair. However, the computational complexity required for such solutions is exponential (see [8]). Therefore, a practical solution to be used among business parties must give up some of these ideal properties. In the business arena, although the ultimate goal of businesses is to increase their gains, optimization of these gains is usually compromised, and stability is commonly a goal of equal importance. This may result from stability being an enabler of gains. Thus, an important property of a mechanism of the type we study is its stability.

One way to arrive at a stable coalition configuration among self-interested agents is to compute a solution such as, e.g., the Kernel [1]. Although stable, the Kernel is computationally hyper-exponential, thus inappropriate for practical use. Relaxation of this complexity was performed in [10], however the relaxed solution (which is polynomially complex) still requires complete knowledge of the values of coalitions, and assumes that each coalition has a unique value. Another way to provide stability is

* This work was supported in part by NSF grant # IIS0222914.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AMMAS '03, July 14-18, 2003, Melbourne, Australia.

Copyright 2003 ACM 1-58113-683-8/03/0007...\$5.00.

to design a computationally feasible stable coalition formation mechanism. This will inevitably require that the utility be compromised, however stability and fairness can be planned for. The goal of this study is to develop such a mechanism.

We design, implement and evaluate a mechanism that allows business parties to form coalitions for performing tasks jointly. An example business domain in which such a mechanism is applicable is the business-to-business (B2B) Request For Proposal (RFP) domain. In the B2B RFP domain, contractors attempt to join together to address complex requests, composed of several subtasks, for products or services. To beneficially address these requests it may be essential to form coalitions. Time is however very limited, and although each request may have a price tag attached to it, its value to each business party is typically different and unknown to the other parties. The mechanism we developed addresses exactly such conditions. We provide a protocol for agents to negotiate the formation of coalitions based on their *estimated* values. We further provide heuristics for proposal preparation and selection as well as adaptation methods, to be used by the agents in conjunction with the protocol. Our experiments demonstrate stability, as deviation from pure strategy profiles proves non-beneficial. They further show satisfying gains, proved to be close to the (experimentally computed) optimum.

2. PROBLEM DESCRIPTION

We consider situations where a set of RFP tasks, $\mathfrak{T} = \{T_1, \dots, T_n\}$ should be satisfied as soon as possible. Each task $T_i \in \mathfrak{T}$ consists of sub-tasks T_{i1}, \dots, T_{in} . There is a set of self-interested agents, $\mathfrak{A} = \{A_1, \dots, A_m\}$ that are able to perform some of these sub-tasks and each tries to maximize its benefits. An agent is capable of performing only a subset of the subtasks of a given task. We assume that there is a boolean function, ϕ from $\mathfrak{T} \times \mathfrak{A}$ to $\{true, false\}$, that associates with an agent A_j and a sub-task T_{ik} the value *true* if the agent is capable of performing T_{ik} and the value *false* otherwise. We further assume that all the agents know this function. That is, each agent knows its own capabilities and the capabilities of all the other agents. An agent A_j has a cost, b'_{ik} , for each subtask T_{ik} it can perform. These costs are private knowledge and an agent, while knowing its own costs, does not know the costs of other agents and may be able only to estimate these costs.

Since an agent cannot perform all of the subtasks of a given task, it must cooperate with other agents in order to satisfy that task. Thus, we assume that agents form coalitions in order to satisfy tasks. A coalition is paid for satisfying a task. A coalition \mathfrak{C}_T for a task T is a tuple $\langle C, alloc, u \rangle$ where C is a set of member agents, $alloc$ is an allocation function that associates with each subtask of T a member of C such that $alloc(T_{ik}) = A_j$ only if $\phi(A_j, T_{ik}) = true$. u is a payoff distribution vector – its elements are the benefits of \mathfrak{C}_T members. As the payment made to a coalition to perform its task increases and the sum of the costs associated with the performance of the subtasks by the coalition members decreases, the efficiency of the coalition increases and its members may be able to increase their benefits. We assume that an agent is rational and will join a coalition only if it believes that joining will increase its benefits.

For simplicity, we assume that an agent cannot participate in more than one coalition at a time and also, as a member of a coalition, it can perform only one sub-task at a time.

3. SOLUTION APPROACH

To allow agents to form coalitions given the special settings presented above, we devise a mechanism that consists of a protocol (Section 3.1) and a set of strategies (Section 3.2). To avoid the unrealistic exponential search complexity for optimal strategies, we suggest strategies, which are based on a set of heuristics. These strategies are later evaluated experimentally (Section 4). Participating agents must adhere to the protocol, and this adherence is enforceable. Of course, agents may choose not to participate in the protocol, yet by that they choose to avoid the potential gains from participation. When participating in the coalition formation and following the protocol, agents are free to select strategies other than those we propose. We nevertheless believe, based on our experimental evaluation, that the proposed strategies are highly beneficial, and given their low computational complexity it is reasonable to assume that agents will use them instead of searching for others. The details of the protocol and the strategies follow.

3.1 Protocol

The coalition formation protocol we propose is a special type of an auction with an extension for coalition formation. The protocol consists of a central manager and multiple agents that can join in. The manager supports two roles – an auctioneer role and a coalition negotiation manager role. Both of these roles are assumed to be neutral trusted third parties that do not discriminate among the participating agents and do not disclose their private information to others.

The auctioneer represents businesses that are in a pressing need for complex products or services and they express their needs by publishing RFPs. At the time of publishing an RFP, its issuer sets its price, however as a result of the urgency, the value of the RFP to its issuer decreases as time elapses. The role of the auctioneer is thus to publish the available RFPs, to collect proposals of coalitions addressing these RFPs, to determine the winning coalition for each RFP, and to discount the price of the RFP over time. The auctioneer is also responsible for distributing the payments to coalition members after they complete the execution of an RFP.

In our model each RFP is a task and its complexity is manifested by its partition into sub-tasks. The auction protocol is used to allocate the tasks in \mathfrak{T} to coalitions. The auction is divided into rounds, r_1, r_2, \dots and it ends when there are no more tasks to auction or no more agents that participate in the auction, or the values of the remaining tasks are all nullified. At the beginning of the auction, the auctioneer announces, for each task $T_i \in \mathfrak{T}$, the price, $P(T_i)$, that will be paid to a coalition that will perform T_i at the first round of the auction. In each round, the prices of unallocated tasks are reduced by a factor δ . This δ is announced by the auctioneer at the beginning of the auction as well. At each round r , for each unallocated task $T_i \in \mathfrak{T}$, there may be zero or more coalitions that propose to perform the task. The auctioneer verifies that the members of the proposing coalitions can jointly perform the task (i.e., they hold all of the required capabilities) and awards each task to the first qualified coalition that proposed to perform that task. When two (or more) qualified coalitions submitted their proposals simultaneously, the auctioneer selects one of them randomly. The winning coalition is paid $P(T_i) \delta^r$ upon performing the entire task. Partial fulfillment of a task yields no payment. Additionally, the submission of a proposal to the

auctioneer is binding, i.e., after a coalition is awarded a given task its members cannot break their contract (this can be enforced by imposing very high fines).

The cost B_C of a coalition $\mathcal{C}_T = \langle C, alloc, u \rangle$ is the sum of the costs of the agents for performing their allocated subtasks within \mathcal{C}_T , i.e., $B_C = \sum_j b_j$, where b_j is the cost of A_j and the sum is over all A_j members of C . Thus, the net benefits of a coalition \mathcal{C}_T that was awarded task T_i at round r is $N = P(T_i) \delta^r - B_C$. The protocol manager is in charge of distributing B_C among the members of C . Each agent A_j is paid its cost b_j plus an equal share of N . That is, A_j will be paid $b_j + N/|C|$. Here, a member of C is uncertain regarding B_C since it is uncertain about the costs of the other members. Yet, the protocol manager must know the exact b_j values for payment calculation, and the protocol dictates that the agents disclose this information to the manager. We are aware that self-interested agents may be motivated to deceitfully inflate their cost to increase the payment allocated to them, however we assume that they do not do that.* As stated in the introduction, our major goal is to provide the coalition formation and task allocation protocols. Payment distribution schemes other than the one we use may be devised however are not the focus of this paper. Prior to submitting proposals to the auctioneer, the agents need to form coalitions. For this, we suggest that the agents negotiate in order to form coalitions. This negotiation is facilitated by our mechanism via its negotiation manager role. During the coalition formation negotiation, agents send and receive proposals for coalitions to be formed. A proposal by an agent A_j specifies a coalition $\mathcal{C}_T = \langle C, alloc, u \rangle$ where $A_j \in C$. The coalition formation negotiation is performed via the negotiation manager. At each auction round the protocol allows only one negotiation round. This is enforced by the manager. At each negotiation round, the manager orders the agents randomly, and the agents perform negotiation actions in that order. Each agent, in its turn, can either send a proposal for forming a coalition C to all of its members or accept such a formation proposal made to it by another. An agent has only one turn in each round. All offers are valid for one round and thus an agent making an offer must wait until it hears from all of the agents to which it proposed. It cannot accept any other proposal in this round. Note that proposals are all sent via the auctioneer and recorded there. As a result, an attempt to bypass the protocol - sending proposals externally and agreeing on coalitions to be formed - will be detected by the manager and can be penalized for. If all the members of a proposed coalition accept the proposal, the coalition is proposed to the auctioneer for performing a specific task. If a coalition is awarded the task then the members of the coalition quit the protocol. The agents that have not joined a coalition in a given round continue to negotiate in the next round.

3.2 Heuristics

An agent that participates in the protocol presented above needs means to decide which coalitions to propose to which other agents. It also needs to be able to decide upon acceptance and

* Our payment distribution scheme was chosen for its simplicity. However businesses, and in particular publicly held ones, are audited by their accountants and later need to reveal this information to shareholders. Deceitful expenses can be exposed in this process and will be penalized for. Therefore, our payment scheme is reasonable for a B2B environment.

rejection of proposals it receives from others. As stated earlier, computing the best strategy to handle such proposals is exponentially complex. Hence, we present a set of heuristics on which these strategies rely. These heuristics will provide the agent with a method for ranking the coalitions it can be part of according to their desirability to it. For this ranking, several criteria may be used.

To rank candidate coalitions, an agent should first inspect the set of RFP tasks and their partition into sub-tasks. It should then examine the other agents and their capability to perform these sub-tasks. (As stated above, these capabilities are assumed to be common knowledge). Following, the agent can compute the coalitions that can jointly address the RFP. Only then can the agent rank these coalitions. Suppose that there are n agents, m tasks and a maximal number of k subtasks per task. For each sub-task s of a task T that the agent can perform, there are $O(n^k)$ possible sets of agents that can perform the other subtasks of T . The number of subtasks that the agent can perform is of order $O(mk)$. Thus, overall the agent faces an exponential complexity. In our solution we assume that k is small. Fortunately, in many real-world RFP environments this assumption holds. Consequently, inspection of all possible coalitions is feasible in spite of the exponential complexity. Otherwise, the search itself would require simplifying heuristics to provide feasibility.

Ranking is affected by the knowledge that agents have. Some heuristics are applicable only in case that an agent has enough knowledge about the costs of other agents. The agents may try to build estimations about the costs of the other agents in the negotiation process. With no prior information, an agent can evaluate the cost of a sub-task to other agents as equal to its own cost for performing that sub-task. The agent may however have a rough estimation, based on “common knowledge” of the cost of other sub-tasks. During the negotiation, an agent may change its estimations regarding the costs of different agents, according to their behavior in the negotiation. A payoff demanded by an agent as part of a proposal is not a good estimator of its cost, because it includes the desired profit as well, and thus it is an over-estimator of the cost.† Nevertheless, combined with other estimators, payoff requests can be useful for cost estimation.

Once an agent has computed the list of candidate coalitions and has acquired some knowledge of the costs the other members of these coalitions should incur when performing the sub-tasks allocated to them within these coalitions, it can rank the coalitions. Following, we provide heuristics to be used by the agent for this ranking. We consider two basic heuristics for ranking coalitions as well as two adaptation methods that can be used in conjunction with the basic heuristics. Our experimental results show (see section 4) that, in spite of the simplicity of the proposed heuristics, they provide good results when compared to a centralized experimental optimum (computed via iterative hill climbing). We also developed, implemented and tested other heuristics. However, the ones presented in the paper were lead to the best results.

† Some agents may try to inflate their demands, either in an attempt to mislead regarding their costs, or out of mere greed. However, an agent that demands too much may hurt itself, since it incurs the risk that proposals the exclude it will be favored.

Marginal heuristic: The first heuristic is based on the marginal profit of the whole coalition (thus called henceforth *marginal*). This heuristic suggests that, in order to rank a coalition, an agent should sum the estimated costs of all the agents participating in the coalition for performing their allocated subtasks, and subtracts the sum from the current value of the coalition. This difference is the agent’s estimation of the net value of the task. To enable comparison of the heuristic with other heuristics, we will normalize the net value by dividing it by the maximal possible value of a coalition. The marginal heuristic guides the agent to choose a coalition with the highest normalized value.

This heuristic represents a simple, common sense, behavioral pattern. Preferring coalitions that have a greater net value is a reasonable decision, as these may provide a higher utility to the member agents as well. In fact, given our protocol, a greater net value of the coalition guarantees (in the case of coalitions of an equal size), a higher individual utility. The only issue here is that, in practice, the agent only estimates the net value of the coalition, this estimation may be far from the actual net value, and therefore utility maximization is not guaranteed.

Expert heuristic: The second heuristic is based on reduction of competition, allowing an agent to capitalize on its expertise (thus called henceforth *expert*). An agent may be considered an expert (with respect to a subtask required for some RFP) when only a few other agents can perform it as well. The expert heuristic directs an agent to seek tasks with a low number of competitors.

The marginal heuristic may be problematic when several agents compete over the same subtask. This might result in many rejected offers, lowering overall utility. The expert heuristic tries to resolve this problem. It suggests that an agent try to form coalitions in which it performs subtasks of which it is an expert. The normalized value of the heuristic for a coalition is 1 minus the number of agents capable of performing the task that the agents perform in the coalition, divided by the total number of agents. The expert heuristic attempts to minimize collisions of the agents, by driving them to a subtask allocation based on their capabilities. It also aims at better matching agents to subtasks, resulting in a wider range of tasks being performed.

Both of the heuristics that we provide yield similar ranking results on successive negotiation rounds. As a result, an offer, which was rejected on one round, may be regenerated, offered again, and rejected again on the proceeding round. Thus, recurring rejected offers may be a major problem of the suggested heuristics. In fact, in our control experiments, we were able to generate this problem. To solve this problem, we applied two methods that allow agents to adapt their proposals and responses based on results of preceding negotiation rounds.

The first adaptation method suggests that an agent will exclude agents that have rejected its current proposals from any future coalition proposal it makes. This inevitably leads to choosing other agents, different from the rejecting ones, thus increasing the chances of successful negotiation. Our experiments show that this method indeed improves the results of the coalition formation process. The first adaptation method, however, may be reacting too severely. Excluding other agents from *any* future proposal may prevent the formation of beneficial coalitions. Hence, we devised another adaptation method, according to which an agent that offered a coalition that was rejected will not offer the same allocation of subtasks to the rejecting agents in the following

rounds. This method increases the variety of proposed coalitions, without being too preventive.

Using the proposed heuristics and the adaptation methods, each agent at its turn (set by the protocol) inspects and ranks possible coalitions. In case it has not received any proposal at the current negotiation round, it sends the coalition with the highest rank as a proposal to the members of the candidate coalition. An agent that received proposals compares its share from the best proposal received to its estimated share from the coalition it plans to form. It takes into consideration the discount rate, by willing to accept an offer even if its proposed share is less than its share in its preferable coalition, but more than the discounted desirable share. If the best proposal is acceptable, it sends an acceptance message to the sender of the proposal. Otherwise, it sends the coalition with the highest rank as a proposal to the members of the candidate coalition.

The heuristics we suggest as well as the adaptation method were evaluated experimentally. The details of this evaluation follow.

4. EXPERIMENTAL EVALUATION

To examine our protocol and to compare between the proposed heuristics and to study how parameters of the environment influence the agents’ behavior we have developed a simulation system and performed a series of experiments. We present below the settings of these experiments, followed by the results.

4.1 Settings

A setting in our experiment consists of the following: (1) the number of agents; (2) the number of tasks and (3) the number of subtasks per task. To create a specific configuration, the following parameters were determined: the value of each task, the capabilities of each agent (i.e., which sub-tasks it can perform) and the cost of a given agent to perform the tasks it is able to do.

In most of the experiments we considered six basic settings: (i) 6 agents, 2 tasks and 3 subtasks/task; (ii) 10 agents, 2 tasks and 5 subtasks/task; (iii) 10 agents, 5 tasks and 5 subtasks/task; (iv) 12 agents, 3 tasks and 4 subtasks/task; (v) 16 agents, 4 tasks and 4 subtasks/task; (vi) 16 agents, 5 tasks and 4 subtasks/task. We chose these settings since they provide a variety of agents/tasks combinations varying the number of agents between 6 and 16. In some environments the number of subtasks to be performed is larger than the number of agents (settings iii and vi) thus, given our assumptions that each agent could perform only one subtask at a time, not all tasks can be performed in these settings. In other settings the number of agents is equal to the number of possible subtasks to be performed (settings i,ii,iv,v), thus the ability to perform all tasks depends on the agents’ capabilities. We also experimented with larger settings: (1) 80 agents, 20 tasks and 3 subtasks/task; and (2) 60 agents, 20 tasks and 3 subtasks/task.

In each experiment, for each of the basic settings, we randomly generated between 100 and 1,500 configurations (depending on the experiment). For each such configuration, we randomly chose 50% of the subtasks to be “specialized tasks” while the rest were “regular tasks”. Each agent had a probability of 0.4 to be able to perform a regular subtask, but only probability of 0.15 to perform a specialized task. For each configuration, the value of the task and the costs for performing the subtasks were determined as following. First, for each subtask, we have randomly selected a mean cost with a uniform distribution between 1 and 99. The

value of a task was determined to be the sum of the means of its subtasks times 1.5, giving an average profit of 50%. The actual cost for a given subtask to be performed by an agent that is capable of doing it was determined using a normal distribution with the mean cost of that we have randomly associated with the given subtask before, and a certain deviation (2 in the basic settings). In most of the experiments, the discount factor δ was set to 0.01. Finally, we consider two settings with respect to the information available to the agents. In the first setting, which we refer to as “complete information” case, each agent knows the costs of the other agents, while in the “incomplete information” case, they know only the mean values of the costs. In both cases each agent knows the capabilities of all the agents.

4.2 Results

When evaluating our proposed protocol and strategies, we are interested in the ratio between the total overall outcome of agent systems acting according to our proposed heuristics, and the optimal centralized outcome of such systems. Because of the high complexity of computing the optimal task allocation, we have computed a near-optimal value by using a hill-climbing algorithm.

4.2.1 Basic experiment

The goal of our first set of experiments was to compare the marginal and expert heuristics described above, allowing agents to rank coalitions by assigning weights to each heuristic. Since the heuristics are destined to influence in different ways, our hypothesis was that some combination of the two heuristics would give better results than each of the pure heuristics alone would. We also assumed that having complete information would assist the marginal heuristic much more than it would assist the expert heuristic, since the former relies directly on the missing information (of the actual subtask costs, and hence, the actual net task value) in order to evaluate the coalitions.

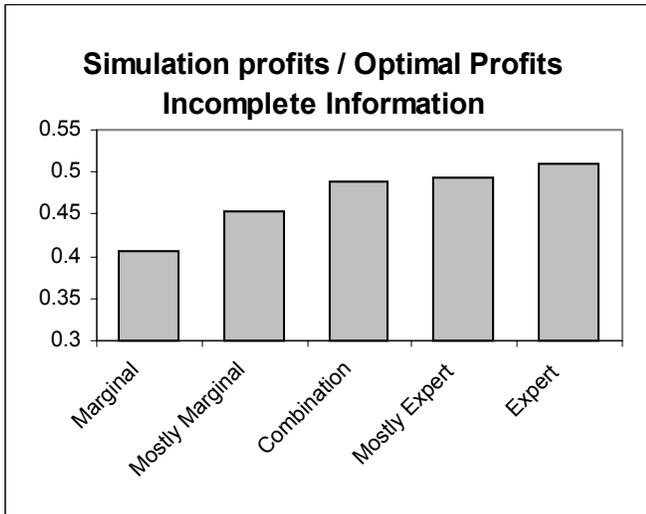
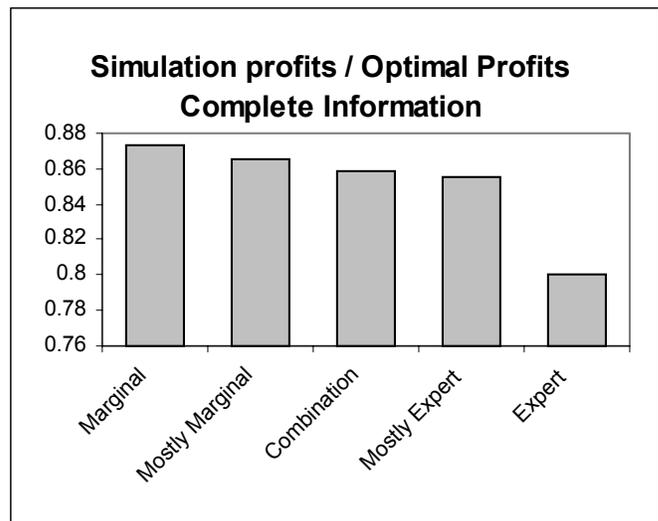


Figure 1. Incomplete information: the expert heuristic is better

The experiments examined coalition formation with no adaptation during the negotiation. We tested several combinations of the two heuristics, from purely marginal heuristic to purely expert heuristic. When ranking possible coalitions, the combined heuristics associates with each coalition a weighted sum of the rankings of the pure heuristics. In particular, the pure marginal

strategy assigns a weight of 1 to the marginal heuristic and 0 to the expert heuristic; The mostly marginal strategy assigns a weight of 3 to the marginal heuristic and 1 to the expert heuristic; The combination strategy allocates an equal weight to each of the two heuristics; The mostly expert strategy assigns a 3:1 ratio in favor of the expert heuristic; and the pure expert strategy assigns a 0 weight to the marginal heuristic and 1 to the expert heuristic.

The results of this set of experiments appear in figures 1 and 2. We measured the ratio between the utilities of the various heuristics, and the utility of the centralized near optimal utility. We were surprised to find out that the pure heuristics provided higher average gains than any combination of strategies did. In the case of incomplete information, the expert method was better than the marginal one; in the case of complete information, the marginal method was better. Moreover, in each case we can see a monotonic increase in the overall utility as a function of the ratio



between marginal and expert weights.

Figure 2. Complete information: the marginal heuristic is better

Both heuristics performed significantly better when having complete information. As we have assumed, the marginal method was assisted by the information much more than the expert one was. To gain a better understanding of this phenomenon, we have inspected the number of coalition contracts signed. From this we learned that, whereas with incomplete information the expert heuristic draws its strength by being able to reach much more contracts than the marginal heuristic, in complete information case the number of signed contracts for the expert and the marginal heuristics is almost the same, though the marginal method performs much better by selecting more beneficial tasks. We also noted that when having complete information, the marginal heuristic arrived at contracts signing faster than the expert heuristic did. The reason for this pattern of contract signing is the reliance on information exhibited by the heuristics. When there is incomplete information, the marginal heuristic uses inaccurate information for decisions. It thus runs into conflicts, eventually closing fewer contracts. The expert heuristic does not use the inaccurate information used by the marginal heuristic and avoids these conflicts. When complete information is available, both heuristics can sign more contracts, thanks to a more accurate payoff estimation and division.

4.2.2 Negotiation with adaptation

The second set of experiments examined the effect of adaptation on negotiation results. We have used the two adaptation methods discussed in Section 3.2. The first and second adaptation methods are referred to here as A1 and A2, respectively. Recall that the aim of both adaptation methods is to reduce collisions and prevent recurring rejected offers. A1 should prevent many more collisions than A2 would, but may cause the missing of potentially beneficial coalitions. We hypothesized that adding either of the adaptation methods to the marginal and expert heuristics will yield higher gains than the pure marginal and expert without adaptation do. We tested the methods in an incomplete information environment similar to that of experiment 1.

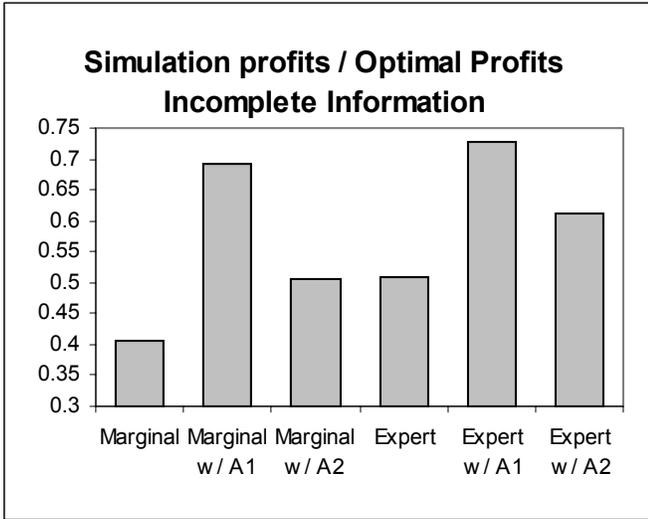


Figure 3. Adaptation increases profits

The results of this experiment set appear in figure 3. As the figure shows, applying both adaptation methods has indeed increased the overall utility, for both the marginal and the expert heuristics. It appears that A1 was substantially better than A2, though. The reason is that the limited reaction of A2 to proposal rejection results in it not handling many collisions, in particular in cases where a new proposal is very similar to a previously rejected one.

We have also conducted large-scale experiments to verify the results of experiments 1 and 2, with 60-80 agents and 20 subtasks. The results were consistent with smaller scale experiments.

4.2.3 Discount rate

In the basic setting, we set the discount rate to 0.01. In order to verify that our results are consistent when the discount rate changes, we performed a third set of experiments in which we varied the discount rate, keeping the other parameters fixed, similar to the parameters of the previous experiments. We measured the ratio between the average utilities of the marginal and expert heuristics, and the centralized near-optimal utility. This was performed in both the incomplete information case (Figure 4) and the complete information cases (Figure 5). We assumed that an increase of the discount rate would cause a decrease in the number of contracts signed, and of the total utility. Although the agents consider the discount rate when answering proposals, if the discount rate is high and the contracts are not signed fast, they may never be signed. We thus expected that most of the contracts would be signed in the first negotiation rounds;

this may reduce or even nullify the expert heuristic's advantage of solving conflicts.

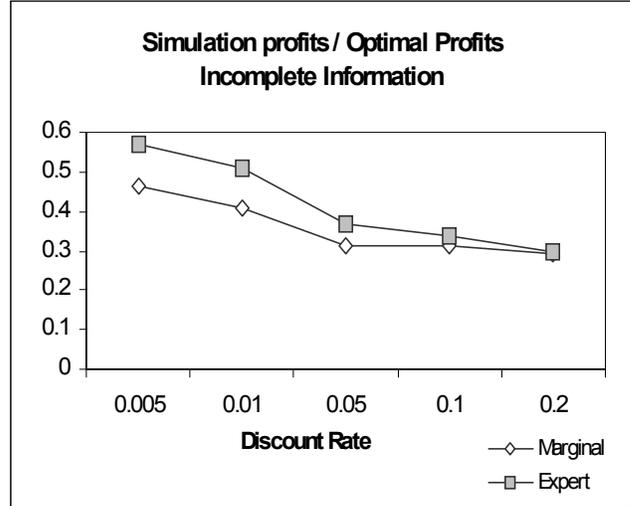


Figure 4. Incomplete Information: utility decreases with increase in discount rate

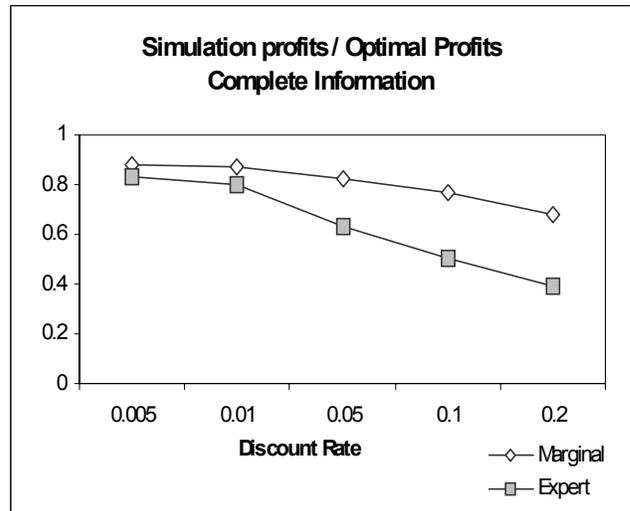


Figure 5. Complete Information: utility decreases with increase in discount rate

As can be seen in figures 4 and 5, in all cases, the overall utility decreased when the discount rate increased. The expert heuristic utility was much more susceptible to the discount rate than the marginal heuristic utility was. When having incomplete information, a high discount rate led both heuristics to yield a similar number of signed contracts in similar time. In the case of complete information, not only had the marginal heuristic yielded many more signed contracts than the expert heuristic did, it also resulted in contracts being signed much earlier.

4.2.4 Cost distribution

As discussed above, the cost of an agent for performing each sub-task is computed based on a distribution function of a "known average" with a specific standard deviation. The average reflects the common knowledge about the usual cost of performing a subtask; the higher the deviation is, the greater will be the dispersion of agents' cost of performing a sub-task. In a sense, in

the case of incomplete information, since the agents know only the mean value of the subtasks, as the dispersion increases the agents' uncertainty also increases. We wanted to examine the effect of different deviations on the performance of the heuristics. For this, we performed an experiment similar to the previous ones, but this time we changed the standard deviation, keeping the other parameters unchanged. We assumed that a higher dispersion will lead to a decrease in the overall utility, compared to its optimal value and a smaller dispersion will be similar to the case of complete information. We expected the optimal value itself to increase when the deviation increases, since an optimal task allocation will allocate an agent with the smallest cost to each subtask; higher deviation from the average will result in higher net value. That is, if the deviation is small, the costs of some agents are slightly above the average cost of certain subtask, and some are slightly below it. In this case, the optimal algorithm will choose an agent that its cost is slightly below the average resulting in a relatively small profit. However, when the deviation is large, the cost of some agents is much lower than the average cost. The optimal allocation will select an agent with a cost that is much lower than the average, and gains a higher profit than in the small deviation case. Similarly, in the case of complete information (Fig. 7), the marginal heuristic results have not suffered much from increasing the deviation. Since the marginal heuristic chooses the most valuable tasks, increasing the deviation, and thus the dispersion, increases the possibility of gaining more from choosing the right tasks. However, the expert heuristic that does not take the task value into account suffers even in the case of the complete information when compared to the optimal solution. As appears in Fig. 6, when having incomplete information, a very small deviation of 0.01 had given results close to the results with complete information (Fig. 7), as we have anticipated. In addition, the marginal heuristics suffered more from increased uncertainty than the expert heuristics that relies less on the other agents' costs.

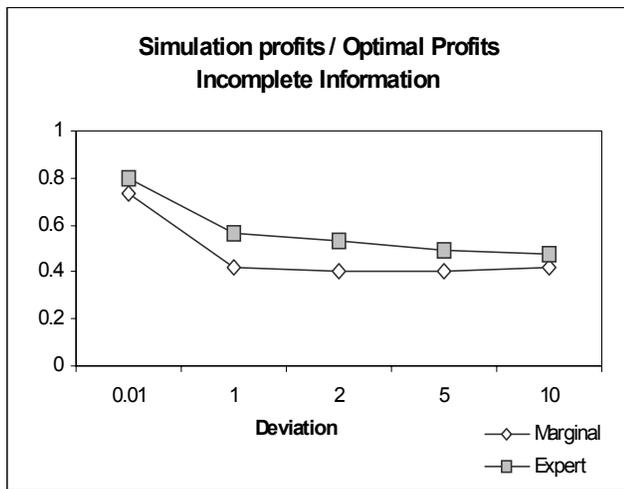


Figure 6. Incomplete information: the utility decreases with an increase in cost deviation

4.2.5 Heterogeneous environments

All the experiments so far have studied the agents' performance in homogenous environments, where all the agents follow the same strategies. In our last experiment, we examined the effect that a deviation from the common strategy will have on the performance

of one agent. It should be more interesting to examine the effect of a deviation from a superior strategy, since it is likely that the majority of agents will implement the superior strategy. Thus, in the case of incomplete information we examined the effect of one agent deviating from the (superior) expert heuristic practiced by the majority and implementing the marginal heuristic. In the case of complete information, we let one agent to implement the expert heuristic and the others to use the marginal heuristic.

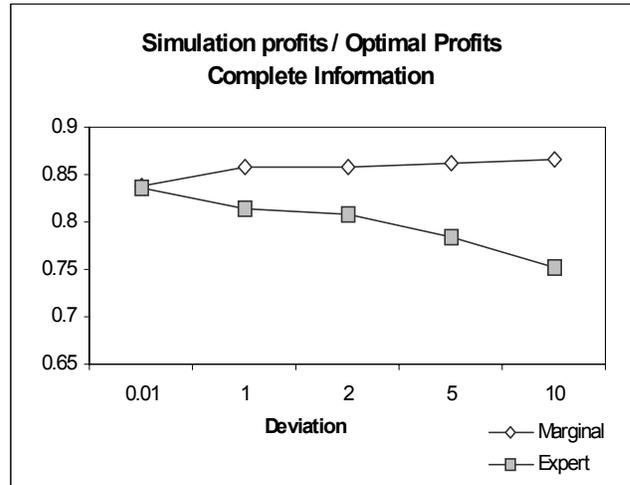


Figure 7. Complete information: the expert utility decreases with an increase in cost deviation, while the marginal utility increases

The basic settings were set as in the previous experiments. Since measuring the performance of one agent can lead to large fluctuations, we have performed 1500 runs, many more than in previous experiments. We assumed that deviating from the common method might help an agent that implements the marginal heuristic, others using the expert heuristic. Such an agent may benefit from its ability to select more profitable tasks, relying on the other agents' behavior for avoiding collisions.

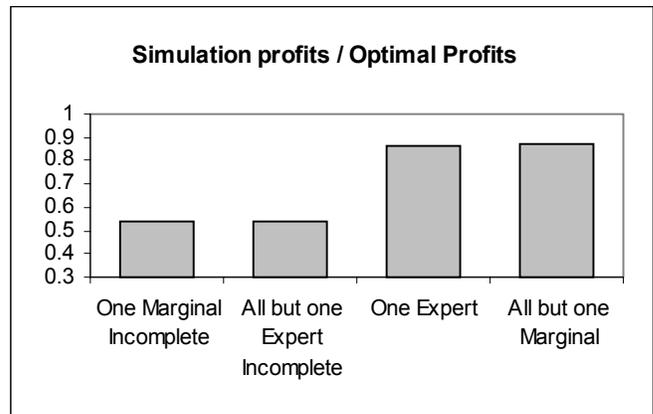


Figure 8. Deviation from the majority strategy is not beneficial.

The results, as appear in Figure 8, clearly show that deviating from the dominant heuristic is not beneficial, nor does it harm the deviating agent. In the incomplete information case, the average utility of the deviated agent was exactly equal to the average of the others. In the complete information case, there was a slight decrease when deviating, but it was so small that it does not stand statistical analysis, in spite of the high number of experiments.

5. RELATED WORK

Game theory provides an analysis of the possible coalitions that shall form as a result of a coalition formation process, and the resulting disbursements to the agents (see, e.g., [3]). However, game theory does not provide algorithms that agents can use to form coalitions. Given a formed coalition configuration (i.e., a partition of the agents to subsets) game theory usually concentrates on checking its stability or its fairness using concepts such as the Kernel [1] and on the calculation of the corresponding payments. Game theory is also not concerned with computational complexity, and the solutions are usually exponentially complex. We provide a coalition formation protocol focusing on feasibility and fairness, and suggest heuristics that provide benefits and maintain stability. There have been several attempts to generalize the stability concepts of coalition formation, such as the core, for situations of asymmetric information [12],[6]. As in the case of complete information, these studies focus on the stability concept, but algorithms for agents' activity are not provided.

Many algorithms for answering the question of group formation in cooperative environments were suggested (e.g., [11],[2]). In [1] the problem is addressed for self-interested agents, but in superadditive environments. In [10], solutions were proposed for non-superadditive environments, but the value of each coalition is known. Sandholm and Lesser [9] present a coalition formation model for bounded-rational agents and a general classification of coalition games. As in [9] we also allow for varying coalitional values, but provide the agents with heuristics that could be computed in polynomial time. However, we assume that time is costly, and that agents take the coalition formation time into consideration when deciding on whether to join a coalition. In [9], the value of a coalition depends on the computation time. However, we consider cases in which the time for computing the coalition values is polynomial. Sandholm et al. [8] offer an algorithm that gives a tight bound of an optimal coalition structure, but in their work they take into account only one value for each coalition; in our case, a coalition may have different values for each task it may perform.

All the works we mentioned assume complete information: each of the agents knows the exact value of each possible coalition. This is not our case, however. In real world situations, rarely do other agents know each agent's exact value and costs of fulfilling each task [4]. Therefore, it is not possible to utilize the techniques presented in the above papers. In particular, the methods used to check the stability of a given state require that all agents will hold the same beliefs about the state. More related to our work is research on fuzzy and stochastic co-operative games [7]. In such games agents faces situations of uncertainty, including, for example, vagueness of expected coalition values and corresponding payoffs. This preliminary research attempts to find formal models to address these problems, while we provide experimental results that present the advantages of using our proposed protocol and heuristics.

6. CONCLUSION

In this paper we consider the problem of coalition formation for cases where groups can only perform tasks cooperatively. In particular, we consider situations where a complex task comprised of sub-tasks; each sub-task should be performed by a different agent; the costs that an agent incurs for performing a specific sub-task may be unknown to other agents; and, time for addressing a

task is limited. To address this problem, we have developed a protocol that enables agents to negotiate and form coalitions, and provided them with two simple heuristics for choosing coalition partners: "marginal" heuristic and "expert" heuristic. Via experiments, we found out that the marginal heuristic is the best when there is complete information, while the expert heuristic is better when the information is not complete. In both cases, the experiments showed that these heuristics are stable and beneficial. Deviation from following the recommended heuristic does not increase the deviated agent benefits.

There are many open issues to address. The most important one is whether it is possible to relax the assumption of equal distribution of a coalition benefits, but to maintain the stability and the efficiency of the agents system. For practical use in B2B RFP situations, the predictability, the stability and the fairness of the mechanism are most important, and these properties were experimentally proven for our solution.

7. REFERENCES

- [1] Davis, M., and Maschler, M. The Kernel of a cooperative game. *Naval Research Logistics Quart.*, 12 (1965), 223-259.
- [2] Jennings, N. R. Controlling cooperative problem solving in industrial multi-agent systems using joint intentions. *Artificial Intelligence Journal*, 75(2) (1995), 1-46.
- [3] Kahan, J. P., and Rapoport A. *Theories of coalition formation*. Lawrence Erlbaum Associates, Hillsdale NJ, 1984.
- [4] Klusch, M. and Gerber, A. Issues of dynamic coalition formation among agents. *Proc. KSCO-2002, France, 2002*.
- [5] Klusch, M., and Shehory, O. A polynomial kernel-oriented coalition formation algorithm for rational information agents, in *Proceedings of ICMAS '96 (Kyoto Japan, 1996)*, 157-164.
- [6] Kwasnica, A. M. *Bayesian Implementable of Efficient and Core Allocations*. Penn. State U. Working Paper, 2002.
- [7] Mares, M. Fuzzy coalition structures, *Fuzzy Sets and Systems*, 114 (2000), 23-33.
- [8] Sandholm, T., Larson, K., Andersson, M., Shehory, O. and Tohme F. Coalition structure generation with worst case guarantees, *Artificial Intelligence Journal*, 111(1-2) (1999), 209-238.
- [9] Sandholm, T., and Lesser, V. Coalition formation among bounded rational agents, *Artificial Intelligence Journal* 94(1-2, (1997), 99-137,
- [10] Shehory, O., and Kraus, S. Feasible formation of coalitions among autonomous agents, *Computational Intelligence*, 15(3) (1999), 218-251.
- [11] Shehory, O., and Kraus, S. Methods for task allocation via agent coalition formation, *Artificial Intelligence Journal*, 101 (1-2) (1998), 165-200.
- [12] Vohra, R. Incomplete Information, Incentive Compatibility and the Core, *J. of Economic Theory*, 86 (1999), 123-147.
- [13] Zlotkin, G., and Rosenschein, J. Coalition, cryptography, and stability: mechanisms for coalition formation in task oriented domains, *Proc. AAAI '94 (Seattle WA, 1994)*, 432-437.