Branch and Bound for Negotiations in Large Agreement Spaces

(Extended Abstract)

Dave de Jonge
Artificial Intelligence Research Institute,
IIIA-CSIC
Campus de la UAB
08193, Bellaterra, Catalonia, Spain
davedejonge@iiia.csic.es

Carles Sierra
Artificial Intelligence Research Institute,
IIIA-CSIC
Campus de la UAB
08193, Bellaterra, Catalonia, Spain
sierra@iiia.csic.es

ABSTRACT

We introduce a new multiagent negotiation algorithm for large and complex domains, called NB^3. It applies Branch & Bound to search for good offers to propose. To analyze its performance we present a new problem called the Negotiating Salesmen Problem. We have conducted some experiments with NB^3 from which we conclude that it manages to decrease the traveling cost of the agents significantly, that it outperforms random search and that it scales well with the complexity of the problem.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems

General Terms

Algorithms, Experimentation

Keywords

Negotiation, Search, Negotiating Salesmen Problem

1. INTRODUCTION

Previously proposed negotiation algorithms have mostly focused on the utility space. They assume that given a utility aspiration level it is always possible to find a proposal that would fit that level. In this paper we focus on complex problems for which these classical continuity assumptions do not apply and thus solutions have to be found directly at domain level. Also, we address a number of realistic assumptions that make the application of current negotiation algorithms unfeasible: the space of solutions is huge, utility is non-linear and therefore difficult to calculate, the environment is only partially observable, decisions have to be made within a limited time frame and solutions may involve many agents, possibly human.

We introduce a new family of Branch and Bound algorithms, namely NB^3 (Negotiation Based Branch & Bound), that use negotiation as the key element in the exploration of the joint space of solutions for multiple agents.

As far as we know there are no algorithms implemented that are comparable with the work we present here. Most work on negotiations involves very simple scenarios with only two agents, a small space of solutions and linearly additive utility functions. Although non-linear utility has been studied in for example [1] they still assume that the utility is simple to calculate. The combination of search and negotiation has been studied in [2], but they assume a mediator that must be trusted and severely limits the freedom of the agents, making it more useful for cooperative scenarios rather than competitive ones. In [3] there is also search, but once again with simple utility functions. In [4] it was suggested to use genetic algorithms to explore the agreement space, but no implementation or results were given.

2. NB^3 BASIC CONCEPT

We assume a multiagent scenario in which each agent has a personal cost function to minimize. Each agent has a set of possible actions that it can execute in order to change the state of the world into a new state for which its personal cost function has a lower value. The result of an action depends on the actions executed by the other agents, so they have to negotiate over which joint plan of actions to execute. Their interests are however conflicting: a certain world state might yield low costs for one agent, but high costs for another agent. The agents are assumed to be selfish: they are only interested in minimizing their personal cost function. This means that the agents must compromise: each agent should propose plans that lower his own cost as much as possible, but that at the same time lower the costs of the other agents sufficiently to make them accept these plans.

We designed an algorithm, which we call NB^3, to run on such an agent. The other agents present might also run this algorithm, or any other negotiation algorithm, or they might be human, but that is irrelevant to us. NB^3 applies a Branch and Bound tree search to explore the space of all possible plans under the above mentioned assumptions. Each node in the tree represents a partial plan, and maintains a lower- and upper-bound for the cost function of the agent as well as an estimation of the lower- and upper-bounds of the costs of all the other agents. Whenever for a certain node the lower bound of one of the agents is higher than the reservation value of that agent, that node can be pruned, because it

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

means that this partial plan can never yield a lower cost for the agent than its reservation value. In order to determine which nodes to expand, NB^3 uses a heuristic that is based on the offers that the other agents have made previously. In this way, the search can be directed towards a solution that is acceptable to all agents. For more details we refer to [5].

3. NEGOTIATING SALESMEN PROBLEM

To test the algorithm we have defined a new problem, called the Negotiating Salesmen Problem (NSP). It is a variant of the Traveling Salesman Problem, but with multiple salesmen, each only interested in minimizing its own path.

The idea is that there is a set of cities and a set of salesmen and each city needs to be visited by at least one agent. There is one home city where each agent should start and finish its trajectory. Every other city is assigned to one salesman that has to visit it. However, the salesmen are allowed to exchange their cities amongst one another, so that the agents can decrease the distances they have to travel. For example: if a city v is assigned to agent α, but α prefers to visit another city v', which is assigned to agent β, then α will propose β to exchange v for v'. If β however doesn't want v he will not accept this deal. And if no other agent wants to accept v either, then α is obliged to travel along city v. However, we impose the restriction that some cities cannot be exchanged. The cities that can be exchanged are called the *interchangeable cities*, while the cities that cannot be exchanged are the *fixed cities*. Each agent therefore prefers to visit cities that are close to any of his fixed cities.

4. EXPERIMENTAL RESULTS

We have implemented an agent that applies NB^3 to the NSP and conducted a number of experiments with this implementation. Ideally, we should test our algorithm against other negotiation algorithms but, as mentioned, we don't know of any such algorithm that could handle the hard conditions we are considering.

Note that comparing the algorithm against existing search algorithms will not work, since they do not apply negotiation. Pure search algorithms might find the most selfish solution, or the socially optimal solution, but are not able to find the best compromise between these two extremes, given the offers made by other agents. Moreover we do not claim that our search algorithm is better than any existing search algorithm, we only claim that we have made the first algorithm that successfully combines search and negotiation in large and complex agreement spaces.

In order to do useful experiments anyway, we have tested the algorithm against a simplified version of itself that expands the search tree randomly, i.e. without using smart heuristics. Also we have done some tests in which all agents were running NB^3 and repeated these tests with different problem sizes to see how the algorithm scales. Finally we have compared the results with the socially optimal solution.

For any agent we determined a score by comparing its path length after negotiations with its path length before the negotiations. So if an agent scores for example 40% it means that its final path length was 40% shorter than its initial path length. The scores presented here are each averaged over all agents and 25 problem instances.

From the results we can conclude that our algorithm significantly outperforms random search. In the NB^3 vs. random search experiments the NB^3 agents were able to decrease their path lengths by 30%, while the random search agents did not score higher than 10% to 20% depending on the number of NB^3 agents present.

From the experiments with varying complexity of the NSP-instances we conclude that NB^3 scales very well with increasing complexity. When increasing the number of agents from 6 to 16 while holding the number of interchangeable cities per agent fixed at 10, the results stay stable between 25% and 30%. Increasing the number of cities per agent from 6 till 16 while holding the number of agents fixed at 10, the average score decreased from 38% to around 25%.

It is impossible for all the agents to decrease their path lengths with 100% because that would mean they are not traveling at all anymore. Therefore we also compared the results with the length of the paths of the socially optimal solution, and it turned out that the agents are able to decrease their costs by 65% relative to the social optimum (so 0% indicates no decrease of path length, while 100% indicates the agents have reached the social optimum). Note however that NB^3 was designed for selfish agents, and not for agents that want to reach the social optimum. Therefore, even if the social optimum is found by some of the agents, they might not propose it because they might try to reach more selfish solutions.

5. CONCLUSIONS

We have developed a new algorithm that successfully combines search with negotiation under hard, realistic conditions. We applied it to the Negotiating Salesmen Problem and conclude that its results are significantly better than random search and scale well with increasing problem size.

6. ACKNOWLEDGMENTS

Supported by the Agreement Technologies CONSOLIDER project, contract CSD2007-0022 and INGENIO 2010 and CHIST-ERA project ACE and the Spanish Ministry of Education and Science TIN2010-16306 project CBIT.

7. REFERENCES

[1] Ivan Marsa-Maestre, Miguel A. Lopez-Carmona, Juan R. Velasco, and Enrique de la Hoz. Effective bidding and deal identification for negotiations in highly nonlinear scenarios. In *AAMAS '09*, pages 1057–1064, 2009.

[2] M. Klein, P. Faratin, H. Sayama, and Y. Bar-Yam. Protocols for negotiating complex contracts. *Intelligent Systems, IEEE*, 18(6):32 – 38, nov-dec 2003.

[3] Peyman Faratin, Carles Sierra, and Nicholas R. Jennings. Using similarity criteria to make negotiation trade-offs. pages 119–126, 2000.

[4] Angela Fabregues and Carles Sierra. An agent architecture for simultaneous bilateral negotiations. In *Proceedings of the 13è Congrés Internacional de l'Associació Catalana d'Intel·ligència Artificial (CCIA 2010)*, pages 29–38, Espluga de Francolí, Tarragona, October 2010.

[5] Dave de Jonge and Carles Sierra. Branch and bound for negotiations in large agreement spaces. Technical report, IIIA-CSIC, Bellaterra (Barcelona), Spain, 2012. `http://www.iiia.csic.es/files/mon_library/TR-IIIA-2012-01.pdf`.