

Subjective Partial Cooperation in Multi-agent Local Search

(Extended Abstract)

Adi Eisen, Lahan Mor and Roie Zivan
Industrial Engineering and Management department, Ben Gurion University
Beer-Sheva, Israel
{eisenad,lahanm,zivanr}@bgu.ac.il

ABSTRACT

A partial cooperative model that was recently proposed offers a balance between the two extreme scenarios commonly assumed in multi-agent systems, either completely competitive or fully cooperative agents. Partial cooperative agents act cooperatively in a distributed search process, as long as the outcome satisfies some threshold on their personal utility, otherwise, they act selfishly. While personal thresholds formalize intentions for cooperation in the existing model, inter agent relationships are not considered, although in realistic scenarios agents are often willing to act more cooperatively in their interactions with specific agents than with others, i.e., they have subjective cooperation intentions.

We propose an extension to the partial cooperative model by introducing subjective intentions for cooperation. When making decisions on whether to perform cooperative actions agents take into consideration the identity of the agents these actions affect. A preliminary experimental evaluation demonstrates the advantage algorithms that consider the model have over algorithms that do not.

Categories and Subject Descriptors

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

General Terms

Algorithms, Experimentation

Keywords

Distributed Search, Cooperation, Self interest

1. INTRODUCTION

Many realistic multi agent applications include partially cooperative agents, e.g., consider a standard working environment in which employees perform tasks for the benefit of the organization they work for and get a pay check in return. In this most common situation agents are naturally self interested and often have the option to increase their own benefit within the organization, however, they are required to act loyally to increase the organizational profits.

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*
Copyright © 2014, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

A partial cooperation model for representing such common scenarios was recently proposed [2]. This model defines a threshold on the personal gain of an agent from an interaction with other agents and assumes that agents act cooperatively as long as their personal benefit is higher than this threshold. In this model, cooperation decisions agents make do not depend on the identity of the agents they cooperate with. However, in many realistic scenarios this is a major parameter in the decision. Consider the working scenario described above. If one worker asks another to replace shifts with her, obviously, the personal relations between the agents will affect the decision whether to agree. Such subjective relationships were previously observed by social scientists and applied in game theory, e.g., in algorithms for multi player games [3].

In this work we extend the partial cooperative model proposed in [2] by allowing the representation of multi-agent scenarios in which partial cooperative agents attempt to solve asymmetric distributed constraint optimization problems (ADCOPs [1]) similar to the examples described above.

The proposed extended model defines thresholds on reductions of personal utilities that are caused as a result of violations of constraints with specific agents. Thus, an agent can be willing to endure a cost incurred following the violation of a constraint in which some agents are involved, while not willing to endure the same cost incurred due to the violation of a constraint involving other agents.

As in [2], we are able to adjust the anytime mechanism proposed in [4] such that in combination with this mechanism, any local search algorithm will produce solutions that satisfy the thresholds of all agents. However, a preliminary experimental study demonstrates that algorithms that are aware of these thresholds have a clear advantage over algorithms that do not.

2. SUBJECTIVE PARTIAL COOPERATION

¹In the partial cooperation model, a parameter λ_i is used to represent for each agent i its intentions for cooperation. An agent i is willing to cooperate in an interaction with other agents as long as the utility it derives from the outcome $U_i(o)$ of the interaction (or cost it incurs) is within λ_i from the utility it would have derived if it would have performed selfishly μ_i , i.e., $\mu_i - U_i(o) \leq \lambda_i$. Thus, we are seeking the solution to the ADCOP that maximizes the global utility (e.g., sum of all personal utilities) while satisfying all personal thresholds.

¹For a formal description of asymmetric DCOPs see [1].

We extend the model by defining for each agent i a specific threshold for each constrained neighbor j , $\lambda_{i(j)}$ and add the requirement that for each of its neighboring agent j , the reduction in utility for agent i incurred from the constraints involving j is not larger than $\lambda_{i(j)}$. In addition we require the following realistic bounds: $\max_j \lambda_{i(j)} \leq \lambda_i \leq \sum_j \lambda_{i(j)}$, i.e., cooperation is composed from cooperative actions that affect neighboring agents, thus, a personal threshold that is larger than the global threshold is meaningless. Similarly, it is impossible to breach a global threshold that is larger than the sum of all personal thresholds.

Validating that algorithms report solutions that satisfy all personal thresholds can be done in an algorithm independent manner by making a small adjustment to the *Anytime mechanism* for DCOP local search algorithms presented in [4]. This mechanism allows agents to hold in each iteration the best solution explored so far. The adjustment suggested in [1] allows agents to reject solutions in which their threshold was not satisfied. A small additional adjustment is required so agents will also validate that all $\lambda_{i(j)}$ thresholds are satisfied as well.

3. LOCAL SEARCH ALGORITHMS

The adjusted anytime mechanism allows the use of any distributed local search algorithm for solving subjective partial cooperative problems. However, in [1] specific designed algorithms that exploit the properties of the partial cooperative model were found to outperform standard distributed local search algorithms when they are used in combination with the adjusted anytime mechanism. Two such algorithms were designed. In the most successful one, *Goods-MGM*, agents send *nogoods* to neighboring agents that their current assignment prevents the sending agents from satisfying their thresholds. When an agent that sent a *nogood* detects that, following assignment replacements of other neighboring agents, this *nogood* is not valid anymore, it sends the neighbor holding the *nogood* a *good* message indicating that the *nogood* can be discarded (for more details see [2]). This algorithm can be adjusted to the subjective partial cooperative model by having agents send *nogoods* to their neighbors when the threshold on the constraints involving them, $\lambda_{i(j)}$, cannot be satisfied.

In order to validate that this adjustment of the algorithm is indeed necessary, we performed an experiment comparing four types of algorithms, each representing different levels of awareness to the model’s details: a standard DCOP algorithm in which agents assume that constraints are symmetric, an ADCOP algorithm in which agents are aware that the problem includes asymmetric constraints, a partial cooperative algorithm (the Goods-MGM algorithm described above) in which agents consider the λ_i thresholds, and a subjective partial cooperative algorithm in which agents consider all of the above and the $\lambda_{i(j)}$ thresholds as well. These algorithms are denoted in Figure 1 by MGM, MGM-ADCOP, Goods-MGM-PC and Goods-MGM-SPC respectively. All algorithms were combined with the adjusted anytime framework that validates that all solutions satisfy all thresholds.

Figure 1 presents the costs of the solutions found by the four algorithms as a function of the number of iterations the algorithm is performed. The problems are random asymmetric minimization DCOPs with 50 agents, density 0.1, tightness 0.5 and costs for violated constraints selected uniformly in the range 1...99. Thresholds were selected randomly from




Figure 1: Average solution costs of the MGM versions when solving random problems

a normal distribution. Results are averaged over 50 runs of each algorithm solving different random problem instances. It is clear that the only algorithm that does not converge immediately is the version of Goods-MGM that is aware of the details of the model. This algorithm improves monotonically and finds solutions with lower costs than the other algorithms. Obviously the other algorithms do not explore valid solutions after the first few iterations.

4. DISCUSSION

Multi-agent combinatorial problems in which each agent has its own valuation for combinations of assignments can naturally be represented by the Asymmetric Distributed Constraint Optimization (ADCOP) model. Solving an ADCOP is a distributed search process that requires the cooperation of all participating agents.

The partial cooperation model represents the intentions for cooperation of partial cooperative agents by defining a threshold on the loss agents are willing to endure while cooperating. The subjective partial cooperation model we propose in this work extends the model by representing the intentions of agents for cooperation with respect to personal relationships with other agents. This is done by specifying subjective thresholds on the reduction in utility an agent is willing to endure as a result of violations of constraints involving each of the agents it is constrained with.

An adjustment of the anytime mechanism proposed in [4] can validate that any local search algorithm reports solutions that satisfy all thresholds. However, we presented an experiment that encourages the design of specific model aware algorithms for subjective partial cooperative distributed problems. Thus, in future work we intend to design and evaluate such algorithms.

5. REFERENCES

- [1] T. Grinshpoun, A. Grubshtein, R. Zivan, A. Netzer, and A. Meisels. Asymmetric distributed constraint optimization problems. *J. Artif. Intell. Res. (JAIR)*, 47:613–647, 2013.
- [2] A. Grubshtein, R. Zivan, and A. Meisels. Partial cooperation in multi-agent local search. In *Proc. ECAI-2012*, Montpellier, France, August 2012.
- [3] B. Wilson, I. Zuckerman, and D. S. Nau. Modeling social preferences in multi-player games. In *AAMAS*, pages 337–344, 2011.
- [4] R. Zivan. Anytime local search for distributed constraint optimization. In *AAAI-2008*, Chicago, USA, 2008.