Socially Motivated Partial Cooperation in Multi-agent Local Search

Extended Abstract

Tal Ze'evi
Ben-Gurion University of the Negev
Beer-Sheva, Israel
talzee@post.bgu.ac.il

Roie Zivan
Ben-Gurion University of the Negev
Beer-Sheva, Israel
zivanr@bgu.ac.il

Omer Lev
Ben-Gurion University of the Negev
Beer-Sheva, Israel
omerlev@bgu.ac.il

ABSTRACT

Partial Cooperation is a paradigm and a corresponding model for representing multi-agent systems in which agents are willing to cooperate in order to achieve a global goal, as long as some minimal threshold on their personal utility is satisfied. Distributed local search algorithms were proposed in order to solve asymmetric distributed constraint optimization problems (ADCOPs) in which agents are partially cooperative.

We contribute by: 1) extending the partial cooperative model to allow it to represent dynamic cooperation intentions, affected by changes in agents' wealth, in accordance with social studies literature. 2) proposing a novel local search algorithm in which agents receive indications of others' preferences on their actions and thus, can perform actions that are socially beneficial. Our empirical study reveals the advantage of the proposed algorithm in multiple benchmarks. Specifically, on realistic meeting scheduling problems it overcomes limitations of standard local search algorithms.

KEYWORDS

Coordination and Cooperation; Distributed Constraints

ACM Reference Format:

Tal Ze'evi, Roie Zivan, and Omer Lev. 2018. Socially Motivated Partial Cooperation in Multi-agent Local Search. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 3 pages.

1 INTRODUCTION

Multi-agent systems commonly seek to reach an optimal state. One approach considers fully cooperative agents that perform actions in order to achieve a common global goal (e.g. [1, 6, 7]), while another explores agents that are self-interested, which take rational actions that increase their personal gains (e.g. [2]).

A partially cooperative model that handles scenarios which do not fall into these two extreme classes was proposed in [3, 10]. In such settings, agents act cooperatively – motivated by a desire to increase global (group) utility – as long as a minimum condition on their personal utility is satisfied. Such scenarios are common in many real-world settings, e.g., in car navigation applications, where in order to avoid generating traffic jams, some vehicles would be directed to slower routs. However, the driver would not follow the system's instructions if the delay is over some tolerance threshold.

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.

The partial cooperation model represents the willingness of agents to cooperate by defining thresholds of minimum requirements for cooperation. Such willingness has support in social science theory [9].

However, previous attempts to design partially cooperative models assumed people have a fixed reference point according to which they determine their cooperation intentions, and they are so altruistic as to give away any profit they gain, even if they themselves brought it about. In real life, people's intentions for cooperation are affected by changes in their wealth and by who brought it about [4, 9]. This dynamic nature of intentions cannot be expressed by the existing partial cooperation model. Instead, we introduce a model in which agents' cooperation is based on the amount of utility they gain or lose. In the navigating system example, a removal of a road block, which shortens the driving time may result in the driver willingness to tolerate a detour.

The adjustment of the partial cooperative model to realistic social behavior of humans allowed us to analyze the results produced by the two local search algorithms proposed in [3] when solving problems that include additional types of agents, i.e., agents whose cooperation willingness changes in different patterns following changes in their personal utility. This investigation revealed a weakness of these distributed local search algorithms - the agents in these algorithms attempt to find feasible solutions (solutions that satisfy the minimum personal utility requirements of all agents) that maximize their own gains according to their own knowledge. Hence, the willingness of agents to cooperate under some conditions was used only to maintain the validity of the solutions rather than reach the desired high social welfare solution.

Following these insights we propose a novel approach towards partial cooperative local search in which agents indicate to their neighbors which value assignments are preferred by them. These indications allow agents to make socially beneficial selections of value assignments. Our empirical results demonstrate the advantage of the proposed algorithm over previously proposed partially cooperative local search algorithms in solving structured, unstructured and realistic Asymmetric DCOPs (Distributed Constraint Optimization Problem).

2 REFERENCE DEPENDENT PARTIAL COOPERATION

The partial cooperative paradigm was designed such that agents determine their intentions for cooperation using a fixed point of reference. The parameter λ bounds the losses that an agent is willing to undertake in order to contribute to the global objective, i.e., agents perform actions only if they do not result in a cost that exceeds the maximum cost they are willing to endure. Formally, an agent A_i will be willing to cooperate in an interaction, as long as the cost it has to pay in the outcome o, denoted as $C_i(o)$ of the interaction, is within λ_i from the costs it would have to pay if all agents would act selfishly, denoted as μ_i (i.e., "baseline cost"): $C_i(o) \leq \mu_i \cdot (1 + \lambda_i)$. Nevertheless, behavioral economics literature indicates that peoples' intentions change with respect to changes in their wealth (i.e., "reference-dependence"). Reference-dependent theories indicate that people are more sensitive to changes in wealth rather than to absolute wealth level [4, 9]. In order to allow the partial cooperative model to represent dynamic reference points, we redefine some parameters of the model:

Definition 1. *Let $\mu_{i,t}$ be the reference cost of agent A_i at iteration t of the algorithm (where $\mu_{i,0}$ is the baseline cost as defined above).*

Definition 2. *A complete assignment S is feasible in iteration t if it satisfies the following condition: $\forall i \in A, c_i(S) \leq \mu_{i,t} \cdot (1 + \lambda_i)$*

The outcome of a distributed algorithm that runs for m iterations is the complete assignment at the end of the m'th iteration (S_m), and it is feasible if the definition above holds for $\mu_{i,m}$ and λ_i.

Next, we present a number of examples of types of agents that can be represented by the extended model:

- **Type 1** Fixed reference parameter, i.e., for each agent A_i, for each iteration t, $\mu_{i,t} = \mu_{i,0}$ This type of agents is identical to the types described in [3].
- **Type 2** Calculation of μ in each iteration as follows: $\mu_{i,t} = \mu_{i,t-1} + Min\{0, \frac{c_i(S_t) - c_i(S_{t-1})}{1+\lambda_i}\}$ This type resembles people that maintain and manage 'mental budgets' for philanthropic giving (based on *mental accounting* mechanisms cf. [5]).
- **Type 3** This type is inspired by *reciprocal altruism*, in which an individual is willing to cooperate and give up personal wealth for others, with the expectation that they will act in a similar manner in the future [8]. Formally this behavior is represented by a calculation of μ in each iteration as follows: $\mu_{i,t} = \mu_{i,t-1} + Min\{0, \Phi_{i,t-1}\left(\frac{c_i(S_t) - c_i(S_{t-1})}{1+\lambda_i}\right)\}$. where $\Phi_{i,t} = 0$ if the change in agent A_i's cost (i.e., $c(S_t) - c(S_{t-1})$) was caused by an action performed by their neighbor; and $\Phi_{i,t} = 1$ if the change was brought about only by agent A_i's own actions.

3 SOCIALLY-MOTIVATED LOCAL SEARCH

In order to allow agents to exploit the cooperative intentions of their neighboring agents, and so to improve the solution's quality (social welfare), we propose a novel approach towards partial cooperative local search, in which agents take an extra step in the interaction process before selecting an assignment. In this new stage each agent shares with her neighbors some information regarding her preferences over her assignment selection. After exchanging this information, each agent attempts to find a variable assignment, taking into consideration her own preferences as well as the indications received from her neighbors. This approach can be integrated with any local search algorithm, i.e., after the sharing preferences phase, agents could act in accordance to the specification of any local search interaction protocol.

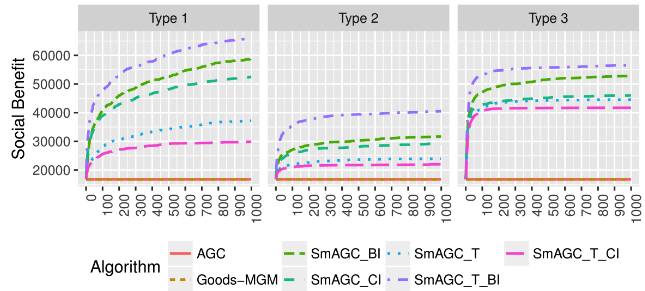

Figure 1: Social benefit for meeting scheduling DCOPs.

We make a distinction between two categories of indications that agents share with their neighbors. The first, we call 'taboo' assignments, i.e., an agent informs her neighbor which of the neighbors' value selections will cause a breach of the current cooperation threshold. The second, which we call a 'vote' allows agents to direct their neighbors to a specific value that they wish the neighbor will select. Such a vote can be binary or weighted.

We combine this approach with the AGC algorithm (cf. [3]) and propose five variants of Socially Motivated (SM) AGC that differ in the type of indications agents share with each other: binary 'votes' (SmAGC_BI), cost 'votes' (SmAGC_CI), 'taboo' (SmAGC_T), 'taboo' + binary 'votes' (SmAGC_T_BI) and 'taboo' + cost 'votes' (SmAGC_T_CI).

Figure 1 presents the aggregated social benefit produced in each iteration of the search for meeting scheduling DCOPs, using $\lambda_i = 0.1, \forall i \in \mathcal{A}$. Results were averaged over 50 random instances, each included 100 agents (n=100) trying to coordinate 50 meetings (m=50), where each agent has 8 optional time slots in their domain (d=8) and the probability of an agent being invited to a meeting is p=0.05. Both AGC and Goods-MGM failed to exploit the willingness of agents to cooperate and produced solutions with very low social welfare. Agents who shared binary votes were able to achieve the highest social welfare and attendance rates with statistical significance of $p < 0.001$, while sharing only 'taboos' or both 'taboos' and 'votes' with exact costs achieved much lower social welfare ($p < 0.01$). Agents of type 3 achieved better results than agents of type 2 for all variants of SM_AGC ($p < 0.001$), and better than agents of type 1 when agents share only 'taboos' or both 'taboos' and 'votes' with exact costs ($p < 0.01$).

4 CONCLUSIONS

We proposed an extension of the partial cooperative paradigm, which allows simulation of realistic scenarios, in which agents intentions for cooperation can change with respect to utility gains. Alongside, we presented a local search algorithm in which the cooperative intentions of agents can be exploited, not only to ensure that the solution obtained is acceptable by all agents , but also in order to select a high quality solution. A significant advantage of the proposed algorithm over the existing partial cooperative algorithms was found even when only insatiability indications ('taboo') were shared. The socially motivated local search algorithm produces high quality solutions on realistic problems with hard equality constraints (meeting scheduling) where standard local search fails.

REFERENCES

[1] A. Gershman, A. Meisels, and R. Zivan. 2009. Asynchronous Forward Bounding. *J. of Artificial Intelligence Research* 34 (2009), 25–46.

[2] J. Grant, S. Kraus, M. Wooldridge, and I. Zuckerman. 2011. Manipulating Boolean Games through Communication. In *IJCAI*. 210–215.

[3] Alon Grubshtein, Roie Zivan, and Amnon Meisels. 2012. Partial Cooperation in Multi-agent Local Search. In *ECAI 2012 - 20th European Conference on Artificial Intelligence. Including Prestigious Applications of Artificial Intelligence (PAIS-2012) System Demonstrations Track, Montpellier, France, August 27-31 , 2012*. 378–383.

[4] Daniel Kahneman and Amos Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the econometric society* (1979), 263–291.

[5] Monica C LaBarge and Jeffrey L Stinson. 2014. The Role of Mental Budgeting in Philanthropic Decision-Making. *Nonprofit and Voluntary Sector Quarterly* 43, 6 (2014), 993–1013.

[6] P. J. Modi, W. Shen, M. Tambe, and M. Yokoo. 2005. ADOPT: asynchronous distributed constraints optimizationwith quality guarantees. *Artificial Intelligence* 161:1-2 (January 2005), 149–180.

[7] Ruben Stranders, Alessandro Farinelli, Alex Rogers, and Nicholas R. Jennings. 2009. Decentralised Coordination of Mobile Sensors Using the Max-Sum Algorithm. In *IJCAI 2009, Proceedings of the 21st International Joint Conference on Artificial Intelligence, Pasadena, California, USA, July 11-17, 2009*. 299–304.

[8] Robert L Trivers. 1971. The evolution of reciprocal altruism. *The Quarterly review of biology* 46, 1 (1971), 35–57.

[9] Pamala Wiepking and Beth Breeze. 2012. Feeling poor, acting stingy: The effect of money perceptions on charitable giving. *International Journal of Nonprofit and Voluntary Sector Marketing* 17, 1 (2012), 13–24.

[10] Roie Zivan, Alon Grubshtein, Michal Friedman, and Amnon Meisels. 2012. Partial cooperation in multi-agent search. In *International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2012, Valencia, Spain, June 4-8, 2012 (3 Volumes)*. 1267–1268.