Polynomial-Time Multi-Agent Pathfinding with Heterogeneous and Self-Interested Agents

Extended Abstract

Manao Machida
NEC Corporation
Kanagawa, 1753 Shimonumabe, Nakahara-ku, Kawasaki
manao-machida@ap.jp.nec.com

ABSTRACT

This paper proposes a polynomial-time strategyproof mechanism that solves multi-agent pathfinding (MAPF) problems with heterogeneous and self-interested agents. In MAPF, agents need to reach their goal destinations while avoiding collisions between them. In this paper, we consider heterogeneous and self-interested MAPF. Agents are heterogeneous if the costs of traversing a given path differ between agents. In particular, we assume each agent has a private linear cost function of travel time. The proposed strategyproof mechanisms aim to make agents truthfully declare the slope of the private linear cost function.

KEYWORDS

agent coordination; multi-agent pathfinding; path planning; mechanism design; self-interested agents

ACM Reference Format:

Manao Machida. 2019. Polynomial-Time Multi-Agent Pathfinding with Heterogeneous and Self-Interested Agents. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 3 pages.

1 INTRODUCTION

In the multi-agent path finding (MAPF) problem, agents on a graph must move from their start vertices to their goal vertices without colliding with each other. An objective of MAPF is to minimize the global cost (e.g., the sum of travel costs of all agents). MAPF has many practical applications in video games, traffic control [7, 13], robotics [5, 9, 16], and vehicle routing [8]. Minimizing the sum of travel costs of MAPF is an NP-complete problem because it is a generalization of the sliding tile puzzle [10].

Most previous work on MAPF focused on homogeneous and cooperative agents who share the global cost function [4, 11–15]. In contrast, agents are heterogeneous if the costs of traversing a given path differ between agents. If agents are self-interested, they seek to minimize their individual travel costs. To the best of our knowledge, only one mechanism for heterogeneous and self-interested MAPF (HSI-MAPF) has been proposed[2].

[2] showed a mapping between MAPF and combinatorial auctions and proposed an iterative combinatorial auction for HSI-MAPF. Their mechanism is a strategyproof mechanism for achieving an optimal solution. However, this mechanism must solve the NP-hard problem to be strategyproof.

In this paper, we assume each agent has a private linear cost function of travel time. This paper proposes a polynomial-time strategyproof mechanism for HSI-MAPF. Strategyproof mechanisms for HSI-MAPF aim to make agents truthfully declare the slope of the private linear cost function. The proposed mechanism is based on a Cooperative A* (CA*) algorithm [13] and a truthful mechanism for one-parameter agents [3].

The contribution of the paper is to propose the polynomial-time strategyproof mechanism for HSI-MAPF. To the best of our knowledge, there has been no research into polynomial-time mechanisms for HSI-MAPF.

2 HETEROGENEOUS MAPF

A MAPF instance M consists of a graph $G = (V, E)$ and a set of agents $K = \{1, \ldots, k\}$. Each agent $i \in K$ has a unique start vertex s_i and a unique goal vertex g_i and needs to go from s_i to g_i while avoiding collisions with other agents. In each time step, agents can move along the edges of the graph or wait in their locations. Agent i conflicts with agent $j \neq i$ if they are situated in the same vertex in the same time step or traverse the same edge in the same time step. Let $T = \{0, 1, \ldots, \}$ be a set of time steps and $P_i \subset V \times T$ be agent i's path which is a set of pairs of vertices and times that agent i passes. The travel time is $n(P_i) = \min\{t \in T | \forall t' \geq t, (g_i, t') \in P_i\}$. Each agent has a private linear cost function of travel time. We call the slope of the function "time step cost." The travel cost of agent i is given by $w_i n(P_i)$, where $w_i \in \mathbf{R}_+$ is the time step cost.

3 STRATEGYPROOF COORDINATED A*

CA* [13] plans each path of the single agent in order. Each path is a shortest path without colliding with paths that have already been planned.

SCA* is modified CA* for heterogeneous and self-interested agents. SCA* decides the sequence of path planning by declarations of time step costs. Moreover, SCA* makes the agents pay monetary costs. Let b be the declarations of agents, $\beta(b)$ be paths that SCA* determines on the basis of b, and $\tau(b)$ be the payment function. The cost of agent i in SCA* is $c_i(b) = w_i n(\beta_i(b)) + \tau_i(b)$.

We assume $b_1 \geq \ldots \geq b_k$ without loss of generality, and if $b_i = b_j$ and $i < j$ then SCA* plans the path of agent i before that of agent j. The SCA* payment is given by

$$\tau_i(b) = \sum_{j>i} b_j \left\{ n(\beta_i(b_{j+1}, b_{-i})) - n(\beta_i(b_j, b_{-i})) \right\} \qquad (1)$$

where $(b'_i, b_{-i}) = (b_1, \ldots, b_{i-1}, b'_i, b_{i+1}, \ldots, b_k)$ and $b_{k+1} = 0$.

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

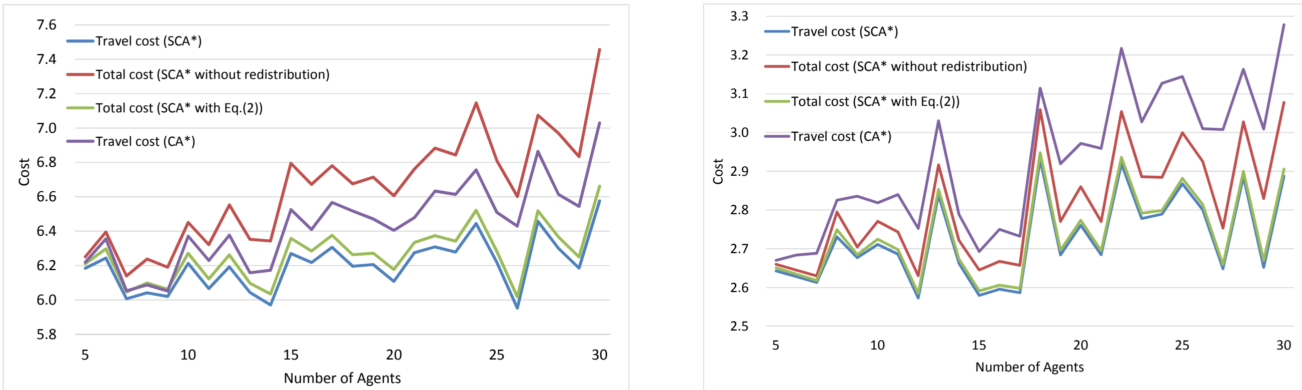

Figure 1: Average costs for each agent on 4-connected 16x16 grid maps with 20% **obstacles, where the time step costs of agents were selected randomly from the uniform distribution on (left)** $[0, 1]$**, (right)** $[0, 0.08] \cup [0.98, 1]$

For example, agents 1, 2 declare time step costs 3.2, 2.5, respectively. If agent 1 declares 0, then the path length of agent 1 is 7, and if agent 1 declares 2.5, then the path length of agent 1 is 5. Hence, the payment of agent 1 is $2.5 * (7 - 5) = 5$.

Theorem 3.1. *If CA* for $M = (G, K)$ in any order has a solution, SCA* for M is strategyproof.*

Proof. If $c_i(w_i, b_{-i}) \leq c_i(b)$ holds for all $b \in \mathbf{R}_+^k$, then agent i has no incentive to lie.

$$c_i(w_i, b_{-i}) - c_i(b) = (b_i - w_i) n(\beta_i(b)) - \int_{w_i}^{b_i} n(\beta_i(c, b_{-i})) dc \leq 0$$

holds because $n(\beta_i(\cdot, b_{-i}))$ is monotonically non-increasing. □

We also shows that SCA* is a polynomial-time mechanism.

Theorem 3.2. *β runs in $O(k^2 |V| |E|)$, and τ runs in $O(k^3 |V| |E|)$.*

Proof. CA* runs in $O(k^2 |V| |E|)$, so β runs in $O(k^2 |V| |E|)$. τ_i must execute CA* without agent i and execute A* to calculate agent i's path at most k times. Hence, τ runs in $O(k^3 |V| |E|)$. □

SCA* can redistribute a part of the payments for reducing the costs of agents. The redistribution h based on redistribution methods [6] is given by

$$h_i(b_{-i}) = \frac{1}{k} \min_{l \in K \cup \{k+1\} \setminus \{i\}} \sum_{j \in K} \tau_j(b_l, b_{-i}). \tag{2}$$

The cost of agent i is $c_i(b) = w_i \beta_i(b) + \tau_i(b) - h_i(b_{-i})$. SCA* with redistribution is also a strategyproof and polynomial-time mechanism.

4 EXPERIMENTAL RESULTS

We conducted an empirical evaluation of SCA*. Table 1 shows the runtime of each function when solving problems on a 4-connected 8×8 grid map, with the number of agents ranging between 8 and 14. All times were measured on an Intel(R) Core(TM) i7-6700 CPU @ 3.40 GHz. The setting of the problem is the same as that in [1] who compared the iterative combinatorial auction for MAPF and other optimal solvers, so we can compare our result and Table 1 in [1]. SCA* scales well with the number of agents. However, the iterative combinatorial auction and other optimal solvers do not scale well with the number of agents.

Table 1: Runtime (ms) of SCA* with Eq. (2) on 4-connected 8x8 grid map with no obstacles

k	β	τ	h
8	0.27	0.68	37.20
9	0.23	0.67	54.34
10	0.23	0.85	86.06
11	0.24	1.01	124.91
12	0.25	1.29	186.97
13	0.26	1.51	257.62
14	0.28	1.77	341.54

Figure 1 (left) shows the average costs of a single agent in SCA* and CA* on 4-connected 16×16 grid maps with 20% obstacles, where the time step costs of agents were selected randomly from the uniform distribution between 0 and 1. SCA* had a lower travel cost than CA*. The total cost is the sum of the travel cost and the payment. SCA* with the redistribution had a lower total cost than CA*.

We show that, where only a few agents have high time step costs, SCA* without the redistribution can have a lower total cost than CA*. Figure 1 (right) shows the average cost of a single agent on maps the same as previous maps, where the time step costs of agents were selected randomly from the uniform distribution on $[0, 0.08] \cup [0.98, 1]$. In the problems, the agents' time step costs were polarized, and only 20% agents had the high time step costs. The total cost of SCA* was lower than that of CA* in these problems, even if SCA* does not conduct the redistribution.

ACKNOWLEDGMENTS

This work was performed of the METI R&D Program Drones and Robots for Ecologically Sustainable Societies Project supported by NEDO.

REFERENCES

[1] Ofra Amir, Guni Sharon, and Roni Stern. 2014. *An Empirical Evaluation of a Combinatorial Auction for Solving Multi-Agent Pathfinding Problems.* Technical Report TR–05–14. Harvard University.

[2] Ofra Amir, Guni Sharon, and Roni Stern. 2015. Multi-agent Pathfinding As a Combinatorial Auction. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI'15)*. AAAI Press, 2003–2009. http://dl.acm.org/citation.cfm?id=2886521.2886599

[3] A. Archer and É. Tardos. 2001. Truthful Mechanisms for One-Parameter Agents. In *Proceedings of the 42Nd IEEE Symposium on Foundations of Computer Science (FOCS '01)*. IEEE Computer Society, Washington, DC, USA, 482–. http://dl.acm.org/citation.cfm?id=874063.875583

[4] Max Barer, Guni Sharon, Roni Stern, and Ariel Felner. 2014. Suboptimal Variants of the Conflict-based Search Algorithm for the Multi-agent Pathfinding Problem. In *Proceedings of the Twenty-first European Conference on Artificial Intelligence (ECAI'14)*. IOS Press, Amsterdam, The Netherlands, The Netherlands, 961–962. https://doi.org/10.3233/978-1-61499-419-0-961

[5] M. Bennewitz, W. Burgard, and S. Thrun. 2002. Finding and Optimizing Solvable Priority Schemes for Decoupled Path Planning Techniques for Teams of Mobile Robots. *Robotics and Autonomous Systems* 41, 2 (2002), 89–99.

[6] Ruggiero Cavallo. 2006. Optimal Decision-Making With Minimal Waste: Strategyproof Redistribution of VCG Payments. In *Proc. of the 5th Int. Joint Conf. on Autonomous Agents and Multi Agent Systems (AAMAS'06)*. Hakodate, Japan, 882–889. http://econcs.seas.harvard.edu/files/econcs/files/cavallo-redis.pdf

[7] Kurt Dresner and Peter Stone. 2008. A Multiagent Approach to Autonomous Intersection Management. *J. Artif. Int. Res.* 31, 1 (March 2008), 591–656. http://dl.acm.org/citation.cfm?id=1622655.1622672

[8] Scott Kiesel, Ethan Burns, Christopher Wilt, and Wheeler Ruml. 2012. Integrating Vehicle Routing and Motion Planning. In *ICAPS*.

[9] L. Pallottino, V. G. Scordio, A. Bicchi, and E. Frazzoli. 2007. Decentralized Cooperative Policy for Conflict Resolution in Multivehicle Systems. *Trans. Rob.* 23, 6 (Dec. 2007), 1170–1183. https://doi.org/10.1109/TRO.2007.909810

[10] Daniel Ratner and Manfred Warrnuth. 1986. Finding a shortest solution for the $N \times N$ extension of the 15-puzzle is intractable. In *AAAI-86*. 168–172.

[11] Guni Sharon, Roni Stern, Ariel Felner, and Nathan R. Sturtevant. 2015. Conflict-based Search for Optimal Multi-agent Pathfinding. *Artif. Intell.* 219, C (Feb. 2015), 40–66. https://doi.org/10.1016/j.artint.2014.11.006

[12] Guni Sharon, Roni Stern, Meir Goldenberg, and Ariel Felner. 2013. The increasing cost tree search for optimal multi-agent pathfinding. *Artificial Intelligence* 195 (2013), 470–495. https://doi.org/10.1016/j.artint.2012.11.006

[13] David Silver. 2005. Cooperative Pathfinding. In *Proceedings of the First AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE'05)*. 117–122.

[14] Trevor Standley. 2012. Independence Detection for Multi-Agent Pathfinding Problems. In *MAPF@AAAI*.

[15] Trevor Standley and Richard Korf. 2011. Complete Algorithms for Cooperative Pathfinding Problems. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume Volume One (IJCAI'11)*. AAAI Press, 668–673. https://doi.org/10.5591/978-1-57735-516-8/IJCAI11-118

[16] Glenn Wagner and Howie Choset. 2011. M*: A complete multirobot path planning algorithm with performance bounds. *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems* (2011), 3260–3267.