Preference-Based Multi-Objective Multi-Agent Path Finding

JAAMAS Track

Florence Ho
National Institute of Advanced Industrial Science and Technology (AIST)
Tokyo, Japan
florence.ho@aist.go.jp

Shinji Nakadai
NEC Corporation,
National Institute of Advanced Industrial Science and Technology (AIST)
Tokyo, Japan
nakadai@nec.com

ABSTRACT

Multi-Agent Path Finding (MAPF) consists in computing a set of collision-free paths for a team of agents on a given graph while minimizing one objective, such as the sum of paths costs or the makespan. However, real-world applications may require the consideration of multiple objectives. Thus, in this work, we propose to address a novel extension of MAPF, Scalarized Multi-Objective MAPF (MOMAPF), that aims to optimize multiple given objectives while computing collision-free paths for all agents and incorporating the preferences of a decision maker over each objective. The preferences of a decision maker are reflected by a weight value associated to each objective and all weighted objectives are combined into one scalar to minimize. We introduce a solver for Scalarized MOMAPF based on Conflict-Based Search (CBS), Scalarized MO-CBS, that incorporates an adapted path planner based on an evolutionary algorithm, the Genetic Algorithm (GA). We also introduce three practical objectives to consider in path planning: efficiency, safety, and smoothness. We evaluate the performance of our proposed method in function of the input parameters of GA on experimental simulations.

KEYWORDS

Multi-Objective Optimization; Multi-Agent Path Finding; Path Planning

ACM Reference Format:

Florence Ho and Shinji Nakadai. 2023. Preference-Based Multi-Objective Multi-Agent Path Finding: JAAMAS Track. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

1 INTRODUCTION

The MAPF problem aims at planning a set of paths for a group of agents from given start locations to goal locations while avoiding collisions. A significant amount of research works have been devoted to MAPF for various real-world applications, such as warehouse logistics [6, 12], and traffic management [1, 5]. In particular, the MAPF problem assumes the optimization of a single objective, such as the sum of paths lengths or makespan. However, in many real-world multi-robot systems [9, 13], several objectives may need to be considered in the generation of a conflict-free solution. Recently, a novel extension of the MAPF problem, Multi-Objective MAPF (MOMAPF) has been proposed [10]. The MOMAPF problem aims to determine conflict-free paths for all agents while optimizing several given objectives. The goal of MOMAPF, with multiple objectives, is to find the set of Pareto-optimal solutions[1].

In our work [2], we propose a novel extension of MAPF to multi-objective optimization, Scalarized MOMAPF, that allows us to incorporate the preferences of a human decision maker (DM) into multi-agent path planning over each given objective. We aim to determine one solution that belongs or is as close as possible to the Pareto-optimal set. This solution represents the preferred solution of a DM. We propose a scalarization of objectives [4, 8] that converts a multi-objective optimization problem into a single objective optimization problem. We integrate a DM's preferences over each objective with the use of weights values assigned to each objective. Moreover, we propose a solver for Scalarized MOMAPF based on the Conflict-Based Search (CBS) algorithm, Scalarized MO-CBS, whereby we incorporate an adapted path planning method based on an evolutionary method, the Genetic Algorithm (GA). In particular, we propose to address three practical objectives for multi-agent path planning: efficiency, safety, and smoothness. We evaluate our approach on benchmark simulations scenarios. Although our solver lacks theoretical guarantees, we empirically measure the trade-offs between solution costs and runtime in function of the different input parameters of GA that influence the quality of the solutions generated by our approach.

2 PROPOSED FRAMEWORK

2.1 Scalarized MOMAPF

Let $A = \{a_1, \dots, a_N\}$ denote a set of N agents in a 2D space. The space is represented by a map that is a uniform grid G. Each agent a_i has a start location s_i and a goal location g_i to reach. For any agent, each action can either be wait or move and requires one unit of time. Let $X = \{x_1, \dots, x_N\}$ represent the set of paths for all the agents, with x_i ($i \in \{1, \dots, N\}$) a path associated to the agent a_i. We define the path x_i that connects the start location s_i to the goal location g_i by a sequence of waypoints $x_i = (s_i, p_1, p_2, \dots, g_i)$. A waypoint is a coordinate location in G. X is also called a solution. We consider K objectives, each defined by a function f_k ($k \in \{1, \dots, K\}$) that computes the value for each objective for each path x_i. The cost of a solution for each objective k is defined as the sum of all the individual path costs over all agents: $f_k(X) = \sum_{x_i \in X} f_k(x_i)$. In this work, we propose to solve a scalarization of the MOMAPF problem,

[1] A solution is Pareto-optimal if there exists no other solution that will provide an improvement in one objective without causing a deterioration in at least one of the other objectives.

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

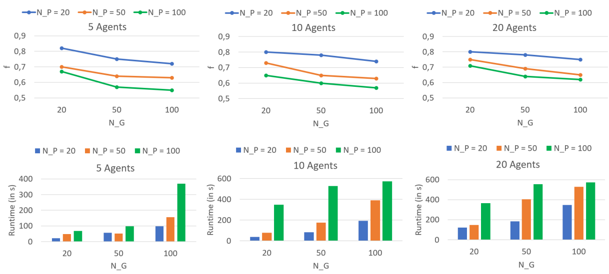

Figure 1: Comparisons of average costs of solutions and runtime obtained by Scalarized MO-CBS with GA for different population sizes N_P, number of generations N_G, and number of agents, in randomly generated 30×30 maps with 10% obstacle density.

whereby a weight $(w_k)_{k \in 1,\ldots,K}$ such that $\sum_{k=1}^{K} w_k = 1$ is associated to each objective. We assume each weight represents a level of preference from a DM for each objective. A simple scalarizing function is the weighted sum as shown in Equation 1. This method has been shown to work well when all the objectives are convex [8]. Another possible scalarization method that works regardless of the convexity of objectives is the Tchebycheff approach [14] as shown in Equation 2.

$$\textit{Weighted sum} \quad \sum_{k=1}^{K} w_k \cdot f_k(X) \quad \textit{with} \quad \sum_{k=1}^{K} w_k = 1 \tag{1}$$

$$\textit{Tchebycheff} \quad \max_{1 \leq k \leq K} \{ w_k \cdot |f_k(X) - z_k^{Utopia}| \} \quad \textit{with} \quad \sum_{k=1}^{K} w_k = 1 \tag{2}$$

Since different objective functions may have different magnitudes, the normalization of objectives is required to obtain a solution consistent with the weights assigned by the DM [7]. In particular, the normalization requires the computation of a Utopia point and a Nadir point. The Utopia point $z_k^{Utopia} = \sum_{a_i \in A} z_{ik}^{Utopia}$ is computed with the optimal solution for each agent a_i with respect to the k^{th} objective, and the Nadir point $z_k^{Nadir} = \sum_{a_i \in A} z_{ik}^{Nadir}$ is computed with the worst solution in the Pareto-optimal set for each agent a_i with respect to the k^{th} objective.

2.2 Scalarized MO-CBS with GA

We propose an extension of CBS [11] to solve Scalarized MOMAPF instances, Scalarized MO-CBS. This approach keeps the same steps as in standard CBS but differs in two points:

(1) **A different cost function**: the computation of the cost value for a constraint tree node takes into account the multiple objectives with the scalarization approach. We hereby consider three objectives: efficiency (minimizing path length), safety (maximizing the distance of each waypoint of a given path to all obstacles), and smoothness (minimizing the number of turns in a given path).
(2) **A different path planner**: we propose a path planner based on the Genetic Algorithm (GA) that allows us to address the multiple objectives in a scalarized form.

We also adapt the steps of the GA [3] to path planning. Each candidate in a population is a feasible path composed of a sequence of waypoints from the start to the goal location of the given agent. The crossover step in GA exchanges parts of paths between selected candidates to generate new candidates (offsprings). The mutation step in GA randomly changes some waypoints of a given candidate.

Due to the probabilistic nature of metaheuristics such as GA, there are no theoretical guarantees on optimality. Consequently, the solution returned by Scalarized MO-CBS is not guaranteed to be Pareto-optimal. So in this work, we focus on the evaluation of our approach by varying the input parameters of the GA-based path planner that influence the quality of a solution.

3 EXPERIMENTS

We evaluate the performance of our proposed approach Scalarized MO-CBS with a GA-based path planner on a benchmark of generated grid maps. We consider 30×30 random maps with different percentages of static obstacles randomly generated, 5%, 10%, and 20%.

We vary the number of agents in the maps from 5 to 20. We set a runtime limit of 10 minutes. We assume that a DM has already given weights values to quantify his or her preferences over each objective. The efficiency of our proposed approach Scalarized MO-CBS relies on the input parameters of GA, namely, the population size, the number of generations, the mutation and crossover rates. In all considered instances, we observed that increasing the population size and the number of generations decrease the costs of the solutions returned by our solver with GA, while leading to an increase in the runtime, as shown in Figure 1 for example. We also observed that setting a small mutation rate and a high crossover rate allow us to generate reasonable improvements in the solutions.

4 CONCLUSIONS

This work presents a novel extension of the MAPF problem, the Scalarized MOMAPF problem, that allows us to incorporate preferences over several given objectives with the use of weights values while generating conflict-free paths for multiple agents. This allows us to determine more realistic conflict-free paths with the consideration of multiple objectives. We propose a novel approach, Scalarized MO-CBS, that incorporates the scalarization of multiple objectives, and an adapted path planner based on the Genetic Algorithm (GA). Since GA is a metaheuristic, our approach does not guarantee Pareto optimality of the solution generated by Scalarized MO-CBS. The solution quality generated by our approach relies mainly on the input parameters of GA, namely, the population size, the number of generations, the crossover and mutation operators. On the other hand, the use of a metaheuristic allows us to reduce the computational effort in obtaining a satisfactory possibly near-optimal result. Thus, we provide an empirical analysis of the performance of our solver on experimental simulations with varying obstacles densities. Overall, our presented framework constitutes a novel perspective in the MAPF field and could help to further expand MAPF to real-world applications.

REFERENCES

[1] Florence Ho, Rúben Geraldes, Artur Goncalves, Bastien Rigault, Benjamin Sportich, Daisuke Kubo, Marc Cavazza, and Helmut Prendinger. 2022. Decentralized Multi-Agent Path Finding for UAV Traffic Management. *IEEE Transactions on Intelligent Transportation Systems* 23, 2 (2022), 997–1008.

[2] Florence Ho and Shinji Nakadai. 2023. Preference-Based Multi-Objective Multi-Agent Path Finding. *Autonomous Agents and Multi-Agent Systems* 37, 1 (2023), 1–22.
[3] John H. Holland. 1992. Genetic Algorithms. *Scientific American* 267, 1 (1992), 66–73.
[4] Xiujuan Lei and Zhong Shi. 2004. Overview of multi-objective optimization methods. *Journal of Systems Engineering and Electronics* 15 (2004), 142–146.
[5] Jiaoyang Li, Zhe Chen, Yi Zheng, Shao-Hung Chan, Daniel Harabor, Peter J. Stuckey, Hang Ma, and Sven Koenig. 2021. Scalable Rail Planning and Replanning: Winning the 2020 Flatland Challenge. *Proceedings of the International Conference on Automated Planning and Scheduling* 31 (2021), 477–485.
[6] Hang Ma, Jiaoyang Li, T. K. Satish Kumar, and Sven Koenig. 2017. Lifelong Multi-Agent Path Finding for Online Pickup and Delivery Tasks. In *International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 837–845.
[7] Helmut Mausser. 2006. Normalization and Other Topics in Multi-Objective Optimization. In *Proceedings of the Fields-MITACS Industrial Problems Workshop*.
[8] Kaisa Miettinen and Marko Makela. 2002. On scalarizing functions in multiobjective optimization. *OR Spectrum* 24 (2002), 193–213.
[9] Justin Montoya, Sivakumar Rathinam, and Zachary Wood. 2014. Multiobjective Departure Runway Scheduling Using Dynamic Programming. *IEEE Transactions on Intelligent Transportation Systems* 15, 1 (2014), 399–413.
[10] Zhongqiang Ren, Sivakumar Rathinam, and Howie Choset. 2021. Multi-objective Conflict-based Search for Multi-agent Path Finding. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. 8786–8791.
[11] Guni Sharon, Roni Stern, Ariel Felner, and Nathan Sturtevant. 2012. Conflict-Based Search For Optimal Multi-Agent Path Finding. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI)*. 563–569.
[12] Roni Stern, Nathan Sturtevant, Ariel Felner, Sven Koenig, Hang Ma, Thayne Walker, Jiaoyang Li, Dor Atzmon, Liron Cohen, T. K. Satish Kumar, Eli Boyarski, and Roman Bartak. 2019. Multi-Agent Pathfinding: Definitions, Variants, and Benchmarks. In *Twelfth Annual Symposium on Combinatorial Search (SOCS)*. 151–159.
[13] Ryoji Tanabe and Hisao Ishibuchi. 2020. An Easy-to-use Real-world Multi-objective Optimization Problem Suite. *Applied Soft Computing* 89 (2020), 106078.
[14] Qingfu Zhang and Hui Li. 2007. MOEA/D: A Multiobjective Evolutionary Algorithm Based on Decomposition. *IEEE Transactions on Evolutionary Computation* 11 (2007), 712 – 731.