Research Project Proposal: Exploiting Task Distribution in Multi-Agent Pickup and Delivery

Andrea Di Pietro, andrea4.dipietro@mail.polimi.it

1. Introduction to the problem

1.1. Research topic

Multi-Agent Pickup and Delivery (MAPD) is a problem in the field of multi-robot systems and consists in completing a set of tasks that require agents to move from a pickup location to a delivery location avoiding conflicts with other agents. The set of tasks can be updated at each time step with new tasks to be executed. This problem combines two different known problems: task assignment and Multi-Agent Path Finding (MAPF). In particular MAPD is considered as an extension of the MAPF problem because it removes the assumptions that each agent executes exactly one task and that the task assignment is defined in advance.

MAPD is a problem of great importance because of its applications in many real-world scenarios such as autonomous vehicles and automated warehouses [1][6]. In fact MAPD represents a more complex and realistic problem than MAPF considering that in real applications agents are usually required to perform more than one task and the list of tasks is not always known in advance but it is updated dynamically, and this forces the system to continuously replan the agents' paths and task assignments accordingly. For these reasons new researches about solutions for MAPD problems are being developed on top of the current works and solutions for MAPF problems.

1.2. Research problem

The problem that will be addressed in the research project concerns the possibility to use the probability distribution of incoming tasks to improve the efficiency of MAPD solutions. Differently from the literature we assume that, even though the exact pickup and delivery locations and the arrival time of tasks in the system are unknown, a probability distribution of these properties across the environment is available and we exploit this information when the system performs task assignment or path planning in order to improve the overall performance of the solution [4].
Existing algorithms assume that nothing about tasks is known until a task is added to the system, but in some real-world scenarios it is reasonable to assume that some information is available in advance. For example, in the context of a warehouse that manages the pickup and delivery of items for orders from customers there will exist periods of the year or times of the day in which some items are in greater demand with respect to some others. These data are usually known or can be reliably estimated and they are already analysed for other purposes (for example supply management) but they could also have a positive impact if considered during the resolution of MAPD instances because they give information about the frequency, the expected arrival time and the pickup location of single tasks.
The probability distribution of one or more of these properties (the pickup location, the delivery location, the task frequency in a given time window) could be exploited when the system performs path planning and when it chooses which agent should be assigned a new task by considering that some locations or zones of the environment are more likely to need free agents in their area in the near future because new tasks are most probably expected

there. This could allow the system to optimize the use of its own agents and to reduce the average service time and the makespan of the solution.

2. Main related works

The works about Multi-Agent Pickup and Delivery have been focusing on the formal definition of the problem and on the proposal of different algorithms that can be classified according to two opposite assumptions:

- Some algorithms assume that the sequence of tasks is known *a priori* when the system starts to execute tasks and they refer to an offline formulation of MAPD by assigning all tasks and computing all the paths only once at the beginning without waiting for the tasks to be added to the system [2].
- Other algorithms adopt an online strategy by considering the new tasks only when they are added to the system and they perform path replannigs according to the properties of the new tasks [3][1]. These algorithms can differ in the way they perform the task assignment and in the way they plan the paths of single agents, but a common aspect is that they can't take advantage of any knowledge about probability distribution of the incoming tasks.

While the former type of algorithms makes the strong assumption that all information about tasks is known from the very beginning, the latter makes an equally strong assumption that any kind of information about tasks can't be available in advance. This research aims to propose a valuable alternative to these methods in the case in which none of the two assumptions is valid for a given MAPD problem and the actual scenario lies in the middle between these two extreme hypotheses.

3. Research plan

3.1. Research goals

The goal of the research is to analyse how knowledge about the probability distribution of the task arrivals and of the task locations across the environment can be exploited to improve the performance of a MAPD solution. Starting from the analysis of the current solutions the goal is to design and test the performance of new MAPD algorithms that use the model of the distribution of tasks to perform tasks assignment and path planning. This means that the research starts from theoretical analysis in order to provide some solutions that can be applied in real-world applications and it includes an experimental phase to analyse and to evaluate empirically the performance of the developed algorithms compared to existing MAPD solutions.

3.2. Research plan decomposition

The research is organized in the following steps:

1. Analysing the State of the Art, the properties of existing MAPF and MAPD solutions and the performance and limits of existing algorithms.
2. Defining a new formal model for MAPD problems which takes into account information about the probability distribution of tasks that is currently ignored and not modelled in existing frameworks.
3. Designing new MAPD algorithms exploiting the probability distribution of tasks to perform a better assignment of tasks to the agents and/or a better path planning. Properties such as the completeness of the solution and the computational complexity bounds of the algorithms will be analysed.
4. Implementing a framework that allows to test the results of the application of the existing and of the new algorithms in different configurations of MAPD problems.

5. Testing the new algorithms that exploit the probability distribution of tasks against the existing solutions that ignore it in different conditions and configurations of the environment. This will allow to compare, according to the classical performance evaluation metrics, the different algorithms. It is also important to evaluate the performance in environments and with assumptions that are realistic for real-world scenarios such as a classical warehouse environment.

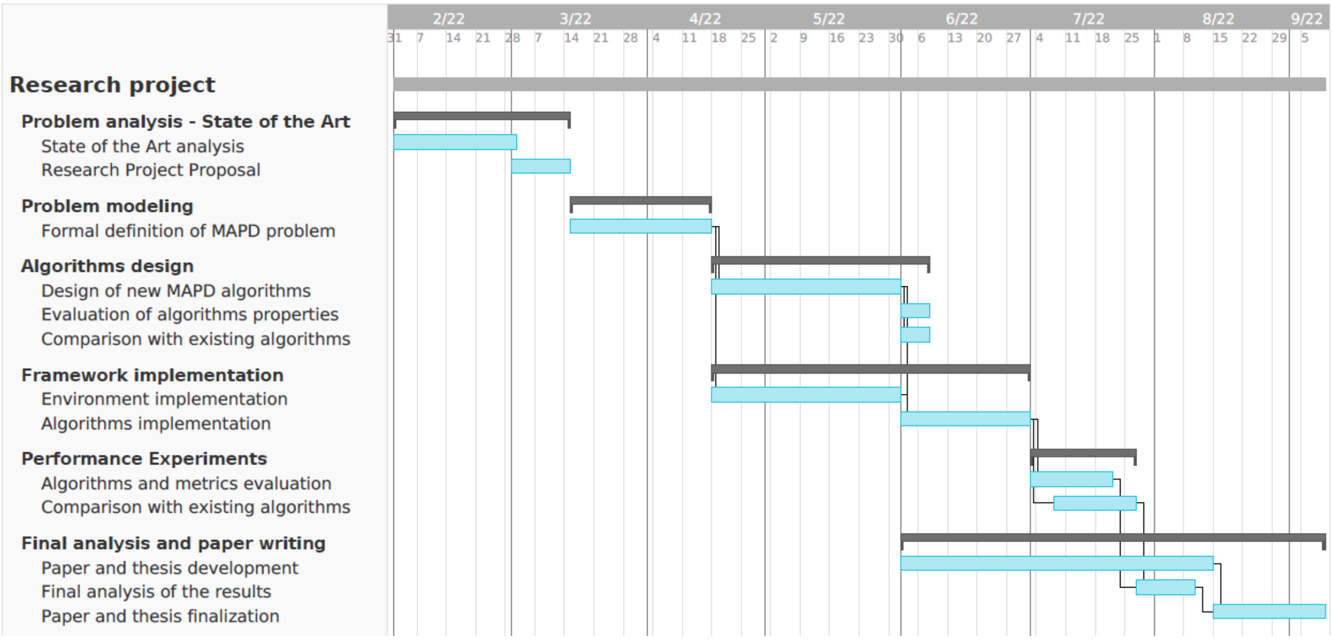

3.3. Evaluation metrics

The two classical objective functions that are used to evaluate MAPD problems and that derive from similar MAPF evaluation metrics should be employed as metrics for the new MAPD algorithms as well [5].

- *Makespan*: the number of time steps after which all the tasks have been executed.
- *Service time*: the average number of time steps between the addition of a task to the system and the completion of the task itself.

Another interesting metric that is commonly employed for the performance analysis in real warehouses is the *throughput*, that is the the average number of completed tasks per time step [1].
Other properties of the algorithms that should be taken into consideration are the space and time computational complexity. Both the theoretical bounds and the results provided by the experiments are fundamental to understand whether the algorithms can scale to systems with a large number of agents and tasks.

References

[1] Li, J., Tinka, A., Kiesel, S., Durham, J., Kumar, S., and Koenig, S. Lifelong multi-agent path finding in large-scale warehouses. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)* (2021), pp. 11272–11281.

[2] Liu, M., Ma, H., Li, J., and Koenig, S. Task and path planning for multi-agent pickup and delivery. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (2019), p. 1152–1160.

[3] Ma, H., Li, J., Kumar, S., and Koenig, S. Lifelong multi-agent path finding for online pickup and delivery tasks. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (2017), pp. 837–845.

[4] Salzman, O., and Stern, R. Research challenges and opportunities in multi-agent path finding and multiagent pickup and delivery problems. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (2020), pp. 662–667.

[5] Stern, R., Sturtevant, N., Felner, A., Koenig, S., Ma, H., Walker, T., Li, J., Atzmon, D., Cohen, L., Kumar, S., Boyarski, E., and Bartak., R. Multi-agent pathfinding: Definitions, variants, and benchmarks. In *Proceedings of the Symposium on Combinatorial Search (SoCS)* (2019), pp. 151–158.

[6] Wurman, P. R., D'Andrea, R., and Mountz, M. Coordinating hundreds of cooperative, autonomous vehicles in warehouses. *AI magazine 29, 1* (2008), 9–20.