Research Project Proposal: Robustness in Multi-Agent Pickup and Delivery

Giacomo Lodigiani, giacomo.lodigiani@polimi.it

1. Introduction to the problem

Multi-Agent Pickup and Delivery is a problem in which agents must assign targets, which may continuously appear at any time step, among themselves and plan paths to those targets, avoiding collisions. This problem, which combines long-term planning and target assignment, presents great practical relevance; we will argue the importance of finding solutions resilient to adverse events during operations and we will propose, in this research project, novel methods to find those solutions.

1.1. Research topic

Multi-Agent Pickup and Delivery (MAPD) derives from the combination of two known problem in the field of Multi-Agent Systems: Multi-Agent Path Finding (MAPF, the planning problem), and task assignment; these two problems must be solved jointly in order to solve MAPD. The MAPD problem is long-term and dynamic: new tasks may be added at any time step. Despite being studied only recently, its application is of great practical importance, since the MAPD formalization is the one that better describes what happens in automated warehouses, where robots must continuously deliver new orders to specific locations. An example is the successful Amazon Robotics system employed by Amazon [7]. The MAPD problem has been proven to be NP-hard [4] so its complexity grows exponentially with the number of agents; so far only few online algorithms have been proposed in theory [5] to address it, but how to solve this problem in a realistic setting, with possible adverse events (like failure or attacks) and maintaining an acceptable level of performance, still remains an open question.

1.2. Research question

The main focus of our research will be on solving the MAPD problem taking into account the fact that the agents may not behave as expected, for example in case of failure or attack. In the former case an agent deviates from its planned path following some probability distribution (for example older robots are more likely to fail), in the latter the behaviour of one or more agents will be adversarial, trying to cause the maximum damage possible. Keeping into account the different information available in these situations, we would like to find for each one algorithms returning long-term robust solutions, i.e., following the definition given by Ma et al. [5], algorithms returning a solution, if one exists, that finishes a finite set of tasks in a finite amount of time. Until now, only few algorithms have been proposed to solve the MAPD problem, and none of them considers the impact on robustness of agents not strictly following the computed plan. This event is not so unlikely in a real world scenario, considering for example a typical warehouse with hundreds of robots and possible human interaction; so, solutions to this problem could have great practical benefits and could help the development of a new generation of more reliable, secure and efficient Multi-Agent Systems.

2. Main related works

Study on the theoretical problem of Multi-Agent Pickup and Delivery has just started, so the literature is not very broad; our research project will start from the work carried out by Ma, Koening et al. [4][5], which clearly defines the framework of the MAPD problem and the concept of long-term robustness. A much wider literature investigates various issues of just one component of MAPD, Multi-Agent Path Finding, but could still give us interesting insights. Li et al. [3] address the problem of *Lifelong* MAPF in a warehouse environment by decomposing it into a sequence of Windowed MAPF (collisions need to be resolved only for the first w time steps) instances solved by continuously re-planning after a set amount of time steps, considering task preassigned to agents; Hoenig et al. [2] study the same problem of MAPF in a warehouse environment and propose as a solution the use of particular graphs, called Action Dependency Graphs, that capture the action precedence relationships of a MAPF solution and can be used to enforce these relationships on real robots with higher-order dynamics. Atzmon et al. [1] analyze a particular type of failure in the MAPF context, delays, and offer a solution based on a probabilistic framework. In conclusion, state of the art on the topic is scarce, especially regarding application of theory in real world scenarios; this means that opportunities for research are broad and advancements in the topic could bring tangible improvements in many implementations of Multi-Agent Systems.

3. Research plan

3.1. Research goals

The main goal of our research is to find solutions to the MAPD problem that can be reliably applied in real contexts, where actions are not always executed as planned and unexpected events may happen. We will study theoretical probabilistic and game theory frameworks to describe these events (like failures and attacks) and we will propose MAPD algorithms to find long-term robust solutions in these situations. Therefore, our research is a blend of theoretical analysis and algorithm design, ultimately aimed at closing the gap that stills exists between theory and applications in the MAPD domain.

3.2. Research plan decomposition

Our research is organized in the following steps:

1. Analyze existing MAPD algorithms in terms of long-term robustness and extend them with simple recovery routines to allow the completion of plans even in case of unexpected events.
2. Formalize frameworks to describe failures and attacks in the MAPD context and use them to evaluate the performance impact on MAPD algorithms with recovery routines.
3. Design and implement new long-term robust MAPD algorithms.
   - Focus on robustness at *planning* time, avoiding time expensive recovery routines.
   - Validate the approach comparing performance (using metrics at Section 3.3) with algorithms implementing naive recovery routines.
   - Estimate theoretical bounds for complexity and plan length.
4. Study the trade-off between robustness and performance, possibly using classical MAPF benchmarks [6] .
   - Compute evaluation metrics at Section 3.3 for the new algorithms and compare them with those of classical algorithms in cases where failures and attacks happen or not.
   - Study the importance of tuning the parameters of the new algorithms.

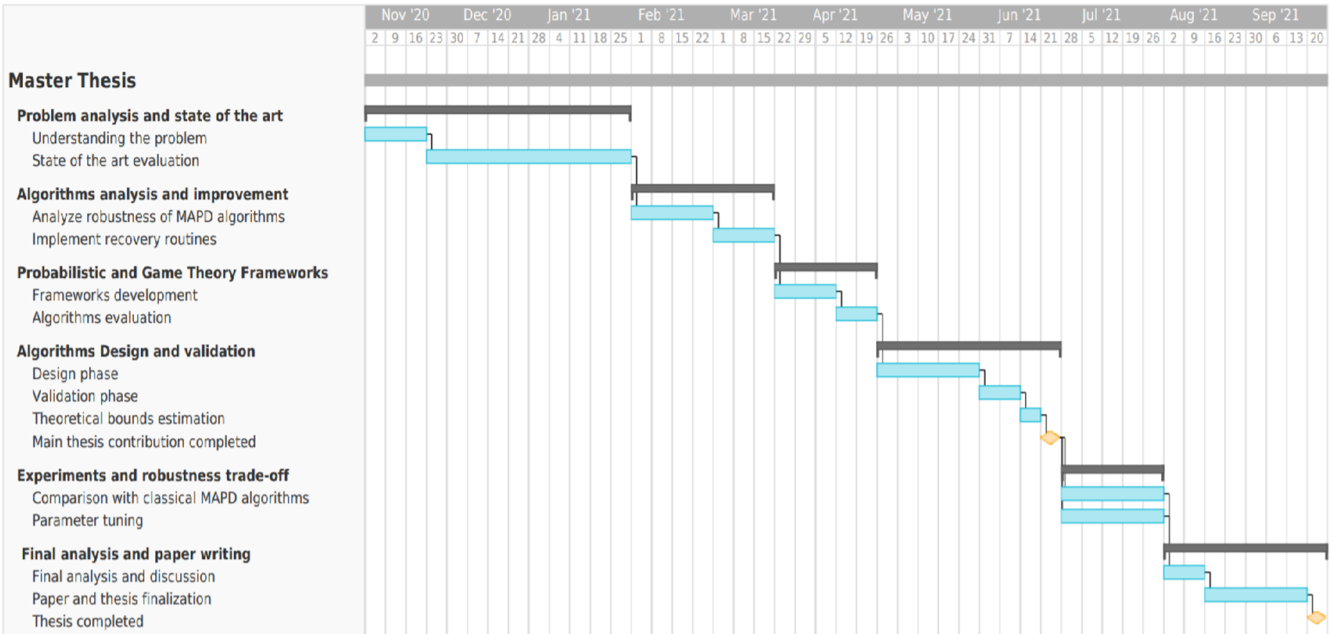

Figure 1: Gantt diagram for our project, which should be concluded by September 2021.

3.3. Evaluation metrics

To evaluate and compare our algorithms, we will use classical metrics employed in the MAPF and MAPD domain.

- Makespan: earliest time step when all tasks are finished.
- Service time: average number of time steps needed to finish executing each task after it was added to the system.

In particular, we will use these metrics to study the impact of adverse events and we will analyze how different algorithms manage to reduce the repercussions on performance caused by an increasing rate of said events. Also we will give importance to the time and space complexity of our new implemented algorithms.

References

[1] Atzmon, D., Stern, R., Felner, A., Sturtevant, N., and Koenigz, S. Probabilistic robust multi-agent path finding. In *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)* (2020), pp. 29–37.

[2] Hoenig, W., Kiesel, S., Durham, A. T. J., and Ayanian, N. Persistent and robust execution of mapf schedules in warehouses. *IEEE Robotics and Automation Letters, 4, (2)* (2019), 1125–1131.

[3] 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 International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (2020).

[4] Ma, H. Target assignment and path planning for navigation tasks with teams of agents. *PhD thesis, Department of Computer Science, University of Southern California, Los Angeles (California)* (2020).

[5] 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.

[6] 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).

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