

Representation, Planning, Execution, and Their Integration in Multiple Moving Agents

Keisuke Okumura¹, Yasumasa Tamura¹, Xavier Défago¹

¹ School of Computing, Tokyo Institute of Technology
okumura.k@coord.c.titech.ac.jp, {tamura, defago}@c.titech.ac.jp

keywords: multi-agent path finding, anytime planning, time-independence, robotic platform

Intelligent multi-agent systems with physically moving agents are attractive for various practical scenarios, e.g., intersection management [Dresner and Stone, 2008], automated warehouse [Wurman *et al.*, 2008], and vehicle parking [Okoso *et al.*, 2019]. Representation, planning, execution, and their smooth integration are essential factors for developing such systems. In this talk, we present three studies coping with these aspects and provide an overall picture.

Planning *Multi-Agent Path Finding (MAPF)* [Stern, 2019] is a planning problem where multiple agents on a graph are assigned a path to their respective goals without collisions. Powerful optimal solvers have been studied [Felner *et al.*, 2017]; however, finding optimal solutions is NP-hard [Yu and LaValle, 2013], thus, we must use sub-optimal solvers to cope with large-scale problems, e.g., with hundreds of agents. Furthermore, optimal solvers are typically inappropriate in online real-time settings where deliberation time is limited. In contrast, this is where *anytime planners* [Zilberstein, 1996] — i.e., get initial solutions quickly, improve them gradually, and provide valid plans whenever interrupted — are particularly attractive.

As the first topic, we introduce a framework of iterative refinement to realize *anytime MAPF*. The proposal uses a sub-optimal MAPF solver to obtain an initial plan quickly, then repeats the two steps: 1) select a subset of agents, 2) use an optimal MAPF solver to refine their paths while keeping other paths unchanged. This scheme rapidly yields “good” enough plans with high scalability because we can use any solver for the initial plan and the optimal solvers are used in small size problems. As a proof-of-concept, we present an algorithm called *IR* which tries to reduce the gap between ideal and actual costs of agents’ paths. Evaluations in various scenarios with MAPF benchmarks yield promising results in that convergence is fast, scalable, and with reasonable quality.

Execution An MAPF plan is bound to be executed in real-world situations with agents (robots). Typical MAPF is defined in discrete time. Agents are assumed to do two kinds of atomic actions synchronously: move to a neighboring location or stay at their current location. Perfect executions for the planning are however difficult to ensure because timing assumptions are inherently uncertain in reality. Even worse, the potential of unexpected interference increases with the number of agents, hence the need to prepare for imperfect executions regarding the timing assumptions.

So far, two policies enforce a robust execution of MAPF plans taken as input: either by forcing agents to synchronize or by executing plans while preserving temporal dependencies [Ma *et al.*, 2017]; however, considering an extreme example where one agent moves very slowly or crashes, it turns out that they are still vulnerable to delays.

As the second topic, we present an alternative approach, *time-independent planning* [Okumura *et al.*, 2020], which is both online and distributed. We represent reality as a transition system that changes configurations according to atomic actions of agents. In this time-independent model, any *a priori* knowledge for timings of atomic actions is unavailable, representing non-deterministic behaviors of the external environment. We propose an example of time-independent planning, *Causal-PIBT*, which extends a recently-developed decoupled approach to solve MAPF iteratively, called Priority Inheritance with Backtracking (PIBT) [Okumura *et al.*, 2019]. We also present how an offline MAPF plan enhances Causal-PIBT. Empirical results in a simulated environment with stochastic delays of agents’ moves, called MAPF-DP (with Delay Probabilities) [Ma *et al.*, 2017], support the validity of our proposal.

Representation, Integration Representation, how to model the world, has been the central issue for AI and robotics fields [Davis *et al.*, 1993]; however, there is room to consider whether

robots themselves should have a representation of their working environment. As famously argued by Brooks [Brooks, 1991], “the world is its own best model”.

Derivative concepts, direct use of the world as the representation, can be seen in navigation tasks. In general, robots rely on internal maps as the representation, however, the use of the internal maps entails a number of serious difficulties derived mostly from the correspondence of external physical objects and internal representation. Using the environment itself as the representation, realized by spatially and distributedly deploying sensors or tags in the environment, has been demonstrated as useful for navigation, e.g., guiding robots by sensor network [Verma *et al.*, 2005] or stigmergic approaches [Khaliq and Saffiotti, 2015].

As the final topic, we introduce a concept that offloads not only the representation but also the planning function to the environment. The rationales are 1) functional separation to focus robots in other tasks, and 2) response to a dynamic environment. As a proof-of-concept, we present *AFADA* and its prototyping; an architecture that consists of mobile robots that evolve over an active environment consisting of flat cells each equipped with a computing unit. Each cell can communicate with its direct neighboring cells and a robot on the cell. Using this local communication, the cells collectively manage environmental representation. The cells do the planning, i.e., navigating robots to their destinations. Robots just follow the instructions from the cell — neither representation nor planning is necessary in robots.

References

- [Brooks, 1991] Rodney A Brooks. Intelligence without representation. *Artificial intelligence*, 47(1-3):139–159, 1991.
- [Davis *et al.*, 1993] Randall Davis, Howard Shrobe, and Peter Szolovits. What is a knowledge representation? *AI magazine*, 14(1):17–17, 1993.
- [Dresner and Stone, 2008] Kurt Dresner and Peter Stone. A multiagent approach to autonomous intersection management. *Journal of artificial intelligence research*, 31:591–656, 2008.
- [Felner *et al.*, 2017] Ariel Felner, Roni Stern, Solomon Eyal Shimony, Eli Boyarski, Meir Goldenberg, Guni Sharon, Nathan Sturtevant, Glenn Wagner, and Pavel Surynek. Search-based optimal solvers for the multi-agent pathfinding problem: Summary and challenges. In *Tenth Annual Symposium on Combinatorial Search*, 2017.
- [Khaliq and Saffiotti, 2015] Ali Abdul Khaliq and Alessandro Saffiotti. Stigmergy at work: Planning and navigation for a service robot on an rfid floor. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1085–1092. IEEE, 2015.
- [Ma *et al.*, 2017] Hang Ma, TK Satish Kumar, and Sven Koenig. Multi-agent path finding with delay probabilities. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [Okoso *et al.*, 2019] Ayano Okoso, Keisuke Otaki, and Tomoki Nishi. Multi-agent path finding with priority for cooperative automated valet parking. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 2135–2140. IEEE, 2019.
- [Okumura *et al.*, 2019] Keisuke Okumura, Manao Machida, Xavier Défago, and Yasumasa Tamura. Priority inheritance with backtracking for iterative multi-agent path finding. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 535–542, 7 2019.
- [Okumura *et al.*, 2020] Keisuke Okumura, Yasumasa Tamura, and Xavier Défago. Time-independent planning for multiple moving agents. *arXiv preprint arXiv:2005.13187*, 2020.
- [Stern, 2019] Roni Stern. Multi-agent path finding—an overview. In *Artificial Intelligence*, pages 96–115. Springer, 2019.
- [Verma *et al.*, 2005] Atul Verma, Hemjit Sawant, and Jindong Tan. Selection and navigation of mobile sensor nodes using a sensor network. In *Third IEEE International Conference on Pervasive Computing and Communications*, pages 41–50. IEEE, 2005.
- [Wurman *et al.*, 2008] Peter R Wurman, Raffaello D’Andrea, and Mick Mountz. Coordinating hundreds of cooperative, autonomous vehicles in warehouses. *AI magazine*, 29(1):9–9, 2008.
- [Yu and LaValle, 2013] Jingjin Yu and Steven M LaValle. Structure and intractability of optimal multi-robot path planning on graphs. In *AAAI*, 2013.
- [Zilberstein, 1996] Shlomo Zilberstein. Using anytime algorithms in intelligent systems. *AI magazine*, 17(3):73–73, 1996.