Proceedings of the Seventeenth International Symposium on Combinatorial Search (SoCS 2024)

Multi-Agent Motion Planning with Bézier Curve Optimization under Kinodynamic Constraints (Extended Abstract)*

Jingtian Yan, Jiaoyang Li
Carnegie Mellon University
jingtianyan@cmu.edu, jiaoyangli@cmu.edu

Abstract

Multi-Agent Motion Planning (MAMP) is a problem that seeks collision-free dynamically-feasible trajectories for multiple moving agents in a known environment while minimizing their travel time. In this paper, we introduce a three-level planner called PSB that combines search-based and optimization-based techniques to address the challenges posed by MAMP. PSB fully considers the kinodynamic capability of the agents and produces solutions with smooth speed profiles. Empirically, we evaluate PSB within the domain of obstacle-rich grid map navigation for mobile robots. PSB shows up to 49.79% improvements in solution cost compared to existing methods while achieving significant improvement in scalability.

Introduction

Multi-Agent Motion Planning (MAMP) is a problem that focuses on finding collision-free dynamically-feasible trajectories for multiple agents in a known environment while minimizing their travel time. We define this problem by a graph $G = (V, E)$ and a set of M agents $\mathcal{A} = \{a_1, ..., a_M\}$. Agents can move from vertex $c_i \in V$ to $c_j \in V$ along edge $(c_i, c_j) \in E$. Each agent a_i initiates its movement from a specified *start (vertex)* $c_i^s \in V$ toward a designated *goal (vertex)* $c_i^g \in V$. We define the *path* of an agent as a sequence of vertices that leads it from start to goal. We define the *spatio-temporal profile* $\ell_i(t)$ as a function that quantifies the distance traversed by an agent over time along a given path. The spatio-temporal profile must satisfy the kinodynamic constraints of each agent, which limit the gradient of the spatio-temporal profile with respect to time:

$$\underline{U_i^k} \leq d^k \ell_i(t)/dt^k \leq \overline{U_i^k}, \forall k \in 1, ..., K \quad (1)$$

where $\underline{U_i^k}$ and $\overline{U_i^k}$ are constant values that define the constraints on the k-th order gradient. The *trajectory* is the combination of a path and its associated spatio-temporal profile. We use arrival time to indicate the time needed for an agent to arrive at its goal. Our task is to plan collision-free trajectories for all agents while minimizing the sum of their arrival time, where a collision happens if the time duration during which two agents occupy the same vertex overlap.

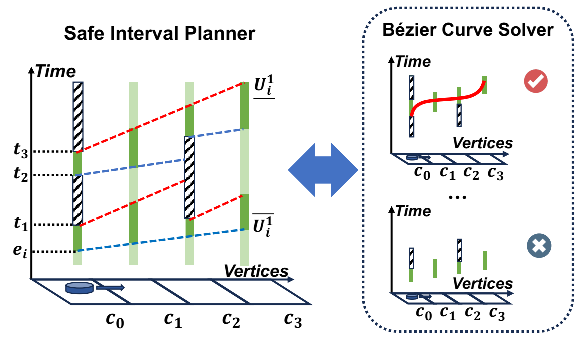

Figure 1: Example for a 1-D scenario. The shadowed strips denote time intervals occupied by other agents, the green segments denote safe intervals. The dark green segments represent the intervals in the open list.

In this paper, we introduce a three-level planner called PSB (**P**BS-**S**IPP-**B**ézier) to address the MAMP problem. Our method combines search-based techniques for fast pathfinding (with corresponding temporal constraints), and optimization-based techniques to generate spatio-temporal profiles exploiting the full kinodynamic capability of agents.

PBS-SIPP-Bézier (PSB)

PSB consists of a three-level planner. Level 1 uses Priority-Based Search (PBS) to resolve collisions among agents through priority ordering searching, where the trajectory of each agent is planned by Level 2 and Level 3. Level 2 uses an extended Safe Interval Path Planning (SIPP) to search for the optimal trajectory for each agent, where the spatio-temporal profile of the trajectory is optimized by Level 3. Given the path (together with temporal constraints) from Level 2, Level 3 uses BCP (**B**ézier-**C**urve-based **P**lanner) to generate the optimal spatio-temporal profile.

Priority-Based Search (PBS)

PBS resolves the collisions among agents by iteratively exploring various agent priority sequences through a search tree. The lower-priority agents need to avoid collision with

*This paper is a short version of (Yan and Li 2024).

higher-priority agents. Each node in the tree represents a different set of priorities, and the search progresses by updating these priorities when a collision happens. For example, if a collision between agents a_i and a_j is detected in a node, we resolve it by expanding it into two child nodes. In one node, a_i is given higher priority than a_j, while in the other node, the priority is reversed.

Safe Interval Path Planner

Level 2 aims to find a trajectory for a given agent that minimizes its arrival time while satisfying the constraints from Level 1. To achieve this, we run an extended SIPP on the safe interval graph generated based on those constraints. This graph associates each vertex with a set of safe intervals, where *safe interval* $[lb, ub)$ is a time duration that the agent can stay at a vertex without colliding with higher-priority agents. The extended SIPP plans a path (along with safe intervals) and calls Level 3 to specify the spatio-temporal profile considering kinodynamic constraints.

We provide a simplified 1-D example of this process in Fig. 1. Since there are two safe intervals $[e_i, t_1)$ and $[t_2, t_3)$ at the start c_0, we generate two SIPP nodes and insert them into the open list. In each iteration, we expand the node from the open list with the smallest f-value (= the lower bound of its safe interval plus the minimum arrival time required to reach the goal from its vertex). In this example, we choose the node with safe interval $[e_i, t_1)$. During node expansion, to speed up the search process, we use relaxed kinodynamic constraints. Since the precise kinodynamic constraints are enforced at Level 3, this relaxation still guarantees completeness. Specifically, we assume that the agent occupies each vertex for only an instant of time while allowing adjustment of its speed between its minimum and maximum speed instantly. In our case, as the agent moves from c_i^0 to c_i^1, a new interval $[e_i+t_{min}, t_1+t_{max})$ is generated at c_i^1, where t_{min} and t_{max} are the time it takes for the agent to finish this movement with its maximum and minimum speed. We iterate the safe intervals at c_1 that overlap with this new interval and insert them into the open list. When a node reaches the goal, we backtrack to retrieve the full path and its associated safe intervals. Then, we call Level 3 to determine the kinodynamically feasible trajectory. If Level 3 finds a trajectory with a smaller arrival time than the best trajectory found so far, we update the current best trajectory. We terminate the search if no node in the open list has a smaller f-value than the arrival time of the current best trajectory.

Bézier-curve-based Planner (BCP)

Level 3 aims to generate a kinodynamically feasible spatio-temporal profile for a given path within given safe intervals while minimizing the arrival time. We use Bézier curve $B^T(t)$, where T is the *arrival time*, to represent the spatio-temporal profile. A Bézier curve is a function parameterized by a set of control points and scale factor T. With a sufficiently large number of control points, it is able to approximate any continuous function $f(t)$ with $t \in [0, T]$. BCP formulates finding the optimal spatio-temporal profile by finding the minimum T and control points P for $B^T(t)$ that satisfy the kinodynamic and temporal constraints. BCP uses a Linear Programming (LP) problem to encode these constraints, then employs binary search to determine the optimal travel time.

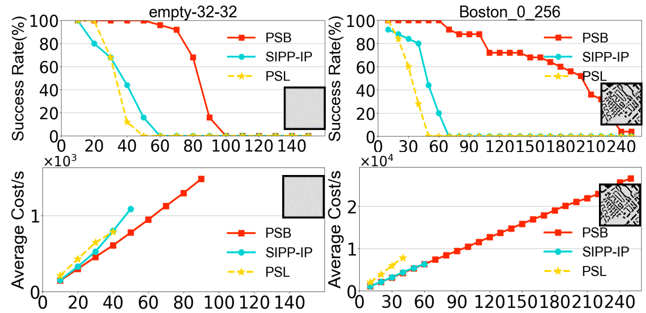

Figure 2: Success rate and solution cost of PSB, PSL, and SIPP-IP on 32 × 32 empty map (left) and a 256 × 256 city map (right). The runtime limit is 5 minutes for each instance.

Experimental Results

We compare PSB with PSL (Li et al. 2023) and SIPP-IP (Ali and Yakovlev 2023) on a 32×32 and a 256×256 map. Both PSL and SIPP-IP use PBS for high-level search. At the low level, PSL assumes agents move at a constant speed, utilizing a MILP model to determine the speed profile. Meanwhile, SIPP-IP discretizes the action space and uses motion primitives for trajectory search. All agents have identical kinodynamic constraints, where the speed is bounded by the range of $[0, 2]$ $grid/s$, while the acceleration is confined to $[-0.5, 0.5]$ $grid/s^2$. Fig. 2 presents the *success rate* and *solution quality* of all three methods. We evaluate the solution quality using the sum of the arrival time of all agents. PSB shows better solution quality and success rate than both PSL and SIPP-IP. We also evaluate the performance of PSB in the traffic intersection coordination domain. Please refer to (Yan and Li 2024) for more details.

Acknowledgements

The research is supported by the National Science Foundation under grant number 2328671 and the CMU Manufacturing Futures Institute, made possible by the Richard King Mellon Foundation.

References

Ali, Z. A.; and Yakovlev, K. 2023. Safe Interval Path Planning with Kinodynamic Constraints. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 12330–12337.

Li, J.; Hoang, T. A.; Lin, E.; Vu, H. L.; and Koenig, S. 2023. Intersection Coordination with Priority-Based Search for Autonomous Vehicles. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 11578–11585.

Yan, J.; and Li, J. 2024. Multi-Agent Motion Planning with Bézier Curve Optimization under Kinodynamic Constraints. *IEEE Robotics and Automation Letters*, 9(3): 3021–3028.