

Sampling-Based Multi-Modal Multi-Robot Multi-Goal Path Planning

Valentin N. Hartmann*, Tirza Heinle*, Yijiang Huang, Stelian Coros

Computational Robotics Lab
ETH Zürich, Zürich, CH

Abstract. In many robotics applications, multiple robots are working in a shared workspace to complete a set of tasks as fast as possible. Such settings can be treated as multi-modal multi-robot multi-goal path planning problems, where each robot has to reach *a set* of goals. Existing approaches to this type of problem often use some type of decomposition, and many are neither optimal nor complete. Given the lack of optimal solvers for this problem, it is unknown how big the optimality gap of current solvers is.

We formalize this problem as a centralized path planning problem and present planners that are probabilistically complete and asymptotically optimal. The planners are modifications of standard sampling-based planners with the required changes to work efficiently in the multi-modal, multi-robot, multi-goal setting. We validate the planners on a diverse range of problems including scenarios with various robots, planning horizons, and collaborative tasks such as handovers, and compare the planners against a suboptimal prioritized planner. With the optimal planners, we answer the question of how close to optimal the suboptimal prioritized planners are. Videos and code for the planners and the benchmark is available at <https://vhartmann.com/mrmg-planning/>.

1 Introduction

As adoption of robots increases and tasks become more and more automated, it will be increasingly important to deploy solutions that are not only able to work longer and cheaper but are also competitive in throughput with humans. In order to achieve this, in many cases, multiple robots need to be used and effectively coordinated in the same workspace: Enabling motion planning for multiple tasks with multiple robots is crucial in order to maximize the usefulness of robots in industrial settings. While workcells with multiple robots exist, the robots typically act independently from each other in order to simplify the programming and avoid dealing with robot-robot interactions.

Most work in continuous multi-robot planning is focusing on single-goal settings where all robots start moving at the same time and reach their respective goals simultaneously [1–3]. Conversely, in most real use cases, robots need to do multiple tasks in sequence, e.g., welding multiple points, or picking and placing multiple things after another in order to sort objects. Even when only considering a single pick and place task per robot, each robot needs to reach two goals: The pick, and the place location. Since

* equal contribution




Fig. 1: Examples of problems with their initial state (left) and goal state (right): (top) two robots cooperating to reorient boxes, (bottom) four robots stack boxes on top of each other.

the robots act in the same environment in these scenarios and thus possibly block each other from doing their tasks, we can not formulate the problem at hand as a *sequence of path planning problems*, but need to solve the multi-robot multi-goal planning problem if we want to find an optimal solution.

Multi-goal planning can also be framed as multi-modal planning: Multi-modal path planning [4–7] finds paths through sequences of *modes*, i.e., through variations of a continuous configuration space that occur, e.g., by grasping an object, or moving an object in the workspace. This means that in each *mode* a different set of constraints is active. Most work on multi-modal path planning for robots only considers single-robot settings.

For the multi-robot, multi-goal planning problem we consider, only suboptimal planners exist [8, 9]. In this paper, we propose complete and optimal planners for the multi-agent, multi-goal, multi-modal planning problem, and with the help of these planners both provide optimal baselines for the problem, and answer the question ‘*How close to optimal are the suboptimal planners?*’. Summarizing, our contributions in this paper are

- a formalization of multi-modal, multi-robot, multi-goal path planning in continuous spaces,
- adaptations of probabilistically complete and almost-surely asymptotically optimal planners to this setting.

We compare the planners to a suboptimal, incomplete, prioritized planner on a multi-modal, multi-robot, multi-goal motion planning benchmark containing various base-scenarios that can be further adjusted in difficulty by changing the number of robots and tasks.

Table 1: Comparison of path planning approaches. Ticks in parentheses mean that the approach is only optimal/complete under some assumptions that do not always hold.

Planning Space	Approach	Examples	Multi-Robot	Task Structure	Multi-Modal	Complete	Opt.
Discrete	MAPF	[10, 11]	✓	Single goal		✓	✓
	MAPD	[12, 13]	✓	Goal sequence		✓	✓
Continuous	Multi-Modal Planning	[4, 5]		General task graphs	✓	✓	✓
	MRMP	[14–17]	✓	Single goal	✓	✓	
	Prioritized Planning	[8, 18]	✓	Goal sequence	✓	(✓)	
	Hybrid Methods	[19]	✓	Unordered goals	✓	✓	(✓)
	Ours		✓	General task graphs	✓	✓	✓

2 Related Work

Most multi-robot planning work is focusing either on the setting where all robots share the same work and configuration space, or on solving single goal problems. It is commonly assumed that the environment stays completely static during planning, therefore excluding manipulation planning settings. We group the different research areas and their respective focus in Table 1. In the following, we discuss single-goal multi-robot planning (MRMP) and multi-agent path finding (MAPF) more in-depth.

2.1 Continuous Multi-Robot Planning

A simple approach to multi-robot planning is planning in the composite space of all robots [16], and applying standard methods such as RRT(*) [20] or PRM(*) [21] in this higher dimensional space. This approach is not scalable to many robots or high degrees of freedom.

There are various approaches to tackle this limitation. Discrete RRT (dRRT*) [15, 22]), builds a roadmap for each robot, and combines them to implicitly form a tensor graph. Instead of planning naively on the implicit graph, an RRT-like strategy is proposed to explore this implicit graph. Fast-dRRT [23] applies this approach sequentially on the subproblems to solve multi-goal problems. M* [1] plans in separate spaces, and plans jointly only where required, i.e., when conflicts between the single-robot plans arise. Similarly, conflict based search (CBS) [24] plans for each robot individually, and introduces constraints to solve conflicts between single agent plans. CBS originates from the MAPF-setting, but can also be applied to continuous multi-robot planning problems [14, 25, 26].

Prioritized planners as proposed in, e.g., [8, 17, 18, 27] achieve more scalable solvers by not requiring completeness and optimality. Similarly, [28] proposes a scalable approach that is based on operator splitting and sampling, and while the approach is complete, it is not optimal.

The prioritized planners can be used to plan for multi-goal problems by incrementally generating a path per agent, and considering it as fixed for the later planning problems as done in [8, 18]. To the best of the authors’ knowledge, there are no multi-modal multi-robot multi-goal planners that are complete and optimal in the general setting.

A slightly different approach is taken in [29], where plans for the separate robots are planned, and then - if conflicts between the paths arise - fixed by introducing pauses.

Compared to the discussed works, we are interested in optimally solving multi-robot multi-goal planning problems without any restricting assumptions. Further, we are interested in manipulation, i.e., problems where the environment changes through the actions of the robot.

2.2 Multi-Agent Path Finding

MAPF [10] generally refers to planning in grid-like 2D environments with disk-robots, as found in warehouses. A common approach to solve this problem is conflict based [24] or priority based search (PBS) [30]. Many works improve upon the initial approach, such as improved CBS [31], or bounded suboptimal CBS [32], scaling to 1000s of agents.

Multi-agent pickup and delivery (MAPD) refers to the setting where we do not only intend to go from point A to point B, but traveling between multiple points of interest, as part of a single planning problem [12, 33]. While MAPD can be framed as lifelong MAPF problem, i.e., once an agent has finished its task of the MAPF problem, a new goal is assigned to the robot, these solvers are suboptimal [34]. Multi-label A* [13] extends the classic A* algorithm to a sequence of multiple goals, and is able to optimally solve MAPD problems. Both MAPF and MAPD planners typically assume that the environment remains unchanged over the planning duration.

3 Multi-Robot Multi-Goal Path Planning

Informally, our objective is to find a collision free path for each robot that passes through a sequence of goals that minimizes a cost (e.g., the latest completion time of all robots). Goals might involve multiple robots: A handover of an object between two robots implies a constraint on two robots. A goal could additionally imply a *mode transition*, i.e., a robot grasping something and thus changing the environment for the remaining planning problem.

3.1 Preliminaries

To formalize the problem, we first introduce the notion of tasks, the concept of a mode, and the space that we plan in. We follow the work from Thomason et al. [4], and generalize it to multiple robots. Compared to [4], task sequencing and allocation happens implicitly in the planning. Refer to Fig. 2 for an informal schematic of what tasks are, and how the task assignment changes over time.

Task A task consists of the robots that are assigned to the task, and a set of goal-constraints g that need to be fulfilled to consider the task done¹. Typically, the constraint

¹ Additionally, a task could include path constraints h during a task, such as, e.g., an orientation constraint on the end-effector of a robot. This requires constrained planning, which we leave for future work.




Fig. 2: An illustration for what tasks are, and how robots move through task assignments and space together. The robots start with the task assignment $[1, 2]^T$ and proceed to the terminal mode. The goals of the tasks are shown with lower opacity and dashed lines. The vertical red dashed lines indicate points where a task is completed, and the task assignment changes. Task t indicates the terminal task, that both robots need to reach jointly.

g is a single goal pose or a goal region, but g could also be a more complex nonlinear constraint such as a grasp constraint. We do not assume anything on how g is defined except that we can sample constraint-satisfying configurations (e.g., using projection methods). A task can have post-conditions that can alter the scene-graph of the environment, i.e., which objects are linked to each other. As example, consider the task ‘robot r_1 grasps object o_1 ’: Here, the goal constraints are that the robot is grasping o_1 , and the post-condition is that o_1 is linked to the end-effector of the robot.

We use $s \in \mathcal{S} = \mathcal{S}_{r_1} \times \dots \times \mathcal{S}_{r_n}$ to denote the *task assignment* of all robots, where \mathcal{S}_{r_i} is the set of tasks that can be assigned to robot r_i .

Modes The constraints of a task can be fulfilled by a set of poses, where the chosen pose might affect the environment that we plan in due to a tasks post-condition. Consider again the grasping-task: The robot is able to fulfill a grasping constraint by grasping from any side, which changes the collision geometry that we need to consider in the rest of the planning problem. We use *mode* $m \in \mathcal{M} = \mathcal{S} \times \mathcal{Q}_o$ to refer to the combination of the discrete task assignment s (which implies a scene-graph, i.e., which objects are linked to which parents), and the poses of all movable objects $q_o \in \mathcal{Q}_o$, defined by the relative transformation to their parent-frames. Importantly, defining q_o as relative transformation to the parent frame means that a mode *only changes* when the scene-graph changes, i.e., upon completion of a task that has post-conditions that change the scene-graph.

Task Order Specification A task order is induced by an oracle $\mathcal{O}(m) \subseteq \mathcal{S}$ which, given the current mode, returns a set of *all* possible successor task assignments. Note that the task assignment (and thus the mode) only changes when the goal constraints of a previous task are fulfilled. Some examples of how a task order can be specified are




Fig. 3: Examples for the different ways in which a task ordering can be specified: **(left)** As a sequence of tasks **(middle)** as a partially ordered sequence and **(right)** as a set of unassigned and unordered tasks (i.e., the order of task groups 1 and 2 is exchangeable).




Fig. 4: Example of a mode-graph $\mathcal{G}(\mathcal{O})$, right) that is implied by the ordering on the left. Modes with full opacity show an example of how a set of modes can be incrementally reached. Each mode can be reached with different parameterizations θ_i^k (e.g., different grasps), as indicated for mode (4, 2). Fig. 2 is one possible realization for this problem specification.

illustrated in Fig. 3. The oracle induces a graph $\mathcal{G}(\mathcal{O})$ of task assignments that we can traverse to go from the start to the end-mode (Fig. 4, right).

Configuration Space We describe a path in the composite space of all robots, all objects, and which tasks are currently assigned to each robot:

$$\mathcal{Q} = \underbrace{\mathcal{Q}_{r_1} \times \cdots \times \mathcal{Q}_{r_N}}_{\mathcal{Q}_R} \times \underbrace{\mathcal{S}_{r_1} \times \cdots \times \mathcal{S}_{r_N}}_S \times \mathcal{Q}_o. \quad (1)$$

Here, \mathcal{Q}_{r_i} is the configuration space of robot r_i , and correspondingly, \mathcal{Q}_R is the composite configuration space of all robots; \mathcal{Q}_o is the composite configuration space of all objects. We use $\mathcal{Q}_{\text{free}} \subseteq \mathcal{Q}$ to denote the part of the configuration space that is collision-free.

Clearly, not all degrees of freedom are actuated. We assume that we can only plan for the robots' degrees of freedom directly, and all others need to be influenced indirectly. Particularly, the object configurations and the task assignment can only be changed through the completion of tasks (i.e., fulfilling the goal constraints g of a task).

3.2 Problem formulation

A multi-modal, multi-robot, multi-goal path planning problem is then given by the tuple $(R, \mathcal{Q}_R, q_{\text{start}}, m_{\text{start}}, \mathcal{O})$, where R is the set of robots, \mathcal{Q}_R is the configuration space of the robots, q_{start} is the initial configuration, m_{start} is the start mode, and \mathcal{O} is the oracle giving us the possible next task assignments from the current mode. In the following, we will use q^{r_i} for the pose of robot r_i .

We want to find a collision free path $\pi(t) : \mathbb{R} \rightarrow \mathcal{Q}_R$ and the task assignment sequence $s(t) : \mathbb{R} \rightarrow \mathcal{S}$, that minimizes the cost function $c(\cdot)$. The path π maps time to the composite robot configuration \mathcal{Q}_R , and the mode sequence $s(t)$ maps time to the task assignment \mathcal{S} , which together imply the scene-graph at a time, and thus the poses of all objects.

Cost functions We are often interested in finding the minimum makespan plan. However, purely minimizing makespan can lead to optimal paths that bring undesired side-effects, e.g., containing unnecessary movement, due to the large nullspace in the makespan cost. Thus, we consider cost functions of the family

$$c(q_1, q_2) = (1-w) \max_r \|q_1^r - q_2^r\|_2 + w \sum_r \|q_1^r - q_2^r\|_2. \quad (2)$$

where if w is small, we get a minimum makespan optimization problem (which is ‘regularized’ by the path length), and if w is 1 the total cost is the sum of all robot path lengths.

4 Planners

We now introduce the necessary adaptations to four sampling-based planners (RRT*, PRM* [35], AIT* and EIT* [36]). We refer to the original works for a detailed explanation of how the algorithms work. Generally, the planners proceed as their standard version, with an adjusted distance function that accounts for the connectivity between the modes. Before discussing any specific adjustments required in the planners in Section 4.1, we discuss mode sampling, the distance function used in the planners, and the concept of the informed set [37].

Mode sampling Instead of sampling the (unactuated) object poses and task assignments, we sample $x = (q, m)$, where $q \in \mathcal{Q}_R$ is a configuration containing all robots’ configurations and m is a mode that determines the task assignment and all object poses. To sample the modes, we do not sample the continuous space $\mathcal{S} \times \mathcal{Q}_O$, but we sample from the list of previously reached modes. To enable that, we store the list of modes M that we reached so far, and expand it whenever we reach a new mode (i.e., whenever we reach a configuration that fulfills the constraints g of a task). To sample m from the set M , we maintain the most recently added mode F_{-1} and the set of modes F that we did not yet expand (the frontier). We then sample according to:

$$m \sim \begin{cases} \mathcal{U}(F) & \text{with probability } p(1 - \epsilon) \\ \mathcal{U}(F_{-1}) & \text{with probability } p\epsilon \\ \mathcal{U}(M \setminus F) & \text{with probability } 1 - p \end{cases}$$

where $p \in [0, 1]$ and $\epsilon \in [0, 1]$ are weighting parameters, and $\mathcal{U}(\cdot)$ denotes a uniform distribution over a finite set. If $p = 1$ and $\epsilon = 1$, we reach a fully *greedy* version, resulting in the most direct progression possible towards the terminal mode.

We refer to the *frontier sampling strategy* if $p > 0.5$ holds. This biases the planner towards the unexpanded modes (i.e., frontiers) compared to the fully greedy approach, while still maintaining a greedy progression toward the terminal state².

Distance function We assume that any node in a different mode is inaccessible except through the transition nodes:

$$d(x_1, x_2) = \begin{cases} d_{\text{mode}}(q_1, q_2) & \text{if } m_1 = m_2, \\ 0 & \text{if } q_1 = q_2 \text{ and } s(m_2) \in \mathcal{O}(m_1) \text{ and } g(q_1, m_1) = 0 \\ \infty & \text{otherwise,} \end{cases}$$

where we use $s(m)$ to denote the task assignment of mode m . For the in-mode distance any metric can be chosen; the work in [38] discusses distance metrics in multi-robot problems in more depth. We use the maximum of the per-robot L_2 distance, as this performed best in our experiments.

Locally informed sampling Once we found a solution, the informed set [37] can be used to restrict the points that we need to consider for planning to the points that can improve the current solution. The informed set in our problem is

$$\mathcal{Q}_{\text{inf}} = \{q \mid c_{\text{lb}}(q_{\text{start}}, q) + c_{\text{lb}}(q, Q_{\text{goal}}) \leq c_{\text{best}}\}, \quad (3)$$

with Q_{goal} denoting the set of possible goal configurations.

In our setting, the optimal paths are long compared to the size of the configuration space. Therefore, the informed set can be larger than the valid configuration space, and thus not helpful.

Given that we do know that a solution needs to pass through some regions in space (the goals that we need to reach), the standard approach of taking the cost Eq. (2) as c_{lb} is not approximating the true path cost well across modes. It is not easy to construct a good lower bound for the path-cost, since the regions that we need to pass through are *subspaces of the full space*: even though we know that robot r_i needs to be at pose q_{g_1} first, robot r_j might be unconstrained, thus resulting in a volume that we need to pass through instead of a single point. Additionally, compared to the Euclidean norm as cost, our cost function does not easily allow for direct sampling from the informed set.

To alleviate these issues, we propose *locally informed sampling*: instead of using the start and the goal in Eq. (3), we randomly sample two points q_i, q_j on the previously found path, and do informed sampling with these points:

$$\mathcal{Q}_{\text{lis}} = \{q \mid c(q_i, q) + c(q, q_j) \leq c^\pi(q_i, q_j)\},$$

where $c^\pi(\cdot)$ is the cost to reach a point on the path along the best found path. We extended this to the sampling of modes by sampling from the set of modes that are in-between the modes that q_i, q_j are in respectively, i.e., the modes in $\mathcal{G}(\mathcal{O})$ that are reachable from m_i and can reach m_j . Locally informed sampling does not influence optimality or completeness guarantees, since we sample the indices randomly, and thus also sample from the standard informed set.

² A fully greedy strategy is not probabilistically complete, as the planners might get stuck in a dead-end.

4.1 Considerations for Multi-Modal, Multi-Goal Planners

RRT* Compared to standard RRT* [20, 35], we maintain a set of previously reached modes M (i.e., a mode-graph similar to Fig. 4), which is initialized with the start mode. In order to sample configurations in the planner, we uniformly sample poses within the limits of the robots, and sample modes from this set. If we previously found a path, we use locally informed sampling to sample the mode and the pose. When using a goal bias, we bias the sampling in the RRT to the *goal in a mode* (i.e., to a configuration satisfying the goal constraints g of a task from the current task assignment).

We also use a bidirectional variant (BiRRT*), which first samples goals in each mode, which are used as starts in the next mode, and then runs bidirectional RRT* in each mode.

PRM* The main change in the PRM* planner is that the roadmap is constructed by first sampling a batch of transition nodes (i.e., nodes that fulfill the constraints of a task), which enlarges the set of reached modes M and then sampling a batch of nodes in the reached modes, until the terminal mode is reached. Then, a standard A* with an edge queue is used for the search over the multi-modal roadmap.

Heuristics Since A* needs a heuristic in order to be effective, and the Euclidean distance is not informative in our setting, we compute a heuristic for the A* search by running a reverse search on the transition nodes to compute a lower bound on the cost to reach a mode. The heuristic is then the sum of the cost to reach a specific transition and the cost to reach the goal from this transition mode:

$$h(x = (q, m)) = \min_i [c(q, q_i^{m, \text{trans}}) + c_{\text{rev}}(q_i^{m, \text{trans}})] \quad (4)$$

where $q^{m, \text{trans}}$ are the poses belonging to the transition nodes in a mode. This heuristic is always underestimating the true cost, and is thus admissible, and since c and c_{rev} are consistent, h is also consistent.

Informed Tree-planners AIT* and EIT* [36, 39] use adaptive problem-specific heuristics to focus the search on promising regions. These heuristics are computed and refined through an *asymmetric bidirectional search* to obtain more informative exploration and improved solution quality.

Search To explore the graph, AIT* and EIT* perform their respective *asymmetric bidirectional search* strategies, as described in [36]. Although each planner utilizes distinct heuristics, both follow an identical approach to handle transitions: When the forward search reaches a transition node, it only expands to valid successor modes, enforcing forward-directed connectivity; the reverse search applies the same logic in opposite direction (i.e., expanding towards previous modes). Both planners need an admissible a priori cost-to-come (i.e., from the start to the state x) heuristic - we use a modified version of Eq. (4) with the lower bound cost to reach a mode computed from the start




Fig. 5: The initial states of a selection of problems that are available in the benchmark. The poses that have to be reached by the robots or objects are drawn with lower opacity or indicated with a marker.

node instead of the goal (i.e., replacing c_{rev} with the cost computed by running a forward search only over the transition nodes of the previous modes). This heuristic is a trade-off between heuristic accuracy and computational efficiency.

The reverse search in EIT* incorporates heuristics approximating the computational effort required to find a valid path. We define the computational effort as the collision checking effort. Following this definition, the effort-to-come is computed analogously to the cost-to-come heuristic in Eq. (4), where cost terms are substituted by effort.

4.2 Postprocessing

When a path is found in any of the planners, we run a postprocessing step to rapidly improve the path cost, which decreases the volume of the informed set. Our approach is a minor modification of partial shortcutting [40]: The main idea of partial shortcutting is that we do not shortcut all dimensions of the path at once, but select a subset of dimensions from the composite space, and try to shortcut only this subset, while keeping the rest of the plan the same. Instead of uniformly sampling from all possible subsets, in practice it is more efficient to choose a robot, and then shortcut all dimensions of this robot at once. Crucially, compared to shortcutting all dimensions at once, this enables improving the unconstrained parts of the path (i.e., the parts that do not need to satisfy goal constraints g), while still keeping the simplicity of the shortcutting approach.

4.3 Proof sketches

All planners sample uniformly within the modes, and uniformly sample the transitions between modes, therefore gradually expanding the set of sampleable modes, and thus eventually discovering all possible modes (since we assume that the oracle \mathcal{O} gives us *all* possible next task assignments, and thus modes). Within the modes, all planners work as their respective base versions, and are therefore probabilistically complete, and will eventually connect to the transitions. If a path exists, the discovered transitions will lead to the terminal mode, which contains the goal pose. Thus, all planners will find a path if one exists and are therefore probabilistically complete.

The proof sketch for almost-sure asymptotic optimality follows a similar logic: The search that we run in the PRM* and EIT*/AIT*-planners is optimal, and therefore, as the number of uniform-random samples approaches infinity, we converge to the optimal

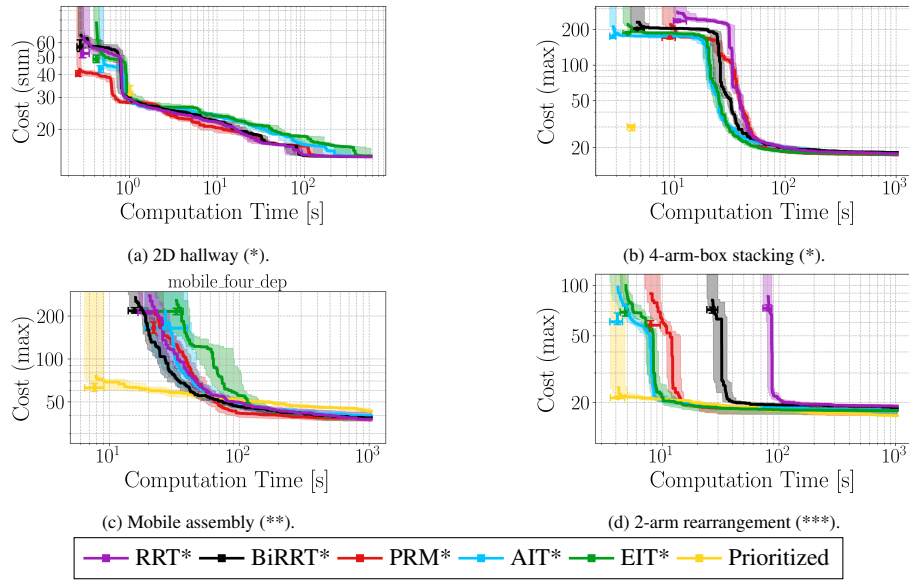


Fig. 6: Evolution of median cost over time along with the 95% non-parametric confidence intervals over 50 runs. We also show the median initial solution time and cost using the square with error bars. (*) indicates a fully given task sequence, (**) indicates a partial task ordering, (***) indicates unassigned and unordered tasks. Note that the prioritized planner only acts as anytime planner and improves its cost in the (**) and (***) settings, where multiple different sequences can be generated and planned for.

paths within the modes, and to the optimal mode-transitions. Similarly, the rewiring-step in our RRT*-planners is the same as in RRT*, and thus inherits the optimality guarantees from RRT* even over multiple modes.

5 Experiments

The main contribution of this work is the formalization of the multi-modal, multi-robot, multi-goal planning problem and the adaptation of standard sampling-based planners to the formulation. We provide implementations of the planners in Python, while the computationally expensive parts, i.e., collision checking and forward kinematics use a more performant backend, which can be easily replaced, allowing for, e.g., GPU parallelization if the backend supports it. In addition to implementations of the planners, we provide a wide range of benchmarks for the multi-modal, multi-robot, multi-goal planning problem. The scenarios in our benchmark range from simple settings where the optimal solution path is known to help validate properties of the planner, to more complex scenarios with up to 5 robots (with a total of 30 degrees of freedom). For many of the problems, there are versions with a complete task ordering and assignment, and others that are either unordered, partially ordered, or where tasks are unassigned. We show a selection of scenarios in Fig. 5.






Fig. 7: The path at different iterations in the process, from the initial to the optimized solution. Color progression indicates time from start to end (rectangle , disk ).

5.1 Experiments

We present a selection of experiments to showcase the different types of scenarios and task sequence specifications and compare the optimal planners to a suboptimal prioritized planner [9] on them³. We show four scenarios: (a) A 2D scenario with two robots that is similar to the classic wall-gap, where the robots have to switch positions and go back to their start poses (6D conf. space, 3 subgoals, Fig. 5a), (b) a scenario with 4 robot arms where 8 boxes have to be stacked (24D conf. space, 17 subgoals, Fig. 1), (c) a scenario involving 4 mobile manipulators rearranging a wall (24D conf. space, 17 subgoals, Fig. 5f), and (d) a scenario with 2 robot arms that have to collaborate with handovers to rearrange and reorient 9 boxes (12D conf. space, 28 subgoals, Fig. 5e). In these scenarios, the task sequence is fully determined for scenarios (a) and (b), *partially given* for scenario (c) and *not assigned or ordered* for scenario (d). To deal with partially ordered problems in the prioritized planner, we generate valid random sequences, and plan for those until the time runs out.

We use the cost function with $w = 1$ (i.e., a path-length cost) in scenario (a), and $w = 0.01$ (i.e., a min-makespan cost with a small path-length regularization) in scenarios (b)-(d). We report the median solution costs over 50 runs with different random seeds. We use a frontier mode sampling strategy ($p = 0.98$, $\epsilon = 0$) until a path is found, and then switch to uniform mode sampling in all planners if not stated otherwise.

The experiments were run with Python 3.10 on a Ryzen 7 5800X (8-core, 4'491 MHz) and 32 GB RAM.

5.2 Results

Figure 6 shows the resulting cost evolution plots. Figure 7 shows the evolution of a path for the hallway problem. Table 2 summarizes a more extensive evaluation of selected planners on a wider range of scenarios. First, we want to point out that the prioritized planner does not find solutions on some problems, highlighting again that the prioritized planner is not complete for some problem-settings.

Comparing the optimal planners introduced in this work and the prioritized planner shows that the suboptimal planner with our addition of postprocessing the path finds

³ Synchronized planners and standard MAPF planners are not applicable to these problems, and are thus not part of the benchmark.




Fig. 8: The optimal paths when using a max (**left**) or sum (**right**) cost function in the 2D hallway scenario, where the robots have to reach a goal on the other side and return. Color indicates time from purple (start) to yellow (end).

reasonable solutions: the cost is between 20% and 70% over the optimal solution, but better than the initial solutions of the optimal planners. The optimal planners generally converge to solutions with lower cost, but do so at a later time. Surprisingly, the planners that plan in composite space find initial solutions faster than the suboptimal planner in many settings. This is likely because (1) finding a *feasible* path is relatively simple in composite space in many scenarios, where the robots can simply be ‘moved out of the way’ and (2) fixing previously planned paths can restrict the feasible space too much.

Influence of the cost function We illustrate the difference of the sum-cost and the max-cost on the hallway example in Fig. 8: Both robots go through the wall gap twice when using the sum-cost, and do not make use of the passage at the top, compared to the (regularized) max-cost, where they make use of the passage at the top.

Convergence in complex problems We can see in Table 2 that the optimal planners do not converge to the same cost for some of the problems in the allocated time. This highlights that BiRRT* has a slower convergence rate compared to AIT*. We did run a subset of the scenarios for longer in order to confirm that they do indeed converge to the same costs.

In scenarios (c) and (d) and the unordered and unassigned scenarios in Table 2, the planners do not reliably converge to the same solution at the end of their runtime but remain 5%-10% above the lowest found solution, showing that some of the found solutions are not optimal. This is because the problem space is extremely large due to the discrete choice of task assignment and ordering in addition to the high dimensionality of the continuous space. Thus, while the planners are theoretically optimal, in practice, we find that the planners tend to get stuck in a local optimum and do not switch the discrete assignment and ordering often. Note that in scenario (d), the median final cost of the prioritized planner is lower than the median final cost of the composite space planners. This is likely due to the better exploration of different sequences.

5.3 Ablations

We present ablations to explain the design decisions that deviate from the standard versions of the planners. We want to emphasize again that while the core ideas of the

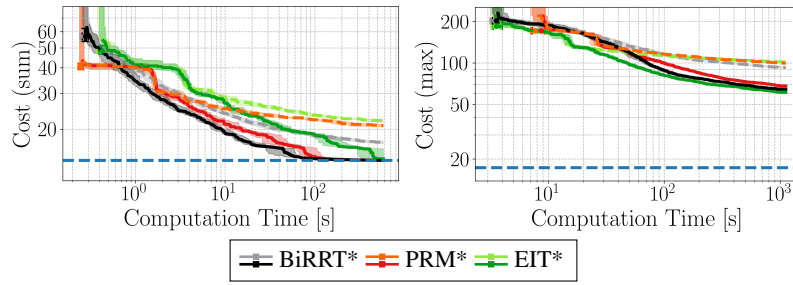


Fig. 9: Local (solid) and global (dashed) informed sampling on the 2D hallway (**left**) and the 4-arm box-stacking (**right**) scenarios. The dashed blue line is the best cost of the non-ablated planners.

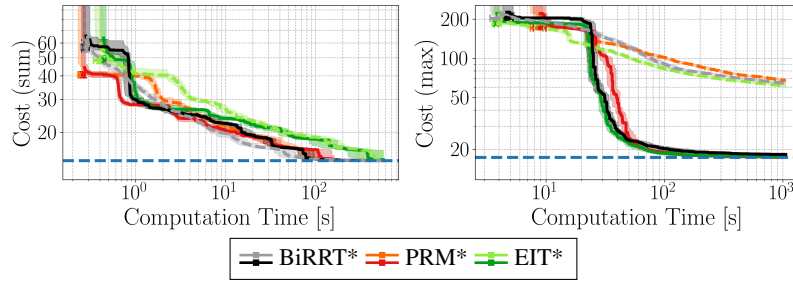


Fig. 10: Planners with (solid) and without (dashed) shortcutting on the 2D hallway (**left**) and the 4-arm box-stacking (**right**) scenarios. All planners are using locally informed sampling. The dashed blue line is the best cost of the non-ablated planners.

planners that we use here are the same as in their original versions, the adaptations are crucial in order to obtain the planning times and costs that we show here. In the following, we do not show the results for all planners as the effects are similar for most, and would crowd the figures.

Locally informed sampling Only using standard informed sampling in the problems we consider here is insufficient, as the informed set is large compared to the configuration space, and thus effectively equivalent to uniform sampling. In Fig. 9, the planners do not use shortcutting in order to better highlight the difference between the ablations. The planners using locally informed sampling outperform the planners that do use standard informed sampling.

Shortcutting in the planner Shortcutting a path after it was found leads to faster convergence than using only informed sampling (Fig. 10) while not affecting the optimality or completeness properties of the planners. This effect is more pronounced in high-dimensional scenarios.

Mode sampling strategy Employing a *frontier sampling strategy* during the initial search significantly accelerates the discovery of a first solution compared to uniform




Fig. 11: Frontier (solid) and uniform (dashed) initial mode sampling strategies on the 2D hallway (**left**) and 4-arm box-stacking (**right**) scenarios. The dashed blue line is the best cost of the non-ablated planners.

sampling (Fig. 11) in high dimensional, long horizon scenarios. While not shown here, a greedy sampling strategy gives a further significant speedup for the initial solution but does sacrifice completeness if not done carefully.

6 Discussion & Limitations

Due to planning in the composite space and treating the problem as a single big planning problem, the planners are not expected to scale well with the number of robots (due to the higher dimensionality) or with the number of goals, as can be seen in the slower convergence in the high dimensional problems. However, we believe that formulating the problem in composite space helps understand how choices such as prioritization affect the continuous space, and this deeper understanding can enable the design of better planners.

We introduced a problem formulation and planners that can deal with all types of task specification. However, we believe that there is a better way to approach the problem than a brute-force search on the task level, which the sampling-based planners effectively do here. Despite these limitations, we find that the planning in composite space is likely good enough for many complex industrial settings, where the number of used robots does not go much beyond four robots. Furthermore, with our initial mode sampling strategy defined in Section 4, we were able to rapidly reduce the decision space by selecting a feasible mode sequence early in the search, allowing us to find an initial solution quickly and thus restricting the search space.

7 Conclusion

We presented a formalization of the multi-modal, multi-robot, multi-goal motion planning problem, and open-source the developed benchmark, which contains diverse problems reaching from simple environments with short goal sequences, to complex long horizon problems requiring intricate coordination of multiple robots and many mode

Table 2: Results for more scenarios from the benchmark problems. We show the median of all the statistics of 25 runs. The cost $c_{t_{\max}}$ is the cost after the time specified. In addition to the solution of the standard prioritized planner, we also shortcut the plan from the prioritized planner using our shortcutting method, and show the cost under $c_{t_{\max}}$.

	#Dim	#Tasks	t_{\max} [s]	t_{init} [s]			c_{init}			$c_{t_{\max}}$		
				Prio	BiRRT*	AIT*	Prio	BiRRT*	AIT*	Prio	BiRRT*	AIT*
simple	6	3	100	1.09	0.17	0.20	4.88	11.94	6.24	3.69	2.27	2.28
single_agent_mover	3	3	100	—	0.09	0.11	—	8.83	4.47	—	3.15	2.90
2d_handover	6	6	500	—	4.13	20.43	—	56.97	53.74	—	17.90	19.21
random_2d	9	13	500	—	7.95	31.29	—	124.49	90.56	—	26.09	27.11
other_hallway	6	3	500	2.92	0.37	0.37	29.41	35.66	28.45	23.62	9.76	10.67
three_agents	9	6	500	34.13	1.70	1.81	24.96	38.79	35.81	18.84	6.35	6.47
triple_waypoints	20	19	1000	10.76	1.00	0.56	36.44	170.02	113.93	26.04	17.34	15.73
welding	24	25	1000	65.82	1.93	1.33	27.49	265.84	185.70	17.82	14.49	12.04
handover	12	4	100	—	0.80	2.85	—	52.30	47.22	—	12.12	12.13
eggcartons	12	19	500	5.12	0.80	0.94	60.05	164.81	158.32	55.51	44.46	43.27
bottles	12	9	500	32.45	16.00	92.72	18.82	86.14	59.69	17.32	15.16	14.65
box_rearrangement	12	27	1000	14.10	3.39	7.19	38.30	223.80	171.97	32.98	28.27	27.11
box_reorientation	12	27	500	28.79	3.76	4.95	114.15	221.58	207.76	85.31	55.74	55.66
pyramid	24	13	500	23.99	3.23	2.10	46.13	142.32	105.08	22.61	18.09	17.65
box_stacking	24	17	1000	14.98	7.03	4.51	53.34	204.13	187.34	26.99	17.90	16.68
mobile_wall_four	24	17	1000	29.34	11.61	13.38	103.01	248.29	220.33	87.20	46.73	43.47
mobile_strut	14	91	1000	75.37	13.67	49.35	345.86	918.65	747.70	343.56	326.51	304.97
three_robot_truss	18	29	1000	24.68	3.92	3.77	89.29	283.28	211.93	76.81	69.99	66.38
spiral_tower	24	31	1000	196.74	26.37	31.00	122.85	410.49	311.52	106.46	100.44	94.00
spiral_tower_two	12	15	500	14.14	1.37	2.02	47.70	124.45	110.74	46.57	44.21	43.87
cube_four	24	49	1000	117.86	41.87	59.44	161.52	623.01	423.71	135.55	135.30	119.76
dep_piano	6	5	500	1.59	0.41	0.61	12.11	22.68	20.92	5.80	3.55	3.55
dep_box_reorientation	12	27	1000	33.86	5.92	7.29	114.49	220.99	206.70	68.15	57.14	55.81
dep_box_stacking	24	17	1000	20.31	5.26	2.93	60.03	205.04	169.28	26.49	22.60	20.05
dep_mobile_wall_four	24	17	1000	34.50	22.70	18.00	75.17	221.79	182.59	50.41	44.18	38.19
unordered_box_reor.	12	14	1000	8.09	16.43	4.80	26.69	88.42	83.23	20.74	21.48	19.73
unordered_bottles	12	10	500	5.57	6.45	7.07	22.80	86.62	67.38	16.53	17.76	17.15
unassigned_tsp	6	6	500	0.97	1.25	0.10	6.32	14.90	6.33	2.60	3.88	2.49
unassigned_cleanup	12	12	1000	8.14	30.36	6.62	20.61	69.82	66.76	15.38	18.47	17.24
unassigned_stacking	24	10	1000	17.49	110.49	10.02	36.76	88.10	93.85	12.32	11.84	10.65

switches. We adapt multiple standard sampling-based planners that are probabilistically complete and asymptotically optimal to this type of problem, and show that the adaptations are crucial in order to obtain the shown performance. We also benchmarked a prioritized planner against these optimal planners and show that the solutions found by such a suboptimal planner can be competitive in some settings.

There are some clear improvements possible, such as exploiting the multi-robot structure of the planning problems, the nature of the mode families [41], or the multi-query-like structure of the problem [42].

There are many extensions to the chosen problem formulation: Both the problem formulation and the planners support extension of the transition logic to full multi-robot task and motion planning by changing the task-assignment-oracle. In the future, we also want to support constrained multi-robot planning, or kinodynamic planning. Finally, the formulation of the problem is amenable to parallelization, which could offer much faster planning times, and possibly better exploration of the problem space.

Acknowledgments

The authors thank Miguel Zamora for insightful discussions and comments. We would like to thank the reviewers for their comments. This work was part of an Innovation project supported by Innosuisse. VH was supported by the Swiss National Science Foundation through the National Centre of Competence in Digital Fabrication (Grant No. 200021 200644). YH was supported by the ETH postdoc fellowship and the SNSF Ambizione Grant (Grant No. 223384).

References

1. G. Wagner and H. Choset, "Subdimensional expansion for multirobot path planning," *Artif. intell.*, vol. 219, pp. 1–24, 2015.
2. R. Shome and L. E. Kavraki, "Asymptotically optimal kinodynamic planning using bundles of edges," in *Proc. of the IEEE Int. Conf. Robot. Automat. (ICRA)*. IEEE, 2021, pp. 9988–9994.
3. S. Lin, A. Liu, J. Wang, and X. Kong, "A review of path-planning approaches for multiple mobile robots," *Machines*, vol. 10, no. 9, p. 773, 2022.
4. W. Thomason, M. P. Strub, and J. D. Gammell, "Task and Motion Informed Trees (TMIT*): Almost-surely asymptotically optimal integrated task and motion planning," *IEEE Robot. and Automat. Lett. (R-AL)*, vol. 7, no. 4, pp. 11 370–11 377, 2022.
5. K. Hauser, "Task planning with continuous actions and nondeterministic motion planning queries," in *Proc. of AAAI Workshop on Bridging the Gap between Task and Motion Planning*, 2010.
6. P. Englert, I. M. R. Fernández, R. K. Ramachandran, and G. S. Sukhatme, "Sampling-Based Motion Planning on Sequenced Manifolds," in *Proc. of Robotics: Science and Systems (R:SS)*, 2021.
7. P. S. Schmitt, W. Neubauer, W. Feiten, K. M. Wurm, G. V. Wichert, and W. Burgard, "Optimal, sampling-based manipulation planning," in *Proc. of the IEEE Int. Conf. Robot. Automat. (ICRA)*. IEEE, 2017, pp. 3426–3432.
8. V. N. Hartmann, A. Orthey, D. Driess, O. S. Oguz, and M. Toussaint, "Long-horizon multi-robot rearrangement planning for construction assembly," *IEEE Trans. Robot.*, 2022.
9. V. N. Hartmann and M. Toussaint, "Towards computing low-makespan solutions for multi-arm multi-task planning problems," in *Int. Conf. on Autom. Plan. and Sched.: Planning and Robotics Workshop (RobPlan)*, 2023.
10. R. Stern, "Multi-agent path finding—an overview," *Artificial Intelligence: 5th RAAI Summer School, Tutorial Lectures*, pp. 96–115, 2019.
11. J. Yu and S. M. LaValle, "Optimal multirobot path planning on graphs: Complete algorithms and effective heuristics," *IEEE Trans. Robot.*, vol. 32, no. 5, pp. 1163–1177, 2016.
12. H. Ma, J. Li, T. S. Kumar, and S. Koenig, "Lifelong multi-agent path finding for online pickup and delivery tasks," in *Int. Conf. on Autonomous Agents and MultiAgent Syst. (AAMAS)*, ser. AAMAS '17, 2017, p. 837–845.
13. F. Grenouilleau, W.-J. Van Hoesve, and J. N. Hooker, "A multi-label A* algorithm for multi-agent pathfinding," in *Int. Conf. on Autom. Plan. and Sched.*, vol. 29, 2019, pp. 181–185.
14. A. Moldagalieva, J. Ortiz-Haro, M. Toussaint, and W. Hönig, "db-cbs: Discontinuity-bounded conflict-based search for multi-robot kinodynamic motion planning," in *Proc. of the IEEE Int. Conf. Robot. Automat. (ICRA)*. IEEE, 2024, pp. 14 569–14 575.
15. R. Shome, K. Solovey, A. Dobson, D. Halperin, and K. E. Bekris, "dRRT*: Scalable and informed asymptotically-optimal multi-robot motion planning," *Autonomous Robots*, vol. 44, no. 3-4, pp. 443–467, 2020.
16. G. Sanchez and J.-C. Latombe, "Using a PRM planner to compare centralized and decoupled planning for multi-robot systems," in *Proc. of the IEEE Int. Conf. Robot. Automat. (ICRA)*, vol. 2, 2002, pp. 2112–2119 vol.2.
17. A. Orthey, S. Akbar, and M. Toussaint, "Multilevel motion planning: A fiber bundle formulation," *Int. J. of Robot. Research*, vol. 43, no. 1, pp. 3–33, 2024.
18. J. Chen, J. Li, Y. Huang, C. Garrett, D. Sun, C. Fan, A. Hofmann, C. Mueller, S. Koenig, and B. C. Williams, "Cooperative task and motion planning for multi-arm assembly systems," *arXiv preprint arXiv:2203.02475*, 2022.

19. J. Motes, T. Chen, T. Bretl, M. M. Aguirre, and N. M. Amato, "Hypergraph-based multi-robot task and motion planning," *IEEE Transactions on Robotics*, vol. 39, no. 5, pp. 4166–4186, 2023.
20. S. M. LaValle and J. J. Kuffner, "Rapidly-exploring random trees: Progress and prospects," *Algorithmic and computational robotics*, pp. 303–307, 2001.
21. L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *Proc. of the IEEE Int. Conf. Robot. Automat. (ICRA)*, vol. 12, no. 4, pp. 566–580, 1996.
22. K. Solovey, O. Salzman, and D. Halperin, "Finding a needle in an exponential haystack: Discrete RRT for exploration of implicit roadmaps in multi-robot motion planning," *Int. J. of Robot. Research*, vol. 35, no. 5, pp. 501–513, 2016.
23. A. Solano, A. Sieverling, R. Giesemann, and A. Orthey, "Fast-dRRT*: Efficient Multi-Robot Motion Planning for Automated Industrial Manufacturing," *arXiv preprint arXiv:2309.10665*, 2023.
24. G. Sharon, R. Stern, A. Felner, and N. R. Sturtevant, "Conflict-based search for optimal multi-agent pathfinding," *Artif. intell.*, vol. 219, pp. 40–66, 2015.
25. I. Solis, J. Motes, R. Sandström, and N. M. Amato, "Representation-optimal multi-robot motion planning using conflict-based search," *IEEE Robot. and Automat. Lett. (R-AL)*, vol. 6, no. 3, pp. 4608–4615, 2021.
26. J. Kottinger, S. Almagor, and M. Lahijanian, "Conflict-Based Search for Multi-Robot Motion Planning with Kinodynamic Constraints," in *Proc. of the IEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2022, pp. 13 494–13 499.
27. J. P. Van Den Berg and M. H. Overmars, "Prioritized motion planning for multiple robots," in *Proc. of the IEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*. IEEE, 2005, pp. 430–435.
28. K. Okumura and X. Défago, "Quick multi-robot motion planning by combining sampling and search," in *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, 8 2023, pp. 252–261.
29. G. Han, J. Park, and C. Nam, "Stop-n-go: Search-based conflict resolution for motion planning of multiple robotic manipulators," in *Proc. of the IEEE Int. Conf. Robot. Automat. (ICRA)*, 2025, pp. 7719–7725.
30. H. Ma, D. Harabor, P. J. Stuckey, J. Li, and S. Koenig, "Searching with consistent prioritization for multi-agent path finding," in *Proc. of the AAAI Conf. on artif. intell.*, vol. 33, no. 01, 2019, pp. 7643–7650.
31. E. Boyarski, A. Felner, R. Stern, G. Sharon, O. Betzalel, D. Tolpin, and E. Shimony, "ICBS: The improved conflict-based search algorithm for multi-agent pathfinding," in *Int. Symp. on Combinatorial Search*, vol. 6, no. 1, 2015, pp. 223–225.
32. J. Li, W. Ruml, and S. Koenig, "EECBS: A bounded-suboptimal search for multi-agent path finding," in *Proc. of the AAAI Conf. on artif. intell.*, vol. 35, no. 14, 2021, pp. 12 353–12 362.
33. O. Salzman and R. Stern, "Research challenges and opportunities in multi-agent path finding and multi-agent pickup and delivery problems," 2020, pp. 1711–1715.
34. G. Mouratidis, B. Nebel, and S. Koenig, "Fools Rush in Where Angels Fear to Tread in Multi-Goal CBS," in *Int. Symp. on Combinatorial Search*, vol. 17, 2024, pp. 243–251.
35. S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *Int. J. of Robot. Research*, vol. 30, no. 7, pp. 846–894, 2011.
36. M. P. Strub and J. D. Gammell, "Adaptively Informed Trees (AIT*) and Effort Informed Trees (EIT*): Asymmetric bidirectional sampling-based path planning," *Int. J. of Robot. Research*, vol. 41, no. 4, pp. 390–417, Apr. 2022.
37. J. D. Gammell, S. S. Srinivasa, and T. D. Barfoot, "Informed RRT*: Optimal sampling-based path planning focused via direct sampling of an admissible ellipsoidal heuristic," in *Proc. of the IEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 14–18 Sept. 2014, pp. 2997–3004.

38. A. Atias, K. Solovey, O. Salzman, and D. Halperin, “Effective metrics for multi-robot motion-planning,” *Int. J. of Robot. Research*, vol. 37, no. 13-14, pp. 1741–1759, 2018.
39. M. P. Strub and J. D. Gammell, “Adaptively Informed Trees (AIT*): Fast asymptotically optimal path planning through adaptive heuristics,” in *Proc. of the IEEE Int. Conf. Robot. Automat. (ICRA)*. IEEE, 2020, pp. 3191–3198.
40. R. Geraerts and M. H. Overmars, “Creating high-quality paths for motion planning,” *Int. J. of Robot. Research*, vol. 26, no. 8, pp. 845–863, 2007.
41. Z. Kingston, C. Chamzas, and L. E. Kavraki, “Using experience to improve constrained planning on foliations for multi-modal problems,” in *Proc. of the IEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2021, pp. 6922–6927.
42. V. N. Hartmann, M. P. Strub, M. Toussaint, and J. D. Gammell, “Effort informed roadmaps (EIRM*): Efficient asymptotically optimal multiquery planning by actively reusing validation effort,” in *Int. Symp. of Robot. Research*, 2022.