

Scalable Multi-Robot Motion Planning for Congested Environments With Topological Guidance

Courtney McBeth , James Motes , Diane Uwacu , Marco Morales , *Member, IEEE*,
and Nancy M. Amato , *Fellow, IEEE*

Abstract—Multi-robot motion planning (MRMP) is the problem of finding collision-free paths for a set of robots in a continuous state space. The difficulty of MRMP increases with the number of robots and is exacerbated in environments with narrow passages that robots must pass through, like warehouse aisles where coordination between robots is required. In single-robot settings, topology-guided motion planning methods have shown improved performance in these constricted environments. In this work, we extend an existing topology-guided single-robot motion planning method to the multi-robot domain to leverage the improved efficiency provided by topological guidance. We demonstrate our method’s ability to efficiently plan paths in complex environments with many narrow passages, scaling to robot teams of size up to 25 times larger than existing methods in this class of problems. By leveraging knowledge of the topology of the environment, we also find higher-quality solutions than other methods.

Index Terms—Path planning for multiple mobile robots or agents, multi-robot systems, motion and path planning.

I. INTRODUCTION

MULTI-ROBOT systems have become ubiquitous in many settings such as autonomous factories and warehouses. When the motions of two or more robots are in conflict, highly coordinated multi-robot motion planning (MRMP) is required to avoid collisions with each other and with the environment.

Existing MRMP approaches perform well in open environments but struggle with narrow passages (e.g. warehouse aisles) due to the difficulty of avoiding collisions with obstacles within tight spaces. With multi-robot teams, narrow passages may be introduced or exacerbated by inter-robot collisions, further complicating planning. There are three classes of MRMP approaches: decoupled, which plan a path for each robot separately and offer the least amount of coordination, coupled, which plan in the *composite space*, which incorporates the degrees

of freedom of all robots, and hybrid, which offer a mix of coupled and decoupled planning. Decoupled approaches (e.g., Decoupled PRM [1]) allow for linear scaling with the number of robots; however, they lack the coordination necessary to resolve complex inter-robot collisions. Coupled methods (e.g., Composite RRT [2] and Composite PRM [1]) provide this coordination but consider a large search space that becomes computationally intractable for large teams. Hybrid methods (e.g., CBS-MP [3] and M* [4]) leverage the scalability of decoupled methods along with an increased level of coordination. However, in environments with narrow passages where inter-robot collisions are likely, their performance is limited by the time spent resolving conflicts. In this work, we propose a coupled method that exploits topological guidance to more intelligently and efficiently search the composite space.

Guided planning methods leverage external information to efficiently find paths. Topology-guided methods [5], [6], [7] exploit representations of the environment to direct motion planning through narrow passages. Topological skeletons are embedded graphs encoding the environment. Prior work [5], [6] has explored guiding planning around a skeleton.

In this work, we extend topological skeleton guidance to multi-robot systems to provide the level of coordination required for large teams in environments with narrow passages. We present *Composite Dynamic Region-biased Rapidly-exploring Random Trees* (CDR-RRT), a coupled MRMP approach that leverages knowledge of the workspace topology, leading to improved planning times while retaining probabilistic completeness. We demonstrate significantly improved scalability compared to existing state-of-the-art methods, successfully finding paths for teams of size up to 25 times larger in complex environments where inter-robot collisions are likely. Our contributions include:

- The development of novel composite-space analogs of workspace skeletons and sampling regions.
- A scalable, probabilistically complete MRMP method that leverages topological guidance to address problems that require high levels of coordination.
- An experimental validation in a variety of congested and open environments with scaling robot team sizes.

II. PRELIMINARIES AND RELATED WORK

In this section, we discuss the motion planning problem and research in both the multi-robot motion planning and guided motion planning domains.

A. Motion Planning Preliminaries

A robot’s *degrees of freedom* (DOFs) include its position and orientation in the *workspace*, the 2D or 3D space within which

Manuscript received 25 May 2023; accepted 17 August 2023. Date of publication 7 September 2023; date of current version 19 September 2023. This letter was recommended for publication by Associate Editor T. Bandyopadhyay and Editor A. Bera upon evaluation of the reviewers’ comments. This work was supported in part by Foxconn Interconnect Technology (FIT) and in part by the Center for Networked Intelligent Components and Environments (C-NICE) at UIUC. (*Corresponding author: Courtney McBeth.*)

Courtney McBeth, James Motes, Marco Morales, and Nancy M. Amato are with the Parasol Lab, Department of Computer Science, University of Illinois Urbana-Champaign, Champaign, IL 61820 USA (e-mail: cmcbeth2@illinois.edu; jmotes2@illinois.edu; moralesa@illinois.edu; namato@illinois.edu).

Diane Uwacu is with the Department of Computer Science and Engineering, Texas A&M University, College Station, TX 77840 USA (e-mail: duwacu@tamu.edu).

This letter has supplementary downloadable material available at <https://doi.org/10.1109/LRA.2023.3312980>, provided by the authors.

Digital Object Identifier 10.1109/LRA.2023.3312980

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see <https://creativecommons.org/licenses/by/4.0/>

TABLE I
OVERVIEW OF RELATED METHODS

Algorithm	Coordination	Scalable	Probabilistic Completeness	Narrow Passages
Decoupled PRM [1]	Decoupled	X		
Composite RRT [2]	Coupled		X	
Composite PRM [1]	Coupled		X	
MRdRRT [15]	Coupled	X	X	
CBS-MP [3]	Hybrid	X	X	
MAPF/C [16]	Hybrid	X	X	
CDR-RRT [Ours]	Coupled	X	X	X

the robot physically exists, as well as other configurable values such as joint angles. A *configuration* is a set of values describing the robot's DOFs. The *configuration space* (C_{space}) is the set of all robot configurations [8]. The *free space* (C_{free}) is the subset of C_{space} that only contains valid configurations (e.g. configurations not in collision with obstacles). The *obstacle space* (C_{obst}) contains all configurations that are not valid. Given a start configuration, q_{start} , and a goal configuration, q_{goal} , the *motion planning problem* strives to find a path from q_{start} to q_{goal} through C_{free} .

Searching the entire C_{space} is intractable [9], [10], resulting in the emergence of sampling-based motion planning algorithms [2], [11]. These algorithms forego completeness guarantees in favor of faster planning and *probabilistic completeness*, meaning that the probability of finding a solution, if one exists, converges to 1 in the upper limit of planning time. Unfortunately, these randomized sampling techniques suffer in constrained environments [12], where they face the *narrow passage problem*. This refers to the difficulty of sampling valid configurations within narrow corridors that by volume make up a small proportion of the freespace.

The underlying sampling-based motion planning algorithm that forms the basis of our method is Rapidly-exploring Random Trees (RRT) [2]. This method iteratively grows a tree, T , from q_{start} to q_{goal} . During each iteration, a random configuration q_{rand} is sampled. We then find q_{near} , the configuration in T closest to q_{rand} . T is then extended a maximum distance Δ from q_{near} in the direction of q_{rand} . RRTs exhibit a Voronoi bias that results in the rapid exploration of the C_{space} and makes RRTs an efficient way to handle single-query motion planning problems. RRT variants have been developed to improve performance in the presence of narrow passages [13], [14].

B. Multi-Robot Motion Planning

Multi-robot motion planning consists of finding valid paths for a set of robots between their respective starts and goals. The composite space is the Cartesian product of each of the n individual robots' C_{spaces} : $C_{composite} = C_1 \times C_2 \times \dots \times C_n$ where C_i represents the C_{space} of robot i . A composite configuration consists of values for each robot's DOFs. The composite free space is made up of all valid configurations such that no robot is in collision with another robot. The MRMP problem can be formulated as finding a continuous path through the composite free space.

Table I compares select MRMP approaches. Decoupled approaches such as Decoupled PRM [1] plan individual robot paths in their own decoupled C_{spaces} and thus do not offer completeness or optimality guarantees. The lack of coordination degrades their performance in narrow passages, where inter-robot collisions may not be possible to avoid along the individual paths only using velocity tuning.

Coupled methods such as Composite PRM [1] and Composite RRT [2] plan directly within the composite C_{space} . Other composite methods (e.g., MRdRRT [15]) build individual robot roadmaps, then search an implicit composite roadmap. These methods maintain the probabilistic completeness of the single-robot methods they use. However, due to the decoupled individual roadmap construction, they lack the level of coordination required to efficiently find paths in congested environments with narrow passages.

Hybrid methods such as CBS-MP [3], MAPF/C [16], and M* [4] seek to leverage the strengths of both coupled and decoupled methods. For example, CBS-MP [3] plans individual robot paths in their decoupled C_{spaces} and then reconciles the paths in the composite C_{space} . In the worst case, these methods will explore the whole composite C_{space} , but on average, the runtimes are comparable to decoupled methods while providing varying levels of probabilistic completeness and representation optimality guarantees. These methods are generally not well suited for environments with narrow passages due to the computational effort expended transitioning the planner from decoupled to the high level of coordination required. Here, we propose a method for composite-space RRT construction while leveraging topological guidance to improve performance in environments with narrow passages.

C. Guided Motion Planning

Topological guidance has not been well explored in the composite space; however, some hybrid methods have explored using topological information to construct decoupled roadmaps. Ryan [17] decomposes the workspace into halls, represented by singly-linked chains of vertices, and open spaces, represented by fully-connected subgraphs. Yu et al. [18] propose a method for roadmap construction via overlaying a lattice structure onto the workspace. They show that this method leads to efficient path planning for large groups of robots in relatively open environments.

Several single-robot motion planning strategies have been proposed to adapt planning to the workspace. The Feature Sensitive Motion Planning Strategy [19] attempts to subdivide the environment into homogeneous workspace regions that are planned in individually, adapting roadmap construction to local features. The individual roadmaps are merged into a complete roadmap of the planning space. This strategy allows the planner to use resources efficiently.

Workspace Decomposition Strategies [20], [21], [22] help concentrate planning in narrow areas of the workspace. SyClop [23] uses an RRT to sample frontier decomposition regions. A User-Guided Planning Strategy [24] allows the user to define and manipulate workspace sampling regions that the planner explores in real-time. The planner relies on the user's intuition to identify narrow passages and find paths faster.

Skeleton-based strategies leverage the topology of the workspace using an embedded graph (Fig. 2(a)) that is homotopy equivalent to the workspace. All points in the workspace can be smoothly collapsed to the skeleton [25]. Skeleton edges describe contiguous volumes of the free workspace (e.g., tunnels or rooms) and vertices represent connections between these volumes. Given the environment, the skeleton is precomputed and may be used for multiple queries and different types of robots. Examples include medial axis skeletons [26] in 2D and mean curvature skeletons [27] in 3D. Skeletons are generally



Fig. 1. We demonstrate our method on a variety of physical and simulated environments. Video: <https://youtu.be/xP7mSWxdYfs>.




Fig. 2. Example workspace skeleton (a) and query skeleton (b). The query skeleton is a directed and pruned skeleton that only contains edges along a path from q_{start} to q_{goal} .

Algorithm 1: Dynamic Region-biased RRT.

Input: Environment e and a Query $\{q_s, q_g\}$

Output: Tree T

- 1: $W \leftarrow \text{COMPUTEWORKSPACE SKELETON}(e)$
 - 2: $S \leftarrow \text{COMPUTEQUERY SKELETON}(W, q_s, q_g)$
 - 3: $T \leftarrow (\emptyset, \emptyset)$
 - 4: $R \leftarrow \text{INITIALREGIONS}(S, q_s)$
 - 5: **while** $\neg(q_g \in T)$ **do**
 - 6: $\text{REGIONBIASEDRRTGROWTH}(T, S, R)$
 - 7: **return** T
-

quick to compute. The medial axis skeleton, for example, can be computed in $O(n \log n)$ time where n is the number of obstacle edges [28].

Dynamic Region Sampling with PRM (DR-PRM) [6], initiates local components at the vertices of a skeleton, expands them along adjacent edges, then merges them to form a complete roadmap. Hierarchical Annotated Skeleton Planning [29] extends DR-PRM by relaxing its reliance on skeleton edges over time. We describe the single-query counterpart of DR-PRM, Dynamic Region-biased RRT (DR-RRT), in detail in Section II-D since we extend our method from it. These methods show the advantage of using workspace information to guide planning in \mathcal{C}_{space} when they are closely related; however, they are constrained to single-robot settings.

D. Dynamic Region-Biased RRT

DR-RRT [5] (Algorithm 1) grows an RRT while constraining sampling within regions that advance along a skeleton.

1) *Query Skeleton:* Algorithm 1 creates two skeletons: one for the entire workspace (line 1), and the *query skeleton* (Fig. 2(b); line 2), which retains only edges that are along a path from the start to the goal in the workspace.

2) *Sampling Regions:* Sampling regions will advance along the edges of the query skeleton, using the skeleton, a solution in the workspace, to guide construction of an RRT, a solution in

Algorithm 2: Region-biased RRT Growth.

Input: Tree T , Query Skeleton S , Regions R , region radius η , maximum for failed extension attempts τ

- 1: $r \leftarrow \text{SELECTREGION}(R)$
 - 2: $q_{rand} \leftarrow r.\text{GETRANDOMCFG}()$
 - 3: $q_{near} \leftarrow \text{NEARESTNEIGHBOR}(T, q_{rand})$
 - 4: $q_{new} \leftarrow \text{EXTEND}(q_{near}, q_{rand}, \Delta)$
 - 5: **if** $r = \emptyset$ **then** \triangleright The entire environment was chosen
 - 6: **return**
 - 7: **if** $q_{new}.\text{EXTENSION SUCCEEDED}()$ **then**
 - 8: $r.\text{INCREMENT SUCSESSES}()$
 - 9: **else**
 - 10: $r.\text{INCREMENT FAILURES}()$
 - 11: **for** $r \in R$ **do**
 - 12: **while** $r.\text{INREGION}(q_{new})$ **do**
 - 13: $r.\text{ADVANCE ALONG SKELETON EDGE}()$
 - 14: **if** $r.\text{AT END OF SKELETON EDGE}()$ **then**
 - 15: $R \leftarrow R \setminus \{r\}$
 - 16: **if** $r.\text{NUM FAILURES}() > \tau$ **then**
 - 17: $R \leftarrow R \setminus \{r\}$
 - 18: **for all** $v \in S.\text{UNEXPLORED VERTICES}()$ **do**
 - 19: **if** $\delta(v, q_{new}) < \eta$ **then**
 - 20: $R \leftarrow R \cup \text{NEWREGION}(v)$
 - 21: $S.\text{MARKEXPLORED}(v)$
-

the \mathcal{C}_{space} . A *region* is a bounded volume in the workspace, e.g., a bounding sphere. Construction of the RRT begins by initializing the first region centered on the skeleton vertex closest to q_{start} (line 4). During each iteration, Algorithm 2 selects a region to guide sampling (line 1). The probability of selecting a region is proportional to its *extension success rate*. To maintain probabilistic completeness, with a small probability, the entire environment is chosen (see Section III-E). A random configuration q_{rand} is selected from this region to grow the tree toward (line 3). The algorithm then proceeds as a general RRT by attempting to extend the tree to q_{new} (line 4). The extension success rate is updated based on the outcome of this attempt (lines 7–10).

3) *Region Advancement:* Once a configuration has been added to the tree, all regions that are in contact with q_{new} are advanced forward along their skeleton edges until they leave q_{new} behind (Algorithm 2, lines 11–13). Any region that reaches the end of its edge or exceeds the maximum number of extension failures is deleted (lines 14–17). Then, new regions are created on each unexplored skeleton vertex that is within a small distance of q_{new} (lines 18–21). This cycle of region selection, tree extension, region advancement, and region creation continues until the tree extends to q_{goal} .

III. COMPOSITE DYNAMIC REGION-BIASED RRT

In this article, we extend DR-RRT to multi-robot systems to propose a new method, Composite Dynamic Region-biased RRT (CDR-RRT). We limit the exploration of the composite space to areas that are likely to yield a solution because of the exponential size of the search space. We do this by developing composite analogs for workspace skeletons and regions that allow for coupled multi-robot motion planning while leveraging topological guidance as in DR-RRT. We leverage lazy construction of the

Algorithm 3: Composite Region-biased RRT Growth.

Input: Tree T , Composite Skeleton S , Region r , Failed Vertex and Edge Constraints C , Maximum Failed Extension Attempts τ , Maximum Failed Composite Skeleton Edges per Vertex ρ

- 1: $Bound \leftarrow \text{SELECTREGIONORWHOLEENVIRONMENT}()$
- 2: $q_{rand} \leftarrow r.\text{GETRANDOMCOMPOSITECFG}(Bound)$
- 3: $q_{near} \leftarrow \text{NEARESTNEIGHBOR}(T, q_{rand})$
- 4: $q_{new} \leftarrow \text{EXTEND}(q_{near}, q_{rand}, \Delta)$
- 5: **if** $Bound = \text{WHOLEENVIRONMENT}$ **then**
- 6: **return**
- 7: **if** $q_{new}.\text{EXTENSIONSSUCCEEDED}()$ **then**
- 8: $r.\text{INCREMENTSUCCESSSES}()$
- 9: **else**
- 10: $r.\text{INCREMENTFAILURES}()$
- 11: **while** $r.\text{INCOMPOSITEREGION}(q_{new})$ **do**
- 12: $r.\text{ADVANCEALONGCOMPOSITESKELETONEDGE}()$
- 13: **if** $r.\text{ATENDOFCOMPOSITESKELETONEDGE}()$ **then**
- 14: $V_t \leftarrow r.\text{EDGE.TARGET} \triangleright$ Get the edge's target vertex
- 15: $r \leftarrow S.\text{GROWCOMPOSITESKELETON}(V_t, C) \triangleright$ Algorithm 4
- 16: **if** $r.\text{NUMFAILURES} > \tau$ **then**
- 17: $C \leftarrow C \cup \{r.\text{EDGE}\}$
- 18: $V_s \leftarrow r.\text{EDGE.SOURCE} \triangleright$ Get the edge's source vertex
- 19: $V_s.\text{INCREMENTFAILURES}()$
- 20: **if** $V_s.\text{NUMFAILURES} > \rho$ **then**
- 21: $C \leftarrow C \cup \{V_s\}$
- 22: $V_s \leftarrow V_s.\text{GETPREDECESSOR}()$ \triangleright Backtrack vertex
- 23: $r \leftarrow S.\text{GROWCOMPOSITESKELETON}(V_s, C) \triangleright$ Algorithm 4




Fig. 3. Example of a composite edge and a composite region. Each of the three robots has its own workspace skeleton. Obstacles are gray, and the workspace skeletons are purple. Three skeleton edges that comprise a composite edge are shown in blue. All possible combinations of such edges make up the composite skeleton. Three individual regions that make up a composite region along this composite edge are shown in green. An RRT grown along the highlighted skeleton edges is shown in orange.

composite skeleton as we exploit a greedy heuristic to search the composite space.

A. Composite Skeleton

As the composite space is the Cartesian product of each robot's C_{space} , the composite skeleton (Fig. 3) is the Cartesian product of the workspace skeleton for each of the n robots. It consists of composite vertices and edges which respectively represent a set of n vertices or edges in the workspace skeleton where each of the n robots lies. We avoid the exponential expansion associated with the computation of the full composite

Algorithm 4: Grow Composite Skeleton.

Input: Composite Skeleton S , Composite Source Vertex V , Failed Vertex and Edge Constraints C

- 1: **while** $V.\text{EXCEEDEDGROWTHATTEMPTS}()$ **do**
- 2: $V \leftarrow V.\text{GETPREDECESSOR}()$
- 3: $Paths \leftarrow \text{MAPFSOLUTION}(V, C)$
- 4: $E \leftarrow \text{EXTRACTFIRSTCOMPOSITEEDGE}(Paths)$
- 5: $S.\text{ADDEDGE}(E)$
- 6: **return** $\text{NEWREGION}(E)$




Fig. 4. Examples of composite region advancement for two point robots in a 1D environment. The individual robots' skeletons are shown on the axes and the composite skeleton between them. Each individual tree is shown in blue and the composite tree is in purple. Intermediates are tick marks along the skeletons. Individual regions advance at the same rate with respect to their intermediates until the composite region leaves q_{new} behind at the position shown in dark green. In (a), RRT extension reaches the newly sampled point. In (b), the sample point cannot be reached due to inter-robot collision, so extension continues until the furthest valid point along the edge.

skeleton graph by using local, on-demand construction. A composite region is made up of n individual sampling regions, one in each robot's workspace.

Computing a composite query skeleton, which is a directed and pruned version of the composite skeleton, requires an explicit computation of the composite skeleton. Instead, we heuristically construct and search the composite skeleton one edge at a time and only consider edges likely to be along a feasible low-cost path from q_{start} to q_{goal} . We discuss our heuristic to capture these edges in Section III-C.

B. Guided Composite RRT Construction

To begin RRT construction, we compute the first composite skeleton edge to explore. Section III-C discusses the construction of composite skeleton edges by growing the composite skeleton from a source composite vertex. CDR-RRT then proceeds as DR-RRT by iteratively performing region-biased sampling, RRT growth, region advancement and deletion, and new region creation until q_{goal} is reached.

During each iteration of CDR-RRT, Algorithm 3 selects a composite region r . After sampling from r and extending the tree, we advance r forward until it leaves q_{new} behind (line 12). In DR-RRT [5], as individual regions advance along skeleton edges, they are centered on intermediate points along the edges. Correspondingly, we create composite intermediates along composite skeleton edges. In composite region advancement, as shown in Fig. 4, all individual regions are advanced forward the minimum amount of intermediates such that the composite region is no longer touching q_{new} .

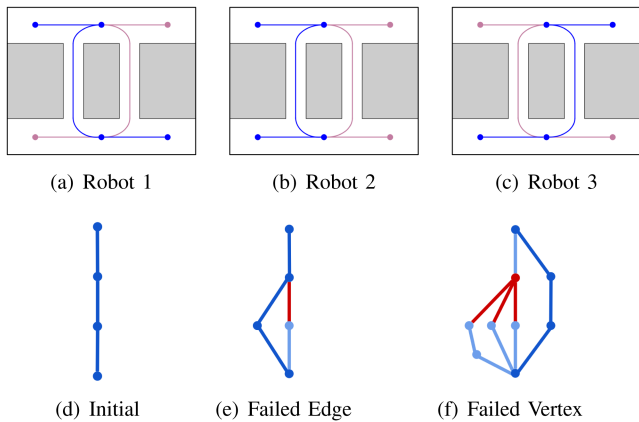


Fig. 5. Example of composite skeleton construction. (a)–(c) Show the initial MAPF solution over the skeleton in blue. These paths are combined into a path through the composite skeleton (d). (e) If a composite skeleton edge (in red) cannot be traversed, we impose a constraint that future MAPF solutions cannot contain that edge and replan the paths from the last vertex reached. (f) If the maximum number of failed outgoing edges from a vertex is exceeded, a constraint is imposed that future MAPF solutions cannot contain that vertex, and paths are replanned from its predecessor.

When a composite region reaches the target vertex at the end of its edge (i.e., it is centered on the target and contains q_{new}) (line 13), we add an outgoing edge to the composite skeleton from this vertex and spawn a new region (line 15). When a region surpasses τ failed extension attempts, it is deleted and replaced with a new region (line 23).

C. Multi-Agent Pathfinding Heuristic

We use a multi-agent pathfinding (MAPF) heuristic to identify the next composite skeleton edges to construct and explore given a source composite vertex, V (Algorithm 4). MAPF is the discrete state space equivalent of the MRMP problem. We use MAPF to generate a path for each robot over the workspace skeleton from V to the vertex closest to each robot’s goal (line 3). Section III-D details our MAPF formulation. We ensure these paths are feasible by accounting for potential collisions between robots. We define the capacity of a skeleton edge as the minimum width between obstacles along the edge. If the total width of the robots traversing that edge exceeds the capacity, a conflict has occurred. These conflicts are resolved by the MAPF algorithm. When multiple robots can traverse an edge, the region size is set large enough to capture all robots.

We extract composite skeleton edges from the produced MAPF solution (Algorithm 4, line 4; Fig. 5) and iteratively create a region to traverse each edge. If a region exceeds the maximum number of extension failures traversing an edge, we consider it failed and impose a constraint that no further MAPF solutions can contain that edge (Algorithm 3, line 17; Fig. 5(e)). We also increment the number of failed growth attempts that each source composite skeleton vertex has seen (Algorithm 3, line 19). To avoid excess computation exploring a region of the composite skeleton unlikely to yield a path, if a vertex exceeds the maximum number of growth failures ρ , we impose a constraint that no further MAPF solutions can contain that vertex (Algorithm 3, line 21) and backtrack to the last vertex below ρ failures (Algorithm 3, line 22; Fig. 5(f)).

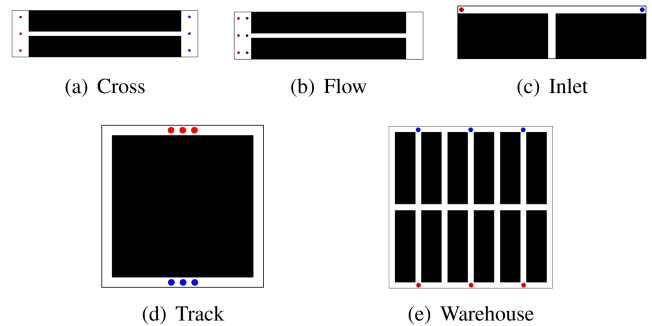


Fig. 6. Cross (a) and Flow (b) scenarios feature a single hallway through which all robots must pass. In Cross, the robots on either side must swap places. In Flow, all robots must move from one side of the hallway to the other. In the Inlet (c) scenario, the robots must swap places. One robot must move into the inlet to allow the other to pass. The Track environment (d) features a ring through which robots must move. The robots on top and bottom must swap places by moving in the same direction to avoid collision. The Warehouse scenario (e) has three variants of aisle widths (1, 2, and 4 meters; the 1 m variant is shown). The robots on top and bottom must swap places with the robot with which they are aligned vertically.




Fig. 7. In the Open Cross (a) environment, the red robots must move from left to right and the blue robots from top to bottom. In the Maze Cross (b) environment, robots starting on the left and right must swap places by moving through the 3D tunnel structure.

D. Implementation Details


To generate MAPF solutions, we adapt Conflict-Based Search (CBS) [30] and Priority-Based Search (PBS) [31]. Both use a hierarchical approach with a low-level search to find individual paths for each agent and a high-level search to resolve conflicts between paths. CBS finds optimal paths with respect to the makespan while PBS has been shown to achieve improved performance in scenarios where the optimal path for one robot blocks the path for other robots.

The size of the full composite skeleton is exponential in the number of robots, so we optimize memory usage by leveraging local construction of the composite skeleton. We can also remove composite vertices and edges when they are no longer useful. A composite edge is no longer useful when the composite region that traverses it has reached the end and been deleted. A composite vertex is no longer useful when all of its incoming and outgoing edges are no longer useful.


E. Theoretical Analysis

Theorem: CDR-RRT is probabilistically complete.


Proof: During each iteration of CDR-RRT, there is a probability $\epsilon > 0$ of sampling from the entire environment rather than within a region (Algorithm 3, line 1). Sampling from the entire environment guarantees probabilistic completeness, ensuring that a valid path from q_{start} to q_{goal} will be found, if one exists,



(a) Hallway



(b) Inlet



(c) Track

Fig. 8. Running time and path cost results for the Hallway (a), Inlet (b), and Track (c) environments.

even if all regions are unable to produce valid configurations. As ϵ increases to 1 or as the size of regions is increased to encompass the entire workspace, knowledge of the workspace topology is utilized less and the method eventually reduces to Composite RRT, which has been shown to be probabilistically complete [32].

IV. VALIDATION

We run scaling MRMP queries designed to highlight the strengths and weaknesses of our approach. We consider both environments with different narrow passage widths and open environments to measure how our method compares to others

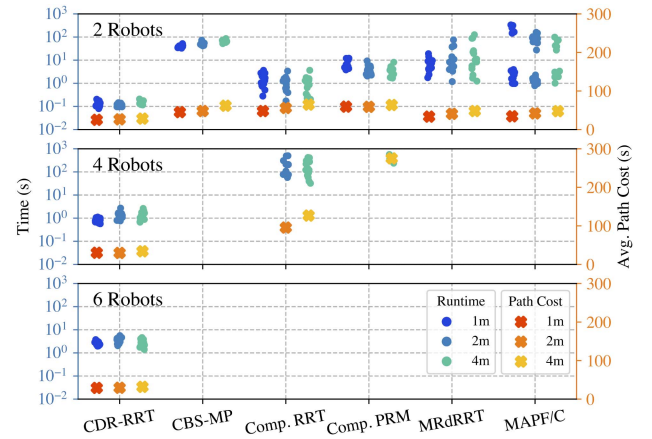


Fig. 9. Running time and path cost results for the Warehouse scenarios.




Fig. 10. Running time results for up to 100 robots for CDR-RRT on the Track Environment including heuristic evaluation times.

when the workspace is and is not informative. We measure each algorithm's performance in scenarios that require various levels of coordination during planning to demonstrate our improved performance when high coordination is required.

A. Experimental Setup

We compare to several state-of-the-art MRMP methods (Table I). We use CBS-MP [3] with DR-PRM [6] to construct the individual roadmaps as a comparison against a hybrid method with workspace guidance. We use the implementation of MAPF/C described in [16] with SPARS [33] roadmap generation. Although MRdRRT [15] was designed primarily for manipulators, we compare to it since its use of a tensor-product roadmap to conduct an RRT-style search of the composite space is similar to our composite skeleton guidance. We pre-compute medial-axis skeletons for 2D environments and mean curvature skeletons for 3D environments.

All methods were implemented in C++ in the Parasol Planning Library. The experiments were run in simulation with holonomic mobile robots using a desktop computer with an Intel Core i9-10900KF CPU at 3.7 GHz and 128 GB of RAM. Each method is given 600 seconds to find a plan or is considered unsuccessful. We report planning times and average path costs given by the makespan.

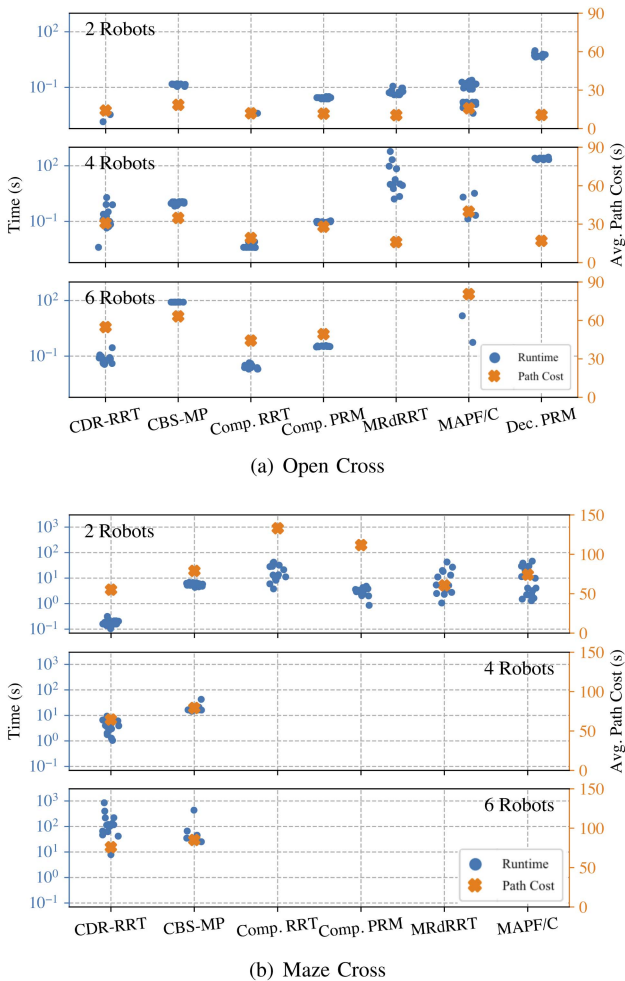


Fig. 11. Results for the Open Cross (a) and Maze Cross (b) environments for all methods. Those that were unable to find a solution within the time limit are omitted.

B. Environments

In this section, we describe our experimental environments and explain why these scenarios highlight the advantages and disadvantages of our approach relative to other methods.

1) *Corridors (Fig. 6(a)–(d))*: We consider the Hallway environment featuring a single tunnel within which only one robot may fit vertically, preventing robots from passing each other. We consider two variants of this scenario, one in which two groups of robots start on opposite ends of the tunnel and swap places (*Cross*, Fig. 6(a)) and one in which all robots start on one side and must move to the other (*Flow*, Fig. 6(b)). In the Inlet scenario (Fig. 6(c)), we show how each method performs when one robot must explicitly move out of the other’s way, requiring a high level of coordination. Similarly, in the Track scenario (Fig. 6(d)), we show how each method performs when all robots must move in the same direction around an obstacle, again requiring a high level of coordination, but for larger robot groups.

2) *Warehouse (Fig. 6(e))*: The Warehouse scenario is designed to imitate the motions required to fetch or place items on shelves in a warehouse. The topology creates several parallel narrow passages, and queries are selected such that conflicting

choices of aisles are likely, thus requiring coordination during planning to avoid inter-robot collisions. It includes a width-wise aisle cutting through the middle of the length-wise aisles creating entry/exit points that can be used to avoid collisions. We scale the number of aisles with the number of robots. We consider three variants with progressively doubled aisle widths.

3) *Open Cross (Fig. 7(a))*: We evaluate our approach on a classic open MRMP scenario [3] to demonstrate how our method compares when the workspace is not informative, limiting the benefit of topological skeleton guidance.

4) *Maze Cross (Fig. 7(b))*: In this 3D environment with a narrow maze tunnel, the number of degrees of freedom of each mobile robot is doubled relative to a 2D workspace, resulting in a very large composite space.

C. Narrow Passages

The Hallway, Inlet, and Track scenario results are given in Fig. 8. In the Hallway scenario, CDR-RRT’s use of skeleton guidance allows it to find considerably lower cost paths than other methods in both variants of the problem, especially with the larger 4 and 6-robot groups. Composite RRT was able to solve the 6-robot scenarios but with significantly decreased success rates and higher path costs. MRdRRT and MAPF/C were unable to find solutions within the time limit.

Considering the Inlet scenario, runtimes for CDR-RRT, Composite RRT, Composite PRM, and MAPF/C were similar; however, the solution quality varies greatly between the methods. CDR-RRT has an average path cost of 15.35 s, while Composite RRT, Composite PRM, and MAPF/C have average path costs of 22.31 s, 22.13 s, and 25.30 s respectively. By sampling along skeleton edges, CDR-RRT finds more direct, lower-cost paths. CBS-MP and MRdRRT were unable to find solutions. We also demonstrate our method on physical robots (Fig. 1). We use Turtlebot3s and integrate ROS with our planning library to follow paths generated by CDR-RRT.

In the Track environment, only CDR-RRT is able to find a solution to the 6-robot scenario due to the difficulty of sampling paths where each robot moves in the same direction. CDR-RRT’s use of a MAPF heuristic to find feasible paths over the composite skeleton allows it to efficiently recognize the robots must move either clockwise or counterclockwise. None of the other methods was able to achieve over a 7% success rate on the 4-robot scenario. Decoupled PRM was unable to find a solution for any scenario.

The Warehouse scenarios demonstrate each method’s performance in with varying widths of narrow passages (results in Fig. 9). CDR-RRT achieves the lowest planning time on all scenarios and its use of skeleton guidance allows it to scale to the more complex 4 and 6-robot scenarios, where the performance of other methods significantly degrades. Decoupled PRM was unable to complete any scenario.

D. Scalability

We ran up to 100-robot Track scenarios with an 1800-second time limit for CDR-RRT to measure its performance on larger problems. Fig. 10 shows that CDR-RRT is able to efficiently find paths for very large robot teams when a high level of coordination is required. We also show that CDR-RRT’s MAPF heuristic evaluation maintains scalability.

E. Robot Crossings

The Open Cross and Maze Cross environments evaluate each method's performance in different robot cross scenarios, in which topological guidance provides varying levels of benefit. The Open Cross environment results are in Fig. 11(a). When the topology is not useful, workspace guidance still biases robot paths along skeleton edges, which, in an environment without obstacles, increases the potential for collision relative to Composite RRT and Composite PRM. As a result, CDR-RRT has a higher average planning time than Composite RRT and Composite PRM. This shows that composite skeleton guidance is most effective when there are narrow passages that robots must pass through.

In the Maze Cross scenarios (Fig. 11(b)), CDR-RRT and CBS-MP are the only methods able to plan for 4 robots. Only CDR-RRT is able to plan for 6 robots with 100% success within the time limit (CBS-MP - 33%). Decoupled PRM failed to solve the 2-robot scenario. In 3D environments, the size of the composite C_{space} increases significantly, boosting the impact of composite skeleton guidance.

V. CONCLUSION AND FUTURE WORK

We present Composite Dynamic Region-biased Rapidly-exploring Random Trees, a scalable workspace-guided multi-robot motion planning approach. We validate our method on a variety of environments, with and without narrow passages, to demonstrate its strengths and weaknesses. We show improved performance in constricted environments. Future work will explore the use of composite skeleton guidance for PRM-based roadmap construction, expanding its utility to multi-query scenarios, as well as extending skeleton guidance to non-holonomic robot teams.

REFERENCES

- [1] G. Sanchez and J.-C. Latombe, "Using a PRM planner to compare centralized and decoupled planning for multi-robot systems," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2002, vol. 2, pp. 2112–2119.
- [2] S. M. Lavalle, "Rapidly-exploring random trees: A new tool for path planning," Iowa State Univ., Ames, IA, USA, Tech. Rep. 9811, 1998.
- [3] I. Solis, J. Motes, R. Sandström, and N. M. Amato, "Representation-optimal multi-robot motion planning using conflict-based search," *IEEE Robot. Automat. Lett.*, vol. 6, no. 3, pp. 4608–4615, Jul. 2021.
- [4] G. Wagner and H. Choset, "Subdimensional expansion for multirobot path planning," *Artif. Intell.*, vol. 219, pp. 1–24, 2015.
- [5] J. Denny, R. Sandström, A. Bregger, and N. M. Amato, "Dynamic region-biased rapidly-exploring random trees," in *Algorithmic Foundations of Robotics XII*. Berlin, Germany: Springer, 2020, pp. 640–655.
- [6] R. Sandstrom, D. Uwacu, J. Denny, and N. M. Amato, "Topology-guided roadmap construction with dynamic region sampling," *IEEE Robot. Automat. Lett.*, vol. 5, no. 4, pp. 6161–6168, Oct. 2020.
- [7] M. Rickert, A. Sieverling, and O. Brock, "Balancing exploration and exploitation in sampling-based motion planning," *IEEE Trans. Robot.*, vol. 30, no. 6, pp. 1305–1317, Dec. 2014.
- [8] T. Lozano-Pérez and M. A. Wesley, "An algorithm for planning collision-free paths among polyhedral obstacles," *Commun. ACM*, vol. 22, no. 10, pp. 560–570, Oct. 1979.
- [9] J. T. Schwartz and M. Sharir, "On the 'piano movers' problem. II. General techniques for computing topological properties of real algebraic manifolds," *Adv. Appl. Math.*, vol. 4, no. 3, pp. 298–351, 1983.
- [10] J. F. Canny, *The Complexity of Robot Motion Planning*. Cambridge, MA, USA: MIT Press, 1988.
- [11] L. E. Kavraki, P. Švestka, J. C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE Trans. Robot. Automat.*, vol. 12, no. 4, pp. 566–580, Aug. 1996.
- [12] D. Hsu, J.-C. Latombe, and H. Kurniawati, "On the probabilistic foundations of probabilistic roadmap planning," *Int. J. Robot. Res.*, vol. 25, pp. 627–643, Jul. 2006.
- [13] S. Rodriguez, X. Tang, J.-M. Lien, and N. M. Amato, "An obstacle-based rapidly-exploring random tree," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2006, pp. 895–900.
- [14] A. Yershova, L. Jaillet, T. Simeon, and S. M. Lavalle, "Dynamic-domain RRTs: Efficient exploration by controlling the sampling domain," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2005, pp. 3856–3861.
- [15] 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. Robot. Res.*, vol. 35, no. 5, pp. 501–513, 2016.
- [16] W. Hönig, J. A. Preiss, T. K. S. Kumar, G. S. Sukhatme, and N. Ayanian, "Trajectory planning for quadrotor swarms," *IEEE Trans. Robot.*, vol. 34, no. 4, pp. 856–869, Aug. 2018.
- [17] N. Ryan, "Constraint-based multi-robot path planning," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2010, pp. 922–928.
- [18] J. Yu and D. Rus, "An effective algorithmic framework for near optimal multi-robot path planning," *Robot. Research.*, Cham, Switzerland: Springer, 2018, pp. 495–511, doi: [10.1007/978-3-319-51532-8_30](https://doi.org/10.1007/978-3-319-51532-8_30).
- [19] M. Morales, L. Tapia, R. Pearce, S. Rodriguez, and N. M. Amato, "A machine learning approach for feature-sensitive motion planning," in *Algorithmic Foundations of Robotics VI.*, Berlin, Germany: Springer, 2005, pp. 361–376.
- [20] H. Kurniawati and D. Hsu, "Workspace importance sampling for probabilistic roadmap planning," in *Proc. IEEE Int. Conf. Intell. Robot. Syst.*, Sep. 2004, vol. 2, pp. 1618–1623.
- [21] J. Berg and M. Overmars, "Using workspace information as a guide to non-uniform sampling in probabilistic roadmap planners," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2004, pp. 453–460.
- [22] H. Kurniawati and D. Hsu, "Workspace-based connectivity oracle - An adaptive sampling strategy for PRM planning," in *Algorithmic Foundations of Robotics VII*. Berlin, Germany: Springer, 2008, pp. 35–51.
- [23] E. Plaku, L. Kavraki, and M. Vardi, "Motion planning with dynamics by a synergistic combination of layers of planning," *IEEE Trans. Robot.*, vol. 26, no. 3, pp. 469–482, Jun. 2010.
- [24] J. Denny, R. Sandström, N. Julian, and N. M. Amato, "A region-based strategy for collaborative roadmap construction," in *Algorithmic Foundations of Robotics XI*. Berlin, Germany: Springer, 2015, pp. 125–141.
- [25] S. Bhattacharya, M. Likhachev, and V. Kumar, "Topological constraints in search-based robot path planning," *Auton. Robots*, vol. 33, no. 3, pp. 273–290, 2012.
- [26] H. Blum, "A transformation for extracting new descriptors of shape," in *Models for Perception of Speech and Visual Form*, W. Wathen-Dunn, Ed. Cambridge, MA, USA: MIT Press, 1967, pp. 362–380.
- [27] A. Tagliasacchi, I. Alhashim, M. Olson, and H. Zhang, "Mean curvature skeletons," *Comput. Graph. Forum*, vol. 31, pp. 1735–1744, 2012.
- [28] D. T. Lee, "Medial axis transformation of a planar shape," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-4, no. 4, pp. 363–369, Jul. 1982.
- [29] D. Uwacu, A. Yammanuru, M. Morales, and N. M. Amato, "Hierarchical planning with annotated skeleton guidance," *IEEE Robot. Automat. Lett.*, vol. 7, no. 4, pp. 11055–11061, Oct. 2022.
- [30] 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.
- [31] H. Ma, D. Harabor, P. J. Stuckey, J. Li, and S. Koenig, "Searching with consistent prioritization for multi-agent path finding," in *Proc. 33rd AAAI Conf. Artif. Intell.*, 2019, pp. 7643–7650, doi: [10.1609/aaai.v33i01.33017643](https://doi.org/10.1609/aaai.v33i01.33017643).
- [32] S. LaValle and J. Kuffner, "Randomized kinodynamic planning," in *Proc. Int. Conf. Robot. Autom.*, 1999, vol. 1, pp. 473–479.
- [33] A. Dobson and K. E. Bekris, "Sparse roadmap spanners for asymptotically near-optimal motion planning," *Int. J. Robot. Res.*, vol. 33, no. 1, pp. 18–47, 2013.