

Discrete-Guided Diffusion for Scalable and Safe Multi-Robot Motion Planning

Jinhao Liang¹, Sven Koenig^{2, 3}, Ferdinando Fioretto¹

¹University of Virginia

²University of California, Irvine

³Örebro University

jliang@email.virginia.edu, sven.koenig@uci.edu, fioretto@virginia.edu

Abstract

Multi-Robot Motion Planning (MRMP) involves generating collision-free trajectories for multiple robots operating in a shared continuous workspace. While discrete multi-agent path finding (MAPF) methods are broadly adopted due to their scalability, their coarse discretization severely limits trajectory quality. In contrast, continuous optimization-based planners offer higher-quality paths but suffer from the curse of dimensionality, resulting in poor scalability with respect to the number of robots. This paper tackles the limitations of these two approaches by introducing a novel framework that integrates discrete MAPF solvers with constrained generative diffusion models. The resulting framework, called *Discrete-Guided Diffusion* (DGD), has three key characteristics: (1) it decomposes the original nonconvex MRMP problem into tractable subproblems with convex configuration spaces, (2) it combines discrete MAPF solutions with constrained optimization techniques to guide diffusion models capture complex spatiotemporal dependencies among robots, and (3) it incorporates a lightweight constraint repair mechanism to ensure trajectory feasibility. The proposed method sets a new state-of-the-art performance in large-scale, complex environments, scaling to 100 robots while achieving planning efficiency and high success rates.

Extended version — <https://arxiv.org/abs/2508.20095>

Project page — <https://discrete-guided-diffusion.github.io>

1 Introduction

Multi-Robot Motion Planning (MRMP) is a fundamental problem in robotics that requires generating collision-free trajectories for multiple robots operating in a shared environment. MRMP arises in diverse applications such as automated warehouses, coordinated drone fleets, and autonomous driving. Despite its importance, efficiently solving MRMP in complex environments remains a significant challenge due to the high dimensionality and combinatorial complexity (Yu and LaValle 2013).

To address these challenges, two major paradigms have emerged. Optimization-based methods formulate the problem as a continuous, often nonconvex trajectory optimization problem and are capable of producing smooth, high-

quality paths (Marcucci et al. 2023). However, their scalability is severely limited as the number of robots and obstacles increases. Another body of work considers a discretized version of MRMP, known as Multi-Agent Path Finding (MAPF). It discretizes the space and time, significantly reducing computational complexity (Li, Ruml, and Koenig 2021; Okumura 2024). MAPF algorithms scale to hundreds of robots, but their reliance on discretized grid-based movement and synchronized time steps limits their applicability in continuous and dynamic environments.

To address these limitations, recent work has explored the use of generative models, particularly diffusion models, to learn distributions over trajectories in continuous spaces (Xiao et al. 2022; Carvalho et al. 2023). These models show promising results in single-robot planning as they enable diverse and high-quality trajectory generation. However, their extension to multi-robot settings introduces significant challenges: *Diffusion models must capture complex inter-robot spatiotemporal dependencies while simultaneously avoiding collisions*, a problem that becomes highly intractable as the number of robots and obstacles grows.

Several recent approaches attempt to extend diffusion models to MRMP via gradient-based guidance (Shaoul et al. 2024; Ding et al. 2025). These methods are able to generate high-quality trajectories, but are limited by two key challenges: First, the difficulty of ensuring the feasibility of the generated trajectories with respect to non-collision and kinodynamic constraints. This is challenging since gradient-based guidance cannot natively guarantee global constraint satisfaction. Second, these models struggle in cluttered or dense environments, where the complexity of the configuration space increases. Recently, Liang et al. (2025) proposed an approach to ensure feasibility through repeated projections within the diffusion process. This is a promising solution, but it comes at a high computational overhead, due to the nonconvexity and high dimensionality of constraints.

Contribution. This paper addresses these challenges through a novel integration of discrete MAPF with continuous denoising diffusion. The proposed *Discrete-Guided Diffusion* (DGD), first decomposes the original MRMP problem into a sequence of local subproblems with convex configuration space, each defined by the spatiotemporal structure obtained from the MAPF solution. On each subproblem, a diffusion model is then used to generate trajectories guided

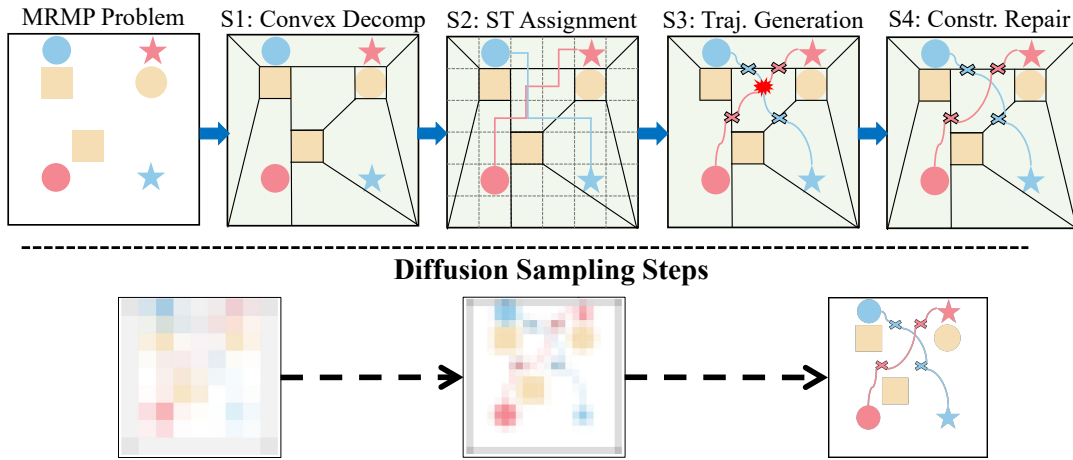


Figure 1: Overview of Discrete-Guided Diffusion. **S1** shows that the nonconvex configuration space is approximated by multiple convex regions. The wheat regions denote the obstacles, and the green regions denote the free space. The robots move from their start \bullet to their goal positions \star . **S2** illustrates that the local start and goal for each robot in each subproblem is determined by the spatiotemporal dependency obtained from MAPF solutions. **S3** shows that the trajectory is generated by diffusion models guided by discrete spatiotemporal dependency. Collision is marked by \star . **S4** shows that the infeasible problem is repaired by constraint-aware diffusion models.

by this spatiotemporal structure. To ensure feasibility, DGD relies on a lightweight projected guidance rendering each step of the diffusion process feasible, and enforced only when constraint violations are detected. The whole framework is schematically illustrated in Figure 1. The effectiveness and efficiency of DGD is demonstrated both theoretically and practically across a range of challenging benchmarks with up to 100 robots.

2 Related Work

Diffusion generative models have recently emerged as a powerful alternative to sampling-based and optimization planners, learning a distribution over continuous trajectories that captures multimodality without hand-tuned costs (Xiao et al. 2022; Carvalho et al. 2023; Luo et al. 2024). While early work focused on single-robot problems, extending to multi-robot motion planning requires enforcing *joint* collision and kinodynamic constraints, which dramatically increases the complexity of the diffusion process.

To address this issue the recent literature has explored three main strategies: (i) *Classifier or gradient guidance* adds an external collision penalty during the reverse process (Ding et al. 2025). Although simple, guidance offers no guarantees on constrained satisfaction and degrades in dense environments. Shaoul et al. (2024) extends this approach by using diffusion models as the single-robot planner in a constraint-based MAPF framework, but this process utilizes diffusion to generate single-robot trajectories without explicitly modeling inter-robot interactions. (ii) *Constrained score estimation* methods learn the score function with added penalty terms (Naderiparizi et al. 2025), but constraints satisfaction still cannot be guaranteed. (iii) Finally, *Projection-based refinement* methods, Liang et al. (2025), interleaves diffusion steps with a nonlinear projection onto

the global feasible set. This approach produces provably collision-free trajectories for a few robots and is therefore central to our discussion and later comparisons. Nonetheless, the need to solve a full high-dimensional projection at every diffusion step incurs substantial computational overhead and limits scalability to larger teams.

To address these limitations, the proposed *Discrete-Guided Diffusion* (DGD) exploits a hybrid approach where a MAPF-derived spatiotemporal skeleton is fused with a constraint-aware diffusion process *inside* locally convex subproblems, avoiding costly global projections while preserving formal feasibility and yielding state-of-the-art performance on large-scale MRMP benchmarks. To the best of our knowledge, DGD is the first to use discrete information to guide diffusion models for scalable and collision-free trajectory generation.

3 Preliminaries

Multi-Robot Motion Planning (MRMP). MRMP involves computing collision-free trajectories for multiple robots navigating a shared environment from designated start positions to goal states. Let $\mathcal{A} = \{a_1, \dots, a_{N_a}\}$ be a set of N_a robots that move in a bounded, planar workspace $\mathcal{W} \subset \mathbb{R}^2$ containing a set of obstacles $\mathcal{O} = \{o_1, \dots, o_{N_o}\}$. The *free configuration space* is defined as the set of all points in the workspace that are not occupied by obstacles: $\mathcal{C}_f := \mathcal{W} \setminus (\bigcup_{o \in \mathcal{O}} \mathcal{B}_r(o))$, where $\mathcal{B}_r(o)$ is the Minkowski sum of obstacle o with a closed ball of radius r , corresponding to the robot’s radius.

Each robot a_i is modeled as a disk of radius $r_i > 0$ and is characterized at discrete time steps $h \in \{0, \dots, H\}$ by a 2-D position $\pi_i^h = (x_i^h, y_i^h) \in \mathcal{C}_f$. A *trajectory* for robot a_i is the sequence $\pi_i := (\pi_i^0, \pi_i^1, \dots, \pi_i^H)$. For each robot, their start and target states are defined by sets $\mathcal{B} = [b_1, b_2, \dots, b_{N_a}]$

and $\mathbf{E} = [e_1, e_2, \dots, e_{N_a}]$.

The goal of MRMP is to compute a set of trajectories $\mathbf{\Pi} = \{\pi_1, \pi_2, \dots, \pi_{N_a}\}$ such that each robot starts at its designated start position and reaches its target position while avoiding collisions with obstacles and other robots. For every time index h we require

$$\pi_i^h \in \mathcal{C}_f, \quad \forall i, \quad (1a)$$

$$\|\pi_i^h - \pi_j^h\|_2 \geq r_i + r_j, \quad \forall i < j, \quad (1b)$$

$$\|\pi_i^{h+1} - \pi_i^h\|_2 \leq v_{\max} \Delta t, \quad \forall i, h < H, \quad (1c)$$

where (1a) enforces obstacle avoidance, (1b) inter-robot separation, and (1c) a first-order kinematic bound with time step Δt and speed limit v_{\max} .

Multi-Agent Path Finding (MAPF). In contrast, Multi-Agent Path Finding assumes the *configuration space and time to be discretized*, which significantly reduces the complexity of the original problem. The workspace is represented by an undirected graph $\mathcal{G} = (V, E)$ whose vertices encode grid cells or roadmap milestones. The input to a MAPF problem is a tuple (\mathcal{G}, S, T) , where $S: [1, \dots, k] \rightarrow V$ maps robots to their source vertices and $T: [1, \dots, k] \rightarrow V$ maps robots to their target vertices. At each integer time step t , robot i occupies a vertex $v_i^t \in V$ and chooses either a wait action or an edge $(v_i^t, v_i^{t+1}) \in E$.

The result of a MAPF algorithm is a set of discrete multi-robot trajectories $\mathbf{\Pi}_M = \{\pi_{m,1}, \pi_{m,2}, \dots, \pi_{m,N_a}\}$, specifying the grid location of each robot at every time step, satisfying vertex conflicts $v_i^t \neq v_j^t \forall i \neq j, t$ and edge conflicts $(v_i^t, v_i^{t+1}) \neq (v_j^{t+1}, v_j^t) \forall i \neq j, t$.

The Discrete-Continuous Gap

The MAPF solution $\mathbf{\Pi}_M$ satisfies the constraints of the discrete problem, but it does not directly translate to a continuous trajectory for the robots. To obtain a continuous trajectory, three key challenges arise:

1. Embedding the vertices V of \mathcal{G} in the continuous space \mathcal{C}_f , which may not be a grid.
2. Time-parameterizing each edge into a collision-free segment that respects constraint (1c). This requires ensuring that the trajectory segments do not violate the kinematic limits of the robots.
3. Although the MAPF solution ensures discrete collision avoidance, it does not exploit constraint (1b) in the continuous space, which can result in suboptimal solutions.

A coarse grid may also lead to unrealistic motions or infeasible timing, whereas a fine grid blows up the state space exponentially, making MAPF itself expensive. To contrast these challenges, this paper proposes a framework that leverages MAPF solely as a *structural guide* while a continuous diffusion model refines the trajectories to respect (1a)–(1c).

4 Discrete-Guided Diffusion Framework

Solving the MRMP problem poses two fundamental challenges. First, MRMP is *NP-hard* and the complexity grows exponentially with the number of robots, making it difficult to scale to large problems while preserving high-quality

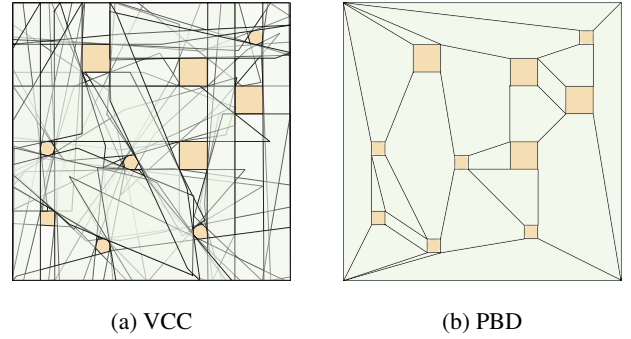


Figure 2: Comparison of Configuration Space Decomposition results between Visibility Clique Cover (VCC) and the proposed priority-based decomposition (PBD). Yellow objects represent obstacles, and each green-filled polygon with a black border indicates a convex region generated by the decomposition methods.

trajectories generation. Second, although generative models offer better scalability and trajectory quality, ensuring the *feasibility* of entire trajectories remains challenging, particularly in high-dimensional planning spaces.

To address these challenges, this section presents an MRMP framework that couples a discrete MAPF backbone with a continuous score-based diffusion model to generate high-quality trajectories for multiple robots. The framework, called *Discrete-Guided Diffusion* (DGD). Specifically, the first two stages mitigate the scalability challenge, and the next two stages address the challenge of constraint satisfaction and trajectory quality in high-dimensional spaces:

- S1 Priority-based Convex Decomposition (PBD):** approximation of the nonconvex free configuration space \mathcal{C}_f using nonoverlapping convex regions to enable problem decomposition (Section 4.1)
- S2 Spatiotemporal assignment:** determination of entry and exit events for each convex region by leveraging spatiotemporal dependencies extracted from a MAPF solution $\mathbf{\Pi}_M$ (Section 4.2).
- S3 Diffusion-based Trajectory Generation:** generation of high-quality trajectories for each subproblem using a diffusion model guided by $\mathbf{\Pi}_M$ (Section 4.3).
- S4 Constraint-aware Diffusion Refinement:** correction of infeasible trajectories through constraint-aware diffusion models by introducing a lightweight projection mechanism to enforce constraint satisfaction (Section 4.4).

Proofs for the theoretical results discussed in this section are provided in the extended version of the paper.

4.1 Priority-Based Convex Decomposition

Recent advancements in single-robot planning replace the original nonconvex configuration space \mathcal{C}_f with a *union of convex polytopes*, enabling single-robot motion planning to be formulated as convex programs (Werner et al. 2024a,b). The current state-of-the-art, the Visibility Clique Cover (VCC) algorithm (Werner et al. 2024a) approximates the free configuration space using a small number of large

convex polytopes. While this eliminates nonconvex obstacle avoidance constraints, the resulting polytopes *overlap* heavily in practice. As shown in Figure 2a, VCC produces substantially more sets than required to cover \mathcal{C}_f , highlighting the inefficiency introduced by overlapping regions. For MRMP, this is problematic for two reasons: **(i)** First, overlap explodes the number of discrete mode switches that must be co-optimized, and **(ii)** Second, a single robot may lie in *several* polytopes at once, making a clean decomposition into independent subproblems impossible. *This motivates us to find more compact and structure-aware representations.*

Priority-based Decomposition (PBD). To address these limitations, we introduce a novel decomposition approach that guarantees non-overlapping convex regions while optimizing for robot traffic patterns. This method builds upon computational geometry techniques, inspired by the Hertel-Mehlhorn algorithm (Hertel and Mehlhorn 1983). It produces a partition $\mathcal{C}_f^c = \{R_1, R_2, \dots, R_k\}$ where $R_i \cap R_j = \emptyset$ for $i \neq j$ and $\bigcup_{i=1}^k R_i \in \mathcal{C}_f$, ensuring each robot belongs to *exactly one region at any time*. This eliminates the critical issue in overlapping methods, where a single point may belong to multiple regions, resulting in larger subproblems due to redundant coverage and making parallelization difficult, as subproblems are no longer independent. The decomposition is performed in two steps:

1. Triangulate \mathcal{C}_f in $O(|\mathcal{V}| \log |\mathcal{V}|)$ time, where \mathcal{V} is a set of all vertices.
2. Iteratively remove diagonals to merge adjacent triangles, but *prioritize* edges that connect regions with the *highest robot traffic* according to the MAPF solution Π_M .

The resulting **Priority-Based Decomposition** keeps the number of regions proportional to where the robots actually need maneuvering, dramatically reducing the downstream dimension required for diffusion, as illustrated in Figure 2b. A detailed presentation of the algorithm is given in the extended version.

Theorem 1 (Sound, compact, and efficient). *PBD outputs a finite set \mathcal{C}_f^c of pairwise disjoint convex polygons satisfying $\bigcup_{R \in \mathcal{C}_f^c} R \in \mathcal{C}_f$. Its runtime is $O(|\mathcal{V}| \log |\mathcal{V}|)$, where \mathcal{V} are the vertices of the triangulation.*

Sketch of Proof. The runtime follows from (1) hole removal and triangulation in $O(|\mathcal{V}| \log |\mathcal{V}|)$ time, and (2) at most $O(|\mathcal{V}|)$ heap-based merge operations, each taking $O(\log |\mathcal{V}|)$ time. \square

Remark 1 (Region count). *Empirically, PBD eliminates up to 32.6% of redundant convex regions across our benchmark scenarios. Compared to the baseline Visibility Clique Cover, it reduces the total number of region sets by over 50%, while also achieving faster runtime. This directly translates to a comparable reduction in the number of diffusion model calls required for trajectory generation.*

4.2 MAPF-Driven Spatiotemporal Assignment

While PBD yields a non-overlapping spatial partition $\mathcal{C}_f^c = \{R_1, \dots, R_k\}$, this alone is insufficient for decomposing MRMP. In multi-robot systems, coordination hinges not

only on spatial separation but also on the *spatiotemporal ordering of robot actions*. Even when robots are assigned to disjoint regions, their trajectories can remain interdependent due to timing constraints. A central challenge lies in determining *when* and *where* each robot should enter or exit specific regions. These decisions introduce temporal coupling between otherwise spatially decoupled subproblems. Without resolving this temporal-spatial assignment, region-based planning cannot guarantee global feasibility.

To address this challenge, this paper makes a key observation: *a discrete MAPF solution $\Pi_M = \{\pi_{m,i}\}_{i=1}^{N_a}$ inherently encodes a coarse spatiotemporal collision schedule*. Thus, rather than leveraging Π_M for its specific paths, we extract and exploit its embedded coordination structure. Although MAPF operates in a discretized space-time domain, it effectively resolves collisions and preserves essential inter-agent dependencies. In particular, for each robot a_i and region R_j , we extract two key events:

- **Exit event:** Time t_{out} and position π_{out} , such that $\pi_{m,i}(t_{\text{out}}) \in R_j$ and $\pi_{m,i}(t_{\text{out}}+1) \notin R_j$.
- **Entry event:** Time t_{in} and position π_{in} , such that $\pi_{m,i}(t_{\text{in}}) \in R_k$ and $\pi_{m,i}(t_{\text{in}}-1) \notin R_k$.

From these, we construct a structured transition set: $\mathcal{T} = \{(a_i, R_j, t_{\text{out}}, \pi_{\text{out}}, R_k, t_{\text{in}}, \pi_{\text{in}})\}$, where each tuple indicates that agent a_i exits region R_j at $(t_{\text{out}}, \pi_{\text{out}})$ and enters region R_k at $(t_{\text{in}}, \pi_{\text{in}})$. These transitions provide start and goal constraints for region-level diffusion models, enabling temporally consistent and spatially decoupled planning.

The transition set \mathcal{T} is constructed by a single forward pass through the MAPF trajectories and the region map (see the extended version for details), as illustrated in **S2** of Figure 1. Its structural utility is formalized as follows.

Proposition 1 (Valid subproblem decomposition). *Fix any transition set \mathcal{T} obtained from Π_M and \mathcal{C}_f^c . Then:*

- (a) *Each robot is uniquely assigned to one region $R \in \mathcal{C}_f^c$ at each time step.*
- (b) *Within each region R , the set of robot trajectories is temporally bounded by entry/exit events in \mathcal{T} .*
- (c) *There are no inter-region constraints at any time step, so each region defines an independent subproblem.*

Sketch of Proof. Disjointness of regions ensures spatial uniqueness. The time-based conditions ensure that each robot’s regional assignment is temporally non-overlapping. Feasibility constraints in MRMP (e.g., inter-robot collision) only apply to robots simultaneously present in the same region. \square

4.3 Diffusion-based Trajectory Generation

By decomposing the original problem into multiple independent subproblems, we are able to operate within smaller, decoupled subspaces. This improves runtime efficiency and sample quality. In this subsection, we describe how diffusion models are used to generate high-quality trajectories within each subproblem, as illustrated in **S3** of Figure 1.

Diffusion models – preliminaries Denoising Diffusion Probabilistic Models (Sohl-Dickstein et al. 2015; Ho, Jain, and Abbeel 2020; Song et al. 2020) define a generative

process by learning to reverse a forward stochastic transformation that progressively corrupts structured data into noise. The generative model then approximates the inverse of this transformation to restore the original structure, allowing sampling from the learned distribution. The diffusion model is thus trained to minimize the difference between the estimated and true scores of the perturbed data. Once trained, the score network s_θ is used to *denoise* random samples $\mathbf{x}_T \sim \mathcal{N}(0, I)$ by iteratively updating \mathbf{x}_t along the *score* direction until the samples resemble the original data. This is also known as the *sampling phase*, which follows the Stochastic Gradient Langevin Dynamics (SGLD) update rule:

$$\mathbf{x}_t = \mathbf{x}_{t+1} + \frac{\epsilon}{2} s_\theta(\mathbf{x}_{t+1}, t+1) + \sqrt{\epsilon} \cdot z, \quad (2)$$

where ϵ denotes the step size and z is standard Normal.

Trajectory Generation using Diffusion Models To facilitate the generation of collision-free trajectories by diffusion models, a straightforward method is to bias the sampling process by incorporating gradient-based guidance to encourage the robot to avoid obstacles and other robots. This can be achieved by adding a penalty term into Eq. (2):

$$\mathbf{x}_t = \mathbf{x}_{t+1} + \frac{\epsilon}{2} (s_\theta(\mathbf{x}_{t+1}, t+1) + \mathcal{J}(\mathbf{x}_{t+1}, \mathcal{O})) + \sqrt{\epsilon} \cdot z,$$

where

$$\mathcal{J}(\mathbf{x}_t, \mathcal{O}) = \sum_{i=1}^{N_a} \nabla_{\mathbf{x}} d_o(\mathbf{x}_t^i, \mathcal{O}) + \sum_{i=1}^{N_a} \nabla_{\mathbf{x}} d_a(\mathbf{x}_t^i, \mathbf{x}_t^{-i}).$$

We define the obstacle and inter-agent penalty functions based on the squared L2-norm distance:

$$d_o(\boldsymbol{\pi}_i, \mathcal{O}) = \sum_{h=0}^H \sum_{o_j \in \mathcal{O}} \max \{0, r_{i,j}^o - \|\boldsymbol{\pi}_i^h - o_j\|_2\},$$

$$d_a(\boldsymbol{\pi}_i, \boldsymbol{\Pi}_{-i}) = \sum_{h=0}^H \sum_{\pi_j \in \boldsymbol{\Pi}_{-i}} \max \{0, r_{i,j}^a - \|\boldsymbol{\pi}_i^h - \pi_j^h\|_2\},$$

where $d_o(\cdot)$ and $d_a(\cdot)$ measure the violation of minimum safety distances $r_{i,j}^o$ and $r_{i,j}^a$ between robot a_i and obstacle o_j , and between robots a_i and a_j , respectively. $\boldsymbol{\Pi}_{-i}$ denotes the set of trajectories of all robots except a_i .

Incorporating the penalty term into the sampling process allows the diffusion model to sample trajectories that follow the learned distribution while softly enforcing task requirements (e.g., collision avoidance). However, the guidance term alone does not guarantee that each trajectory remains within its assigned convex region, which is essential for maintaining independence across subproblems.

Fortunately, since each subproblem is defined over a convex domain, we can efficiently enforce this constraint by projecting each intermediate sample back onto the feasible set. To this end, we integrate a convex projection operator \mathcal{P}_R into the SGLD update rule with minimal computational overhead:

$$\mathbf{x}_t = \mathcal{P}_R(\mathbf{x}_t),$$

where $\mathcal{P}_R(\mathbf{x}) = \arg \min_{\mathbf{y} \in R} \|\mathbf{x} - \mathbf{y}\|_2^2$ denotes the Euclidean projection onto the feasible convex region $R \in \mathcal{C}_f^c$ associated with the specific subproblem.

Theorem 2 (Obstacle avoidance). *Let \mathbf{x}_0 denotes the trajectory generated by diffusion models with projection operator \mathcal{P}_R . For arbitrary small ξ , there exist t such that $\sum_{i=1}^{N_a} \|\nabla_{\mathbf{x}} d_o(\mathbf{x}_0, \mathcal{O})\|_2 \leq \xi$.*

Sketch of Proof. Since R is a convex set, the projection operator \mathcal{P}_R guarantees that the generated trajectory lies within this feasible set (Christopher, Baek, and Fioretto 2024). Moreover, since $\bigcup_{R \in \mathcal{C}_f^c} R \subseteq \mathcal{C}_f$, and \mathcal{C}_f denotes the free configuration space, the generated trajectories satisfy the obstacle avoidance constraints. \square

Additionally, beyond extracting timing structure, the MAPF plan $\boldsymbol{\Pi}_M$ provides a strong *trajectory prior*. Thus, inspired by recent work on Diffusion Models Inversion (Meng et al. 2022; Rout et al. 2025), which shows that even an imperfect or suboptimal initial solution can lead to better performance than sampling from pure noise. Motivated by this, we adopt a similar strategy to improve sampling efficiency and quality. Unlike standard diffusion models that start from Gaussian noise $\mathbf{x}_T \sim \mathcal{N}(0, I)$, we leverage a valid but suboptimal solution from a MAPF solver to initialize the diffusion model’s reverse process $\mathbf{x}_T = \boldsymbol{\Pi}_M$, which provides a strong structural prior for the sampling process.

4.4 Constraint-aware Diffusion Refinement

The previous stages ensure that each trajectory remains within its assigned convex region and satisfies obstacle avoidance constraints. For inter-robot collision avoidance, although multiple guidance methods have been introduced to encourage generating collision-free trajectories, these mechanisms do not provide strict guarantees, especially with a large number of robots.

To address this limitation, we examine each region separately and introduce a *Constraint-aware Diffusion* refinement step to repair the infeasible subproblems, as shown in **S4** of Figure 1. If any constraint in (1a)–(1b) is violated, we run a constraints-aware diffusion model to obtain the feasible solution. Specifically, we introduce another *projection operator* $\mathcal{P}_{\boldsymbol{\Pi}}$ into the SGLD update rule:

$$\mathbf{x}_t = \mathcal{P}_{\boldsymbol{\Pi}}(\mathbf{x}_t),$$

where

$$\mathcal{P}_{\boldsymbol{\Pi}}(\mathbf{x}) = \arg \min_{\mathbf{y}} \|\mathbf{x} - \mathbf{y}\|_2^2 \quad \text{s.t.} \quad \text{Eq. (1b), (1c)}.$$

To solve $\mathcal{P}_{\boldsymbol{\Pi}}$ more efficiently, we use a Lagrangian dual method, where the constraints are incorporated into a relaxed objective by using Lagrange multipliers $\boldsymbol{\nu}$:

$$\mathcal{L}(\mathbf{y}, \boldsymbol{\nu}) = \|\mathbf{x} - \mathbf{y}\|_2^2 + \boldsymbol{\nu} \sum_{i=1}^{N_a} \nabla_{\mathbf{y}} d_a(\mathbf{y}_t^i, \mathbf{y}_t^{-i}),$$

To accelerate the convergence, a quadratic penalty term ρ is introduced, and the augmented Lagrangian function is defined as:

$$\mathcal{L}_{\text{alm}}(\mathbf{y}, \boldsymbol{\nu}) = \mathcal{L}(\mathbf{y}, \boldsymbol{\nu}) + \rho \left\| \sum_{i=1}^{N_a} \nabla_{\mathbf{y}} d_a(\mathbf{y}_t^i, \mathbf{y}_t^{-i}) \right\|^2.$$

A dual ascent strategy is used to solve the relaxed problem, where the primal variable \mathbf{x}' is updated via gradient descent on the augmented Lagrangian $\mathcal{L}_{\text{alm}}(\mathbf{y}, \boldsymbol{\nu})$, and




Figure 3: Examples of benchmark used for MRMP experiments, with increasing complexity. Colorful spheres and plates denote the start and goals of robots. White objects indicate obstacles.

the dual variable ν is updated according to $\nu \leftarrow \nu + \rho \left\| \sum_{i=1}^{N_a} \nabla_{\mathbf{y}} d_a(\mathbf{y}_t^i, \mathbf{y}_t^{-i}) \right\|^2$. The penalty coefficient ρ is gradually increased to enforce the constraints more strictly. This iterative process continues until the residual of the collision avoidance constraints falls below a predefined threshold, at which point the algorithm returns a feasible trajectory.

Remark 2. *Although inter-robot collision avoidance constraints are generally difficult to handle in MRMP due to their nonconvexity and high dimensionality, our approach focuses only on resolving such constraints within infeasible subproblems. These subproblems operate in significantly lower-dimensional spaces compared to the global problem. Moreover, by applying a Lagrangian relaxation, we further simplify the optimization landscape, making the constraint repair process highly efficient in practice.*

5 Experiments

5.1 Experimental Setup

Settings. To evaluate the performance of diffusion-based MRMP algorithms, we design two experimental settings targeting different aspects of the problem. **(1) Feasibility, Efficiency, Quality:** We adopt a diverse MRMP benchmark introduced in (Liang et al. 2025), which includes four representative map types: random maps with varying obstacle densities (basic and dense), and structured maps mimicking real-world environments (room and shelf layouts). For each map type, we consider six robot counts (3, 6, 9, 12, 15, and 18), and evaluate across 25 instances with different configurations. Illustrations of the benchmark maps are included in the Figure 3. **(2) Scalability:** To evaluate the scalability of our method, we use a large-scale map containing over 100 heterogeneous obstacles. This setup supports scenarios with up to 100 robots, providing a challenging setting to examine the method’s ability to generate feasible plans in high-dimensional and cluttered environments.

Evaluation Metrics. Different algorithms are empirically evaluated based on Success rate, Running time, Path Length, and Acceleration. The *Success rate* indicates the proportion of test cases solved without collisions and within the time limit (900 seconds), *Running time* measures the computational efficiency required to generate a solution, *Path Length* and *Acceleration* evaluate the quality of generated trajectories.

Map	Robots	DGD			MPD			MMD			SMD		
		S	T	P	S	T	P	S	T	P	S	T	P
Basic	6	100	12.2	1.21	76	9.1	1.12	100	27.1	1.12	100	254.3	1.10
	12	96	25.0	1.28	52	8.7	1.10	100	56.0	1.12	N/A	N/A	N/A
	18	96	65.7	1.33	8	9.4	1.09	96	86.3	1.13	N/A	N/A	N/A
Dense	6	100	12.3	1.26	4	9.2	1.11	40	36.5	1.15	100	287.3	1.13
	12	100	32.9	1.35	0	N/A	N/A	8	62.5	1.15	N/A	N/A	N/A
	18	76	47.8	1.37	0	N/A	N/A	8	87.1	1.18	N/A	N/A	N/A
Room	6	100	38.9	1.44	0	N/A	N/A	24	55.5	1.16	100	181.0	1.19
	12	100	84.0	1.50	0	N/A	N/A	4	122.6	1.14	N/A	N/A	N/A
	18	100	219.0	1.55	0	N/A	N/A	0	N/A	N/A	N/A	N/A	N/A
Shelf	6	100	64.6	1.30	0	N/A	N/A	32	48.5	1.18	100	274.3	1.16
	12	100	131.5	1.33	0	N/A	N/A	4	95.1	1.22	N/A	N/A	N/A
	18	84	324.8	1.38	0	N/A	N/A	0	N/A	N/A	N/A	N/A	N/A

Table 1: Success rate (S, in percentage), running time (T, in seconds), and path length (P).

Competing Methods. The proposed approach is compared against three SOTA diffusion-based baseline methods:

- Motion Planning Diffusion (MPD):** A state-of-the-art diffusion model for single-robot motion planning (Carvalho et al. 2023), which we extend to the multi-robot setting for comparison.
- Multi-robot Motion Planning Diffusion (MMD):** A method that integrates diffusion models with a classical search-based MAPF algorithm to generate MRMP solutions under soft collision constraints (Shaoul et al. 2024).
- Simultaneous MRMP Diffusion (SMD):** The current SOTA method recently introduced in (Liang et al. 2025). It integrates diffusion models with constrained optimization to generate collision-free trajectories for MRMP.

Implementation details are available in the extended version.

5.2 Comparison across Methods

We evaluate the performance of all methods on the standard MRMP benchmark, comparing success rates, runtimes, and path lengths across scenarios of increasing complexity and robot count (6, 12, and 18 agents). Full results are summarized in Table 1, and additional metrics and ablations are provided in the extended version.

Baseline Performance

Motion Planning Diffusion (MPD) achieves the lowest runtimes and path lengths when successful. This is the case for the simplest problems (Basic maps with 6 robots). However, MPD fails to scale, consistently failing to generate feasible solutions beyond three robots in any more complex scenario. This underscores the limitations of unguided diffusion in multi-agent settings with tight coupling.

Multi-robot Motion Planning Diffusion (MMD) improves substantially over MPD by jointly modeling all agents. It maintains high success rates on Basic maps with up to 18 robots. However, its performance sharply degrades in more constrained environments. For instance, on Room and Shelf maps with 12 or more agents, MMD fails consistently. This

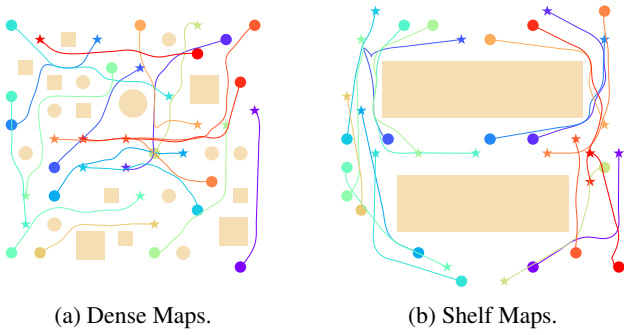


Figure 4: Trajectories generated by DGD.

breakdown is attributed to its lack of explicit coordination mechanisms, which leads to congestion and deadlocks in narrow passages.

Simultaneous MRMP Diffusion (SMD) shows high-quality planning with 100% success rates for up to 6 robots. However, its computational cost is orders of magnitude higher than other methods. For 12 or more robots, the running time becomes prohibitively long, and no solution is obtained within the imposed time limit.

Discrete Guided Diffusion (DGD)

DGD outperforms all baselines in both success rate and computational efficiency. By decomposing the MRMP problem into multiple region-level subproblems using a structured MAPF schedule, DGD solves each subproblem independently and in parallel, significantly reducing planning overhead. DGD attains over 92% success rate across all settings except the most challenging (Dense and Shelf maps with 18 robots), where it still leads with a success rate of 76%. Additionally, it achieves such results with lowest runtimes for the non-trivial benchmarks. Notably, on Dense maps with 6 robots, DGD matches SMD’s perfect success rate while using only 4% of its runtime.

DGD’s path lengths are comparable to those of MPD and MMD. While the average path length appears slightly higher, this is primarily because the averages are computed only over successful trials. In more challenging scenarios, where MPD and MMD often fail, DGD still succeeds, albeit with longer paths due to the complexity of the environment. Example trajectories generated by DGD are shown in Figure 4 and the corresponding figures for the baselines are available in the extended version.

5.3 Scalability Analysis

To evaluate the scalability of DGD, we test it on large-scale environments containing a diverse mix of obstacle types within a single map. These environments feature **104** obstacles and **100** robots, substantially exceeding the complexity and scale of standard MRMP benchmarks.

As in earlier experiments, DGD decomposes the global MRMP problem into a set of independent subproblems, each confined to a convex region. This decomposition enables efficient parallelized planning, and allows DGD to scale gracefully to high-dimensional instances. Representative trajec-




Figure 5: Trajectories generated by DGD on large maps. The figure is horizontally cropped for readability, with a dashed line on the bottom indicating continuation. The complete figure is provided in the extended version.

ries generated by DGD in this large-scale setting are shown in Figure 5 (figure cropped due to large size).

To the best of our knowledge, DGD is the first diffusion-based MRMP method capable of generating feasible solutions in environments with over 100 obstacles and 100 robots. In contrast to prior diffusion approaches, which have been demonstrated only on up to 40 robots in obstacle-free settings (Shaoul et al. 2024), *DGD achieves a 2.5× improvement in robot count while simultaneously addressing significantly more complex environments*. This is significant: scalability is a key impediment to apply generative models to real-world multi-robot systems, where both geometric complexity and agent coordination pose significant challenges, and thus, this work makes a significant step towards practical diffusion-based solvers for multi-robot planning.

6 Conclusion

This work introduces *Discrete Guided Diffusion (DGD)*, a method that leverages a decomposition strategy to enable diffusion models to efficiently generate collision-free trajectories for multi-robot systems. By decomposing the original MRMP problem into multiple tractable subproblems, DGD alleviates the significant computational burden. To address the spatiotemporal dependencies across multiple robots, MAPF solutions are employed to guide diffusion models in generating solutions for each subproblem. In addition, an efficient constraint repair mechanism ensures the feasibility of the generated trajectories.

Extensive experiments across varying obstacle densities and increasing robot counts demonstrate that DGD achieves significantly higher success rates and lower running times than competing methods. Experiments on large-scale maps further validate its scalability, making it a promising approach for real-world applications.

While DGD achieves remarkable performance on MRMP tasks, there are several promising directions to further enhance its capabilities. One opportunity lies in improving the decomposition process to produce more flexible subproblem partitions. In addition, extending the current framework to incorporate interactions between adjacent subregions could help better preserve consistency across the full trajectory.

Acknowledgments

This research at the University of Virginia was partially supported by NSF grants 2334936, 2334448, NSF CAREER Award 2401285, and DARPA Contract No. #HR0011252E005. The research at the University of California, Irvine was supported by NSF grants 2544613, 2434916, 2321786, 2112533, as well as gifts from Amazon Robotics and the Donald Bren Foundation. The authors acknowledge Research Computing at the University of Virginia for providing computational resources that have contributed to the results reported within this paper. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the sponsoring organizations, agencies, or the U.S. government.

References

- Carvalho, J.; Le, A. T.; Baierl, M.; Koert, D.; and Peters, J. 2023. Motion planning diffusion: Learning and planning of robot motions with diffusion models. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1916–1923. IEEE.
- Christopher, J. K.; Baek, S.; and Fioretto, N. 2024. Constrained synthesis with projected diffusion models. *Advances in Neural Information Processing Systems*, 37: 89307–89333.
- Ding, K.; Jiao, C.; Hu, Y.; Zhou, K.; Wu, P.; Mu, Y.; and Liu, C. 2025. SwarmDiff: Swarm Robotic Trajectory Planning in Cluttered Environments via Diffusion Transformer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 4164–4173.
- Hertel, S.; and Mehlhorn, K. 1983. Fast triangulation of simple polygons. In *Foundations of Computation Theory: Proceedings of the 1983 International FCT-Conference Borgholm, Sweden, August 21–27, 1983 4*, 207–218. Springer.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Li, J.; Ruml, W.; and Koenig, S. 2021. Eecbs: A bounded-suboptimal search for multi-agent path finding. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 12353–12362.
- Liang, J.; Christopher, J. K.; Koenig, S.; and Fioretto, F. 2025. Simultaneous Multi-Robot Motion Planning with Projected Diffusion Models. *arXiv preprint arXiv:2502.03607*.
- Luo, Y.; Sun, C.; Tenenbaum, J. B.; and Du, Y. 2024. Potential Based Diffusion Motion Planning. In *Forty-first International Conference on Machine Learning*.
- Marcucci, T.; Petersen, M.; von Wrangel, D.; and Tedrake, R. 2023. Motion planning around obstacles with convex optimization. *Science Robotics*, 8(84): eadf7843.
- Meng, C.; He, Y.; Song, Y.; Song, J.; Wu, J.; Zhu, J.-Y.; and Ermon, S. 2022. SDEdit: Guided Image Synthesis and Editing with Stochastic Differential Equations. In *International Conference on Learning Representations*.
- Naderiparizi, S.; Liang, X.; Zwartsenberg, B.; and Wood, F. 2025. Constrained Generative Modeling with Manually Bridged Diffusion Models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 19607–19615.
- Okumura, K. 2024. Engineering LaCAM*: Towards Real-time, Large-scale, and Near-optimal Multi-agent Pathfinding. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, 1501–1509.
- Rout, L.; Chen, Y.; Ruiz, N.; Caramanis, C.; Shakkottai, S.; and Chu, W. 2025. Semantic Image Inversion and Editing using Rectified Stochastic Differential Equations. In *The Thirteenth International Conference on Learning Representations*.
- Shaoul, Y.; Mishani, I.; Vats, S.; Li, J.; and Likhachev, M. 2024. Multi-Robot Motion Planning with Diffusion Models. *arXiv preprint arXiv:2410.03072*.
- Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, 2256–2265. pmlr.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*.
- Werner, P.; Amice, A.; Marcucci, T.; Rus, D.; and Tedrake, R. 2024a. Approximating robot configuration spaces with few convex sets using clique covers of visibility graphs. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 10359–10365. IEEE.
- Werner, P.; Cohn, T.; Jiang, R. H.; Seyde, T.; Simchowitz, M.; Tedrake, R.; and Rus, D. 2024b. Faster Algorithms for Growing Collision-Free Convex Polytopes in Robot Configuration Space. *arXiv preprint arXiv:2410.12649*.
- Xiao, X.; Liu, B.; Warnell, G.; and Stone, P. 2022. Motion planning and control for mobile robot navigation using machine learning: a survey. *Autonomous Robots*, 46(5): 569–597.
- Yu, J.; and LaValle, S. 2013. Structure and intractability of optimal multi-robot path planning on graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 27, 1443–1449.