

GRACE: A Unified 2D Multi-Robot Path Planning Simulator & Benchmark for Grid, Roadmap, And Continuous Environments

Chuanlong Zang^{1,2}, Anna Mannucci¹, Isabelle Barz¹, Philipp Schillinger¹, Florian Lier¹, Wolfgang Hönig²

Abstract—Advancing Multi-Agent Pathfinding (MAPF) and Multi-Robot Motion Planning (MRMP) requires platforms that enable transparent, reproducible comparisons across modeling choices. Existing tools either scale under simplifying assumptions (grids, homogeneous agents) or offer higher fidelity with less comparable instrumentation. We present *GRACE*, a unified 2D simulator+benchmark that instantiates the same task at multiple abstraction levels (grid, roadmap, continuous) via explicit, reproducible operators and a common evaluation protocol. Our empirical results on public maps and representative planners enable commensurate comparisons on a shared instance set. Furthermore, we quantify the expected representation-fidelity trade-offs (MRMP solves instances at higher fidelity but lower speed, while grid/roadmap planners scale farther). By consolidating representation, execution, and evaluation, *GRACE* thereby aims to make cross-representation studies more comparable and provides a means to advance multi-robot planning research and its translation to practice.

I. INTRODUCTION

As autonomous systems become increasingly prevalent, coordinating multiple robots in shared environments is critical for unlocking their full potential. This challenge centers on finding collision-free paths or trajectories for agents in a known environment, while optimizing objectives such as makespan or sum-of-costs (SoC). This task is addressed by research areas including Multi-Agent Pathfinding (MAPF) and Multi-Robot Motion Planning (MRMP) [1]. Typically, MAPF has focused on discrete-time grid-based formulations with homogeneous agents, allowing scalability and algorithmic guarantees [2], [3]. MRMP operates in continuous space-time with kinematic/kinodynamic constraints and realistic footprints, capturing operational fidelity [4], [5].

Both views are valuable: discrete abstractions enable rapid design-space exploration and algorithmic benchmarking, while continuous models expose execution-level feasibility and performance limits. To advance research and ensure reliable deployment, fair and reproducible comparisons across these representations are crucial.

The current ecosystem fragments this goal. Grid-centric MAPF simulators and benchmarks excel in large-scale comparisons but often omit agent dynamics, continuous time, and heterogeneity, limiting conclusions about real robots [6]. In contrast, general-purpose robotics simulators provide high fidelity, but make cross-representation and cross-planner benchmarking cumbersome: identical tasks are hard to express at multiple abstraction levels; and instrumentation




Fig. 1. **GRACE: Single scenario, multiple planning representations.** Identical agents/scenarios, distinct trajectories from representation-specific constraints: (a) occupancy-grid (4-connected discrete actions), (b) graph abstraction (single integrator), and (c) continuous space (double integrator, heterogeneous footprint). Individual path lengths highlight these representation-specific constraints.

is inconsistent between methods [7]. This fragmentation complicates answering basic questions such as: *When is a roadmap abstraction good enough? When do roadmap or continuous models materially change conclusions? How do MAPF and MRMP compare on an equal footing under the same tasks and metrics?*

This paper introduces *GRACE*, a unified 2D simulation and benchmarking platform designed to address this critical gap. *GRACE* uniquely supports grid, roadmap, and continuous environment models within a single API and evaluation protocol, enabling consistent evaluation across abstraction levels (as demonstrated in Fig. 1). Its flexible design allows seamless integration and comparative analysis of diverse MAPF and MRMP planners (classical, learning-based, or hybrid) under standardized scenario definitions and metrics.

Contributions. 1) *GRACE*, a unified simulator and benchmark for grid/roadmap/continuous planning, which systematically defines and implements explicit abstraction operators and cross-representation conversion mechanisms to ensure reproducible abstractions and standardized evaluation; and 2) an empirical study using *GRACE* that (i) compares representations on shared instance set, (ii) benchmarks planner families on grids, and (iii) reports determinism and observed scalability (up to $\sim 2k$ agents in our setup), highlighting representation-fidelity trade-offs.

II. RELATED WORK

This section reviews efforts to solve, simulate and benchmark MAPF and MRMP across grid, roadmap, and continuous representations. We first discuss existing benchmarks and platforms (Sec. II-A), followed by a survey of planner families across those representations (Sec. II-B), as algorithmic assumptions critically influence fair benchmark evaluation.

¹Robert Bosch GmbH, Corporate Research, Stuttgart, Germany. Chuanlong.Zang@de.bosch.com

²Technical University of Berlin, Berlin, Germany.

Code: <https://github.com/boschresearch/GRACE>

TABLE I

COMPARISON OF 2D MULTI-AGENT PATH PLANNING PLATFORMS.

Platform	Rep.	CST Env.	CST Plan.	Dyn.	Het.	Scale
MovingAI-related [8], [9], [10], [11]	G	×	×	×	×	1k
asprilo [12]	G	×	×	×	×	50
Flatland [13]	G	×	×	×	×	6k
LoRR [14]	G	×	×	×	×	10k
POGEMA [6]	G	×	×	×	×	1M
MRP-Bench [7]	C	✓	×	×	×	10
mrmg-planning [15]	C	✓	✓	✓	✓	10
REMROC [16]	G+C	✓	✓	✓	×	10
SMART [17]	G,C	✓	×	✓	×	2k
GRACE (Ours)	G,R,C	✓	✓	✓	✓	2k

Rep. \in {G=grid,R=roadmap,C=continuous}: see Sec. III-A; G+C: hierarchical global + local planner. **CST Env./Plan.**=continuous space-time environment and planning; **Dyn.**=agent dynamics; **Het.**=heterogeneous agents; **Scale**=max number of agents in one experiment.

A. Benchmarks and Platforms in MAPF and MRMP

Most MAPF benchmarks traditionally evaluate search-based [10], [11], [34], learning-based [35], and reduction-based algorithms [9], [36] in 2D grid environments with discrete time and uniform agents. Such benchmarks popularized by the MovingAI community [8] similarly emphasize grid-based MAPF, enabling standardized evaluation with simplified assumptions.

Recent work, including asprilo [12], SkyRover [37], MRP-Bench [7], REMROC [16], has moved towards more expressive settings. These include intralogistic domains, real-life robotics use cases, human-shared environments, cross-domain coordination and platforms with agent size or kinematic constraints [38]. As a realistic testbed in a continuous 2D environment, SMART [17] leverages the robust discretized planning of the action dependency graph [39].

However, many benchmarks still assume grid-based path planners and environments. To address this, [40] introduces MAPF with generalized Conflicts (MAPF/C), which overcomes the grid limit by defining conflicts in continuous space. [24] proposes a continuous-time MAPF benchmark that works on both grids and roadmaps. These capture complexities such as continuous travel times, heterogeneous agents, and arbitrarily placed objects. In kinodynamic MRMP, more realistic comparisons are offered, although often under different assumptions [31]. In other related domains, such as multi-goal multi-robot path planning in continuous space, mrmg-planning [15] formulates a single robot path planning approach to solve multi-robot problem, facing scalability issue. Our work builds on these efforts to support the rigorous evaluation of diverse agents across three environmental representations: grid-based environments, roadmap graphs, and shared continuous spaces.

TABLE II

COMPARISON OF 2D MULTI-AGENT PATH PLANNING APPROACHES.

Planners	Rep.	Dyn.	Het.	Opt. [†]	Comp.	OSS	Eval.
CBSH2-RTC [18]	G	×	×	opt.	✓	✓	✓
EECBS [19]	G	×	×	subopt.	✓	✓	✓
LaCAM [2]	G	×	×	subopt.	✓	✓	✓
MAPF-LNS2 [20]	G	×	×	subopt.	✓	✓	✓
SCRIMP [21]	G	×	×	subopt.	✓	✓	✓
MAPF-GPT-DDG [22]	G	×	×	subopt.	✓	✓	✓
CBS-MP [23]	R	×	✓	repr.	prob.	×	×
CCBS [24]	R	×	×	*	*	✓	✓
PSIPP/CTC [25]	R	×	×	subopt.	×	✓	✓
CCBS+DK [26]	R	×	×	*	*	×	×
SST* [27]	C	✓	✓	asym.	prob.	✓	✓
MRRP [28]	C	×	✓	subopt.	×	✓	×
K-CBS [29]	C	✓	✓	-	prob.	✓	×
SSSP [30]	C	×	✓	-	prob.	✓	×
db-CBS [31]	C	✓	✓	asym.	prob.	✓	✓
(E)CHPBS [32]	C	✓	✓	subopt.	×	×	×

[†] All listed MAPF planners optimize the sum-of-costs (SoC).

Dyn.=agent dynamics; **Het.**=heterogeneous agents; **Opt.**=optimality; **Comp.**=completeness; **OSS**=open source; **Eval.**=evaluated in GRACE; **repr.**=representation(optimal); **prob.**=probabilistically(complete); **asym.**=asymptotically(optimal); *: ongoing discussions [33]; -: no proof.

Beyond static benchmarks, community challenges and competitions have further driven advancements in the research field. The League of Robot Runners (LoRR) provides standardized MAPF competitions with new domains (e.g., fulfillment, sortation) [14]. Flatland targets large-scale railway traffic on grid graphs [13]. At the system level, RoboCup Logistics League (RCLL) tackles smart-factory intralogistics with real robots, evaluating far beyond 2D planning alone [41]. GRACE complements these efforts by isolating and unifying the multi-robot planning layer across grid/roadmap/continuous representations, while remaining compatible with grid-style challenge maps (e.g., LoRR/Flatland) and clearly scoping out hardware-centric aspects of RCLL.

We summarize the most widely used MAPF and MRMP simulation benchmarks in Tab. I.

B. Planner Families Across Representations

A detailed comparison of 2D MAPF and MRMP solvers is shown in Tab. II. A concise summary of the main planning families is reported below.

1) *Grid-centric MAPF (Rep-G)*: On occupancy grids with discrete time, MAPF planners extensively exploit graph structure and uniform moves. Conflict-Based optimal and suboptimal Search (CBS) minimize SoC under vertex / edge conflict constraints [42], [43], [19], [44]. For scalability to larger problem instances, methods like Large-Nearhood Search (LNS) and related metaheuristics trade optimality for

computational efficiency [20], [45]. Priority-based schemes offer strong speed with weaker optimality and completeness guarantees [46], [47]. Recently, learning-driven approaches, including decentralized MARL policies [48], [49], and search guided by learned heuristics [50], aim to improve responsiveness in dense regimes. Furthermore, hybrid models combine learned proposals with classical repair [51].

2) *Roadmap-based MAPF (Rep-R)*: When continuous environments are abstracted to directed graphs, timing becomes non-uniform and continuous. Safe Interval Path Planning (SIPP) [52] underpins Continuous Conflict-Based Search (CCBS) variants [24] and prioritized formulations with pre-computed conflict intervals, scaling to large teams [25]. Most roadmap methods assume disc robots with single-integrator kinematics, enabling richer timing than grids without full kinodynamics.

3) *Continuous-space MAMP/MRMP (Rep-C)*: Multi-Agent Motion Planning (MAMP) generalizes MAPF to continuous spaces and agent dynamics. Hybrid frameworks bridge MAPF and MAMP: MAPF-POST [53] refines discrete plans into temporally coordinated trajectories. Other approaches combine a discrete search with continuous sampling [30]. Fully continuous MRMP operates directly in continuous state–time with footprints and dynamics. Sampling-based planners address feasibility under differential constraints [54], [55]. Recent kinodynamic MRMP [56], [31] combines discrete conflict reasoning with dynamic feasibility checks, offering a logical middle ground.

III. PROBLEM FORMULATION

This section formalizes the multi-agent planning problem within GRACE, outlining its foundational environment, agent, and planner models (Sec.III-A–III-C). The goal is to find feasible (Sec.III-D) per-agent plans from specified start-goal pairs, adhering to these models. A summarizing table (Table III) at the end of this section provides an overview of key relationships, with all symbols and specialized terminology formally defined in Sec.III-A–III-F.

A. Environment Model and Agent’s Goals

Let the workspace be a closed, bounded 2D region $\Omega \subset \mathbb{R}^2$ with polygonal obstacles $\mathcal{O} \subset \Omega$. The free space is defined as $X_{\text{free}} = \Omega \setminus \mathcal{O}$. GRACE employs abstraction operators $\Phi_{\text{env}}^\rho(\Omega, \mathcal{O})$, $\rho \in \{\text{cont}, \text{road}, \text{grid}\}$, to transform the detailed workspace into simplified representations. These operators produce three levels of abstraction:

- **Continuous**: The environment remains X_{free} (“as is”); goals (g_i herein) are locations in continuous space. Thus, $\Phi_{\text{env}}^{\text{cont}}(\Omega, \mathcal{O}) = X_{\text{free}}$.
- **Roadmap**: A directed graph $G_r = (V_r, E_r)$ is extracted from X_{free} ; goals are mapped to vertices in V_r . Thus, $\Phi_{\text{env}}^{\text{road}}(\Omega, \mathcal{O}) = G_r$.
- **Grid**: An M -connected occupancy grid $G_g = (V_g, E_g)$ discretizes Ω ; goals snap to containing cells. Thus, $\Phi_{\text{env}}^{\text{grid}}(\Omega, \mathcal{O}) = G_g$.

B. Agent Model

Let N be the total number of agents. Each agent $\{i\}_1^N$ has a configuration space $Q_i \in \{\mathbb{R}^2, \text{SE}(2)\}$, corresponding to translation-only or full planar rigid-body pose (with $\text{SE}(2) = \mathbb{R}^2 \times \text{SO}(2)$). Its properties include a start location $s_i \in Q_i$, a convex footprint $S_i \subset \mathbb{R}^2$, and an action model specifying its discrete, kinematic or kinodynamic motion profile f_i .

Let $B(c, r) = \{z \in \mathbb{R}^2 : \|z - c\| \leq r\}$ be the closed Euclidean ball with center c and radius r . For a compact robot footprint S_i (in its body frame), we define the minimum enclosing circle

$$c_{\text{eff}}(S_i), r_{\text{eff}}(S_i) := \underset{c \in \mathbb{R}^2}{\text{argmin}} \max_{z \in S_i} \|z - c\|$$

with center c_{eff} and radius r_{eff} . To compare heterogeneous footprints with a single-disc model, we define the minimal common effective radius $r_{\text{com}} := \max_{i=1, \dots, N} r_{\text{eff}}(S_i)$, that is, the smallest r such that $S_i \subseteq B(0, r)$ for all agents.

We take $c_{\text{eff}}(S_i)$ as the body-fixed reference point and denote its world position by $p_i(t) \in \mathbb{R}^2$ (so if $Q_i = \text{SE}(2)$, the agent’s state is $x_i(t) = (p_i(t), \theta_i(t))$).

- **Continuous**: Agents retain their arbitrary convex footprints S_i and heterogeneous motion profile f_i . Thus, we define $\Phi_{\text{ag}}^{\text{cont}}$ as the identity (no abstraction).
- **Roadmap**: For comparability, agents are abstracted to r_{com} with stop-go single-integrator: $\Phi_{\text{ag}}^{\text{road}}(S_i, f_i) = (r_{\text{com}}, \dot{x}_i = u_i, \|u_i^{\text{linear}}\| \in \{0, \nu\})$, where $\nu > 0$ is the common linear speed.
- **Grid**: Agents are abstracted to homogeneous, tile-sized discs of radius r_{com} with motion defined by M -connected moves with discrete unit time steps: $\Phi_{\text{ag}}^{\text{grid}}(S_i, f_i) = (\text{cell}, \mathcal{A}_{M\text{-conn}} \cup \{(0, 0)\}, \Delta t)$. Here, $\mathcal{A}_{M\text{-conn}} \subset \mathbb{Z}^2$ denotes the allowed neighbor offsets (e.g., for 4-connectivity), $(0, 0)$ encodes the “wait” action, and $\Delta t > 0$ is the discrete time step (often 1). The grid cell side length $a(M)$ is determined by the maximum displacement allowed by M -connectivity $L_{\text{max}}(M) := \max_{d \in \mathcal{A}_{M\text{-conn}}} \|d\|_2$ such that $a(M) = 2L_{\text{max}}(M)r_{\text{com}}$. For example, $a(4) = 2r_{\text{com}}$ and $a(8) = 2\sqrt{2}r_{\text{com}}$.

C. Planner Model

Let \mathcal{P}_ρ be a registry of multi-agent planners compatible with representation $\rho \in \{\text{cont}, \text{road}, \text{grid}\}$. Given an environment representation, agent abstractions, and start–goal pairs $\{(s_i, g_i)\}_{i=1}^N$, if successful, a planner $P \in \mathcal{P}_\rho$ returns plans $\Pi = \{\pi_i\}_{i=1}^N$ for each agent:

- **Continuous**: π_i is a (kino)dynamically feasible trajectory (into X_{free} in its position component).
- **Roadmap**: π_i is a time-parameterized path along nodes/edges of G_r (continuous time).
- **Grid**: π_i is a sequence of time-stamped cell moves on G_g (discrete time).

D. Feasibility

A plan $\pi_i : [0, T_i] \rightarrow Q_i$ with a finite horizon T_i , is considered feasible if it satisfies all constraints related to

TABLE III
FEATURE MAPPING.

Aspect	Continuous	Roadmap	Grid
Time	continuous	continuous	discretized
Footprint	S_i (convex)	disc r_{com}	disc r_{com}
Orient.	if $Q_i = \text{SE}(2)$	—	—
Kinematics	f_i	SI(ν)	unit moves
Collision	swept S_i	cap/edges/ints	MAPF rules
Goals	invariant set	remain at node	remain at cell

SI=single-integrator; cap=capacity; ints=geometric intersections. Refer to Sec. III for definitions of concepts and symbols.

obstacle avoidance and inter-agent collision prevention, as defined by the chosen representation.

- **Continuous:** Let $R(\theta)$ be the planar rotation by angle θ , and let \oplus be the Minkowski sum operator. The occupied region of agent i at time t is

$$X_i(t) = \begin{cases} p_i(t) \oplus S_i, & \text{if } Q_i = \mathbb{R}^2, \\ p_i(t) \oplus (R(\theta_i(t)) S_i), & \text{if } Q_i = \text{SE}(2). \end{cases}$$

A plan is feasible iff $X_i(t) \subseteq X_{\text{free}}$ for all i and $\forall t \in [0, T_i]$ and $X_i(t) \cap X_j(t) = \emptyset$ for all $i \neq j$ and $\forall t$; additionally, π_i satisfies f_i .

- **Roadmap:** We enforce space–time feasibility for discs of radius r_{com} . Feasibility requires $\|p_i(t) - p_j(t)\| \geq 2r_{\text{com}}$ for all t , enforced via (i) vertex capacity 1, (ii) no head-on edge swaps, (iii) same-edge headway $\geq 2r_{\text{com}}/\nu$, and (iv) treating geometric edge intersections as virtual vertices with capacity 1.
- **Grid:** With cell side $a(M) = 2L_{\text{max}}(M)r_{\text{com}}$ and linear interpolation over Δt , we forbid co-occupancy, opposing edge swaps [1], diagonal crossings, and corner-meets; these imply $\|p_i(t) - p_j(t)\| \geq 2r_{\text{com}}$ throughout each unit move.

Let Π_{feasible} denote the subset of feasible plans.

E. Objectives

While feasibility is paramount, GRACE also supports optimization for common objectives. These include, but are not limited to, makespan and SoC:

$$\min_{\Pi_{\text{feasible}}} J(\Pi), \text{ with } J \in \{\text{makespan}, \text{SoC}, \dots\}.$$

The framework does not fix J ; metrics are computed consistently across representations.

F. Assumptions and Scope

The problem formulation operates under the following assumptions: environments are static and fully known; tasks are single-shot with fixed start-goal pairs. All agent shapes are fixed (no articulated robots) and convex. Motion is constrained to a planar domain; orientation is included where relevant. A parked-at-goal model is used: upon arrival at time T_i , agents remain at their goal for all $t \geq T_i$. The choice of representation is user-controlled; GRACE provides reproducible mappings and consistent instrumentation for all abstraction levels.

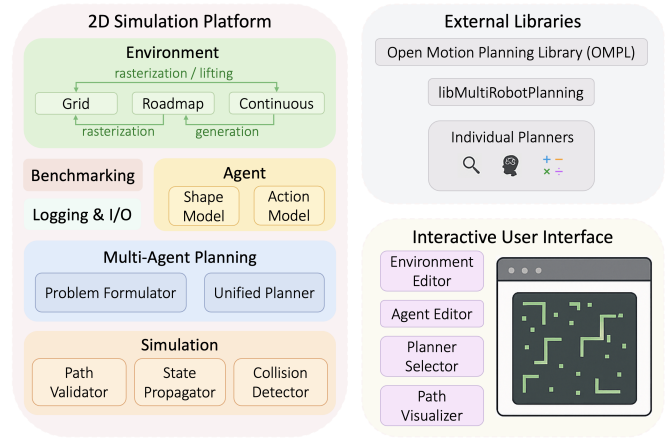


Fig. 2. Architectural Diagram of GRACE.

G. Problem statement

Given a base world $(\Omega, \mathcal{O}, \{S_i, f_i, (s_i, g_i)\}_{i=1}^N)$ and a representation $\rho \in \{\text{cont}, \text{road}, \text{grid}\}$, GRACE instantiates a planner-ready problem via Φ_{env}^ρ and Φ_{ag}^ρ , then re-embeds returned plans to compute common metrics. Cross-representation comparisons are restricted to tasks that admit multiple ρ (the overlapping regime); outside this regime we report representation-specific trade-offs rather than excluding such cases.

IV. GRACE ARCHITECTURE

Fig. 2 provides an overview of GRACE’s architecture: a C++20 core based on Box2D [57] for deterministic, continuous-time simulation, an interactive UI (SFML + ImGui), and external planners connected via a unified interface. A single physical scenario can be realized at multiple abstraction levels, planned by different algorithms, and executed under one simulator.

A. Representation Operators and Consistency

We implement the environment and agent abstractions from Sec. III-A and III-B with explicit, reproducible parameters, so cross-representation studies do not require re-authoring.

Continuous \rightarrow *Roadmap*. Given $X_{\text{free}} = \Phi_{\text{env}}^{\text{cont}}(\Omega, \mathcal{O})$, we construct a directed roadmap $G_r = (V_r, E_r)$ using OMPL [58] and pin starts/goals $\{(s_i, g_i)\}$ to V_r . Edge weights store Euclidean length ℓ_e ; with the common speed $v > 0$ from $\Phi_{\text{ag}}^{\text{road}}$, traversal times are $\tau_e = \ell_e/v$.

Continuous \rightarrow *Grid*. We rasterize Ω with cell side $a(M) = 2L_{\text{max}}(M)r_{\text{com}}$. Occupancy is computed by polygon coverage of obstacles. We use $\Delta t = 1$.

Roadmap \rightarrow *Grid*. We rasterize edges of G_r onto the grid overlay to obtain a graph-constrained grid that preserves roadmap topology and enables discrete-time scheduling.

Grid \rightarrow *Continuous*. We embed polylines through cell centers in world coordinates and reconstruct obstacles directly from the occupancy mask.

Roadmap \rightarrow *Continuous*. We inflate obstacles per robot by the disk surrogate $B(0, r_{\text{eff}}(S_i))$:

$$\mathcal{O}^{(i)} = \mathcal{O} \oplus B(0, r_{\text{eff}}(S_i)), \quad X_{\text{free}}^{(i)} = \Omega \setminus \mathcal{O}^{(i)}.$$

Given polyline vertices $(p_k)_{k=0}^K$ defining robot i 's path on the roadmap, each path segment is validated obstacle-free passage via the swept set:

$$[p_k, p_{k+1}] \oplus B(0, r_{\text{eff}}(S_i)) \subseteq X_{\text{free}}^{(i)}.$$

B. Multi-agent Planning Pipeline

Problem Formulator. Given $\rho \in \{\text{cont, road, grid}\}$, the formulator compiles planner-ready instances that respect the time model, capacity/avoidance rules (Sec. III-D), and objective hooks (Sec. III-E). It supplies graphs/costs for roadmap/grid and collision predicates/state propagators for continuous space.

Unified Planner Our unified planner normalizes inputs (instance, budgets) and outputs (paths/trajectories with timestamps, status, and costs) across grid MAPF, roadmap MAPF, and continuous MRMP planners. Planners run as shared libraries or subprocesses with wall-clock/node budgets.

C. Simulation Core

Box2D serves as the cross-platform deterministic, continuous-time substrate. We advance at 60 Hz ($\Delta t_{\text{sim}} = 1/60$ s) with substepping (default $n = 4$) for contact stability. Each agent tracks its reference trajectory kinematically; collision checking uses broad-phase dynamic AABB trees and narrow-phase contact manifolds. We log contact events, nearest-approach distances, and tracking error. Given the same instance, seeds, and step parameters, the state trace is bitwise identical.

D. Interactive User Interface

The UI supports (i) environment edits (obstacles, resolution/roadmap parameters), (ii) agent edits (footprints, dynamics, starts/goals), (iii) planner/budget selection, and (iv) plan/execution visualization. Overlays include occupancy/roadmap layers, time-expanded reservations, clearance heatmaps, and per-agent timelines. Edits are transactional: dependent artifacts are invalidated and recomputed so UI and headless runs remain consistent.

V. EXPERIMENTS AND RESULTS

We evaluate GRACE against three claims: generality, its capacity to yield new research insights, and practicality.

Experimental Conditions. All studies are conducted headless with fixed seeds and recorded manifests to ensure reproducibility. Grid experiments always use 4-connected moves; roadmap experiments use disk agents with single-integrator dynamic; continuous experiments execute time-parameterized trajectories with collision checks in Box2D. All experiments are performed on a laptop with an Intel(R) Core(TM) i7-13850HX CPU and 32GB of memory. For learning-based grid planners, we download trained models from official repositories and evaluate inference-time only.

We leverage GPU on the same laptop with NVIDIA RTX 3500 Ada (12 GB).

Planners. Tab. II lists the planners and representations they operate on and whether we evaluate them in GRACE (**Eval.**). We integrate official implementations via shared libraries or subprocesses (links in **OSS**). All run under identical instance definitions and budgets: 1 min (grid/roadmap) and 5 min (continuous) with default hyperparameters.

Metrics. We implement various metrics to support the benchmark, including success rate (%), SoC (path length in meter), makespan (s), and planning time (s). In addition, we also record the simulation speed, quantified as the Real-Time Factor (RTF), which is the ratio of simulated time to wall time (RTF = simulated time / wall time).

A. Cross-Environment-Model Benchmarking

Purpose. We assess *generality* by expressing a fixed task set at multiple abstraction levels, including continuous (MRMP), roadmap (MAPF), and grid (MAPF). We measure how each representation trades SoC, makespan, and planning time.

Setup. We use a 5×5 m² workspace with four map families: `empty`, `maze`, `warehouse`, and `window`. Teams contain up to $N=8$ robots with heterogeneous convex footprints (maximum effective radius $r_{\text{eff}}^{\text{max}}=0.08$ m). To isolate representation effects, we use identical double-integrator bounds for continuous representation: $v \in [-0.5, 0.5]$ m/s, $a \in [-2, 2]$ m/s², consistent with [31]. Roadmaps are instantiated at three densities (S=sparse, D=dense, SD=superdense); grids use cell size derived from r_{com} (32×32).

Results. As shown in Fig. 3, we observe clear representation-induced trade-offs. Roadmap consistently achieves near-continuous SoC with shorter makespans and much faster planning; higher roadmap density tightens SoC and makespan at a small runtime cost. Grid further accelerates planning and often matches roadmap-level makespans, but incurs a larger SoC penalty. Makespan across MRMP and MAPF is *not* directly comparable because the time/kinematic models differ (double-integrator vs. single-integrator/unit-step); thus, makespan values should be read qualitatively and relative to the per-map continuous baseline.

Observations. (1) Roadmap is “good enough” in these tasks. Across maps, roadmap SoC stays within ≈ 108 – 118% of continuous MRMP, while makespan drops to ≈ 52 – 61% and planning time to ≈ 0.16 – 1.6% of db-CBS. (2) Grid trades more SoC for similar makespan and extreme speed. Grid averages $\approx 123\%$ SoC, $\approx 61\%$ makespan, and $\approx 2.4\%$ planning time. (3) Density matters on roadmaps. Superdense roadmaps tighten SoC (108%) and makespan (52%) at modest runtime (1.6%), indicating a controllable speed–quality knob.

Takeaway. Moving from continuous to roadmap trades kinodynamic fidelity for speed, yielding faster planning and shorter makespans with a small SoC penalty. Further abstraction to grids amplifies planning speed but often increases makespan and SoC. However, environmental structure (e.g., aisles) can mitigate these penalties, sometimes favoring discrete abstractions. These findings underscore GRACE’s

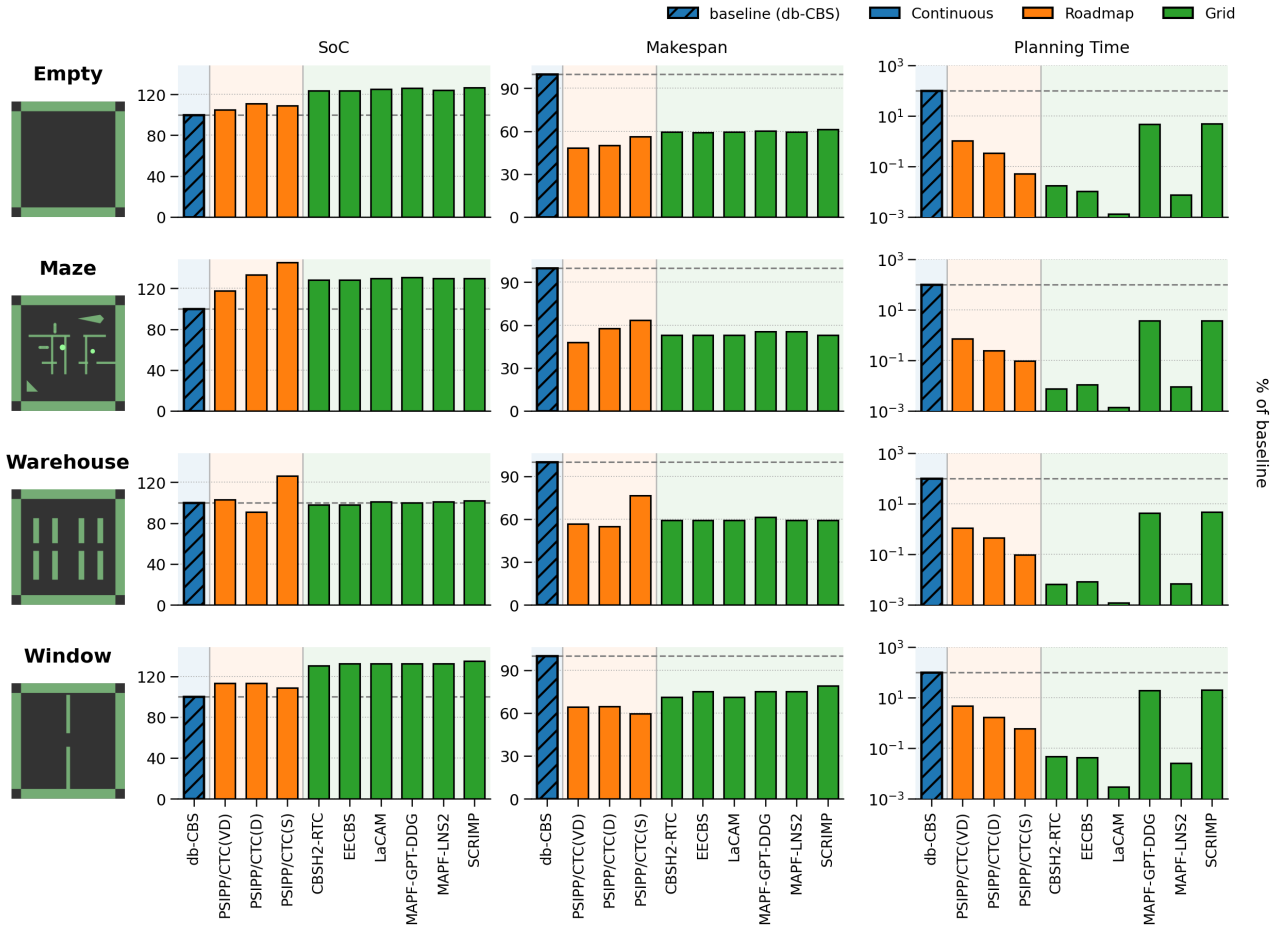


Fig. 3. Cross-Environment Comparison of MRMP and MAPF Solvers: SoC, Makespan, Planning Time (db-CBS=100%)

ability to expose representation-induced trade-offs, revealing when discrete abstractions are adequate and when kinodynamic fidelity becomes critical.

B. Cross-MAPF-Planner Benchmarking on Grids

Purpose. We evaluate GRACE’s value in *facilitating rigorous, comparative analysis* by demonstrating its capability as a reproducible, drop-in grid benchmark also for single-representation (e.g., grid-only) studies. This test shows direct compatibility with established MAPF community evaluations, while preserving our unified protocol for instance specification, budgets, logging, and evaluation.

Setup. We evaluate six planners on four canonical grid maps (random-64-64-10, maze-32-32-4, warehouse-10-20-10-2-1, den312d) with 25 scenarios each, all from [8]. The tile side length is 1 meter. Agent counts (N) are swept, applying fixed wall-clock budgets and the standard success criterion (Sec. III-D).

Success Rate. Fig. 4 reports success rate vs. N . Three consistent trends: (i) LaCAM achieves 100% success across the entire sweep on all four maps in our setting; (ii) EECBS is strongest at low-mid N but degrades earlier on den312d/maze-32-32-4, while MAPF-LNS2 degrades more gradually at larger N ; (iii) learned/hybrid methods

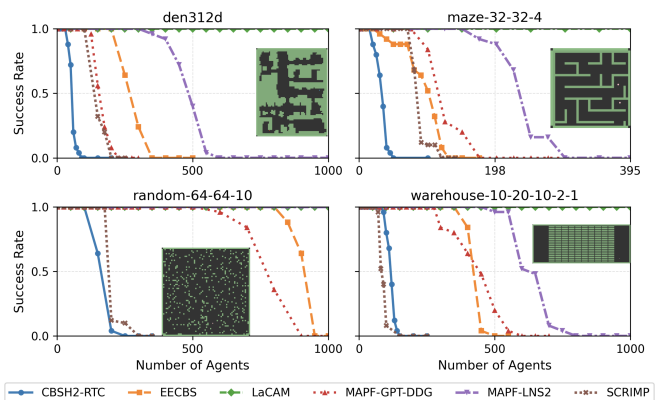


Fig. 4. Grid MAPF Comparison: Success Rate with Number of Agents.

(MAPF-GPT-DDG, SCRIMP) can match or exceed search on random/warehouse but show higher across-map variance.

SoC and Planning Time. To compare path quality fairly, we compute average SoC and planning time only on the subset of instances where *all* planners succeed. CBSH2-RTC achieves the lowest SoC, followed by EECBS; LNS and learned/hybrid trade slightly higher SoC for improved scale tolerance; LaCAM prioritizes robustness at larger N .

TABLE IV

SoC AND PLANNING TIME AVERAGED OVER THE COMMON-SUCCESS SUBSET ACROSS ALL PLANNERS.

Planner	Mean SoC (m) ↓	Mean planning time (s) ↓
CBSH2-RTC	3827 ± 286	4.61 ± 1.84
EECBS	3940 ± 288	0.63 ± 0.08
MAPF-LNS2	4206 ± 321	0.45 ± 0.05
MAPF-GPT-DDG	4257 ± 333	5.66 ± 0.26 †
LaCAM	4511 ± 358	0.01 ± 0.00
SCRIMP	5290 ± 422	20.86 ± 1.56 †

† Includes Python driver and model loading overhead (Docker).
gray: 95 % Confidence Interval (CI).

Observation and follow-up. Because LaCAM already achieves 100% success on all four maps in the nominal sweep, we run an additional large-scale stress test, as shown in Sec. V-C; LaCAM maintained 100% success until $N = 2500$ in our 1 minute budget in that extended regime.

Takeaway. This grid-only benchmark confirms GRACE’s utility as a reproducible platform for MAPF planner evaluation, fully compatible with existing community standards. The results further illuminate the trade-offs between planner robustness, optimality, and efficiency across map types and agent densities within GRACE’s unified framework.

C. Scalability and Software Resources

Purpose. We evaluate *practicality* through the runtime scalability, resource usage, and execution determinism of GRACE’s simulator and benchmarking stack, and we report capacity-style metrics that make comparisons across hardware and existing simulators straightforward.

Setup. We use the same integration settings as Sec. IV-C. For each run, we record OS-reported peak CPU utilization and process resident peak memory. To facilitate cross-paper comparison, we summarize (i) the largest team size with a real-time factor $\text{RTF} \geq 1$, (ii) the team size at which the median curve crosses a reference level $\text{RTF} = 10$ (as used in [17]), and (iii) peak-memory scaling via a linear fit; we report the slope in MB/agent and coefficient of determination R^2 . Collision checking remains enabled.

Results. We stress-test the grid representation by replaying LaCAM plans on `warehouse_large` (size: 540m × 140m) with team sizes up to $N = 2500$. GRACE sustains real time up to the largest tested team size at $N = 2500$ and remained above the reference level at approximately $N = 550$. Peak memory grows linearly with team size; a least-squares fit yielded a slope of ~ 0.179 MB/agent with $R^2 = 1.00$, indicating a small per-agent footprint and modest constant overhead. Peak CPU stabilizes around 100–110% in this configuration, consistent with a lightly threaded but memory-efficient workload. Table V reports representative points on the median curves.

GRACE uses the same waypoint-tracking and contact model for all environment representations (Sec. IV-C), so differences observed come from the plans and not from execution.

TABLE V

KEY POINTS ACROSS N DURING PATH-PLAYBACK ON GRID REPRESENTATION `WAREHOUSE_LARGE` USING LACAM.

N	RTF (median)	Peak CPU (%)	Peak Mem (GB)
2	1288.21	9.9	0.039
550	10.59	108.6	0.134
1000	5.41	107.4	0.213
2000	2.39	107.1	0.387
2500	1.84	107.2	0.476

We report median Real-Time Factor (RTF), peak CPU, and peak memory at selected team sizes. Peak CPU denotes the OS-aggregated utilization during playback and indicates system load rather than hardware-independent cost.

VI. CONCLUSION

We present GRACE, a unified simulator and benchmark that enables commensurate evaluation of discrete MAPF and kinodynamically faithful MAMP across grid, roadmap, and continuous representations under a single protocol. Experiments verify our claims of generality (convertible environments; heterogeneous agents and obstacles), usefulness (fair, reproducible comparisons that expose representation–algorithm trade-offs), and practicality (fast-speed playback and deterministic logs on modest hardware).

Limitations include a 2D focus, static maps, fixed convex geometries, and agent abstractions for roadmap/grid. Future work will explore uncertainty and dynamic changes, lifelong/multi-goal settings, richer motion dynamics and objectives, automated abstraction/roadmap selection, and connectors to operational stacks, including ROS 2 and Fleet Management Systems (FMS), moving toward deployment.

REFERENCES

- [1] R. Stern, N. R. Sturtevant *et al.*, “Multi-agent pathfinding: Definitions, variants, and benchmarks,” in *International Symposium on Combinatorial Search (SOCS)*, 2019, pp. 151–158.
- [2] K. Okumura, “Engineering lacam*: Towards real-time, large-scale, and near-optimal multi-agent pathfinding,” in *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2024, pp. 1501–1509.
- [3] J. Ren, E. Ewing *et al.*, “Map connectivity and empirical hardness of grid-based multi-agent pathfinding problem,” in *International Conference on Automated Planning and Scheduling (ICAPS)*, 2024, pp. 484–488.
- [4] M. Debor, W. Hönig *et al.*, “Trajectory planning for heterogeneous robot teams,” in *International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 7924–7931.
- [5] J. Gao, Y. Li *et al.*, “A review of graph-based multi-agent pathfinding solvers: From classical to beyond classical,” *Knowledge-based Systems*, vol. 283, p. 111121, 2024.
- [6] A. Skrynnik, A. Andreychuk *et al.*, “POGEMA: A benchmark platform for cooperative multi-agent pathfinding,” in *International Conference on Learning Representations (ICLR)*, 2025, pp. 5576–5606.
- [7] S. Schaefer, L. Palmieri *et al.*, “A benchmark for multi-robot planning in realistic, complex and cluttered environments,” in *International Conference on Robotics and Automation (ICRA)*, 2023, pp. 9231–9237.
- [8] N. R. Sturtevant, “Benchmarks for grid-based pathfinding,” *IEEE Transactions on Computational Intelligence and AI in Games (TCI-AIG)*, vol. 4, no. 2, pp. 144–148, 2012.
- [9] O. Kaduri, E. Boyarski *et al.*, “Experimental evaluation of classical multi agent path finding algorithms,” in *International Symposium on Combinatorial Search (SOCS)*, 2021, pp. 126–130.
- [10] B. Shen, “Tracking progress in multi-agent path finding (student abstract),” in *International Symposium on Combinatorial Search (SOCS)*, 2023, pp. 200–201.

- [11] J. Tan, Y. Luo *et al.*, “Benchmarking large neighborhood search for multi-agent path finding,” *CoRR*, vol. abs/2407.09451, 2024.
- [12] M. Gebser, P. Obermeier *et al.*, “Experimenting with robotic intralogistics domains,” *Theory and Practice of Logic Programming (TPLP)*, vol. 18, no. 3-4, pp. 502–519, 2018.
- [13] S. P. Mohanty, E. Nygren *et al.*, “Flatland-rl : Multi-agent reinforcement learning on trains,” *CoRR*, vol. abs/2012.05893, 2020.
- [14] S.-H. Chan, Z. Chen *et al.*, “The League of Robot Runners Competition: Goals, Designs, and Implementation,” in *ICAPS 2024 System’s Demonstration track*, 2024.
- [15] V. N. Hartmann, T. Heinle *et al.*, “A benchmark for optimal multimodal multi-robot multi-goal path planning with given robot assignment,” *CoRR*, vol. abs/2503.03509, 2025.
- [16] L. Heuer, L. Palmieri *et al.*, “Benchmarking multi-robot coordination in realistic, unstructured human-shared environments,” in *International Conference on Robotics and Automation (ICRA)*, 2024, pp. 14 541–14 547.
- [17] J. Yan, Z. Li *et al.*, “Advancing MAPF towards the real world: A scalable multi-agent realistic testbed (SMART),” *CoRR*, vol. abs/2503.04798, 2025.
- [18] T. T. Walker, N. R. Sturtevant *et al.*, “Extended increasing cost tree search for non-unit cost domains,” in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2018, pp. 534–540.
- [19] J. Li, W. Ruml *et al.*, “EECBS: A bounded-suboptimal search for multi-agent path finding,” in *AAAI Conference on Artificial Intelligence (AAAI)*, 2021, pp. 12 353–12 362.
- [20] J. Li, Z. Chen *et al.*, “MAPF-LNS2: fast repairing for multi-agent path finding via large neighborhood search,” in *AAAI Conference on Artificial Intelligence (AAAI)*, 2022, pp. 10 256–10 265.
- [21] Y. Wang, B. Xiang *et al.*, “SCRIMP: scalable communication for reinforcement- and imitation-learning-based multi-agent pathfinding,” in *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2023, pp. 2598–2600.
- [22] A. Andreychuk, K. S. Yakovlev *et al.*, “Advancing learnable multi-agent pathfinding solvers with active fine-tuning,” in *International Conference on Intelligent Robots and Systems (IROS)*, 2025, pp. 10 564–10 571.
- [23] I. Solis, J. Motes *et al.*, “Representation-optimal multi-robot motion planning using conflict-based search,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 6, no. 3, pp. 4608–4615, 2021.
- [24] A. Andreychuk, K. S. Yakovlev *et al.*, “Multi-agent pathfinding with continuous time,” *Artificial Intelligence (AIJ)*, vol. 305, p. 103662, 2022.
- [25] K. Kasaura, M. Nishimura *et al.*, “Prioritized safe interval path planning for multi-agent pathfinding with continuous time on 2d roadmaps,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 7, no. 4, pp. 10 494–10 501, 2022.
- [26] T. T. Walker, N. R. Sturtevant *et al.*, “Cliques analysis and bypassing in continuous-time conflict-based search,” in *International Symposium on Combinatorial Search (SOCS)*, 2024, pp. 152–160.
- [27] Y. Li, Z. Littlefield *et al.*, “Asymptotically optimal sampling-based kinodynamic planning,” *International Journal of Robotics Research (IJRR)*, vol. 35, no. 5, pp. 528–564, 2016.
- [28] B. Binder, F. Beck *et al.*, “Multi robot route planning (MRRP): extended spatial-temporal prioritized planning,” in *International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 4133–4139.
- [29] J. Kottinger, S. Almagor *et al.*, “Conflict-based search for explainable multi-agent path finding,” in *International Conference on Automated Planning and Scheduling (ICAPS)*, 2022, pp. 692–700.
- [30] K. Okumura and X. Défago, “Quick multi-robot motion planning by combining sampling and search,” in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2023, pp. 252–261.
- [31] A. Moldagalieva, J. O. de Haro *et al.*, “db-cbs: Discontinuity-bounded conflict-based search for multi-robot kinodynamic motion planning,” in *International Conference on Robotics and Automation (ICRA)*, 2024, pp. 14 569–14 575.
- [32] W. Lin, W. Song *et al.*, “Multi-agent path finding with heterogeneous geometric and kinematic constraints in continuous space,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 10, no. 1, pp. 492–499, 2025.
- [33] A. Li, Z. Chen *et al.*, “CBS with continuous-time revisit,” *CoRR*, vol. abs/2501.07744, 2025.
- [34] S. K. Moghadam, M. Ebrahimi *et al.*, “Guards: Benchmarks for weighted grid-based pathfinding,” *Expert Systems With Applications*, vol. 249, p. 123719, 2024.
- [35] J. Alkazzi and K. Okumura, “A comprehensive review on leveraging machine learning for multi-agent path finding,” *IEEE Access*, vol. 12, pp. 57 390–57 409, 2024.
- [36] J. Svancara, D. Atzmon *et al.*, “Which objective function is solved faster in multi-agent pathfinding? it depends,” in *International Conference on Agents and Artificial Intelligence (ICAART)*, 2024, pp. 23–33.
- [37] W. Ma, W. Li *et al.*, “Skyrover: A modular simulator for cross-domain pathfinding,” in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2025, pp. 11 086–11 090.
- [38] W. Hönig, T. K. S. Kumar *et al.*, “Path planning with kinematic constraints for robot groups,” *CoRR*, vol. abs/1704.07538, 2017.
- [39] W. Hönig, S. Kiesel *et al.*, “Persistent and robust execution of MAPF schedules in warehouses,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 4, no. 2, pp. 1125–1131, 2019.
- [40] W. Hönig, J. A. Preiss *et al.*, “Trajectory planning for quadrotor swarms,” *IEEE Transactions on Robotics (T-RO)*, vol. 34, no. 4, pp. 856–869, 2018.
- [41] E. Barros, J. P. Hanna *et al.*, Eds., *RoboCup 2024: Robot World Cup XXVII*, ser. Lecture Notes in Computer Science, vol. 15570, 2025.
- [42] G. Sharon, R. Stern *et al.*, “Conflict-based search for optimal multi-agent pathfinding,” *Artificial Intelligence (AIJ)*, vol. 219, pp. 40–66, 2015.
- [43] J. Li, D. Harabor *et al.*, “Pairwise symmetry reasoning for multi-agent path finding search,” *Artificial Intelligence (AIJ)*, vol. 301, p. 103574, 2021.
- [44] S. Chan, J. Li *et al.*, “Flex distribution for bounded-suboptimal multi-agent path finding,” in *AAAI Conference on Artificial Intelligence (AAAI)*, 2022, pp. 9313–9322.
- [45] S. Chan, Z. Chen *et al.*, “Anytime multi-agent path finding using operation parallelism in large neighborhood search,” in *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2024, pp. 2183–2185.
- [46] S. Chan, R. Stern *et al.*, “Greedy priority-based search for suboptimal multi-agent path finding,” in *International Symposium on Combinatorial Search (SOCS)*, 2023, pp. 11–19.
- [47] K. Okumura and H. Nagai, “Lightweight and effective preference construction in pibt for large-scale multi-agent pathfinding,” in *International Symposium on Combinatorial Search (SOCS)*, vol. 18, 2025, pp. 110–117.
- [48] G. Sartoretti, J. Kerr *et al.*, “PRIMAL: pathfinding via reinforcement and imitation multi-agent learning,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 4, no. 3, pp. 2378–2385, 2019.
- [49] R. Veerapaneni, Q. Wang *et al.*, “Improving learnt local MAPF policies with heuristic search,” in *International Conference on Automated Planning and Scheduling (ICAPS)*, 2024, pp. 597–606.
- [50] H. Tang, F. Berto *et al.*, “Himap: Learning heuristics-informed policies for large-scale multi-agent pathfinding,” in *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2024, pp. 2498–2500.
- [51] Y. Wang, T. Duhan *et al.*, “LNS2+RL: combining multi-agent reinforcement learning with large neighborhood search in multi-agent path finding,” in *AAAI Conference on Artificial Intelligence (AAAI)*, 2025, pp. 23 343–23 350.
- [52] M. Phillips and M. Likhachev, “SIPP: safe interval path planning for dynamic environments,” in *International Conference on Robotics and Automation (ICRA)*, 2011, pp. 5628–5635.
- [53] W. Hönig, T. K. S. Kumar *et al.*, “Multi-agent path finding with kinematic constraints,” in *International Conference on Automated Planning and Scheduling (ICAPS)*, 2016, pp. 477–485.
- [54] R. Shome, K. Solovey *et al.*, “drrt*: Scalable and informed asymptotically-optimal multi-robot motion planning,” *Autonomous Robots*, vol. 44, no. 3, pp. 443–467, 2020.
- [55] D. Dayan, K. Solovey *et al.*, “Near-optimal multi-robot motion planning with finite sampling,” *IEEE Transactions on Robotics (T-RO)*, vol. 39, no. 5, pp. 3422–3436, 2023.
- [56] J. Kottinger, S. Almagor *et al.*, “Conflict-based search for multi-robot motion planning with kinodynamic constraints,” in *International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 13 494–13 499.
- [57] “Box2D — box2d.org,” <https://box2d.org/>, [Accessed 12-09-2025].
- [58] I. A. Sucan, M. Moll *et al.*, “The open motion planning library,” *IEEE Robotics & Automation Magazine*, vol. 19, no. 4, pp. 72–82, 2012.