Combining multi-robot motion planning and goal allocation using roadmaps

Salvado, João; Mansouri, Masoumeh; Pecora, Federico

DOI:

[10.1109/ICRA48506.2021.9560861](https://doi.org/10.1109/ICRA48506.2021.9560861)

License:

Other (please specify with Rights Statement)

Document Version

Peer reviewed version

Citation for published version (Harvard):

Salvado, J, Mansouri, M & Pecora, F 2021, Combining multi-robot motion planning and goal allocation using roadmaps. in *2021 IEEE International Conference on Robotics and Automation (ICRA)*., 9560861, IEEE International Conference on Robotics and Automation, IEEE, pp. 10016-10022, 2021 IEEE International Conference on Robotics and Automation (ICRA), Xi'an , China, 30/05/21.
<https://doi.org/10.1109/ICRA48506.2021.9560861>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

J. Salvado, M. Mansouri and F. Pecora, "Combining Multi-Robot Motion Planning and Goal Allocation using Roadmaps," 2021 IEEE International Conference on Robotics and Automation (ICRA), 2021, pp. 10016-10022, doi: 10.1109/ICRA48506.2021.9560861.

© 2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Combining Multi-Robot Motion Planning and Goal Allocation using Roadmaps

João Salvado, Masoumeh Mansouri, Federico Pecora

Abstract—This paper addresses the problem of automating fleets of robots with non-holonomic dynamics. Previously studied methods either specialize in facets of this problem, that is, one or a combination of multi-robot goal allocation, motion planning, and coordination, and typically sacrifice optimality and completeness for scalability. We propose an approach that constructs an abstract multi-robot roadmap in a reduced configuration space, where we account for environment connectivity and interference cost between robots occupying the same polygons. Querying the road-map results in a robot-goal assignment and abstract multi-robot trajectory. This is then exploited to de-compose the original problem into smaller problems, each of which is solved with a multi-robot motion planner that accounts for kinodynamic constraints. We validate the approach experimentally to demonstrate the advantage of considering task assignment and motion planning holistically, and explore some methods for balancing solution quality and computational efficiency.

I. INTRODUCTION

Planning for multi-robot systems subsumes four problems: task allocation, which is the problem of determining how tasks should be assigned to robots in order to maximize fleet performance; coordination, which is to decide how robots should avoid each other while traversing the shared environment; motion planning, i.e., determining the trajectories robots take in completing tasks; and the problem of determining control inputs that bring about the realization of these trajectories under given constraints. These problems overlap in non-trivial ways. Loosely-coupled integrated approaches trade completeness for efficiency, and no such approach has thus far been developed that tackles all four aspects of the overall problem; trajectory optimization based methods employ constrained optimization formulations of the overall problem, which typically leads to near-optimal solutions, but at the cost of impractical computational overhead; graph-based methods leverage abstractions of the environment, robots and tasks to achieve scalability, but these abstractions are typically unrealistic.

In this paper, we propose a method which exploits key insights from all three types of approaches. A graph-based abstraction, which we call *multi-robot roadmap*, is used to determine how robots should traverse a partition of navigable space into convex polygons. Heuristic search is used to simultaneously determine task assignments and polygons to be traversed, while accounting for both individual robot performance and the cost of interference among robots (the

balance of which can be regulated with a single parameter). Search results in an abstract multi-robot trajectory, which is then used to decompose the multi-robot motion planning problem into a partial order of smaller trajectory optimization problem instances. The solutions of these are then combined to obtain a kinodynamically feasible multi-robot trajectory.

We show empirically that our approach scales to tens of robots, and that considering task assignment and motion planning holistically leads to high quality solutions according to common metrics for measuring the performance of multi-robot systems. As for other roadmap-based approaches, our method allows some of the computation to be performed offline, i.e., without knowing the initial and final conditions of the robot fleet. We show empirically that the amount of computation that is done offline can be regulated by changing the granularity of the roadmap. We also show that a more fine-grained roadmap leads to better problem decomposition, which in turn improves both solution quality and query times. Finally, we demonstrate how increasing the weight of interference cost in the heuristic search leads to more efficient problem decompositions and improves the makespan of the fleet's trajectories.

II. RELATED WORK

Considering task allocation, motion planning, coordination and control jointly can lead to highly-optimized solutions. However, this incurs significant computational overhead. In industrial practice, the problem is simplified by re-engineering the environment and/or the robots [1]. Others sacrifice completeness and optimality for scalability, e.g., by combining single-robot motion planning and coordination. Examples of such approaches include prioritized planning [2], [3], and coordination via velocity profile adjustment [4], [5]. None of these account for task allocation.

Task allocation has often been considered separately from the other problems [6]. For instance, Nam and Shell [7] formulate the optimal task assignment problem as a constrained optimization problem where interference among robots is considered as a cost. This and other similar approaches [8] rely on loosely-coupled integrations with other solvers, e.g., for coordination and motion planning [9].

Approaches based on trajectory optimization deal with the complexity of multi-robot planning by relaxing/convexifying certain constraints to exploit efficient sequential convex programming solutions [10]. These methods have been applied for holonomic robots [11], non-holonomic robots in obstacle-free environments [12], and for robots with non-convex

João Salvado and Federico Pecora are with the AASS Research Centre, Örebro University, <name>.<surname>@oru.se

Masoumeh Mansouri is with the School of Computer Science at the University of Birmingham, m.mansouri@bham.ac.uk

dynamics [13]. Roadmaps have also been utilized to warm-start non-linear solvers [14]. Again, though, none of these approaches spans the gamut of all four aspects of the problem without imposing restrictive assumptions.

Graph-based methods are known to scale to large fleets of robots, however, these make further unrealistic assumptions. Some exploit efficient network flow algorithms for solving multi-robot goal allocation and motion planning jointly. However, robots are required to move in “grid-like” environments represented as a graph, where edges/paths are non-colliding and vertices/states are occupied by at most one robot at a time [15]. Conflicts can be solved by imposing constraints over the paths of the robots (e.g., computed via a binary search tree [16]), or by constructing the resulting configuration free space [17], possibly also allocating tasks to heterogeneous robots [18]. Yet other approaches consider the joint configuration space of multiple robots only when conflicts cannot be solved locally [19]. Another class of graph-based methods considers disc-shaped robots as pebbles moving on a graph [20], [21] assuming robots do not overlap on start and goal configurations [22] (an assumption which can be relaxed by generating multiple pebble graphs [23]).

Some graph-based methods handle complexity by dividing the overall problem into multiple levels of abstraction. Honig et al. [24] use realistic dynamic models to post-process solutions derived from a high-level graph search. Overmars et al. [25] generate roadmap abstractions by partitioning a single-robot roadmap into regions, though this approach does not consider interference between robots in the same region. A probabilistically complete roadmap-based approach for multi-robot motion planning is proposed by Le and Plaku [26], in which sampling of the combined multi-robot configuration space is guided by a search in a discrete space with holonomic point robots. Similarly, Du et al. [27] exploit a pre-computed roadmap for single robots as a source of states to be found closer to the newly sampled state in the multi-robot configuration space. None of these approaches consider the task allocation problem.

III. PROBLEM STATEMENT

A. Nomenclature

In this article, we refer to tuples, sets and vectors as $\langle \dots \rangle$, $\{ \dots \}$, and $[\dots]$, respectively. Non-bold uppercase letters are used for constants. Bold letters denote matrices (uppercase) and vectors (lowercase), and the notation $\mathbf{x}[i]$ indicates the i -th element of vector $\mathbf{x}$. To simplify exposition, we will consistently use $r \in \{1, \dots, R\}$ to index robots, $p \in \{1, \dots, P\}$ for polygons, $g \in \{1, \dots, G\}$ for goals, and $k \in \{1, \dots, N\}$ for discrete time. Continuous time is normalized w.r.t. the planning horizon, i.e., $t \in [0, 1]$, and discrete time k corresponds to continuous time $\frac{k}{N}$.

B. Preliminary Definitions

a) *Robots and World*: Consider a bounded 2D-world setting $\mathcal{W} \subset \mathbb{R}^2$, together with an obstacle region $\mathcal{O} \subset \mathcal{W}$. A robot r has a geometry $\mathcal{A}_r(\mathbf{q}_r)$ in configuration $\mathbf{q}_r \in \mathcal{C}_r$, where $\mathcal{C}_r$ is the configuration space of r . The

robot’s configuration-free space is $\mathcal{C}_r^{\text{free}} = \{\mathbf{q}_r \in \mathcal{C}_r \mid \mathcal{A}_r(\mathbf{q}_r) \cap \mathcal{O} = \emptyset\}$, that is, the region of the world where r ’s geometry intersects no obstacle. Each robot has state $\mathbf{x}_r = [\mathbf{q}_r, \dot{\mathbf{q}}_r]$, a control input $\mathbf{u}_r$, and a dynamic model $\dot{\mathbf{x}}_r = \mathbf{f}_r(\mathbf{x}_r(t), \mathbf{u}_r(t))$. A single robot is therefore fully defined by $\mathcal{R}_r = \langle \mathcal{A}_r, \mathbf{x}_r, \mathbf{u}_r, \mathbf{f}_r \rangle$. The overall state and controls of a fleet of R robots are $\mathbf{x} = [\mathbf{x}_1 \dots \mathbf{x}_R]$ and $\mathbf{u} = [\mathbf{u}_1 \dots \mathbf{u}_R]$; similarly, the overall configuration space and geometry of the fleet are, respectively, $\mathcal{C}^{\text{free}} = \prod_{r=1}^R \mathcal{C}_r^{\text{free}}$ and $\mathcal{A} = \cup_{r=1}^R \mathcal{A}_r$.

b) *Trajectory*: A single robot trajectory between a start configuration $\mathbf{s}_r \in \mathcal{C}_r^{\text{free}}$ and goal configuration $\mathbf{g}_r \in \mathcal{C}_r^{\text{free}}$ is a continuous mapping $\tau_r : [0, 1] \mapsto \mathcal{C}_r^{\text{free}}$, where $\tau_r(0) = \mathbf{s}_r$ and $\tau_r(1) = \mathbf{g}_r$. We say that a trajectory is *free* when $\tau_r(t) \in \mathcal{C}_r^{\text{free}}$ for all t . Note that, for a robot to navigate along the trajectory $\tau_r(t)$ while adhering to kinodynamic constraints, a *kinematically feasible control policy* $\pi_r(\mathbf{f}_r, \tau_r) = \mathbf{u}_r(\mathbf{x}_r(t), t)$ must be followed. We say that a *trajectory is feasible* if such a policy exists and the trajectory τ_r is free.

c) *Obstacle Free Space*: Let the obstacle free space be the set $\mathcal{W}^{\text{free}} = \mathcal{W} \setminus \mathcal{O}$. This can be approximated with a set $\{\mathcal{P}_1, \dots, \mathcal{P}_P\}$ of adjacent non-overlapping convex polygons. We say that polygons are *adjacent* if they intersect on their frontier as, e.g., in a Voronoi tessellation [28] with convex cells. Given such a set of adjacent convex polygons, one can construct the *connectivity graph* of the obstacle free space $G_{\mathcal{W}^{\text{free}}} = \langle V_{\mathcal{W}^{\text{free}}}, E_{\mathcal{W}^{\text{free}}} \rangle$, where the set of vertices $V_{\mathcal{W}^{\text{free}}} = \{\mathcal{P}_1, \dots, \mathcal{P}_P\}$ contains one vertex per polygon and $E_{\mathcal{W}^{\text{free}}} = \{(\mathcal{P}_i, \mathcal{P}_j) \in V_{\mathcal{W}^{\text{free}}}^2 \mid \exists \mathbf{q}_r \in \mathcal{C}_r^{\text{free}} \text{ s.t. } \mathbf{q}_r \subset \partial \mathcal{P}_i \cap \partial \mathcal{P}_j\}$, having $\partial \mathcal{P}$ as the boundary of a polygon, i.e., an edge exist if one can “place” a robot on the intersection of the two polygons. Figure 1(a) shows an example where the connectivity graph has set of four vertices $V_{\mathcal{W}^{\text{free}}} = \{\mathcal{P}_1, \mathcal{P}_2, \mathcal{P}_3, \mathcal{P}_4\}$ and edges $E_{\mathcal{W}^{\text{free}}} = \{(\mathcal{P}_1, \mathcal{P}_2), (\mathcal{P}_1, \mathcal{P}_4), (\mathcal{P}_2, \mathcal{P}_3), (\mathcal{P}_3, \mathcal{P}_4)\}$.

d) *Robot Goal Allocation*: Let $\mathcal{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_R\}$ and $\mathcal{G} = \{\mathbf{g}_1, \dots, \mathbf{g}_G\}$ be, respectively, a set of start and goal configurations for a fleet of R robots. Moreover, let $\sigma : \mathcal{S} \mapsto \mathcal{G}$, be a mapping from robot starting poses to goals. Hence, finding an allocation of robots to goals means computing σ .

C. Problem Formulation

We define a *Multi-robot Motion Planning and Goal Allocation Problem (MMP-GA)* as a tuple $\mathcal{M} = \langle \mathcal{W}^{\text{free}}, \mathcal{R}, \mathcal{S}, \mathcal{G} \rangle$. The task is to find function σ and a joint fleet trajectory $\tau(t) = [\tau_1(t) \dots \tau_R(t)]$ such that:

- 1) $\tau_r(0) = \mathbf{s}_r \forall r \in \{1, \dots, R\}$ (start configuration)
- 2) $\tau_r(1) = \sigma(\mathbf{s}_r) \forall r \in \{1, \dots, R\}$ (goal configuration)
- 3) $\tau_r(t) \in \mathcal{C}_r^{\text{free}}$ and $\exists \pi_r(\mathbf{f}_r, \tau_r) = \mathbf{u}_r(\mathbf{x}_r, t) \forall r \in \{1, \dots, R\}, \forall t \in [0, 1]$ (all trajectories are feasible),
- 4) $\mathcal{A}_i(\tau_i(t)) \cap \mathcal{A}_j(\tau_j(t)) = \emptyset \forall t \in [0, 1], \forall i, j \neq i \in \{1, \dots, R\}$ (robots do not collide with each other),

while maximizing some performance criterion $L(t, \mathbf{x}, \mathbf{u})$ (we omit the subscript r when referring to the entire fleet).

IV. APPROACH

This paper solves a specific instance of the MMP-GA problem, where robots are homogeneous, meaning that they

Fig. 1: Running example of a MMP-GA problem with two robots and two goals: (a) shows a partial representation of the graph G^K and trajectory τ^K connecting abstract start and goal configurations; (b) extends the representation of the abstract trajectory τ^K , illustrating how step-cost and heuristic are computed; (c) demonstrates how this problem is de-composed.

have the same geometry, control and state space, and kinodynamics. The configuration space $\mathcal{C}_r$ of each robot is $SE(2) = \mathbb{R}^2\mathbb{S}^1$ (2D rigid bodies). Hence, the configuration of r can be defined by Euclidean coordinates and orientation, $\mathbf{q}_r = [x_r \ y_r \ \theta_r]$. Additionally, we assume that robots have car-like non-holonomic kinodynamics. For example,

$$\dot{x}_r = v_r \cos(\theta_r), \dot{y}_r = v_r \sin(\theta_r), \dot{\theta}_r = k_r v_r, \dot{v}_r = a_r \quad (1)$$

$$f_r(\mathbf{x}_r, \mathbf{u}_r) = \dot{\mathbf{x}}_r = [\dot{x}_r \ \dot{y}_r \ \dot{\theta}_r \ \dot{v}_r], \quad \mathbf{u}_r = [k_r \ a_r], \quad (2)$$

where k_r is the curvature, v_r is the speed, and a_r the acceleration. In summary, the input of the MMP-GA is: (1) a partition of $\mathcal{W}^{\text{free}}$ into adjacent polygons $\{\mathcal{P}_1, \dots, \mathcal{P}_P\}$; (2) a set of homogeneous robots $\{\mathcal{R}_1, \dots, \mathcal{R}_R\}$; (3) start configurations $[s_1 \dots s_R] \in SE(2)^R$; (4) goal configurations $[g_1 \dots g_G] \in SE(2)^G$; (5) a cost function $L(\cdot)$. A solution to this problem is an assignment $\sigma(\cdot)$ of robots to goals, and a fleet trajectory $\tau(t) = [\tau_1(t) \dots \tau_R(t)]$.

Our method can be briefly described as follows. (A) First, we build a discrete abstraction of the MMP-GA problem, where the robots in the fleet are considered jointly in a reduced configuration space, $\mathcal{C}^K$. Each configuration in this abstract space represents the number of robots per polygon in $\mathcal{W}^{\text{free}}$. Kinodynamics are not considered and collisions between robots are considered as a cost. (B) The previously constructed abstraction is queried in order to find an abstract trajectory τ^K . This trajectory is a sequence of abstract configurations in $\mathcal{C}^K$, where the number of robots in each polygon changes in a discrete manner. At this level of abstraction, it is not possible to know which particular robot is entering or exiting a polygon, thus robots are *anonymous*. (C) An augmented trajectory τ^τ in continuous configuration space is obtained by associating each configuration in τ^K to a representative continuous configuration in $\mathcal{C}$. This is further augmented by connecting the actual start configurations of all robots, and exploring the possible de-anonymizations of robots that reach the desired goal configurations along τ^τ . (D) We exploit both τ^K and τ^τ to de-compose the problem in time and space, based on the observation that robots navigating through different polygons do not require coordination. At this point, the overall MMP-GA problem is decomposed into a partial order of simpler MMP-GA instances, one for each subset of robots for which joint motion planning is necessary. (E) Finally, a multi-robot motion planner is employed to obtain a feasible trajectory for each simpler

MMP-GA problem instance. The composite configuration space of subsets of robots and car-like kinodynamics are considered in each MMP-GA instance, and the combination of the resulting trajectories in $\mathcal{C}$ yields the overall fleet trajectory τ .

We call the abstract space constructed in phase (A) a *multi-robot roadmap*. As for other roadmap-based approaches to motion planning [29], [30], the idea is to allow some of the computation to be performed offline, i.e., without knowing the initial and final conditions (respectively, $\mathcal{S}$ and $(\sigma, \mathcal{G})$ in our approach). As a result, we can see phase (A) as the roadmap-building phase; phases (B–E) constitute the online query phase, in which the state of the actual fleet is connected to the roadmap and task assignment is performed (B–C), and kinematically feasible trajectories are computed for the robots in the fleet (D–E). A key intuition in our approach is that the roadmap also helps identify which subsets of robots need to be involved in multi-robot motion planning, and when. Next, we describe each stage in detail.

A. Multi-Robot Roadmap: a Discrete Abstraction

At the highest level of abstraction, the MMP-GA is seen as the problem of finding a sequence of polygon transitions from the fleet’s starting set of polygons to the set of goal polygons. We define an abstract fleet configuration space $\mathcal{C}^K = \{0, \dots, R\}^P$ representing the number of robots present in each polygon. Let $\lambda : \mathcal{C} \mapsto \mathcal{C}^K$ define the mapping of fleet configurations to the abstract space, that is, $\lambda(\mathbf{q}) = \mathbf{c}^K$ represents the polygon occupancy by a fleet in configuration $\mathbf{q}$. An example of this can be found in figure 1(a), where $\lambda(s) = \mathbf{c}_5^K = [1 \ 1 \ 0 \ 0]$, meaning the starting configuration has a robot in $\mathcal{P}_1$ and another robot in $\mathcal{P}_2$.

We construct offline a complete and weighted multi-robot road-map $G^K = \langle V^K, E^K \rangle$ as follows:

- $V^K = \{\mathbf{c}^K \mid \exists \mathbf{q} \in \mathcal{C}^{\text{free}} \text{ s.t. } \lambda(\mathbf{q}) = \mathbf{c}^K\}$, i.e., a vertex exists for each assignment of all robots to polygons such that their geometries are contained within the polygons;
- $(\mathbf{c}_i^K, \mathbf{c}_j^K) \in E^K$ iff exactly one robot moves from one polygon to an adjacent polygon, that is,
 - $\|\mathbf{c}_i^K - \mathbf{c}_j^K\|_1 = 2$, and
 - the polygons $\arg \max_{p \in \{1, \dots, P\}} (\mathbf{c}_j^K - \mathbf{c}_i^K)[p]$ and $\arg \min_{p \in \{1, \dots, P\}} (\mathbf{c}_j^K - \mathbf{c}_i^K)[p]$ are adjacent;

- The weights of vertices and edges are given by functions $w_v : V^\kappa \mapsto \mathbb{R}$ and $w_e : E^\kappa \mapsto \mathbb{R}$.

In the following, given an edge $(c_i^\kappa, c_j^\kappa) \in E^\kappa$, we indicate with $\text{out}(c_i^\kappa, c_j^\kappa) = \arg \min_{p \in \{1, \dots, P\}} (c_j^\kappa - c_i^\kappa)[p]$ the polygon from which the robot transitions along the edge, and with $\text{in}(c_i^\kappa, c_j^\kappa) = \arg \max_{p \in \{1, \dots, P\}} (c_j^\kappa - c_i^\kappa)[p]$ the polygon to which the robot transitions. Hence, we can indicate the border (line or point) through which the robot transitions between the two polygons with $\mathcal{P}_{\text{out}(c_i^\kappa, c_j^\kappa)} \cap \mathcal{P}_{\text{in}(c_i^\kappa, c_j^\kappa)}$.

Given the original MMP-GA problem $\langle \mathcal{W}^{\text{free}}, \mathcal{R}, \mathcal{S}, \mathcal{G} \rangle$, we can define a corresponding abstract problem where $\mathcal{W}^{\text{free}} \approx G^\kappa$, $\mathcal{R}$ is represented as number of robots per polygon, $\mathcal{S} \approx \lambda(s)$, and $\mathcal{G} \approx \lambda(g)$ (see Figure 1(a)).

B. Multi-Robot Roadmap Query with Anonymous Robots

An A^* search is conducted on graph G^κ to compute a discrete path $\tau^\kappa : \{0, \dots, N\} \mapsto \mathcal{C}^\kappa$ connecting start configuration $\lambda(s) \in \mathcal{C}^\kappa$ and goal configuration $\lambda(g) \in \mathcal{C}^\kappa$. The step-cost function used in the search is

$$c(c_i^\kappa, c_j^\kappa) = \alpha(1 - \gamma) w_v(c_j^\kappa) + \gamma w_e(c_i^\kappa, c_j^\kappa), \quad (3)$$

where function α regularizes and factor $\gamma \in [0, 1]$ prioritizes edge cost w_e versus vertex cost w_v . To account for the geometry of the polygons in computing this cost (and the heuristic used by the A^* search, as shown below), we compute the centroid $\text{ctd}(\mathcal{P})$ of a polygon $\mathcal{P}$; similarly, let $\text{ctd}(\mathcal{P}_i \cap \mathcal{P}_j)$ be the center point of the border of two connected polygons $\mathcal{P}_i$ and $\mathcal{P}_j$. We are interested in making the step cost function penalize solutions requiring complex fleet coordination or the traversal of highly constrained regions. This can be done by making w_v penalize configurations where many robots are in the same polygon, e.g., $w_v(c^\kappa) = \|c^\kappa\|_\infty - 1$.

Euclidean distances between polygons are introduced via the edge cost. Since we consider holonomic point robots at this level of abstraction, the cost of an edge is the Euclidean distance between the centroids of the polygons between which the robot moves and the centroid of their border:

$$w_e(c_i^\kappa, c_j^\kappa) = \left\| \text{ctd}(\mathcal{P}_{\text{out}(c_i^\kappa, c_j^\kappa)}) - \text{ctd}(\mathcal{P}_{\text{in}(c_i^\kappa, c_j^\kappa)} \cap \mathcal{P}_{\text{out}(c_i^\kappa, c_j^\kappa)}) \right\|_2 + \left\| \text{ctd}(\mathcal{P}_{\text{in}(c_i^\kappa, c_j^\kappa)}) - \text{ctd}(\mathcal{P}_{\text{in}(c_i^\kappa, c_j^\kappa)} \cap \mathcal{P}_{\text{out}(c_i^\kappa, c_j^\kappa)}) \right\|_2. \quad (4)$$

We can also define the heuristic function in terms of Euclidean distances. However, as G^κ is a multi-robot roadmap, we need to consider the transitions of multiple robots when estimating the distance to the goal configuration $c_g^\kappa = \lambda(g)$.

Let $\zeta^-(c^\kappa) = \{p \in \{1, \dots, P\} \mid (c^\kappa - c_g^\kappa)[p] < 0\}$ be the set of goal polygons that are missing robots in a given abstract configuration c^κ . Dually, let $\zeta^+(c^\kappa) = \{p \in \{1, \dots, P\} \mid (c^\kappa - c_g^\kappa)[p] > 0\}$ be the set of polygons with a robot surplus, i.e., more robots than goals. We can then

define the heuristic estimator as

$$h(c^\kappa) = \sum_{p^+ \in \zeta^+(c^\kappa)} \left(\gamma (c^\kappa - c_g^\kappa)[p^+] \times \min_{p^- \in \zeta^-(c^\kappa)} \left(\left\| \text{ctd}(\mathcal{P}_{p^-}) - \text{ctd}(\mathcal{P}_{p^+}) \right\|_2 \right) \right), \quad (5)$$

that is, the sum of distances between polygons with surplus robots and the closest goal polygon with missing robots, multiplied by that surplus.

The A^* search yields¹ an abstract trajectory τ^κ where robots are *anonymous*: at each discrete transition $\tau^\kappa(k)$ to $\tau^\kappa(k+1)$ a single robot exits $\text{out}(\cdot)$ and enters $\text{in}(\cdot)$, but it is not decided which robot this is. Figure 1(b) details the trajectory τ^κ as a sequence of the abstract configurations $c_5^\kappa, c_9^\kappa, c_{10}^\kappa$, that would be obtained in our running example.

C. De-Anonymization of Robots

In this phase, we de-anonymize robots along the abstract trajectory τ^κ obtained as a result of the query. To do this, we construct a weighted tree $\mathcal{T} = (V^\tau, E^\tau)$ whose root is the fleet's starting configuration $s \in \mathcal{C}$ (recall that $\lambda(s) = \tau^\kappa(0)$). The children of a node q at depth k of the tree are all configurations $q' \in \mathcal{C}$ such that $\lambda(q') = \tau^\kappa(k+1)$ and $\lambda(q) = \tau^\kappa(k)$, and each robot r is either in a start configuration ($q'[r] \in \mathcal{S}$), goal configuration ($q'[r] \in \mathcal{G}$), or a polygon centroid ($q'[r] = [\text{ctd}(\mathcal{P}_p) \ 0]$ for some $\mathcal{P}_p \in \mathcal{W}^{\text{free}}$). Expansion ceases at depth N , since $\lambda(g) = \tau^\kappa(N)$.

Let edge $(q, q') \in E^\tau$, and let r_t be the robot that transitions from polygon $\mathcal{P}_{\text{out}(\lambda(q), \lambda(q'))}$ to polygon $\mathcal{P}_{\text{in}(\lambda(q), \lambda(q'))}$. Then, the edge weight is

$$w_e^\tau(q, q') = \left\| q[r_t] - \text{ctd}(\mathcal{P}_{\text{in}(\lambda(q), \lambda(q'))} \cap \mathcal{P}_{\text{out}(\lambda(q), \lambda(q'))}) \right\|_2 + \left\| q'[r_t] - \text{ctd}(\mathcal{P}_{\text{in}(\lambda(q), \lambda(q'))} \cap \mathcal{P}_{\text{out}(\lambda(q), \lambda(q'))}) \right\|_2 + \sum_{r \neq r_t} \|q'[r] - q[r]\|_2, \quad (6)$$

which is the distance as computed in w_e for robot r_t plus the distance of robots moving inside their initial polygon.

A path from root to leaf in $\mathcal{T}$ is a sequence of transitions for all robots in the fleet which leads all goals to be reached by a robot. This path is therefore a fleet trajectory $\tau^\tau = \prod_{r=1}^R \tau_r^\tau$. Each τ_r^τ is a trajectory for a single robot, starting in $\tau_r^\tau(0) = s_r$, followed by a sequence of polygon centroids for $k = \{1, \dots, N-1\}$, and ending in a goal configuration $\tau_r^\tau(N)$. The overall fleet trajectory τ^τ implies a specific robot-goal assignment, that is, $\sigma(s_r) = \tau_r^\tau(N)$.

D. Decoupling the Problem in Time and Space

The fleet trajectory τ^τ tells us the sequence of polygons traversed by each robot. We exploit this information to decompose the original MMP-GA problem into independent sub-problems. We first partition the domain $\{0, \dots, N\}$ of τ^τ into intervals $[k_i, k_j[$ as follows:

¹Since the heuristic is admissible (as it assumes that goals are reached by the robots that are closest, and that the lengths of the corresponding paths are straight line distances between polygon centroids), τ^κ is also optimal. However, this does not entail optimality of the final fleet trajectory τ .

- within an interval $[k_i, k_j]$, all robots must move unless they are already in a goal, i.e., $\tau_r^\tau(k_i) = \tau_r^\tau(k_j) \implies \tau_r^\tau(k_j) = \sigma(s_r)$.
- every pair of robots that are in the same polygon at k_i and k_j must not be co-located in the polygon centroid.

With the resulting intervals $\mathcal{K} = \{[k_1, k_2], \dots, [k_{m-1}, k_m]\}$ we can decompose the original MMP-GA problem into a sequence of MMP-GA problems. Each problem is defined as $\mathcal{M}(k_i, k_j, \mathcal{S}) = \langle \mathcal{W}_{k_i, k_j}^{\text{free}}, \prod_{r \in \mathcal{S}} \mathcal{R}_r, \prod_{r \in \mathcal{S}} \tau_r^\tau(k_i), \prod_{r \in \mathcal{S}} \tau_r^\tau(k_j) \rangle$, where $\mathcal{W}_{k_i, k_j}^{\text{free}} \subseteq \mathcal{W}^{\text{free}}$ is the union of polygons traversed in the interval $[k_i, k_j]$ by the subset of robot $\mathcal{S} \subseteq \{1, \dots, R\}$.

We can further partition $\mathcal{M}(k_i, k_j, \mathcal{S})$ into sub-problems that involve subsets of robots traversing distinct polygons. Specifically, we construct a graph with a vertex for each robot and an edge for each pair of robots traversing the same polygon. The strongly connected components of this graph define a partition $\mathcal{S}(k_i, k_j)$ of set $\{1, \dots, R\}$. Each element $\mathcal{S}$ of this partition defines the set of robots $\bigcup_{r \in \mathcal{S}} \mathcal{R}_r$ that need to be considered jointly in motion planning. Overall, we have therefore decomposed the original problem $\mathcal{M}$ into

$$\bigoplus_{[k_i, k_j] \in \mathcal{K}} \left(\bigotimes_{\mathcal{S} \in \mathcal{S}(k_i, k_j)} \mathcal{M}(k_i, k_j, \mathcal{S}) \right), \quad (7)$$

where the symbol $\bigoplus$ refers sequential problem resolution and $\bigotimes$ refers to parallel problem resolution.

Figure 1(c) shows an example where we decomposed the original MMP-GA problem into two MMP-GAs: $\mathcal{M}([0, 2], r_2)$ where robot r_2 traverses polygons $\{\mathcal{P}_2, \mathcal{P}_3\}$, and $\mathcal{M}([0, 2], r_1)$ where robot r_1 traverses polygons $\{\mathcal{P}_1, \mathcal{P}_4\}$, both with the same horizon $[0, 2]$.

E. Multi-robot Motion Planning

For a given (k_i, k_j) , each $\mathcal{M}(k_i, k_j, \mathcal{S})$ can now be solved in parallel via some multi-robot motion planner which can compute a dynamically feasible trajectory τ for the multi-robot fleet. The planner will have as input a sequence of polygons, a start and goal configuration, kinodynamic constraints $\cup_{i \in \mathcal{S}} f_i$, continuous-time planning horizon $[\frac{k_i}{N}, \frac{k_j}{N}]$, and boundary conditions $\tau^\tau(k_i)$ and $\tau^\tau(k_j)$. As output, it will return a trajectory τ where non-holonomic robots traverse those polygons from start to goal configuration while not colliding with each other.

V. EXPERIMENTS AND RESULTS

The following empirical evaluation aims to (1) assess the scalability of the approach in terms of number of robots and polygons, and (2) assess the quality of the obtained fleet trajectories in terms of overall fleet performance.

A. Setup

The approach is tested on a PC running Ubuntu 16.04 equipped with an 8-thread Intel Core i7-6820HQ CPU @ 2.70GHz. $\mathcal{W}^{\text{free}}$ was constructed manually via the Boost Polygon Library (BPL) [31] and $G_{\mathcal{W}^{\text{free}}}$ was computed via the connectivity extraction algorithm. Both $G^{\mathcal{K}}$ and $\mathcal{T}$ were

Fig. 2: An example solution is shown for scenario 1 (a) and scenario 2 (b). On the right, the computation time of the query phase (B–E) vs. number of robots for each scenario.

generated via the Boost Graph Library (BGL) [32], and off-the-shelf A* and strongly connected component finding algorithms were utilized. The multi-robot motion planner used (Section IV-E) is the trajectory optimization based planner described in [13].

B. Scalability

R	1	2	4	6	8	10
$ V_{\mathcal{K}} $	10	55	715	4985	23860	88135
$ E_{\mathcal{K}} $	24	210	5280	47928	270540	1120136
time (s)	0.0002	0.0019	0.085	2.45	46	698

TABLE I: Roadmap building phase (Phase A) of scenario 1.

We test scalability using the scenarios depicted in Figure 2. Robots are non-holonomic with car-like kinodynamics, see eq. (1). Figure 2 shows the computation time of the query phase (B–E) vs. the number of robots in the scenario. Each asterisk represents the average solution computation time of a set of 100 runs. The top and bottom lines indicate, respectively, the 90- and 10-percentile. Table I summarizes how the size and time required to compute the roadmap scale with fleet size in scenario 1. Note that $|V_{\mathcal{K}}| \leq \binom{P+R}{P}$.

C. Solution Quality

To evaluate the quality of the computed trajectories, we simulate the execution of obtained fleet trajectories via the multi-robot coordination framework described in [5]. We measure the following metrics: sum of completion times (m_1), maximum completion time (aka makespan, m_2), sum of path lengths (m_3), and maximum path length (m_4).

1) *Goal Allocation and Obstacle Free Space $\mathcal{W}^{\text{free}}$* : As observed in Section IV-C, finding a robot-goal allocation σ is a consequence of computing τ^τ . To assess the quality of the assignment σ , we compare it to the assignment σ_d that minimizes the sum of Euclidean distances between robots and goals, i.e., minimize $\sum_{r=1}^R \sum_{g=1}^G d_{rg} \cdot \|s[r] - g[g]\|_2$, where $d_{rg} \in \{0, 1\}$, $\sigma_d(s_r) = g_g$ iff $d_{rg} = 1$, and each robot (goal) is constrained to be assigned to only one goal (robot). Since σ_d is not aware of the environment, for scenarios such as the one depicted in Figure 3 we expect to obtain better solution quality with our method. This is confirmed by our experimental results: as shown in Figure 3(b), we obtain a significant improvement in terms of m_1 , m_2 and m_4 . The naive assignment and our approach score similarly on metric m_3 (sum of path length). This is not surprising — for instance, if we assume that there are two robots $\mathcal{R}_1$

Fig. 3: [(a) Left] trajectories obtained with our approach (σ); [(a) Right] trajectories obtained with σ_d . N.b. the yellow robot is assigned the dark blue robot goal in σ_d , which is closer to the yellow’s start configuration if the $\mathcal{W}^{\text{free}}$ is not considered. (b) Average metrics of a set of 100 experiments with a fleet of 5 robots on the scenario in (a).

and $\mathcal{R}_2$ initially in poses s_1 and s_2 , two goals g_1 and g_2 , and that these are all located along the X axis in the order s_1, s_2, g_1, g_2 ; then, the assignment $\sigma(s_1) = g_2$ and $\sigma(s_2) = g_1$, which requires $\mathcal{R}_1$ to overtake $\mathcal{R}_2$, will measure similarly to other assignments on m_3 , but better on m_1 , m_2 and m_4 .

Fig. 4: (a) Average metrics of a set of 100 experiments with a fleet of 4 robots. (b) Run time divided by each phase of the algorithm as described in sections in IV, note that y-axis is broken at 3% since phase of section E will take the remaining time. (c) query phase B–E run time in both discretizations.

2) *Polygon Discretization of $\mathcal{W}^{\text{free}}$* : A further set of 100 experiments was conducted in Scenario 1 (see Figure 2) to assess the effect of the number of polygons into which $\mathcal{W}^{\text{free}}$ is discretized. We obtained a more granular discretization by subdividing each of the original polygons into two equally-sized polygons along the longer axis. The results for the same setting of start and goal configurations are shown in Figure 4(a). We observe that higher granularity of the discretization (more polygons) leads to an improvement in metric m_3 (sum of path lengths). Furthermore, it is interesting to notice that more polygons lead to faster queries, as shown in Figure 4(c). This is because, on one hand, the higher granularity leads to an increase in the computation time of phase B compared to other phases, as shown in Figure 4(b); but on the other, this leads to the abstract trajectory $\tau^{\mathcal{K}}$ having more steps, which in turn leads to more problem decomposition in phase E. This comes at the expense of more offline computation and bigger multi-robot road-maps.

Fig. 5: Figures (a) and (b) report the average metrics improvement for $\gamma = \{20\%, 40\%, 60\%, 80\%\}$. Figure (c) reports on the average run time for the different gammas and also shows the 95-percentile and the 5-percentile lines. Figure (d) shows an example of the trajectories when $\gamma = 100\%$ on the top, and on the bottom when $\gamma = 40\%$.

3) *Cost of Coordination γ* : The factor γ in the step cost function in eq. (3) is used to balance Euclidean distance and the cost incurred by robots occupying the same polygon (interference cost). To evaluate the effect of γ , we conducted 50 experiments per γ value for a set of 5 robots in the scenario shown in Figure 5(d). Note that we are using Scenario 2 with two additional corridors, in order to allow multiple non-congested path options with longer detours. The results indicate that when interference cost is given more weight, longer paths might be chosen, see Figure 5(a). This decreases the overall time spent yielding, since robots traverse less congested regions, thus improving the makespan as shown in Figure 5(b). Note also that, for $\gamma = \{40\%, 60\%\}$, the computation time of the more complex problems is reduced, as shown by the lower 95-percentile in Figure 5(c). This occurs because robots traversing separate regions do not have to be considered jointly, thus making the computation of motion plans easier.

VI. CONCLUSIONS

We have presented a scalable roadmap based approach for solving the MMP-GA. In the future, we intend to exploit our spatio-temporal decomposition for achieving a receding-horizon fleet controller. Furthermore, given the fact that the suitability of a γ value depends on the shape of the environment robots operate in, we intend to investigate learning techniques that can appropriately estimate this value. Finally, we will investigate alternative graph search methods to A^* for obtaining the abstract multi-robot trajectory $\tau^{\mathcal{K}}$, with the aim of further reducing the computational overhead of queries to the roadmap.

Acknowledgments. This work is supported by the EU H2020 project ILIAD, the Swedish Knowledge Foundation (KKS) research profile Semantic Robots, and Vinnova projects iQMobility and AutoHauler.

REFERENCES

- [1] P. R. Wurman, R. D’Andrea, and M. Mountz, “Coordinating hundreds of cooperative, autonomous vehicles in warehouses,” *AI magazine*, vol. 29, no. 1, pp. 9–9, 2008.

- [2] M. Bennewitz, W. Burgard, and S. Thrun, "Optimizing schedules for prioritized path planning of multi-robot systems," in *Proceedings 2001 ICRA. IEEE Int. Conference on Robotics and Automation (Cat. No. 01CH37164)*, vol. 1. IEEE, 2001, pp. 271–276.
- [3] M. Erdmann and T. Lozano-Perez, "On multiple moving objects," *Algorithmica*, vol. 2, no. 1–4, p. 477, 1987.
- [4] D. Bareiss and J. van den Berg, "Generalized reciprocal collision avoidance," *Int. J. Rob. Res. (IJRR)*, vol. 34, no. 12, pp. 1501–1514, 2015.
- [5] F. Pecora, H. Andreasson, M. Mansouri, and V. Petkov, "A loosely-coupled approach for multi-robot coordination, motion planning and control," in *Proc. Int. Conf. on Automated Planning and Scheduling (ICAPS)*, 2018.
- [6] B. P. Gerkey and M. J. Mataric, "A formal analysis and taxonomy of task allocation in multi-robot systems," *Int. J. Rob. Res. (IJRR)*, vol. 23, no. 9, pp. 939–954, 2004.
- [7] C. Nam and D. A. Shell, "Assignment algorithms for modeling resource contention in multirobot task allocation," *IEEE Transactions on Automation Science and Engineering*, vol. 12, no. 3, pp. 889–900, 2015.
- [8] J. Guerrero and G. Oliver, "Physical interference impact in multi-robot task allocation auction methods," in *IEEE Workshop on Distributed Intelligent Systems: Collective Intelligence and Its Applications (DIS'06)*. IEEE, 2006, pp. 19–24.
- [9] P. Forte, A. Mannucci, H. Andreasson, and F. Pecora, "Online task assignment and coordination in multi-robot fleets," *IEEE Robotics and Automation Letters*, 2021, to appear.
- [10] J. Schulman, Y. Duan, J. Ho, A. Lee, I. Awwal, H. Bradlow, J. Pan, S. Patil, K. Goldberg, and P. Abbeel, "Motion planning with sequential convex optimization and convex collision checking," *Int. J. Rob. Res. (IJRR)*, vol. 33, no. 9, pp. 1251–1270, 2014.
- [11] T. Schouwenaars, B. De Moor, E. Feron, and J. How, "Mixed integer programming for multi-vehicle path planning," in *Proc. European Control Conf. (ECC)*. IEEE, 2001, pp. 2603–2608.
- [12] F. Augugliaro, A. P. Schoellig, and R. D'Andrea, "Generation of collision-free trajectories for a quadcopter fleet: A sequential convex programming approach," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*. IEEE, 2012, pp. 1917–1922.
- [13] J. Salvado, R. Krug, M. Mansouri, and F. Pecora, "Motion planning and goal assignment for robot fleets using trajectory optimization," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*. IEEE, 2018, pp. 7939–7946.
- [14] N. Mansard, A. DelPrete, M. Geisert, S. Tonneau, and O. Stasse, "Using a memory of motion to efficiently warm-start a nonlinear predictive controller," in *2018 IEEE Int. Conf. on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 2986–2993.
- [15] J. Yu and S. M. LaValle, "Multi-agent path planning and network flow," in *Algorithmic foundations of robotics X*. Springer, 2013, pp. 157–173.
- [16] G. Sharon, R. Stern, A. Felner, and N. R. Sturtevant, "Conflict-based search for optimal multi-agent pathfinding," *Artificial Intelligence*, vol. 219, pp. 40–66, 2015.
- [17] H. Ma, S. Koenig, N. Ayanian, L. Cohen, W. Hönig, T. Kumar, T. Uras, H. Xu, C. Tovey, and G. Sharon, "Overview: Generalizations of multi-agent path finding to real-world scenarios," *arXiv preprint arXiv:1702.05515*, 2017.
- [18] H. Ma and S. Koenig, "Optimal target assignment and path finding for teams of agents," in *Proceedings of the Int. Conf. on Autonomous Agents & Multiagent Systems*. Int. Foundation for Autonomous Agents and Multiagent Systems, 2016, pp. 1144–1152.
- [19] G. Wagner and H. Choset, "Subdimensional expansion for multirobot path planning," *Artificial Intelligence*, vol. 219, pp. 1–24, 2015.
- [20] R. Luna and K. E. Bekris, "Push and swap: Fast cooperative pathfinding with completeness guarantees," in *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI)*, 2011, pp. 294–300.
- [21] D. M. Kornhauser, G. Miller, and P. Spirakis, "Coordinating pebble motion on graphs, the diameter of permutation groups, and applications," Master's thesis, M. I. T., Dept. of Electrical Engineering and Computer Science, 1984.
- [22] A. Adler, M. De Berg, D. Halperin, and K. Solovey, "Efficient multi-robot motion planning for unlabeled discs in simple polygons," in *Algorithmic foundations of robotics XI*. Springer, 2015, pp. 1–17.
- [23] K. Solovey and D. Halperin, "k-color multi-robot motion planning," *The International Journal of Robotics Research*, vol. 33, no. 1, pp. 82–97, 2014.
- [24] W. Hönig, T. S. Kumar, L. Cohen, H. Ma, H. Xu, N. Ayanian, and S. Koenig, "Multi-agent path finding with kinematic constraints," in *Proc. Int. Conf. on Automated Planning and Scheduling (ICAPS)*, 2016, pp. 477–485.
- [25] P. Švestka and M. H. Overmars, "Coordinated path planning for multiple robots," *Robotics and autonomous systems*, vol. 23, no. 3, pp. 125–152, 1998.
- [26] D. Le and E. Plaku, "Multi-robot motion planning with dynamics via coordinated sampling-based expansion guided by multi-agent search," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1868–1875, 2019.
- [27] 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. Rob. Res. (IJRR)*, vol. 35, no. 5, pp. 501–513, 2016.
- [28] Q. Du, V. Faber, and M. Gunzburger, "Centroidal voronoi tessellations: Applications and algorithms," *SIAM review*, vol. 41, no. 4, pp. 637–676, 1999.
- [29] L. E. Kavraki, P. Švestka, J.-C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE transactions on Robotics and Automation*, vol. 12, no. 4, pp. 566–580, 1996.
- [30] P. Bhattacharya and M. L. Gavrilova, "Roadmap-based path planning—using the voronoi diagram for a clearance-based shortest path," *IEEE Robotics & Automation Magazine*, vol. 15, no. 2, pp. 58–66, 2008.
- [31] L. Simonson and G. Suto, "Geometry template library for stl-like 2d operations," *Colorado: GTL Boostcon*, vol. 4, 2009.
- [32] K. Beevers and J. Peng, "A* graph search within the bgl framework," *Boost Graph Library 1.33.0* (http://www.cs.rpi.edu/~beevek/research/astar_bgl04.pdf October 2004), 2003.