

Collision Avoidance Based on Robust Lexicographic Task Assignment

Tony A. Wood, Mitchell Khoo, Elad Michael, Chris Manzie, and Iman Shames

Abstract—Traditional task assignment approaches for multi-agent motion control do not take the possibility of collisions into account. This can lead to challenging requirements for path planning. We derive an assignment method that not only minimises the largest distance between an agent and its assigned destination but also provides local constraints for guaranteed collision avoidance. To this end, we introduce a sequential bottleneck optimisation problem and define a notion of robustness of an optimising assignment to changes of individual assignment costs. Conditioned on a sufficient level of robustness in relation to the size of the agents, we construct time-varying position bounds for every individual agent. These local constraints are a direct byproduct of the assignment procedure and only depend on the initial agent positions, the destinations that are to be visited, and a timing parameter. We prove that no agent that is assigned to move to one of the target locations collides with any other agent if all agents satisfy their local position constraints. We demonstrate the method in a illustrative case study.

I. INTRODUCTION

For autonomous systems with multiple agents and multiple tasks the first decision to be made is how to allocate the tasks to the agents. This type of decision is referred to as an assignment problem and may involve any of a variety of objectives. In [1], [2] the sum of individual costs incurred for assigning robotic agents to tasks is minimised in a so called Linear Sum Assignment Problem (LSAP). An overview of assignment problems is given in [3]–[5].

We focus here on a specific type of assignment problem where the largest incurred individual agent-to-task assignment cost is minimised. This objective is referred to as the Bottleneck Assignment Problem (BAP). It commonly applies to multi-agent problems where tasks are completed in parallel and the overall completion time is of interest, as in [6] for instance. A Lexicographic Bottleneck Assignment Problem (LexBAP) describes a subclass of the BAP where not only the largest assignment cost but also all other assignment costs are minimised with a sequence of decreasing hierarchy [7]. A review of the state-of-the-art methods to solve the BAP is provided in [5]. In [8] an algorithm to solve the BAP with distributed computation is introduced. Global sensitivity of the bottleneck optimising assignment with respect to changes in the assignment costs is investigated in [9].

In applications where the tasks involve agents moving towards desired locations, collisions between agents may

occur. Motion control with inter-agent collision avoidance is a heavily researched problem and many different strategies exist, e.g., reciprocal collision avoidance [10], [11], barrier certificates [12], buffered Voronoi cells [13], and multi-agent rapidly-exploring random trees [14]. In [15] independent trajectories for all possible agent-destination pairs are designed, tasks are allocated by selecting an assignment that minimises the lexicographic cost of the chosen trajectories, and collision are avoided by subsequently delaying the starting times of certain agents. Collision avoidance awareness criteria are included in a task assignment problem in [16] by designing a quadratic assignment cost function that incentivises collision-free paths.

Similar to [17]–[19], we investigate intrinsic properties of assignment problems that provide conditions for which inter-agent collisions are avoided. In [17] it was first shown that straight lines connecting agents to allocated destinations do not intersect if the allocation corresponds to the solution of an LSAP, with Euclidean distances between initial agent positions and target destinations as assignment costs. Conditioned on sufficiently large distances between initial positions and target locations, a time parametrisation of straight-line trajectories is derived in [19] such that agents with finite extent do not collide when following an assignment that solves an LSAP, with squared distances as assignment costs. We consider a case where the largest distance between any agent and its assigned destination is to be minimised. The non-intersection properties of the LSAP do not hold for the BAP. However, assuming agents are point-masses that travel with constant velocity on straight paths towards their allocated destinations, it is shown in [18] that no collisions occur if the task allocation is based on either a LexBAP, with distances as assignment costs, or a designated assignment problem which combines a BAP, with distances as costs, and an LSAP, with squared distances as costs. We assume that agents have finite extent, consider arbitrary distance metrics, and do not limit the trajectories to straight lines.

The main contribution of this paper is the derivation of local position constraints for every agent. Rather than determining trajectories for all agents, the method proposed here provides time-varying sets of positions that guarantee collision avoidance but leave some degree of freedom for low-level path planning and motion control. In particular, we propose a sequential bottleneck task assignment approach and show that it produces the unique optimising solution of a LexBAP if robustness conditions are satisfied. We construct the local constraints that prevent collisions by quantifying the robustness to changes of the distances to the destinations. The constraints for an individual agent do not explicitly

This work was supported by Defence Science and Technology Group through research agreements MyIP:7558, MyIP:7562, and MyIP:9156.

The authors are with the Control and Signal Processing Group, Department of Electrical and Electronic Engineering, University of Melbourne, Parkville, VIC 3010, Australia. Email: {wood.t, manziec, iman.shames}@unimelb.edu.au and {khoom1, eladm}@student.unimelb.edu.au.

depend on the positions of the other agents. Information on other agents only gets accounted for via the assignment costs. We prove that it is sufficient for every agent to satisfy its local constraints in order to guarantee that no agent that is assigned to one of the tasks will collide with any other agent.

The paper is structured as follows. We formulate the assignment and collision avoidance objectives in Section II. In Section III we introduce the assignment procedure and investigate its properties. These properties are applied to derive sufficient conditions for collision avoidance in Section IV. We illustrate the resulting agent position constraints in a case study in Section V before concluding in Section VI.

II. PROBLEM FORMULATION

We consider m agents, $\mathcal{A} := \{\alpha_l\}_{l=1}^m$, with initial positions, $p_i(0) \in \mathbb{R}^{n_p}$, $n_p \in \mathbb{N}$, $i \in \mathcal{A}$, and n target destinations, $g_j \in \mathbb{R}^{n_p}$, $j \in \mathcal{T}$, where $\mathcal{T} := \{\tau_l\}_{l=1}^n$ and τ_l represents the task of going to destination g_l . Without loss of generality we assume there are at least as many agents as tasks, $m \geq n$, and consider only non-trivial cases where there are multiple agents, $m > 1$, and at least one task, $n \geq 1$.

We define a set of binary variables, $\Pi := \{\pi_{i,j} \in \{0, 1\} \mid (i, j) \in \mathcal{A} \times \mathcal{T}\}$, that indicates agent $i \in \mathcal{A}$ being assigned to task $j \in \mathcal{T}$ if and only if $\pi_{i,j} = 1$. We call such a set an *assignment* and denote the set of all assignments by $\mathbb{B}_{\mathcal{A}, \mathcal{T}}$. We assume that a cost is incurred when an agent proceeds to visit an assigned location.

Assumption 1: The cost of assigning an agent, $i \in \mathcal{A}$, to complete a task, $j \in \mathcal{T}$, is given by the distance of the initial agent position to the corresponding destination, $d(p_i(0), g_j)$, where $d : \mathbb{R}^{n_p} \times \mathbb{R}^{n_p} \rightarrow [0, \infty)$ is an arbitrary distance function that satisfies the triangle inequality,

$$d(p, p') \leq d(p, p'') + d(p'', p'), \quad (1)$$

for all $p, p', p'' \in \mathbb{R}^{n_p}$.

At any given time, t , the centroid location of agent $i \in \mathcal{A}$ is given by a point, $p_i(t)$. However, the body of the agent occupies a finite volume in \mathbb{R}^{n_p} . If two agents lie in close proximity, they may collide.

Definition 1 (Inter-agent safety distance): Agents $i, i' \in \mathcal{A}$, $i \neq i'$, do not collide with each other at time t if

$$d(p_i(t), p_{i'}(t)) > s_{i,i'}, \quad (2)$$

where $s_{i,i'} = s_{i',i} \geq 0$ is known and named the *safety distance* between i and i' .

We note that the collision avoidance condition in (2) couples the motion control problem of agents i and i' . Satisfying this condition introduces a non-convex constraint with respect to the positions of the agents, $p_i(t)$ and $p_{i'}(t)$, and is therefore challenging. The objective considered in this paper is to derive local position constraints for each individual agent, that sufficiently guarantee avoidance of collisions of agents that are assigned to tasks such that the maximum agent-to-destination distance, the so-called *bottleneck*, is minimised.

Problem 1: Suppose that one agent is assigned to every target destination such that the largest individual cost is

minimised according to the BAP,

$$\min_{\Pi \in \mathbb{B}_{\mathcal{A}, \mathcal{T}}} \max_{(i,j) \in \mathcal{A} \times \mathcal{T}} \pi_{i,j} d(p_i(0), g_j) \quad (3a)$$

$$\text{s.t.} \quad \sum_{i \in \mathcal{A}} \pi_{i,j} = 1 \quad \forall j \in \mathcal{T}, \quad (3b)$$

$$\sum_{j \in \mathcal{T}} \pi_{i,j} \leq 1 \quad \forall i \in \mathcal{A}. \quad (3c)$$

Find time-varying sets of safe positions, $\mathcal{L}_i(t) \subset \mathbb{R}^{n_p}$, for all agents, $i \in \mathcal{A}$, such that no agent assigned to a task collides with any other agent over a time interval, $t \in [0, T]$, i.e., if $p_i(t) \in \mathcal{L}_i(t)$ for all $i \in \mathcal{A}$, $t \in [0, T]$, then $d(p_{i^*}(t), p_{i'}(t)) > s_{i^*, i'}$, for all $i^* \in \{i \in \mathcal{A} \mid \sum_{j \in \mathcal{T}} \pi_{i,j} = 1\}$, $i' \in \mathcal{A}$, $t \in [0, T]$.

III. TASK ASSIGNMENT

In this section we introduce task assignment criteria and derive properties that will be applied to obtain collision avoidance conditions in Section IV. We illustrate details of the assignment procedure with an example in Section III-C.

Given a set of agents, \mathcal{A} , and a set of tasks, \mathcal{T} , where $|\mathcal{A}| = m > 1$ and $m \geq |\mathcal{T}| = n \geq 1$, we define the complete bipartite *assignment graph* $\mathcal{G} := (\mathcal{A}, \mathcal{T}, \mathcal{E})$, with vertex set, $\mathcal{V} := \mathcal{A} \cup \mathcal{T}$, and edge set, $\mathcal{E} := \mathcal{A} \times \mathcal{T}$. We also define the set of *assignment weights*, $\mathcal{W} := \{w_{i,j} \geq 0 \mid (i, j) \in \mathcal{E}\}$. An assignment, $\Pi \in \mathbb{B}_{\mathcal{A}, \mathcal{T}}$, is an *admissible* allocation of a subset of tasks, $\bar{\mathcal{T}} \subseteq \mathcal{T}$, to a subset of agents, $\bar{\mathcal{A}} \subseteq \mathcal{A}$, with respect to edge subset, $\hat{\mathcal{E}} \subseteq \mathcal{E} := \bar{\mathcal{A}} \times \bar{\mathcal{T}}$, if all considered tasks in $\bar{\mathcal{T}}$, are assigned to one agent in $\bar{\mathcal{A}}$ and all these agents are assigned to at most one task. The set of admissible assignments for subgraph $\hat{\mathcal{G}} := (\bar{\mathcal{A}}, \bar{\mathcal{T}}, \hat{\mathcal{E}})$ is given by,

$$\mathcal{P}_{\bar{\mathcal{A}}, \bar{\mathcal{T}}}(\hat{\mathcal{E}}) := \left\{ \Pi \in \mathbb{B}_{\mathcal{A}, \mathcal{T}} \mid \forall j \in \bar{\mathcal{T}} \sum_{i \in \{\bar{\mathcal{A}} \mid (i, j) \in \hat{\mathcal{E}}\}} \pi_{i,j} = 1, \right. \\ \left. \forall i \in \bar{\mathcal{A}} \sum_{j \in \{\bar{\mathcal{T}} \mid (i, j) \in \hat{\mathcal{E}}\}} \pi_{i,j} \leq 1 \right\}.$$

A. Bottleneck Assignment

The BAP formulated in (3) can be solved efficiently in polynomial time, see [5]. Independent of how the BAP is solved, we define operators associated to it in the following.

Definition 2 (Bottleneck assignment): We consider a subgraph of the assignment graph, $\hat{\mathcal{G}} = (\bar{\mathcal{A}}, \bar{\mathcal{T}}, \hat{\mathcal{E}})$, the assignment weights, \mathcal{W} , and an assignment, $\Pi \in \mathbb{B}_{\mathcal{A}, \mathcal{T}}$. We define a function that returns the largest value among the weights corresponding to assigned edges in $\hat{\mathcal{E}}$,

$$b(\Pi, \hat{\mathcal{E}}, \mathcal{W}) := \max_{(i,j) \in \hat{\mathcal{E}}} \pi_{i,j} w_{i,j}, \quad (4)$$

a function that returns the *bottleneck weight* of subgraph $\hat{\mathcal{G}}$,

$$\mathcal{B}_{\bar{\mathcal{A}}, \bar{\mathcal{T}}}(\hat{\mathcal{E}}, \mathcal{W}) := \min_{\Pi \in \mathcal{P}_{\bar{\mathcal{A}}, \bar{\mathcal{T}}}(\hat{\mathcal{E}})} b(\Pi, \hat{\mathcal{E}}, \mathcal{W}), \quad (5)$$

a map that returns the *bottleneck minimising assignments*,

$$\mathcal{B}_{\bar{\mathcal{A}}, \bar{\mathcal{T}}}(\hat{\mathcal{E}}, \mathcal{W}) := \arg \min_{\Pi \in \mathcal{P}_{\bar{\mathcal{A}}, \bar{\mathcal{T}}}(\hat{\mathcal{E}})} b(\Pi, \hat{\mathcal{E}}, \mathcal{W}),$$

and the set of edges with weight equal to the bottleneck

$$E_{\bar{\mathcal{A}}, \bar{\mathcal{T}}}(\hat{\mathcal{E}}, \mathcal{W}) := \{(i, j) \in \hat{\mathcal{E}} \mid w_{i,j} = B_{\bar{\mathcal{A}}, \bar{\mathcal{T}}}(\hat{\mathcal{E}}, \mathcal{W})\}. \quad (6)$$

Next, we quantify how sensitive an optimising assignment is to changes of the bottleneck weight. Specifically, we define the robustness margin that captures the difference between the optimal bottleneck cost and the cost of the next best assignment when a specific bottleneck edge is removed. In Section IV we use this information to determine by how much the positions of agents can vary such that the optimality of the assignment is maintained and collisions do not occur.

Definition 3 (Robustness margin): Given the assignment weights, \mathcal{W} , we consider the complete bipartite graph formed by a subset of agents, $\bar{\mathcal{A}}$, and a subset of tasks, $\bar{\mathcal{T}}$, i.e. subgraph $\bar{\mathcal{G}} := (\bar{\mathcal{A}}, \bar{\mathcal{T}}, \bar{\mathcal{E}})$, with edge set $\bar{\mathcal{E}} = \bar{\mathcal{A}} \times \bar{\mathcal{T}}$. For $|\bar{\mathcal{E}}| > 1$, we define the set of so-called *maximum-margin bottleneck edges* as,

$$e_{\bar{\mathcal{A}}, \bar{\mathcal{T}}}(\mathcal{W}) := \arg \max_{(i,j) \in E_{\bar{\mathcal{A}}, \bar{\mathcal{T}}}(\bar{\mathcal{E}}, \mathcal{W})} B_{\bar{\mathcal{A}}, \bar{\mathcal{T}}}(\bar{\mathcal{E}} \setminus \{(i, j)\}, \mathcal{W}),$$

and the corresponding *robustness margin*,

$$r_{\bar{\mathcal{A}}, \bar{\mathcal{T}}}(\mathcal{W}) := \max_{(i,j) \in E_{\bar{\mathcal{A}}, \bar{\mathcal{T}}}(\bar{\mathcal{E}}, \mathcal{W})} B_{\bar{\mathcal{A}}, \bar{\mathcal{T}}}(\bar{\mathcal{E}} \setminus \{(i, j)\}, \mathcal{W}) - w_{i,j}.$$

For $|\bar{\mathcal{A}}| = |\bar{\mathcal{T}}| = |\bar{\mathcal{E}}| = 1$, the maximum-margin bottleneck edge is set to be the singleton edge $e_{\bar{\mathcal{A}}, \bar{\mathcal{T}}}(\mathcal{W}) = \bar{\mathcal{E}}$ and the robustness margin is assumed to be infinity, $r_{\bar{\mathcal{A}}, \bar{\mathcal{T}}}(\mathcal{W}) = \infty$.

The following proposition links the robustness margin to uniqueness properties of bottleneck minimising assignments. While a strictly positive robustness margin does not imply that there is a unique bottleneck minimising assignment, it does ensure that all bottleneck minimising assignments have in common that they assign all agent-task pairs in the set of maximum-margin bottleneck edges.

Proposition 1 (Proof in Appendix I): Given the assignment weights, \mathcal{W} , we consider the complete bipartite graph, $\bar{\mathcal{G}} = (\bar{\mathcal{A}}, \bar{\mathcal{T}}, \bar{\mathcal{E}})$, formed by a subset of agents, $\bar{\mathcal{A}} \in \mathcal{A}$, and a subset of tasks, $\bar{\mathcal{T}} \in \mathcal{T}$, with edge set $\bar{\mathcal{E}} = \bar{\mathcal{A}} \times \bar{\mathcal{T}}$, where $|\bar{\mathcal{E}}| > 1$. If the robustness margin is strictly positive, $r_{\bar{\mathcal{A}}, \bar{\mathcal{T}}}(\mathcal{W}) > 0$, then all bottleneck minimising assignments for subgraph $\bar{\mathcal{G}}$ assign all maximum-margin bottleneck edges, i.e., $\pi_{i^*, j^*} = 1, \forall (i^*, j^*) \in e_{\bar{\mathcal{A}}, \bar{\mathcal{T}}}(\mathcal{W}), \forall \Pi = \{\pi_{i,j} \in \{0, 1\} \mid (i, j) \in \bar{\mathcal{E}}\} \in \mathcal{B}_{\bar{\mathcal{A}}, \bar{\mathcal{T}}}(\bar{\mathcal{E}}, \mathcal{W})$.

B. Sequential Bottleneck Assignment

The BAP does not fully determine all agent-task pairings. Given the primary assignment criterion introduced in (3), we define a sequence of criteria of decreasing hierarchy.

Definition 4 (Sequential bottleneck assignment): We consider the agents, \mathcal{A} , the tasks, \mathcal{T} , and the assignment weights, \mathcal{W} . An assignment, Π^* , is *sequential bottleneck optimising* if it is bottleneck minimising for the assignment graph, $\mathcal{G} = (\mathcal{A}, \mathcal{T}, \mathcal{E})$, and the sequence of subgraphs $\bar{\mathcal{G}}_2, \bar{\mathcal{G}}_3, \dots, \bar{\mathcal{G}}_n$, where $\bar{\mathcal{G}}_k = (\bar{\mathcal{A}}_k, \bar{\mathcal{T}}_k, \bar{\mathcal{E}}_k)$ is the complete bipartite graph of the subset of agents, $\bar{\mathcal{A}}_k \subset \mathcal{A}$, and the subset of tasks, $\bar{\mathcal{T}}_k \subset \mathcal{T}$, obtained by removing a maximum-margin bottleneck agent and task from $\bar{\mathcal{G}}_{k-1}$, i.e., $\Pi^* \in \mathcal{S}_{\mathcal{A}, \mathcal{T}}(\mathcal{W})$,

$$\mathcal{S}_{\mathcal{A}, \mathcal{T}}(\mathcal{W}) := \{\Pi \in \mathbb{B}_{\mathcal{A}, \mathcal{T}} \mid \forall k \in \{1, \dots, n\} \Pi \in \mathcal{B}_{\bar{\mathcal{A}}_k, \bar{\mathcal{T}}_k}(\bar{\mathcal{E}}_k, \mathcal{W})\},$$

where $\bar{\mathcal{E}}_k = \bar{\mathcal{A}}_k \times \bar{\mathcal{T}}_k, \bar{\mathcal{A}}_1 = \mathcal{A}, \bar{\mathcal{T}}_1 = \mathcal{T}$,

$$\bar{\mathcal{A}}_k = \mathcal{A} \setminus \{i_1^*, \dots, i_{k-1}^*\}, \quad (7a)$$

$$\bar{\mathcal{T}}_k = \mathcal{T} \setminus \{j_1^*, \dots, j_{k-1}^*\}, \quad (7b)$$

with so-called *k-th order bottleneck edge*,

$$(i_k^*, j_k^*) \in e_{\bar{\mathcal{A}}_k, \bar{\mathcal{T}}_k}(\mathcal{W}), \quad (8)$$

and *k-th order robustness margin*,

$$\mu_k = r_{\bar{\mathcal{A}}_k, \bar{\mathcal{T}}_k}(\mathcal{W}). \quad (9)$$

Given a sequential bottleneck assignment, $\Pi^* \in \mathcal{S}_{\mathcal{A}, \mathcal{T}}(\mathcal{W})$, the set of all assigned agents is defined by the set of *k-th order bottleneck agents* of all orders, $\{i_1^*, \dots, i_n^*\} \subseteq \mathcal{A}$. For the case where there are more agents than tasks, $m > n$, the set of unassigned agents is $\mathcal{A} \setminus \{i_1^*, \dots, i_n^*\}$. The following two propositions show that the assignment of the *k-th order bottleneck edge* is invariant to an additive weight perturbation that is smaller than the *k-th order robustness margin*. We use this property to link the robustness margins to bounds on the deviation of the agent positions that guarantee collision avoidance in Section IV.

Proposition 2 (Proof in Appendix II): Given the agents, \mathcal{A} , with $|\mathcal{A}| = m > 1$, the tasks, \mathcal{T} , with $m \geq |\mathcal{T}| = n \geq 1$, and the assignment weights, $\mathcal{W} = \{w_{i,j} \geq 0 \mid (i, j) \in \mathcal{A} \times \mathcal{T}\}$, we consider a sequential bottleneck optimising assignment, $\Pi^* \in \mathcal{S}_{\mathcal{A}, \mathcal{T}}(\mathcal{W})$, with *k-th order bottleneck edges*, (i_k^*, j_k^*) , and robustness margins, μ_k , for $k \in \{1, \dots, n\}$, introduced in Definition 4. We have

$$w_{i_a^*, j_a^*} + \mu_a \leq \max\{w_{i_a^*, j_b^*}, w_{i_b^*, j_a^*}\},$$

for all $a \in \{1, \dots, n-1\}$ and $b \in \{a+1, \dots, n\}$.

Proposition 3 (Proof in Appendix III): Given the agents, \mathcal{A} , with $|\mathcal{A}| = m > 1$, the strictly less tasks, \mathcal{T} , with $m > |\mathcal{T}| = n \geq 1$, and the assignment weights, $\mathcal{W} = \{w_{i,j} \geq 0 \mid (i, j) \in \mathcal{A} \times \mathcal{T}\}$, we consider a sequential bottleneck optimising assignment, $\Pi^* \in \mathcal{S}_{\mathcal{A}, \mathcal{T}}(\mathcal{W})$, with *k-th order bottleneck edges*, (i_k^*, j_k^*) , and robustness margins, μ_k , for $k \in \{1, \dots, n\}$, introduced in Definition 4. We have

$$w_{i_a^*, j_a^*} + \mu_a \leq w_{i', j_a^*},$$

for all $a \in \{1, \dots, n\}$ and $i' \in \mathcal{A} \setminus \{i_1^*, \dots, i_n^*\}$.

We note that the *k-th order bottleneck edge* defined in (8) may not be unique. However, if the robustness margin is strictly positive for all orders, it follows from Proposition 1 that all possible *k-th order bottleneck edges* lead to a unique sequentially bottleneck optimising assignment.

Corollary 1: We consider a sequential bottleneck optimising assignment, $\Pi^* \in \mathcal{S}_{\mathcal{A}, \mathcal{T}}(\mathcal{W})$. If the *k-th order robustness margin* is strictly positive, $r_{\bar{\mathcal{A}}_k, \bar{\mathcal{T}}_k}(\mathcal{W}) > 0$, for all $k \in \{1, \dots, n\}$, then the sequence of bottleneck weight values, $(w_{i_1^*, j_1^*}, \dots, w_{i_n^*, j_n^*})$, is unique and lexicographically optimal, i.e., for all other assignments, $\Pi' \in \mathcal{P}_{\mathcal{A}, \mathcal{T}}(\mathcal{E}) \setminus \{\Pi^*\}$, either $w_{i_1^*, j_1^*} < w_{i'_1, j'_1}$ or there exists an order, $l \in \{2, \dots, n\}$, such that $w_{i_k^*, j_k^*} = w_{i'_k, j'_k}$ for $k \in \{1, \dots, l-1\}$ and $w_{i_l^*, j_l^*} < w_{i'_l, j'_l}$, where $(w_{i'_1, j'_1}, \dots, w_{i'_n, j'_n})$ is a sequence of weights associated to assigned edges according to Π' that are ordered with non-increasing values.




Fig. 1: Subgraphs, $\bar{\mathcal{G}}_k = (\bar{\mathcal{A}}_k, \bar{\mathcal{T}}_k, \bar{\mathcal{E}}_k)$, with edge weights, \mathcal{W} , for sequential bottleneck assignment of tasks, $\mathcal{T} = \{\tau_l\}_{l=1}^3$ [green nodes], to agents, $\mathcal{A} = \{\alpha_l\}_{l=1}^4$ [blue nodes]. For each order, $k \in \{1, 2, 3\}$, a maximum-margin bottleneck edge, $(i_k^*, j_k^*) \in e_{\bar{\mathcal{A}}_k, \bar{\mathcal{T}}_k}(\mathcal{W})$ [red lines], and a critical edge associated with the robustness margin, $(i_k^c, j_k^c) \in E_{\bar{\mathcal{A}}_k, \bar{\mathcal{T}}_k}(\bar{\mathcal{E}}_k \setminus \{(i_k^*, j_k^*)\}, \mathcal{W})$ [orange lines], are shown. A bottleneck minimising assignment, $\hat{\Pi} \in \mathcal{B}_{\mathcal{A}, \mathcal{T}}(\mathcal{E}, \mathcal{W})$ [dotted], and the sequential bottleneck optimising assignment, $\Pi' \in \mathcal{S}_{\mathcal{A}, \mathcal{T}}(\mathcal{W}) \subset \mathcal{B}_{\mathcal{A}, \mathcal{T}}(\mathcal{E}, \mathcal{W})$ [dashed], are highlighted.

We call a sequential bottleneck optimising assignment with strictly positive robustness margins for all orders a *robust lexicographic assignment*. Such an assignment is a special case solution of a LexBAP with a unique optimiser. Finding the k -th order bottleneck and the corresponding robustness margin requires solving two BAPs that can each be computed with a complexity of $\mathcal{O}(|\bar{\mathcal{E}}_k| |\bar{\mathcal{T}}_k|)$, see [8]. From (7) it follows that $|\bar{\mathcal{E}}_k| = (m - k + 1)(n - k + 1)$ and $|\bar{\mathcal{T}}_k| = n - k$. Finding the k -th order bottleneck edges and robustness margins for all $k \in \{1, \dots, n\}$ can therefore be achieved with an overall complexity of $\mathcal{O}(mn^3)$.

Remark 1: The algorithm introduced in [8] solves the BAP without centralised decision making. Agents only have knowledge of weights associated to them, i.e., agent $i \in \mathcal{A}$ has access to the subset of weights $\mathcal{W}_i := \{w_{i,j} \mid j \in \mathcal{T}\}$, and communicates local estimates of maximal and minimal edge weights to other agents. Because solving the sequential bottleneck assignment consists of solving $2n$ nested BAPs, it follows that a sequential bottleneck optimising assignment and the corresponding robustness margins can also be obtained with distributed computation.

C. Assignment Example

We consider the sequential bottleneck assignment of $n = 3$ tasks, $\mathcal{T} = \{\tau_l\}_{l=1}^3$, to $m = 4$ agents, $\mathcal{A} = \{\alpha_l\}_{l=1}^4$, with assignment weights, \mathcal{W} , illustrated in Fig. 1. The sequence of BAPs is initialised with the full assignment graph, $\bar{\mathcal{G}}_1 = \mathcal{G} = (\mathcal{A}, \mathcal{T}, \mathcal{E})$, with $\mathcal{E} = \mathcal{A} \times \mathcal{T}$, as shown in Fig. 1a. The bottleneck weight is $B_{\mathcal{A}, \mathcal{T}}(\mathcal{E}, \mathcal{W}) = 4$. There are two edges that have weight equal to the bottleneck, $E_{\mathcal{A}, \mathcal{T}}(\mathcal{E}, \mathcal{W}) = \{(\alpha_1, \tau_1), (\alpha_2, \tau_2)\}$, and one of them is the unique maximum-margin bottleneck edge, $(i_1^*, j_1^*) = (\alpha_2, \tau_2) = e_{\mathcal{A}, \mathcal{T}}(\mathcal{W})$. That is, $B_{\mathcal{A}, \mathcal{T}}(\mathcal{E} \setminus \{(\alpha_1, \tau_1)\}, \mathcal{W}) = 4$, $B_{\mathcal{A}, \mathcal{T}}(\mathcal{E} \setminus \{(\alpha_2, \tau_2)\}, \mathcal{W}) = 7$, and the first order robustness

margin is therefore $\mu_1 = r_{\mathcal{A}, \mathcal{T}}(\mathcal{W}) = 3$. Because the robustness margin is greater than zero, all bottleneck minimising assignments allocate α_2 to τ_2 . We note that there exist multiple valid assignments with maximum weight equal to the bottleneck. Assignment $\hat{\Pi}$, corresponding to the task-agent pairings $\{(\alpha_1, \tau_1), (\alpha_2, \tau_2), (\alpha_4, \tau_3)\}$, and assignment Π' , corresponding to $\{(\alpha_1, \tau_3), (\alpha_2, \tau_2), (\alpha_4, \tau_1)\}$, both result in a bottleneck weight of 4, i.e., $\hat{\Pi}, \Pi' \in \mathcal{B}_{\mathcal{A}, \mathcal{T}}(\mathcal{E}, \mathcal{W})$. These bottleneck minimising assignments differ in the higher order criteria of the sequential bottleneck assignment however.

By removing the first order bottleneck agent and task, we obtain the second order subgraph, $\bar{\mathcal{G}}_2 = (\bar{\mathcal{A}}_2, \bar{\mathcal{T}}_2, \bar{\mathcal{E}}_2)$, with $\bar{\mathcal{A}}_2 = \mathcal{A} \setminus \{\alpha_2\}$, $\bar{\mathcal{T}}_2 = \mathcal{T} \setminus \{\tau_2\}$, $\bar{\mathcal{E}}_2 = \bar{\mathcal{A}}_2 \times \bar{\mathcal{T}}_2$, shown in Fig. 1b. The second order bottleneck weight is $B_{\bar{\mathcal{A}}_2, \bar{\mathcal{T}}_2}(\bar{\mathcal{E}}_2, \mathcal{W}) = 2$. There are two edges with weight equal to the bottleneck, $E_{\bar{\mathcal{A}}_2, \bar{\mathcal{T}}_2}(\bar{\mathcal{E}}_2, \mathcal{W}) = \{(\alpha_1, \tau_3), (\alpha_4, \tau_1)\}$, and both are maximum-margin bottlenecks, $e_{\bar{\mathcal{A}}_2, \bar{\mathcal{T}}_2}(\bar{\mathcal{E}}_2, \mathcal{W}) = E_{\bar{\mathcal{A}}_2, \bar{\mathcal{T}}_2}(\bar{\mathcal{E}}_2, \mathcal{W})$, with $B_{\bar{\mathcal{A}}_2, \bar{\mathcal{T}}_2}(\bar{\mathcal{E}}_2 \setminus \{(\alpha_1, \tau_3)\}, \mathcal{W}) = B_{\bar{\mathcal{A}}_2, \bar{\mathcal{T}}_2}(\bar{\mathcal{E}}_2 \setminus \{(\alpha_4, \tau_1)\}, \mathcal{W}) = 4$. From Proposition 1 it follows that all sequential bottleneck optimal assignments involve the agent-task pairings corresponding to both edges in $e_{\bar{\mathcal{A}}_2, \bar{\mathcal{T}}_2}(\bar{\mathcal{E}}_2, \mathcal{W})$. From these two edges, we arbitrarily select the second order bottleneck edge, $(i_2^*, j_2^*) = (\alpha_4, \tau_1)$, with robustness margin, $\mu_2 = r_{\bar{\mathcal{A}}_2, \bar{\mathcal{T}}_2}(\mathcal{W}) = 2$. In the final step we consider the subgraph, $\bar{\mathcal{G}}_3 = (\bar{\mathcal{A}}_3, \bar{\mathcal{T}}_3, \bar{\mathcal{E}}_3)$, with $\bar{\mathcal{A}}_3 = \bar{\mathcal{A}}_2 \setminus \{\alpha_4\}$, $\bar{\mathcal{T}}_3 = \bar{\mathcal{T}}_2 \setminus \{\tau_1\}$, $\bar{\mathcal{E}}_3 = \bar{\mathcal{A}}_3 \times \bar{\mathcal{T}}_3$, shown in Fig. 1c. The third order bottleneck weight is, $B_{\bar{\mathcal{A}}_3, \bar{\mathcal{T}}_3}(\bar{\mathcal{E}}_3, \mathcal{W}) = 2$, with unique bottleneck edge, $(i_3^*, j_3^*) = e_{\bar{\mathcal{A}}_3, \bar{\mathcal{T}}_3}(\bar{\mathcal{E}}_3, \mathcal{W}) = e_{\bar{\mathcal{A}}_3, \bar{\mathcal{T}}_3}(\mathcal{W}) = (\alpha_1, \tau_3)$, and robustness margin, $\mu_3 = r_{\bar{\mathcal{A}}_3, \bar{\mathcal{T}}_3}(\mathcal{W}) = 6$.

The resulting sequential bottleneck optimising assignment, $\Pi^* = \Pi' \in \mathcal{S}_{\bar{\mathcal{A}}_3, \bar{\mathcal{T}}_3}(\mathcal{W})$, is unique and a robust lexicographical assignment because the robustness margins are strictly positive for all orders, i.e., every other admissible assignment, $\Pi \in \mathcal{P}_{\mathcal{A}, \mathcal{T}}(\mathcal{E}) \setminus \{\Pi^*\}$, has a lexicographically larger weight sequence than $(w_{i_1^*, j_1^*}^*, w_{i_2^*, j_2^*}^*, w_{i_3^*, j_3^*}^*) = (4, 2, 2)$. Furthermore, we observe that $w_{i_1^*, j_1^*}^* + \mu_1 < \max\{w_{i_1^*, j_2^*}^*, w_{i_2^*, j_1^*}^*\} = 8$, $w_{i_1^*, j_1^*}^* + \mu_1 < \max\{w_{i_1^*, j_3^*}^*, w_{i_3^*, j_1^*}^*\} = 9$, and $w_{i_2^*, j_2^*}^* + \mu_2 = \max\{w_{i_2^*, j_3^*}^*, w_{i_3^*, j_2^*}^*\} = 4$ in accordance with Proposition 2. We also see that $w_{i_1^*, j_1^*}^* + \mu_1 < w_{\alpha_3, j_1^*}^* = 9$, $w_{i_2^*, j_2^*}^* + \mu_2 < w_{\alpha_3, j_2^*}^* = 7$, and $w_{i_3^*, j_3^*}^* + \mu_3 = w_{\alpha_3, j_3^*}^* = 8$ in agreement with Proposition 3.

IV. COLLISION AVOIDANCE

In this section we address inter-agent collision avoidance. We first investigate sufficient conditions related to the sequential bottleneck assignment in Section IV-A. Then, in Section IV-B we introduce time-dependent position constraints for the individual agents.

A. Sufficient Conditions for Collision Avoidance

Assumption 2: The assignment weights, $\mathcal{W} = \{w_{i,j} = d(p_i(0), g_j) \mid (i, j) \in \mathcal{A} \times \mathcal{T}\}$, are the distances between initial agent positions and destinations, defined in Assumption 1. Agents, \mathcal{A} , are allocated to destinations, \mathcal{T} , with a sequential bottleneck optimising assignment, $\Pi^* \in \mathcal{S}_{\mathcal{A}, \mathcal{T}}(\mathcal{W})$,

as in Definition 4, with k -th order bottleneck agent-task pair, (i_k^*, j_k^*) , and robustness margin, μ_k , for $k \in \{1, \dots, n\}$.

Using the safety distances from Definition 1, we provide a first condition which guarantees that an assigned agent does not collide with any other agent at a particular time.

Lemma 1: Given Assumptions 1 and 2. The k -th order bottleneck agent, i_k^* , $k \in \{1, \dots, n\}$, does not collide with any other agent, $i' \in \mathcal{A} \setminus \{i_k^*\}$, at time t if

$$d(p_{i'}(0), p_{i'}(t)) + d(p_{i_k^*}(t), g_{j_k^*}) < w_{i', j_k^*} - s_{i', i_k^*}. \quad (10)$$

Proof: Considering the triangle inequality in (1), the distance between the initial position of agent i' and the target destination assigned to agent i_k^* is bounded, $d(p_{i'}(0), g_{j_k^*}) \leq d(p_{i'}(0), p_{i'}(t)) + d(p_{i'}(t), p_{i_k^*}(t)) + d(p_{i_k^*}(t), g_{j_k^*})$. By applying (10), with $w_{i', j_k^*} = d(p_{i'}(0), g_{j_k^*})$, we obtain $d(p_{i'}(t), p_{i_k^*}(t)) > s_{i', i_k^*}$ as required from Definition 1. ■

For guaranteed collision avoidance among assigned agents we combine the concept of robustness margin with the condition given in Lemma 1.

Proposition 4: Given Assumptions 1 and 2. The k -th order bottleneck agent, i_k^* , $k \in \{1, \dots, n-1\}$, does not collide with a higher order bottleneck agent, i_l^* , $l \in \{k+1, \dots, n\}$, at time t if both following conditions are satisfied,

$$d(p_{i_l^*}(0), p_{i_l^*}(t)) + d(p_{i_k^*}(t), g_{j_k^*}) < w_{i_l^*, j_k^*} + \mu_k - s_{i_l^*, i_k^*}, \quad (11a)$$

$$d(p_{i_k^*}(0), p_{i_k^*}(t)) + d(p_{i_l^*}(t), g_{j_l^*}) < w_{i_k^*, j_l^*} + \mu_k - s_{i_l^*, i_k^*}. \quad (11b)$$

Proof: For an arbitrary order, $k \in \{1, \dots, n-1\}$, let $l \in \{k+1, \dots, n\}$ be an arbitrary higher order. If $w_{i_l^*, j_k^*} > d(p_{i_l^*}(0), p_{i_l^*}(t)) + d(p_{i_k^*}(t), g_{j_k^*}) + s_{i_l^*, i_k^*}$, agents i_k^* and i_l^* do not collide at time t as shown in Lemma 1, with agent i_l^* playing the role of the 'other agent'.

It remains to consider the case where $w_{i_l^*, j_k^*} \leq d(p_{i_l^*}(0), p_{i_l^*}(t)) + d(p_{i_k^*}(t), g_{j_k^*}) + s_{i_l^*, i_k^*}$. Assume for the sake of contradiction that the condition in (10) also does not hold from perspective of agent i_l^* , where agent i_k^* is the 'other agent', i.e., $w_{i_k^*, j_l^*} \leq d(p_{i_k^*}(0), p_{i_k^*}(t)) + d(p_{i_l^*}(t), g_{j_l^*}) + s_{i_k^*, i_l^*}$. Then, from (11a) we have $w_{i_l^*, j_k^*} < w_{i_k^*, j_k^*} + \mu_k$ and from (11b) we have $w_{i_k^*, j_l^*} < w_{i_l^*, j_l^*} + \mu_k$. This however contradicts Proposition 2. It follows that $w_{i_l^*, j_k^*} > d(p_{i_l^*}(0), p_{i_l^*}(t)) + d(p_{i_k^*}(t), g_{j_k^*}) + s_{i_l^*, i_k^*}$. Thus, according to Lemma 1, agents i_k^* and i_l^* do not collide at time t . ■

B. Local Constraints for Guaranteed Collision Avoidance

We now derive individual position constraints for every agent such that if all agents satisfy their associated constraints, collisions involving assigned agents are avoided. The constraints rely on the robustness of the sequential bottleneck assignment as formalised in the following assumption.

Assumption 3: There exists an upper bound on the safety distance, $s \geq s_{i, i'}$, between all agents, $i, i' \in \mathcal{A}$, $i \neq i'$, that is smaller than the k -th order bottleneck robustness margin for all $k \in \{1, \dots, n\}$, i.e.,

$$s < \mu := \min_{k \in \{1, \dots, n\}} \mu_k. \quad (12)$$

The agent position constraints are composed of up to two components. The first component is a bound on the distance

of an agent from its initial position. This bound is limited by the same time-varying parameter, $a(t)$, for all agents up to a saturation value, A_k , that is agent dependent.

Assumption 4: The distance of an agent i , to its initial position is bounded by $a_i(t)$, i.e.,

$$d(p_i(0), p_i(t)) < a_i(t),$$

for all $i \in \mathcal{A}$ and $t \in [0, T]$, with

$$a_{i_k^*}(t) = \begin{cases} a(t) & \text{if } a(t) \leq A_k, \\ A_k & \text{otherwise,} \end{cases}$$

for all assigned agents i_k^* , $k \in \{1, \dots, n\}$, and $a_{i'}(t) = a_{i_n^*}(t)$, for all unassigned agents, $i' \in \mathcal{A} \setminus \{i_1^*, \dots, i_n^*\}$, where

$$A_k := \min_{l \in \{1, \dots, k\}} w_{i_l^*, j_l^*} + \mu_l - \frac{1}{2}(\mu + s) \quad (13)$$

and $a(t) \geq \frac{1}{2}(\mu - s)$.

The second component of the position constraints applies only to assigned agents. It consists of a bound on the distance of an agent to its target destination that decreases when the bound on the distance from the initial position increases.

Assumption 5: The distance of an assigned agent, i_k^* , to its destination is bounded by $b_{j_k^*}(t)$, i.e.,

$$d(p_{i_k^*}(t), g_{j_k^*}) < b_{j_k^*}(t),$$

for all assigned agents i_k^* , $k \in \{1, \dots, n\}$, and $t \in [0, T]$, with

$$b_{j_k^*}(t) = A_k - a_{i_k^*}(t) + \frac{1}{2}(\mu - s).$$

Satisfaction of these local bounds provides a sufficient condition for collision avoidance as shown in the following.

Theorem 1: Given n target destinations allocated to m agents with a sequential bottleneck optimising assignment, as specified in Assumption 2, with robustness margins satisfying Assumption 3, based on distance weights introduced in Assumption 1. No assigned agent, i_k^* , $k \in \{1, \dots, n\}$, collides with any other agent, $i' \in \mathcal{A}$, at any time, $t \in [0, T]$, if all agents satisfy the position bounds of Assumption 4 and all assigned agents additionally satisfy the position bounds of Assumption 5.

Proof: By construction, the bounds in Assumptions 4 and 5 satisfy $a_{i_k^*}(t) \geq a_{i_l^*}(t)$, $b_{j_k^*}(t) \geq b_{j_l^*}(t)$, and $a_{i_k^*}(t) + b_{j_k^*}(t) \leq w_{i_k^*, j_k^*} + \mu_k - s$ for all $k \in \{1, \dots, n\}$, $l \in \{k, \dots, n\}$, and $t \in [0, T]$. With Assumption 3 the conditions in Proposition 4 are satisfied for all $k \in \{1, \dots, n-1\}$ and $t \in [0, T]$. Thus, no assigned agents collide with each other. Because of Proposition 3, Lemma 1 is satisfied for all $k \in \{1, \dots, n\}$, $i' \in \mathcal{A} \setminus \{i_1^*, \dots, i_n^*\}$, and $t \in [0, T]$. Thus, assigned agents also do not collide with unassigned agents. ■

We note that there exist positions, $p_{i_k^*}(t)$, for every assigned agent, i_k^* , $k \in \{1, \dots, n\}$, at every time, $t \in [0, T]$, that satisfy Assumptions 4 and 5 if Assumption 3 holds. That is, Assumption 4 bounds the position, $p_{i_k^*}(t)$, to lie within a ball centred at the initial position, $p_{i_k^*}(0)$, with radius $a_{i_k^*}(t) \in [\frac{1}{2}(\mu - s), A_k]$ and Assumption 5 bounds the position, $p_{i_k^*}(t)$, to lie within a ball centred at the target

destination, $g_{j_k^*}$, with radius $b_{j_k^*}(t) \in [\frac{1}{2}(\mu - s), A_k]$. The two balls intersect for all $t \in [0, T]$ if $\mu > s$ because the sum of the radii is larger than the distance between the centres,

$$a_{i_k^*}(t) + b_{j_k^*}(t) \geq w_{i_k^*, j_k^*} + \mu - s > w_{i_k^*, j_k^*} = d(p_{i_k^*}(0), g_{j_k^*}).$$

It follows that the constrained sets of safe positions constructed from the bounds in Assumptions 4 and 5, i.e.,

$$\mathcal{L}_{i_k^*}(t) = \{p \in \mathbb{R}^{n_p} \mid d(p_{i_k^*}(0), p) < a_{i_k^*}(t), \quad (14a) \\ d(p, g_{j_k^*}) < b_{j_k^*}(t)\},$$

for all assigned agents i_k^* , $k \in \{1, \dots, n\}$, and

$$\mathcal{L}_{i'}(t) = \{p \in \mathbb{R}^{n_p} \mid d(p_{i'}(0), p) < a_{i'}(t)\}, \quad (14b)$$

for unassigned agents, $i' \in \mathcal{A} \setminus \{i_1^*, \dots, i_n^*\}$, are all non-empty if Assumption 3 holds.

For any agent, $i \in \mathcal{A}$, the safe set, $\mathcal{L}_i(t)$, depends on the timing parameter, $a(t)$, the k -th order bottleneck weights and robustness margins obtained in the sequential bottleneck assignment, $w_{i_k^*, j_k^*}$ and μ_k respectively, for $k \in \{1, \dots, n\}$, the initial position, $p_i(0)$, and the location of its assigned destination, g_j , if $\pi_{i,j}^* = 1$. We see that the larger the robustness margins are, the less the agent positions need to be constrained in (14). These local position constraints provide a solution for the problem outlined in Section II.

Corollary 2: The agent position constraints given in (14) can be obtained with a complexity $\mathcal{O}(mn^3)$ via a sequential bottleneck assignment. The resulting local safe sets are a solution for Problem 1 if Assumption 3 holds.

An optimising assignment, $\Pi^* \in \mathcal{S}_{\mathcal{A}, \mathcal{T}}(\mathcal{W})$, and the robustness margins are determined only based on information of the initial agent positions relative to the target destinations. Knowledge of the absolute positions of the other agents is not required to compute the individual constraints.

Remark 2: The collision avoidance constraints given in (14) can be obtained with distributed computation. Agents need to coordinate through the shared scheduling variable, $a(t)$, and by exchanging estimates of all the k -order bottleneck distances and robustness margins, for all $k \in \{1, \dots, n\}$, during the assignment as suggested in Remark 1.

The local agent position constraints proposed in (14) can be incorporated in many different motion control or trajectory planning applications. For instance, because the position bounds provide collision avoidance guarantees without fully specifying position trajectories, they can be included in predictive optimisation approaches with objective functions that do not consider the coordination among the agents. The resulting optimisation problems can incorporate additional constraints such as avoidance of other objects. The sets in (14) are convex as they are constructed from distance functions. This allows to bypass an extra procedure for convex approximation of a safe region which is typically required in model predictive motion control, see [20] for instance. For specific choices of the applied distance functions, e.g., the 1-norm or the infinity-norm distances, the constraints are linear in the position variables and can be efficiently encoded. The time-varying position bounds can

also be used to verify that motion control strategies derived from simplified assumptions do not result in collisions in practice or in higher fidelity simulations. We investigate an example in the following Section.

V. CASE STUDY

We consider $m = 8$ mobile robots, $\mathcal{A} = \{\alpha_l\}_{l=1}^8$, operating on a plane. The location of each agent, $i \in \mathcal{A}$, at time t is represented by a position point, $p_i(t) = (x_i(t), y_i(t)) \in \mathbb{R}^{n_p}$ with $n_p = 2$. We assume there are $n = 6$ tasks, $\mathcal{T} = \{\tau_l\}_{l=1}^6$, corresponding to the actions of visiting target destinations, $\{g_j \in \mathbb{R}^{n_p} \mid j \in \mathcal{T}\}$, that are each to be completed by one agent given the initial positions $\{p_i(0) \in \mathbb{R}^{n_p} \mid i \in \mathcal{A}\}$. Fig. 2a illustrates the configuration at time $t = 0$. The tasks are allocated to the agents according to a sequential bottleneck assignment, as proposed in Assumption 2, with the distance function in Assumption 1 specified to be the Euclidean distance, $d(p, p') = \|p' - p\|_2$. The resulting optimising assignment, $\Pi^* \in \mathcal{S}_{\mathcal{A}, \mathcal{T}}(\mathcal{W})$, is described in Table I, with k -order bottleneck edges, defined in (8), and robustness margins, defined in (9), for all $k \in \{1, \dots, 6\}$. We note that for every order the set of maximum-margin bottleneck edges is a singleton, all robustness margins are strictly positive, and that Π^* is therefore a robust lexicographic assignment.

TABLE I: Sequential bottleneck optimising assignment of target destinations to agents with initial positions and colour coding shown in Fig. 2a.

Order k	Bottleneck edge (i_k^*, j_k^*)	Bottleneck weight $w_{i_k^*, j_k^*}$	Robustness margin μ_k	Bound limit A_k
1	(α_5, τ_1)	87.95m	10.78m	92.72m
2	(α_4, τ_5)	78.67m	9.99m	82.64m
3	(α_1, τ_3)	73.53m	9.02m	76.54m
4	(α_2, τ_2)	64.56m	27.82m	76.54m
5	(α_3, τ_6)	60.53m	21.30m	75.83m
6	(α_7, τ_4)	59.08m	23.38m	75.83m

The agents are of varying size but we assume that no agent occupies any space outside of a radius of 1.5m from its position point, i.e., at time t the body of agent, $i \in \mathcal{A}$, lies within a safety circle with diameter $s = 3\text{m}$ centred at $p_i(t)$. The smallest robustness margin of the assignment, defined in (12), is $\mu = 9.02\text{m}$. Because $s < \mu$, Assumption 3 is satisfied and the safe sets, $\mathcal{L}_i(t) \subset \mathbb{R}^2$, in (14) are non-empty for all agents $i \in \mathcal{A}$. Figs. 2b to 2e shows the time-varying areas that satisfy the position constraints for every agent given the time-dependent coordination parameter,

$$a(t) = v^{\text{ref}}t + \frac{1}{2}(\mu - s),$$

evolving with constant rate, $v^{\text{ref}} = 10\text{m/s}$. We observe that for any assigned agent, i_k^* , $k \in \{1, \dots, 6\}$, there is always at least a distance of 3m between any point in the safe set, $\mathcal{L}_{i_k^*}(t)$, and any point in the safe of any other agent, $\mathcal{L}_{i'}(t)$, $i' \in \mathcal{A} \setminus \{i_k^*\}$. The only safe sets that intersect are those of the unassigned agents. For the assigned agents, $\{\alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5, \alpha_7\}$, the local constraints consist of the

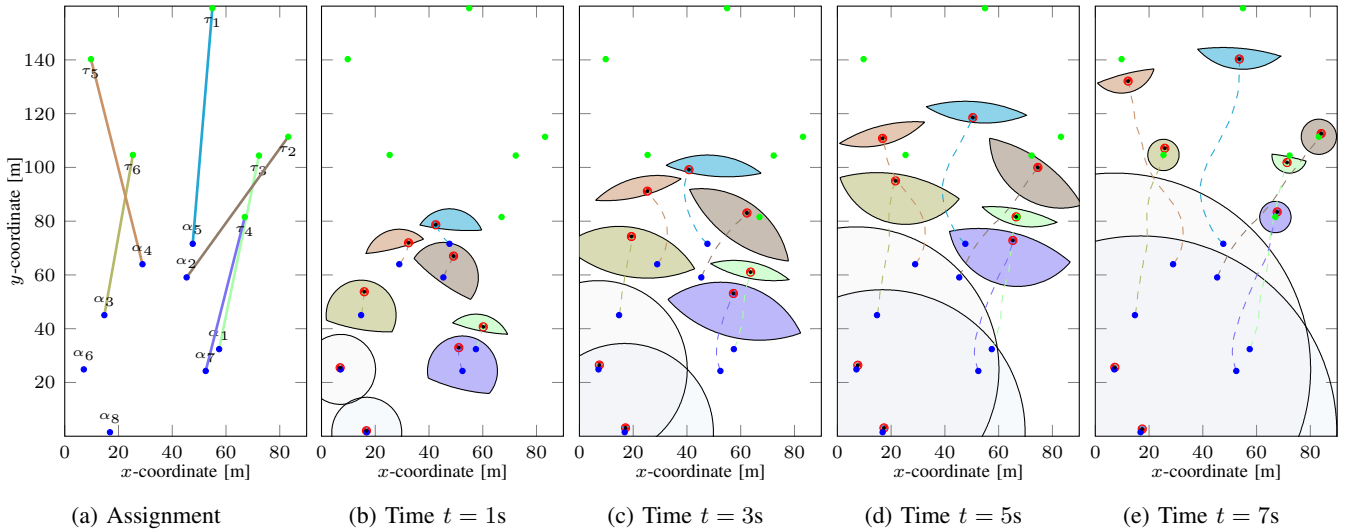


Fig. 2: Case study with agents $\mathcal{A} = \{\alpha_l\}_{l=1}^8$ and tasks $\mathcal{T} = \{\tau_l\}_{l=1}^6$. The sequential bottleneck optimising assignment is illustrated in (a). The positions satisfying the constraints, $\mathcal{L}_i(t)$, are shown in (b)-(e) as shaded areas for every agent, $i \in \mathcal{A}$, with initial positions, $p_i(0)$ [blue dots], positions at the time t , $p_i(t)$ [black dots], surrounded by safety circle with diameter s [red circles], past trajectories, $p_i(t')$, $t' \in [0, t]$ [dashed lines], and with destinations, g_j [green dots], for all tasks $j \in \mathcal{T}$.

intersection of both bounds introduced in Assumptions 4 and 5. The two bounds overlap by a different amount for every agent, i_k^* , $k \in \{1, \dots, 6\}$, depending on the difference of the corresponding saturation value, A_k defined in (13), and the associated weight, $w_{i_k^*, j_k^*}$. We observe from Table I that the saturation values of agents α_2 and α_7 , corresponding to bottleneck orders $k = 4$ and $k = 6$ respectively, are not given by their own associated robustness margins but rather by the value of a lower order robustness margin. That is, the sequence of saturation values, (A_1, \dots, A_6) , is set to be non-increasing in agreement with (13). The constraints for the unassigned agents, $\{\alpha_6, \alpha_8\}$, only consist of limits on the distances from their initial positions that are bound by $A_6 = 75.83\text{m}$. If the positions of all agents remain inside the individual safe sets at all considered times, it is guaranteed from Theorem 1 that none of the six assigned agents collide with any agent, including the two unassigned agents.

We consider a scenario where the motion of the robotic agents is governed by nonlinear unicycle models. The agents are guided from their initial positions to their assigned destinations with decentralised feedback controllers. The controllers are obtained from an optimal control approach in which the models are linearised around reference straight-line trajectories, with constant reference speed, $v(t) = v^{\text{ref}}$, and where input constraints are neglected. Figs. 2b to 2e show a sample simulation in which the controllers are applied to the nonlinear models with input constraints, steering rate disturbances, and where the initial heading angles are not aligned with the targets. The resulting agent position trajectories clearly deviate from the straight-line references. However, because the positions of all agents never leave the local safe sets, collisions involving the assigned agents are guaranteed not to occur. It is enough to tune or design the controllers for every agent individually such that the

deviations from the straight-line paths do not violate the local constraints. This demonstrates how the motion control of the agents can be decoupled while still guaranteeing collision avoidance. We note that the same analysis and planning can be applied to robots moving in three dimensions, $n_p = 3$, without additional complexity.¹

VI. CONCLUSIONS

We derived local time-varying position constraints for agents that are assigned to move to different target destinations. We specified conditions for which satisfaction of the constraints guarantees that collisions are avoided. The introduced method allocates targets to agents such that the largest distance between an agent and its destination is minimised. There may be agents that are not assigned to any task but all agents are constrained in order to maintain a minimal safety distance from any assigned agent. The parameters of the local constraints are obtained from a sequential bottleneck assignment procedure and depend only on the distances between agent initial positions and destinations. The local constraints define non-empty convex sets of safe positions for every agent if the optimising assignment is sufficiently robust to changes in the distances relative to the size of the agents. We defined a method to quantify the robustness and observed that the more robust the assignment is, the larger the regions of safe positions are.

The constraints can be constructed in polynomial time and with distributed computation where agents only access limited information on other agents. The procedure consists of solving multiple coupled BAPs and involves structure which may be exploited for faster computation in future work. The proposed constraints provide sufficient but not

¹A supplementary video available at <http://ieeexplore.ieee.org> contains animations of this case study and other scenarios in 2D and 3D.

necessary conditions for collision avoidance. In cases where not all conditions are satisfied, other strategies may exist to avoid collisions. Incorporating other strategies and investigating alternative local constraints derived from the properties of the assignment motivates further research on this topic.

APPENDIX I

PROOF OF PROPOSITION 1

Proof: From (6) we know that $w_{i^*,j^*} = B_{\bar{\mathcal{A}},\bar{\mathcal{T}}}(\bar{\mathcal{E}},\mathcal{W})$ for any $(i^*,j^*) \in e_{\bar{\mathcal{A}},\bar{\mathcal{T}}}(\mathcal{W})$. Given $r_{\bar{\mathcal{A}},\bar{\mathcal{T}}}(\mathcal{W}) > 0$, we assume for the sake of contradiction that there exist $\Pi \in \mathcal{B}_{\bar{\mathcal{A}},\bar{\mathcal{T}}}(\bar{\mathcal{E}},\mathcal{W})$ and $(i^*,j^*) \in e_{\bar{\mathcal{A}},\bar{\mathcal{T}}}(\mathcal{W})$ with $\pi_{i^*,j^*} = 0$. From (5) it follows that $b(\Pi,\bar{\mathcal{E}},\mathcal{W}) = B_{\bar{\mathcal{A}},\bar{\mathcal{T}}}(\bar{\mathcal{E}},\mathcal{W})$. Because $\pi_{i^*,j^*}w_{i^*,j^*} = 0$ in (4), we have $b(\Pi,\bar{\mathcal{E}},\mathcal{W}) = b(\Pi,\bar{\mathcal{E}} \setminus \{(i^*,j^*)\},\mathcal{W})$. From (5) it follows that $b(\Pi,\bar{\mathcal{E}} \setminus \{(i^*,j^*)\},\mathcal{W}) \geq B_{\bar{\mathcal{A}},\bar{\mathcal{T}}}(\bar{\mathcal{E}} \setminus \{(i^*,j^*)\},\mathcal{W})$ and thus $B_{\bar{\mathcal{A}},\bar{\mathcal{T}}}(\bar{\mathcal{E}} \setminus \{(i^*,j^*)\},\mathcal{W}) \leq B_{\bar{\mathcal{A}},\bar{\mathcal{T}}}(\bar{\mathcal{E}},\mathcal{W})$. We know that $B_{\bar{\mathcal{A}},\bar{\mathcal{T}}}(\bar{\mathcal{E}} \setminus \{(i^*,j^*)\},\mathcal{W}) \geq B_{\bar{\mathcal{A}},\bar{\mathcal{T}}}(\bar{\mathcal{E}},\mathcal{W})$ from (5) because $\{\Pi \in \mathcal{P}_{\bar{\mathcal{A}},\bar{\mathcal{T}}}(\bar{\mathcal{E}} \setminus \{(i^*,j^*)\}) \mid \pi_{i^*,j^*} = 0\} \subset \mathcal{P}_{\bar{\mathcal{A}},\bar{\mathcal{T}}}(\bar{\mathcal{E}})$. It follows that $B_{\bar{\mathcal{A}},\bar{\mathcal{T}}}(\bar{\mathcal{E}} \setminus \{(i^*,j^*)\},\mathcal{W}) = B_{\bar{\mathcal{A}},\bar{\mathcal{T}}}(\bar{\mathcal{E}},\mathcal{W})$ which contradicts $r_{\bar{\mathcal{A}},\bar{\mathcal{T}}}(\mathcal{W}) = B_{\bar{\mathcal{A}},\bar{\mathcal{T}}}(\bar{\mathcal{E}} \setminus \{(i^*,j^*)\},\mathcal{W}) - B_{\bar{\mathcal{A}},\bar{\mathcal{T}}}(\bar{\mathcal{E}},\mathcal{W}) > 0$. ■

APPENDIX II

PROOF OF PROPOSITION 2

Proof: For an arbitrary order, $a \in \{1, \dots, n-1\}$, and an arbitrary higher order, $b \in \{a+1, \dots, n\}$, we consider the assignment, $\hat{\Pi} = \{\hat{\pi}_{i,j} \in \{0,1\} \mid (i,j) \in \mathcal{A} \times \mathcal{T}\}$, obtained from $\Pi^* = \{\pi_{i,j}^* \in \{0,1\} \mid (i,j) \in \mathcal{A} \times \mathcal{T}\}$ by switching the agent-task pairings of (i_a^*,j_a^*) and (i_b^*,j_b^*) to (i_a^*,j_b^*) and (i_b^*,j_a^*) , i.e.,

$$\hat{\pi}_{i,j} = \begin{cases} 0 & \text{if } (i,j) = (i_a^*,j_a^*) \text{ or } (i,j) = (i_b^*,j_b^*), \\ 1 & \text{if } (i,j) = (i_a^*,j_b^*) \text{ or } (i,j) = (i_b^*,j_a^*), \\ \pi_{i,j}^* & \text{otherwise.} \end{cases}$$

We note that $(i_a^*,j_b^*), (i_b^*,j_a^*) \in \bar{\mathcal{E}}_a$ by definition in (7). Because (i_a^*,j_b^*) and (i_b^*,j_a^*) are not assigned according to Π^* , it follows from Definition 4 that $\max\{w_{i_a^*,j_b^*}, w_{i_b^*,j_a^*}\} \geq b(\Pi^*,\bar{\mathcal{E}}_a,\mathcal{W})$. By construction of $\hat{\Pi}$, we know that $b(\hat{\Pi},\hat{\mathcal{E}},\mathcal{W}) = b(\Pi^*,\hat{\mathcal{E}},\mathcal{W}) \leq b(\Pi^*,\bar{\mathcal{E}}_a,\mathcal{W})$, with $\hat{\mathcal{E}} := \bar{\mathcal{E}}_a \setminus \{(i_a^*,j_a^*), (i_b^*,j_b^*), (i_a^*,j_b^*), (i_b^*,j_a^*)\}$. It follows that $b(\hat{\Pi},\bar{\mathcal{E}}_a,\mathcal{W}) = \max\{w_{i_a^*,j_b^*}, w_{i_b^*,j_a^*}\}$. Since (i_a^*,j_a^*) is not assigned according to $\hat{\Pi}$, we know $b(\hat{\Pi},\bar{\mathcal{E}}_a \setminus \{(i_a^*,j_a^*)\},\mathcal{W}) = \max\{w_{i_a^*,j_b^*}, w_{i_b^*,j_a^*}\}$. Finally, from Definitions 2 and 3 we have that $w_{i_a^*,j_a^*} + \mu_a = B_{\bar{\mathcal{A}},\bar{\mathcal{T}}_a}(\bar{\mathcal{E}}_a \setminus \{(i_a^*,j_a^*)\},\mathcal{W}) \leq b(\hat{\Pi},\bar{\mathcal{E}}_a \setminus \{(i_a^*,j_a^*)\},\mathcal{W})$. ■

APPENDIX III

PROOF OF PROPOSITION 3

Proof: For an arbitrary order, $a \in \{1, \dots, n\}$, and an arbitrary unassigned agent, $i' \in \mathcal{A} \setminus \{i_1^*, \dots, i_n^*\}$, consider the assignment, $\Pi' = \{\pi'_{i,j} \in \{0,1\} \mid (i,j) \in \mathcal{A} \times \mathcal{T}\}$, obtained from $\Pi^* = \{\pi_{i,j}^* \in \{0,1\} \mid (i,j) \in \mathcal{A} \times \mathcal{T}\}$ by switching the agent assigned to task j_a^* from i_a^* to i' , i.e.,

$$\pi'_{i,j} = \begin{cases} 0 & \text{if } (i,j) = (i_a^*,j_a^*), \\ 1 & \text{if } (i,j) = (i',j_a^*), \\ \pi_{i,j}^* & \text{otherwise.} \end{cases}$$

We note that $(i',j_a^*) \in \bar{\mathcal{E}}_a$ by definition in (7). Because (i',j_a^*) is not assigned according to Π^* , it follows from Definition 4 that $w_{i',j_a^*} \geq b(\Pi^*,\bar{\mathcal{E}}_a,\mathcal{W})$. By construction of Π' , we know that $b(\Pi',\mathcal{E}',\mathcal{W}) = b(\Pi^*,\mathcal{E}',\mathcal{W}) \leq b(\Pi^*,\bar{\mathcal{E}}_a,\mathcal{W})$, with $\mathcal{E}' := \bar{\mathcal{E}}_a \setminus \{(i',j_a^*), (i_a^*,j_a^*)\}$. It follows that $b(\Pi',\bar{\mathcal{E}}_a,\mathcal{W}) = w_{i',j_a^*}$. Since (i_a^*,j_a^*) is not assigned according to Π' , we know that $b(\Pi',\bar{\mathcal{E}}_a \setminus \{(i_a^*,j_a^*)\},\mathcal{W}) = w_{i',j_a^*}$. Finally, from Definitions 2 and 3 we have that $w_{i_a^*,j_a^*} + \mu_a = B(\bar{\mathcal{E}}_a \setminus \{(i_a^*,j_a^*)\},\mathcal{W}) \leq b(\Pi',\bar{\mathcal{E}}_a \setminus \{(i_a^*,j_a^*)\},\mathcal{W})$. ■

REFERENCES

- [1] S. Chopra, G. Notarstefano, M. Rice, and M. Egerstedt, "A distributed version of the hungarian method for multirobot assignment," *IEEE Transactions on Robotics*, vol. 33, no. 4, pp. 932–947, 2017.
- [2] H. L. Choi, L. Brunet, and J. P. How, "Consensus-based decentralized auctions for robust task allocation," *IEEE Transactions on Robotics*, vol. 25, no. 4, pp. 912–926, 2009.
- [3] D. W. Pentico, "Assignment problems: A golden anniversary survey," *European Journal of Operational Research*, vol. 176, no. 2, pp. 774–793, 2007.
- [4] B. P. Gerkey and M. J. Mataric, "A formal analysis and taxonomy of task allocation in multi-robot systems," *International Journal of Robotics Research*, vol. 23, no. 9, pp. 939–954, 2004.
- [5] R. E. Burkard, M. Dell'Amico, and S. Martello, *Assignment problems, revised reprint*. Siam, 2012.
- [6] I. Shames, A. Dostovalova, J. Kim, and H. Hmam, "Task allocation and motion control for threat-seduction decoys," in *Conference on Decision and Control*, 2017, pp. 4509–4514.
- [7] R. E. Burkard and F. Rendl, "Lexicographic bottleneck problems," *Operations Research Letters*, vol. 10, no. 5, pp. 303–308, 1991.
- [8] M. Khoo, T. A. Wood, C. Manzie, and I. Shames, "Distributed algorithm for solving the bottleneck assignment problem," *Conference on Decision and Control*, pp. 1850–1855, 2019.
- [9] E. Michael, T. A. Wood, C. Manzie, and I. Shames, "Uncertainty intervals for robust bottleneck assignment," in *European Control Conference*. IEEE, 2019, pp. 4204–4209.
- [10] J. van den Berg, S. J. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Robotics research*. Springer, 2011, pp. 3–19.
- [11] J. Alonso-Mora, P. Beardsley, and R. Siegwart, "Cooperative collision avoidance for nonholonomic robots," *IEEE Transactions on Robotics*, vol. 34, no. 2, pp. 404–420, 2018.
- [12] L. Wang, A. D. Ames, and M. Egerstedt, "Safety barrier certificates for collisions-free multirobot systems," *IEEE Transactions on Robotics*, vol. 33, no. 3, pp. 661–674, 2017.
- [13] D. Zhou, Z. Wang, S. Bandyopadhyay, and M. Schwager, "Fast, on-line collision avoidance for dynamic vehicles using buffered voronoi cells," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 1047–1054, 2017.
- [14] V. R. Desaraju and J. P. How, "Decentralized path planning for multi-agent teams with complex constraints," *Autonomous Robots*, vol. 32, no. 4, pp. 385–403, 2012.
- [15] M. Turpin, K. Mohta, N. Michael, and V. Kumar, "Goal assignment and trajectory planning for large teams of interchangeable robots," *Autonomous Robots*, vol. 37, no. 4, pp. 401–415, 2014.
- [16] F. Wu, V. S. Varadharajan, and G. Beltrame, "Collision-aware task assignment for multi-robot systems," in *International Symposium on Multi-Robot and Multi-Agent Systems*, 2019, pp. 30–36.
- [17] I. Shames, B. Fidan, and B. D. Anderson, "Close target reconnaissance with guaranteed collision avoidance," *International Journal of Robust and Nonlinear Control*, vol. 21, no. 16, pp. 1823–1840, 2011.
- [18] P. MacAlpine, E. Price, and P. Stone, "Scram: Scalable collision-avoiding role assignment with minimal-makespan for formational positioning," in *AAAI Conference on Artificial Intelligence*, 2015, pp. 2096–2102.
- [19] M. Turpin, N. Michael, and V. Kumar, "Capt: Concurrent assignment and planning of trajectories for multiple robots," *International Journal of Robotics Research*, vol. 33, no. 1, pp. 98–112, 2014.
- [20] A. Bemporad and C. Rocchi, "Decentralized linear time-varying model predictive control of a formation of unmanned aerial vehicles," *Conference on Decision and Control and European Control Conference*, pp. 7488–7493, 2011.