

Near-Optimal Min-Sum Motion Planning for Two Square Robots in a Polygonal Environment*

Pankaj K. Agarwal[†] Dan Halperin[‡] Micha Sharir[§] Alex Steiger[¶]

Abstract

Let $\mathcal{W} \subset \mathbb{R}^2$ be a planar polygonal environment (i.e., a polygon potentially with holes) with a total of n vertices, and let A, B be two robots, each modeled as an axis-aligned unit square, that can translate inside $\mathcal{W}$. Given source and target placements $s_A, t_A, s_B, t_B \in \mathcal{W}$ of A and B , respectively, the goal is to compute a *collision-free motion plan* π^* , i.e., a motion plan that continuously moves A from s_A to t_A and B from s_B to t_B so that A and B remain inside $\mathcal{W}$ and do not collide with each other during the motion. Furthermore, if such a plan exists, then we wish to return a plan that minimizes the sum of the lengths of the paths traversed by the robots. Given $\mathcal{W}, s_A, t_A, s_B, t_B$ and a parameter $\varepsilon > 0$, we present an $n^2 \varepsilon^{-O(1)} \log n$ -time $(1 + \varepsilon)$ -approximation algorithm for this problem. We are not aware of any polynomial-time algorithm for this problem, nor do we know whether the problem is NP-Hard. Our result is the first polynomial-time $(1 + \varepsilon)$ -approximation algorithm for an optimal motion-planning problem involving two robots moving in a polygonal environment.

1 Introduction

The basic motion-planning problem is to decide whether a robot (i.e., a rigid or multi-link moving object) can move from a given start position to a given target position without colliding with obstacles on its way, and avoiding collision of different parts of the robot. If the answer is positive, we also want to plan such a motion. With the advancement of robotics, we witness the growing deployment of *teams* of robots in logistics, wildlife monitoring, buildings and bridges inspection and more. Motion planning for many robots requires that, in addition to not colliding with obstacles, the robots should not collide with one another, which in turn necessitates studying the problem in high-dimensional *configuration spaces*. Furthermore, we wish to ensure a good quality of the motion, such as being short or having a small makespan. Already for two simple robots, such as unit squares or discs, translating in a planar polygonal environment, little is known when it comes to optimizing the robot motion. Although polynomial-time algorithms are known for computing a collision-free motion plan of two simple robots [35], no polynomial-time algorithm is known for computing a plan such that the sum (or the maximum) of the path lengths of the two robots is minimized, nor is the problem known to be NP-hard. Even a polynomial-time constant-factor approximation algorithm is not known for this problem (without further restrictions).

Problem statement. Let $\square = \{x \in \mathbb{R}^2 \mid \|x\|_\infty \leq 1\}$ denote the unit-radius axis-aligned square centered at the origin, referred to as a *unit square* for short. For a point $p \in \mathbb{R}^2$ and a real value $\lambda \geq 0$, we use $p + \lambda\square$ to denote the axis-parallel square of radius λ centered at p . Let A and B be two robots, each modeled as a unit square, that can translate inside the same closed planar polygonal environment (a connected polygon possibly with holes) $\mathcal{W}$ with n vertices. A placement of A or B is represented by a point in $\mathcal{W}$ — the position of its center. For such a placement to be free of collision with $\partial\mathcal{W}$, the boundary of $\mathcal{W}$, the representing point should be at L_∞ -distance at least 1 from $\partial\mathcal{W}$. We denote by $\mathcal{F}$, the *free space* of a single robot, the subset of $\mathcal{W}$ consisting of such points. Note that the robots may be at L_∞ -distance 1 from $\partial\mathcal{W}$ and hence they are allowed to make contact with the

*The full version of the paper can be accessed at <https://arxiv.org/abs/2310.20615>.

Work by P.A. and A.S. has been partially supported by IIS-1814493, CCF-2007556, and CCF-2223870. Work by D.H. has been supported in part by the Israel Science Foundation (grant nos. 1736/19 and 2261/23), by NSF/US-Israel-BSF (grant no. 2019754), by the Israel Ministry of Science and Technology (grant no. 103129), by the Blavatnik Computer Science Research Fund, and by the Yandex Machine Learning Initiative for Machine Learning at Tel Aviv University. Work by M.S. has been supported by Israel Science Foundation Grants 260/18 and 495/23.

[†]Department of Computer Science, Duke University, Durham, NC 27708-0129, USA; pankaj@cs.duke.edu.

[‡]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel; danha@tauex.tau.ac.il.

[§]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel; michas@tauex.tau.ac.il.

[¶]Department of Computer Science, Duke University, Durham, NC 27708-0129, USA; asteiger@cs.duke.edu.

obstacles. A (joint) *configuration* of A and B is represented as a pair $(p_A, p_B) \in \mathcal{W} \times \mathcal{W}$, where p_A (resp., p_B) is the placement of A (resp., B). We also represent a configuration as a point $p \in \mathbb{R}^4$, where the first (resp., last) pair of coordinates represent the placement of A (resp., B). The *configuration space*, called *C-space* for short, namely the set of all configurations, is thus represented as $\mathcal{W} \times \mathcal{W} \subset \mathbb{R}^4$. A configuration $\mathbf{p} = (p_A, p_B) \in \mathbb{R}^4$ is called *free* if $p_A, p_B \in \mathcal{F}$, that is, $p_A + \square, p_B + \square \subseteq \mathcal{W}$, and $\|p_A - p_B\|_\infty \geq 2$. Such a free configuration is called a *kissing configuration* if $\|p_A - p_B\|_\infty = 2$, i.e., the robots touch each other (but their interiors remain disjoint). Let $\mathbf{F} := \mathbf{F}(\mathcal{W})$ denote the (four-dimensional) *free space*, namely the set of all free configurations. Clearly, $\mathbf{F} \subset \mathcal{F} \times \mathcal{F}$.

Two free configurations $\mathbf{s}, \mathbf{t} \in \mathbf{F}$ are *reachable* if they lie in the same connected component of $\mathbf{F}$, i.e., there is a path contained in $\mathbf{F}$ from $\mathbf{s}$ to $\mathbf{t}$. For two reachable free configurations $\mathbf{s} := (s_A, s_B), \mathbf{t} := (t_A, t_B) \in \mathbf{F}$, a path $\pi \subseteq \mathbf{F}$ from $\mathbf{s}$ to $\mathbf{t}$ is called a (*feasible*) *plan* of A and B from $\mathbf{s}$ to $\mathbf{t}$, or an $(\mathbf{s}, \mathbf{t})$ -*plan* for brevity. With a slight abuse of notation, we also use π as a (continuous) parameterization $\pi : [0, 1] \rightarrow \mathbf{F}$, with $\pi(0) = \mathbf{s}$ and $\pi(1) = \mathbf{t}$. For a path $\pi \subseteq \mathbf{F}$, let π_A (resp., π_B) be the projection of π onto the two-dimensional plane spanned by the first (resp., last) two coordinates, which specifies the path followed by A (resp., B) that π induces; we have $\pi_A, \pi_B \subset \mathcal{F}$. Let $\mathfrak{c}(\pi_A), \mathfrak{c}(\pi_B)$ denote the (Euclidean) arc length of the paths π_A, π_B , respectively, in $\mathbb{R}^2$. We define $\mathfrak{c}(\pi)$, the *cost* of π , to be the sum of the lengths of π_A and π_B , i.e., $\mathfrak{c}(\pi) = \mathfrak{c}(\pi_A) + \mathfrak{c}(\pi_B)$. Let $\pi^*(\mathbf{s}, \mathbf{t})$ denote an *optimal* $(\mathbf{s}, \mathbf{t})$ -plan, i.e., a plan that minimizes the sum of the lengths of the two paths.¹ If $\mathbf{s}$ and $\mathbf{t}$ are not reachable, i.e., they lie in different connected components of $\mathbf{F}$, then $\pi^*(\mathbf{s}, \mathbf{t})$ does not exist. We refer to the problem of computing $\pi^*(\mathbf{s}, \mathbf{t})$ as the (*optimal*) *min-sum motion-planning problem*. In this paper we study the min-sum motion-planning problem for two translating axis-aligned unit squares, and present a $(1 + \varepsilon)$ -approximation algorithm that runs in $n^2 \varepsilon^{-O(1)} \log n$ time.

Related work. Algorithmic motion planning has been studied for well over fifty years in computer science and beyond. The rigorous study of algorithmic motion planning dates back to the work of Schwartz and Sharir [33] and Canny [10]. See [17, 18, 27, 29] for a review of key relevant results. We mention here only a small sample of these results—the ones that are most closely related to the problem at hand.

When only one square robot translates, or more generally when only one convex polygonal robot of a constant description complexity (that is, with a constant number of vertices) translates, the problem is equivalent—through C-space formulation—to moving a point robot amid polygonal obstacles with $O(n)$ vertices, and it can be solved in $O(n \log n)$ time [11, 20, 43]. Interestingly, the analogous problem in 3D, namely finding the shortest path for a point robot amid polyhedral obstacles, is NP-hard [9] and fast $(1 + \varepsilon)$ -approximation algorithms are known [9, 34]. Note that this hard problem has only three degrees of freedom of motion, and there are other optimal motion-planning problems for robots with three degrees of freedom that are NP-hard [5, 6]. Our two-square problem has four degrees of freedom, which suggests it might be NP-hard as well, though, as we have remarked earlier, this is an open problem.

Computing a feasible (not necessarily optimal) plan for a team of translating unit square robots in a polygonal environment is PSPACE-hard [38] (see also [7, 8, 19, 21, 40, 46] for related intractability results). Notwithstanding a rich literature on multi-robot motion planning in both continuous and discrete setting (robots moving on a graph in the latter setting), see, e.g., [12, 23, 24, 32, 36, 41, 42], little is known about algorithms producing paths with provable quality guarantees. Approximation algorithms for minimizing the total path-length are given in [2, 37, 39] for a set of unit-disc robots assuming a certain separation between the start and goal positions, as well as from the obstacles. The separation assumption makes the problem considerably easier. A feasible plan always exists, and one can first compute an optimal path for each robot independently, ignoring other robots and then locally modify them so that the robots do not collide with each other during their motion. An $O(1)$ -approximation algorithm was proposed in [14] for computing a plan that minimizes the makespan for a set of unit discs (or squares) in the plane without obstacles, again assuming some separation. Computing the min-sum motion plan for two unit squares/discs even in the absence of obstacles is non-trivial [15, 25]. We are unaware of any constant-factor approximation algorithms for the min-sum motion-planning problem even for two unit squares/discs in a planar polygonal environment without any assumptions on the work environment or on the start/final configurations.

Quite a few of the algorithmic results for teams of robots distinguish between the labeled and unlabeled versions: In the labeled version, like in the two-square problem studied here, each robot is designated its own unique target position. In the unlabeled case, each robot can finish at any of the (collective) target positions, as long as at the end of the motion all the target position are occupied by robots. For a team of unlabeled unit discs,

¹The existence of π^* can be proved using a simple compactness argument, since $\mathcal{F}$ and $\mathbf{F}$ are closed.

an approximate solution for the minimum total path length is given in [39], assuming a certain separation between the start and goal positions of the robots, as well as from the obstacles. A similar result has also been obtained for a team of labeled unit discs in [37], using the slightly more relaxed requirement of the existence of *revolving areas* around the start and target positions. In both cases the approximation bounds are crude, and we omit them here. The latter result for labeled unit discs has recently been improved [2], to give an $O(1)$ -approximation of the optimal total length of the paths, under exactly the same conditions as in [37].

The central and prevalent family of practical motion-planning techniques in robotics is based on sampling of the underlying C-space; see, e.g., [22, 24, 27, 28, 30], and [32] for a recent review. Another major line of work on optimizing multi-robot motion plans addresses a discrete version of the problem, where robots are moving on graphs. In this setting the robots are often referred to as *agents*, and the problem is called *Multi Agent Path Finding (MAPF)*. There is a rich literature on MAPF, and we refer the reader to the recent survey [41]. A commonly used optimization criterion (particularly in the study of MAPF, but elsewhere as well) is *makespan*, where we wish to minimize the time by which all the robots reach their destination, assuming they move in some prespecified maximum speed; see, e.g., [14, 46].

There are a variety of additional optimization criteria in robot motion planning. A common one, related to motion safety, is requiring high *clearance*, namely, requiring that the robot stays far from the obstacles in its environment—this can be obtained using Voronoi diagrams (e.g., [30]). In the context of multi-robot planning we may also require that the robots stay sufficiently far from one another (e.g., [12]). A natural requirement is to produce paths that are at once short and far away from obstacles, which is a more intricate task even for a single robot translating in the plane; see, e.g., [1, 44, 45].

Our contributions. We consider the following simple case of min-sum motion-planning for two unit-square robots. Let $\mathcal{W}$ be a polygonal environment, i.e., a polygon possibly with holes. As already stated, we assume that the two robots A and B are axis-parallel squares of side-length 2. Given a source and a target free configurations $(s_A, s_B), (t_A, t_B) \in \mathcal{W}$, the goal is to compute a collision-free motion plan for A from s_A to t_A and B from s_B to t_B , such that the sum of the lengths of the two tours traversed by the robots is minimized, or otherwise report that there is no such collision-free motion plan. Our main result is the following theorem, which provides an efficient ε -approximation algorithm for this problem.²

THEOREM 1.1. *Let $\mathcal{W}$ be a closed polygonal environment with n vertices, let A, B be two axis-parallel unit-square robots translating inside $\mathcal{W}$, and let $\mathbf{s}, \mathbf{t}$ be source and target configurations of A, B . For any $\varepsilon \in (0, 1)$, a motion plan π from $\mathbf{s}$ to $\mathbf{t}$ with $\mathfrak{c}(\pi) \leq (1 + \varepsilon)\mathfrak{c}(\pi^*)$, if there exists a such a motion, can be computed in $n^2\varepsilon^{-O(1)}\log n$ time, where π^* is an optimal $(\mathbf{s}, \mathbf{t})$ -plan.*

Although our result falls short of answering whether the min-sum problem for two robots is in P, it is a significant contribution to the theory of optimal multi-robot motion planning. First, as mentioned above, a polynomial-time algorithm was not known, even for constant-factor approximation, and we present an FPTAS for this problem. Second, we prove several structural properties of an optimal plan, which could lead to a polynomial-time algorithm in some special cases, e.g., when $\mathcal{W}$ is rectilinear and we consider the L_1 -length of a path. Note that our FPTAS does not rule out the possibility of the problem being NP-hard because, as in other NP-hard optimal motion-planning problems, the construction might use a polynomial number of bits (see, e.g., [9]). Finally, our algorithm is very simple and follows the widely-used sampling paradigm. More precisely, we sample a finite set $\mathcal{V} \subset \mathbf{F}$ of free configurations that contains $\mathbf{s}, \mathbf{t}$. We connect a pair of configurations $\mathbf{p} := (p_A, p_B)$ and $\mathbf{q} := (q_A, q_B)$ in $\mathcal{V}$ by an edge if there is a *simple* (feasible) plan from $\mathbf{p}$ to $\mathbf{q}$, namely, we can move A from p_A to q_A (not necessarily along a straight segment) while keeping B parked at p_B and then move B from p_B to q_B while A is parked at q_A , or vice-versa. The cost of the edge $(\mathbf{p}, \mathbf{q})$ is the minimum cost of such a plan. We then compute a shortest path in this graph. The question is, of course, how we (efficiently) choose a small number of free configurations (linear in n) so that the resulting graph is guaranteed to contain a path from $\mathbf{s}$ to $\mathbf{t}$ that corresponds to a near-optimal $(\mathbf{s}, \mathbf{t})$ -plan. Most of this paper is about answering this question. We note that the runtime of our algorithm nearly matches that of the best known algorithm for finding *any* $(\mathbf{s}, \mathbf{t})$ -plan for two unit squares in a planar polygonal environment, which takes $O(n^2)$ time [35].

²In principle, our approach extends to two identical centrally-symmetric regular convex polygons, but the analysis becomes even more technical, so for simplicity we only focus on unit squares.

There are four main technical contributions of this paper. First, we prove a few key properties of an optimal plan (Section 3). Concretely, we show that there is always an optimal plan in which only one robot moves at any given time while the other robot is *parked* (remains stationary). Thus an optimal plan can be represented as a sequence of *moves*, where each move is specified as a 3-tuple (R, π, p) , where $R \in \{A, B\}$ is the robot that is moving along a path $\pi \subseteq \mathcal{W}$ and the other robot is parked at $p \in \mathcal{F}$, where $\pi \times \{p\}$ or $\{p\} \times \pi$ is in $\mathbf{F}$ (π also encodes the starting and terminating placements of R in this move). We refer to such a plan as a *decoupled plan*.³

Second, we show that among all decoupled plans, there exists one in which for each move (R, π, p) , except possibly the first and the last moves, there is a point $q \in \pi$ such that (p, q) (or (q, p) as the case might be) is a kissing configuration. We refer to such a plan as a *kissing plan*. We use the kissing property to prove that there exists an optimal, kissing plan π^* composed of $O(\mathfrak{c}(\pi^*) + 1)$ moves. Our usage of kissing configurations is different from earlier work (see, e.g., [4, 16, 22]) in a few ways. First, the focus of these works is on motion in contact. For example, Aronov *et al.* [4] use a continuum of kissing configurations to reduce the dimension of the underlying joint configuration space of a pair or of a triple of robots, under various extra conditions. In contrast, kissing configurations in this paper arise as part of individual robot moves, often a singular/discrete configuration, in a (possibly long) alternating sequence of moves. Second, earlier work deals with feasible motion, while we show that there exists optimal plans in which almost every move contains a kissing configuration.

Finally, we prove that there is always a kissing plan in which neither of the robots is ever parked deep inside *corridors*. A formal definition of corridors is given in Section 2, but intuitively a corridor is a (narrow) region of $\mathcal{F}$ bounded by two of its edges that is far from all vertices of $\mathcal{F}$ and not wide enough to let one robot pass the other.

Next, using these three properties of an optimal plan, we show that we can deform an optimal kissing plan to a *tame plan*, at a slight increase of its cost, in which (roughly speaking) a robot is always parked near a vertex of $\mathcal{W}$ or of a corridor at each move. Furthermore, the deformed plan $\tilde{\pi}$ is composed of $O(\mathfrak{c}(\pi) + 1)$ moves and remains a kissing plan (Section 5). Ensuring the kissing property in this deformation is delicate and requires a rather involved argument, so we first prove the existence of a tame plan without ensuring the kissing property (Section 5). This weaker property already leads to an $n^3 \varepsilon^{-O(1)} \log n$ -time $(1 + \varepsilon)$ -approximation algorithm. A key ingredient in computing these deformations is the notion of *revolving areas* within $\mathcal{F}$, the two-dimensional free space with respect to one robot, roughly a unit square inside $\mathcal{F}$ (again see below for a precise definition). We can show that if each of s_A, s_B, t_A, t_B lies in a revolving area, then there is an (s, t) -plan π composed of $O(1)$ moves with cost $\mathfrak{c}(\pi) \leq \varrho(s_A, s_B) + \varrho(t_A, t_B) + O(1)$, where $\varrho(\cdot, \cdot)$ is the geodesic distance between two points in $\mathcal{F}$. The notion of revolving areas was used in [2, 37] to make a strong separation assumption on each of the start and target configurations, which was exploited to compute a near-optimal plan. Here, we prove the existence of revolving areas in the neighborhood of a non-tame plan and use them for auxiliary parking spots to convert the plan into a near-optimal tame plan.

The existence of an kissing, tame, near-optimal (s, t) -plan π^* enables us to choose a set $\mathcal{V}$ of $n\varepsilon^{-O(1)}$ (nearly) kissing configurations and to build a graph $\mathcal{G}$ over them so that π^* can be retracted to a path in $\mathcal{G}$ at a slight increase in its cost, thereby reducing the problem to computing a shortest path in $\mathcal{G}$. Ensuring that the two robots do not collide with each other in the retracted path requires care and thus the retraction map is somewhat involved. This retraction step introduces $O(\varepsilon)$ additive error, so we need a separate procedure to handle the case when $\mathfrak{c}(\pi^*)$ is small, say, at most $1/4$. By exploiting the topology of $\mathbf{F}$, we describe an $O(n \log^2 n)$ -time $O(1)$ -approximation algorithm for computing an optimal (s, t) -plan when $\mathfrak{c}(\pi^*) \leq 1/4$ (Section 8). We then plug it into the above algorithm to obtain a $(1 + \varepsilon)$ -approximation algorithm for all values of $\mathfrak{c}(\pi^*)$.

2 Preliminaries

Definitions. Let $\mathcal{F}$ be the free space of one robot as defined above. Throughout the paper, we regard s_A, s_B, t_A, t_B as additional vertices of $\mathcal{F}$. For a point $p \in \mathcal{W}$, let $\mathcal{F}[p] := \{x \in \mathcal{F} \mid \|x - p\|_\infty \geq 2\}$ be the set of all placements $x \in \mathcal{F}$ of A such that A does not collide with B if B is placed at p , i.e., $\text{int}(x + \square) \cap \text{int}(p + \square) = \emptyset$. It is well known that $\mathcal{F}$ and $\mathcal{F}[p]$ are polygonal and have $O(n)$ vertices, and that they can be computed in $O(n \log^2 n)$ time [13]. See Figure 1. For $p, q \in \mathcal{F}$, let $\varrho(p, q)$ denote the geodesic distance between p and q in $\mathcal{F}$. We call a configuration $(a, b) \in \mathbf{F}$ *x-separated* if $|x(a) - x(b)| \geq 2$ and *y-separated* if $|y(a) - y(b)| \geq 2$. (a, b) is always *x-separated* or *y-separated* (or both) since $\|a - b\|_\infty \geq 2$.

³We note that the notion of *decoupled* has been used in multiple ways in the context of multi-robot motion planning [26].

Given source and target configurations $\mathbf{s}, \mathbf{t} \in \mathbf{F}$, we call an $(\mathbf{s}, \mathbf{t})$ -plan $\pi : [0, 1] \rightarrow \mathbf{F}$ *decoupled* if only one robot moves at any time while the other robot is *parked* at some point in $\mathcal{F}$, and if there is only a finite number of switches between the moving and parking robots. A decoupled plan can be represented as a finite sequence

$$(R_1, \pi_1, p_1), (R_2, \pi_2, p_2), \dots, (R_k, \pi_k, p_k).$$

For a path $\pi \in \mathcal{F}$ and two values $\lambda, \lambda' \in [0, 1]$, $\lambda < \lambda'$, we denote by $\pi(\lambda, \lambda')$ the *pathlet* of π between times λ and λ' , which itself is a path (with a suitable reparameterization). It will be convenient to specify the portion of a path π between two points $p, p' \in \pi$ using the notation $\pi[p, p']$. We define the distance between closest points in a pair of sets using either the L_2 -distance or the L_∞ -distance. For any pair of subsets $X, Y \subset \mathbb{R}^2$, set

$$d_\ell(\mathbf{X}, \mathbf{Y}) := \min_{x \in \mathbf{X}, y \in \mathbf{Y}} \|x - y\|_\ell, \quad \text{for } \ell \in \{2, \infty\}.$$

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited




Figure 4. Two examples of corridors K (shaded) with blockers e_i, e_j that contains squares $S \subset \mathcal{F}$ with radii strictly less than 2 since their centers lie in the interiors of the corridors. The left (resp., right) corridor has direction vector u_{ij} with angle $\pi/4$ (resp., 0). Both examples are maximal since the portals σ_L have L_∞ -length 2 and σ_R contain a vertex g of $\mathcal{F}$.

Optimal plan for $\mathcal{W} = \mathbb{R}^2$. Suppose the work environment is the entire plane $\mathbb{R}^2$, i.e., there are no obstacles. In this case, Esteban *et al.* [15] proved that an optimal plan is a piecewise-linear decoupled plan consisting of at most three moves, and each move consists of at most three line segments. See Figure 2 for an example. Note that the parking position in some cases (such as the one in Figure 2) is not necessarily near the initial/final placements, which is one of the challenges in developing an efficient algorithm for computing an optimal plan.

Revolving area. A *revolving area* is a unit(-radius) square $p + \square$, for some $p \in \mathcal{F}$, that is contained in $\mathcal{F}$; we denote it by $\text{RA}(p)$ (Figure 3). For $p_A, p_B \in \partial \text{RA}(p)$ with $\|p_A - p_B\|_\infty = 2$, (p_A, p_B) is a *kissing configuration*, and we say that this kissing configuration lies in the revolving area $\text{RA}(p)$. In Section 4 we give useful lemmas regarding revolving areas, which play a key role in deforming an optimal path into a near-optimal *tame* plan (defined later in Section 5) that is easier to compute.




Figure 3. Example of a kissing configuration $(q, r) \in \mathbf{F}$ with $q, r \in \partial \text{RA}(p)$.

Corridor and sanctum. Intuitively, a corridor K is a (narrow) trapezoid in $\mathcal{F}$ bounded by two edges of $\mathcal{F}$, so that if one robot is parked inside K , the other one cannot pass around it (within K). This implies that when both robots are in the same corridor, their motions are constrained in ways that we will later explore. We now give a formal definition. Let e_i, e_j be a pair of edges of $\mathcal{F}$ that support an axis-aligned square (of any size) contained in $\mathcal{F}$, i.e., there exists an axis-aligned square $S \subset \mathcal{F}$ such that e_i (resp., e_j) touches a vertex of S , say v_i (resp., v_j), but does not intersect $\text{int}(S)$. Let $u_{ij} \in [0, \pi)$ be a direction normal to the segment $v_i v_j$; $u_{ij} = k\pi/4$ for some $0 \leq k \leq 3$.⁴ A *corridor* K bounded by e_i, e_j is a trapezoid such that (i) two of the edges of K are portions of e_i, e_j , called *blockers*; (ii) the other edges of K , called *portals*, are normal to the direction u_{ij} ; (iii) the L_∞ -length of each portal (i.e., the L_∞ -distance between its endpoints) is at most 2; and (iv) no vertex of $\mathcal{F}$ (including s_A, t_A, s_B, t_B) lies in the interior of K . See Figure 4. We refer to u_{ij} as the direction of the corridor. The following lemma directly follows from condition (iii).

LEMMA 2.1. *Let K be a corridor with direction vector u . For any segment $vw \subset \text{int}(K)$ normal to u , $\|v - w\|_\infty < 2$. Furthermore, $\|v - w\|_2 < 2$ if u is axis-parallel, otherwise $\|v - w\|_2 < 2\sqrt{2}$.*

A corridor K is *maximal* if there is no other corridor that contains K . If K is maximal, condition (iv) is “tight” for at least one portal σ of K in the sense that there is a vertex of $\mathcal{F}$ (not necessarily an endpoint of e_i or e_j) on σ . In particular, there is a vertex of $\mathcal{F}$ on the shorter portal of K ; if both portals have the same length, both contain such vertices. Let $\mathcal{K}$ be the set of all maximal corridors in $\mathcal{F}$, and let $\mathbf{X}$ be the set of vertices of $\mathcal{F}$ (including s_A, s_B, t_A, t_B) and the vertices of corridors in $\mathcal{K}$. We charge each corridor $K \in \mathcal{K}$ to a vertex of $\mathcal{F}$ or s_A, t_A, s_B, t_B on its shorter portal. It can be shown that each such vertex is charged $O(1)$ times. Since there are $O(n)$ vertices of $\mathcal{F}$, we have $|\mathcal{K}|, |\mathbf{X}| = O(n)$.

Let ℓ_L, ℓ_R be the lines supporting the portals σ_L, σ_R of K , and let $\text{len}(K)$ be the L_∞ -distance between ℓ_L, ℓ_R . Let u_L (resp., u_R) be the inner normal of σ_L (resp., σ_R), i.e., pointing toward the interior of K ; $u_L = -u_R$. For $D = L, R$ and any value $\tau \geq 0$, let $\ell_D^{(\tau)}$ be the line ℓ_D shifted in direction u_D at L_∞ -distance τ from ℓ_D , let $\sigma_D^{(\tau)}$

⁴If v_i or v_j is not unique, i.e., when e_i or e_j are axis-aligned, we can choose v_i or v_j (or both) so that $u_{ij} \in \{0, \pi/2\}$.




Figure 5. Illustrations of various portal-parallel lines supporting segments in K , and the sanctum K^S of K .




Figure 6. Illustration of the proof of Lemma 2.2.

be the segment $K \cap \ell_D^{(\tau)}$, and let $K^{(\tau)} \subseteq K$ be the (possibly empty) trapezoid bounded by the blockers of K and segments $\sigma_L^{(\tau)}, \sigma_R^{(\tau)}$. (We assume here that τ is sufficiently small so as to guarantee the shifts from ℓ_L to $\ell^{(\tau)}$ and from ℓ_R to $\ell_R^{(\tau)}$ do not collide.) Note that $K = K^{(0)}$. Similarly, we define portal-parallel lines and segments by points that they contain: For any point $p \in K$, let ℓ_p be the line normal to u_L (and u_R) containing p , and let $\sigma_p := K \cap \ell_p$. For any corridor $K \in \mathcal{K}$ with $\text{len}(K) \geq 20$, we define its *sanctum* to be $K^S := K^{(10)} \subset K$. See Figure 5. A corridor $K \in \mathcal{K}$ with $\text{len}(K) < 20$ has an empty sanctum. The following two lemmas capture the essence of a corridor.

LEMMA 2.2. *Let $K \in \mathcal{K}$ be a maximal corridor, and let u its direction, i.e., one of the unit vectors normal to the portals of K . Let I be a time interval in a plan π of A and B , during which both robots are in K , i.e., $\pi_A(\lambda), \pi_B(\lambda) \in K$ for all $\lambda \in I$. Then the sign of $g(\lambda) := \langle \pi_A(\lambda) - \pi_B(\lambda), u \rangle$ is the same for all $\lambda \in I$, where $\langle \cdot \rangle$ is the inner product.*

Proof. Suppose to the contrary that there exist two time instances $\lambda_1, \lambda_2 \in I$, with $\lambda_1 < \lambda_2$, such that $g(\lambda_1) < 0$ and $g(\lambda_2) > 0$ (or the other way around). Since π_A, π_B are continuous functions, there exists $\lambda_0 \in (\lambda_1, \lambda_2)$ with $g(\lambda_0) = 0$. But then $\pi_A(\lambda_0)$ and $\pi_B(\lambda_0)$ lie on a segment parallel to the portals of K and thus $\|\pi_A(\lambda_0) - \pi_B(\lambda_0)\|_\infty < 2$, which means that the robots intersect at these placements, contradicting the assumption that π is a feasible plan. Hence, $g(\cdot)$ has the same sign over the entire interval I . See Figure 6. $\square$

The following lemma describes a crucial relationship between revolving areas and corridors, whose proof is found in the full version [3].

LEMMA 2.3. *Suppose $p \in \mathcal{F}$ is a point such that p does not lie in any corridor of $\mathcal{K}$ and $d_\infty(p, \mathcal{X}) \geq 2$. Then there is a revolving area $q + \square \subseteq \mathcal{F}$, for some $q \in \mathcal{F}$, that contains p .*

3 Well-structured Optimal Plans

We present a sequence of transformations for optimal plans, which leads to the existence of an optimal plan with certain desirable properties. A few auxiliary lemmas and the proofs of the stated lemmas are found in the full version [3]. Using the easily established fact that $\mathbf{F}$ is polyhedral, it can be shown that an optimal plan is piecewise linear with its breakpoints lying on 2-faces of $\partial \mathbf{F}$. We show that there always exists a piecewise-linear, decoupled plan such that a robot is never parked in the sanctum of a corridor, and the moving robot *kisses* the parked robot in each move, except possibly in the first and the last moves.




Figure 7. Example of three moves, π_i and π_{i+2} of robot A and π_{i+1} of robot B , in a plan π . By modifying π_i, π_{i+2} so that A parks at p'_{i+1} instead of p_{i+1} , B kisses A during move π_i .

We first observe that each facet (three-dimensional face) of $\partial \mathbf{F}$ corresponds to a maximal connected set of placements at which some vertex (resp., edge) of one of the robots touches some edge (resp., vertex) of $\partial \mathcal{W}$ or of the other robot. This implies that each connected component of $\partial \mathbf{F}$ is a polyhedral region in $\mathbb{R}^4$. The distance between two points $a, b \in \mathbf{F}$ is the sum of the Euclidean length of the projections of $b - a$ onto the 2-planes formed by the first and the last pairs of coordinates, so it is the L_1 -distance of two L_2 -distances. Still, we claim that an optimal path (in $\mathbf{F}$) must be piecewise linear, with bends only at 2-faces (or faces of lower dimension) of $\partial \mathbf{F}$. This follows since both the L_2 and L_1 -distances satisfy the triangle inequality, and since paths that bend at the relative interior of some 3-face of $\mathbf{F}$ can be shortened. Hence, from now on we only consider piecewise-linear plans.

3.1 Decoupled optimal plans

We begin by proving that there always exists an optimal (piecewise-linear) plan that is decoupled, i.e., only one robot moves at any given time. Such *decoupled* plans are desirable, as during the motion of the moving robot, the parked robot can be treated as an additional obstacle that is part of the environment. Thus, given the start and target placements, s and t , of the moving robot, at some single move in the plan, and the position p of the parked robot, the optimal motion for the moving robot is the shortest path from s to t in $\mathcal{F}[p]$.

LEMMA 3.1. *Given reachable configurations $s, t \in \mathbf{F}$, there is always a decoupled, optimal (s, t) -plan.*

We sketch the proof here and refer the reader to the full version [3] for the rest of the details. We begin with a (piecewise-linear) optimal (s, t) -plan $\pi = \langle s = x^0, x^1, \dots, x^k = t \rangle$ in $\mathbf{F}$, where $\pi^i = x^{i-1}x^i$, for $1 \leq i \leq k$, is a line segment in $\mathbf{F}$. Let π_A^i (resp., π_B^i) be the line segment in $\mathcal{F}$ along which A (resp., B) moves during π^i . We show that π^i can be decoupled by either moving A along π_A^i and then moving B along π_B^i , or vice-versa. In particular, we show that if neither of these decoupled plans were feasible, there would exist a time $\lambda^* \in [0, 1]$ for which $\|\pi_A(\lambda^*) - \pi_B(\lambda^*)\|_\infty < 2$, i.e., $\pi^i(\lambda^*) \notin \mathbf{F}$, so π^i is not a feasible plan, which is a contradiction.

3.2 Kissing plans

We call a decoupled plan π a *kissing plan* if the robots kiss on all but possibly the first and the last moves. Formally, let $\langle \pi \rangle = (R_1, \pi_1, p_1), \dots, (R_k, \pi_k, p_k)$ be the move sequence of π . Then π is a kissing plan if, for all $1 < i < k$, there exists a point $q_i \in \pi_i$ such that (p_i, q_i) is a kissing configuration. We show that a decoupled plan can be converted into a kissing plan, without changing the images of the paths traveled by A and B in the plan, by reducing the number of moves and adjusting the parking places (Figure 7). We obtain the following:

LEMMA 3.2. *Let π be a decoupled, optimal plan with the minimum number of moves. There exists a decoupled, kissing, optimal plan π' with the same number of moves, such that the first move is made by the same robot as in π , and the pathlet of the first move in π' contains that of π .*

Proof. The proof is by induction on k . Let $\langle \pi \rangle = (R_1, \pi_1, p_1), \dots, (R_k, \pi_k, p_k)$. If $k = 2$, the claim holds trivially, that is, vacuously, so assume $k > 2$. Without loss of generality, A moves first, i.e., $R_1 = A$. Then $(p_1 = s_B), p_3, p_5, \dots$ are the parking placements of B ; $(p_0 = s_A), p_2, p_4, \dots$ are the parking placements of A ; $p_k = t_A, p_{k+1} = t_B$ if k is odd, and $p_k = t_B, p_{k+1} = t_A$ if k is even; $\pi_1 := \pi_A[s_A, p_2]$ is the motion of A in the first move and $\pi_2 := \pi_B[s_B, p_3]$ is the motion of B in its first move. There are two cases to consider.

1. If $(\pi_3 \oplus \square) \cap (\pi_2 \oplus \square) = \emptyset$ then

$$\pi' := \begin{cases} (A, \pi_1 \parallel \pi_3, p_1), (B, \pi_2 \parallel \pi_4, p_4), (R_5, \pi_5, p_5), \dots, (R_k, \pi_k, p_k) & \text{if } k > 3 \\ (A, \pi_1 \parallel \pi_3, s_B), (B, \pi_2, t_A) & \text{if } k = 3 \end{cases}$$

is a decoupled, optimal plan with fewer than k moves, which contradicts the assumption that π has the fewest moves among all decoupled, optimal plans.

2. If $(\pi_3 \oplus \square) \cap (\pi_2 \oplus \square) \neq \emptyset$, let p' be the first point reached on π_3 such that $(p' + \square) \cap (\pi_2 \oplus \square) \neq \emptyset$; note that p' may be p_2 . By the choice of p' , the interior of $p' + \square$ is disjoint from $\pi_2 \oplus \square$, so B kisses A at that placement when moving along π_2 . Define $\pi_{3<} := \pi_3[p_2, p']$ and $\pi_{3>} := \pi_3[p', p_4]$. Again, the choice of p' also implies that $\pi_{3<} \oplus \square$ is interior disjoint from $\pi_2 \oplus \square$. Then

$$\pi' := (A, \pi_1 \parallel \pi_{3<}, s_B), (B, \pi_2, p'), (A, \pi_{3>}, p_3), (R_4, \pi_4, p_4), \dots, (R_k, \pi_k, p_k)$$

is a decoupled, optimal (s, t) -plan in which B kisses A , parked at p' , as it moves along π_2 .

Set $s' := (p', s_B)$. Let π'_0 be the decoupled (s', t) -plan composed of all but the first move of π' . Then $\alpha(\pi'_0) = \alpha(\pi') - 1 = \alpha(\pi) - 1$. Furthermore, π'_0 is a decoupled, optimal (s', t) -plan. We apply the induction hypothesis to π'_0 to obtain a decoupled, kissing, optimal (s', t) -plan π''_0 satisfying the lemma, with B making the first move (B, π'_2, p') . Since the lemma guarantees that $\pi_2 \subset \pi'_2$, B kisses A (parked at p') during the first move of π''_0 . Set $\pi'' := (A, \pi_1 \parallel \pi_{3<}, s_A) \parallel \pi''_0$. Then the robots kiss on all moves of π'' except possibly in the first and the last moves. Furthermore

$$\alpha(\pi'') = \alpha(\pi''_0) + 1 = \alpha(\pi'_0) + 1 = \alpha(\pi),$$

and $\pi_1 \subseteq \pi_1 \parallel \pi_{3<}$. Hence π'' satisfies the lemma, which establishes the induction step and thus completes the proof of the lemma.

□

3.3 Bounding alternations

In a sequence of lemmas, we show that for any $s, t \in \mathbf{F}$, there exists a decoupled, kissing, optimal (s, t) -plan π with $\alpha(\pi) = O(\mathfrak{C}(\pi) + 1)$. We begin with a simple observation whose proof is omitted.

LEMMA 3.3. *Let e be a horizontal or vertical segment of length at most 2. Then $e \cap \mathcal{F}$ is a connected (possibly empty) interval.*

For any region $\nabla \subseteq \mathcal{F}$ and any two points $p, q \in \nabla$, let $\varrho_{\nabla}(p, q)$ be the length of the shortest (p, q) -path in $\nabla \cap \mathcal{F}$. Note that $\varrho(p, q) = \varrho_{\mathcal{F}}(p, q)$. The proof of the following lemma is found in the full version [3].

LEMMA 3.4. *Let S be any axis-aligned unit-radius square. (i) $S \cap \mathcal{F}$ is composed of xy -monotone components (without holes). (ii) At most two components intersect ∂S . (iii) For any p, q that lie in the interior of a common component of $S \cap \mathcal{F}$, there exists an xy -monotone (p, q) -path P such that $|P| = \varrho_S(p, q) = \varrho(p, q)$.*

The next lemma shows that there is a simple optimal motion between configurations as long as they are sufficiently close and both x -separated or both y -separated.

LEMMA 3.5. *Let Q_A, Q_B be axis-aligned unit-radius squares. For $s = (s_A, s_B), t = (t_A, t_B) \in \mathbf{F}$ such that s_A, t_A (resp., s_B, t_B) lie in a common component of $\text{int}(Q_A) \cap \mathcal{F}$ (resp., $\text{int}(Q_B) \cap \mathcal{F}$) and s and t are both x -separated or both y -separated, there exists an optimal (and trivially kissing) plan π with $\mathfrak{C}(\pi) = \varrho(s_A, t_A) + \varrho(s_B, t_B)$ and $\alpha(\pi) \leq 2$.*

Proof. Without loss of generality, the configurations are x -separated. Using standard transformations as necessary, we can assume $x(s_A) - x(s_B) \geq 2$. Then $x(s_A) - 2 < x(t_A) < x(s_A) + 2$ (resp., $x(s_B) - 2 < x(t_B) < x(s_B) + 2$) since s_A, t_A (resp., s_B, t_B) lie in the interior of Q_A (resp., Q_B). (t_A, t_B) is x -separated so $|x(t_A) - x(t_B)| \geq 2$. If




Figure 8. Examples of x -separated configurations $\mathbf{s}, \mathbf{t}$ and squares Q_A, Q_B that satisfy Lemma 3.5, s_A is left of s_B . (left) s_A, s_B are both left of their respective target placements, t_A, t_B . B moves first from s_B to t_B and then A moves from s_A to t_A . (right) s_A is left of t_A but s_B is right of t_B , so both 2-move plans are feasible.

$x(t_A) - x(t_B) \leq -2$ then $x(t_A) < x(s_B) \leq x(s_A) - 2$, which is a contradiction. Hence $x(t_A) - x(t_B) \geq 2$. Let P_A be the xy -monotone (s_A, t_A) -path in $Q_A \cap \mathcal{F}$ and let P_B be the xy -monotone (s_B, t_B) -path in $Q_B \cap \mathcal{F}$ from Lemma 3.4. There are two cases.

First, suppose $x(s_A) - x(t_A)$ and $x(s_B) - x(t_B)$ are zero or their signs are the same, say, non-negative for concreteness. See Figure 8(left). Then P_B lies to the right of line $x = s_A + 2$ and hence $P_B \subset \mathcal{F}[s_A]$. Similarly, P_A lies to the left of line $x = t_B - 2$ and hence $P_A \subset \mathcal{F}[t_B]$.

Otherwise, $x(s_A) - x(t_A)$ and $x(s_B) - x(t_B)$ are non-zero and their signs are different; for concreteness, suppose $x(s_A) - x(t_A) < 0 < x(s_B) - x(t_B)$. See Figure 8(right). Then $x(s_A) < x(t_A) \leq x(t_B) - 2 < x(t_A) - 2$. See Figure 8. Then P_B lies to the right of line $x = x(t_A) + 2$, and hence right of line $x = x(s_A) + 2$, so $P_B \subset \mathcal{F}[s_A]$. Similarly, P_A lies to the left of line $x = x(t_B) - 2$ so $P_A \subset \mathcal{F}[t_B]$.

Thus, in either case, the desired plan π is to first move B along P_B while A is parked at s_A , then move A along P_A while B is parked at t_B , which is trivially kissing since it has at most two moves. The other cases are symmetric. $\square$

The previous lemma allows us to shortcut kissing plans and to use a packing argument to establish a useful upper bound on the number of moves in an optimal plan.

LEMMA 3.6. *Given reachable configurations $\mathbf{s}, \mathbf{t} \in \mathbf{F}$, there exists a decoupled, kissing, optimal $(\mathbf{s}, \mathbf{t})$ -plan $\pi = (\pi_A, \pi_B)$ with $\alpha(\pi) \leq c(\min\{\phi(\pi_A), \phi(\pi_B)\} + 1)$, for some global constant $c \geq 1$.*

Proof. Without loss of generality, assume $\phi(\pi_A) \leq \phi(\pi_B)$. Let $\mathbb{G}$ be the axis-aligned uniform grid with square cells of radius 1 such that all parking places lie in the interior of grid cells and π does not pass through a vertex of $\mathbb{G}$. Let $\mathcal{G} \subset \mathbb{G}$ be the set of grid cells that contain at least one parking place of A . It is easily seen that $|\mathcal{G}| \leq 4\phi(\pi_A)$. We will show that we can shortcut π to obtain a new plan π' if necessary so that $\phi(\pi') \leq \phi(\pi)$, A is parked only $O(1)$ times in each cell of $\mathcal{G}$, the parking places of A in π' are a subset of those in π and π' is also a kissing plan. For a cell $g \in \mathbb{G}$, let $N(g) \subset \mathbb{G}$ be the set of cells $g' \in \mathbb{G}$ such that there exists a pair of points $p \in g, q \in g'$ with $\|p - q\|_\infty = 2$, i.e., (p, q) is a kissing configuration. Note that $|N(g)| \leq 25$.

Fix a cell $g \in \mathcal{G}$. Let C be a connected component of $g \cap \mathcal{F}$ that contains a parking place of A . Recall that π is a kissing plan so B kisses A at each parking place of A . For each parking place ξ of A in C , we label it with cell $\tau \in \mathbb{G}$ if B was in cell τ when it kissed A at ξ . If there are more than one such cell, we arbitrarily choose one of them. If C contains more than two parking places of A with the same label τ such that all of them are x -separated or all of them are y -separated, then we shortcut π as follows. Let λ^- (resp., λ^+) be the first (resp., last) time instance such that $\pi(\lambda^-)$ (resp., $\pi(\lambda^+)$) is a x -separated kissing configuration with $\pi_A(\lambda^-) \in \xi, \pi_B(\lambda^-) \in \tau$ (resp., $\pi_A(\lambda^+) \in \xi, \pi_B(\lambda^+) \in \tau$). We replace $\pi(\lambda^-, \lambda^+)$ with the $(\pi(\lambda^-), \pi(\lambda^+))$ -plan described in Lemma 3.5 of cost $\varrho(\pi_A(\lambda^-), \pi_A(\lambda^+)) + \varrho(\pi_B(\lambda^-), \pi_B(\lambda^+))$. We repeat this procedure in C until there are no such parking places of A in g . We repeat this step for all cells $g \in \mathcal{G}$. Let π' be the resulting plan. By construction, $\phi(\pi') \leq \phi(\pi)$ and π' is a kissing plan.

We now bound $\alpha(\pi')$. First note that π intersects at most two components of $g \cap \mathcal{F}$ for each cell $g \in \mathcal{G}$. For each such connected component, the plan π' has at most $4|N(g)| \leq 100$ parking places. Therefore g contains at most 200 parkings of A in the plan π' . Summing over all cells of $\mathcal{G}$, we obtain that A is parked $O(|\pi_A|) = O(\phi(\pi_A) + 1)$ times in the plan π' . Since A and B park alternately, $\alpha(\pi') = O(\phi(\pi_A) + 1)$. $\square$




Figure 9. (left) Example of a plan from $\pi(\lambda_A^-) = (\xi_A, \xi_B)$ to $\pi(\lambda_A^+) = (\zeta_A, \zeta_B)$ that satisfies Lemma 4.2, i.e., $\pi_A(I_A)$ (blue) does not enter K^S , where A first moves from ξ_A to γ_A , then B moves from ξ_B to ζ_B , and then A moves to ζ_A from γ_A . (right) Example of an optimal (ξ, ζ) -plan for ξ, ζ as defined in Lemma 4.3. Figures are not drawn to scale.

4 Paths Inside a Corridor

In this section we prove the existence of a decoupled, kissing, optimal plan in which neither of the two robots is ever parked in the sanctum of a corridor. We prove this result by introducing some convenient notations and establishing a few properties of a decoupled path inside a corridor. The proof is quite involved, so we only provide a sketch here and refer to the full version [3] for a full proof.

Suppose π is an optimal, decoupled, and kissing (s, t) -plan, and let $K \in \mathcal{K}$ be a corridor such that one of the robots, say, A , enters K and parks inside the sanctum K^S of K . We have that $s_A, s_B, t_A, t_B \notin \text{int}(K)$ since no point in X lies in the interior any corridor by definition. Let $I_A := [\lambda_A^-, \lambda_A^+]$ be a maximal time interval during which (the center of) A is inside $K^{(2)} \subset K$, which contains the time λ at which $\pi_A(\lambda) \in K^S$. Let σ_0, σ_1 be the (not necessarily distinct) portals of K last crossed in $\pi_A(0, \lambda_A^-)$ and first crossed in $\pi_A(\lambda_A^+, 1)$, respectively. Then $\pi_A(\lambda_A^-), \pi_A(\lambda_A^+)$ lie on the edges $\sigma_0^{(2)}$ and $\sigma_1^{(2)}$ of $K^{(2)}$, respectively, which again are not necessarily distinct. We first prove a few properties of π and then show that π can be transformed to another decoupled (s, t) -plan without increasing the cost, so that neither A nor B parks inside the sanctum K^S during the interval I_A . First, we argue that B also enters K during the interval I_A :

LEMMA 4.1. *There is a maximal interval I_B such that $I_A \cap I_B \neq \emptyset$ and B is in K during I_B , i.e., $\pi_B(\lambda) \in K$ for all $\lambda \in I_B$. Furthermore B enters and exits K during I_B through the same portals as A .*

Proof. If $\pi_B(\lambda) \notin K$ for all $\lambda \in I_A$, then $K^S \subseteq \bigcap_{\lambda \in I_A} \mathcal{F}[\pi_B(\lambda)]$ and there is no need to park A inside K^S . That is, we can first move A along $\pi_A(I_A)$ while B is parked at $\pi_B(\lambda_A^-)$, then park A at the portal σ_1 and move B along $\pi_B(I_A)$, and then follow the rest of the plan, $\pi(\lambda_A^-, 1)$. So we assume that $\pi_B(\lambda) \in K$ for some $\lambda \in I_A$.

Let $I_B := [\lambda_B^-, \lambda_B^+]$ be a maximal interval with $I_A \cap I_B \neq \emptyset$ during which B is in K . Let u be a vector normal to σ_0 . By Lemma 2.2, the sign of $\langle \pi_A(\lambda) - \pi_B(\lambda), u \rangle$ is the same for all $\lambda \in I_A \cap I_B$. If A and B enter through different portals of K , then we claim that π is not an optimal plan. Indeed, if A enters and exits at the same portal, (i.e., $\sigma_0 = \sigma_1$), then we can shortcut $\pi_A(I_A)$ along $\sigma_0^{(2)}$ to obtain a cheaper (s, t) -plan, and if A exits at the other portal (i.e., $\sigma_0 \neq \sigma_1$), we can shortcut $\pi_B(I_B)$ along that portal (at which B entered) to obtain a cheaper (s, t) -plan. A similar short-cutting argument holds if B does not exit through the same portal as A . This completes the proof of the lemma. $\square$

By Lemma 4.1, we can assume that B also enters K through the portal σ_0 and exits K through σ_1 .

LEMMA 4.2. *If $\sigma_0 = \sigma_1$, i.e., A enters and exits K from the same portal during interval I_A then A does not enter the sanctum K^S during I_A . Similarly B does not enter K^S during the interval $I_A \cap I_B$ (Figure 9(left)).*

In light of the previous lemma, we further assume that A and B both exit K at the same portal, in addition to the assumption that they both enter K through its other portal. Next, we prove⁵ a key property of optimal plans inside a corridor, which relies on the characterization of optimal plans without obstacles [15, 31].

⁵We only need a weaker version of Lemma 4.3 which we prove in [3].

LEMMA 4.3. Let ξ_A (resp. ζ_A) be a point on $\sigma_0^{(2)}$ (resp. $\sigma_1^{(2)}$). Let $\xi_B \in \mathcal{F}[\xi_A]$ be a point lying between σ_0 and $\sigma_0^{(2)}$, and let $\zeta_B \in \mathcal{F}[\zeta_A]$ be a point lying between $\sigma_1^{(4)}$ and $\sigma_1^{(2)}$. Set $\xi := (\xi_A, \xi_B)$ and $\zeta := (\zeta_A, \zeta_B)$. There exists a decoupled, optimal (ξ, ζ) -plan $\psi = (\psi_A, \psi_B) \subset K \times K$ that consists of two moves: first move A along the shortest (ξ_A, ζ_A) -path in $\mathcal{F}[\xi_B]$ while B is being parked at ξ_B , and then move B along the shortest (ξ_B, ζ_B) -path in $\mathcal{F}[\zeta_A]$ while A is being parked at ζ_A (Figure 9(right)).

Using Lemma 4.3, we prove that A does not park inside K during the interval I_A .

LEMMA 4.4. Suppose a robot, say A , is parked inside the sanctum of a corridor K at time $\lambda \in [0, 1]$ in a decoupled, kissing, optimal (s, t) -plan π , and let I_A be a maximal time interval with $\lambda \in I_A$ during which A is inside $K^{(2)}$. Let I_B be the maximal time interval with $I_A \cap I_B \neq \emptyset$ as given by Lemma 4.1. Then there exists a decoupled, optimal, and kissing (s, t) -plan π' and an interval $I \supseteq I_A \cup I_B$ such that neither A nor B parks inside the sanctum K^S of K during $\pi'(I)$ and $\pi(\lambda) = \pi'(\lambda)$ for all $\lambda \notin I$.

By applying Lemma 4.4 repeatedly, we obtain the following corollary.

COROLLARY 4.1. For any reachable configurations $s, t \in \mathbf{F}$, there exists a decoupled, optimal, and kissing (s, t) -plan in which no robot parks inside the sanctum of a corridor of $\mathcal{K}$.

5 Near-Optimal Tame Plans

Recall that $\mathbf{X}$ is the set of vertices of $\mathcal{F}$ (including s_A, s_B, t_A, t_B) and the vertices of all maximal corridors in $\mathcal{K}$, i.e., the endpoints of their portals. In this section, we show that a kissing, decoupled, optimal plan can be deformed by paying a fixed (constant) cost so that all robots are parked near a point of $\mathbf{X}$. We sketch the proof here and refer to the full version [3] for the rest of the details. For two parameters Δ^-, Δ^+ with $0 \leq \Delta^- \leq \Delta^+$, we say that a point $p \in \mathcal{F}$ is (Δ^-, Δ^+) -close (to X) if $d_\infty(p, X) \in [\Delta^-, \Delta^+]$. Often we will be interested in only one of Δ^- and Δ^+ , so we say is Δ -close (resp., Δ -far, Δ -tight) if $d_\infty(p, X) \leq \Delta$ (resp., $d_\infty(p, X) \geq \Delta, d_\infty(p, X) = \Delta$). A decoupled (s, t) -plan $\pi = (\pi_A, \pi_B)$ is called Δ -tame (or tame if the value of Δ is clear from the context) if every parking place on π_A, π_B is Δ -close. The following lemma is the main result of this section and one of the crucial properties on which our algorithm relies. Throughout this section, we set $\Delta_0 := 30$, which is simply a constant that is sufficiently large for our needs.

LEMMA 5.1. Given reachable configurations $s, t \in \mathbf{F}$, let π be a decoupled, kissing (s, t) -plan. For any parameter $\Delta \geq \Delta_0$, there exists a decoupled, kissing, Δ -tame (s, t) -plan π' such that $\pi' = \pi$ if $\phi(\pi) \leq \Delta$, and $\phi(\pi') \leq \phi(\pi) + c_1$ and $\alpha(\pi') \leq \alpha(\pi) + c_2$ otherwise, where $c_1 \geq \Delta_0$ and $c_2 > 0$ are absolute constants that do not depend on Δ .

For any $\varepsilon \in (0, 1]$ and optimal plan π^* , if $\phi(\pi^*) \leq c_1/\varepsilon$, then π^* is obviously (c_1/ε) -tame (recalling that s_A, s_B, t_A, t_B are in $\mathbf{X}$). Otherwise, by Lemma 5.1, there exists a (c_1/ε) -tame (s, t) -plan of cost at most $\phi(\pi^*) + c_1 \leq (1 + \varepsilon)\phi(\pi^*)$. Hence, using Lemma 3.6 to bound the number of moves, we obtain:

COROLLARY 5.1. Given reachable configurations $s, t \in \mathbf{F}$ and $\varepsilon \in (0, 1]$, there exists a decoupled, kissing, (c_1/ε) -tame (s, t) -plan π with $\phi(\pi) \leq (1 + \varepsilon)\phi(\pi^*)$ and $\alpha(\pi) \leq c_2(\phi(\pi^*) + 1)$, where $c_1, c_2 > 0$ are absolute constants that do not depend on ε .

Let π be an optimal, decoupled, kissing (s, t) -plan. By Corollary 4.1, we can assume that no robot is parked inside the sanctum of a corridor. Let $\ell := \alpha(\pi)$ and let $(R_1, \pi_1, p_1), \dots, (R_\ell, \pi_\ell, p_\ell)$ be the sequence of moves of π . Let i (resp., j), $1 < i \leq j < \ell$, be the smallest (resp., largest) index such that p_i, p_j are $(\Delta - 4)$ -far, i.e., p_i (resp. p_j) is the first (resp. last) $(\Delta - 4)$ -far parking place in π . If there are no such indices, then π is Δ -tame and we are done. So suppose i, j exist. Note that it can be that $i = j$. By the definitions of corridors and sanctums, p_i and p_j do not lie inside a corridor $K \in \mathcal{K}$ because any point in $K \setminus K^S$ is $(\Delta_0 - 4)$ -close, p_i, p_j are $(\Delta - 4)$ -far, and $\Delta \geq \Delta_0$. Therefore there is a revolving area around each of p_i and p_j by Lemma 2.3.

The proof of Lemma 5.1 is based on the following observation, which is proved in the full version [3]. Let $s = (s_A, s_B), t = (t_A, t_B) \in \mathbf{F}$ be reachable kissing configurations with the property that there exist $r^-, r^+ \in \mathcal{F}$ such that $s_A, s_B \in \text{RA}(r^-), t_A, t_B \in \text{RA}(r^+)$, and r^-, r^+ are 3-far. Then there exists a decoupled, kissing (s, t) -plan $\tilde{\pi}$ with $\phi(\tilde{\pi}) \leq \varrho(s_A, t_A) + \varrho(s_B, t_B) + O(1)$ and $\alpha(\tilde{\pi}) = O(1)$, and all parking places in $\tilde{\pi}$ lie in $\text{RA}(r^-)$ or $\text{RA}(r^+)$. Since π is a kissing plan, there are kissing configurations $q = (q_A, q_B)$ and $q' = (q'_A, q'_B)$ on moves i and j . If q_A, q'_A, q_B, q'_B each is $(\Delta - 2)$ -close and lies in a revolving area then Lemma 5.1 follows from this observation but

we may not be so lucky— q_A or q'_A may not be $(\Delta - 2)$ -close or may not lie in revolving areas, so the proof is much more involved and deferred to the full version [3].

The proof of the previous observation, however, can be slightly adapted to prove the following variant:

LEMMA 5.2. *Let $\mathbf{u} = (u_A, u_B), \mathbf{v} = (v_A, v_B) \in \mathbf{F}$ be two configurations such that there exist four points $\bar{u}_A, \bar{u}_B, \bar{v}_A, \bar{v}_B \in \mathcal{F}$ with $u_A \in \text{RA}(\bar{u}_A), u_B \in \text{RA}(\bar{u}_B), v_A \in \text{RA}(\bar{v}_A), v_B \in \text{RA}(\bar{v}_B)$, and $\bar{u}_A, \bar{u}_B, \bar{v}_A, \bar{v}_B$ are 3-far, then there exists a decoupled $(\mathbf{u}, \mathbf{v})$ -plan $\tilde{\pi}$ with $\mathfrak{c}(\tilde{\pi}) \leq \varrho(u_A, u_B) + \varrho(v_A, v_B) + 78$, $\alpha(\tilde{\pi}) \leq 40$, and all parking places of $\tilde{\pi}$ lie in the four revolving areas.*

Using Lemma 5.2, we can prove the following weaker version of Lemma 5.1, which guarantees that $\tilde{\pi}$ is Δ -tame but does not guarantee the kissing property.

LEMMA 5.3. *Given reachable configurations $\mathbf{s}, \mathbf{t} \in \mathbf{F}$, let π be a decoupled, kissing $(\mathbf{s}, \mathbf{t})$ -plan. For any parameter $\Delta \geq \Delta_0$, there exists a decoupled Δ -tame $(\mathbf{s}, \mathbf{t})$ -plan $\tilde{\pi}$ such that $\tilde{\pi} = \pi$ if $\mathfrak{c}(\pi) \leq \Delta$ and $\mathfrak{c}(\tilde{\pi}) \leq \mathfrak{c}(\pi) + c_1$ and $\alpha(\tilde{\pi}) \leq \alpha(\pi) + c_2$ otherwise, for some absolute constants $c_1, c_2 > 0$ independent of Δ .*

Proof. Let p_i, p_j be as defined above. Suppose $R_i = B$, i.e., A moves from p_{i-2} to p_i in the $(i - 1)$ -st move along π_{i-1} and is parked at p_i , then B moves from p_{i-1} to p_{i+1} along π_i in the i -th move. Let u_A be the last point along π_{i-1} that is $(\Delta - 4)$ -close, i.e., $d_\infty(\pi_{i-1}[u_A, p_i], X) \geq \Delta - 4$. Recall that p_{i-2} is $(\Delta - 4)$ -close. Note that u_A may be p_{i-2} or p_i , and u_A is $(\Delta - 4)$ -tight. Since p_i does not lie in a corridor, we claim that u_A also does not lie inside a corridor. Indeed if $u_A \in K$ for some $K \in \mathcal{K}$, then A exits K at some point $\xi \in \pi_{i-1}[u_A, p_i]$ but then $d_\infty(\xi, X) < \Delta_0 - 4 \leq \Delta - 4$, contradicting that u_A is the last $(\Delta - 4)$ -close point on π_{i-1} . Since u_A does not lie in a corridor, by Lemma 2.3, there is a $(\Delta - 5, \Delta - 3)$ -close point $\bar{u}_A \in \mathcal{F}$ such that $u_A \in \text{RA}(\bar{u}_A)$.

Next, B kisses A parked at p_i during the i -th move. Since p_i is $(\Delta - 4)$ -far, π_i contains a $(\Delta - 6)$ -far point. If $\pi_i \cap (u_A + 2\Box) = \emptyset$, let u_B be the last $(\Delta - 6)$ -close point on π_i if there exists one and $u_B = p_{i-1}$ otherwise (i.e., all points on π_i are $(\Delta - 6)$ -far). Then u_B is $(\Delta - 6)$ -tight. On the other hand, if $\pi_i \cap (u_A + 2\Box) \neq \emptyset$, let u_B be the first intersection point of π_i with $u_A + 2\Box$, i.e., $\pi_i[p_{i-1}, u_B] \cap \text{int}(u_A + 2\Box) = \emptyset$. Since u_A is $(\Delta - 4)$ -tight, u_B is $(\Delta - 6, \Delta - 2)$ -close. Since p_i and u_A are not inside a corridor, a similar argument as above implies that u_B is also not in a corridor. Therefore there exists a $(\Delta - 7, \Delta - 1)$ -close point $\bar{u}_B$ such that $u_B \in \text{RA}(\bar{u}_B)$. Set $\mathbf{u} = (u_A, u_B)$.

Without loss of generality, assume that $R_j = B$. Then using a symmetric argument, we find points $v_A \in \pi_{j+1}$ such that $v_A \in \text{RA}(\bar{v}_A)$ and v_A is $(\Delta - 4)$ -close, and $v_B \in \pi_j$ such that v_B is $(\Delta - 6, \Delta - 4)$ -close and $v_B \in \text{RA}(\bar{v}_B)$, for some $(\Delta - 7, \Delta - 1)$ -close points $\bar{v}_A, \bar{v}_B \in \mathcal{F}$. Set $\mathbf{v} = (v_A, v_B)$.

Since $\bar{u}_A, \bar{u}_B, \bar{v}_A, \bar{v}_B$ each is $(\Delta - 7)$ -far and $\Delta - 7 \geq \Delta_0 - 7 \geq 3$, each is $(3, \Delta - 1)$ -close. Let $\psi = (\psi_A, \psi_B)$ be the decoupled $(\mathbf{u}, \mathbf{v})$ -plan according to Lemma 5.2, with $\langle \psi \rangle = (S_1, \psi_1, q_1), \dots, (S_h, \psi_h, q_h)$. We obtain a new $(\mathbf{s}, \mathbf{t})$ -plan $\tilde{\pi}$ by replacing $\pi_A[u_A, v_A]$ and $\pi_B[u_B, v_B]$ with ψ_A and ψ_B , respectively. More precisely,

$$\begin{aligned} \langle \tilde{\pi} \rangle = & (R_1, \pi_1, p_1), \dots, (R_{i-2}, \pi_{i-2}, p_{i-2}), (A, \pi_{i-1}[p_{i-2}, u_A], p_{i-1}), (B, \pi_i[p_{i-1}, u_B], u_A) \\ & \circ \langle \psi \rangle \circ \\ & (B, \pi_j[v_B, p_{j+1}], v_A), (A, \pi_{j+1}[v_A, p_{j+2}], p_{j+1}), (R_{j+2}, \pi_{j+2}, p_{j+2}), \dots, (R_\ell, \pi_\ell, p_\ell). \end{aligned}$$

It is easily seen that $\tilde{\pi}$ is a (feasible) $(\mathbf{s}, \mathbf{t})$ -plan. By Lemma 5.2, all parking places in ψ and thus in $\tilde{\pi}$ are Δ -close, $\mathfrak{c}(\tilde{\pi}) \leq \mathfrak{c}(\pi) + 78$, and $\alpha(\tilde{\pi}) \leq \alpha(\pi) + 40$. $\square$

A similar argument as for Corollary 5.1, but using Lemma 5.3, implies the following corollary.

COROLLARY 5.2. *Given reachable configurations $\mathbf{s}, \mathbf{t} \in \mathbf{F}$ and $\varepsilon \in (0, 1]$, there exists a decoupled (c_1/ε) -tame $(\mathbf{s}, \mathbf{t})$ -plan π with $\mathfrak{c}(\pi) \leq (1 + \varepsilon)\mathfrak{c}(\pi^*)$ and $\alpha(\pi) \leq c_2(\mathfrak{c}(\pi^*) + 1)$, where $c_1, c_2 > 0$ are absolute constants that do not depend on ε .*

Returning to the proof of Lemma 5.1, we first briefly sketch the idea. Let $\lambda_i \in [0, 1]$ (resp., $\lambda_j \in [0, 1]$) be the earliest (resp., latest) time during the move i (resp., j) such that $\pi(\lambda_i)$ (resp., $\pi(\lambda_j)$) is a kissing configuration; there exists such a value since π is kissing. If $R_i = B$ then $\pi_A(\lambda_i) = p_i$ and $\pi_B(\lambda_i) \in \pi_i$, and $\pi_A(\lambda_i) \in \pi_i$ and $\pi_B(\lambda_i) = p_i$ otherwise; the same holds for λ_j . We similarly define λ_{i-1} (resp., λ_{j+1}) to be the latest (resp., earliest) time during the move $i - 1$ (resp., $j + 1$) such that $\pi(\lambda_{i-1})$ (resp., $\pi(\lambda_{j+1})$) is a kissing configuration;

if no such configuration exists, then $i - 1 = 1$ (resp., $j + 1 = \ell$) and we set $\lambda_{i-1} = 0$ (resp., $\lambda_{j+1} = 1$). Then $0 \leq \lambda_{i-1} \leq \lambda_i \leq \lambda_j \leq \lambda_{j+1} \leq 1$. If $i = j$ then $\pi_A(\lambda_i, \lambda_j)$ and $\pi_B(\lambda_i, \lambda_j)$ are points. For $0 \leq r \leq 3$, let $a_r := \pi_A(\lambda_{i-1+r})$ and $b_r := \pi_B(\lambda_{i-1+r})$. Without loss of generality, $R_{i-1} = A$ and $R_i = B$, so A moves first from $a_0 = p_{i-2}$ to $a_1 = p_i$ then B moves from b_0 to b_1 in the given motion plan $\pi(\lambda_i, \lambda_j)$. The proof of Lemma 5.1, found in the full version [3], is divided into two cases:

- (i) There exists a $(\Delta - 6)$ -close point on $\pi_A[\lambda_i, \lambda_j]$ or $\pi_B[\lambda_i, \lambda_j]$, say, on $\pi_A[\lambda_i, \lambda_j]$. In this case, we find two $(\Delta - 6)$ -close points q^-, q^+ on $\pi_A[\lambda_i, \lambda_j]$ and modify $\pi_A[\lambda_{i-1}, \lambda_{j+1}]$ and $\pi_B[\lambda_{i-1}, \lambda_{j+1}]$, using the above observation, so that A and B are parked at Δ -close points near $a_0, b_0, a_3, b_3, q^-,$ or q^+ and they lie in revolving areas. The surgery on π_A, π_B increases their lengths by $O(1)$ and adds $O(1)$ new alternations.
- (ii) There is no $(\Delta - 6)$ -close point on $\pi_A[\lambda_i, \lambda_j]$ or $\pi_B[\lambda_i, \lambda_j]$. In this case, we find Δ -close parking places in the vicinity of $\pi_A[\lambda_{i-1}, \lambda_i]$, $\pi_B[\lambda_{i-1}, \lambda_i]$, $\pi_A[\lambda_j, \lambda_{j+1}]$, and $\pi_B[\lambda_j, \lambda_{j+1}]$ and again modify $\pi_A[\lambda_{i-1}, \lambda_{j+1}]$ and $\pi_B[\lambda_{i-1}, \lambda_{j+1}]$. We cannot always guarantee the existence of revolving areas that contain parking places. Therefore the surgery as well as the analysis is more involved. Nevertheless, we are able to argue that the increase in the cost of the plan and in the number of alternations is $O(1)$.

6 Discretizing the Free Space

We next describe how the near-optimal tame plans described in Section 5 can be retracted to a path in a graph constructed over a discrete set of points. Let π be a decoupled, kissing, Δ -tame (s, t) -plan, and let $\varepsilon \in (0, 1)$ be a parameter. We can assume that $\alpha(\pi) = O(\mathfrak{c}(\pi) + 1)$. Let $\mathbb{G}$ be the axis-aligned uniform grid with square cells of radius ε such that all parking places lie in the interior of grid cells and π does not pass through a vertex of $\mathbb{G}$. Let $\mathcal{F}^\#$ be the overlay of $\mathbb{G}$ and $\mathcal{F}$, restricted to $\mathcal{F}$. Each face of $\mathcal{F}^\#$ is a connected component of $\mathcal{F} \cap g$ for some grid cell g of $\mathbb{G}$. Let $\mathcal{V}$ be the set of vertices of $\mathcal{F}^\#$. Our goal is to “retract” the parking places of π to the points of $\mathcal{V}$, i.e., robots are parked at the points of $\mathcal{V}$ instead of their original parking places. Furthermore, since π is kissing, we want to ensure that the retracted path is ε -nearly-kissing, i.e., whenever a robot moves, it comes within L_∞ -distance 4ε of the boundary of the other robot (parked at a vertex of $\mathcal{V}$). However, if for a parking place q , say, of A , we pick only one point in $\mathcal{V}$ to park A at instead of q , B may collide with A during its next move, especially if B kisses A at q during the next move. Hence, we choose multiple points of $\mathcal{V}$ (in the neighborhood of q) and move A between them during the next move of B to ensure that A and B do not collide. Furthermore, we want to maintain the property of being decoupled (i.e., only one robot moves at a time), which means when we move A between nearby points of $\mathcal{V}$ to make way for B , we must first park B somewhere, also in $\mathcal{V}$. These technical constraints make the retraction rather involved. The description and analysis of the retraction are given in the full version [3], which imply the following lemma.

LEMMA 6.1. *Let $\varepsilon \in (0, 1)$ be a parameter, and let π be a decoupled, Δ -tame (s, t) -plan. There exists a decoupled, $(\Delta + 2 + 4\varepsilon)$ -tame, (s, t) -plan π' such that $\mathfrak{c}(\pi') \leq \mathfrak{c}(\pi) + \varepsilon\alpha(\pi)$ and $\alpha(\pi') = \alpha(\pi)$, and every parking place of π' is in $\mathcal{V}$, for some constant $c > 0$ that does not depend on ε, Δ . If π is kissing, then π' is ε -nearly-kissing.*

7 Algorithm

We are now ready to describe our algorithm to compute an (s, t) -plan π with $\mathfrak{c}(\pi) \leq (1 + \varepsilon)\mathfrak{c}(\pi^*)$ for any $\varepsilon \in (0, 1]$. We first describe an $n^3\varepsilon^{-O(1)} \log n$ -time algorithm (Lemma 7.1) under the assumption that $\mathfrak{c}(\pi^*) > 1/4$. With further efforts, we present a near-quadratic time algorithm (Lemma 7.2) and how to remove the assumption (Section 7.2).

The algorithm consists of three stages. First, we choose a set $\tilde{\mathcal{V}}$ of $O(n/\varepsilon^4)$ points so that a robot is always parked at one of the points in $\tilde{\mathcal{V}}$. Next, we construct a graph $\mathcal{G} = (\mathcal{C}, \mathcal{E})$ where $\mathcal{C} \subseteq \tilde{\mathcal{V}} \times \tilde{\mathcal{V}}$ is a set of (feasible) configurations and each edge is a (decoupled) plan between a pair of configurations of $\mathcal{C}$ with one move. We compute a shortest path in $\mathcal{G}$, which corresponds to an (s, t) -plan $\hat{\pi}$ with $\mathfrak{c}(\hat{\pi}) \leq (1 + \varepsilon)\mathfrak{c}(\pi^*) + O(\varepsilon) \leq \mathfrak{c}(\pi) \leq (1 + O(\varepsilon))\mathfrak{c}(\pi^*)$ for $\mathfrak{c}(\pi^*) \geq 1/4$.

Set $\bar{\varepsilon} := \varepsilon/c_0$ and $\Delta := c_1/\bar{\varepsilon}$ where $c_0, c_1 \geq 1$ are sufficiently large constants (independent of ε) to be chosen later. Let $\mathbb{G}$, $\mathcal{F}^\#$, and $\mathcal{V}$ be the same as in Section 6 but using $\bar{\varepsilon}$ for ε . Let F_ε be the set of faces of $\mathcal{F}^\#$ that contain a $(\Delta + 2 + 4\bar{\varepsilon})$ -close point; any point in a face $C \in F_\varepsilon$ is $(\Delta + 2 + 6\bar{\varepsilon})$ -close. Let $\tilde{\mathcal{V}} \subseteq \mathcal{V}$ be the set of vertices of F_ε ; $|\tilde{\mathcal{V}}| = O(n\Delta^2/\bar{\varepsilon}^2) = O(n/\varepsilon^4)$. We now describe the weighted graph

$\mathcal{G} = (\mathcal{C}, \mathcal{E})$. We set $\mathcal{C} := \{(a, b) \in \tilde{\mathcal{V}} \times \tilde{\mathcal{V}} \mid \|a - b\|_\infty \geq 2\}$. Note that $\mathbf{s}, \mathbf{t} \in \mathcal{C}$ and $\mathcal{C} \subset \mathbf{F}$. We construct $\mathcal{E}$ as follows: For every ordered triple $(u, v, p) \in \tilde{\mathcal{V}} \times \tilde{\mathcal{V}} \times \tilde{\mathcal{V}}$ with $u \neq v$ and $\|p - u\|_2, \|p - v\|_2 \geq 2$, we set $\omega((u, p) \rightarrow (v, p)) = \omega((p, u) \rightarrow (p, v)) := \varrho_{\mathcal{F}[p]}(u, v)$, and if this value is not ∞ we add edges $(u, p) \rightarrow (v, p)$ and $(p, u) \rightarrow (p, v)$ to $\mathcal{E}$ with $\omega((u, p) \rightarrow (v, p)) = \omega((p, u) \rightarrow (p, v))$ as their weight, which corresponds to moving A (resp., B) from u to v along a shortest path in $\mathcal{F}[p]$ while B (resp., A) is parked at p . Then $|\mathcal{E}| = |\tilde{\mathcal{V}}|^3 = O(n^3/\varepsilon^{12})$.

Finally, we compute a shortest path (by weight) Φ in $\mathcal{G}$ from $\mathbf{s}$ to $\mathbf{t}$. After having computed Φ , the $(\mathbf{s}, \mathbf{t})$ -plan corresponding to Φ can be retrieved in a straightforward manner, and the cost of the resulting plan is the same as the weight of the path. We conclude by stating the following lemma:

LEMMA 7.1. *Given $\mathbf{s}, \mathbf{t} \in \mathbf{F}$, and $\varepsilon \in (0, 1)$, there exists a path Φ from $\mathbf{s}$ to $\mathbf{t}$ in $\mathcal{G}$, if $\mathbf{s}, \mathbf{t}$ are reachable, whose weight is at most $(1 + \varepsilon)\mathfrak{c}(\pi^*) + O(\varepsilon)$, which is $(1 + O(\varepsilon))\mathfrak{c}(\pi^*)$ if $\mathfrak{c}(\pi^*) > 1/4$, where π^* is a decoupled, optimal $(\mathbf{s}, \mathbf{t})$ -plan. Conversely, a path Φ from $\mathbf{s}$ to $\mathbf{t}$ in $\mathcal{G}$ corresponds to an $(\mathbf{s}, \mathbf{t})$ -plan $\hat{\pi}$ of cost $\omega(\Phi)$. Furthermore, a shortest path from $\mathbf{s}$ to $\mathbf{t}$ in $\mathcal{G}$ can be computed in $O(n^3\varepsilon^{-12} \log n)$ time.*

Proof. By Corollary 5.2, there exists a decoupled $(\Delta = c_1/\bar{\varepsilon})$ -tame plan π with $\mathfrak{c}(\pi) \leq (1 + \bar{\varepsilon})\mathfrak{c}(\pi^*)$ and $\alpha(\pi) \leq c_2(\alpha(\pi^*) + 1)$ for some constants $c_1, c_2 > 0$. (We make use of the stronger Corollary 5.1 that guarantees π is kissing when we improve the algorithm in the next subsection.) Then, by Lemma 6.1 with $\bar{\varepsilon}$ as parameter ε , there exists a $(\Delta + 2 + 4\bar{\varepsilon})$ -tame, decoupled plan π' such that all parking places belong to $\tilde{\mathcal{V}}$ and

$$\mathfrak{c}(\pi') \leq \mathfrak{c}(\pi) + \bar{\varepsilon}\alpha(\pi) \leq (1 + \bar{\varepsilon})\mathfrak{c}(\pi^*) + c_2\bar{\varepsilon}(\mathfrak{c}(\pi^*) + 1) \leq (1 + \bar{\varepsilon}(1 + c_2))\mathfrak{c}(\pi^*) + c_2\bar{\varepsilon}.$$

At this point, we have additive error $O(\varepsilon)$. Here we make use of our assumption that $\mathfrak{c}(\pi^*) > 1/4$ and have

$$\mathfrak{c}(\pi') \leq (1 + \bar{\varepsilon}(1 + 5c_2))\mathfrak{c}(\pi^*).$$

Then by choosing $c_0 := 1 + 5c_2$, we have $\bar{\varepsilon} = \varepsilon/(1 + 5c_2)$ and hence

$$\mathfrak{c}(\pi') \leq (1 + \varepsilon)\mathfrak{c}(\pi^*).$$

Let $\langle \pi' \rangle = (R_1, \pi_1, p_1), \dots, (R_\ell, \pi_\ell, p_\ell)$. Without loss of generality, assume that $R_1 = A$. Then we map $\hat{\pi}$ to a path from $\mathbf{s}$ to $\mathbf{t}$ in $\mathcal{G}$ as follows. For each $1 \leq i \leq \ell$, π_i is a path followed by one of the robots from p_{i-1} to p_{i+1} while the other is parked at p_i , so $\|p_{i+1} - p_i\|_\infty \geq 2$ and $\varrho_{\mathcal{F}[p_i]}(p_{i-1}, p_{i+1}) \leq \mathfrak{c}(\pi_i)$. Therefore $(p_{i-1}, p_i) \rightarrow (p_{i+1}, p_i), (p_i, p_{i-1}) \rightarrow (p_i, p_{i+1}) \in \mathcal{E}$ with their weights being at most $\mathfrak{c}(\pi_i)$. Hence $\mathbf{s} = (p_0, p_1) \rightarrow (p_2, p_1) \rightarrow (p_2, p_3) \rightarrow \dots \rightarrow \mathbf{t}$ is a path in $\mathcal{G}$ of weight at most $\mathfrak{c}(\pi)$.

Converting Φ to a decoupled $(\mathbf{s}, \mathbf{t})$ -plan of cost at most $\omega(\Phi)$ is straightforward and omitted from here. It remains to analyze the runtime of the algorithm. $\mathcal{F}$ and $\tilde{\mathcal{V}}$ can be computed in $O(n \log^2 n + |\tilde{\mathcal{V}}|) = O(n(\log^2 n + 1/\varepsilon^4))$ time [13]. For any ordered pair $(u, p) \in \tilde{\mathcal{V}} \times \tilde{\mathcal{V}}$, $\mathcal{F}[p]$ can be computed from $\mathcal{F}$ in $O(n \log n)$ time and processed [20] in $O(n \log n)$ time into a data structure that answers $O(\log n)$ -time shortest-path queries from u to any query point $v \in \mathcal{F}[p]$. So we can compute $\omega((u, p) \rightarrow (v, p)) = \omega((p, u) \rightarrow (p, v))$ in $O(n \log n + |\tilde{\mathcal{V}}| \log n) = O((n/\varepsilon^4) \log n)$ time, for all $v \in \tilde{\mathcal{V}}$. Repeating this process for all $O((n/\varepsilon^4)^2)$ pairs $(u, p) \in \tilde{\mathcal{V}} \times \tilde{\mathcal{V}}$, we compute $\mathcal{G}$ and its edge weights in $O(|\mathcal{E}| \log n) = O((n/\varepsilon^4)^3 \log n)$ time. Finally, computing the shortest path Φ in $\mathcal{G}$ and reporting its corresponding plan takes $O(|\mathcal{E}| + |\mathcal{C}| \log |\mathcal{C}|)$ time using Dijkstra's algorithm, which is dominated by the $O(|\mathcal{E}| \log n)$ time to build $\mathcal{G}$. Therefore the overall running time is $O(n^3\varepsilon^{-12} \log n)$. $\square$

7.1 Reducing the runtime

Now we describe how to reduce the runtime to $O(n^2\varepsilon^{-O(1)} \log n)$ using Corollary 5.1 (instead of Corollary 5.2). The high-level idea is to reduce the number of vertices, $|\mathcal{C}|$, from $O(n^3 \text{poly}(\log n, 1/\varepsilon))$ to $O(n^2 \text{poly}(\log n, 1/\varepsilon))$ while maintaining the $O(|\tilde{\mathcal{V}}|)$ degree of each node. The effect is that the size of each of $|\mathcal{C}|, |\mathcal{E}|$ reduces by a factor of n , which reduces the overall runtime by a factor of n .

We first describe the graph $\mathcal{G} = (\mathcal{C}, \mathcal{E})$. We set $\mathcal{C} := \{(a, b) \in \tilde{\mathcal{V}} \times \tilde{\mathcal{V}} \mid 2 \leq \|a - b\|_\infty \leq 2(1 + \bar{\varepsilon})\}$. Note the new condition that $\|a - b\|_\infty \leq 2(1 + \bar{\varepsilon})$. For a pair of nearby configurations $\mathbf{u} = (u_A, u_B), \mathbf{v} = (v_A, v_B) \in \mathcal{C}$, we consider two possible $(\mathbf{u}, \mathbf{v})$ -plans: (i) keep A parked at u_A while B moves from u_B to v_B along a shortest path in $\mathcal{F}[u_A]$, then park B at v_B and move A from u_A to v_A along a shortest path in $\mathcal{F}[v_B]$, and (ii) keep B parked at

u_B while A moves from u_A to v_A along a shortest path in $\mathcal{F}[u_B]$, then park A at v_A and move B from u_B to v_B along a shortest path in $\mathcal{F}[v_A]$. Set

$$\omega(\mathbf{u}, \mathbf{v}) := \min\{\varrho_{\mathcal{F}[u_B]}(u_A, v_A) + \varrho_{\mathcal{F}[v_A]}(u_B, v_B), \varrho_{\mathcal{F}[u_A]}(u_B, v_B) + \varrho_{\mathcal{F}[v_B]}(u_A, v_A)\}.$$

If $\omega(\mathbf{u}, \mathbf{v}) < \infty$, we add $\mathbf{u} \rightarrow \mathbf{v}$ to $\mathcal{E}$ with $\omega(\mathbf{u}, \mathbf{v})$ as its weight. Then $|\mathcal{E}| = |\tilde{\mathcal{V}}|^2 = O(n^2/\varepsilon^8)$. For a fixed configuration $\mathbf{u} := (u_A, u_B) \in \mathcal{C}$, we compute the shortest path from u_A to all points of $\tilde{\mathcal{V}}$ within $\mathcal{F}[u_B]$, using the same data structure as before [20], and do the same for u_B to all points of $\tilde{\mathcal{V}}$ in $\mathcal{F}[u_A]$. After repeating this step for all configurations in $\mathcal{C}$, we have all the information to compute $\omega(\mathbf{u}, \mathbf{v})$ for all $(\mathbf{u}, \mathbf{v}) \in \mathcal{C} \times \mathcal{C}$. The overall runtime can be shown to be $O(|\mathcal{E}| \log n)$ as before, which is $O(n^2 \varepsilon^{-8} \log n)$ here.

A similar argument for Lemma 7.1 that uses Corollary 5.1 instead of Corollary 5.2 proves the following lemma, which is the same as Lemma 7.1, except that the plan $\hat{\pi}$ is $\bar{\varepsilon}$ -nearly-kissing (and hence ε -nearly-kissing).

LEMMA 7.2. *Given $\mathbf{s}, \mathbf{t} \in \mathbf{F}$, and $\varepsilon \in (0, 1)$, there exists a path Φ from $\mathbf{s}$ to $\mathbf{t}$ in $\mathcal{G}$, if $\mathbf{s}, \mathbf{t}$ are reachable, whose weight is at most $(1 + \varepsilon)\mathfrak{c}(\pi^*) + O(\varepsilon)$, which is bounded by $(1 + O(\varepsilon))\mathfrak{c}(\pi^*)$ if $\mathfrak{c}(\pi^*) > 1/4$, where π^* is a decoupled, kissing, optimal $(\mathbf{s}, \mathbf{t})$ -plan. Conversely, a path Φ from $\mathbf{s}$ to $\mathbf{t}$ in $\mathcal{G}$ corresponds to a decoupled, ε -nearly-kissing $(\mathbf{s}, \mathbf{t})$ -plan $\hat{\pi}$ of cost $\omega(\Phi)$. Furthermore, a shortest path from $\mathbf{s}$ to $\mathbf{t}$ in $\mathcal{G}$ can be computed in $O(n^2 \varepsilon^{-8} \log n)$ time.*

7.2 Handling nearby configurations

We now describe how we compute an $(\mathbf{s}, \mathbf{t})$ -plan of cost at most $(1 + \varepsilon)\mathfrak{c}(\pi^*)$ even when $\mathfrak{c}(\pi^*) \leq 1/4$. The algorithm described in the following Section 8 (cf. Lemma 8.2) either reports an 8-approximation $\gamma \leq 2$ of $\mathfrak{c}(\pi^*)$, i.e., $\mathfrak{c}(\pi^*) \leq \gamma \leq 8\mathfrak{c}(\pi^*)$, or it reports that $\mathfrak{c}(\pi^*) > 1/4$. So we first run this algorithm. If it reports $\mathfrak{c}(\pi^*) > 1/4$, we run the algorithm above (with improved runtime). Otherwise, we have $\gamma \leq 2$ and $\mathfrak{c}(\pi^*) \leq \gamma \leq 8\mathfrak{c}(\pi^*)$. Then $\gamma/8 \leq \mathfrak{c}(\pi^*) \leq \gamma \leq 2$. In this case, we simply run the above algorithm except we set $\bar{\varepsilon} := \gamma\varepsilon/c_0$ for a parameter $c_0 > 0$ to be chosen later and set $\Delta := \gamma$.

Then $\tilde{\mathcal{V}}$ contains $(\gamma + 2\bar{\varepsilon})$ -close points and $|\tilde{\mathcal{V}}| = O(n\Delta^2/\bar{\varepsilon}^2) = O(n\gamma^2/(\gamma\varepsilon)^2) = O(n/\varepsilon^2)$. Following the same argument as in the proof of Lemma 7.1, we claim that c_0 can be chosen so that there exists a $(\Delta + 2 + 4\bar{\varepsilon})$ -tame plan π' with $\mathfrak{c}(\pi') \leq (1 + \varepsilon)\mathfrak{c}(\pi^*)$ and all parking places of π' are in $\tilde{\mathcal{V}}$.

To prove the claim, note that π^* is trivially $(\Delta = \gamma)$ -tame since $\mathfrak{c}(\pi^*) \leq \gamma$. By Lemma 3.6, we have

$$\alpha(\pi^*) \leq c_2(\mathfrak{c}(\pi^*) + 1) \leq 3c_2$$

for a constant $c_2 > 0$. Then, by Lemma 6.1 with $\bar{\varepsilon}$ as parameter ε , there exists a decoupled, $(\Delta + 2 + 4\bar{\varepsilon})$ -tame, $\bar{\varepsilon}$ -nearly-kissing plan π' with all parking places of π' in $\tilde{\mathcal{V}}$ and

$$\mathfrak{c}(\pi') \leq \mathfrak{c}(\pi^*) + \bar{\varepsilon}\alpha(\pi^*) \leq \mathfrak{c}(\pi^*) + 3c_2\bar{\varepsilon} = \mathfrak{c}(\pi^*) + 3c_2\gamma\varepsilon/c_0 \leq (1 + 24c_2\varepsilon/c_0)\mathfrak{c}(\pi^*),$$

where the last inequality follows by $\mathfrak{c}(\pi^*) \geq \gamma/8$. So we choose $c_0 := 1/(24c_2)$. This proves the claim. The rest of the analysis follows from the previous algorithm, including the runtime analysis, since the algorithm from Lemma 8.2 only takes $O(n \log^2 n)$ additional time.

8 $O(1)$ -Approximate Plans for Close Configurations

In this section describe a procedure that, in $O(n \log^2 n)$ time, either compute an 8-approximation γ of π^* or detects that $\mathfrak{c}(\pi^*) > 1/4$, where π^* is an optimal $(\mathbf{s}, \mathbf{t})$ -plan. First, we introduce some notations.

If all moves of a plan π are xy -monotone, we say π is xy -monotone. For a (piecewise-linear) xy -monotone $(\mathbf{s}, \mathbf{t})$ -plan π , $\mathbf{s}, \mathbf{t} \in \mathbf{F}$, let $\$(\pi)$ be the L_1 -cost of π , i.e., if $\langle u_1, u_2, \dots, u_g \rangle$ (resp., $\langle v_1, v_2, \dots, v_h \rangle$) is the sequence of vertices of π_A (resp., π_B), then

$$\$(\pi) = \sum_{i=1}^{g-1} \|u_i - u_{i+1}\|_1 + \sum_{i=1}^{h-1} \|v_i - v_{i+1}\|_1.$$

Recall that we say a configuration $(a, b) \in \mathbf{F}$ is x -separated if $|x(a) - x(b)| \geq 2$ and is y -separated if $|y(a) - y(b)| \geq 2$. The following lemma states the key property of an optimal plan, which forms the basis of our algorithm.




Figure 10. Illustration of s_A, t_A, s_B, t_B positioned as assumed in Step (III) of the algorithm in Lemma 8.2, where $s_A + \square, s_B + \square$ are solid and $t_A + \square, t_B + \square$ are dashed.

LEMMA 8.1. *Let $\mathbf{s}, \mathbf{t}$ be two configurations such that $\mathfrak{c}(\pi^*(\mathbf{s}, \mathbf{t})) \leq 1/4$. Then there exists an optimal $(\mathbf{s}, \mathbf{t})$ -plan π^* with $\alpha(\pi^*) \leq 4$. Furthermore if both $\mathbf{s}$ and $\mathbf{t}$ are x -separated or both of them are y -separated then $\alpha(\pi^*) \leq 2$.*

Algorithm. Let $\square_A := (s_A + (1/4)\square) \cap (t_A + (1/4)\square)$ and $\square_B := (s_B + (1/4)\square) \cap (t_B + (1/4)\square)$. The algorithm searches for a $(\mathbf{s}, \mathbf{t})$ -plan $\pi = (\pi_A, \pi_B)$ contained in $\square_A \times \square_B$ with $\alpha(\pi) \leq 4$ and minimum L_1 -cost. As we will prove, the search only needs to be successful at finding such a plan when $\mathfrak{c}(\pi^*) \leq 1/4$, so the algorithm is described assuming that is true. There are three main steps.

Step (I). We first do a simple check. Let C_A (resp., C_B) be the component of $\square_A \cap \mathcal{F}$ (resp., $\square_B \cap \mathcal{F}$) containing s_A (resp., s_B). If $t_A \notin C_A$ (resp., $t_B \notin C_B$) then s_A, t_A (resp., s_B, t_B) lie in different components of $\mathcal{F} \cap \square_A$ (resp., $\mathcal{F} \cap \square_B$) and we report that $\mathfrak{c}(\pi^*) > 1/4$. Otherwise, we proceed to Step (II).

Step (II). Now let $C'_A \subseteq C_A$ (resp., $C'_B \subseteq C_B$) be the component of $C_A \cap \mathcal{F}[s_B]$ (resp., $C_B \cap \mathcal{F}[t_A]$) containing s_A (resp., t_B). It is possible that $C_A = C'_A$ or $C_B = C'_B$. We next check if there exists a plan with at most two moves in which A moves first: We first check if $t_A \in C'_A$ and $s_B \in C'_B$. If so, there exists an xy -monotone path π_A from s_A to t_A in C'_A , i.e., while B is parked at s_B , and an xy -monotone path from s_B to t_B in C'_B , i.e., while A is parked at t_A , by Lemma 3.4. Then we report the cost $\mathfrak{c}(\pi)$ of the corresponding xy -monotone plan π . Otherwise, we check if there is a plan with at most two moves in which B moves first in a similar fashion, then report its L_2 -cost if yes. If no two-move plan is found, we proceed to the next step, Step (III).

We will later prove that if $\mathfrak{c}(\pi^*) \leq 1/4$ and $\mathbf{s}, \mathbf{t}$ are both x -separated or both y -separated, then Step (II) must find and report a plan π . Hence $\mathbf{s}$ is only x -separated and $\mathbf{t}$ is only y -separated, or vice-versa. So, we continue our search for a plan π assuming, without loss of generality, that $\mathbf{s}$ is only x -separated and $\mathbf{t}$ is only y -separated in Step (III).

Step (III). For concreteness, assume that

$$(8.1) \quad x(s_B) \leq x(s_A) - 2 \text{ and } y(t_B) \leq y(t_A) - 2.$$

Under the assumption $\mathfrak{c}(\pi^*) \leq 1/4$ it can be shown that

$$(8.2) \quad y(s_A) - 2 < y(s_B) \leq y(s_A) - 7/4 \text{ and } x(t_A) - 2 < x(t_B) \leq x(t_A) - 7/4.$$

See Figure 10. For any configuration $\mathbf{p} = (p_A, p_B) \in C_A \times C_B$ which is both x -separated and y -separated, let $\Pi(\mathbf{p})$ be the decoupled $(\mathbf{s}, \mathbf{t})$ -plan which is the concatenation of the decoupled, optimal $(\mathbf{s}, \mathbf{p})$ -plan with at most two moves and the decoupled, optimal $(\mathbf{p}, \mathbf{t})$ -plan with at most two moves, each implied by Lemma 3.5. We define a set $\mathbf{P}$ of $O(n)$ candidate free configurations in $C'_A \times C'_B \subset \mathbf{F}$. Among the candidate pairs of $\mathbf{P}$, we return the L_2 -cost of the plan $\Pi(\mathbf{p})$ which minimizes its L_1 -cost, $\$(\Pi(\mathbf{p}))$. If $\mathbf{P} = \emptyset$ we conclude that $\mathfrak{c}(\pi^*) > 1/4$. Assuming (8.1) holds, it is convenient to compute the position of the bottom-left (resp., top-right) corner vertex of A (resp., B) at the candidate parked positions and then use them to compute positions of the centers of A and B . Let $\mathbf{1} = (1, 1)$. Define $\tilde{C}'_A := C'_A - \mathbf{1}$ and $\tilde{C}'_B := C'_B + \mathbf{1}$, i.e., $\tilde{C}'_A - \mathbf{1}$ (resp., $\tilde{C}'_B + \mathbf{1}$) is the set of positions of the bottom-left (resp., top-right) corner of A (resp., B) while it is placed in C'_A (resp., C'_B). Let $\tilde{\mathcal{V}}$ be the vertices of

the arrangement $\mathcal{A}(\{\tilde{C}'_A, \tilde{C}'_B\})$, i.e., the vertices of the overlay of the two polygons. Let $\tilde{\mathcal{L}}_X$ be the set of all points on the vertical lines through vertices of $\tilde{\mathcal{V}} \cup \{s_A - 1, t_A - 1, s_B + 1, t_B + 1\}$ and let $\tilde{\mathcal{L}}_Y$ be the set of all points on the horizontal lines through $\{s_A - 1, t_A - 1, s_B + 1, t_B + 1\}$. Finally, let $\tilde{P}$ be the vertices of the arrangement $\tilde{\mathcal{A}} := \mathcal{A}(\{\tilde{C}'_A, \tilde{C}'_B, \tilde{\mathcal{L}}_X, \tilde{\mathcal{L}}_Y\})$, i.e., the vertices of the overlay of the two polygons and specified axis-parallel lines. $\tilde{P}$ is the set of positions for the bottom-left and top-right corners of A and B , respectively, at the candidate parking places. Recall, we are interested in parking places at which the two corners lie on the same vertical line and B lies below A . Hence, we define the set $\mathbf{P}$ of parking candidates as

$$\mathbf{P} := \{(\tilde{p}_A + 1, \tilde{p}_B - 1) \mid (\tilde{p}_A, \tilde{p}_B) \in (\tilde{P} \cap \tilde{C}'_A) \times (\tilde{P} \cap \tilde{C}'_B), x(\tilde{p}_A) = x(\tilde{p}_B), y(\tilde{p}_A) \geq y(\tilde{p}_B)\}.$$

This concludes the description of the algorithm.

Correctness. It is easy to verify that if the algorithm succeeds to find a plan π and reports its L_2 -cost $\mathfrak{c}(\pi)$ in Step (II) or Step (III) that $\pi \subset \square_A \times \square_B$, $\alpha(\pi) \leq 4$, and π is feasible. If the algorithm reports $\mathfrak{c}(\pi^*) > 1/4$ in Step (I), then s_A, t_A (resp., s_B, t_B) lie in different components of C_A (resp., C_B) and hence the path π_A (resp., π_B) in any feasible (s, t) -plan (π_A, π_B) must exit $\square_A$ (resp., $\square_B$). So the algorithm behaves correctly in this case. If the algorithm reports a plan π in Steps (II) or (III), all parking places of A (resp., B) are contained in $\square_A$ (resp., $\square_B$) and hence the L_2 -cost of each (xy) -monotone move is at most $1/2$. It follows that $\mathfrak{c}(\pi) \leq \frac{1}{2}\alpha(\pi) \leq 2$.

First suppose $\mathfrak{c}(\pi^*) > 1/4$. If the algorithm fails in both Step (II) and Step (III) to find any plan and report its cost, it correctly reports $\mathfrak{c}(\pi^*) > 1/4$. Otherwise, the algorithm reports the cost $\mathfrak{c}(\pi)$ of a plan π , where $\mathfrak{c}(\pi) \leq 2$ by the discussion above. Then

$$\mathfrak{c}(\pi^*) \leq \mathfrak{c}(\pi) \leq 2 \leq 8\mathfrak{c}(\pi^*).$$

In either case, the algorithm behaves as claimed.

Next, suppose $\mathfrak{c}(\pi^*) \leq 1/4$. We now prove Lemma 8.1.

Proof. Let $(\pi_A^*, \pi_B^*) = \pi^*$. Since π_A^*, π_B^* are continuous, there is a time instance $\lambda \in (0, 1)$ such that $(p_A, p_B) = \pi^*(\lambda)$ is both x -separated and y -separated, in particular, $|x(p_A) - x(p_B)| = 2$. Then $p_A \in \square_A$, $p_B \in \square_B$, and $x(p_B) = x(p_A) - 2$ since $\mathfrak{c}(\pi^*) \leq 1/4$ and $x(s_B) < x(s_A) - 2$. By Lemma 3.4, there exists an xy -monotone optimal (s, q) -plan π_0 with at most two moves, since s, q are both x -separated, and there exists an xy -monotone optimal (q, t) -plan π_1 with at most two moves, since q, t are both y -separated. Then $\langle \pi_0 \rangle \circ \langle \pi_1 \rangle$ is an xy -monotone optimal (s, t) -plan, which has at most four moves.

Next, suppose s, t are both, say, x -separated. Then Lemma 3.5 implies there exists an (optimal xy -monotone) (s, t) -plan with at most two moves. $\square$

In view of Lemma 8.1, if Step (II) fails, then s is only x -separated and t is only y -separated, or vice-versa. Henceforth, we assume s, t are oriented as assumed in the algorithm, i.e., s is only x -separated with s_A right of s_B , t is only y -separated with t_A above t_B , and (8.1) and (8.2) are satisfied.

To finish the proof, we prove that Step (III) succeeds to find a plan $\pi \subset \square_A \times \square_B$, under the assumption that $\mathfrak{c}(\pi^*) \leq 1/4$, with $\mathfrak{c}(\pi) \leq 8\mathfrak{c}(\pi^*)$. Assume, without loss of generality, that A moves first in π^* . The proof of Lemma 8.1 implies that $\langle \pi^* \rangle = (A, \pi_1, s_B), (B, \pi_2, q_A), (A, \pi_3, q_B), (B, \pi_4, t_A)$ where $\pi_1 \subset C'_A$ and $\pi_4 \subset C'_B$. Hence $q_A \in C'_A$ and $q_B \in C'_B$. Let $\tilde{q}_A := q_A - 1$ (resp., $\tilde{q}_B := q_B + 1$) be the position of the bottom-left (resp., top-right) corner of A (resp., B) and let $\tilde{g}_A$ (resp., $\tilde{g}_B$) be the cell of $\tilde{\mathcal{A}}$ containing $\tilde{q}_A$ (resp., $\tilde{q}_B$). Since $q_A \in C'_A$ (resp., $q_B \in C'_B$), we have $\tilde{g}_A \subseteq \tilde{C}'_A$ (resp., $\tilde{g}_B \subseteq \tilde{C}'_B$). See that $x(\tilde{q}_B) = x(\tilde{q}_A)$ and $y(\tilde{q}_B) \leq y(\tilde{q}_A)$ since $x(q_B) = x(q_A) - 2$ and $y(q_B) \leq y(q_A) - 2$. That is, $\tilde{q}_A, \tilde{q}_B$ lie on the same vertical line with $\tilde{q}_A$ above $\tilde{q}_B$. For a configuration $(\tilde{p}_A, \tilde{p}_B) \in \times(\tilde{g}_A \times \tilde{g}_B) \cap \mathbf{P}$, the L_1 -cost $\$(\Pi(\tilde{p}_A, \tilde{p}_B))$ is

$$\begin{aligned} \$(\Pi(\tilde{p}_A + 1, \tilde{p}_B - 1)) &= |x(s_A) - (x(\tilde{p}_A) + 1)| + |(x(\tilde{p}_A) + 1) - x(t_A)| + \\ &\quad |y(s_A) - (y(\tilde{p}_A) + 1)| + |(y(\tilde{p}_A) + 1) - y(t_A)| + \\ &\quad |x(s_B) - (x(\tilde{p}_B) + 1)| + |(x(\tilde{p}_B) + 1) - x(t_B)| + \\ &\quad |y(s_B) - (y(\tilde{p}_B) + 1)| + |(y(\tilde{p}_B) + 1) - y(t_B)|. \end{aligned}$$

By construction of $\tilde{\mathcal{A}}$, the cells of $\tilde{\mathcal{A}}$ contained in $\tilde{C}'_A$ or $\tilde{C}'_B$ are trapezoids with two vertical edges (which may be points); in particular, $\tilde{g}_A, \tilde{g}_B$ are such cells. Therefore, by construction, the vertical edges of $\tilde{g}_A$ and $\tilde{g}_B$ lie on

the same lines of $\tilde{\mathcal{L}}_X$. Moreover, for any point $r \in \{s_A - \mathbf{1}, t_A - \mathbf{1}, s_B + \mathbf{1}, t_B + \mathbf{1}\}$, any cell $\tilde{g}$ of $\tilde{\mathcal{A}}$ is contained in a quadrant of $\mathbb{R}^2$ with r as the origin, due to the inclusion of $\tilde{\mathcal{L}}_A, \tilde{\mathcal{L}}_B$ in the definition of $\tilde{\mathcal{A}}$. In particular, the signs of $x(p) - x(r)$ are the same for all points $p \in \tilde{g}$ and the signs of $y(p) - y(r)$ are the same for all points $p \in \tilde{g}$. It follows that the L_1 -cost $\$(\Pi(\tilde{p}_A + \mathbf{1}, \tilde{p}_B - \mathbf{1}))$ over the configurations $(\tilde{p}_A, \tilde{p}_B) \in (\tilde{g}_A \times \tilde{g}_B) \cap \mathbf{P}$ is a linear function in the coordinates of the points in the configurations; i.e., it can be written as

$$\$(\Pi(\tilde{p}_A + \mathbf{1}, \tilde{p}_B - \mathbf{1})) = a_0 + a_1x(\tilde{p}_A) + a_2y(\tilde{p}_A) + a_3x(\tilde{p}_B) + a_4x(\tilde{p}_B),$$

where $a_0, \dots, a_4 \in \mathbb{R}$. Since $\tilde{g}_A, \tilde{g}_B$ are trapezoids whose left and right edges are vertical and bounded by the same lines in $\tilde{\mathcal{L}}_X$, it can be shown that there exists a configuration $(\tilde{p}_A, \tilde{p}_B) \in \tilde{g}_A \times \tilde{g}_B$ where $\tilde{p}_A$ (resp., $\tilde{p}_B$) is a vertex of $\tilde{g}_A$ (resp. $\tilde{g}_B$) such that

$$\$(\Pi(\tilde{p}_A + \mathbf{1}, \tilde{p}_B - \mathbf{1})) \leq \$(\Pi(\tilde{q}_A + \mathbf{1}, \tilde{q}_B - \mathbf{1})) = \$(\pi^*)$$

and the constraints $x(\tilde{p}_A) = x(\tilde{p}_B)$ and $y(\tilde{p}_A) \geq y(\tilde{p}_B)$ hold. That is, there exists $\mathbf{p} := (\tilde{p}_A + \mathbf{1}, \tilde{p}_B - \mathbf{1}) \in \mathbf{P}$ such that $\$(\Pi(\mathbf{p})) \leq \(π^*) . Then the plan π whose cost is reported is such that $\mathfrak{c}(\pi) \leq \$(\pi) \leq \$(\Pi(\mathbf{p})) \leq \$(\pi^*) \leq \sqrt{2}\mathfrak{c}(\pi^*)$, which completes the proof.

Runtime analysis. We first compute the components C_A, C'_A, C_B, C'_B in $O(n \log^2 n)$ time [13]. Then Steps (I) and (II) take $O(n)$ time. Consider Step (III). Since C'_A and C'_B are xy -monotone by Lemma 3.4, any segment in $\mathbb{R}^2$ intersects each of $\tilde{C}'_A, \tilde{C}'_B$ $O(1)$ times. Then $|\tilde{\mathcal{V}}| = O(n)$ and $\tilde{\mathcal{V}}$ is obtained in $O(n \log n)$ time by computing the arrangement $\mathcal{A}(\{\tilde{C}'_A, \tilde{C}'_B\})$. Furthermore, the axis-parallel lines in $\tilde{\mathcal{L}}_X \cup \tilde{\mathcal{L}}_Y$ each intersect $O(1)$ segments of $\tilde{C}'_A, \tilde{C}'_B$. $\tilde{\mathcal{L}}_Y$ consists of four horizontal lines. Hence there are $O(1)$ vertices in $\tilde{P}$ that lie on any vertical line. So $|\tilde{P}| = O(n)$, which we obtain by computing the arrangement $\tilde{\mathcal{A}} = \mathcal{A}(\{\tilde{C}'_A, \tilde{C}'_B, \tilde{\mathcal{L}}_X, \tilde{\mathcal{L}}_Y\})$ in $O(n \log n)$ time. As we compute $\tilde{\mathcal{A}}$, we mark each of its cells $\tilde{g}$ with each of $\tilde{C}'_A, \tilde{C}'_B$ (possibly both) that contain $\tilde{g}$. By the constraints $x(p_A) = x(p_B)$ and $y(p_A) \geq y(p_B)$ in the definition of $\mathbf{P}$, we have $|\mathbf{P}| = O(n)$, and hence $\mathbf{P}$ is computed in $O(n)$ time.

If $\mathbf{P} = \emptyset$ we report $\mathfrak{c}(\pi^*) > 1/4$, otherwise we find the configuration $\mathbf{P} \in \mathbf{P}$ for which $\$(\Pi(\mathbf{p}))$ is minimized in $O(1)$ time per configuration. Then we compute and report the L_2 -cost $\mathfrak{c}(\Pi(\mathbf{p}))$ in $O(n \log n)$ time [20]. Overall, the algorithm takes $O(n \log^2 n)$ time. Putting everything together, we conclude with the following lemma.

LEMMA 8.2. *There is an algorithm that, given a polygonal environment $\mathcal{W}$ with n vertices, two robots A, B , each modeled as a unit square, and $\mathbf{s} = (s_A, s_B), \mathbf{t} = (t_A, t_B) \in \mathbf{F}$, reports a value $\gamma \leq 2$ with $\mathfrak{c}(\pi^*) \leq \gamma \leq 8\mathfrak{c}(\pi^*)$ or reports that $\mathfrak{c}(\pi^*) > 1/4$, where π^* is an optimal $(\mathbf{s}, \mathbf{t})$ -plan; when both such a value γ exists and $\mathfrak{c}(\pi^*) > 1/4$, it reports either outcome arbitrarily. Its runtime is $O(n \log^2 n)$.*

9 Conclusion

We have described a $(1 + \varepsilon)$ -approximation algorithm for the min-sum motion planning problem for two congruent square robots in a planar polygonal environment with running time $n^2 \varepsilon^{-O(1)} \log n$, i.e., our algorithm is an FPTAS. We also describe an $O(n \log^2 n)$ -time 8-approximation algorithm for the problem when the cost of the optimal plan is less than $1/4$, which is used as a subroutine in our FPTAS. We conclude with some questions for future work. Can our techniques be extended

- (i) to obtain a $(1 + \varepsilon)$ -approximation algorithm for min-sum motion planning for $k > 2$ robots with running time $(n/\varepsilon)^{O(k)}$?
- (ii) to work for translating robots with congruent shapes other than squares, such as other centrally-symmetric regular polygons, disks, or convex polygons?
- (iii) to optimize both *clearance* and the total lengths of the paths in some fashion, where clearance is the minimum distance from any robot to any other robot or obstacle during the plan?

Acknowledgement. We thank Vera Sacristán and Rodrigo Silveira for helpful discussions.

References

- [1] P. K. Agarwal, K. Fox, and O. Salzman. An efficient algorithm for computing high-quality paths amid polygonal obstacles. *ACM Trans. Alg.*, 14(4):46:1–46:21, 2018.
- [2] P. K. Agarwal, T. Geft, D. Halperin, and E. Taylor. Multi-robot motion planning for unit discs with revolving areas. *Comput. Geom.*, 114:102019, 2023.
- [3] P. K. Agarwal, D. Halperin, M. Sharir, and A. Steiger. Near-optimal min-sum motion planning for two square robots in a polygonal environment. *CoRR*, abs/2310.20615, 2023.
- [4] B. Aronov, M. de Berg, A. F. van der Stappen, P. Svestka, and J. Vleugels. Motion planning for multiple robots. *Discret. Comput. Geom.*, 22(4):505–525, 1999.
- [5] T. Asano, D. G. Kirkpatrick, and C.-K. Yap. d_1 -optimal motion for a rod (extended abstract). In *Proc. 12th Annual Sympos. Comput. Geom.*, pages 252–263. ACM, 1996.
- [6] T. Asano, D. G. Kirkpatrick, and C.-K. Yap. Minimizing the trace length of a rod endpoint in the presence of polygonal obstacles is NP-hard. In *Proc. 15th Canadian Conf. Comput. Geom.*, pages 10–13, 2003.
- [7] T. Brocken, G. W. van der Heijden, I. Kostitsyna, L. E. Lo-Wong, and R. J. A. Surtel. Multi-robot motion planning of k -colored discs is PSPACE-hard. In *FUN*, volume 157 of *LIPIcs*, pages 15:1–15:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- [8] J. Brunner, L. Chung, E. D. Demaine, D. H. Hendrickson, A. Hesterberg, A. Suhl, and A. Zeff. 1 X 1 rush hour with fixed blocks is PSPACE-complete. In *10th International Conference on Fun with Algorithms*, volume 157, pages 7:1–7:14, 2021.
- [9] J. Canny and J. Reif. New lower bound techniques for robot motion planning problems. In *Proceedings IEEE Symposium on Foundations of Computer Science*, pages 49–60, 1987.
- [10] J. F. Canny. *The Complexity of Robot Motion Planning*. MIT Press, Cambridge, MA, 1988.
- [11] D. Z. Chen and H. Wang. Computing shortest paths among curved obstacles in the plane. *ACM Trans. Alg.*, 11(4):1–46, 2015.
- [12] D. Dayan, K. Solovey, M. Pavone, and D. Halperin. Near-optimal multi-robot motion planning with finite sampling. *IEEE Trans. Robotics*, 39(5):3422–3436, 2023.
- [13] M. de Berg, O. Cheong, M. J. van Kreveld, and M. H. Overmars. *Computational Geometry: Algorithms and Applications, 3rd Edition*. Springer, 2008.
- [14] E. D. Demaine, S. P. Fekete, P. Keldenich, H. Meijer, and C. Scheffer. Coordinated motion planning: Reconfiguring a swarm of labeled robots with bounded stretch. *SIAM J. Comput.*, 48(6):1727–1762, 2019.
- [15] G. Esteban, D. Halperin, V. Ruíz, V. Sacristán, and R. I. Silveira. Shortest coordinated motions for square robots. In *Algorithms and Data Structures - 18th International Symposium WADS*, volume 14079 of *Lecture Notes in Computer Science*, pages 430–443. Springer, 2023.
- [16] S. Fortune, G. T. Wilfong, and C.-K. Yap. Coordinated motion of two robot arms. In *Proc. 1986 Int. Conf. Robotics and Automation*, pages 1216–1223. IEEE, 1986.
- [17] D. Halperin, L. Kavraki, and K. Solovey. Robotics. In J. E. Goodman, J. O’Rourke, and C. Tóth, editors, *Handbook of Discrete and Computational Geometry*, chapter 51, pages 1343–1376. Chapman & Hall/CRC, 3rd edition, 2018.
- [18] D. Halperin, M. Sharir, and O. Salzman. Algorithmic motion planning. In J. E. Goodman, J. O’Rourke, and C. Tóth, editors, *Handbook of Discrete and Computational Geometry*, chapter 50, pages 1311–1342. Chapman & Hall/CRC, 3rd edition, 2018.
- [19] R. A. Hearn and E. D. Demaine. PSPACE-completeness of sliding-block puzzles and other problems through the nondeterministic constraint logic model of computation. *Theor. Comput. Sci.*, 343(1-2):72–96, 2005.
- [20] J. Hershberger and S. Suri. An optimal algorithm for Euclidean shortest paths in the plane. *SIAM J. Comput.*, 28(6):2215–2256, 1999.
- [21] J. E. Hopcroft, J. T. Schwartz, and M. Sharir. On the complexity of motion planning for multiple independent objects; PSPACE-hardness of the “warehouseman’s problem”. *Int. J. Robotics Res.*, 3(4):76–88, 1984.
- [22] J. E. Hopcroft and G. T. Wilfong. Reducing multiple object motion planning to graph searching. *SIAM J. Comput.*, 15(3):768–785, 1986.
- [23] S. Karaman and E. Frazzoli. Sampling-based algorithms for optimal motion planning. *Int. J. Robotics Res.*, 30(7):846–894, 2011.
- [24] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Trans. Robotics & Automation*, 12(4):566–580, June 1996.
- [25] D. G. Kirkpatrick and P. Liu. Characterizing minimum-length coordinated motions for two discs. In *Proc. 28th Canadian Conf. Comput. Geom.*, pages 252–259, 2016.
- [26] J.-C. Latombe. *Robot Motion Planning*. Kluwer, Boston, MA, 1991.
- [27] S. M. LaValle. *Planning Algorithms*. Cambridge University Press, Cambridge, U.K., 2006. Available at <http://planning.cs.uiuc.edu/>.
- [28] S. M. LaValle and J. J. Kuffner. Rapidly-exploring random trees: Progress and prospects. In *Proceedings Workshop on the Algorithmic Foundations of Robotics*, 2000.

- [29] J. S. Mitchell. Shortest paths and networks. In J. E. Goodman, J. O'Rourke, and C. Tóth, editors, *Handbook of Discrete and Computational Geometry*, chapter 51, pages 811–848. Chapman & Hall/CRC, 3rd edition, 2018.
- [30] C. Ó'Dúnlaing and C.-K. Yap. A retraction method for planning the motion of a disc. *J. Algorithms*, 6:104–111, 1985.
- [31] V. Ruíz Herrero. Optimal coordinated motions for two square robots. Master's thesis, Universitat Politècnica de Catalunya, 2019.
- [32] O. Salzman. Sampling-based robot motion planning. *Commun. ACM*, 62(10):54–63, 2019.
- [33] J. T. Schwartz and M. Sharir. On the piano movers' problem: III. Coordinating the motion of several independent bodies: The special case of circular bodies moving amidst polygonal barriers. *Int. J. Robotics Res.*, 2(3):46–75, 1983.
- [34] J. Sellen, J. Choi, and C.-K. Yap. Precision-sensitive Euclidean shortest path in 3-space. *SIAM J. Comput.*, 29(5):1577–1595, 2000.
- [35] M. Sharir and S. Sifrony. Coordinated motion planning for two independent robots. *Ann. Math. Artif. Intell.*, 3(1):107–130, 1991.
- [36] R. Shome, K. Solovey, A. Dobson, D. Halperin, and K. E. Bekris. dRRT^{*}: Scalable and informed asymptotically-optimal multi-robot motion planning. *Auton. Robots*, 44(3-4):443–467, 2020.
- [37] I. Solomon and D. Halperin. Motion planning for multiple unit-ball robots in $\mathbb{R}^d$. In *Workshop on the Algorithmic Foundations of Robotics, WAFR*, pages 799–816, 2018.
- [38] K. Solovey and D. Halperin. On the hardness of unlabeled multi-robot motion planning. *Int. J. Robotics Res.*, 35(14):1750–1759, 2016.
- [39] K. Solovey, J. Yu, O. Zamir, and D. Halperin. Motion planning for unlabeled discs with optimality guarantees. In *Robotics: Science and Systems*, 2015.
- [40] P. G. Spirakis and C.-K. Yap. Strong NP-hardness of moving many discs. *Inf. Process. Lett.*, 19(1):55–59, 1984.
- [41] R. Stern, N. R. Sturtevant, A. Felner, S. Koenig, H. Ma, T. T. Walker, J. Li, D. Atzmon, L. Cohen, T. K. S. Kumar, R. Barták, and E. Boyarski. Multi-agent pathfinding: Definitions, variants, and benchmarks. In *Proc. 12th Int. Sympos. Combinatorial Search*, pages 151–159, 2019.
- [42] M. Turpin, K. Mohta, N. Michael, and V. Kumar. Goal assignment and trajectory planning for large teams of interchangeable robots. *Auton. Robots*, 37(4):401–415, 2014.
- [43] H. Wang. A new algorithm for Euclidean shortest paths in the plane. In *Proc. 53rd Annual Sympos. Theory of Comput.*, pages 975–988. ACM, 2021.
- [44] R. Wein, J. P. van den Berg, and D. Halperin. The visibility-Voronoi complex and its applications. *Comput. Geom.*, 36(1):66–87, 2007.
- [45] R. Wein, J. P. van den Berg, and D. Halperin. Planning high-quality paths and corridors amidst obstacles. *I. J. Robotics Res.*, 27(11-12):1213–1231, 2008.
- [46] J. Yu and S. M. LaValle. Structure and intractability of optimal multi-robot path planning on graphs. In *Proc. 27th Conf. Artificial Intelligence*, 2013.