# Efficient Multi-Robot Motion Planning for Unlabeled Discs in Simple Polygons 

Aviv Adler, Mark de Berg, Dan Halperin, and Kiril Solovey


#### Abstract

We consider the following motion-planning problem: we are given $m$ unit discs in a simple polygon with $n$ vertices, each at their own start position, and we want to move the dises to a given set of $m$ target positions. Contrary to the standard (labeled) version of the problem, each disc is allowed to be moved to any target position, as long as in the end every target position is occupied. We show that this unlabeled version of the problem can be solved in $O\left(n \log n+m n+m^{2}\right)$ time, assuming that the start and target positions are at least some minimal distance from each other. This is in sharp contrast to the standard (labeled) and more general multi-robot motion planning problem for dises moving in a simple polygon, which is known to be strongly NPhard.


## I. Introduction

The multi-robot motion-planning problem is to plan the motions of several robots operating in a common workspace. In its most basic form, the goal is to move each robot from its start position to some designated target position, while avoiding collision with obstacles in the environment and with other robots. Besides its obvious relevance to robotics, the problem has various other applications, for example in the design of computer games or crowd simulation. Multi-robot motion planning is a natural extension of the single-robot motion planning problem, but it is much more complex due to the high number of degrees of freedom that it entails, even when the individual robots are as simple as discs.

## A. Related work

One of the first occurrences of the multi-robot motionplanning problem in the computational-geometry literature can be found in the series of papers on the Piano Movers' Problem by Schwartz and Sharir. They first considered the problem in a general setting [1] and then narrowed it down to the case of disc robots moving amidst polygonal obstacles [2]. In the latter work an algorithm was presented for the case of two and three robots, with running time of $O\left(n^{3}\right)$ and $O\left(n^{13}\right)$, respectively, where $n$ is the complexity of the workspace.

[^0]Later Yap [3] used the retraction method to develop a more efficient algorithm, which runs in $O\left(n^{2}\right)$ and $O\left(n^{3}\right)$ time for the case of two and three robots, respectively. Several years afterwards, Sharir and Sifrony [4] presented a general approach based on cell decomposition, which is capable of dealing with various types of robot pairs and which has a running time of $O\left(n^{2}\right)$. Moreover, several techniques that reduce the effective number of degrees of freedom of the problem have been proposed [5, 6].

When the number of robots is no longer a fixed constant, the multi-robot motion-planning problem becomes hard. Hopcroft et al. [7] showed that even in the relatively simple setting of $n$ rectangular robots moving in a rectangular workspace, the problem is already PSPACE-hard. Moreover, Spirakis and Yap [8] showed that the problem is strongly NPhard for disc robots in a simple polygon.

In recent years, multi-robot planning has attracted a great deal of attention from the robotics community. This can be mainly attributed to two reasons. First, it is a problem of practical importance. Second, the emergence of the samplingbased techniques, which are relatively easy to implement, yet are highly effective. These techniques attempt to capture the connectivity of the configuration space through random sampling [9, 10]. Although sampling-based algorithms are usually incomplete-they are not guaranteed to find a solution-they tend to be very efficient in practice. Hence, they are considered the method-of-choice for motion-planning problems that involve many degrees of freedom. While sampling-based tools for a single robot can be applied directly to the multirobot problem by considering the group of robots as one large composite robot [11], there is a large body of work that attempts to exploit the unique properties of the multi-robot problem [12, 13, 14, 15, 16, 17].

The aforementioned results deal exclusively with the classical formulation of the multi-robot problem, where the robots are distinct and every robot is assigned a specific target position. The unlabeled variant of the problem, where all the robots are assumed to be identical and thus interchangeable, was first considered by Kloder and Hutchinson [18], who devised a sampling-based algorithm for the problem. Recently a generalization of the unlabeled problem-the $k$-color motionplanning problem-has been proposed, in which there are several groups of interchangeable robots [19]. Turpin et al. [20] considered a special setting of the unlabeled problem with disc robots, namely where the collection of free configurations surrounding every start or target position is star-shaped. This condition allows them to devise an efficient algorithm that computes a solution in which the maximum path length
is minimized. Unfortunately the star-shapedness condition is quite restrictive, and in general it will not be satisfied.

Other related work includes papers that study the number of moves required to move a set of discs between two sets of positions in an unbounded workspace, when a move consists of sliding a single disc-see for example the paper by Bereg et al. [21] which provides upper and lower bounds for the unlabeled case, or the paper by Dumitrescu and Jiang [22] who show that deciding whether a collection of labeled or unlabeled discs can be moved between two sets of positions within $k$ steps is NP-hard. Finally, we mention the problem of pebble motions on graphs, which can be considered as a discrete variant of the multi-robot motion planning problem. In this problem, pebbles need to be moved from one set of vertices of a graph to another, while following a certain set of rules-see for example [23, 24, 25, 26, 27, 28].

## B. Our contribution

Surprisingly, the unlabeled version of the multi-robot motion-planning problem has hardly received any attention in the computational-geometry literature. Indeed, we don't know of any papers that solve the problem in an exact and complete manner, except in a restricted setting studied by Turpin et al. [20] that we mentioned above. We therefore study the following basic variant of the problem: given $m$ unit discs in a simple polygon with $n$ vertices, each at their own start position, and $m$ target positions, find collision-free motions for the discs such that at the end of the motions each disc occupies a target position. We make the additional assumption that the given start and target positions are well-separated. More precisely, any two of the given start and target positions should be at distance at least 4 from each other. Notice that we only assume this extra separation between the robots in their static initial and goal placements; we do not assume any extra separation (beyond non-collision) between a robot and the obstacles, nor do we enforce any extra separation between the robots during the motion. Even this basic version of the problem turns out to have a rich structure and poses several difficulties and interesting questions.

By carefully examining the various properties of the problem we show how to transform it into a discrete pebblemotion problem on graphs. A solution to the pebble problem, which can be generated with rather straightforward techniques, can then be transformed back into a solution to our continuous motion-planning problem. We mention that a similar transformation was used in [19] in the context of a sampling-based method. Using this transformation we are able to devise an efficient algorithm whose running time is $O\left(n \log n+m n+m^{2}\right)$, where $m$ is the number of robots and $n$ is the complexity of the workspace. To be precise, we show that our algorithm runs in $O\left(n \log n+m^{2}\right)$ time, and the overall description length of all the paths to be carried out by the robots has complexity $O\left(m n+m^{2}\right)$. As already mentioned, this is in sharp contrast to the standard (labeled) and more general multi-robot motion planning problem for discs moving in a simple polygon, which is known to be strongly NP-hard [8].

## II. Preliminaries

We consider the problem of $m$ indistinguishable unit-disc robots moving in a simple polygonal workspace $\mathcal{W} \subset \mathbb{R}^{2}$ with $n$ edges. We define $\mathcal{O} \triangleq \mathbb{R}^{2} \backslash \mathcal{W}$ to be the complement of the workspace, and we call $\mathcal{O}$ the obstacle space. Since our robots are discs, a placement of a robot is uniquely specified by the location of its center. Hence, we will sometimes refer to points $x \in \mathcal{W}$ as configurations, and we will say that a robot is at configuration $x$ when its center is placed at the point $x \in W$. For given $x \in \mathbb{R}^{2}$ and $r \in \mathbb{R}_{+}$, we define $\mathcal{D}_{r}(x)$ to be the open disc of radius $r$ centered at $x$.

We consider the unit-disc robots to be open sets. Thus a robot avoids collision with the obstacle space if and only if its center is at distance at least 1 from $\mathcal{O}$, that is, when it is at a configuration located in the free space $\mathcal{F} \triangleq\left\{x \in \mathbb{R}^{2}\right.$ : $\left.\mathcal{D}_{1}(x) \cap \mathcal{O}=\emptyset\right\}$. We require the robots to avoid collisions with each other, so if a robot is at configuration $x$ then no other robot can be at a configuration $y \in \operatorname{Int}\left(\mathcal{D}_{2}(x)\right)$; here $\operatorname{Int}(X)$ denotes the interior of the set $X$. Furthermore, the notation $\partial(X)$ will be used to refer the boundary of $X$. We call $\mathcal{D}_{2}(x)$ the collision disc of the configuration $x$.

Besides the simple polygon $\mathcal{W}$ forming the workspace, we are also given sets $S=\left\{s_{1}, s_{2}, \ldots, s_{m}\right\}$ and $T=$ $\left\{t_{1}, t_{2}, \ldots, t_{m}\right\}$, such that $S, T \subset \mathcal{F}$. These are respectively the sets of start and target configurations of our $m$ identical disc robots. We assume that the configurations in $S$ and $T$ are well-separated:

For any two distinct configurations $x, y \in S \cup T$ we have $\|x-y\| \geqslant 4$.
The problem is now to plan a collision-free motion for our $m$ unit-disc robots such that each of them starts at a configuration in $S$ and ends at a configuration in $T$. Since the robots are indistinguishable (or: unlabeled), it does not matter which robot ends up at which target configuration. Formally, we wish to find paths $\pi_{i}:[0,1] \rightarrow \mathcal{F}$, for $1 \leqslant i \leqslant m$, such that $\pi_{i}(0)=s_{i}$ and $\bigcup_{i=1}^{m} \pi_{i}(1)=T$. Additionally, we require that the robots do not collide with each other: for every $1 \leqslant i \neq j \leqslant m$ and every $\xi \in(0,1)$, we require $\mathcal{D}_{1}\left(\pi_{i}(\xi)\right) \cap \mathcal{D}_{1}\left(\pi_{j}(\xi)\right)=\emptyset$. Note that the requirement that the robots do not collide with the obstacle space $\mathcal{O}$ is implied by the paths $\pi_{i}$ being inside the free space $\mathcal{F}$.

## III. Basic properties of the free space

Recall that the free space $\mathcal{F} \subset \mathcal{W}$ is the set of configurations at which a robot does not collide with the obstacle space. The free space may consist of multiple connected components. We denote these components by $F_{1}, \ldots, F_{q}$, where $q$ is the total number of components. For any $i \in\{1,2, \ldots, q\}$, we let $S_{i} \triangleq S \cap F_{i}$ and $T_{i} \triangleq T \cap F_{i}$. We assume from now on that $\left|S_{i}\right|=\left|T_{i}\right|$ for all $1 \leqslant i \leqslant q$-if this is not the case, then the problem instance obviously has no solution-and we define $m_{i} \triangleq\left|S_{i}\right|=\left|T_{i}\right|$ to be the number of robots in $F_{i}$.

Before we proceed, we need one more piece of notation. For any $x \in \mathcal{W}$, we define $\operatorname{obs}(x)$, the obstacle set of $x$, as $\operatorname{obs}(x) \triangleq\{y \in \mathcal{O}:\|x-y\|<1\}$. In other words, $\operatorname{obs}(x)$ contains the points in the obstacle space overlapping with $\mathcal{D}_{1}(x)$. Note that $\operatorname{obs}(x)=\emptyset$ for $x \in \mathcal{F}$.

In the remainder of this section we prove several crucial properties of the free space, which will allow us to transform our problem to a discrete pebble problem. We start with some properties of individual components $F_{i}$, and then consider the interaction between robots in different components.

## A. Properties of a single connected component of $\mathcal{F}$

We start with a simple observation, for which we provide a proof for completeness.
Lemma 1. Each component $F_{i}$ is simply connected.
Proof: Suppose for a contradiction that $F_{i}$ contains a hole. Then $\operatorname{Compl}(\mathcal{F}) \triangleq \mathbb{R}^{2} \backslash \mathcal{F}$, the complement of the free space, has multiple connected components. One of these is $C_{\mathcal{O}}$, the unbounded component containing $\mathcal{O}$. Let $C$ be another component of $\operatorname{Compl}(\mathcal{F})$, and let $x \in C$. Since $x \notin \mathcal{F}$, there is a point $y \in \mathcal{O}$ with $\|x-y\|<1$. But then $\left\|x^{\prime}-y\right\|<1$ for any point $x^{\prime}$ on the segment $\overline{x y}$, which implies $\overline{x y} \subset \operatorname{Compl}(\mathcal{F})$ and thus contradicts that $C$ and $C_{\mathcal{O}}$ are different components.

Now consider any $x$ in $\mathbb{R}^{2}$. Recall that $\mathcal{D}_{2}(x)$ denotes the collision disc of $x$, that is, $\mathcal{D}_{2}(x)$ is the set of all configurations $y$ for which another robot placed at $y$ collides with a robot at $x$. We now define $D^{*}(x)$ to be the part of $\mathcal{D}_{2}(x)$ that is in the same free-space component as $x$, that is, $D^{*}(x) \triangleq \mathcal{D}_{2}(x) \cap F_{i}$ where $F_{i}$ is the free-space component such that $x \in F_{i}$.

The following three lemmas constitute the theoretical basis on which the correctness and efficiency of our algorithm relies.
Lemma 2. For any $x \in \mathcal{F}$, the set $D^{*}(x)$ is connected.
Proof: Assume for a contradiction that $D^{*}(x)$ is not connected. Let $F_{i}$ be the free-space component containing $x$. Since by definition $x \in D^{*}(x)$, we can find some $y \in D^{*}(x)$ that is in a different connected component of $D^{*}(x)$ from $x$. Since $y \in D^{*}(x) \subset \mathcal{D}_{2}(x)$, the distance between $x$ and $y$ is at most 2 . Hence, any point on the line segment $\overline{x y}$ is within a distance of 1 of either $x$ or $y$. Since $x, y \in F_{i}$, we know that $\overline{x y} \subset \mathcal{W}$, otherwise either $x$ or $y$ would not be in $\mathcal{F}$. We also know that $\overline{x y} \not \subset F_{i}$, since otherwise $x$ and $y$ would not be in different connected components of $D^{*}(x)$. Because $x, y \in F_{i}$, by definition there exists a simple path $\pi \subset F_{i}$ from $x$ to $y$. Since the workspace is a polygon with finite description complexity, we may assume that $\pi$ has finite complexity as well, which implies that $\pi \cap \overline{x y}$ is composed of finitely many isolated points and closed segments. See Figure 1 (a) for an illustration.

We now define $x^{\prime}, y^{\prime}$ as the points on $\pi \cap \overline{x y} \subset D^{*}(x)$ such that $x^{\prime}, y^{\prime}$ are in different connected components of $D^{*}(x)$ and $\left\|x^{\prime}-y^{\prime}\right\|$ is minimized given the first condition. Let $\pi^{\prime}$ be the subpath of $\pi$ joining $x^{\prime}$ to $y^{\prime}$. Notice that $\pi \cap \overline{x^{\prime} y^{\prime}}=\left\{x^{\prime}, y^{\prime}\right\}$. Indeed, if there exists a point $z \in \pi \cap \operatorname{Int}\left(\overline{x^{\prime} y^{\prime}}\right)$, then $z$ must be in a different connected component of $D^{*}(x)$ than either $x^{\prime}$ or $y^{\prime}$, and $\left\|x^{\prime}-y^{\prime}\right\|$ would not be the minimum. Since $\pi$ is a simple path, this means that $\lambda \triangleq \pi^{\prime} \cup \overline{x^{\prime} y^{\prime}}$ is a simple closed curve. The area enclosed by $\lambda$ (including $\lambda$ ) will be referred to as $A$. We note that $\lambda \subset \mathcal{W}$ since $\pi^{\prime} \subset \mathcal{F} \subset \mathcal{W}$ and $\overline{x^{\prime} y^{\prime}} \subset \overline{x y} \subset \mathcal{W}$. This immediately implies that $A \subset \mathcal{W}$, since $\mathcal{W}$ is a simple polygon.

Let $A^{*} \triangleq A \backslash \mathcal{F}$. We claim that $A^{*} \subset \operatorname{Int}\left(\mathcal{D}_{2}(x)\right)$, which implies that there exists a path in $F_{i}$ from $x^{\prime}$ to $y^{\prime}$ that goes along $\partial\left(A^{*}\right)$ and is fully contained in $\mathcal{D}_{2}(x)$. But this contradicts that $x^{\prime}$ and $y^{\prime}$ are in different components of $D^{*}(x)$ and, hence, proves the lemma. It thus remains to prove the claim that $A^{*} \subset \operatorname{Int}\left(\mathcal{D}_{2}(x)\right)$.

Note that for any point $z \in A^{*}$ and any $w \in \operatorname{obs}(z)$ we have $\overline{z w} \cap \pi^{\prime}=\emptyset$, since $\pi^{\prime} \subset \mathcal{F}$. Furthermore, for any $v \in$ $\pi^{\prime}$ we have $\|w-v\| \geqslant 1$, and as $x^{\prime}, y^{\prime} \in \pi^{\prime}$ it follows that $\left\|w-x^{\prime}\right\| \geqslant 1,\left\|w-y^{\prime}\right\| \geqslant 1$. Assume without loss of generality that $\overline{x^{\prime} y^{\prime}}$ is vertical and that locally $A$ lies to the right of $\overline{x^{\prime} y^{\prime}}$, as in Figure 1 (a). Let $K$ be the circle of radius 1 that passes through $x^{\prime}$ and $y^{\prime}$, and whose center lies to the left of $\overline{x^{\prime} y^{\prime}}-$ such a circle always exists since $\left\|x^{\prime}-y^{\prime}\right\| \leqslant\|x-y\| \leqslant 2$. (If $\left\|x^{\prime}-y^{\prime}\right\|=2$ then the center of the circle lies on $\overline{x^{\prime} y^{\prime}}$.) Let $\zeta$ be the arc of this circle lying to the right of $\overline{x^{\prime} y^{\prime}}$; note that this is the shorter of the two arcs joining $x^{\prime}$ and $y^{\prime}$ if they are of different lengths. Then $A^{*}$ is a contained entirely within the area enclosed by $\zeta$ and $\overline{x^{\prime} y^{\prime}}$. Furthermore, $A^{*} \subset$ $\operatorname{Int}(A) \cup \operatorname{Int}\left(\overline{x^{\prime} y^{\prime}}\right)$ since $\pi^{\prime} \subset \mathcal{F}$. Therefore, since $\overline{x^{\prime} y^{\prime}}$ is a subsegment of $\overline{x y}$ and $\zeta$ cannot cross $\partial\left(\mathcal{D}_{2}(x)\right)$, it follows that $A^{*} \subset\left(\operatorname{Int}(A) \cup \operatorname{Int}\left(\overline{x^{\prime} y^{\prime}}\right)\right) \cap \operatorname{Int}\left(\mathcal{D}_{2}(x)\right) \subset \operatorname{Int}\left(\mathcal{D}_{2}(x)\right)$,
which finishes the proof of the lemma.

## B. Interference between different connected components of $\mathcal{F}$.

Let $F_{i}, F_{j}$ be two distinct components of $\mathcal{F}$, and let $x \in F_{i}$ be such that $\mathcal{D}_{2}(x) \cap F_{j} \neq \emptyset$. We then call $x$ an interference configuration from $F_{i}$ to $F_{j}$, and define the interference set from $F_{i}$ to $F_{j}$ as $I_{(i, j)} \triangleq\left\{x \in F_{i}: \mathcal{D}_{2}(x) \cap F_{j} \neq \emptyset\right\}$. We also define the mutual interference set of $F_{i}, F_{j}$ as $I_{\{i, j\}} \triangleq$ $I_{(i, j)} \cup I_{(j, i)}$. Intuitively, an interference configuration from $F_{i}$ to $F_{j}$ is a configuration for a robot in $F_{i}$ which could block a path in $F_{j}$, and the interference set is the set of all such points. The mutual interference set of $F_{i}, F_{j}$ is the set of all single-robot configurations in either component which might block a valid single-robot path in the other component.

Lemma 3. For any mutual interference set $I_{\{i, j\}}$ and any two configurations $x_{1}, x_{2} \in I_{\{i, j\}}$ we have $\mathcal{D}_{2}\left(x_{1}\right) \cap \mathcal{D}_{2}\left(x_{2}\right) \neq \emptyset$.

Proof: The proof is similar in spirit to the proof of Lemma 2 albeit slightly more involved. Assume for a contradiction that $x_{1}, x_{2} \in I_{\{i, j\}}$ and $\mathcal{D}_{2}\left(x_{1}\right) \cap \mathcal{D}_{2}\left(x_{2}\right)=\emptyset$. By definition there exist $y_{1} \in \mathcal{D}_{2}\left(x_{1}\right)$ and $y_{2} \in \mathcal{D}_{2}\left(x_{2}\right)$ such that each pair $\left\{x_{1}, y_{1}\right\},\left\{x_{2}, y_{2}\right\}$ contains one point in $F_{i}$ and one point in $F_{j}$. As shown in the proof for Lemma 2, the segments $\overline{x_{1} y_{1}}, \overline{x_{2} y_{2}}$ are entirely contained in $\mathcal{W}$. We may assume that $\overline{x_{1} y_{1}}$ does not cross $\overline{x_{2} y_{2}}$, since if it did the crossing point would be in $\mathcal{D}_{2}\left(x_{1}\right) \cap \mathcal{D}_{2}\left(x_{2}\right)$ and we would be done. Therefore, there exists a simple closed curve $\lambda \subset \mathcal{W}$ composed of the union of two simple curves $\pi_{i}, \pi_{j}$ and two line segments $\ell_{1}, \ell_{2}$ such that $\pi_{i} \subset F_{i}$ and $\pi_{j} \subset F_{j}$, and $\ell_{1} \subset \overline{x_{1} y_{1}}, \ell_{2} \subset \overline{x_{2} y_{2}}$. Note that both $\ell_{1}$ and $\ell_{2}$ have one endpoint in $F_{i}$ and the other in $F_{j}$; see Figure 1 (b) for an illustration. The end points of $\ell_{1}$ consist of $x_{1}^{\prime}, y_{1}^{\prime}$, such that $x_{1}, x_{1}^{\prime}$ and $y_{1}, y_{1}^{\prime}$ belong to the same connected components, and minimize the distance $\left\|x_{1}^{\prime}-y_{1}^{\prime}\right\|$ ( $\ell_{2}$ is similarly defined).

We refer to the region enclosed by $\lambda$ (including $\lambda$ ) as $A$. Because $\lambda \subset \mathcal{W}$ and $\mathcal{W}$ is a simple polygon, we know that $A \subset \mathcal{W}$. Furthermore, since $\pi_{i}, \pi_{j} \subset \mathcal{F}$, for any $x \in \operatorname{Int}(A)$ and $y \in \operatorname{obs}(x)$ (by definition, $y \in \mathbb{R}^{2} \backslash \mathcal{W}$ so $y \notin A$; thus, $\overline{x y} \cap \lambda \neq \emptyset$ ), we know that $\overline{x y} \cap \pi_{i}=\overline{x y} \cap \pi_{j}=\emptyset$. Thus, $\overline{x y} \cap \operatorname{Int}\left(\ell_{1}\right) \neq \emptyset$ or $\overline{x y} \cap \operatorname{Int}\left(\ell_{2}\right) \neq \emptyset$, or both. Let $A^{*} \triangleq A \backslash \mathcal{F}$ and denote by $A_{1}^{*}$ the set of configurations $x \in A^{*}$ for which there exists $y \in \operatorname{obs}(x)$ such that $\overline{x y} \cap \operatorname{Int}\left(\ell_{1}\right) \neq \emptyset$; the set $A_{2}^{*}$ is defined in a similar manner, only that now $\overline{x y} \cap \operatorname{Int}\left(\ell_{2}\right) \neq \emptyset$. Note that $A^{*}=A_{1}^{*} \cup A_{2}^{*}$.

We claim that $A_{1}^{*} \cap A_{2}^{*} \neq \emptyset$. Indeed, if $A_{1}^{*} \cap A_{2}^{*}=\emptyset$ then there is a path from $x_{1}$ to $y_{1}$ along $\partial\left(A_{1}^{*}\right)$ that stays in $A \backslash A^{*}$ and, hence, stays in $\mathcal{F}$, which would contradict that $x_{1} \in F_{i}$ and $y_{1} \in F_{j}$ for $i \neq j$. Thus, there exists a point $x^{*} \in A_{1}^{*} \cap A_{2}^{*}$. We define the unit circles $K_{1}, K_{2}$ whose boundaries lie on the endpoints of $\ell_{1}, \ell_{2}$ respectively, and whose centers are located

(a)

(b)

Fig. 1. (a) An illustration of Lemma 2. The disc $\mathcal{D}_{2}(x)$ is drawn in green. The closed curve $\lambda$, which consists of the curve $\pi^{\prime}$ and the straight-line $\overline{x^{\prime} y^{\prime}}$, is drawn in blue, and $A$ represents the area that is bounded by $\lambda$. The disc $K$ of radius 1 that touches $x^{\prime}, y^{\prime}$ is drawn in pink. Note that the area $A^{*}$, which is drawn in red, is contained in $A$. The dashed black lines represent $\pi \backslash \pi^{\prime}$. (b) An illustration of Lemma 3 and in particular, the case where $A_{1}^{*} \cap A_{2}^{*} \neq \emptyset$. For simplicity of presentation, we assume that $\ell_{1}=\overline{x_{1} y_{1}}$ and $\ell_{2}=\overline{x_{2} y_{2}}$.
outside $A$. Thus, we have $A_{1}^{*} \subset K_{1}$ and $A_{2}^{*} \subset K_{2}$. Hence, $x^{*} \in K_{1} \cap K_{2}$, which implies $x^{*} \in \mathcal{D}_{2}\left(x_{1}\right) \cap \mathcal{D}_{2}\left(x_{2}\right)$, so $\mathcal{D}_{2}\left(x_{1}\right) \cap \mathcal{D}_{2}\left(x_{2}\right) \neq \emptyset$, contradicting our initial assumption.

The next lemma is a generalization of the previous one. Intuitively, instead of considering a cycle of length 2 among interacting free-space components, we now consider larger cycles.
Lemma 4. Let $\{\phi(1), \phi(2), \ldots, \phi(h)\} \subset\{1,2, \ldots, q\}$, and let $x_{1}, x_{2}, \ldots, x_{h}$ be points such that for all $i, x_{i} \in I_{\{\phi(i), \phi(i+1)\}}$, where $\phi(h+1) \equiv \phi(1)$. (Thus the list is circular with respect to its index). Then there exists some $i \neq j$ such that $\mathcal{D}_{2}\left(x_{i}\right) \cap$ $\mathcal{D}_{2}\left(x_{j}\right) \neq \emptyset$.

Proof: This can be proved in a manner completely analogous to the proof of Lemma 3, we will outline the proof here. We assume for a contradiction that $\mathcal{D}_{2}\left(x_{i}\right) \cap \mathcal{D}_{2}\left(x_{j}\right)=\emptyset$ for all $i \neq j$. We can argue that we can construct a simple closed curve $\lambda \subset \mathcal{W}$ passing through $F_{\phi(1)}, F_{\phi(2)}, \ldots, F_{\phi(h)}$ (in that order), which is composed of simple closed curves $\pi_{i} \subset F_{\phi(i)}$ and line segments $\ell_{i} \subset \mathcal{W}$ with endpoints in $F_{\phi(i)}$ and $F_{\phi(i+1)}$. We then consider the area $A$ enclosed by $\lambda$ and note that $A \subset \mathcal{W}$. Let $A^{*} \triangleq A \backslash \mathcal{F}$. If there exists some simple curve $\pi^{*} \subset A^{*}$ connecting $\ell_{i}$ to $\ell_{j}$ for some $i \neq j$, we can show that there exists some $k$ such that $\mathcal{D}_{2}\left(x_{i}\right) \cap \mathcal{D}_{2}\left(x_{k}\right) \neq \emptyset$, contradicting our assumption. Therefore no such $\pi^{*}$ exists for any $i \neq j$. But this means that there exists some simple path $\pi^{\prime} \subset A \cap \mathcal{F}$ which joins $\pi_{i}$ and $\pi_{j}$ for some $i \neq j$, which contradicts the fact that $\pi_{i}$ and $\pi_{j}$ belong to different components of $\mathcal{F}$.

## IV. ALGORITHM FOR A SINGLE COMPONENT

In this section we consider a single component $F_{i}$ of $\mathcal{F}$. We present an algorithm that solves the problem within $F_{i}$, ignoring the possibility that robots in $F_{i}$ might collide with robots in other components $F_{j}$. In the next section we will show how to avoid such collisions without changing the motion plans within the individual components. As before we set $S_{i} \triangleq S \cap F_{i}$ and $T_{i} \triangleq T \cap F_{i}$, and assume $\left|S_{i}\right|=\left|T_{i}\right|$.

## A. The motion graph

The motion graph $G_{i}$ of $F_{i}$ is a graph whose vertices represent start or target configurations, and whose edges represent "adjacencies" between these configurations, as defined more precisely below.

Recall that for any $x \in F_{i}$ we defined $D^{*}(x) \triangleq \mathcal{D}_{2}(x) \cap F_{i}$ as the part of the collision disc of $x$ inside $F_{i}$, and recall from Lemma 2 that $D^{*}(x)$ is connected. Moreover, for any two distinct configurations $x_{1}, x_{2} \in S_{i} \cup T_{i}$ we have $D^{*}\left(x_{1}\right) \cap D^{*}\left(x_{2}\right)=\emptyset$, because $\mathcal{D}_{2}\left(x_{1}\right) \cap \mathcal{D}_{2}\left(x_{2}\right)=\emptyset$ by the assumption that the start and target positions are wellseparated. The vertices of our motion graph $G_{i}$ correspond to the start and target configurations in $S_{i} \cup T_{i}$. From now on, and with a slight abuse of notation we will not distinguish between configurations in $S_{i} \cup T_{i}$ and their corresponding vertices in $G_{i}$.

Now consider $F_{i}^{*} \triangleq F_{i} \backslash \bigcup_{x \in S_{i} \cup T_{i}} D^{*}(x)$, the complement of the collision discs of the given start and target configurations in $F_{i}$. This complement consists of several connected
components, which we denote by $F_{i}^{1}, F_{i}^{2}, \ldots$ If the motion graph $G_{i}$ contains an edge $\left(x_{1}, x_{2}\right)$ then there is a component $F_{i}^{\ell}$ that is adjacent to both $D^{*}\left(x_{1}\right)$ and $D^{*}\left(x_{2}\right)$. In other words, two configurations $x_{1}$ and $x_{2}$ are connected in $G_{i}$ if there is a path from $x_{1}$ to $x_{2}$ that stays inside $F_{i}$ and does not cross the collision disc of any other configuration $x_{3} \in S_{i} \cup T_{i}$. Figure 2 illustrates the definition of $G_{i}$. The following observation summarizes the main property of the motion graph.
Observation 1. Suppose all robots in $F_{i}$ are located at a start or target configuration in $S_{i} \cup T_{i}$, and let $\left(x_{1}, x_{2}\right)$ be any edge in $G_{i}$. Then a robot located at $x_{1}$ can move to $x_{2}$ without colliding with any of the other robots.
Remark. We could also work with the dual graph of the partitioning of $F_{i}$ into cells induced by the collision discs. This dual graph would, in addition to vertices representing start and target configurations, also have vertices for the regionss $F_{i}^{\ell}$. For the pebble-motion problem discussed below it is easier to work with the graph as we defined it. This graph may have many more edges, but in the implementation of our algorithm described in Section VI we avoid computing it explicitly.

## B. The unlabeled pebble-motion problem

Using the motion graph $G_{i}$ we can view the motionplanning problem within $F_{i}$ as a pebble-motion problem. (A similar approach was taken in [19], where a sampling-based algorithm for multi-robot motion planning produces multiple pebble problems by random sampling of the configuration space.) To this end we represent a robot located at configuration $x \in S_{i} \cup T_{i}$ by a pebble on the corresponding vertex in $G_{i}$. The pebbles are indistinguishable, like the robots, and they can move along the edges of the graph. At the start of the pebble-motion problem for a graph with vertex set $S_{i} \cup T_{i}$, with $\left|S_{i}\right|=\left|T_{i}\right|$, there is a pebble on every vertex $x \in S_{i}$. The goal is to move the pebbles such that each pebble ends up in vertex in $T_{i}$, under the following conditions: (1) no two pebbles may occupy the same vertex at the same time, and (2) pebbles can only halt at vertices, and (3) at most one pebble may move (that is, be in transit along an edge) at any given time. We call this problem the unlabeled pebble-motion problem. The following lemma follows immediately from Observation 1 .


Fig. 2. (a) A partition of a maximal connected component $F$. The start and target positions consist of the elements $S^{\prime}=\left\{s_{1}, s_{2}, s_{3}, s_{4}\right\}, T^{\prime}=$ $\left\{t_{1}, t_{2}, t_{3}, t_{4}\right\}$, respectively, where the areas $D^{*}(s)$ for $s \in S^{\prime}$ are drawn in green and $D^{*}(t)$ for $t \in T^{\prime}$ are drawn in purple. $F^{*}$ consists of the parts $F^{1}, F^{2}, F^{3}$. (b) A motion graph of $F$.

Lemma 5. Any solution to the unlabeled pebble-motion problem on $G_{i}$ can be translated into a valid collision-free motion plan for the robots in $F_{i}$.

Kornhauser [25, Section 3, first lemma] proved that the unlabeled pebble-motion problem is, in fact, always solvable, and he gave an algorithm to find a solution. Since he did not analyze the running time of his algorithm, we sketch the solution in the proof of the lemma below.
Lemma 6. [25] For any graph $G$ with vertex set $S \cup T$ where $|S|=|T|$, there exists a solution to the unlabeled pebblemotion problem. Moreover, a solution can be found in $O\left(|S|^{2}\right)$ time.

Proof: Let $\mathcal{T}_{G}$ be a spanning tree of $G$. The algorithm performs $O(|S|)$ phases. In each phase, one or more pebbles may be moved and one leaf will be removed from $\mathcal{T}_{G}$, possibly with a pebble on it. After the phase ends, the algorithm continues with the next phase on the modified tree $\mathcal{T}_{G}$, until all pebbles have been removed and the problem has been solved. A phase proceeds as follows.

If there are leaves $v$ that are target vertices then we select such a leaf $v$. If $v$ does not yet contain a pebble, we find a pebble closest to $v$ in $\mathcal{T}_{G}$-this can be done by a simple breadth-first search-and move it to $v$ along the shortest path in $G$. Note that the vertices on the shortest path cannot contain other pebbles, since we took a closest pebble. We now remove the leaf $v$, together with the pebble occupying it, and end the phase. If all leaves in $\mathcal{T}_{G}$ are start vertices, then let $v$ be such a leaf. If $v$ is not occupied by a pebble it can be removed from $\mathcal{T}_{G}$, and the phase ends. Otherwise a pebble resides in $v$, which we move away, as follows. We find the closest unoccupied vertex $w$ to $v$ of $\mathcal{T}_{G}$ and move all pebbles on this shortest path (including the pebble on $v$ ) one step closer to $w$, in order of decreasing distance from $w$. After we evacuated $v$ we remove it from $\mathcal{T}_{G}$ to end the phase.

The algorithm produces paths of total length $O\left(|S|^{2}\right)$, and it can easily be implemented to run in $O\left(|S|^{2}\right)$ time. In some cases $\Omega\left(|S|^{2}\right)$ moves are required, e.g., when $G$ is a single path with all start positions in the first half of the path and all target positions in the second half.
Lemma 7. Suppose we have an instance of our multi-robot path planning problem where $\left|S_{i}\right|=\left|T_{i}\right|$ for every component $F_{i}$ of the free space $\mathcal{F}$. Then for each $F_{i}$ there exists a motion plan $\Pi_{i}$ that brings the robots in $F_{i}$ from $S_{i}$ to $T_{i}$, such that they do not collide with the obstacle space nor with the other robots in $F_{i}$.

## V. Combining single-component plans

We now consider possible interactions between robots contained in different components $F_{i}$ and $F_{j}$ of $\mathcal{F}$. As before, we assume that $\left|S_{i}\right|=\left|T_{i}\right|$ for all $i$. We will show that there exists a permutation $\sigma:\{1,2, \ldots, \ell\} \rightarrow\{1,2, \ldots, \ell\}$ such that we can independently execute the single-component motion plans for each component $F_{i}$ as long as we do so in the order $F_{\sigma(1)}, F_{\sigma(2)}, \ldots, F_{\sigma(\ell)}$.

To obtain this order, we define a directed graph representing the structure of $\mathcal{F}$, which we call the directed-interference
forest $\mathcal{G}=(\mathcal{V}, \mathcal{E})$, where the nodes in $\mathcal{V}$ correspond to the components $F_{i}$. We add the directed edge $\left(F_{i}, F_{j}\right)$ to $\mathcal{E}$ if either there exists a start position $s \in S$ such that $s \in I_{(i, j)}$, or there exists a target position $t \in T$ such that $t \in I_{(j, i)}$. For any $i \in\{1,2, \ldots, \ell\}$, we additionally define $N^{+}(i)$ to be the set of indices of the vertices in the out-neighborhood of $v_{i}$; similarly, $N^{-}(i)$ is defined as the set of indices of the vertices in the in-neighborhood of $v_{i}$.

Note that by Lemma 3 and since $S, T$ are well separated, we cannot have more than one start or target position in $I_{\{i, j\}}$. This implies that $\mathcal{E}$ cannot contain both $\left(v_{i}, v_{j}\right)$ and $\left(v_{j}, v_{i}\right)$. Lemma 4 and the well-separatedness condition additionally imply that $\mathcal{G}$ cannot have an undirected cycle. Thus, $\mathcal{G}$ is a directed forest.

We now produce the desired ordering using $\mathcal{G}$. Consider $F_{i} \in \mathcal{V}$, and suppose that for all $j \in N^{+}(i)$, every robot in $F_{j}$ is at a start position, and for all $j \in N^{-}(i)$, every robot in $F_{j}$ is at a target position. Additionally, suppose that for all $j \notin$ $N^{+}(i) \cup N^{-}(i)$, every robot in $F_{j}$ is at a start or target position. Then, by the definition of $\mathcal{G}$, no robot is at a configuration in $I_{\{i, j\}}$ for any $j \neq i$; thus any motion plan for the robots in $F_{i}$, such as the one described in Section IV, can be carried out without being blocked by the robots not in $F_{i}$. Hence, if we have an ordering $\sigma:\{1,2, \ldots, \ell\} \rightarrow\{1,2, \ldots, \ell\}$ such that for all (directed) edges $\left(v_{i}, v_{j}\right) \in \mathcal{E}, \sigma^{-1}(i)<\sigma^{-1}(j)$, where $\sigma^{-1}$ is the inverse permutation of $\sigma$, then we can execute the motion plans for the robots in $F_{\sigma(1)}, F_{\sigma(2)}, \ldots, F_{\sigma(\ell)}$ in that order. Since $\mathcal{G}$ is a directed forest such an ordering can be produced using topological sorting on the vertices of $\mathcal{G}$. Thus, combining this result with Lemma 7 we obtain:

Theorem 8. Let there be a collection of $m$ unlabeled unit-disc robots in a simple polygonal workspace $\mathcal{W} \subset \mathbb{R}^{2}$, with start and target configurations $S, T$ that are well-separated. Then if for every maximal connected component $F_{i}$ of $\mathcal{F}$ (where $\mathcal{F}$ is the free space for a single unit-disc robot in $\mathcal{W})\left|S \cap F_{i}\right|=$ $\left|T \cap F_{i}\right|$, there exists a collision-free motion plan for these robots starting at $S$ which terminates with every position of $T$ occupied by a robot.

## VI. Algorithmic details

In this section we fill in a few missing details in the description of our algorithm. Specifically, we present an efficient method for generating motion graphs and describe a technique for generating configuration-space paths that correspond to edges in the motion graphs. We also consider the complexity of the various subsets of $\mathcal{F}$ used throughout the algorithm.

## A. Partitioning $\mathcal{F}$

We analyze the combinatorial complexity of $\mathcal{F}^{*} \triangleq \mathcal{F} \backslash$ $\bigcup_{x \in S \cup T} D^{*}(x)$ and $\mathbb{D} \triangleq \bigcup_{x \in S \cup T} D^{*}(x)$.
Lemma 9. The combinatorial complexity of $\mathcal{F}^{*}$ is $O(m+n)$.
Proof: We decompose the complement of the workspace polygon into $O(n)$ trapezoids-this is doable by standard vertical decomposition. We define a set $X$, which consists of the trapezoids, and in addition a collection of $O(m)$ unit discs
that are centered at the start and target positions. We now observe that the regions in $X$ are pairwise interior disjoint (and convex). Hence, it is known [29] that the complexity of the union of the regions in $X$, each Minkowski-summed with a unit disc, is linear in the number of regions plus the sum of the complexities of the original regions. As the result of the Minkowski sum operation of $X$ with a unit disc is the the area $\mathcal{F}^{*}$, we conclude that that the complexity of $\mathcal{F}^{*}$ is $O(m+n)$.

Note that this upper bound still holds if we consider instead of $\mathcal{F}^{*}$ the union of $F_{i}^{*} \triangleq F_{i} \backslash \bigcup_{x \in S_{i} \cup T_{i}} D^{*}(x)$, for all $1 \leqslant$ $i \leqslant q$.

Lemma 10. The combinatorial complexity of $\mathbb{D} \triangleq$ $\bigcup_{x \in S \cup T} D^{*}(x)$, is $O(m+n)$.

Proof: Denote by $d \triangleq\left\{d_{1}, d_{2}, \ldots\right\}$ the segments that define $\partial(\mathbb{D})$. Additionally, denote by $f \triangleq\left\{f_{1}, f_{2}, \ldots\right\}$ and $f^{*} \triangleq\left\{f_{1}^{*}, f_{2}^{*}, \ldots\right\}$ the segments that define $\partial(\mathcal{F}), \partial\left(\mathcal{F}^{*}\right)$, respectively. Note that $\partial(\mathbb{D})$ consists of segments that are elements of $f, f^{*}$ and in addition segments that are subsegments of the elements of $f$, denoted by $f^{\prime} \triangleq\left\{f_{1}^{\prime}, f_{2}^{\prime}, \ldots\right\}$. Obviously the complexity of the segments of $d$, that are elements of $f$ or $f^{*}$, is bounded by $O(m+n)$. It might happen that the segments of $f$ will be split into many subsegments in $f^{\prime}$. However, notice that whenever a segment of $f$ is split the endpoints of each subsegment consist of vertices of $\partial(\mathcal{F})$ or $\partial\left(\mathcal{F}^{*}\right)$. Moreover, exactly two segments in $\partial(\mathbb{D})$ share an endpoint. Thus, the complexity of $\mathbb{D}$ is $O(m+n)$.

## B. Generating motion graphs

We consider a specific component $F$ of $\mathcal{F}$ and construct its motion graph $G$. Denote $F^{*} \triangleq F \backslash \bigcup_{x \in(S \cup T) \cap F} D^{*}(x)$. Note that by the analysis in Section $V$ we can ignore the influence of $\mathcal{D}_{2}(x)$ on connected components in $\mathcal{F}$ that do not contain $x$. We assume that $F^{*}$ breaks into $k$ maximal connected components $F^{1}, \ldots, F^{k}$. The construction of $G$, along with the paths in $F$ that correspond to the edges of $G$, is carried out in two steps. First, for every $F^{i}$ we generate the portion of $G$, denoted by $G^{i}$, whose vertices represent start and target positions that touch the boundary of $F^{i}$. Then, we connect between the various parts of $G$.

We consider a specific connected component $F^{i}$ of $F^{*}$ and describe how the respective portion of the motion graph, namely $G^{i}$, is generated. We split the start and target positions that share a boundary with $F^{i}$ into two subsets: $B^{i}$ are those positions for which the collision disc intersects the boundary of $F$ and $H^{i}$ are those positions for which the collision disc floats inside $F$. See Figure 3. We first handle positions in $B^{i}$. Consider the outer boundary $\Gamma^{i}$ of $F^{i} \backslash \bigcup_{x \in S \cup T} D^{*}(x)$. We argue that each $x \in B^{i}$ can contribute exactly one piece to $\Gamma^{i}$.
Lemma 11. If $x \in B^{i}$ then $\partial\left(\mathcal{D}_{2}(x)\right) \cap \partial\left(F^{i}\right)$ consists of $a$ single component.

Proof: By contradiction, assume that the intersection consists of two maximal connected components. Denote by $y, y^{\prime}$ two configurations on the two components. As $F^{i}$ consists of a single connected component of $F$ there exists a path $\pi_{y y^{\prime}} \subset F$
from $y$ to $y^{\prime}$. Additionally, as $y, y^{\prime}, x$ belong to the same connected component of $\mathcal{F}$ there exist two paths- $\pi_{x y}$ from $x$ to $y$ and $\pi_{x y^{\prime}}$ from $x$ to $y^{\prime}$-that lie entirely in $D^{*}(x)$. Thus, the area that is bounded by the three paths $\pi_{y y^{\prime}}, \pi_{x y}, \pi_{x y^{\prime}}$ contains a patch of forbidden space, which contradicts the fact the our workspace is a simple polygon.

For every $x \in B^{i}$ we arbitrarily select a representative point $\beta^{i}(x) \in \partial\left(\mathcal{D}_{2}(x)\right) \cap F^{i}$. We order the points $\beta^{i}(x)$ clockwise around $\Gamma^{i}$, and store them in a circular list $\mathcal{L}^{i}$. We now incorporate the remaining start and target positions $H^{i}$, namely those positions $x$ for which $D_{2}(x) \cap \partial(F)=\emptyset$. Each position in $H^{i}$ will be connected either to $\Gamma^{i}$ or to the boundary of a collision disc of another position in $H^{i}$ as follows. For each $x \in H^{i}$ we shoot a vertical ray upwards until it hits $\partial\left(F^{i}\right)$. Denote the point where the ray hits $\partial\left(F^{i}\right)$ by $c$. If $c \in \partial\left(\mathcal{D}_{2}\left(x^{\prime}\right)\right)$ for some $x^{\prime} \in H^{i}, x^{\prime} \neq x$ then an edge between $x$ and $x^{\prime}$ is added to $G^{i}$. Otherwise, we let $\beta^{i}(x) \triangleq c$ and insert it into the circular list $\mathcal{L}^{i}$ representing the points $\beta^{i}(x)$ along $\Gamma^{i}$ collected so far. After all positions in $H^{i}$ have been handled in this manner, for each pair of consecutive points $\beta^{i}\left(x^{\prime}\right), \beta^{i}\left(x^{\prime \prime}\right)$ in $\mathcal{L}^{i}$ (along $\Gamma^{i}$ ) we add an edge in $G^{i}$ between the vertices $x^{\prime}$ and $x^{\prime \prime}$. (Notice that some of the positions $x$ whose $\beta^{i}(x)$ appear in $\mathcal{L}^{i}$ belong to $H^{i}$; for example $s_{3}$ in Figure 3) Finally, the connection between portions of the motion graph that represent different parts of $F^{*}$ is achieved through positions shared between two sets $B^{i}, B^{j}$, for $i \neq j$.

## C. Transforming graph edges into paths in the free space

There are three different types of transformations depending on how the edge was created. Let $\left(x, x^{\prime}\right)$ be an edge in $G_{i}$. Consider Figure 3 for an illustration. (i) If both $x$ and $x^{\prime}$ belong to $H^{i}$ (see $\left(s_{4}, s_{3}\right)$ in the figure) then the path simply consists of the two straight-line segments $\overline{x c}$ and $\overline{c x^{\prime}}$. For the remaining two cases we note that if either vertex, say $x$, is in $B^{i}$, then part of the path is a simple curve connecting $x$ to $\beta^{i}(x)$ within $D^{*}(x)$ (see the red curves from $s_{1}$ and $t_{2}$ in the figure). We denote this curve by $\delta_{x}$. (ii) $x, x^{\prime} \in B^{i}$ and the points $\beta^{i}(x)$ and $\beta^{i}\left(x^{\prime}\right)$ are consecutive along $\Gamma^{i}$ (see $\left(t_{1}, t_{2}\right)$ in the figure). The path corresponding to the edge $\left(x, x^{\prime}\right)$ in this case is a concatenation of three sub-paths: $\delta_{x}$, the portion of $\Gamma^{i}$ between $\beta^{i}(x)$ and $\beta^{i}\left(x^{\prime}\right)$ (not passing though the boundary of any other collision disc), and $\delta_{x^{\prime}}$. (iii) $x \in H^{i}$ and $x^{\prime} \in B^{i}$ (see $\left(s_{3}, s_{2}\right)$ in the figure). The path is again a concatenation of three paths: the line segment $\overline{x \beta^{i}(x)}$, the portion of $\Gamma^{i}$ between $\beta^{i}(x)$ and $\beta^{i}\left(x^{\prime}\right)$ (not passing though the boundary of any other collision disc), and $\delta_{x^{\prime}}$.

Notice that for all path types above if a robot $r$ moves from $x$ to $x^{\prime}, x^{\prime}$ is not occupied, and all other robots occupy positions only at $S \cup T \backslash\left\{x, x^{\prime}\right\}, r$ will not collide with any other robot during the motion.

## D. Complexity Analysis

We provide complexity analysis of our algorithm and show that a solution to the problem can be produced within $O\left((m+n) \log (m+n)+m n+m^{2}\right)$ operations, which can be rewritten as $O\left(n \log n+m n+m^{2}\right)$.


Fig. 3. (a) An illustration of a component $F^{3}$ of $F^{*}$ and the structures used for generating the relevant portion of the motion graph. The boundary positions of $F^{3}$ consist of $B^{3}:=\left\{s_{2}, t_{1}, t_{2}, t_{4}\right\}$, while the hole positions consist of $H^{3}:=\left\{s_{3}, s_{4}\right\}$. For every $x \in H^{3}$ its boundary representative $\beta^{3}(x)$ is illustrated as a large black dot. A path between $t_{1}$ and $t_{2}$ is illustrated in red. (b) The motion graph $G^{3}$ induced by $F^{3}$.

Recall that the pebble problem solver (Section IV) operates in $O(m)$ phases, where in each phase a leaf node is removed from the spanning tree of $G$. We show, using Lemma 9 and Lemma 10, that each phase can be transformed into a set of movements for the robots whose combinatorial complexity is $O(m+n)$. The crucial observation is that in one phase each edge of the motion graph is used at most once. Thus the set of robot movements in one phase is bounded by the complexity of the movements corresponding to all the edges in the graph together. These comprise $O(m)$ line segments, portions of the boundaries $\Gamma^{i}$ (whose complexity is $O(m+n)$ by Lemma 9), and the paths $\delta_{x}$ inside the $D^{*}(x)^{\prime} s$ (whose complexity is $O(m+n)$ by Lemma 10 . A path of the latter type, $\delta_{x}$, might be traversed twice: once for reaching $x$ and once for leaving $x$. However asymptotically all the movements together have complexity $O(m+n)$.

We note that the cost of generating $\mathcal{F}$, along with its partitions $\mathcal{F}^{*}$ and $\mathbb{D}$, is bounded by $O((m+n) \log (m+n))$, due to [29]. We also note that deciding whether a solution exists for a certain collection of start and target positions can be carried out in $O((m+n) \log n)$ as follows. We first compute $\mathcal{F}$ in $O(n \log n)$ time, and within the same time preprocess it for efficient point location. Then we query the resulting structure with the $m$ points in $S \cup T$, in $O(\log n)$ time each, and verify that in every component $F_{i}$ of $\mathcal{F}$ it holds that $\left|S_{i}\right|=\left|T_{i}\right|$. Thus, we have the following theorem.

Theorem 12. Let $\mathcal{W}$ be a simple polygon with $n$ vertices and let $S=\left\{s_{1}, \ldots, s_{m}\right\}, T=\left\{t_{1}, \ldots, t_{m}\right\}$ be two sets of $m$ points in $\mathcal{W}$. Additionally, assume that for every two distinct element $x, x^{\prime}$ of $S \cup T$ it holds that $\left\|x-x^{\prime}\right\| \geqslant 4$. Then, given $m$ unlabeled unit disc robots, our algorithm can determine whether a path moving the $m$ robots from $S$ to $T$ exists in $O((m+n) \log n)$ time. If a path exists the algorithm finds it in $O\left(n \log n+m n+m^{2}\right)$ time.

## VII. SEPARATION AND SOLVABILITY

Our results from the previous section imply that a separation distance of 4 ensures that the problem always has a solution,
assuming that each connected component contains the same number of start and target positions. However, for smaller separation values this does not have to be the case.

We define a magnitude $\lambda$ to be the minimum separation between start and target positions of unlabeled discs in a simple polygon, that guarantees that a solution always exists assuming that each connected component contains the same number of start and target positions. Our work in the previous sections shows that $\lambda \leqslant 4$. The following proposition provides a lower bound for the value of $\lambda$.

Proposition 13. $\lambda \geqslant 4 \sqrt{2}-2(\approx 3.646)$.
Proof: We describe a concrete example where the value of $\lambda$ has to be greater than $4 \sqrt{2}-2$ to guarantee solvability. See Figure 4 The two robots $r_{1}, r_{2}$, placed at $s_{1}, s_{2}$, respectively, need to leave through the corridor of width 2 located below their initial positions in order to reach $t_{1}, t_{2}$. Denote the separation between the start positions by $2+\rho$. We wish to find the maximal value $\rho$ for which the problem does not have a solution.

It is clear that the two robots cannot enter the corridor simultaneously. Therefore, if a solution exists there also exists one where $r_{2}$ stays put in $s_{2}$ until $r_{1}$ is well down the corridor. We describe a path of $r_{1}$ in which it maintains a maximal distance from $s_{2}$ : first $r_{1}$ has to slide along the edge $A B$, when arriving at $B$ it revolves around it and then slides down along the other edge connected to $B$. This path has a circular arc which is induced by the rotation about $B$ (depicted by the red arrow). The dashed red arc depicts the trace of the boundary of $r_{1}$ throughout this motion.

Denote by $M$ the tangent point between $r_{2}$ at $s_{2}$ and the edge $C D$. Note that the robots would collide when $r_{1}$ is rotated around $B$ in case that $\left\|B-s_{2}\right\|<3$. Thus, it must be that $\| B-$ $M \| \geqslant 2 \sqrt{2}$. As the segment $A D$ consists of the subsegments $A B, B M, M D$, it follows that

$$
\begin{aligned}
4+\rho & =\|A-D\| \\
& =\|A-B\|+\|B-M\|+\|M-D\| \\
& \geqslant 1+\frac{\rho}{2}+2 \sqrt{2}+1 \\
& =2+2 \sqrt{2}+\frac{\rho}{2}
\end{aligned}
$$

and we conclude that $\rho \geqslant 4 \sqrt{2}-4$, and also $\lambda \geqslant \rho+2=$ $4 \sqrt{2}-2$.

## VIII. OpEN PROBLEMS AND FUTURE WORK

We have studied a basic variant of the multi-robot motionplanning problem, where the goal is to find collision-free motions that bring a given set of indistinguishable unit discs in a simple polygon to a given set of target positions. Under the condition that the start and target positions are separated from each other by a distance of at least 4 , we developed an algorithm that solves the problem in time polynomial in the complexity of the polygon as well as in the number of discs: quadratic in the number of robots and near-linear in the complexity of the polygon.

Our result should be contrasted with the labeled counterpart of the problem, which is NP-hard [8]. In the NP-hardness


Fig. 4. It follows from our paper that when the start and goal positions are well separated, then there is always a solution when each free-space component has the same number of start and goal positions. However, it is not true when the separation condition is not met. In the given example, the two robots need to leave the through a corridor of width 2 located below their initial positions. This is only possible when $\rho \geqslant 4 \sqrt{2}-2$, as otherwise the robots will not be able to get to the corridor without colliding with each other. See proof of Proposition 13 for more details.
proof the discs have different radii, however, and there is no restriction on the separation of the start and target position. Thus one of the main open questions resulting from our study is to settle the complexity of the unlabeled problem without this extra separation condition. Very recently, we have made progress in this direction, by showing that the general unlabeled problem is PSPACE-hard [30]. It should be noted that this proof applies to a slightly different setting that consists of unit-square robots translating amidst polygonal obstacles.

A natural question that arises is what happens when the separation of start or target positions is equal to $2+\rho$, where $0<\rho<2$ ? Would it be possible to design an algorithm, whose running time is polynomial in $m$ and $n$, and also depends in some manner on $\rho$ ? We believe that the work by Alt et al. [31] would be a good starting point for this line of work. Following our discussion in Section VII it will also be interesting to find the exact threshold above which a problem is always solvable.

## REFERENCES

[1] J. T. Schwartz and M. Sharir, "On the piano movers problem: II. General techniques for computing topological properties of real algebraic manifolds," Advances in applied Mathematics, vol. 4, no. 3, pp. 298-351, 1983.
[2] ——, "On the piano movers problem: III. Coordinating the motion of several independent bodies," International Journal of Robotics Research, vol. 2, no. 3, pp. 46-75, 1983.
[3] C.-K. Yap, "Coordinating the motion of several discs," Courant Institute of Mathematical Sciences, Michigan State University, New York, Tech. Rep., 1984.
[4] M. Sharir and S. Sifrony, "Coordinated motion planning for two independent robots," Annals of Mathematics and Artificial Intelligence, vol. 3, no. 1, pp. 107-130, 1991.
[5] B. Aronov, M. de Berg, A. F. van der Stappen, P. Švestka, and J. Vleugels, "Motion planning for multiple robots," Discrete \& Computational Geometry, vol. 22, no. 4, pp. 505-525, 1999.
[6] J. van den Berg, J. Snoeyink, M. C. Lin, and D. Manocha, "Centralized path planning for multiple robots: Optimal decoupling into sequential plans," in Robotics: Science and Systems (RSS), 2009.
[7] 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"," International Journal of Robotics Research, vol. 3, no. 4, pp. 76-88, 1984.
[8] P. G. Spirakis and C.-K. Yap, "Strong NP-hardness of moving many discs," Information Processing Letters, vol. 19, no. 1, pp. 55-59, 1984.
[9] L. E. Kavraki, P. Švestka, J.-C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high dimensional configuration spaces," IEEE Transactions on Robotics and Automation, vol. 12, no. 4, pp. 566-580, 1996.
[10] J. J. Kuffner and S. M. Lavalle, "RRT-Connect: An efficient approach to single-query path planning," in International Conference on Robotics and Automation (ICRA), 2000, pp. 995-1001.
[11] G. Sanchez and J.-C. Latombe, "Using a PRM planner to compare centralized and decoupled planning for multirobot systems," in International Conference on Robotics and Automation (ICRA), 2002.
[12] S. Hirsch and D. Halperin, "Hybrid motion planning: Coordinating two discs moving among polygonal obstacles in the plane," in Workshop on the Algorithmic Foundations of Robotics (WAFR). Springer, 2002, pp. 239-255.
[13] O. Salzman, M. Hemmer, and D. Halperin, "On the power of manifold samples in exploring configuration spaces and the dimensionality of narrow passages," to appear, Workshop on the Algorithmic Foundations of Robotics (WAFR), 2012.
[14] K. Solovey, O. Salzman, and D. Halperin, "Finding a needle in an exponential haystack: Discrete RRT for exploration of implicit roadmaps in multi-robot motion planning," CoRR, vol. abs/1305.2889, 2013.
[15] P. Švestka and M. H. Overmars, "Coordinated path planning for multiple robots," Robotics and Autonomous Systems, vol. 23, pp. 125-152, 1998.
[16] G. Wagner and H. Choset, "M*: A complete multirobot path planning algorithm with performance bounds," in International Conference on Intelligent Robots and Systems (IROS), 2011, pp. 3260-3267.
[17] G. Wagner, M. Kang, and H. Choset, "Probabilistic path planning for multiple robots with subdimensional expansion," in International Conference on Robotics and Automation (ICRA), 2012, pp. 2886-2892.
[18] S. Kloder and S. Hutchinson, "Path planning for permutation-invariant multi-robot formations," in ICRA, 2005, pp. 1797-1802.
[19] K. Solovey and D. Halperin, " $k$-color multi-robot motion
planning," International Journal of Robotics Research, 2013, in press (already appeared on-line).
[20] M. Turpin, N. Michael, and V. Kumar, "Concurrent assignment and planning of trajectories for large teams of interchangeable robots," in International Conference on Robotics and Automation (ICRA), 2013, pp. 842-848.
[21] S. Bereg, A. Dumitrescu, and J. Pach, "Sliding disks in the plane," International Journal of Computational Geometry and Applications, vol. 18, no. 5, pp. 373-387, 2008.
[22] A. Dumitrescu and M. Jiang, "On reconfiguration of disks in the plane and related problems," Computational Geometry: Theory and Applications, vol. 46, no. 3, pp. 191-202, 2013.
[23] O. Goldreich, "Shortest move-sequence in the generalized 15-puzzle is NP-hard," Manuscript, Laboratory for Computer Science, MIT, vol. 1, 1984.
[24] G. Goraly and R. Hassin, "Multi-color pebble motion on graphs," Algorithmica, vol. 58, no. 3, pp. 610-636, 2010.
[25] D. Kornhauser, "Coordinating pebble motion on graphs, the diameter of permutation groups, and applications," M.Sc. Thesis, Department of Electrical Engineering and Computer Scienec, Massachusetts Institute of Technology, 1984.
[26] A. Krontiris, R. Luna, and K. E. Bekris, "From feasibility tests to path planners for multi-agent pathfinding," in Symposium on Combinatorial Search, 2013.
[27] C. H. Papadimitriou, P. Raghavan, M. Sudan, and H. Tamaki, "Motion planning on a graph," in Foundations of Computer Science, 1994, pp. 511-520.
[28] J. Yu and S. M. LaValle, "Distance optimal formation control on graphs with a tight convergence time guarantee," in IEEE International Conference on Decision and Control, 2012, pp. 4023-4028.
[29] K. Kedem, R. Livne, J. Pach, and M. Sharir, "On the union of jordan regions and collision-free translational motion amidst polygonal obstacles," Discrete \& Computational Geometry, vol. 1, pp. 59-70, 1986.
[30] K. Solovey and D. Halperin, "PSPACE-hardness of unlabeled motion planning and variants," CoRR, vol. abs/1408.2260, 2014.
[31] H. Alt, R. Fleischer, M. Kaufmann, K. Mehlhorn, S. Näher, S. Schirra, and C. Uhrig, "Approximate motion planning and the complexity of the boundary of the union of simple geometric figures," Algorithmica, vol. 8, no. 16, pp. 391-406, 1992.


[^0]:    A. Adler is with the Electrical Engineering and Computer Science department at MIT, Massachusetts, USA. The work has been carried out in part during his visit to Tel Aviv University, enabled by the generous Melvin M. Goldberg Fellowship for Research in Israel.
    M. de Berg is with the Department of Mathematics and Computing Science, TU Eindhoven, the Netherlands.
    D. Halperin and K. Solovey are with the Blavatnik School of Computer Science, Tel-Aviv University, Israel. Their work has been supported in part by the 7th Framework Programme for Research of the European Commission, under FET-Open grant number 255827 (CGL-Computational Geometry Learning), by the Israel Science Foundation (grant no. 1102/11), by the German-Israeli Foundation (grant no. 1150-82.6/2011), and by the Hermann Minkowski-Minerva Center for Geometry at Tel Aviv University.

