Space-Aware Reconfiguration*

Dan Halperin[1], Marc van Kreveld[2], Golan Miglioli-Levy[1], and Micha Sharir[1]

[1] The Blavatnik School of Computer Science, Tel Aviv University, Israel
{danha@post.tau.ac.il,golanlevy@mail.tau.ac.il,michas@post.tau.ac.il}
[2] Dept. of Information and Computing Sciences, Utrecht University, Utrecht, the Netherlands
{M.J.vanKreveld@uu.nl}

Abstract. We consider the widely studied problem of *reconfiguring* a set of physical objects into a desired target configuration, a typical (sub)task in robotics and automation, arising in product assembly, packaging, stocking store shelves, and more. In this paper we address a variant, which we call *space-aware reconfiguration*, where the goal is to minimize the physical space needed for the reconfiguration, while obeying constraints on the allowable collision-free motions of the objects. Since for given start and target configurations, reconfiguration may be impossible, we translate the entire target configuration rigidly into a location that admits a valid sequence of moves, so that the physical space required by the start and the translated target configurations is minimized.

We investigate two variants of space-aware reconfiguration for the often examined setting of n *unit discs* in the plane, depending on whether the discs are distinguishable (labeled) or indistinguishable (unlabeled). For the labeled case, we propose a representation of size $O(n^4)$ of the space of all feasible translations, and use it to find, in $O(n^6)$ time, a shortest valid translation, or one that minimizes the enclosing circle or axis-aligned box of both the start and target configurations. For the significantly harder unlabeled case, we show that for almost every direction, there exists a translation that makes the problem feasible. We use this to devise heuristic solutions, where we optimize the translation under stricter notions of feasibility. We present an implementation of such a heuristic, which solves unlabeled instances with hundreds of discs in seconds.

Keywords: Computational geometry, Motion planning, Disc reconfiguration, Smallest enclosing disc

* Work by D.H. and G.M.L. has been supported in part by the Israel Science Foundation Grants 825/15, 1736/19, by the Blavatnik Computer Science Research Fund, and by grants from Yandex and from Facebook. Work by M.v.K. has been supported by the the Netherlands Organisation for Scientific Research under Grant 612.001.651. Work by M.S. has been supported by Grant 260/18 from the Israel Science Foundation, by Grant G-1367-407.6/2016 from the German-Israeli Foundation for Scientific Research and Development, and by the Blavatnik Computer Science Research Fund.

1 Introduction

Consider a set of n objects in the plane or in three-dimensional space and two configurations of these objects, a start configuration S and a target configuration T, where in each configuration the objects are pairwise interior disjoint. A typical *reconfiguration* problem asks to efficiently move the objects from S to T, subject to constraints on the allowable motions, the most notable of which is that all the moves be collision free, where efficiency is measured according to certain prescribed criteria.

In the specific problem studied in this paper, we are given n unit discs in the plane and we wish to move them from some start configuration to a target configuration. A valid move is a translation of one disc in a fixed direction from one placement to another without colliding with the other discs. The goal in earlier works on this problem was to minimize the number of moves, and the goal in the present study is to minimize the size of the physical space needed for the reconfiguration, under the constraint that each disc moves exactly once. This problem, like most problems in the domain of reconfiguration, comes in (at least) two flavors: *labeled* and *unlabeled*. In the labeled version, each object has a unique label, which marks its start placement and its unique target placement. In the unlabeled version the objects are indistinguishable, and we do not care which object finally gets to any specific target placement, as long as all the target placements are occupied at the end of the process; in particular all the objects are isothetic (as are the unit discs in our study). For the unlabeled case, Abellanas et al. [1] have shown that $2n - 1$ moves are always sufficient. Dumitrescu and Jiang [10] have shown that $\lfloor 5n/3 \rfloor - 1$ moves are sometimes necessary, and that finding the minimum number of moves is NP-Hard. For the labeled case, Abellanas et al. [1] have shown that $2n$ moves are always sufficient and sometimes necessary. These are several examples of many reconfiguration problems, which have been studied in discrete and computational geometry; see, e.g., [4,5,8,9]. Varying the type of objects, the ambient space, the constraints on the motion and the optimization criteria, we get a wide range of problems, many of which are hard.

Similar problems arise in robotics. For example, such problems arise when a robot needs to arrange products on a shelf in a store, or when a robot needs to move objects around in order to access a specific product that needs to be picked up; see, e.g., [14,15,16]. In robotics, these problems are often referred to as object *rearrangement* problems. In this paper, though, we will stick to the term reconfiguration, which is also in common use.

Another prominent example from robotics and automation is the *assembly planning* problem (see, e.g., [13]), in which the target configuration of the objects comprises their positions in the final product. The goal of assembly planning is to design a sequence of motions that will bring the parts together to form the final product, and we want this sequence to be (collision-free and) optimal according to various criteria [11,12].

We address here a certain criterion, which, to the best of our knowledge, has hardly been studied earlier: *minimizing the physical space* needed to carry out the desired assembly or reconfiguration. Abellanas et al. [1] did study a similar set of problems, in which the discs are placed inside different types of confined spaces. Their technique shows how to minimize the number of moves, given a prescribed size for a bounding rectangle of S and T. We adopt a different approach, in which we regard T as a rigid configuration that can be placed anywhere in the workspace, and the goal is to find a placement for T for which there exists a feasible (collision-free) sequence of moves, and so that the region occupied by S and by T in its translated location, is minimal according to various possible criteria. In this paper we consider the setup where we only allow T to be translated. We call this problem *space-aware reconfiguration* and we study it in this paper for the case of unit discs in the plane.

We say that a disc is placed at a point p, if its center is placed at p. To avoid confusion between placeholder positions for discs (start or target) and the actual discs placed at these positions, we define a *valid configuration* P to be a finite set of points, such that every pair of points in the set lie at distance $\geqslant 2$ from one another, that is, we can place a unit disc at each point of P, so that the discs are pairwise interior disjoint. For any point p, we denote $D_r(p)$ as the disc centered at p with radius r. If r is not specified, then $D(p)$ is a unit disc ($r = 1$). For any valid configuration P, we denote $D(P) = \{D(p) \mid p \in P\}$.

Let S and T be two valid configurations, of n points each. We look for a sequence of n moves that bring the discs from S to T. A move consists of a single translation of one disc from $D(S)$ to a position in T, such that the disc does not collide with any other (stationary) disc on its way—neither with a disc in a start position, which has not been moved yet, nor with a disc that has already been moved to a target position. Each disc has to perform exactly one such move. We call such a sequence of moves an *itinerary.* We say that an itinerary is *valid* if all of its moves are collision-free. We denote such an Unlabeled Single Translation instance of the problem by UST(S,T), and a Labeled Single Translation instance by LST(S,T,M), where M is the matching between S and T induced by the labels; that is, each position in S is matched by M to the position in T with the same label. We call an instance of the problem *feasible* if it has a valid (collision-free) itinerary.

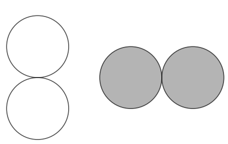

Fig. 1. An infeasible instance for the unlabeled (or labeled) version. The discs of $D(S)$ are drawn as empty discs while the discs of $D(T)$ are drawn as shaded discs. All the figures in this paper were produced with GeoGebra: `www.geogebra.org`.

It is easy to see (consider Figure 1) that even in the unlabeled version, this problem may not have a solution. If the shaded discs in the figure were placed higher (so that their centers were collinear with the center of the top empty disc, say), the problem would have been feasible. We now look for a vector $\vec{v}$ for which there exists a valid itinerary from S to $T + \vec{v}$ (i.e., T translated by $\vec{v}$).[3] In the labeled case, the translated targets retain their labels after the translation. Observe that the initial location of T (when $\vec{v} = (0, 0)$) is now meaningless. From now on, we assume that the input location of T is placed to overlap with S as much as possible, e.g., S and T share their centers of mass or the centers of their smallest enclosing discs. A placement of this kind is ideal for the space-aware paradigm used in this work.

In the space-aware variant studied in this paper, we look for a translation $\vec{v}$ such that (a) S and $T + \vec{v}$ admit a valid itinerary, and (b) some measure of 'nearness' of $T + \vec{v}$ to S is minimized. We denote these space-aware variants as SA-UST(S,T), and SA-LST(S,T,M), for the unlabeled and labeled variants, respectively.

Contribution. We present, in Section 2, a polynomial-time algorithm for the labeled case that runs in $O(n^6)$ time, for constructing the space of all valid translations. We then show, in Section 3, how to find a valid translation that minimizes some measure of space-aware optimality, for example minimizing the area of the smallest enclosing disc of $D(S) \cup D(T + \vec{v})$. All the variants that we study can be solved by algorithms that run in $O(n^6)$ time.

The unlabeled case is much harder (see Appendix A). We first show, in Section 4, that we can find a valid translation in almost any prescribed direction, if we translate T sufficiently far away. We study in Section 5.1 practical heuristic techniques that aim to find shorter valid translations, at the cost of further restricting the notion of validity, and in Section 5.2 we present bounds on how well this heuristic approach works. Finally, in Section 5.3, we present some experimental results of an implementation of the heuristic algorithm. Various technical details and additional comments are delegated to the appendices.

2 Labeled Version: Analysis of the Translation Plane

In this section we consider the labeled version LST of the problem. We are given two valid configurations S and T of n points each, and a one-to-one matching M between the positions of S and those of T, which is the set of pairs $\{(s, M(s)) \mid s \in S\}$, where s and $M(s)$ share the same label, for each $s \in S$. Our goal is to find a translation $\vec{v} \in \mathbb{R}^2$ such that there is a valid collision-free itinerary of n unit discs from S to $T + \vec{v}$. That is, the goal is to define an ordering on the elements of M, denoted by $(s_1, M(s_1)), (s_2, M(s_2)), \ldots, (s_n, M(s_n))$, so that, for each $i = 1, \ldots, n$ in this order, we can translate the disc placed at s_i to the position $M(s_i) + \vec{v}$, so that it does not collide with any still unmoved discs, placed at $s_{i+1}, \ldots, s_n$, or with any of the already translated discs, placed at $M(s_1) + \vec{v}, \ldots, M(s_{i-1}) + \vec{v}$.

[3] The translations from T to $T + \vec{v}$ do not count as moves.

We call a translation $\vec{v}$ a *valid translation* if it yields at least one valid itinerary. In the labeled version, which is easier to solve, we show how to compute the set of all valid translations in $O(n^6)$ time. We then present, in Section 3, three algorithms, each of which finds a valid translation $\vec{v}$ that minimizes a certain measure of proximity between S and $T + \vec{v}$, as reviewed in the introduction.

We first address the subproblem in which $\vec{v}$ is fixed and our goal is to order M so as to obtain a valid itinerary, if at all possible. Let $A = (s, M(s))$ be a pair in the matching. For convenience, we denote s and $M(s)$ by A^S and A^T, respectively. Define the *hippodrome* of two unit discs D, D' to be the convex hull of their union. Observe that the hippodrome is exactly the area that a unit disc will cover while moving from D to D' along a straight trajectory. Define $H_{\vec{v}}(A)$ to be the hippodrome of $D(A^S)$ and $D(A^T + \vec{v})$. Denote by $k_{\vec{v}}$ the overall number of intersecting pairs of hippodromes $(H_{\vec{v}}(A), H_{\vec{v}}(B))$, for all $A \neq B \in M$. See Figure 2 for an illustration.

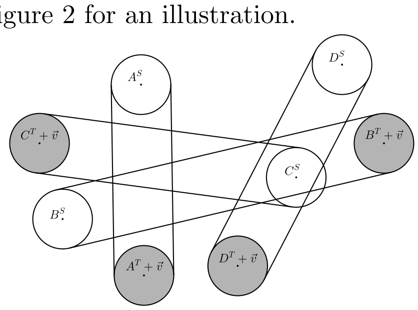

Fig. 2. The hippodromes $H_{\vec{v}}$ for four pairs $A, B, C, D \in M$ and some fixed $\vec{v}$. Notice that even though $H_{\vec{v}}(A) \cap H_{\vec{v}}(B) \neq \emptyset$, there is no restriction that the motion of A must precede or succeed the motion of B. Such restrictions do exist for many other pairs, such as B and C (C has to perform a motion before B).

Theorem 1 (Abellanas et al. [1]). *Let S and T be two valid configurations of n points each, and let $\vec{v}$ be a fixed translation. Let $M : S \to T$ be a bijection between the configurations. Then one can compute, in $O(n \log n + k_{\vec{v}})$ time, a valid itinerary for S and $T + \vec{v}$ with respect to M, if one exists.*

We review the proof of the theorem, adapting it to our notations, and exploit later the ingredients of the analysis for the general problem (where we allow T to be translated). The constraints that the positions of the discs impose on the problem are as follows. We say that a pair $A = (A^S, A^T)$ (in M) has to perform a motion (of $D(A^S)$ to $D(A^T + \vec{v})$) before another pair B, for a given translation $\vec{v}$, if in any ordering Π of M that yields a valid itinerary, the index of A in Π is smaller than the index of B. In other words, for any two pairs $A, B \in M$, A has to perform a motion before B if either the disc $D(A^S)$ blocks the movement of $D(B^S)$ to the position $B^T + \vec{v}$, or the disc $D(B^T + \vec{v})$ blocks the movement of $D(A^S)$ to the position $A^T + \vec{v}$. Formally, we have:

Lemma 1. *Given pairs A, B in the matching and a fixed translation $\vec{v}$, A has to perform a motion before B (with respect to $\vec{v}$) if and only if at least one of the following conditions holds:*

1. $D(A^S) \cap H_{\vec{v}}(B) \neq \emptyset$.
2. $D(B^T + \vec{v}) \cap H_{\vec{v}}(A) \neq \emptyset$.

We next create a digraph whose vertices are the pairs of M, and whose edges are all the ordered pairs $(A, B) \in M^2$, for $A \neq B$, that satisfy (1) or (2). Borrowing a similar notion from assembly planning [13], we call the graph, for a fixed translation $\vec{v}$, the *translation blocking graph* (TBG), and denote it as $G_{\vec{v}}$. Denote the number of edges in $G_{\vec{v}}$ as $m_{\vec{v}}$, and observe that $m_{\vec{v}} \leqslant k_{\vec{v}}$. Indeed, for every edge $(A, B) \in G_{\vec{v}}$ the hippodromes $H_{\vec{v}}(A)$ and $H_{\vec{v}}(B)$ intersect, as is easily verified, but not every pair of intersecting hippodromes necessarily induce an edge; see the pairs A, B in Figure 2. As proved in [1], and as is easy to verify, the subproblem for a fixed $\vec{v}$ is feasible if and only if $G_{\vec{v}}$ is acyclic.

The circular arcs of a hippodrome can be split into two arcs, each of which is x-monotone. This allows us to construct $G_{\vec{v}}$ in $O(n \log n + k_{\vec{v}})$ time, by the sophisticated sweep-line algorithm of Balaban [2], which applies to any collection of well-behaved x-monotone arcs in the plane. (A standard sweeping algorithm would take $O(n \log n + k_{\vec{v}} \log n)$ time.) Checking whether $G_{\vec{v}}$ is acyclic, and, if so, performing topological sorting on $G_{\vec{v}}$, takes $O(n + m_{\vec{v}})$ time. By definition, any topological order of the vertices of $G_{\vec{v}}$, that is of M, defines a valid itinerary. If $G_{\vec{v}}$ has cycles, no valid itinerary exists for $\vec{v}$.

We now consider the translation plane $\mathbb{R}^2$, each of whose points corresponds to a translation vector $\vec{v}$. We say that a point $\vec{v}$ in the translation plane is valid, if the corresponding translation vector is valid, i.e., admits a valid itinerary from S to $T + \vec{v}$. We say that a set of points is valid, if each of its points is valid. Our goal is to construct the region Q of all the valid points (translations), and to partition Q into maximal connected cells, so that all translations in the same cell have common valid itineraries. We actually construct a finer subdivision of the plane, in which all the points in the same cell have the same TBG. Thus, for each cell, either all its points are valid (with at least one common valid itinerary) or all its points are invalid.

Remark. For some instances, Q is empty, as in the scenario depicted in Figure 3. In that case, our algorithms will report that no valid translation exists. We conjecture that every instance of the problem in which the discs in $D(S)$ and in $D(T)$ do not touch one another is feasible (has a valid translation), and leave the settling of this conjecture for future research. A partial support for this conjecture is given in Section 5.2.

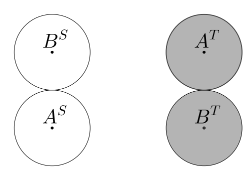

Fig. 3. No itinerary exists between S and $T + \vec{v}$, for any translation $\vec{v}$.

We first fix two pairs $A, B \in M$, and consider the region $\mathcal{V}_{AB}$, which is the locus of those $\vec{v}$ for which the edge AB is present in $G_{\vec{v}}$. We can write $\mathcal{V}_{AB} = \mathcal{V}_{AB}^{(1)} \cup \mathcal{V}_{AB}^{(2)}$, where $\mathcal{V}_{AB}^{(1)}$ (resp., $\mathcal{V}_{AB}^{(2)}$) is the locus of all $\vec{v}$ for which condition (1) (resp., (2)) in Lemma 1 holds. We thus have

$$\mathcal{V}_{AB}^{(1)} = \{\vec{v} \in \mathbb{R}^2 \mid D(A^S) \cap H_{\vec{v}}(B) \neq \varnothing\},$$
$$\mathcal{V}_{AB}^{(2)} = \{\vec{v} \in \mathbb{R}^2 \mid D(B^T + \vec{v}) \cap H_{\vec{v}}(A) \neq \varnothing\}.$$

We call $\mathcal{V}_{AB}^{(1)}$ (resp., $\mathcal{V}_{AB}^{(2)}$) the *vippodrome* of (A, B) of the first (resp., second) type.

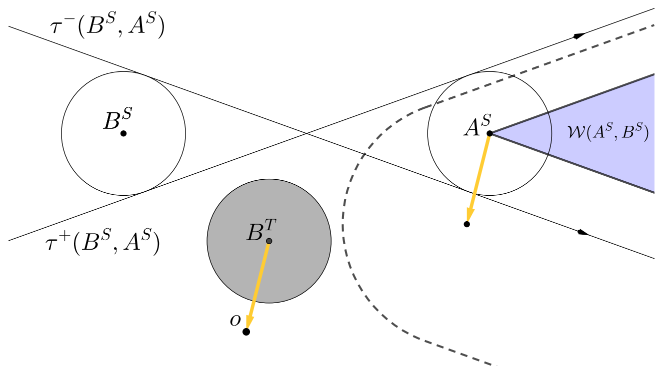

Fig. 4. The vippodrome $\mathcal{V}_{AB}^{(1)}$, which is the region to the right of the dashed curve in the translation plane. The wedge $\mathcal{W}(A^S, B^S)$ is colored in blue. The vippodrome is obtained by first expanding $\mathcal{W}(A^S, B^S)$ by $D_2(o)$, and then by shifting by the vector $-B^T$ (in orange).

To construct $\mathcal{V}_{AB}^{(1)}$, we proceed as follows; see Figure 4. For given pairs $A, B \in M$, consider the two inner tangent lines, $\tau^-(B^S, A^S)$ and $\tau^+(B^S, A^S)$, to $D(B^S)$ and $D(A^S)$, and assume that they are both directed from B^S to A^S, so that B^S lies to the right of $\tau^-(B^S, A^S)$ and to the left of $\tau^+(B^S, A^S)$, and A^S lies to the left of $\tau^-(B^S, A^S)$ and to the right of $\tau^+(B^S, A^S)$. Let $\mathcal{W}(A^S, B^S)$ denote the wedge whose apex is at A^S and whose rays are parallel to (and directed in the same direction as) $\tau^-(B^S, A^S)$ and $\tau^+(B^S, A^S)$. Denote the origin as o. We then have the following representation.

Lemma 2.

$$\mathcal{V}_{AB}^{(1)} = \mathcal{W}(A^S, B^S) \oplus D_2(o) - B^T = (\mathcal{W}(A^S, B^S) - B^T) \oplus D_2(o) \qquad (1)$$
$$\mathcal{V}_{AB}^{(2)} = -\mathcal{W}(B^T, A^T) \oplus D_2(o) + A^S = -(\mathcal{W}(B^T, A^T) - A^S) \oplus D_2(o).$$

Proof. $\mathcal{V}_{AB}^{(1)}$ is the locus of all translations $\vec{v}$ at which $D(A^S)$ intersects $H_{\vec{v}}(B)$. Equivalently, $\mathcal{V}_{AB}^{(1)}$ is the locus of all translations $\vec{v}$ at which $D_2(A^S)$ intersects the segment $e_{\vec{v}} = (B^S,\ B^T + \vec{v})$. The boundary of $\mathcal{V}_{AB}^{(1)}$ thus consists of all translations $\vec{v}$ for which either $e_{\vec{v}}$ is tangent to $D_2(A^S)$ or $B^T + \vec{v}$ touches $D_2(A^S)$. That is, $\partial\mathcal{V}_{AB}^{(1)}$ consists of all translations $\vec{v}$ for which $B^T + \vec{v}$ lies on $\partial\mathcal{W}(A^S, B^S) \oplus D_2(o)$, from which the claim easily follows. The claim for $\mathcal{V}_{AB}^{(2)}$ follows by a symmetric argument, switching between S and T and reversing the direction of the translation. $\square$

Note that the boundary of a vippodrome $\mathcal{V}^{(1)}_{AB}$ is the smooth concatenation of two rays and a circular arc, where the rays are parallel to the rays of $\mathcal{W}(A^S, B^S)$, and where the arc is an arc of the disc $D_2(A^S - B^T)$, of central angle $\pi - \theta$, where θ is the angle of $\mathcal{W}(A^S, B^S)$. The same holds for $\mathcal{V}^{(2)}_{AB}$, with the same disc $D_2(A^S - B^T)$. Hence the boundaries of $\mathcal{V}^{(1)}_{AB}$ and of $\mathcal{V}^{(2)}_{AB}$ (more precisely, the circular portions of these boundaries) might overlap. See Figure 5 for an illustration.

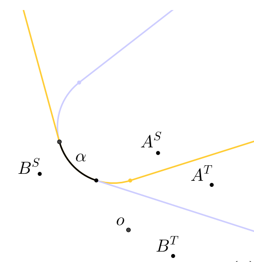

Fig. 5. The vippodromes $\mathcal{V}^{(1)}_{AB}$ and $\mathcal{V}^{(2)}_{AB}$, colored in blue and orange, respectively. The arc α is the overlap portion of the two vippodrome boundaries.

Claim 2. *Let X be a connected set of valid translations. Then every $\vec{v} \in \partial X$ is also a valid translation.*

Proof. Follows from the fact that we allow the translating disc to touch other discs (without penetrating into them). □

Let $\mathcal{V}^\partial = \{\partial \mathcal{V}^{(i)}_{AB} \mid i \in \{1, 2\}, A \neq B \in M\}$ and observe that $|\mathcal{V}^\partial| = 2n(n-1)$. Define the vippodrome arrangement $\mathcal{A}(\mathcal{V}^\partial)$, induced by M, to be the arrangement formed by the curves $\mathcal{V}^\partial$; it is an arrangement of $O(n^2)$ rays and circular arcs. Assuming general position, the only overlaps between features of the arrangement are between circular arcs of the two vippodromes of the same ordered pair A, B (see Figure 5 again). To avoid these overlaps, we partition each of these arcs into two subarcs at the point where the overlap begins (so each circle $\partial D_2(A^S - B^T)$ contributes at most three arcs to the arrangement). It is worth mentioning that it is fairly easy to handle instances that are not in general position, where degeneracies may appear, such as collinear rays, or overlapping circular arcs, of two unrelated vippodromes. In the rest of the paper, we assume general position, to simplify the presentation.

We now observe that any pair of features (rays and circular arcs) of the (modified) arrangement intersect at most twice. Hence the number of vertices in $\mathcal{A}(\mathcal{V}^\partial)$ is at most $O(n^4)$, and so the overall complexity of the arrangement is also $O(n^4)$. For simplicity, we assume from now on that $\mathcal{A}(\mathcal{V}^\partial)$ is clipped to within a sufficiently large rectangle that meets all the faces of the arrangement and contains all its bounded edges. Consider a face f of $\mathcal{A}(\mathcal{V}^\partial)$. We showed that for every ordered pair of pairs $(A, B) \in M^2$, $A \neq B$, the edge AB is either in every graph $G_{\vec{v}}$, for $\vec{v} \in f$, or in none of these graphs. Hence all the graphs $G_{\vec{v}}$, for $\vec{v} \in f$, are identical, and we denote this common graph as G_f.

We can construct $\mathcal{A}(\mathcal{V}^\partial)$ either in $O(n^4 \log n)$ time using a plane-sweep procedure, or in $O(n^2 \lambda_4(n^2))$ time, using the incremental procedure described in [18, Theorem 6.21, p. 172].

After $\mathcal{A}(\mathcal{V}^\partial)$ has been constructed, we traverse its faces and construct the graphs G_f, over all faces f, noting that when we cross from a face f to an adjacent face f', the graph changes by the insertion or deletion of a single edge.[4] We then test each graph G_f for acyclicity. The union of (the closure of) all the faces f for which G_f is acyclic is the desired region Q of valid translations. For each face f that participates in Q, we run a linear-time procedure for topological sorting of G_f, and the order that we obtain[5] defines a valid itinerary for all translations $\vec{v} \in f$. The running time, for a fixed face f, is $O(n + m_f)$, where m_f is the number of edges of G_f. In the worst case we have $m_f = O(n^2)$, so the overall cost of the algorithm is $O(n^6)$.

Remark. It would be interesting to see whether dynamic algorithms for maintaining acyclicity in a directed graph, under insertions and deletions of edges (namely, fully dynamic cycle detection algorithms), could be applicable when we traverse the faces of $\mathcal{A}(\mathcal{V}^\partial)$. Such algorithms can be found in [17], but they do not seems to improve the asymptotic running time of our algorithm. The only efficient algorithms that we are aware of, those with low total update time, only support insertions of edges but not deletions; see [6,7]. If a fully dynamic algorithm of this kind were available, with sublinear update time for each insertion and deletion, it would clearly improve the total running time of our procedure of finding Q, as well as the space-aware optimizations that are presented next.

3 Labeled Version: Space-Aware Optimization

We examine three different variants of the optimization criterion for the SA-LST problem: SA-LST$_{|\vec{v}|}(S,T,M)$ for minimizing the length of the translation vector $\vec{v}$; SA-LST$_{\mathsf{AABR}}(S,T,M)$ for minimizing the area of the axis-aligned bounding rectangle of (the union of the discs of) $D(S) \cup D(T+\vec{v})$, denoted as $\mathsf{AABR}(D(S) \cup D(T+\vec{v}))$; and SA-LST$_{\mathsf{SED}}(S,T,M)$ for minimizing the area of the smallest enclosing disc of (the union of the discs of) $D(S) \cup D(T+\vec{v})$, denoted as $\mathsf{SED}(D(S) \cup D(T+\vec{v}))$.

Solving SA-LST$_{|\vec{v}|}(S,T,M)$ is fairly straightforward by observing that the function $|\vec{v}|$ has a single global minimum at the origin o and no other local minima. For lack of space we defer the details of its solution as well as the details of the only-slightly more involved solution of SA-LST$_{\mathsf{AABR}}(S,T,M)$ to Appendix B. SA-LST$_{\mathsf{SED}}(S,T,M)$ however seems to be much more challenging, and we review here some key ideas in its solution, deferring most of the details to Appendix B. We denote the smallest enclosing disc of a set P as $\mathsf{SED}(P)$,

---

[4] Although we assume general position, circular arcs of vippodromes may still overlap (see Figure 5). Notice, however, that the overlapping arcs bound vippodromes that induce the same constraint on the itinerary and hence crossing the overlapping arcs still incurs insertion or deletion of a single edge.

[5] In general, G_f can have exponentially many topological orders, each of which yields a valid itinerary.

and its radius as $r(P)$. Our goal is to minimize $r(S \cup (T + \vec{v}))$,[6] over all valid translations $\vec{v} \in Q$.

Consider the farthest-neighbor Voronoi diagrams $\mathsf{FVD}(S)$ of S and $\mathsf{FVD}(T)$ of T. Note that $\mathsf{FVD}(T+\vec{v}) = \mathsf{FVD}(T)+\vec{v}$. If we fix $\vec{v}$, the smallest enclosing disc $D = \mathsf{SED}(S\cup(T+\vec{v}))$ is centered either at a Voronoi vertex ξ of $\mathsf{FVD}(S\cup(T+\vec{v}))$, in which case ∂D passes through the three points that lie farthest from (i.e., define) ξ, or at a Voronoi edge e of this diagram, in which case D is the diametral disc formed by the two points that lie farthest from (i.e., define) e. In the former case either (i) the three farthest points belong to the same set (S or $T + \vec{v}$), or (ii) two points belong to one set and the third belongs to the other set. In the latter case, either (iii) the two farthest points belong to the same set, or (iv) they belong to different sets.

In Case (i) we have $O(n)$ candidates for the center of $\mathsf{SED}(S \cup (T + \vec{v}))$, each of which is either a stationary vertex of $\mathsf{FVD}(S)$ or a vertex of $\mathsf{FVD}(T)$ shifted by $\vec{v}$. By the symmetry of the setup, it suffices to focus, without loss of generality, on vertices of $\mathsf{FVD}(S)$. Let ξ be such a vertex. Our desired disc has to be at least as large as the smallest disc $D^*(\xi)$ that is centered at ξ and contains S. Shifting T by $\vec{v}$ is equivalent to shifting ξ by $-\vec{v}$ into $\mathsf{FVD}(T)$: the farthest neighbor(s) of ξ in $T+\vec{v}$ are those that determine the cell, edge or vertex of $\mathsf{FVD}(T)$ containing $\xi - \vec{v}$. We thus need to take the shifted and reflected diagram $\xi - \mathsf{FVD}(T)$ and overlay it with the arrangement $\mathcal{A}(\mathcal{V}^{\partial})$ of the vippodrome boundaries (or rather with its valid portion Q). Each valid face of the overlay has a unique point $t \in T$ such that $t+\vec{v}$ is the farthest neighbor of ξ in $T+\vec{v}$. It is then easy to find, in time proportional to the complexity of the overlay, a valid translation that minimizes the smallest enclosing disc. Since the complexity of the overlay is still $O(n^4)$ (its new vertices are intersections of edges of $\xi - \mathsf{FVD}(T)$ with edges of $\mathcal{A}(\mathcal{V}^{\partial})$, and there are only $O(n^3)$ such intersections), this takes $O(n^4)$ time. Multiplying by the number of vertices ξ, we get a total of $O(n^5)$ time.

The other cases (ii)–(iv) are more involved. Again, for lack of space, we defer further details to Appendix B, where we provide a detailed analysis of each of the three problems and prove the following theorem (split in the appendix into three separate theorems, one theorem for each problem).

Theorem 3. *Problems* SA-LST$_{|\vec{v}|}(S,T,M)$, SA-LST$_{\mathsf{AABR}}(S,T,M)$, *and* SA-LST$_{\mathsf{SED}}(S,T,M)$, *can each be solved in* $O(n^6)$ *time.*

4 Unlabeled Version: Preliminary Analysis

In this section we study the reconfiguration problem for unlabeled discs. The main result of the section is summarized the following theorem.

[6] This is indeed an equivalent formulation to the one given earlier: The smallest enclosing disc of $D(S) \cup D(T + \vec{v})$ has the same center as the disc that we find, and its radius is larger by 1.

Theorem 4. *Let S and T be two valid configurations, of n points each. For every direction δ, except possibly for finitely many special directions, there exists a translation $\vec{v} \in \mathbb{R}^2$ in direction δ such that the unlabeled problem* UST*$(S,T+\vec{v})$ is feasible.*

Remark. Note that the theorem implies that we can always move the discs from S to T in $2n$ moves: first move the discs from S to $T + \vec{v}$, using the valid itinerary provided by the theorem, and then move the discs from $T + \vec{v}$ to T by translating each of them by $-\vec{v}$, in the order of their centers in direction $\vec{v}$. This almost reproduces the result of [1], already mentioned in the introduction, where the bound is $2n - 1$, for the case where we are not allowed to shift the target locations. In some sense, our result is stronger, in that in the second step, all the discs are translated by the same vector $-\vec{v}$.

Proof. Let C be a valid configuration of n points in the plane. Let c, c' be two points in C and let $b(c, c')$ be their perpendicular bisector. Put $\mathcal{B}(C) = \{b(c, c') \mid c, c' \in C, dist(c, c') = 2\}$, which is the set of all perpendicular bisectors (common inner tangents) of any two touching discs of $D(C)$. We say that a direction is *generic* for C if it is not parallel to any line in $\mathcal{B}(C)$. (Note that there are only $O(n)$ non-generic directions.) We fix a generic direction δ for both S and T. Observe that δ is also generic for $T+\vec{v}$, for any vector $\vec{v}$. Without loss of generality, assume that δ is horizontal and points to the right. We define $\Pi_\delta(C)$ to be the reverse lexicographical order of the points in C, that is, $\Pi_\delta(C) = (c_1, c_2, \ldots, c_n)$, so that, for any $1 \leqslant i < j \leqslant n$, c_i is to the right (or at the same x-coordinate but above) c_j. We now fix a matching M_δ according to the orders $\Pi_\delta(S)$ and $\Pi_\delta(T)$, by aligning both orders, i.e., $M_\delta(s_i) = t_i$, for $i = 1, \ldots, n$. The matching M_δ transforms the problem to the labeled version LST(S,T,M_δ). We claim that this specific instance is always feasible, and, moreover, admits valid translations in direction δ. Order M_δ in the same order of $\Pi_\delta(S)$ and $\Pi_\delta(T)$, i.e., $(s_1, t_1), \ldots, (s_n, t_n)$, and denote this order as $\Pi(M_\delta)$. We claim that one can always choose $\vec{v}$, in direction δ, such that $(s_1, t_1 + \vec{v}), \ldots, (s_n, t_n + \vec{v})$ is a valid itinerary.

We apply a simpler variant of the techniques developed in Section 2. Since we have already assigned the fixed order $\Pi(M_\delta)$ to M_δ, we do not need to take into consideration all the vippodromes, but only the ones that impose constraints that violate $\Pi(M_\delta)$. Let A_i be the pair $(s_i, t_i) \in M_\delta$ (so $D(s_i)$ is the disc that moves in step i). Let

$$\mathcal{V}_{bad}(\delta) = \{\mathcal{V}^{(i)}_{A_k A_l} \mid i \in \{1, 2\}, A_k, A_l \in M_\delta, k > l\},$$

and observe that $|\mathcal{V}_{bad}(\delta)| = n(n-1)$. In other words, $\mathcal{V}_{bad}(\delta)$ is the subset of all the vippodromes V, such that, for each $\vec{v} \in V$, the constraint that V represents violates the itinerary according to $\Pi(M_\delta)$ between S and $T + \vec{v}$. Thus, in order to find a valid translation $\vec{v}$ in direction δ, it suffices to show that the ray ρ from the origin in direction δ (the positive x-axis by assumption) is not fully contained in the union of the vippodromes in $\mathcal{V}_{bad}(\delta)$. We will prove a stronger claim.

Lemma 3. *There exists a ray $\rho' \subseteq \rho$ such that $\rho' \cap V = \varnothing$ for every $V \in \mathcal{V}_{bad}(\delta)$.*

Proof. Let $V = \mathcal{V}^{(1)}_{BA}$, such that $A, B \in M_\delta$ and A performs a motion before B according to $\Pi(M_\delta)$; that is, $V \in \mathcal{V}_{bad}(\delta)$. Let $\tau^-(A^S, B^S), \tau^+(A^S, B^S)$ be the rays of $\mathcal{W}(B^S, A^S)$ (recall the setup discussed in Section 2, depicted in Figure 4, and note that here A and B change roles). By construction, if A performs a motion before B according to the itinerary, then A^S is lexicographically larger than B^S, and so B^S is to the left of (or has the same x-coordinate and below) A^S. Since the positive x-direction δ is generic, $D(A^S)$ and $D(B^S)$ cannot lie vertically above one another and have a common inner tangent. Let σ be the ray that emanates from B^S in the direction from A^S to B^S. By our assumption, σ points either directly downwards, or else to the left (contained in the open left halfplane that contains B^S on its right boundary). Note that σ is the mid-ray of the wedge $\mathcal{W}(B^S, A^S)$ (see Figure 4). In the former case, the opening angle of $\mathcal{W}(B^S, A^S)$ is strictly smaller than π, and in the latter case, it is at most π. In either case, $\mathcal{W}(B^S, A^S)$ is disjoint from the rightward-directed ray from B^S (the ray in direction δ). This implies that $\mathcal{W}(B^S, A^S)$ cannot fully contain any rightward-directed ray. Since $V = \mathcal{W}(B^S, A^S) \oplus D_2(o) - A^T$, the same claims hold for V as well. The argument for $\mathcal{V}^{(2)}_{BA}$ is similar. In conclusion, ρ must exit from every vippodrome of $\mathcal{V}_{bad}(\delta)$, which implies the lemma. □

Hence, there are infinitely many translations $\vec{v}$ in ρ, that do not belong to any vippodrome $V \in \mathcal{V}_{bad}(\delta)$. By construction, this implies that $\mathrm{UST}(S, T+\vec{v})$ is feasible for every such $\vec{v}$. Furthermore, the above holds for every generic direction δ. This completes the proof of Theorem 4. □

It is now fairly simple to devise an algorithm for finding a valid translation $\vec{v}$ and for constructing a valid itinerary from S to $T + \vec{v}$. First, choose a generic direction δ, in $O(n \log n)$ time, and assume it to point at the positive x-direction. Calculate $\Pi_\delta(S), \Pi_\delta(T)$ and M_δ in $O(n \log n)$ time. Construct $\mathcal{V}_{bad}(\delta)$ according to M_δ, in $O(n^2)$ time. Intersect all the vippodromes of $\mathcal{V}_{bad}(\delta)$ with the positive x-axis, and find the rightmost intersection point $\vec{v}_{\max}$ with these vippodromes, which can be done in $O(n^2)$ time. Any translation $\vec{v}$ to the right of $\vec{v}_{\max}$ has a valid itinerary from S to $T + \vec{v}$, given by the order $\Pi(M_\delta)$. The overall running time of this algorithm is $O(n^2)$.

5 Unlabeled Version: Space-Aware Practical Solutions

5.1 Heuristics for Short Valid Translations

The analysis in Section 4, while providing an abundance of valid translations, has the disadvantage that the valid translations that it yields are potentially too long (one needs to go sufficiently far away in the δ-direction to get out of all the 'bad' vippodromes). This is undesirable with our space-aware objective in mind, where we seek short valid translations. In this subsection, we provide a heuristic for finding shorter valid translations, thereby obtaining shorter heuristic solutions to SA-UST$_{|\vec{v}|}(S,T)$. For more heuristics, also for SA-UST$_{\mathsf{AABR}}(S,T)$ and SA-UST$_{\mathsf{SED}}(S,T)$, see Appendix C. The resulting algorithms are faster than those obtained for the labeled case (at the cost of not guaranteeing optimality). As the unlabeled problem is much harder than the labeled case (and we believe it to be NP-hard), we make no attempt at solving it exactly.

In the next subsection we will show how to choose a good direction δ for which we can give reasonable upper bounds on the length of the shortest valid translation, or the valid translation $\vec{v}$ that minimizes the smallest enclosing disc of $S \cup (T + \vec{v})$. For now, fix a generic direction δ for S and T. In practice, one might want to choose a sufficiently dense set of generic directions, in the hope of improving the quality of the following solution.

Recall the algorithm for finding a valid translation at the end of Section 4. It is likely that $\vec{v}_{\max}$ is not the shortest valid translation in direction δ, according to $\Pi(M_\delta)$. In order to find the shortest such valid translation, we again construct the bad vippodromes and intersect them with the positive x-axis. Instead of finding the rightmost intersection point, we sort the resulting valid intervals along the x-axis (the ones that are free of all bad vippodromes), and output the leftmost valid point, which is clearly the shortest valid translation, under the present restricted setup. This process is mildly slower than the original algorithm (see the end of Section 4), since we need to sort $O(n^2)$ intervals, and can be carried out in $O(n^2 \log n)$ time.

5.2 Bounding the Heuristic Solutions

The analysis in the preceding subsection provides heuristics for obtaining short valid translations, but gives no guarantees on how well we approximate the optimal valid translation. In this subsection we show how to choose a good direction δ for which we can give a reasonable upper bound on the length of the shortest valid translation in direction δ, or of a valid translation $\vec{v}$ in direction δ that minimizes $\mathsf{SED}(S \cup (T + \vec{v}))$. Recall that we denote by $r(P)$ the radius of the smallest enclosing disc of some set of points P.

Theorem 5. *Let S and T be two valid configurations, of n points each, such that S and T share the centers of their smallest enclosing discs. There exists a translation $\vec{v}$ such that* $\mathrm{UST}(S, T + \vec{v})$ *is feasible and $|\vec{v}| = O((r(S) + r(T))n)$. This also bounds $r(S \cup (T + \vec{v}))$.*

Proof. See Appendix D.

The physical space needed for the reconfiguration is worse by the factor $O(n)$, compared to the ideal bound $O(r(S) + r(T))$, since this is (asymptotically) the

Table 1. Different input types of UST. For each input type, the configurations are presented in their initial positions (sharing the centers of their smallest enclosing discs) and the translated position (of T) according to an approximate shortest valid translation (in red), produced by our heuristic algorithm.

Conf.	Sources/Targets	Initial	Translated	n	$r(S)+r(T)$	$\lvert\vec{v}\rvert$
Circle	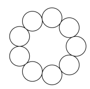	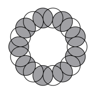	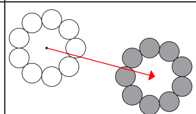	100	65.67	190.19
				200	129.32	376.24
				500	320.31	913.79
				1,000	638.60	1,757.26
Packing[7]	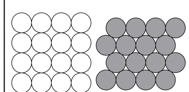	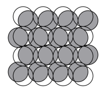	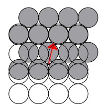	100	27.01	5.53
				210	38.88	2.18
				506	60.76	3.25
				1,024	87.22	19.75
Cross	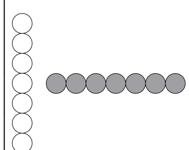	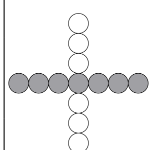	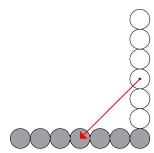	100	200	140.07
				200	400	281.43
				500	1,000	706.15
				1,000	2,000	1,413.47
Random[8]	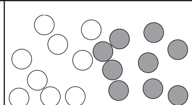	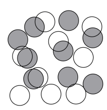	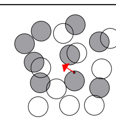	100	36.78	16.96
				200	51.70	34.01
				500	82.26	78.49
				1,000	116.06	147.61

minimum value of $r(S \cup (T + \vec{v}))$, over all translations $\vec{v}$. Interestingly, we can attain this ideal asymptotic bound if the discs of $D(S)$ are sufficiently separated, by at least some fixed constant, and so are the discs of $D(T)$. That is, we have:

Theorem 6. *Let S and T be two valid configurations, of n points each, such that S and T share the centers of their smallest enclosing discs. Assume that there exists a fixed constant $\varepsilon > 0$, so that the distance between any pair of points in S, or any pair of points in T, is at least $2 + \varepsilon$. Then, for* any *direction δ, there exists a translation $\vec{v}$ in direction δ, such that $\mathrm{UST}(S, T + \vec{v})$ is feasible and $|\vec{v}| = O((r(S) + r(T))/\sqrt{\varepsilon})$. The same bound also holds for $r(S \cup (T + \vec{v}))$.*

Proof. See Appendix D.

Note that Theorem 6 is stronger than Theorem 5 also in that it holds for *every* direction δ, whereas Theorem 5 only holds for restricted values of δ.

---

[7] The numbers n are chosen so that the source configuration will be square-like.

[8] The results are averaged over 10 different instances, one of which is depicted.

5.3 Implementation of the Heuristic Algorithm

We implemented the heuristic algorithm for finding an approximate shortest valid translation for UST instances, as outlined in Section 5.1. Our program is written in Python 3.7, and the experiments that we report below were carried out on an Intel Core i7-7500U CPU clocked at 2.9 GHz with 24 GB of RAM.

Table 1 shows the results obtained with our implementation for four different types of input, with the number of discs per type ranging between 100 and $1,024$.

For each instance, we tried $1,000$ different directions δ, and in the table we compare the shortest valid translation that the algorithm produced (over all different directions) with the asymptotically optimal value $r(S) + r(T)$. See Appendix E for a detailed description of each of the input types. As expected, the running time of the implementation is slightly super-quadratic—see Figure 9. Our program runs in about 25 seconds on inputs with a $1,000$ discs; notice that the number of bad vippodromes in such instances is $999,000$.

6 Further Research

Our research can be extended in multiple ways within the space-awareness framework. We could allow two translations per disc while using minimal physical space according to the minimal bounding box or enclosing circle size. These problems can be studied with and without a global rigid translation of the target configuration. All intermediate positions then contribute to the space usage as well. Alternatively, we could allow a translation and rotation of the target configuration, or allow other motion paths than straight line paths. Since these problems are more general, they seem harder to solve with optimal space usage. We may also study space-aware reconfiguration for discs of varying sizes (the labeled version), or for other, more complicated shapes.

Viewing assembly planning from the space-aware perspective raises many challenging problems. We aim to find the smallest space (e.g., a round tabletop of minimum radius) where we can put the separate parts that need to be assembled into the final product, and such that the entire assembly process can occur within this space. The problem is complex since we need to store intermediate subassemblies, such that we can bring together some subassemblies into their relative placement in the final product, while avoiding other subassemblies, all within the same space.

Acknowledgement. The authors thank Mikkel Abrahamsen, Michael Hoffmann and Bettina Speckmann for helpful discussions. They also thank the Lorentz Center in Leiden, where part of the research took part.

References

1. M. Abellanas, S. Bereg, F. Hurtado, A. G. Olaverri, D. Rappaport, and J. Tejel. Moving coins. *Comput. Geom.*, 34(1):35–48, 2006.
2. I. J. Balaban. An optimal algorithm for finding segments intersections. In *Proc. 11th Annu. Symposi. Comput. Geom.*, pages 211–219, 1995.
3. A. Banik, B. B. Bhattacharya, and S. Das. Minimum enclosing circle of a set of fixed points and a mobile point. *Comput. Geom.*, 47(9):891–898, 2014.
4. S. Bereg and A. Dumitrescu. The lifting model for reconfiguration. *Discrete Comput. Geom.*, 35(4):653–669, 2006.
5. S. Bereg, A. Dumitrescu, and J. Pach. Sliding disks in the plane. *Int. J. Comput. Geom. Appl.*, 18(5):373–387, 2008.
6. A. Bernstein and S. Chechi. Incremental topological sort and cycle detection in expected total time. In *Proc. 29th Annu. ACM-SIAM Symposi. Discrete Algorithms*, pages 21–34. SIAM, 2018.
7. S. Bhattacharya and J. Kulkarni. An improved algorithm for incremental cycle detection and topological ordering in sparse graphs. In *Proc. 31st Annu. ACM-SIAM Symposi. Discrete Algorithms*, pages 2509–2521, 2020.
8. G. Călinescu, A. Dumitrescu, and J. Pach. Reconfigurations in graphs and grids. *SIAM J. Discrete Math.*, 22(1):124–138, 2008.
9. A. Dumitrescu. Mover problems. In J. Pach, editor, *Thirty Essays in Geometric Graph Theory*, pages 185–211. Springer, Berlin-Heidelberg, 2013.
10. A. Dumitrescu and M. Jiang. On reconfiguration of disks in the plane and related problems. *Comput. Geom.*, 46(3):191–202, 2013.
11. T. Geft, A. Tamar, K. Goldberg, and D. Halperin. Robust 2d assembly sequencing via geometric planning with learned scores. In *15th Proc. IEEE Internat. Conf. on Automation Science and Engineering (CASE)*, pages 1603–1610, 2019.
12. M. Goldwasser, J.-C. Latombe, and R. Motwani. Complexity measures for assembly sequences. In *Proc. IEEE Internat. Conf. on Robotics and Automation*, volume 2, pages 1851–1857, 1996.
13. D. Halperin, J. Latombe, and R. H. Wilson. A general framework for assembly planning: The motion space approach. *Algorithmica*, 26(3-4):577–601, 2000.
14. S. D. Han, N. M. Stiffler, A. Krontiris, K. E. Bekris, and J. Yu. High-quality tabletop rearrangement with overhand grasps: Hardness results and fast methods. In *Proc. 13th Conf. Robotics: Science and Systems*, 2017.
15. G. Havur, G. Ozbilgin, E. Erdem, and V. Patoglu. Geometric rearrangement of multiple movable objects on cluttered surfaces: A hybrid reasoning approach. In *Proc. IEEE Internat. Conf. on Robotics and Automation (ICRA)*, pages 445–452, 2014.
16. M. Levihn, T. Igarashi, and M. Stilman. Multi-robot multi-object rearrangement in assignment space. In *Proc. IEEE/RSJ Internat. Conf. on Intelligent Robots and Systems (IROS)*, pages 5255–5261, 2012.
17. D. J. Pearce and P. H. J. Kelly. A dynamic topological sort algorithm for directed acyclic graphs. *ACM Journal of Experimental Algorithmics*, 11, 2006.
18. M. Sharir and P. K. Agarwal. *Davenport-Schinzel Sequences and Their Geometric Applications.* Cambridge University Press, New York, 1995.

A Related Work: Hardness of Reconfiguration

To the best of our knowledge, it has not been proven that deciding the existence of a valid itinerary for the unlabeled version (assuming T is stationary, and allowing one move per disc) is NP-hard, but similar reconfiguration problems, such as those called OMC [1] and U-TRANS-RP [10], have been shown to be NP-hard. Both problems seek to find a valid itinerary of translations of discs from a start configuration to a target configuration. Specifically, in OMC (one move per coin), each of the discs is given a subset of possible final targets, and can move (as in our model) exactly once. In U-TRANS-RP, the goal is to decide whether a valid itinerary of at most k moves exists, for the unlabeled variant. The NP-hardness of U-TRANS-RP is shown for the case where k is smaller than n (some discs are already at the target locations, and some may move more than once). Although at the moment we do not know whether it is possible to reduce any of these problems to our setting, we strongly believe that the unlabeled version of our setting is indeed NP-hard, and so minimizing some criterion of optimality for SA-UST(S,T) is most likely even harder.

B Labeled Space-Aware Optimization: Detailed Proofs

B.1 Minimizing the Translation Vector

We define $\text{SA-LST}_{|\vec{v}|}(S,T,M)$ as the following problem: Given an LST instance, of two configurations S and T, and a matching M between S and T, find a valid translation $\vec{v} \in \mathbb{R}^2$, with respect to M, such that $|\vec{v}|$ is minimized.[9]

Theorem 7. $\text{SA-LST}_{|\vec{v}|}(S,T,M)$ *can be solved in* $O(n^6)$ *time.*

Proof. Let $\psi(\vec{v}) = |\vec{v}|$ and observe that it has a single global minimum at the origin o and no other local minima. It thus follows that ψ attains its minimum over any connected region K at a point on ∂K, unless $o \in K$. By Claim 2, for every valid face f of $\mathcal{A}(\mathcal{V}^{\partial})$, ∂f is also valid. Thus, a valid translation of minimum length $\vec{v}$ is either o or a point (in the valid portion) of some curve of $\mathcal{A}(\mathcal{V}^{\partial})$. This suggests the following procedure. Construct the portion Q of the valid translations in $\mathcal{A}(\mathcal{V}^{\partial})$, as explained in Section 2, in $O(n^6)$ time. Check if o is a valid point. If so, report it as the desired valid translation of minimum length. If not, iterate over all the valid edges e of the arrangement. For each such e, find the point $\vec{v} \in e$ that minimizes $\psi(\vec{v})$, in constant time. Taking the shortest $\vec{v}$ among all the translations that we have obtained yields the desired translation. Hence, $\text{SA-LST}_{|\vec{v}|}(S,T,M)$ can be solved in $O(n^6)$ time, as asserted. □

Note that, once Q has been computed, the running time of the optimization step is only $O(n^4)$.

[9] Recall that we assume that at translation $\vec{v} = \vec{0}$ the centers of mass of S and $T + \vec{v} = T$, or the centers of their smallest enclosing discs, coincide.

B.2 Minimizing the Area of the Axis-Aligned Bounding Rectangle

We define SA-LST$_{\mathsf{AABR}}$(S,T,M) as the following problem: Given an LST instance, of two configurations S and T, and a matching M between S and T, find a valid translation $\vec{v} \in \mathbb{R}^2$, with respect to M, such that the area of $\mathsf{AABR}(D(S) \cup D(T + \vec{v}))$ is minimized.

Theorem 8. SA-LST$_{\mathsf{AABR}}$(S,T,M) *can be solved in* $O(n^6)$ *time.*

Proof. Denote $\mathsf{AABR}(D(S))$ by R_1 and $\mathsf{AABR}(D(T))$ by R_2. Note that $\mathsf{AABR}(D(S) \cup D(T + \vec{v}))$ is the axis-aligned bounding box of $R_1 \cup (R_2 + \vec{v})$. Write $R_1 = [a_S, b_S] \times [c_S, d_S]$ and $R_2 = [a_T, b_T] \times [c_T, d_T]$. Then, putting $\vec{v} = (x, y)$, we have that $\mathsf{AABR}(D(S) \cup D(T + \vec{v}))$ is the axis-aligned bounding box of

$$\Big([a_S, b_S] \times [c_S, d_S]\Big) \cup \Big([x + a_T, x + b_T] \times [y + c_T, y + d_T]\Big),$$

so it is the rectangle $[a^*(x), b^*(x)] \times [c^*(y), d^*(y)]$, where

$$a^*(x) = \min\{a_S, x + a_T\}, \qquad b^*(x) = \max\{b_S, x + b_T\},$$
$$c^*(y) = \min\{c_S, y + c_T\}, \qquad d^*(y) = \max\{d_S, y + d_T\}.$$

Let $\varphi(\vec{v}) = \varphi(x, y)$ denote the area of $\mathsf{AABR}(D(S) \cup D(T + \vec{v}))$. That is,

$$\varphi(x, y) = (b^*(x) - a^*(x))\,(d^*(y) - c^*(y))\,.$$

The function $b^*(x) - a^*(x)$ is piecewise linear in x, with the two breakpoints $a_S - a_T$, $b_S - b_T$. Similarly, the function $d^*(y) - c^*(y)$ is piecewise linear in y, with the two breakpoints $c_S - c_T$, $d_S - d_T$ (the breakpoints of either function may appear in any order). Each function is constant over the interval between the breakpoints, has slope -1 to the left of the interval, and slope $+1$ to the right of the interval.

Consider the vertical lines through the breakpoints of $b^*(x) - a^*(x)$ and the horizontal lines through the breakpoints of $d^*(y) - c^*(y)$, and denote the set of these four lines by L. L partitions the plane into nine rectangular (bounded and unbounded) regions. The function $\varphi(x, y)$ is constant over the center region R_0, is a linear function in x over the regions to the left and to the right of R_0, is a linear function in y over the regions above and below R_0, and is a hyperbolic paraboloid, of the form $\pm(x - \alpha)(y - \beta)$, over each of the other four regions. Observe that φ is continuous. Since none of the expressions for φ has any local minimum, it thus easily follows that, for any connected region K, φ attains its minimum over K at a point on ∂K, unless R_0 is fully contained in K, in which case the minimum is attained at all the translations in R_0.

Hence, we proceed exactly as in the previous problem. We check if an arbitrary point in R_0 is valid. If so, report it as the desired valid translation that minimizes $\varphi(\vec{v})$. If not, it means that R_0 is not contained in any cell of Q, and so it suffices to check for the minimum on the boundaries of the faces of Q. We compute Q, in $O(n^6)$ time, iterate over its valid edges, and for each such edge e, we minimize φ over e, which can be done in constant time. (Note that an edge e may cross some lines of L. In such a case, we split e into $O(1)$ subedges, each fully contained in one of the rectangular regions of the partition, and minimize φ separately over each subedge.) We output the translation $\vec{v}$ that minimizes $\varphi(\vec{v})$, over all the translations that we have obtained. Altogether, the procedure runs in $O(n^6)$ time. Hence, SA-LST$_{\mathsf{AABR}}(S,T,M)$ can be solved in $O(n^6)$ time, as asserted. □

Here too, the cost of the optimization step is only $O(n^4)$.

B.3 Minimizing the Area of the Smallest Enclosing Disc

The problem here, denoted as SA-LST$_{\mathsf{SED}}(S,T,M)$, is to find a valid translation $\vec{v}$ that minimizes the radius of the smallest enclosing disc of $S \cup (T + \vec{v})$.[10] We denote the smallest enclosing disc of a set P as $\mathsf{SED}(P)$, and its radius as $r(P)$. A similar problem was studied in [3], in which a characterization of the locus of the center of the smallest enclosing disc, and its radius, are given for a static set of points and only one mobile point, moving along a straight line. Here, we study a more intricate problem, as our mobile points are more numerous and are not moving along a line, but are moving rigidly according to the valid translations of the region Q.

Theorem 9. SA-LST$_{\mathsf{SED}}(S,T,M)$ *can be solved in* $O(n^6)$ *time.*

Proof. For completeness, we repeat some of the details already provided in Section 3. Consider the farthest-neighbor Voronoi diagram $\mathsf{FVD}(S)$ of S and the farthest-neighbor Voronoi diagram $\mathsf{FVD}(T)$ of T. Note that $\mathsf{FVD}(T + \vec{v}) = \mathsf{FVD}(T) + \vec{v}$. If we fix $\vec{v}$, the smallest enclosing disc $D = \mathsf{SED}(S \cup (T + \vec{v}))$ is centered either at a Voronoi vertex ξ of $\mathsf{FVD}(S \cup (T + \vec{v}))$, in which case ∂D passes through the three points that lie farthest from (i.e., define) ξ, or at a Voronoi edge e of this diagram, in which case D is the diametral disc formed by the two points that lie farthest from (i.e., define) e. In the former case either (i) the three farthest points belong to the same set (S or $T + \vec{v}$), or (ii) two points belong to one set and the third belongs to the other set. In the latter case, either (iii) the two farthest points belong to the same set, or (iv) they belong to different sets.

Case (i) is the simplest. In this case we have $O(n)$ candidates for the center of $\mathsf{SED}(S \cup (T + \vec{v}))$, each of which is either a stationary vertex of $\mathsf{FVD}(S)$ or

[10] This is indeed an equivalent formulation to the one given earlier: The smallest enclosing disc of $D(S) \cup D(T + \vec{v})$ has the same center as the disc that we find, and its radius is larger by 1.

a vertex of $\mathsf{FVD}(T)$ shifted by $\vec{v}$. By the symmetry of the setup, it suffices to focus, without loss of generality, on vertices of $\mathsf{FVD}(S)$. Let ξ be such a vertex. Our desired disc has to be at least as large as the smallest disc $D^*(\xi)$ that is centered at ξ and contains S. Shifting T by $\vec{v}$ is equivalent to shifting ξ by $-\vec{v}$ into $\mathsf{FVD}(T)$: the farthest neighbor(s) of ξ in $T+\vec{v}$ are those that determine the cell, edge or vertex of $\mathsf{FVD}(T)$ containing $\xi-\vec{v}$. We thus need to take the shifted and reflected diagram $\xi-\mathsf{FVD}(T)$ and overlay it with the arrangement $\mathcal{A}(\mathcal{V}^{\partial})$ of the vippodrome boundaries (or rather with its valid portion Q). Each valid face of the overlay has a unique point $t \in T$ such that $t+\vec{v}$ is the farthest neighbor of ξ in $T+\vec{v}$. It is then easy to find, in time proportional to the complexity of the overlay, a valid translation that minimizes the smallest enclosing disc. Since the complexity of the overlay is still $O(n^4)$ (its new vertices are intersections of edges of $\xi-\mathsf{FVD}(T)$ with edges of $\mathcal{A}(\mathcal{V}^{\partial})$, and there are only $O(n^3)$ such intersections), this takes $O(n^4)$ time. Multiplying by the number of vertices ξ, we get a total of $O(n^5)$ time.

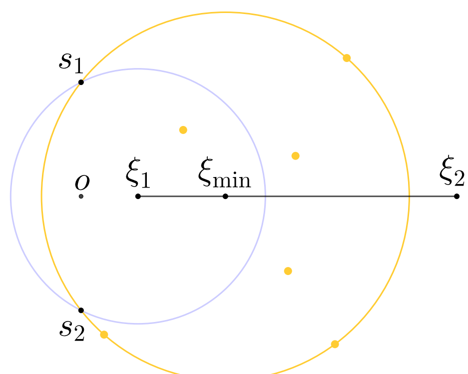

Fig. 6. The origin o is the midpoint of s_1s_2. The points of $T+\vec{v}$ are drawn in orange. D_{ξ_1} (in blue) is the smallest enclosing disc of S. $D_{\xi_{\min}}$ (in orange) is the smallest enclosing disc of $S \cup (T+\vec{v})$ along the farthest-neighbor Voronoi edge $[\xi_1, \xi_2]$.

Consider next the case where the boundary of the smallest enclosing disc passes through only two points of S. (Handling the case where it passes through two points of $T+\vec{v}$ is fully symmetric.) This covers both cases (ii) and (iii). Any such pair of points must define an edge of $\mathsf{FVD}(S)$, and this diagram has only $O(n)$ edges. Fix such an edge, defined by two points $s_1, s_2 \in S$, and denote it as $e = e_{s_1,s_2}$; see Figure 6. Assume, without loss of generality, that the perpendicular bisector $b(s_1, s_2)$ of s_1s_2 is the x-axis, with the origin placed at the midpoint of s_1s_2. Then we can write e as an interval $[\xi_1, \xi_2]$. Without loss of generality, assume that $0 \leqslant \xi_1 < \xi_2$. (If 0 is an interior point of e, split it into two subintervals at 0 and handle each of them separately. Note also that e may be unbounded, in which case one of its endpoints is $\pm\infty$.) For each $\xi \in e$, let D_ξ denote the disc centered at ξ so that its bounding circle passes through s_1 and

s_2. Clearly, for a translation $\vec{v}$, D_ξ is the smallest enclosing disc of $S \cup (T + \vec{v})$, with respect to placements in e, if and only if $T + \vec{v} \subset D_\xi$ and D_ξ is the smallest such disc (meaning that ξ is closest to the origin); by definition, $S \subset D_\xi$ for every such D_ξ. Let $V(\xi)$, for $\xi \in e$, denote the set of all translations $\vec{v}$ for which $T + \vec{v} \subset D_\xi$. We have $V(\xi) = \bigcap_{t \in T}(D_\xi - t)$. Define

$$V = \{(\vec{v}, \xi) \mid \xi \in e,\ \vec{v} \in V(\xi)\},$$

which is a subset of $\mathbb{R}^3$. Clearly, each horizontal cross-section of V (at a fixed ξ) is convex, and it is also easy to check that each vertical line (at a fixed $\vec{v}$) intersects V in a connected interval. From these one can show that V is a connected set.

If the global minimum $\xi_{\min}$ of V is attained at a valid translation then $D_{\xi_{\min}}$ is the desired disc. Otherwise, it is easily checked that the minimum is attained at a (valid) translation $\vec{v}$ that lies on the boundary of some vippodrome. We thus take each of the $O(n^2)$ such boundaries γ, and mark on it the maximal subarcs of valid translations (each delimited at intersection points of γ with other vippodrome boundaries). As is easily checked, for each $\vec{v} \in \gamma$ and for each $t \in T$, the set of those ξ for which $t + \vec{v} \in D_\xi$ is a ray of the form $[g_t^-(\vec{v}), +\infty)$ or of the form $(-\infty, g_t^+(\vec{v})]$. Hence $\vec{v} \in V(\xi)$ if and only if ξ lies in the intersection of these rays, that is

$$\max_{t \in T} g_t^-(\vec{v}) \leqslant \xi \leqslant \min_{t \in T} g_t^+(\vec{v}).$$

We therefore compute the sandwich region between the upper envelope of the functions g_t^- and the lower envelope of the functions g_t^+, and look for the ξ-lowest *valid* point in the region (which necessarily lies on the upper envelope). This takes $O(\lambda_s(n) \log n)$ time, for some constant parameter s (see, e.g., [18]).

Finding the global minimum $\xi_{\min}$ of V (and checking whether it is attained at a valid translation) is simpler. Observe that the global minimum $\xi_{\min}$ is achieved at either ξ_1 or at ξ' such that $D_{\xi'}$ is of radius $r(T)$. The procedure fails to find a translation in $V(\xi_{\min})$ if none of the vippodrome boundaries intersect $V(\xi_{\min})$, and so either every translation in $V(\xi_{\min})$ is valid, or every translation is invalid. That said, it suffices to check for the validity of only one translation in $V(\xi_{\min})$, which can be done in $O(n^2)$ time.

For the running time of this procedure, we need to repeat it for the $O(n)$ pairs of neighbors (s_1, s_2) in $\mathsf{FVD}(S)$ (and similarly for $\mathsf{FVD}(T)$). For each pair we run $O(n^2)$ one-dimensional minimization procedures, over the vippodrome boundaries, each costing $O(n^2)$ time (just to mark on it the valid portions). Thus the total running time is $O(n^5)$, which is subsumed in the bound $O(n^6)$ on finding the valid translations.

Consider finally Case (iv), in which the smallest enclosing disc is the diametral disc of two point $s \in S$ and $t + \vec{v} \in T + \vec{v}$. There are $O(n^2)$ pairs (s, t) to test. Fix one such pair (s_0, t_0), and denote the corresponding diametral disc as $D(\vec{v})$.

A point $s \in S \backslash \{s_0\}$ is in $D(\vec{v})$ if and only if $\angle(s_0, s, t_0 + \vec{v}) \geqslant \pi/2$, which is equivalent to $t_0 + \vec{v}$ lying in the halfplane $H_{s_0,s}$, defined so that it does not

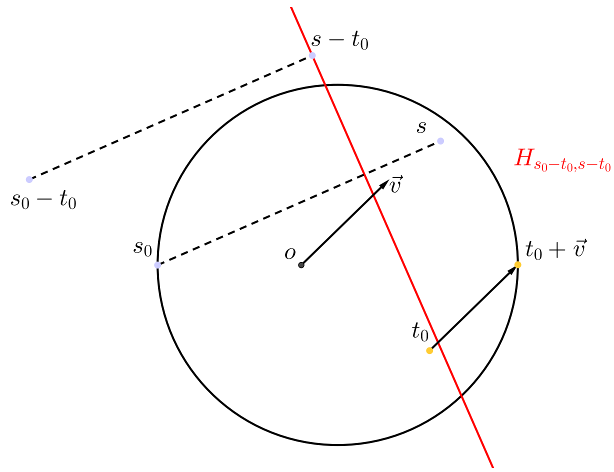

Fig. 7. The points s_0, s (in blue), t_0 (in orange), and their respective translated points according to $-t_0$ and $\vec{v}$. The halfplane $H_{s_0-t_0,s-t_0}$, lying to the right of its bounding line, are drawn in red. Observe that $\vec{v} \in H_{s_0-t_0,s-t_0}$ and so s is inside the diametral disc of s_0 and $t_0 + \vec{v}$.

contain s_0, its bounding line passes through s, and is orthogonal to $s_0 s$. This in turn is equivalent to $\vec{v}$ lying in the similarly defined halfplane $H_{s_0-t_0,s-t_0}$; see Figure 7. This has to hold for every s, and, symmetrically, $\vec{v}$ also has to lie in each halfplane $H_{s_0-t_0,s_0-t}$, for each $t \in T \backslash \{t_0\}$. We thus form the intersection of these $2(n-1)$ halfplanes, in $O(n \log n)$ time, to obtain a convex polygon K_{s_0,t_0} with $O(n)$ edges. We then overlay K_{s_0,t_0} with the vippodrome arrangement $\mathcal{A}(\mathcal{V}^{\partial})$, and seek a valid translation $\vec{v}$ inside K_{s_0,t_0} for which $|\vec{v} + t_0 - s_0|$ (that is, the diameter of our disc) is minimized. This can easily be done in $O(n^4)$ time, by intersecting $\partial K_{s_0,t_0}$ with each vippodrome boundary, and by tracing the valid portions of the vippodrome boundaries within the polygon, similar to the procedure outlined earlier. Altogether this takes $O(n^6)$ time. □

C More Heuristics for the Unlabeled version

In the following, we provide more heuristics for the space-aware optimization problems in the unlabeled version. We exploit Theorem 4, which shows the existence of at least one valid translation for any generic direction δ, and the algorithm, presented in Section 4, for finding such translations. These results motivate the following algorithms. Fix a generic direction δ for S and T, and recall that it takes $O(n \log n)$ time to compute $\Pi_\delta(S)$, $\Pi_\delta(T)$, M_δ, and $\Pi(M_\delta)$. This transforms the problem to a labeled instance, according to M_δ, with the additional constraint that we require the discs to move according to the order $\Pi(M_\delta)$. Since the order is now fixed, it suffices to consider, as in Section 4, only the vippodromes in $\mathcal{V}_{bad}(\delta)$. We form the arrangement $\mathcal{A}_{bad}$ of their boundaries, and collect all the faces of $\mathcal{A}_{bad}$ that lie outside the union of the bad vippodromes. This is the set of all the valid translations, with respect to the matching M_δ and the order $\Pi(M_\delta)$.

Constructing and collecting the faces outside the union of the bad vippodromes, takes $O(n^4 \log n)$, or $O(n^2\lambda_4(n^2))$ time, as shown in Section 2. Finding an optimal translation where we either minimize its length or the area of the axis-aligned bounding rectangle, can be done in $O(n^4)$ additional time (see Sections B.1 and B.2).[11] Optimizing the size of the smallest enclosing disc, using our current algorithm (see Section B.3), is more expensive, and takes $O(n^6)$ time.

As already noted, the only role of the direction δ is to define the orders $\Pi_\delta(S)$ and $\Pi_\delta(T)$ (and thus also the order $\Pi(M_\delta)$ of M_δ), as the reverse lexicographical order, induced by δ and its orthogonal direction, of the disc centers. The algorithms mentioned above will find an optimal valid translation, under this restricted notion of validity, within the entire set of valid translations, not necessarily in direction δ.

We can do better, at the cost of further constraining the set of valid translations, considering only the translations $\vec{v}$ in direction δ, as explained in Section 5.1, in $O(n^2 \log n)$ time.

The above performance bounds also hold for finding the translation $\vec{v}$ that minimizes the area of the axis-aligned bounding rectangle of $D(S) \cup D(T + \vec{v})$: here, at each valid interval I along the ray, we need to compute the minimum of a univariate function of constant complexity over I, which takes constant time.

For minimizing the smallest enclosing disc of $D(S) \cup D(T+\vec{v})$, a suitable variant of our technique, restricted to the one-dimensional problem of translations in direction δ, the running time becomes $O(n^4)$, although we strongly believe that it can be improved to roughly $O(n^3)$.

D Bounds on the Heuristic Solutions: Proofs

Theorem 5. *Let S and T be two valid configurations, of n points each, such that S and T share the centers of their smallest enclosing discs. There exists a translation $\vec{v}$ such that* $\mathrm{UST}(S, T + \vec{v})$ *is feasible and* $|\vec{v}| = O((r(S) + r(T))n)$. *This also bounds* $r(S \cup (T + \vec{v}))$.

Proof. Consider the discs in $D(S)$ (the case of T is essentially identical). We define the following set Δ_S of directions. For each pair $s_1, s_2 \in S$ of neighbors

[11] The running time $O(n^6)$ that is stated there for these optimizations is only a consequence of the (expensive) need to test each face of $\mathcal{A}(\mathcal{V}^\partial)$ for acyclicity, which is no longer needed under our new notion of validity. The other steps of either of the two algorithms only take $O(n^4)$ time.

in the (nearest-neighbor) Delaunay triangulation of S, we include in Δ_S the directions of the two inner tangents of $D(s_1)$ and $D(s_2)$, obtaining at most $12n-24$ directions (each tangent is directed both ways). We obtain a similar set Δ_T of at most $12n-24$ directions from the pairs of neighbors in the Delaunay triangulation of T.

Lemma 4. *For any pair of unit discs in $D(S)$ whose centers are not Delaunay neighbors, the angle between their inner tangents in the wedges that contain neither of the two discs is at least $\pi/2$. The same holds for $D(T)$.*

Proof. Let p_1 and p_2 be two points in S so that p_1 and p_2 are not Delaunay neighbors. Then the segment p_1p_2 intersects a Delaunay edge q_1q_2, where q_1 and q_2 are two other points in S. Since the discs of $D(S)$ are interior-disjoint, p_1 and p_2 are at least $2\sqrt{2}$ apart (since by Thales' theorem, one of the angles $\angle p_1q_1p_2$ or $\angle p_1q_2p_2$ is at least $\pi/2$). The argument for T is identical, and the lemma follows. □

Let δ be a direction whose smallest angle from any direction in $\Delta_S \cup \Delta_T$ is as large as possible. In particular, we can choose δ such that the angle it forms with the direction of any separating line is $\Omega(1/n)$; actually, the reasoning above implies a lower bound of at least $\pi/(24n)$. Observe that δ is generic for S and T, and so, by Theorem 4, there exists a valid translation in direction δ. By rotating the plane, we may assume that δ is the positive x-direction. We compute the reverse lexicographical orders $\Pi_\delta(S)$, $\Pi_\delta(T)$, as defined in Section 4, and obtain the corresponding matching M_δ and its order $\Pi(M_\delta)$.

Lemma 5. *The translation $\vec{v} = ((r(S) + r(T) + 2)(1 + 8n), 0)$ is a valid translation (in direction δ) with respect to the matching M_δ and its order $\Pi(M_\delta)$.*

Proof. By the analysis of Section 4, it suffices to show that $\vec{v}$ does not belong to any bad vippodrome. For this, it suffices to show that any bad vippodrome intersects the positive x-axis to the left of $\vec{v}$. Consider such a bad vippodrome, say $\mathcal{V}_{AB}^{(1)}$, where $A, B \in M_\delta$, such that A appears after B in $\Pi(M_\delta)$. Symmetric arguments apply to bad vippodromes of the form $\mathcal{V}_{AB}^{(2)}$. Assume, without loss of generality, that B^S is above, or at the same height, as A^S. Since A appears after B in $\Pi(M_\delta)$, A^S is lexicographically smaller than B^S, so the direction of $\overrightarrow{B^S A^S}$ points to the left (or vertically down); recall the proof of Lemma 3. Let r_1, r_2 be the two rays of $\mathcal{V}_{AB}^{(1)}$, and assume that r_1 intersects the x-axis to the right of the intersection point of r_2 with the x-axis (r_1 is in direction $\tau^+(B^S, A^S)$); see Figure 8. Let p be the point that r_1 emanates from, and let $q = (q_x, 0)$ be the intersection point of r_1 and the x-axis. We observe that the y-component of the direction of r_1 is negative. Indeed, r_1 forms an angle of at most $\pi/2$ with the direction of $B^S A^S$, which is leftward-directed or points down; in the latter case the angle is strictly smaller than $\pi/2$, as easily follows from our choice of δ. It follows that q_x increases as p moves either to the right or up. By the vippodrome construction, and the locations of S and T, both the x and the y coordinates of p cannot exceed $K := r(S) + r(T) + 2$. We may thus assume that $p = (K, K)$; for any other point, q_x is smaller.

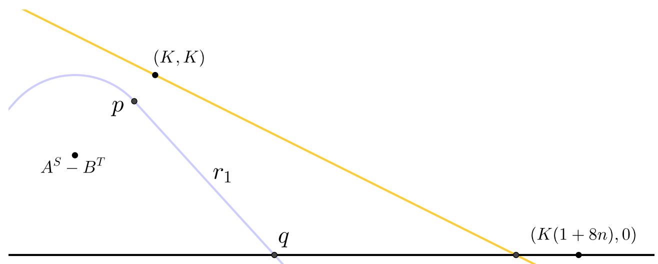

Fig. 8. The x-axis is drawn in black, a bad vippodrome $\mathcal{V}_{AB}^{(1)}$ in blue. The orange line, with slope $-\tan\frac{\pi}{24n}$, meets the x-axis to the right of any bad vippodrome of the first type, which is directed downwards (and to the left), like $\mathcal{V}_{AB}^{(1)}$.

If the x-component of the direction of r_1 is non-positive, $q_x \leqslant K$, so we may assume that it is positive. If A^S and B^S are not neighbors in the Delaunay diagram of S, the angle of $\mathcal{W}(A^S, B^S)$ is at most $\frac{\pi}{2}$, by Lemma 4. Thus, the slope of r_1 is at most -1. If A^S and B^S are neighbors in the Delaunay diagram of S, then the directions of their common inner tangents are in Δ_S, and so the slope of r_1 is at most $-\tan\frac{\pi}{24n}$. We then again may assume that the slope of r_1 is $-\tan\frac{\pi}{24n}$, as for any smaller slope, q_x is smaller.

The supporting line of r_1, according to our upper bounding assumptions, is

$$y + x\tan\frac{\pi}{24n} - K\left(1 + \tan\frac{\pi}{24n}\right) = 0,$$

and thus

$$q_x = K\left(1 + \frac{1}{\tan\frac{\pi}{24n}}\right) < K(1 + 8n),$$

where one can verify that the inequality holds for any $n \geqslant 1$. □

Repeating a symmetric argument for bad vippodromes of the second type, Theorem 5 now follows readily, both for the length of $\vec{v}$ and for $r(S \cup (T + \vec{v}))$, as both of them are clearly $O((r(S) + r(T))n)$. □

Theorem 6. *Let S and T be two valid configurations, of n points each, such that S and T share the centers of their smallest enclosing discs. Assume that there exists a fixed constant $\varepsilon > 0$, so that the distance between any pair of points in S, or any pair of points in T, is at least $2 + \varepsilon$. Then, for* any *direction δ, there exists a translation $\vec{v}$ in direction δ, such that $\mathrm{UST}(S, T + \vec{v})$ is feasible and $|\vec{v}| = O((r(S) + r(T))/\sqrt{\varepsilon})$. The same bound also holds for $r(S \cup (T + \vec{v}))$.*

Proof. Observe that the separation property guarantees that every direction is generic. As is easily checked, the angle between the inner tangents of any pair of discs in $D(S)$, within the wedge containing none of the discs, is at least $c\sqrt{\varepsilon}$, for a suitable constant c. Hence, the opening angle of any vippodrome is at most $\pi - c\sqrt{\varepsilon}$. This implies that, for any direction δ, and for any bad vippodrome, with respect to $\Pi(M_\delta)$, the angle that its ray r_1, in the notation of the proof of Lemma 5, forms with the δ-direction is at least $\frac{c}{2}\sqrt{\varepsilon}$. Following the same analysis as in the proof of Lemma 5, we can show that there exists a valid translation $\vec{v}$ in direction δ, whose length is $O((r(S) + r(T))/\sqrt{\varepsilon})$. This also bounds $r(S \cup (T + \vec{v}))$. □

E Heuristic Implementation Instances

In the experimental results reported in Section 5.3, we consider the following input types, each of which is dense in a different way (see also Table 1):

1. Circle: the configurations are placed on the circumference of a circle. The target discs are slightly rotated (by $\frac{\pi}{n}$) in order to avoid an easy matching.
2. Packing: the source discs are placed in a grid packing. The target discs are placed in a Kepler's packing.
3. Cross: the source (resp., target) discs are placed along a vertical (resp., horizontal) line.
4. Random: both configurations are sampled uniformly at random from a square of size $(2.6\sqrt{n} \times 2.6\sqrt{n})$. For each configuration size, we average the results over 10 different runs.

In Figure 9 we present the running times of the implementation for random input instances. For $n = 100i$ discs ($i = 1, \ldots, 10$), we choose 10 random configurations of n discs (see Table 1), and for each of them, we optimize in 10 random directions δ, for a total of 100 runs per input size n .

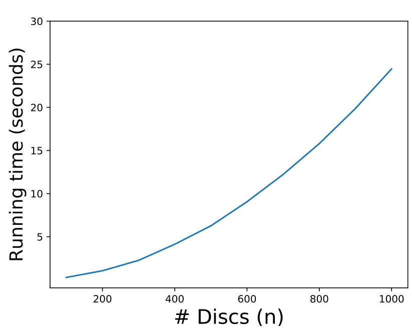

Fig. 9. Running time of the heuristic as a function of the number of discs in the start (and hence also the target) configuration, for the random input type. Each entry (#Discs= $100, 200, \ldots, 1,000$) in the graph is the average running time of 100 instances.