On reconfiguration of disks in the plane and related problems*

Adrian Dumitrescu[†] Minghui Jiang[‡]

June 11, 2012

Abstract

We revisit two natural reconfiguration models for systems of disjoint objects in the plane: translation and sliding. Consider a set of n pairwise interior-disjoint objects in the plane that need to be brought from a given start (initial) configuration S into a desired goal (target) configuration T, without causing collisions. In the translation model, in one move an object is translated along a fixed direction to another position in the plane. In the sliding model, one move is sliding an object to another location in the plane by means of a continuous rigid motion (that could involve rotations). We obtain various combinatorial and computational results for these two models:

(I) For systems of n congruent unlabeled disks in the translation model, Abellanas et al. showed that $2n-1$ moves always suffice and $\lfloor 8n/5 \rfloor$ moves are sometimes necessary for transforming the start configuration into the target configuration. Here we further improve the lower bound to $\lfloor 5n/3 \rfloor - 1$, and thereby give a partial answer to one of their open problems.

(II) We show that the reconfiguration problem with congruent disks in the translation model is NP-hard, in both the labeled and unlabeled variants. This answers another open problem of Abellanas et al.

(III) We also show that the reconfiguration problem with congruent disks in the sliding model is NP-hard, in both the labeled and unlabeled variants.

(IV) For the reconfiguration with translations of n arbitrary labeled convex bodies in the plane, $2n$ moves are always sufficient and sometimes necessary.

Keywords: Disk reconfiguration, movable separability, translation model, sliding model, NP-hardness.

1 Introduction

A *body* (or *object*) in the plane is a compact connected set in $\mathbb{R}^2$ with nonempty interior. Two initially disjoint bodies *collide* if they share an interior point at some time during their motion. Consider a set of n pairwise interior-disjoint objects in the plane that need to be brought from a given start (initial) configuration S into a desired target (goal) configuration T, without causing collisions. The *reconfiguration* problem for such a system is that of computing a sequence of object motions (a schedule, or motion plan) that achieves this task. Depending on the existence of such a sequence of motions, we call that instance of the problem *feasible* and respectively, *infeasible*.

*An extended abstract of this paper appeared in the Proceedings of the 23rd International Symposium on Algorithms and Data Structures, (WADS 2009), Banff, August 2009, Vol. 5664 of LNCS, Springer, pp. 254–265.

[†]Department of Computer Science, University of Wisconsin–Milwaukee, WI 53201-0784, USA. Email: `ad@cs.uwm.edu`. Supported in part by NSF CAREER grant CCF-0444188 and NSF grant DMS-1001667.

[‡]Department of Computer Science, Utah State University, Logan, UT 84322-4205, USA. Email: `mjiang@cc.usu.edu`. Supported in part by NSF grant DBI-0743670.

Our reconfiguration problem is a simplified version of the multi-robot motion planning problem [16], in which a system of robots are operating together in a shared workplace and once in a while need to move from their initial positions to a set of target positions. The workspace is often assumed to extend throughout the entire plane, and has no obstacles other than the robots themselves. In many applications, the robots are indistinguishable (unlabeled), so each of them can occupy any of the specified target positions. Beside multi-robot motion planning, another application which permits the same abstraction is moving around large sets of heavy objects in a warehouse. Typically, one is interested in minimizing the number of moves and designing efficient algorithms for carrying out the motion plan. There are several types of moves that make sense to study, as dictated by specific applications. In this paper we focus on systems of convex bodies.

Next we formulate these models for systems of disks, since they are simpler and most of our results are for this class of convex bodies. These rules can be extended (not necessarily uniquely) to convex bodies in the plane. The decision problems we refer to below, pertaining to various reconfiguration problems we discuss here, are in standard form, and concern systems of (arbitrary or congruent) disks. For instance, the *Reconfiguration Problem* U-SLIDE-RP for congruent disks is: Given a start configuration and a target configuration, each with n unlabeled congruent disks in the plane, and a positive integer k, is there a reconfiguration motion plan with at most k sliding moves? It is worth clarifying that for the unlabeled variant, if the start and target configuration contain subsets of congruent disks, there is freedom is choosing which disks will occupy target positions. However in the labeled variant, this assignment is uniquely determined by the labeling; of course a valid labeling must respect the size of the disks.

1. *Sliding model*: one move is sliding a disk to another location in the plane without colliding with any other disk, where the disk center moves along an arbitrary continuous curve. This model was introduced in [3]. The labeled and unlabeled variants are L-SLIDE-RP and U-SLIDE-RP, respectively.
2. *Translation model*: one move is translating a disk to another location in the plane along a fixed direction without colliding with any other disk. This is a restriction imposed to the sliding model above for making each move as simple as possible. This model was introduced in [1]. The labeled and unlabeled variants are L-TRANS-RP and U-TRANS-RP, respectively.
3. *Lifting model*: one move is lifting a disk and placing it back in the plane anywhere in the free space, that is, at a location where it does not intersect (the interior of) any other disk. This model was introduced in [2]. The labeled and unlabeled variants are L-LIFT-RP and U-LIFT-RP, respectively. We don't present any results for this model here and only mention it for completeness.

In the main part of this paper (Sections 2, 3, and 4), we restrict ourselves to the case of disks in the plane, and study the reconfiguration problem in the translation model and the sliding model. In the last part of this paper (Sections 5 and 6), we study the reconfiguration problem in the translation model for arbitrary convex bodies in the plane (in particular, axis-parallel squares), and observe some interesting differences. Our main results are the following:

(I) For any n, there exist pairs of start and target configurations each with n congruent unlabeled disks, that require $\lfloor 5n/3 \rfloor - 1$ translation moves for reconfiguration (Theorem 1 in Section 2). This improves the previous bound of $\lfloor 8n/5 \rfloor$ due to Abellanas et al. and thereby gives a partial answer to their first open problem regarding the translation model [1].

(II) The reconfiguration problem with congruent disks in the translation model, in both the labeled and unlabeled variants, is NP-hard. That is, L-TRANS-RP and U-TRANS-RP are NP-hard (Theorem 2 and Theorem 3 in Section 3). This answers the second open problem of Abellanas et al. regarding the translation model [1].

(III) The reconfiguration problem with congruent disks in the sliding model, in both the labeled and unlabeled variants, is NP-hard. That is, L-SLIDE-RP and U-SLIDE-RP are NP-hard (Theorem 4 and Theorem 5 in Section 4).

(IV) For the reconfiguration with translations of n arbitrary labeled convex bodies in the plane, $2n$ moves are always sufficient and sometimes necessary (Theorem 6 in Section 5). For the special case of n axis-parallel unit squares, n translations suffice for moving them from the third quadrant to the first quadrant (Theorem 8 in Section 5).

For the class of disks, any instance the reconfiguration problem is always feasible in each of the three models [1, 2, 3, 4, 12, 14]. This follows essentially from the feasibility in the translation model. Our result for convex bodies in the plane (Theorem 6 in Section 5) implies that for this class of geometric objects, any instance of the reconfiguration problem is always feasible in the translation model, and hence in the sliding model and the lifting model too.

Related work. The movable separability of sets of geometry objects under various kinds of motions and various definitions of separations has been extensively studied in discrete and computational geometry [4, 6, 10, 12, 14, 17, 19, 20, 21]. See also the (older) survey by Toussaint [22] on this topic, and two other surveys on related topics [11, 16]. For instance, given a set of disjoint polygons in the plane, may each be moved "to infinity" (in some order) in a continuous motion in the plane without colliding with the others? Often constraints are imposed on the types of motions allowed, e.g., only translations, or only translations in a fixed set of directions. Sometimes only one object is permitted to move at a time, but this may not be enough to allow separation, thus a joint maneuver is required. Without the convexity assumption on the objects, one can show examples where the objects are interlocked and could only be moved "together" in the plane; however they could be easily separated one by one using the third dimension, i.e., in the lifting model.

For the movable separability problem in the plane, Sack and Toussaint [20] have shown how to compute all directions of movability for two simple polygons with non-intersecting interiors. Dehne and Sack [10] presented an algorithm for finding all directions of uni-directional and respectively multi-directional separability for sets of disjoint simple polygons; see also [15, 23]. The problem becomes more challenging in higher dimensions. Let $\mathcal{F}$ be a finite set of disjoint geometric objects in $\mathbb{R}^d$, $d \geq 2$. An object $B \in \mathcal{F}$ is *movable* if there exists a continuous rigid motion that moves B to infinity without colliding with the other objects in $\mathcal{F}$. The objects in $\mathcal{F}$ are *separable in direction* α if they can be separated by a sequence of collision-free translations all in the common direction α. Dawson [7] proved that in any set of n disjoint balls in $\mathbb{R}^d$, at least $\min\{n, d+1\}$ balls are sequentially movable (in possibly different directions). He also showed that any collection of star-shaped bodies in $\mathbb{R}^d$, $d \geq 2$, can be separated by n simultaneous translations, in n possibly different directions. A set of criteria for separability have been studied by the same author in [8, 9]. Fejes-Tóth and Heppes [12] have given a configuration of 13 interior-disjoint convex bodies in which no single body can move by any infinitesimal rigid motion. Snoeyink and Stolfi [21] gave a configuration of 12 tetrahedra with the same property. They also gave a configuration of 30 interior-disjoint convex bodies that cannot be taken apart with two hands, i.e., by using any two rigid motions on two complementary subsets.

2 A new lower bound for translating unlabeled congruent disks

In this section we consider the problem of moving n disks of unit radius, here also referred to as coins, to n target positions using translation moves. Abellanas et al. [1] have shown that $\lfloor 8n/5 \rfloor$ moves are sometimes necessary. Their lower bound construction is shown in Figure 1. Here we further improve this bound to $\lfloor 5n/3 \rfloor - 1$.

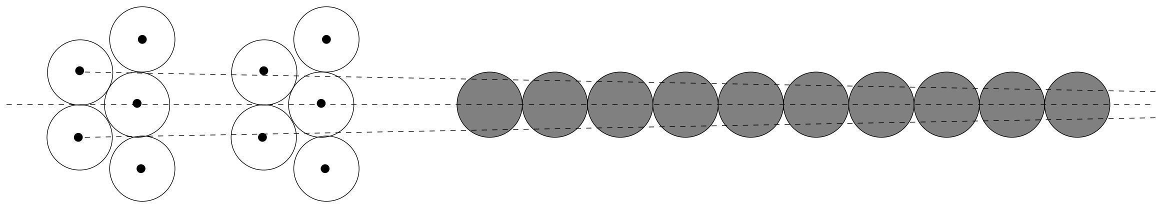

Figure 1: Two groups of five disks with their targets: part of the old $\lfloor 8n/5 \rfloor$ lower bound construction for translating disks. The disks are white and their targets are shaded.

Theorem 1. *For every $m \geq 1$, there exist pairs of start and target configurations each with $n = 3m + 2$ disks, that require $5m + 3$ translation moves for reconfiguration. Consequently, for any n, we have pairs of configurations that require $\lfloor 5n/3 \rfloor - 1$ translation moves.*

Proof. A move is a *target move* if it moves a disk to a final target position. Otherwise, it is a *non-target* move. We also say that a move is a *direct target move* if it moves a disk from its start position directly to its target position.

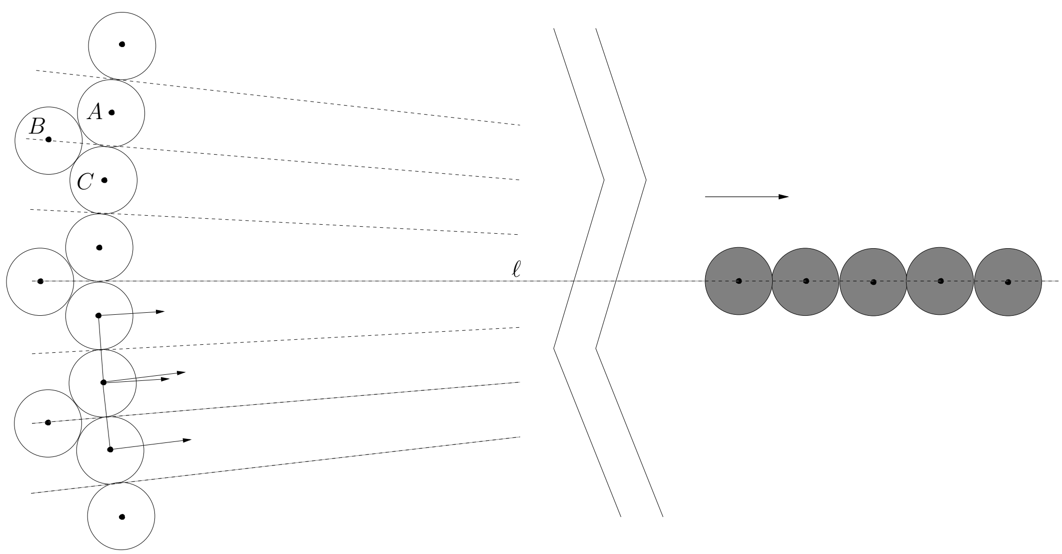

Figure 2: Illustration of the lower bound construction for translating congruent unlabeled disks, for $m = 3$, $n = 11$. The disks are white and their targets are shaded. Two consecutive partially overlapping parallel strips of width 2 are shown.

Let $n = 3m + 2$. The start and target configurations, each with n disks, are shown in Figure 2. The n target positions are all on a horizontal line ℓ, with the disks at these positions forming a horizontal chain, $T_1, \ldots, T_n$, consecutive disks being tangent to each other. Let o denote the center of the median disk, $T_{\lfloor n/2 \rfloor}$. Let $r > 0$ be very large. The start disks are placed on two very slightly convex chains (two concentric circular arcs):

- $2m+2$ disks in the first layer (chain). Their centers are $2m+2$ equidistant points on a circular arc of radius r centered at o.
- m disks in the second layer. Their centers are m equidistant points on a concentric circular arc of radius $r\cos\alpha + \sqrt{3}$. Each pair of consecutive points on the circle of radius r subtends an angle of 2α from the center of the circle (α is very small).

The parameters of the construction are chosen to satisfy: $\sin\alpha = 1/r$ and $2n \sin n\alpha \leq 2$. Set for instance $\alpha = 1/n^2$, which results in $r = \Theta(n^2)$.

Alternatively, the configuration can be viewed as consisting of m groups of three disks each, plus two disks, one at the higher and one at the lower end of the chain along the circle of radius r. Denote the three pairwise tangent start disks in a group by A, B and C, with their centers making an equilateral triangle, and the common tangent of A and C passing through o. Disks A and C are on the first layer, and the "blocking" disk B on the second layer. The groups are numbered from the top. We therefore refer to the three start disks of group i by A_i, B_i and C_i, where $i = 1, \ldots, m$.

For each pair of tangent disks on the first chain, consider the open strip of width 2 of parallel lines orthogonal to the line segment connecting their centers. By selecting r large enough we ensure the following crucial property of the construction: the intersection of all these $2m+1$ parallel strips contains the set of n centers of the targets in its interior. More precisely, let a be the center of T_1 and let b be the center of T_n. Then the closed segment ab of length $2n-2$ is contained in the intersection of all the $2m+1$ open parallel strips of width 2. Observe that for any pair of adjacent disks in the first layer, if both disks are still in their start position, neither can move so that its center lies in the interior of the strip of width 2 determined by their centers. As a consequence for each pair of tangent disks on the first chain at most one of the two disks can have a direct target move, provided its neighbor tangent disks have been already moved away from their start positions. See also Figure 3.

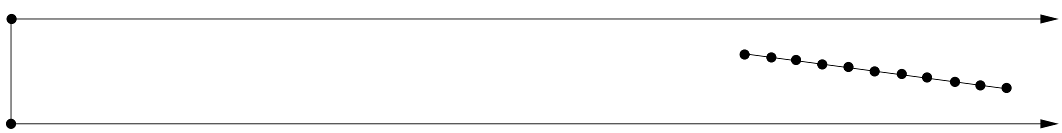

Figure 3: The n disks at the target positions as viewed from a parallel strip of a pair of start positions below the horizontal line ℓ in Figure 2. The targets are shown denser than they are: the chain of targets is in fact longer.

Recall that there are $2m+2$ disks in the first layer and m disks in the second layer. We show that the configuration requires at least $2m+1$ non-target moves, and consequently, at least $3m+2+2m+1 = 5m+3$ moves are required to complete the reconfiguration. Throughout the process let:

- k be the number of disks in the first layer that are in their start positions,
- c be the number of connected components in the intersection graph of *these* disks, i.e., disks in the first layer that are still in their start positions,
- x be the number of non-target moves executed so far.

Let t denote the number of moves executed. Consider the value $\Phi = k - c$ after each move. Initially, $k = 2m+2$ and $c = 1$, so the initial value of $k - c$ is $\Phi_0 = 2m+1$. In the end, $k = c = 0$, hence the final value of $k - c$ is $\Phi_t = 0$, and $x = x_t$ represents the total number of non-target moves executed for reconfiguration. Consider any reconfiguration schedule. It is enough to show that after any move that reduces the value of Φ by some amount, the value of x increases by at least the same amount. Since the reduction of Φ equals $2m+1$, it implies that $x \geq 2m+1$, as desired.

Observe first that a move of a coin in the second layer does not affect the values of k and c, and therefore leaves Φ unchanged. Consider now any move of a coin in the first layer, and examine the way Φ and x are modified as a result of this move and possibly some preceding moves. The argument is essentially a charging scheme that converts the reduction in the value of Φ into non-target moves.

Case 0. If the moved coin is the only member of its component, then k and c decrease both by 1, so the value of Φ is unchanged.

Assume now that the moved coin is from a component of size at least two (in the first layer). We distinguish two cases:

Case 1. The coin is an *end* coin, i.e., one of the two coins at the upper or the lower end of the component (chain) in the current step. Then k decreases by 1 and c is unchanged, thus Φ decreases by 1. By the property of the construction, the current move is a non-target move, thus x increases by 1.

Case 2. The coin is a *middle* coin, i.e., any other coin in the first layer in its start position that is not an end coin in the current step. By the property of the construction, this is necessarily a non-target move (from a component of size at least 3). As a result, k decreases by 1 and c increases by 1, thus Φ decreases by 2. Before the middle coin (A_i or C_i) can be moved by the non-target move, its blocking coin B_i in the second layer must have been moved by a previous non-target move. Observe that this previous non-target move is uniquely assigned to the current non-target move of the middle coin, because the middle coins of different moves cannot be adjacent! Indeed, as soon a middle coin is moved, it breaks up the connect component, and its two adjacent coins cannot become middle coins in subsequent moves. We have therefore found two non-target moves uniquely assigned to this middle coin move, which contribute an increase by 2 to the value of x.

This exhausts the possible cases, and thereby completes the analysis. The lower bounds for values of n other than $3m+2$ are immediately obtainable from the above: for $n = 3m$, at least $5m-1$ moves are needed, while for $n = 3m+1$, at least $5m+1$ moves are needed. This completes the proof of the theorem. □

Remarks. The choice of target positions in a straight chain is only for simplicity. Any configuration of target positions contained in an axis-parallel rectangle R of width $2n$ and height smaller than 2 (say, 1.9) works in the same way, since we can demand that R is contained in the intersection of all parallel open strips of width 2. Moreover the start configuration of disks can be modified so that no two disks touch: Take the current construction and slightly shrink each (start and target) disk with respect to its center. If the reduction is small enough, a similar argument can be made that from two consecutive start disks, at most one allows a direct target move. Even if the disks in the second layer are not touching the disks in the first layer, the blocking effect is still achieved.

3 Hardness results for translating congruent disks

Theorem 2. *The unlabeled version of the disk reconfiguration problem with translations* U-TRANS-RP *is NP-hard even for congruent disks.*

Proof. Here we adapt for our purpose the reduction in [5] showing that the reconfiguration problem with unlabeled chips in graphs is NP-complete. We reduce 3-SET-COVER to U-TRANS-RP. The problem 3-SET-COVER is a restricted variant of SET-COVER. An instance of SET-COVER consists of a family $\mathcal{F}$ of subsets of a finite set U, and a positive integer k. The problem is to decide whether there is a set cover of size k for $\mathcal{F}$, i.e., a subset $\mathcal{F}' \subseteq \mathcal{F}$, with $|\mathcal{F}'| \leq k$, such that every element in U belongs to at least one member of $\mathcal{F}'$. In the variant 3-SET-COVER the size of each set in $\mathcal{F}$ is bounded from above by 3. Both the standard and the restricted variants are known to be NP-hard [13].

Consider an instance of 3-SET-COVER represented by a bipartite graph $(B \cup C, E)$, where $B = \mathcal{F}$, $C = U$, and E describes the membership relation. First construct a "broom" graph G with vertex set $A \cup B \cup C$, where $|A| = |C|$, as shown in Figure 4. Place a start (unlabeled) chip at each element of $A \cup B$, and let each element of $B \cup C$ be a target position. A move in the graph is defined as shifting a chip from v_1 to v_2 ($v_1, v_2 \in V(G)$) along a "free" path in G, so that no intermediate vertices are occupied; see also [5]. Positions in B are called *obstacles*, since any obstacle position that becomes free during reconfiguration must be finally filled by one of the chips. Write $m = |\mathcal{F}|$, and $n = |U|$. Then $|B| = m$ and $|A| = |C| = n$.

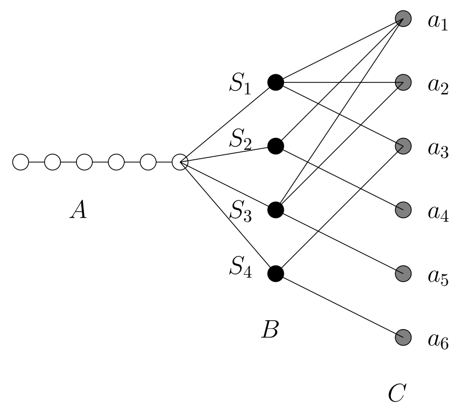

Figure 4: The "broom" graph G corresponding to a 3-SET-COVER instance with $m = |B| = |\mathcal{F}| = 4$ and $n = |A| = |C| = |U| = 6$. Chips-only: white; obstacles: black; target-only: gray.

Now construct a start and target configuration, each with $O((m+n)^8)$ disks, that represents G in a suitable way. The start positions are (correspond to) $S = A \cup B$ and the target positions are $T = B \cup C$. The positions in B are also called *obstacles*, since any obstacle position that becomes free during reconfiguration must be finally filled by one of the disks. Let z be a parameter used in the construction, to be set later large enough. Consider an axis-parallel rectangle R of width $2z \cdot \max\{m+1, n\}$ and height $z \cdot \max\{m+1, n\}$. Figure 5 shows a scaled-down version for a smaller value of z ($z = 10$). Initially place an obstacle disk centered at each grid point in R. The obstacle chips in B from the graph G are represented by m obstacles disks, denoted by $S_1, \ldots, S_m$ (the m sets), whose centers are on on the top side of R at distances $2z, 4z, 6z, \ldots$ from the left side of R. Next we (i) delete some of the obstacle disks in R, (ii) add a set of target-only disks in n connected areas (legs) below R, and (iii) change the positions of some of the obstacle disks in B, as described below:

(i) Consider the obstacles whose centers are on on the bottom side of R at distances $z, 3z, 5z, \ldots$ from the left side of R. Let these be denoted by $a_1, \ldots, a_n$ in Figure 5 and Figure 6. For each edge $S_i a_j$ in the bipartite graph $(B \cup C, E)$, consider the convex hull H_{ij} of the two disks S_i and a_j; see Figure 6 (middle). We refer to these H_{ij}s as *roads*. Delete now from R any obstacle disk D, except the disks S_i, that intersects some H_{ij} in its interior (the disks $a_1, \ldots, a_n$ are also deleted).

(ii) The target-only chips in C from the graph G are represented by n^2 target-only disks located

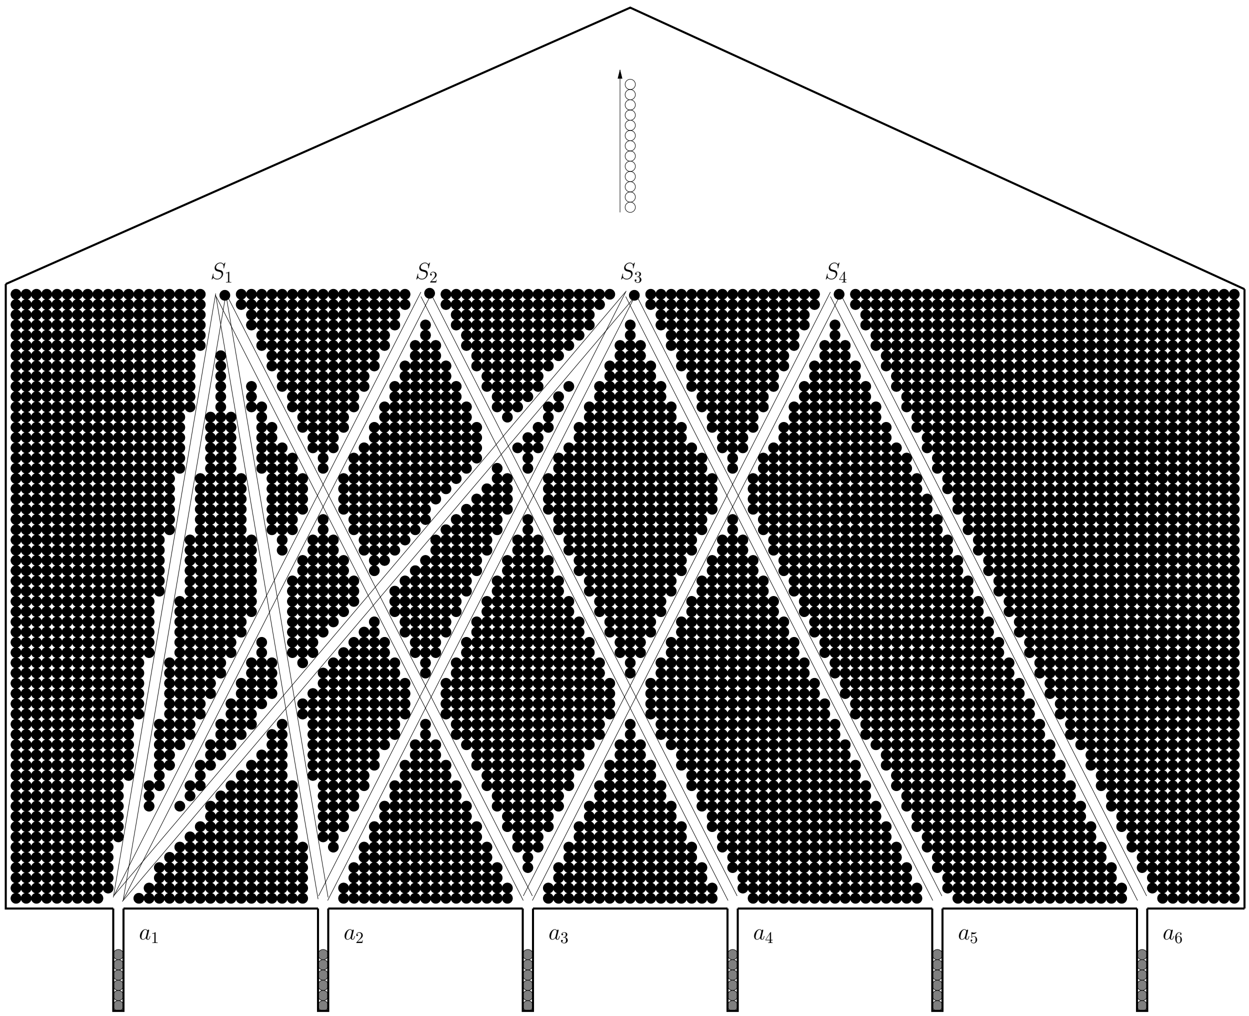

Figure 5: Reduction from 3-SET-COVER to U-TRANS-RP implementing a disk realization of the "broom" graph G in Figure 4 ($z = 10$). $|A| = |C| = n = 6$, and $|B| = m = 4$. The start-only disks are white, the obstacle disks are black, and the target-only disks are gray. Obstacle disks on the side of the roads are pushed tangent to the roads (not shown here). The pentagon with 6 legs that encloses the main part of the construction is enclosed by another thick layer of obstacles. Only 13 out of the 36 start-only disks are shown. Note the displaced final obstacle positions for the disks S_i; detail in Figure 6 (left). An optimal reconfiguration takes $3 \times 36 + 3 = 111$ translation moves (S_2, S_3, and S_4 form an optimal set cover).

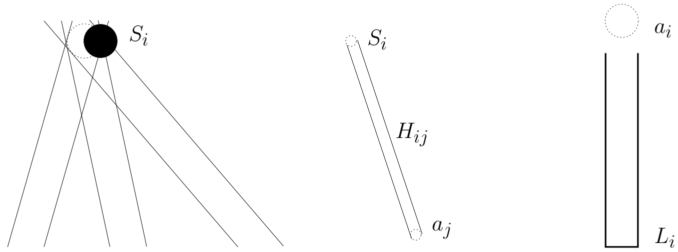

Figure 6: Left: the actual position of the disk S_i (black) partially blocks the three outgoing roads H_{ij} from S_i. The white dotted disk (not part of the construction) is contained in the intersection of the three incident roads H_{ij}. Middle: a road H_{ij}. Right: a white dotted disk (not part of the construction) representing a_i. This element is used in road construction, for the placement of the incident incoming roads. The corresponding leg is L_i.

in n connected areas (legs) extending rectangle R from below. Each leg is a thin rectangle of unit width. These legs extend vertically below the bottom side of R at distances $z, 3z, 5z, \ldots$ from the left side of R, exactly below the obstacles $a_1, \ldots, a_n$ that were previously deleted. Let these legs be denoted by $L_1, \ldots, L_n$ in Figure 6. In each of the legs we place n target-only disk positions. Since each leg is vertical, its vertical axis is not parallel to any of the road directions H_{ij}.

(iii) For each road H_{ij}, push the disk obstacles next to the road sides closer to the roads, so that they become tangent to the road sides. Displace now each of the obstacle disks S_i to a nearby position that partially blocks any of the (at most 3) roads H_{ij} incident to S_i. See Figure 6. The new obstacle positions prohibit any S_i to reach any a_j position in one translation move. This special position of these obstacles is important, as the reduction wouldn't work otherwise (that is, if the obstacle would be placed at the intersection of the outgoing roads), at least not in this way.

The start-only disks in A form a vertical chain of n^2 disks placed on the vertical line ℓ, which is a vertical symmetry axis of R. The position of start disks is such that no start disk is on any of the road directions H_{ij}. Finally enclose all the above start-only disks, obstacle disks, and target-only disks by a closed pentagonal shaped chain with n legs of tangent (or near tangent) obstacle disks, as shown in Figure 5. Surround all the above construction by another thick layer of obstacle disks; a thickness of z will suffice. This concludes the description of the reduction. Clearly, G and the corresponding disk configuration can be constructed in polynomial time. The reduction is complete once we establish the following claim.

Claim. There is a set cover consisting of at most q sets if and only if the disk reconfiguration can be done using at most $3n^2 + q$ translations.

Proof of Claim: Let $z = \Omega((m+n)^3)$. By the choice of z, the construction has the property that even after any set of at most $3n^2 + n$ obstacle disks (in R) are removed, for filling a target position, either with one of the obstacles S_i or with a start disk, still requires at least 3 translation moves. Observe that a disk S_i could be moved to one of the outgoing road directions, but then, this move is part of the 3 necessary moves previously mentioned. Let $q \leq n$ be the size of an optimal set cover. The direct implication is clear from the following motion plan:

1. Translate the disks S_i in B corresponding to an optimal cover to nearby positions on one of the roads implementing the (at most three) edges adjacent to S_i in G.
2. Move these q disks to suitable target positions in the n legs, using another 2 translation moves for each of them.
3. Move $n^2 - q$ start disks to fill the remaining $n^2 - q$ target positions below R, using 3 translation moves for each of them.
4. Fill the emptied obstacle positions among the S_i (corresponding to an optimal cover) with the remaining q start disks, using one translation move for each of them. The total number of moves is $3q + 3(n^2 - q) + q = 3n^2 + q$, as required.

For the converse implication, assume that there is a reconfiguration sequence M with fewer than $3n^2 + q$ translation moves. Observe that a disk S_i needs at least 3 translation moves to occupy a target position in any of the legs: one move to get to a road direction, and two more to finish. The same holds for disks on the vertical line ℓ. Now, if at least q obstacle disks S_i are moved out of their positions, since filling any leg-target requires at least 3 translation moves, and obstacles need to be filled back in, it follows that at least $3n^2 + q$ translation moves must be executed. We can therefore assume that strictly fewer than q obstacle disks S_i are moved out of their positions (and also that this number is not zero) in this sequence. Then at least n target-only elements (in some leg) require 4 translations for filling their positions. Since any other target-only element requires at least 3 translations for filling its position we conclude that overall at least $3n^2 + n$ translation moves are required in M. On the other hand, we have $3n^2 + q \leq 3n^2 + n$, which contradicts our initial assumption. □

A similar reduction can be made for the labeled version by adapting the idea used in [5] to show that the labeled version for reconfiguration of chips in graphs is NP-hard.

Theorem 3. *The labeled version of the disk reconfiguration problem with translations* L-TRANS-RP *is NP-hard even for congruent disks.*

Proof. We reduce 3-SET-COVER to L-TRANS-RP. We use a similar reduction as that shown in Figure 5. The labels are so that all the obstacles in R (including $S_1, \dots, S_m$) must remain in the same position. (The special displaced position of the obstacles S_i is not important, the initial positions are also fine.) There are $2n^2$ start disks, labeled 1 to $2n^2$ from top down, and placed on the vertical line ℓ as in Figure 5. There are n legs placed as in the previous construction. Each leg has $2n$ target positions. The targets in the ith leg L_i, $i = 1, \dots, n$, are labeled (top-down) from $(i-1) \cdot 2n + 1$ to $i \cdot 2n$.

The reduction follows from the next claim. Its proof is similar to that for the previous claim, using the fact that an obstacle S_i representing a selected set must move twice, once to clear the roads, and once to come back.

Claim. There is a set cover consisting of at most q sets if and only if the disk reconfiguration can be done using at most $6n^2 + 2q$ translations.

Proof of Claim: An optimal motion plan is: (1) Translate the disks S_i in B corresponding to an optimal cover to temporary positions above R. (2) Move all $2n^2$ start disks to their targets, using 3 translation moves for each of them. (3) Bring back the obstacle disks S_i from their temporary positions. The total number of moves is $q + 6n^2 + q = 6n^2 + 2q$, as required.

For the converse implication, assume that there is a reconfiguration sequence with fewer than $6n^2 + 2q$ translation moves. If at least q obstacle disks S_i are moved out of their positions, since filling any target requires at least 3 translation moves, and obstacles have to be filled back in, it follows that at least $6n^2 + 2q$ translation moves must be executed. Assume now that strictly fewer than q obstacle disks S_i are moved out of their positions. Then at least $2n$ target-only elements (in some leg) require 4 translations for filling their positions, hence overall at least $6n^2 + 2n$ translation moves are required for reconfiguration. On the other hand, we have $6n^2 + 2q \leq 6n^2 + 2n$, which contradicts our initial assumption. This concludes the proof of Theorem 3. □

4 Hardness results for sliding congruent disks

We start with the unlabeled variant, and adapt for our purpose the reduction in [5] showing that the reconfiguration problem with unlabeled chips in an infinite grid is NP-complete. We reduce the *Rectilinear Steiner Tree* problem R-STEINER to U-SLIDE-RP. An instance of R-STEINER consists of a set S of n points in the plane, and a positive integer bound q. The problem is to decide whether there is a rectilinear Steiner tree (RST), that is, a tree with only horizontal and vertical edges that includes all the points in S, along with possibly some extra *Steiner points*, of total length at most q. For convenience the points can be chosen with integer coordinates. R-STEINER is known to be Strongly NP-complete [13], so we can assume that all point coordinates are given in unary.

Theorem 4. *The unlabeled version of the sliding disks reconfiguration problem in the plane* U-SLIDE-RP *is NP-hard even for congruent disks.*

Proof. Assume the disks have unit diameter. In our construction each (start and target) disk will be centered at an integer grid point. A disk position (i.e., the center of a disk) that is both a start and a target position is called an *obstacle*. We have four types of grid points: free positions, start positions, target positions, and obstacles.

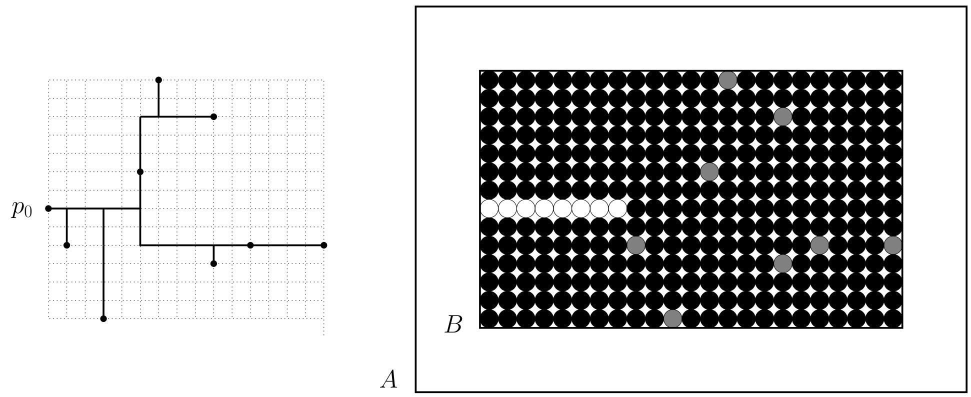

Figure 7: Left: an instance of R-STEINER with $n = 9$ points, and a rectilinear Steiner tree for it. Right: the configuration of start positions (white), target positions (gray), and obstacle positions (black).

Consider an instance $P = \{p_0, p_1, \ldots, p_{n-1}\}$ of R-STEINER with n points. Assume that $p_0 = (0, 0)$ is a leftmost point in P, see Figure 7 (left). The instance of U-SLIDE-RP is illustrated in Figure 7 (right). Choose $n-1$ start positions with zero y-coordinate and x-coordinates $0, -1, -2, \ldots, -(n-2)$, i.e., in a straight horizontal chain extending from p_0 to the left. Choose $n-1$ target positions at the remaining $n-1$ points $\{p_1, \ldots, p_{n-1}\}$ of the R-STEINER instance. Let B be a smallest axis-parallel rectangle containing the $2n-2$ disks at the start and target positions, and Δ be the length of the longer side of B. Consider a sufficiently large axis-parallel rectangle A enclosing B: the boundary of A is at distance $2n\Delta$ from the boundary of B. Place obstacle disks centered at each of the remaining grid points in the rectangle A. The number of disks in the start and target configurations is $O(n^2\Delta^2)$. This construction is done in time polynomial in Δ, which is polynomial in the size of the R-STEINER instance since the coordinates are given in unary. The reduction is complete once we establish the following claim.

Claim. There is a rectilinear Steiner tree of length at most q for P if and only if the disk reconfiguration can be done using at most q sliding moves.

Proof of Claim: We start with the direct implication. Let T be a Steiner tree of length q connecting all points in P (p_0 and the $n-1$ target positions). We show that there is a reconfiguration sequence M of q moves. Pick a leaf of T which is a target position, and fill it with a disk (obstacle or start disk) that is closest to it in T by sliding it along the corresponding collision-free path in T. If the disk comes from an obstacle, a new target-only position results in place of the obstacle. After this move, the length of the resulting tree connecting all target-only positions and p_0 is one less than the length of T. Continue by repeating this step until all targets have been filled (note that there is a leaf target-only position at each step); it follows that q moves are performed in total.

We continue with the converse implication. Let M be a reconfiguration sequence of q moves, where we can also assume that q is minimum. We show that there is a rectilinear Steiner tree of length at most q for P. Since the disks are centered at grid points in both start and target configurations, the center of each disk moved by the sequence M traces a continuous curve between two grid points (the two grid points may coincide, and the curve may have zero length). Let $s_1, \dots, s_{n-1}$ be the $n-1$ start positions. Then there exist $n-1$ continuous curves C_i connecting each start position s_i to a distinct target position t_i, for each $1 \leq i \leq n-1$, and for a suitable permutation $t_1, \dots, t_{n-1}$ of the target positions. Each curve C_i connects two different grid points s_i and t_i, and might be the concatenation of multiple continuous curves traced by the centers of different disks. Since $q \leq 2n\Delta$, we can assume that the curves C_i are in the rectangle A. Let S_i be the set of grid points consisting of the start position s_i, the target position t_i, and the centers of the subset of disks in the construction that must be moved because of the reconfigurations along the curve C_i. We have the following sub-claim:

Sub-Claim. For each i, $1 \leq i \leq n-1$, there is rectilinear path P_i from s_i to t_i consisting of consecutively adjacent grid points in S_i.

To verify this, we proceed as follows. Increase the radius of each disk in the construction from 1/2 to 1. By a standard Minkowski-sum type argument, the set S_i includes the centers of all enlarged disks that intersect the curve C_i. Observe that every grid point in the rectangle B is covered by exactly one enlarged disk, and every non-grid point in B is covered by at least 2 and at most 4 enlarged disks. Here covered means covered by the interior of the enlarged disk. For each point a on the curve C_i, denote by $S_i(a)$ the subset of (at least 1 and at most 4) grid points in S_i whose corresponding disks cover a. We have the following two properties: (1) for any point a on the curve C_i, the grid points in $S_i(a)$ are connected in the grid graph; (2) for any two points a and b of distance less than $2 - \sqrt{2}$ on the curve C_i, the intersection $S_i(a) \cap S_i(b)$ is not empty: observe that the intersection of the two disks of radius 1 centered at a and b contains a disk of diameter more than $\sqrt{2}$, which properly contains a unit square.

Now subdivide the curve C_i into subcurves of length less than $2 - \sqrt{2}$ by a sequence of k points $a_1, \dots, a_k$, where $a_1 = s_i$ and $a_k = t_i$. By property (2), there exists a sequence of $k+1$ grid points $b_0, b_1, \dots, b_k$ such that $b_0 = s_i$, $b_j \in S_i(a_j) \cap S_i(a_{j+1})$ for $1 \leq j \leq k-1$, and $b_k = t_i$. By property (1), for each pair of consecutive points b_{j-1} and b_j in the sequence, $1 \leq j \leq k$, there is a rectilinear path of consecutively adjacent grid points in $S_i(a_j)$ from b_{j-1} to b_j. The concatenation of these rectilinear paths forms a rectilinear path P_i from s_i to t_i consisting of consecutively adjacent grid points in S_i.

To complete the proof of the claim, we apply an argument similar to that in [5, Theorem 5]. Consider the set U of grid points in the union of the $n-1$ rectilinear paths P_i. Each disk centered at one of the grid points in $U \setminus \{t_1, \dots, t_{n-1}\}$ must be moved at least once during the reconfiguration. Therefore the total number of grid points in U is at most $q+n-1$. These grid points, in particular the start positions s_i and the target positions t_i, are connected in the grid graph. Since the $n-1$ start positions s_i are connected to each other without going through the other grid points, we can modify the set U into a smaller set U' of at most $(q+n-1)-(n-2) = q+1$ grid points connecting the right-most start position p_0, and the $n-1$ target positions $p_1, \dots, p_{n-1}$, that is, the points in P. It follows that the points in P can be connected by a grid graph of length at most q, hence there is a rectilinear Steiner tree of length at most q for P. This proves the claim and also concludes the proof of Theorem 4. □

A similar reduction can be made for the labeled version by adapting the idea used in [5] to show that the labeled version for reconfiguration of chips in grids is NP-hard.

Theorem 5. *The labeled version of the sliding disks reconfiguration problem in the plane* L-SLIDE-RP *is NP-hard even for congruent disks.*

Proof. The proof again uses a reduction from R-STEINER. The argument and the construction is similar to that in the proof of our Theorem 4. Consider an instance $P = \{p_0, p_1, \dots, p_{n-1}\}$ of

R-STEINER with n grid points and let p_0 be a left-most point in P. The instance of L-SLIDE-RP is illustrated in Figure 8. Choose the points in P as the n target positions labeled $1, \ldots, n$

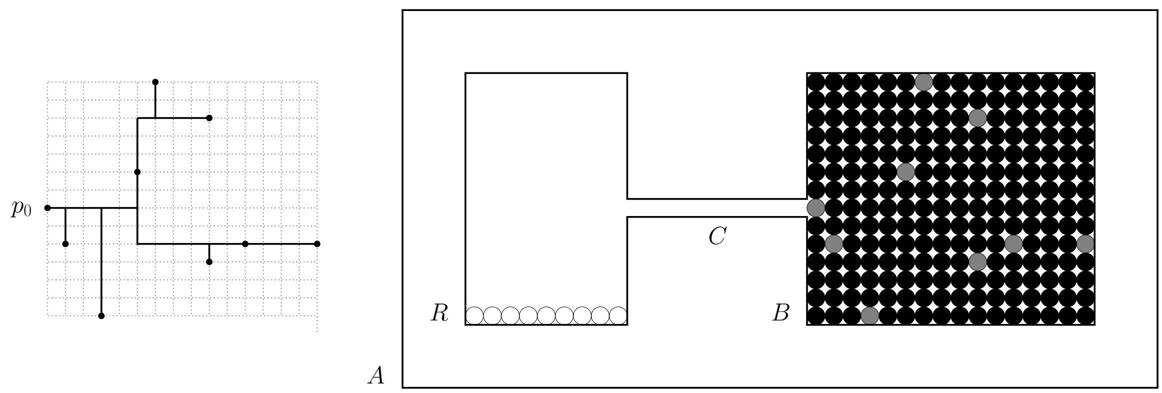

Figure 8: Left: an instance of R-STEINER with $n = 9$ points, and a rectilinear Steiner tree for it. Right: the configuration of start positions (white), target positions (gray), and obstacle positions (black).

in some arbitrary way. Let B be the smallest axis-parallel rectangle containing the n disks at the target positions. Let R be a sufficiently large axis-parallel rectangle containing n disks at the start positions, labeled $1, \ldots, n$ also in some arbitrary way. R is located to the left of B, and is connected to p_0 by a sufficiently long horizontal corridor C of unit width. Let A be a sufficiently large axis-parallel rectangle enclosing R, C, and B. Place an obstacle at each other grid point inside B but not in P, and at each grid point inside A but outside R, C, and B. The start and target labels are the same for each obstacle (so that each obstacle needs to fill in its original position at the end).

Informally, the $q + 1 - n$ obstacles that form a minimum rectilinear Steiner tree of length q for the n points must go out of B into the free area R to make space for reconfiguration, and then come back in to fill their original positions. More precisely, the reduction is complete once we establish the following claim, which is analogous to the claim in the proof of Theorem 4.

Claim. There is a rectilinear Steiner tree of length at most q for P if and only if the disk reconfiguration can be done using at most $2(q + 1 - n) + n = 2q + 2 - n$ sliding moves.

This concludes the proof of Theorem 5. □

5 Translating convex bodies

In this section we consider the general problem of reconfiguration of convex bodies with translations. When the convex bodies have different shapes, sizes, and orientations, assume that the correspondence between the start positions $\{S_1, \ldots, S_n\}$ and the target positions $\{T_1, \ldots, T_n\}$, where T_i is a translated copy of S_i, is given explicitly. Refer to Figure 9. In other words, we deal with the labeled variant of the problem. Our result can be easily extended to the unlabeled variant by first computing a valid correspondence by shape matching.

We first extend the $2n$ upper bound for translating arbitrary disks to arbitrary convex bodies in the plane:

Theorem 6. *For the reconfiguration with translations of n labeled disjoint convex bodies in the plane, $2n$ moves are always sufficient and sometimes necessary.*

Our method for moving convex bodies resembles the method by Abellanas et al. [1] for moving disks. The only difference is that our moves follow a more elaborate order than the simple

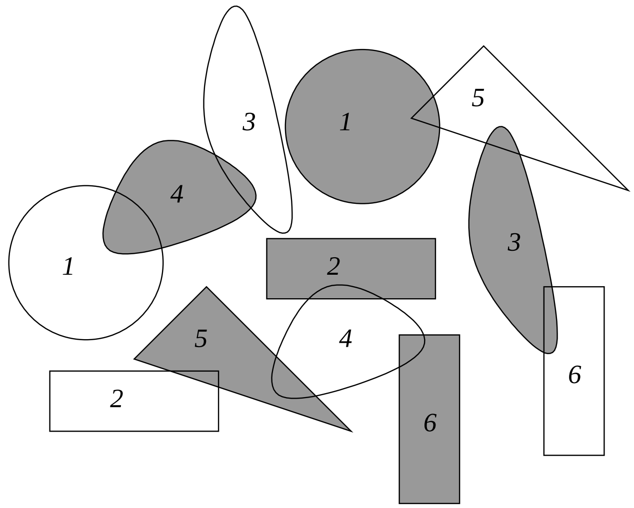

Figure 9: Reconfiguration of convex bodies with translations; the target positions are shaded.

lexical order by coordinates. Our proof of Theorem 6 uses a classical result on the uni-directional separability of disjoint convex sets [14]; see also [19, Theorem 8.7.2]:

Lemma 1 (Guibas and Yao [14])**.** *For any set of n disjoint convex bodies in the plane, and for any direction α, there is an ordering $C_1, \ldots, C_n$ of the convex bodies such that C_i can be moved continuously to infinity in the direction α without colliding with the convex bodies C_j, $1 \le i < j \le n$.*

A *common supporting line* of a set of convex bodies is a line tangent to at least two bodies; a direction is *special* if it is parallel to a common supporting line. The following lemma shows that, by slightly restricting α to non-special directions, we can obtain more freedom in moving the convex bodies:

Lemma 2. *For any set of n disjoint convex bodies in the plane, and for any non-special direction α, there exist a positive angle ϵ and an ordering $C_1, \ldots, C_n$ of the convex bodies such that C_i can be moved continuously to infinity in any direction at most an angle of ϵ away from α without colliding with the convex bodies C_j, $1 \le i < j \le n$.*

Proof. By Lemma 1, there is an ordering $C_1, \ldots, C_n$ of the convex bodies such that C_i can be moved continuously to infinity in the direction α without colliding with the convex bodies C_j, $1 \le i < j \le n$. Let ϵ_i be the smallest angle between the non-special direction α and a special direction parallel to a common supporting line tangent to C_i. Then C_i can be moved continuously to infinity in any direction at most an angle of ϵ away from α without colliding with the convex bodies C_j, $1 \le i < j \le n$. Setting $\epsilon = \min_i \epsilon_i$ completes the proof. □

Given n disjoint convex bodies in the plane, the number of special directions is at most $4 \cdot \binom{n}{2} = O(n^2)$ because each pair of convex bodies determines at most four common supporting lines. Thus in Lemma 2 we can choose the non-special direction α such that $\epsilon = \Omega(1/n^2)$.

We now proceed with the proof of Theorem 6. The lower bound follows easily for instance from the construction with disks of arbitrary radii given in [1]. We now prove the upper bound. Let α be a direction not parallel to any common supporting line of the $2n$ start and target positions of the n convex bodies. Let ℓ be a line perpendicular to the direction α. When ℓ is placed sufficiently far away from the start and target positions in the non-special direction α, it follows by Lemma 2 that, using the freedom of ϵ in moving the convex bodies, we can translate the convex bodies in a certain order, from the start positions directly to some intermediate positions that intersect ℓ, without causing any collision, and we can further restrict the intermediate positions to any desired order along ℓ. Then, by a symmetric argument, we can translate the convex bodies back, from the intermediate positions to the target positions, in a certain order. This completes the proof of Theorem 6.

5.1 Translating unlabeled axis-parallel unit squares

Throughout a translation move, the moving square remains axis-parallel, however the move can be in any direction. We have the following bounds:

Theorem 7. *For the reconfiguration with translations of n unlabeled axis-parallel unit squares in the plane, $2n-1$ moves are always sufficient, and $\lfloor 3n/2 \rfloor$ moves are sometimes necessary.*

Proof. We start with the upper bound. We adapt the general method for translating convex bodies. Obtain an ordering $S_1, \ldots, S_n$ of the start positions and an ordering $T_1, \ldots, T_n$ of the target positions. First move $n-1$ squares, from the start positions $S_1, \ldots, S_{n-1}$ to intermediate positions far away. Next move a square *directly* from S_n to T_n. Then move the $n-1$ squares back, from the intermediate positions to the target position $T_{n-1}, \ldots, T_1$.

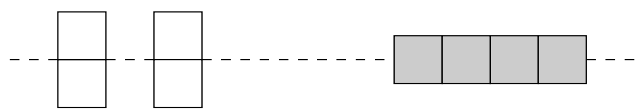

Figure 10: A lower bound of $\lfloor 3n/2 \rfloor$ for translating axis-parallel unit squares. The start positions (grouped in pairs) are tangent to the x-axis, which intersects the target positions (shaded). Each of the target squares is symmetric about the x-axis.

The construction in Figure 10 gives a lower bound of $\lfloor 3n/2 \rfloor$. ☐

Next we prove a better upper bound for a special case. Fix a Cartesian coordinate system, and consider a set n of axis-parallel rectangles. We need the following result of Guibas and Yao [14]:

Lemma 3 (Guibas and Yao [14])**.** *For any set of n disjoint axis-parallel rectangles in the plane, there is an ordering $R_1, \ldots, R_n$ of the rectangles such that R_i can be moved continuously to infinity in any direction between 0 and $\pi/2$ without colliding with the rectangles R_j, $1 \le i < j \le n$.*

Using Lemma 3, we establish a tight upper bound for a special case of reconfiguration, namely from the third quadrant to the first quadrant:

Theorem 8. *For the reconfiguration with translations of n unlabeled axis-parallel unit squares in the plane from the third quadrant to the first quadrant, n moves are always sufficient and sometimes necessary.*

Proof. Obtain an ordering $S_1, \ldots, S_n$ of the start positions in the third quadrant by Lemma 3. Symmetrically, by a rotation of angle π about the origin, obtain an ordering $T_1, \ldots, T_n$ of the target positions in the first quadrant. Then, for $i = 1, \ldots, n$, move the square from S_i to T_{n-i+1}. The tightness of the bound is trivial. ☐

6 Concluding remarks

In this paper we have obtained several results regarding the complexity of reconfiguration in two natural models: sliding and translation. We conclude with some observations and open problems.

Our upper bound of $2n - 1$ for translating unlabeled axis-parallel unit squares (Theorem 7) is the same as the current best upper bound for translating unlabeled unit disks [1]. However, the result in Theorem 8 for squares does not hold for disks: observe that any of the three constructions for the lower bounds of $\lfloor 3n/2 \rfloor$, $\lfloor 8n/5 \rfloor$, or $\lfloor 5n/3 \rfloor - 1$ can be placed with the start disks in the third quadrant and with the target disks in the first quadrant. It remains a challenging problem whether a $(2 - \delta)n$ upper bound is possible for translating n unlabeled axis-parallel unit squares or unit disks, where δ is a positive constant.

The type of construction in Figure 10 has been used previously for disks to obtain the first lower bound of $\lfloor 3n/2 \rfloor$ for translating unit disks [18]. It is interesting to note that neither of the two subsequent improved constructions, $\lfloor 8n/5 \rfloor$ of Abellanas et al. [1], or ours $\lfloor 5n/3 \rfloor - 1$ in Theorem 1, do not seem to work for squares.

A remaining question is: What is the computational complexity of the reconfiguration problem in the lifting model? Are L-LIFT-RP and U-LIFT-RP NP-hard for unit disks?

References

[1] M. Abellanas, S. Bereg, F. Hurtado, A. G. Olaverri, D. Rappaport, and J. Tejel: Moving coins, *Computational Geometry: Theory and Applications*, **34** (2006), 35–48.

[2] S. Bereg and A. Dumitrescu: The lifting model for reconfiguration, *Discrete & Computational Geometry*, **35** (2006), 653–669.

[3] S. Bereg, A. Dumitrescu, and J. Pach: Sliding disks in the plane, *International Journal of Computational Geometry & Applications*, **15** (2008), 373–387.

[4] N. G. de Bruijn: Aufgaben 17 and 18 (in Dutch), *Nieuw Archief voor Wiskunde*, **2** (1954), 67.

[5] G. Călinescu, A. Dumitrescu, and J. Pach: Reconfigurations in graphs and grids, *SIAM Journal on Discrete Mathematics*, **22** (2008), 124–138.

[6] B. Chazelle, T. A. Ottmann, E. Soisalon-Soininen, and D. Wood: The complexity and decidability of separation, in *Proceedings of the 11th International Colloquium on Automata, Languages and Programming*, LNCS 172, Springer, 1984, pp. 119–127.

[7] R. J. Dawson: On removing a ball without disturbing the others, *Mathematics Magazine*, **57** (1984), 27–30.

[8] R. J. Dawson: On the mobility of bodies in $\mathbb{R}^n$, *Math. Proc. Camb. Phil. Soc.*, **98** (1985), 403–412.

[9] R. J. Dawson: Corrigenda volume 98 (1985), 403–412, 'On the mobility of bodies in $\mathbb{R}^n$', *Math. Proc. Camb. Phil. Soc.*, **99** (1986), 377–379.

[10] F. Dehne and J. Sack: Translation separability of sets of polygons, *The Visual Computer*, **3** (1987), 227–235.

[11] A. Dumitrescu: Mover problems, in *Thirty Essays in Geometric Graph Theory*, J. Pach (editor), Springer, New York, August 2012, to appear.

[12] L. Fejes-Tóth and A. Heppes: Über stabile Körpersysteme (in German), *Compositio Mathematica*, **15** (1963), 119–126.

[13] M. Garey and D. Johnson: *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman and Company, 1979.

[14] L. Guibas and F. F. Yao: On translating a set of rectangles, in *Computational Geometry*, F. Preparata (ed.), pp. 61–67, Vol. 1 of *Advances in Computing Research*, JAI Press, London, 1983.

[15] C. Haldar and L. M. Patnaik: On movable separability and isotheticity, *Information Sciences*, **62** (1992), 87–102.

[16] Y. Hwang and N. Ahuja: Gross motion planning — a survey, *ACM Computing Surveys*, **25** (1992), 219–291.

[17] T. Ottmann and P. Widmayer: On translating a set of line segments, *Computer Vision, Graphics, and Image Processing*, **24** (1983), 382–389.

[18] D. Rappaport: Communication at the *16th Canadian Conference on Computational Geometry*, August 2004.

[19] J. O'Rourke: *Computational Geometry in C*, Cambridge University Press, 2nd edition, 1998.

[20] J.-R. Sack and G. T. Toussaint: Translating polygons in the plane, in *Proceedings of the 2nd Symposium of Theoretical Aspects of Computer Science* (STACS 1985), LNCS 182, Springer, 1985, pp. 310–321.

[21] J. Snoeyink and J. Stolfi: Objects that cannot be taken apart with two hands, *Discrete & Computational Geometry*, **12** (1994), 367–384.

[22] G. T. Toussaint: Movable separability of sets, in *Computational Geometry*, G. T. Toussaint (editor), North-Holland, Amsterdam, 1985, pp. 335–375.

[23] G. T. Toussaint: On separating two simple polygons by a single translation, *Discrete & Computational Geometry*, **4** (1989), 265–278.