Discrete Comput Geom 35:653–669 (2006)
DOI: 10.1007/s00454-006-1239-x

Discrete & Computational Geometry

The Lifting Model for Reconfiguration*

Sergey Bereg[1] and Adrian Dumitrescu[2]

[1]Department of Computer Science, University of Texas at Dallas,
P.O. Box 830688, Richardson, TX 75083, USA
besp@utdallas.edu

[2]Department of Computer Science, University of Wisconsin–Milwaukee,
3200 N. Cramer Street, Milwaukee, WI 53211, USA
ad@cs.uwm.edu

Abstract. Given a pair of start and target configurations, each consisting of n pairwise disjoint disks in the plane, what is the minimum number of moves that suffice for transforming the start configuration into the target configuration? In one move a disk is lifted from the plane and placed back in the plane at another location, without intersecting any other disk. We discuss efficient algorithms for this task and estimate their number of moves under different assumptions on disk radii. We then extend our results for arbitrary disks to systems of pseudodisks, in particular to sets of homothetic copies of a convex object.

1. Introduction

Consider a set (system) of n pairwise disjoint objects in the plane that need to be brought from a given start (initial) configuration S into a desired goal (target) configuration T. The *motion planning* problem for such a system is that of computing a sequence of object motions (schedule) that achieves this task. If such a sequence of motions exists, we say that the problem is *feasible* and say that it is *infeasible* otherwise. In the model we are considering in this paper, each move is lifting an object and placing it back in the plane anywhere in the free space, that is, at a position where it does not intersect (the interior of) any other object. We refer to this set of motion rules (moves) as the *lifting model*. In this model feasibility is always guaranteed, and we concentrate on estimating the number of moves that are necessary for reconfiguration. We first examine systems of disks and then extend our results to systems of pseudodisks. It should be mentioned here that it is

---

* A preliminary version of this paper appeared in the *Proceedings of the 21st Annual Symposium on Computational Geometry* (*SOCG* '05), Pisa, Italy, June 2005.

the union of the start and target configurations that must be a pseudodisk arrangement, and not the individual configurations (which are assumed to consist of disjoint pieces!).

Other reconfiguration rules (models) for systems of disks have been examined recently, for example: in [6] moves are restricted so that a disk can only be placed in a position where it is adjacent to at least two other disks; in [1] moves are translations along a fixed direction at each step; in [3] a move is sliding an object along a continuous curve in the plane. Classical puzzles with coins can be found in Gardner's mathematical games article on penny puzzles [8]. Reconfiguration for modular systems acting in a grid-like environment, and where moves must maintain connectivity of the whole system has been recently addressed in [7].

The following simple *universal* algorithm does $2n$ moves for reconfiguration of n objects. In the first step (n moves) move all the objects away anywhere in the free space. In the second step (n moves) bring the objects "back" to target positions. For the class of segments (or rectangles as objects), it is easy to construct examples that require $2n - 1$ moves for reconfiguration, even for congruent segments. We improve the $2n$ upper bound for systems of pseudodisks. Furthermore, for congruent disks we show that approximately n moves are enough (for large n).

A move is a *target move* if it moves an object to a final target position. Otherwise, it is a *non-target* move. Our lower bounds use the the following argument: if no target object coincides with a start object (so each object must move at least once), a schedule with x non-target moves consists of at least $n + x$ moves.

Our results are:

(1) We estimate the number of moves necessary for the reconfiguration of systems of (i) congruent disks (in Section 2, Theorem 1), (ii) arbitrary disks (in Section 3, Theorem 4), and (iii) pseudodisks (in Section 4, Theorem 5).
(2) We obtain along the way several interesting results of independent interest: (i) given two sets each with n pairwise disjoint unit disks, there exists a binary space partition of the plane into polygonal regions each containing roughly the same small number ($\approx n^{2/3}$) of disks and such that the total number of disks intersecting the boundaries of the regions is small ($\approx n^{2/3}$) (in Section 2, Lemma 2); (ii) we prove the planarity of a red–blue intersection graph for systems of pseudodisks (in Section 4, Lemma 5); (iii) we prove a lower bound on the maximum order of induced acyclic subgraphs of a directed graph (in Section 3, Theorem 3).

2. Congruent Disks

In this section we consider the problem of moving n disks of unit radius to n target positions and prove the following result.

Theorem 1. *Given a pair of start and target configurations S and T, each with n congruent disks, one can move disks from S to T using $n + O(n^{2/3})$ moves in the lifting model. The entire motion can be computed in $O(n \log n)$ time. On the other hand, for each n, there exist pairs of configurations which require $n + \Omega(n^{1/2})$ moves for this task.*

In other words, Theorem 1 says that $O(n^{2/3})$ non-target moves always suffice, and $\Omega(n^{1/2})$ non-target moves are sometimes necessary for reconfiguring systems with n congruent disks.

2.1. *Partition*

A *center point* of a finite point set $P \subset \mathbb{R}^d$ is a point c such that every hyperplane passing through c partitions P into two subsets each of size at most $nd/(d+1)$. The existence of a center point for any set of points follows from a classical theorem due to Helly [5]. Jadhav and Mukhopadhyay [9] found a linear-time algorithm for finding a center point using a clever prune-and-search approach.

Lemma 1. *Let S be a set of s pairwise disjoint red unit disks, let T be a set of t pairwise disjoint blue unit disks, and let P be the set of centers of disks in $S \cup T$. Let c be a center point of P. Then there is a line through c which intersects only $O(\sqrt{k})$ disks, where $k = s + t$.*

Proof. Let $m = \lceil \sqrt{k} \rceil$, $\varphi = \pi/m$, and $r = 1/\sin(\varphi/2) = O(\sqrt{k})$. We construct the circle C with center at c and radius r and m lines passing through c with slopes $\pi i/m$, $i = 0, 1, \ldots, m-1$, see Fig. 1(a). The space outside C is partitioned into $2m$ regions by these lines. By pairing opposite regions we define m regions $Q_1, Q_2, \ldots, Q_m$, see Fig. 1(a).

There is a region $Q = Q_i$ with at most k/m points of P. Let l be the centerline of Q, see Fig. 1(b). We show that l intersects at most $O(\sqrt{k})$ disks. It suffices to prove that

(i) l intersects at most $\sqrt{k}$ disks whose centers are outside C, and
(ii) l intersects at most $O(\sqrt{k})$ disks whose centers are inside C.

Centers outside C. Let p be a point from $\mathbb{R}^2 - (C \cup Q)$ nearest to the line l, see Fig. 1(b). The distance from p to l is $r \sin\varphi/2 = 1$. This implies that a disk with center q outside C can intersect l only if q lies in Q. The number of such disks is at most $k/m \leq \sqrt{k}$.

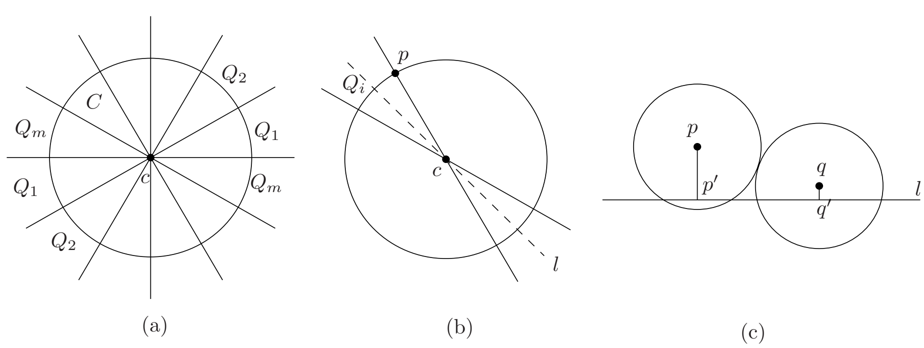

Fig. 1. (a) Example of partition for $k = 30$, (b) the line l, and (c) two disks in S whose centers lie on the same side of ℓ.

Centers inside C. It suffices to prove the claim for the disks in S and disks in T independently. Without loss of generality we assume that l is horizontal. Let p and q be the centers of two disks in S intersecting l such that p and q lie on the same side of l, see Fig. 1(c). Let p' and q' be the projections of p and q onto l, respectively. We show that the distance between p' and q' is at least $\sqrt{3}$. Indeed, $|p'q'|^2 = |pq|^2 - (p_y - q_y)^2 \geq 4 - 1 = 3$. Thus the number of disks in S on one side of l whose centers are inside C is at most $2r/\sqrt{3} + 1$, hence the number of disks in S with centers inside C intersecting l is at most $4r/\sqrt{3} + 2$. Similarly, the number of disks in T with centers inside C intersecting l is at most $4r/\sqrt{3} + 2$, thus the total number of disks with centers inside C intersecting l is at most $8r/\sqrt{3} + 4 = O(\sqrt{k})$. □

Next, we construct a partition of the plane into convex regions such that each region contains a small number of start and target disks.

Lemma 2. *Let S be a set of n pairwise disjoint red unit disks, and let T be a set of n pairwise disjoint blue unit disks. Let $m = \lceil n^{2/3} \rceil$. There exists a (binary space) partition of the plane into $O(n^{1/3})$ convex polygonal (bounded or unbounded) regions such that*

(i) *each region contains between m and $3m$ disks in $S \cup T$ lying completely in the region, and*
(ii) *the total number of disks intersecting the boundaries of all regions is $O(n^{2/3})$.*

Proof. We obtain the desired partition in a recursive manner. Let S be a set of s pairwise disjoint red unit disks, and let T be a set of t pairwise disjoint blue unit disks where $s + t = k$. Initially $s = t = n$ and $k = 2n$. By Lemma 1, there is a line ℓ passing through a center point of their centers that intersect only $O(\sqrt{k})$ disks. Let $S_0 \cup T_0$ be the set of disks intersecting ℓ, where $S_0 \subset S$ and $T_0 \subset T$. Let $S_1 \cup T_1$ and $S_2 \cup T_2$ be the sets of disks on each side of ℓ, where $S_1 \subset S$, $S_2 \subset S$ and $T_1 \subset T$, $T_2 \subset T$. Write $k_i = |S_i \cup T_i|$, for $i = 0, 1, 2$. Since ℓ passes through a center point and $k_0 + k_1 + k_2 = k$, we have $k_1, k_2 \leq 2k/3$ and, by Lemma 1, $k_0 = O(\sqrt{k})$.

We apply Lemma 1 recursively for the sets $S_1 \cup T_1$ and $S_2 \cup T_2$. The recursion stops when $k \leq 3m$. Let $R_1, \ldots, R_r$ be the regions generated by this procedure. By the property of center points, the number of disks (in $S \cup T$) contained in each region is between m and $3m$. Let $B(k)$ be the set of disks intersecting the region boundaries formed when starting with a set of k disks, and write $b(k) = |B(k)|$. We want to obtain an upper bound on $b(2n)$. $b(k)$ satisfies the following recurrence:

$$b(k) \leq \begin{cases} 0 & \text{if} \quad k \leq 3m, \\ c\sqrt{k} + b(k_1) + b(k_2) & \text{if} \quad k > 3m. \end{cases}$$

We can assume that n is sufficiently large, in particular $n \geq (100c^2)^{3/2}$, where c is the constant hidden in the big Oh notation in Lemma 1 (i.e., the partitioning line intersects at most $c\sqrt{k}$ disks). We prove by induction on k that for $k \geq m$ we have

$$b(k) \leq 5c\left(\frac{k}{\sqrt{m}} - \sqrt{k}\right). \tag{1}$$

Note that the right-hand side is non-negative thus the basis of the induction $k \in [m, 3m]$ is satisfied. Assume now $k \geq 3m$. By the induction hypothesis $b(k_i) \leq 5c(k_i/\sqrt{m} - \sqrt{k_i})$ for $i = 1, 2$. Then

$$\begin{aligned} b(k) &\leq k_0 + 5c\frac{k_1 + k_2}{\sqrt{m}} - 5c(\sqrt{k}_1 + \sqrt{k}_2) \\ &\leq c\sqrt{k} + 5c\frac{k}{\sqrt{m}} - 5c(\sqrt{k}_1 + \sqrt{k}_2). \end{aligned}$$

It suffices to prove that

$$\sqrt{k_1} + \sqrt{k_2} \geq \tfrac{6}{5}\sqrt{k},$$

or, equivalently,

$$\sqrt{\frac{k_1}{k}} + \sqrt{\frac{k_2}{k}} \geq \tfrac{6}{5}. \tag{2}$$

Assume that $k_1 \leq k_2$. By our assumption on n, $k \geq 100c^2$, and by the property of center points, $k_1 \geq k/3 - c\sqrt{k}$, thus

$$\frac{k_1}{k} \geq \frac{1}{3} - \frac{c}{\sqrt{k}} \geq \frac{1}{3} - \frac{1}{10} = \frac{7}{30}.$$

Since

$$\frac{k_1}{k} + \frac{k_2}{k} = 1 - \frac{k_0}{k} \geq 1 - \frac{c}{\sqrt{k}} \geq \frac{9}{10},$$

the minimum of $\sqrt{k_1/k} + \sqrt{k_2/k}$ is attained for $k_1/k = \frac{7}{30}$ and $k_2/k = \frac{20}{30}$. Consequently

$$\sqrt{\frac{k_1}{k}} + \sqrt{\frac{k_2}{k}} \geq \sqrt{\frac{7}{30}} + \sqrt{\frac{2}{3}} \geq \frac{6}{5},$$

which proves inequality (2). Since we selected $m = \lceil n^{2/3} \rceil$, inequality (1) implies $b(2n) \leq 10c \cdot n/n^{1/3} = O(n^{2/3})$. □

Remark. The bound on the above recurrence is tight: in the case where the partition is defined by $k_0 = \Theta(\sqrt{k})$ and $k_1 = k_2 = (k - k_0)/2$ at each step.

2.2. *Moving Disks*

By Lemma 2 there exists a partition $\mathcal{R}$ of the plane into at most $r = O(n^{1/3})$ regions $R_1, \ldots, R_r$ such that the number of disks and targets in each region is in the range $[m, 3m]$ and the total number of disks (in $S \cup T$) intersecting the boundaries of the regions is at most $c_1 m$, where c_1 is an absolute constant. Without loss of generality, we can assume that n is sufficiently large, thus $n \geq (c_1 + 6)m$. Let $S_i = S \cap R_i$ and $T_i = T \cap R_i$, where $i = 1, 2, \ldots, r$, be the set of (start) disks and target disks in region R_i. Let $s_i = |S_i|$ and $t_i = |T_i|$ for $i = 1, 2, \ldots, m$. We have $s_i + t_i \in [m, 3m]$. We denote by S_δ and T_δ

the set of start/target disks intersecting the boundaries of the regions. Let $s_\delta = |S_\delta|$ and $t_\delta = |T_\delta|$. We have $\sum_{i=1}^{r} s_i + s_\delta = \sum_{i=1}^{r} t_i + t_\delta = n$, and $s_\delta + t_\delta \le c_1 m$.

We sort the regions R_i in decreasing order of $t_i - s_i$. We first move away to the free space the disks (in S_δ) intersecting the boundaries of the regions. The number of moved disks is $s_\delta \le c_1 m$. The general idea of our *shifting algorithm* is as follows: we move away all the disks from some regions at the front of the list ordered by $t_i - s_i$ until the number of free targets in these regions is at least $3m$. This allows for repeatedly moving (shifting) all the (start) disks from another region to the free targets available, and such that after each step the number of free targets is still at least $3m$ and the regions which are "completed" are free of disks in S. In the end the target positions on the boundaries of the regions are filled and the reconfiguration is complete.

Let $k \ge 1$ be the minimum positive integer such that $\sum_{i=1}^{k} s_i \ge (c_1 + 3)m$. Since, for each i, $s_i \le 3m$, we have $\sum_{i=1}^{k} s_i \le (c_1 + 6)m \le n$.

Claim. $\sum_{i=1}^{k} t_i \ge 3m$.

Proof. We distinguish two cases.

Case 1: $t_k \ge s_k$. By the order condition on $t_i - s_i$ (which holds initially), $t_{k-1} \ge s_{k-1}, \dots, t_1 \ge s_1$, hence

$$\sum_{i=1}^{k} t_i \ge \sum_{i=1}^{k} s_i \ge (c_1 + 3)m \ge 3m.$$

Case 2: $t_k < s_k$. By the order condition on $t_i - s_i$ for $i \ge k+1$, $t_{k+1} < s_{k+1}, \dots, t_r < s_r$. Recall that initially we have $\sum_{i=1}^{r} s_i + s_\delta = \sum_{i=1}^{r} t_i + t_\delta$, and $t_\delta \le c_1 m$, thus

$$\sum_{i=1}^{r} t_i = \sum_{i=1}^{r} s_i + s_\delta - t_\delta \ge \sum_{i=1}^{r} s_i - c_1 m.$$

Therefore,

$$\begin{aligned} \sum_{i=1}^{k} t_i &= \sum_{i=1}^{r} t_i - \sum_{i=k+1}^{r} t_i \ge \left(\sum_{i=1}^{r} s_i - c_1 m\right) - \sum_{i=k+1}^{r} t_i \\ &\ge \sum_{i=1}^{r} s_i - c_1 m - \sum_{i=k+1}^{r} s_i = \sum_{i=1}^{k} s_i - c_1 m \ge 3m. \end{aligned}$$

The last inequality in the above chain follows from the definition of k. □

We continue now with the description of the algorithm. Move away the $\sum_{i=1}^{k} s_i$ disks in the regions $R_1, \dots, R_k$, and still denote by $s_1, \dots, s_r$ and $t_1, \dots, t_r$ the resulting number of disks and targets in the r regions. All remain the same except $s_1, \dots, s_k$ which are now 0. Since $s_1 = \cdots = s_k = 0$, by the above claim,

$$\sum_{i=1}^{k} (t_i - s_i) \ge 3m. \tag{3}$$

Note that the order of $t_i - s_i$ is still maintained for $i \geq k+1$, i.e.,

$$t_{k+1} - s_{k+1} \geq \cdots \geq t_r - s_r. \tag{4}$$

We also have

$$\sum_{i=1}^{r}(t_i - s_i) \geq 3m. \tag{5}$$

To verify (5), note that $\sum_{i=1}^{r} t_i = n - t_\delta$, and $\sum_{i=1}^{r} s_i \leq n - s_\delta - (c_1+3)m$. Therefore,

$$\sum_{i=1}^{r}(t_i - s_i) \geq n - t_\delta - n + s_\delta + (c_1+3)m \geq (c_1+3)m - t_\delta \geq 3m.$$

Up to now, all moves that have been executed are non-target moves, and their number is at most $c_1 m + (c_1+6)m = (2c_1+6)m = O(n^{2/3})$. From this point on, all moves executed will be target moves. We split $\mathcal{R}$ into two groups $G_1 = \{R_1, \ldots, R_k\}$ and $G_2 = \{R_i \mid i \geq k+1\}$. Note that all regions in G_1 are free of start disks and there are at least $3m$ free targets in G_1. This invariant will be maintained through to the end, see below. Repeat the following step as long as G_2 is non-empty. Move the s_{k+1} disks from R_{k+1} to (any) free targets in G_1 and move R_{k+1} from G_2 to G_1, i.e., now $G_1 = \{R_1, \ldots, R_k, R_{k+1}\}$. The number of free targets in G_1 is $\sum_{i=1}^{k+1}(t_i - s_i)$. Clearly,

$$\sum_{i=1}^{k+1}(t_i - s_i) = \sum_{i=1}^{k}(t_i - s_i) + (t_{k+1} - s_{k+1}). \tag{6}$$

We need to show that there are at least $3m$ free targets in G_1. Assume that $\sum_{i=1}^{k+1}(t_i - s_i) < 3m$. By our induction hypothesis, $\sum_{i=1}^{k}(t_i - s_i) \geq 3m$, so $t_{k+1} - s_{k+1} < 0$ must hold. However, then, by (3), $t_i - s_i < 0$ for $i = k+2, \ldots, r$, and thus $\sum_{i=k+2}^{r}(t_i - s_i) \leq 0$ (the inequality is not strict, since the sum could be empty, when $r = k+1$). Consequently,

$$\sum_{i=1}^{r}(t_i - s_i) = \sum_{i=1}^{k+1}(t_i - s_i) + \sum_{i=k+2}^{r}(t_i - s_i) < 3m,$$

by our assumption. This is in contradiction to inequality (5). Since all regions in G_1 are free of start disks, and there are at least $3m$ free targets in G_1, the invariant is maintained after each step. After the current step is completed set $k \leftarrow k+1$.

When G_2 becomes empty, we continue to fill in target positions in G_1 and then on the boundary of the regions using the disks removed at the beginning. The total number of disk moves is $n + O(n^{2/3})$.

First, we show that the partitioning line provided by Lemma 1 can be found in linear time as follows. The center of P can be computed in linear time [9]. For every disk in $S \cup T$ with center p outside C, we decide what region Q_j contains p in $O(1)$ time using the slope of cp. Once the counting for the regions is done, we select Q, the region with the minimum number of points from P.

The computation of the regions R_i takes $O(n \log n)$ time since there are $O(\log n)$ levels of the recursion and the total complexity of the problems on any level is $O(n)$. The sorting of the regions R_i takes $O(n^{1/3} \log n) = O(n)$ time. We maintain a list of free target disks in the regions from G_1. This can be done in $O(1)$ time per move. The total time taken by the reconfiguration algorithm is $O(n \log n)$.

2.3. *Lower Bound*

$$\text{Let } m = \begin{cases} \lfloor\sqrt{n}\rfloor & \text{if } \lfloor\sqrt{n}\rfloor \text{ is odd,} \\ \lfloor\sqrt{n}\rfloor - 1 & \text{if } \lfloor\sqrt{n}\rfloor \text{ is even.} \end{cases}$$

Note that m is odd. Let $T = T' \cup T''$ be the set of target disks and let $S = S' \cup S''$ be the set of start disks, where $|T'| = m^2$ and $|S'| = m^2 - 1$. We place the disks of T' onto a grid $X \times X$ of size $m \times m$ where $X = \{2, 4, \ldots, 2m\}$, see Fig. 2 (top). Let $S' = S'_a \cup S'_b$, where $|S'_a| = (m-1)^2$ and $|S'_b| = 2m - 2$. We place S'_a onto a grid of size $(m-1) \times (m-1)$ so that they overlap with the disks from T', as in Fig. 2 (top). The grid of target disks contains $4m - 4$ disks on its boundary. We "block" them with start disks in S'_b by placing them so that each start disk overlaps with two boundary target disks, see Fig. 2 (top). Note that $S' \cup T'$ is symmetric with respect to rotation by 90°, 180°, and 270°. We place the remaining $n - m^2$ target disks in T'' and the $n - m^2 + 1$ start disks in S'' on a line as in Fig. 2 (bottom).

The m columns of T' are numbered $1, \ldots, m$ from left to right and the m rows are numbered $1, \ldots, m$ from bottom to top. The column of S' in between columns i and $i+1$ of T' is denoted S_i, $i = 1, \ldots, m-1$, and the $\lfloor m/2 \rfloor$ disks in S' intersecting target disks in column 1 of T' to the left are denoted by S_0.

We now show that $n + \lfloor m/2 \rfloor$ is a lower bound on the number of moves. Consider a sequence of moves that transform S into T. After every move, we consider n_x, the number of distinct x-coordinates of occupied target disks from T', and n_y, the number of distinct y-coordinates of occupied target disks from T'. Put $k = m - 1$. We consider the intermediate configuration at the first moment when $\max(n_x, n_y) = k$. It suffices to prove that $k/2$ non-target moves were performed up to this moment.

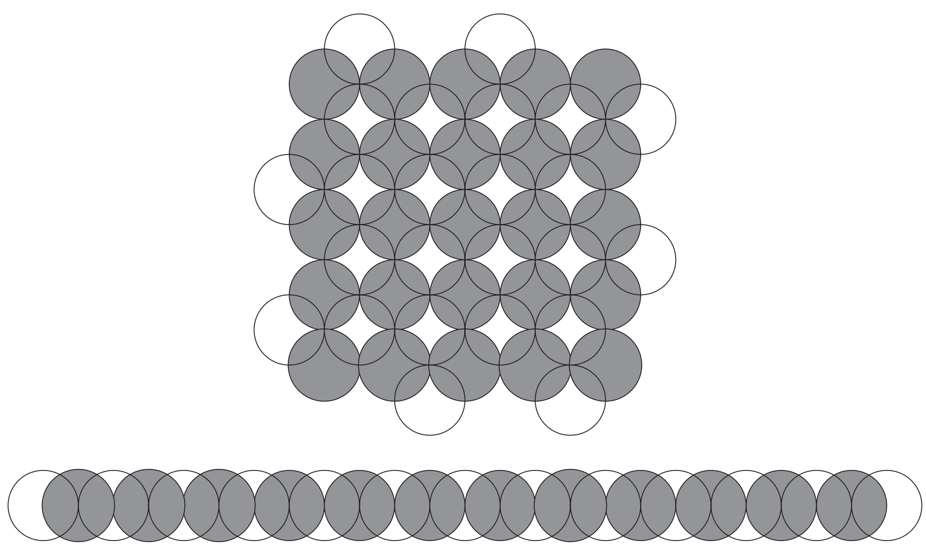

Fig. 2. Constructing a pair of start and target configurations with $n = 37$ disks ($m = 5$): place m^2 target disks in T' and $m^2 - 1$ start disks in S' (top); place $n - m^2$ target disks in T'' and $n - m^2 + 1$ start disks in S'' (bottom). The start disks are white and the target disks are shaded.

Let Y be the set of occupied target disks from T' and let Z be the set of start disks from S' intersecting them. Write $t' = |Y|$ and $s' = |Z|$. Denote by T_i the set of occupied target disks in column i of T', $i = 1, \dots, m$: we refer to T_i as a column of Y. Let t'' be the number of occupied target disks from T'' and let s'' be the number of start disks intersecting them. We want to prove that $s' + s'' - t' - t'' \geq k/2$. Since $s'' \geq t''$, it suffices to prove that $s' - t' \geq k/2$.

Without loss of generality (by the symmetry of $S' \cup T'$) we can assume that $n_x = k$ and $n_y \leq k$. Y consists of one or more blocks, where each *block* is a (maximal) set of consecutive columns of T' each with occupied targets. Denote by q the number of blocks in Y, and by k_j the number of columns of block j, $j = 1, \dots, q$. We have $\sum_{j=1}^{q} k_j = k$. Since any two consecutive blocks are separated by at least one "empty" column of T' (i.e., with no occupied target disks) we have $q \in \{1, 2\}$. First note that by our choice of k, no single block can have target disks from both the first and the last column of T' (any block can span at most $k < m$ columns). Thus each of the at most two blocks must have occupied targets from either column 1 or column m of T', but not both. Without loss of generality (by symmetry) assume that the first (left) block B_1 has occupied targets from column 1 of T', thus $B_1 = T_1 \cup \dots \cup T_{k_1}$.

We claim that the number of disks in S' intersecting target disks in B_1 is at least

$$|B_1| + \left\lceil \frac{k_1}{2} \right\rceil.$$

For each occupied target disk in B_1 we uniquely assign a start disk that intersects it by the following scheme. For i odd (resp. even), assign to a target disk with center at (x, y) the start disk with center at $(x + 1, y + 1)$ (resp. $(x + 1, y - 1)$). In addition, for each pair of consecutive columns T_i, T_{i+1} of B_1 we will specify (next) a start disk in $S_i \cup S_{i+1}$ which intersects $T_i \cup T_{i+1}$ and which is unaccounted by the above scheme.

Assume that i is odd. If T_i contains a target disk ω with center at (x, y) such that there is a target disk immediately below it which is not in T_i, then the start disk (in S_i) centered at $(x + 1, y - 1)$ intersects ω and is unaccounted by the above scheme. Similarly, if T_{i+1} contains a target disk ω with center at (x, y) such that there is a target disk immediately above it which is not in T_{i+1}, then the start disk (in S_{i+1}) centered at $(x + 1, y + 1)$ intersects ω and is unaccounted by the above scheme. If neither of the previous two cases occurs, it means that all the elements of T_i occupy bottom rows $1, 2, \dots$ in column i and all the elements of T_{i+1} occupy top rows $m, m - 1, \dots$ in column $i + 1$. Let $\omega \in T_{i+1}$ be the lowest target disk in column T_{i+1}, with center at (x, y). Since $|T_i| + |T_{i+1}| \leq n_y \leq k = m - 1$, the start disk in S_i with center at $(x - 1, y - 1)$ intersects ω and is unaccounted by our scheme. The argument for even i is analogous.

If k_1 is even then the number of start disks intersecting target disks in B_1 is at least $|B_1| + k_1/2$, as required: we apply the above pairing argument for $i = 1, 3, \dots, k_1 - 1$, to get $k_1/2$ start disks unaccounted by our scheme. Suppose now that k_1 is odd. We apply the pairing argument for $i = 2, 4, \dots, k_1 - 1$ and obtain $(k_1 - 1)/2$ start disks unaccounted by our scheme. We now show that there is yet another start disk ω that intersects a target disk in T_1 and is unaccounted by our scheme. Either (i) a disk in S_0 intersects a target disk in T_1, and can be chosen as ω, or (ii) T_1 consists of the left-top corner target disk of T' and the second disk from the top in S_1 can be chosen as ω.

Therefore in both cases (k_1 even or k_1 odd), at least $\lceil k_1/2 \rceil$ start disks unaccounted by our scheme intersect occupied target disks in B_1. This proves our claim.

If $q = 2$, let B_2 be the second block (which has target disks from the last column of T'). We only have to observe that by taking the origin of the coordinate axes at $(2m + 2, 2m + 2)$ and reversing orientation for the axes, the same scheme can be used (by symmetry) for this block (if it exists), and the same bound holds: the number of disks in S' intersecting target disks in B_2 is at least

$$|B_2| + \left\lceil \frac{k_2}{2} \right\rceil.$$

Putting things together we obtain

$$s' - t' \geq \sum_{j=1}^{q} \left\lceil \frac{k_j}{2} \right\rceil \geq \frac{1}{2} \sum_{j=1}^{q} \frac{k_j}{2} = \frac{k}{2}.$$

This concludes the proof of Theorem 1.

Remark. We note that our lower bound is best possible for *this* construction: indeed, the reader could verify that $n + \lfloor m/2 \rfloor$ moves are enough for reconfiguring S into T.

3. Arbitrary Disks

In this section we consider the problem of moving n disks to n target positions. Let $S = \{s_1, \ldots, s_n\}$ and $T = \{t_1, \ldots, t_n\}$ be the start and target configurations. We assume that for each i, disk s_i is congruent to disk t_i, i.e., t_i is the target position of s_i; if the correspondence $s_i \to t_i$ is not given (but only the two sets of disks), it can be easily computed by sorting both S and T by radius.

In an undirected graph $G = (V, E)$, let d_v denote the degree of vertex v, and let $\alpha(G)$ be the *independence number* of G (i.e., the maximum size of an independent set of vertices). The following result was proved by Caro and Wei [2].

Theorem 2 (Caro and Wei).

$$\alpha(G) \geq \sum_{v \in V} \frac{1}{d_v + 1}.$$

In a directed graph $D = (V, E)$, let $d_v = d_v^+ + d_v^-$ denote the degree of vertex v, where d_v^+ is the out-degree of v and d_v^- is the in-degree of v. Let $\beta(D)$ be the maximum size of a subset V' of V such that $D[V']$, the subgraph induced by V', is acyclic.

Using a similar argument as in the proof of Theorem 2 given in [2], we prove the following result for directed graphs:

Theorem 3.

$$\beta(D) \geq \max \left(\sum_{v \in V} \frac{1}{d_v^+ + 1} \, , \, \sum_{v \in V} \frac{1}{d_v^- + 1} \right).$$

Proof. Write $N^+(v) = \{w\colon (v, w) \in E\}$ and $N^-(v) = \{w\colon (w, v) \in E\}$. Let $<$ be a uniformly chosen total ordering on V, and think of V laid out left to right in this order. Define

$$V' = \{v \in V\colon (v, w) \in E \Rightarrow v < w\}.$$

Let X_v be the indicator random variable for the event $v \in V'$ and $X = \sum_{v\in V} X_v = |V'|$. For each v,

$$E[X_v] = \Pr[v \in V'] = \frac{1}{d_v^+ + 1},$$

since $v \in V'$ if and only if v is the least element among v and $N^+(v)$. Hence

$$E[X] = \sum_{v\in V} \frac{1}{d_v^+ + 1}$$

so there exists a specific ordering $<$ with

$$|V'| \geq \sum_{v\in V} \frac{1}{d_v^+ + 1}.$$

Assume that $D[V']$, the subgraph induced by V', contains a cycle C, and now let v denote the largest element in this cycle. Since $v \in V'$, all vertices in $N^+(v)$ lie right of v with respect to $<$, a contradiction to the fact that C is a cycle. Thus the subgraph induced by V' is acyclic and $\beta(D) \geq |V'|$.

The second inequality can be proven in a similar way, or just by reversing the graph. □

For a disk ω, let $\mathring{\omega}$ denote the interior of ω. Let S be a set of k pairwise disjoint red disks, and let T be a set of l pairwise disjoint blue disks. Consider the bipartite red–blue disk intersection graph $G = (S, T, E)$, where $E = \{(s, t)\colon s \in S, t \in T, \mathring{s} \cap \mathring{t} \neq \emptyset\}$. The next lemma follows from Lemma 5 (in Section 4), nevertheless we include a simple proof for this special case.

Lemma 3. *Any red–blue disk intersection graph $G = (S, T, E)$ is planar.*

Proof. Consider the standard embedding of G in the plane, still denoted by G, such that each disk is mapped to its center, and edges are straight line segments. We denote by $R(p)$ (resp. $B(p)$) the unique red (resp. blue) disk centered at point p. We have to show that no pair of segment edges properly cross. Consider two edges s_1t_1 and s_2t_2 that cross, where s_1, s_2 are the centers of two disks $R(s_1), R(s_2) \in S$ and t_1, t_2 are the centers of two disks $B(t_1), B(t_2) \in T$. It is easy to check degenerate cases, thus one may assume that s_1, s_2, t_1, t_2 are four points in convex position in counterclockwise order. Write r_i for the radius of $R(s_i)$, and q_i for the radius of $B(t_i)$, $i = 1, 2$.

The intersection condition can be written as

$$|s_1t_1| < r_1 + q_1 \quad \text{and} \quad |s_2t_2| < r_2 + q_2.$$

The disjointness condition for disks of the same color can be written as

$$|t_1 t_2| \geq q_1 + q_2 \quad \text{and} \quad |s_1 s_2| \geq r_1 + r_2.$$

Putting these together with the triangle inequality yields

$$r_1 + r_2 + q_1 + q_2 \ > \ |s_1 t_1| + |s_2 t_2| > |s_1 s_2| + |t_1 t_2| \geq r_1 + r_2 + q_1 + q_2,$$

a contradiction. □

Theorem 4. *Given a pair of start and target configurations S and T, each with n disks with arbitrary radii, $9n/5$ moves always suffice for transforming the start configuration into the target configuration. On the other hand, for each n, there exist pairs of configurations which require $\lfloor 5n/3 \rfloor$ moves for this task.*

Proof. We first prove the lower bound. We use disks of different radii (although the radii can be chosen very close to the same value if desired). Since all disks have distinct radii, one can think of the disks as being labeled. Consider the set of three disks, labeled 1, 2, and 3 in Fig. 3. The two start and target disks labeled i are congruent, for $i = 1, 2, 3$. To transform the start configuration into the target configuration takes at least two non-target moves, thus five moves in total. By repeatedly using groups of three (with different radii), we get a lower bound of $5n/3$ moves, when n is a multiple of three, and $\lfloor 5n/3 \rfloor$ in general.

We now prove the upper bound. A short outline of the algorithm is as follows. Move out (far away) a set of disks of size at most αn, where $\alpha < 1$ is a constant. Move to targets in a suitable order the remaining disks so that all moves are target moves. Then fill the targets of the far away disks in any order. The total number of moves is at most $(1 + \alpha)n$, and we show how to do this with $\alpha = \frac{4}{5}$.

We now provide the details. By Lemma 3 applied to the sets of start disks and target disks S and T, where $|S| = |T| = n$, we get

$$|E| \leq 2(|S| + |T|) - 4 = 4n - 4.$$

We think of the start and target disks being labeled from 1 to n, so that the target of start disk i is target disk i. Consider the directed *blocking graph* $D = (S, F)$ on the set S of n start disks, where

$$F = \{(s_i, s_j)\colon\ i \neq j \text{ and } \mathring{s}_i \cap \mathring{t}_j \neq \emptyset\}.$$

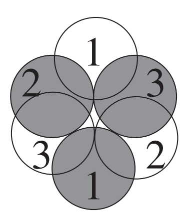

Fig. 3. A group of three disks with their targets: part of the lower bound construction for lifting disks of arbitrary radii. The disks are white and their targets are shaded.

If $(s_i, s_j) \in F$ we say that disk i *blocks* disk j. (Note that $s_i \cap t_i \neq \emptyset$ does not generate any edge in D.) Since if $(s_i, s_j) \in F$ then $(s_i, t_j) \in E$, we obtain $|F| \leq |E| \leq 4n - 4$. Let

$$d^+ = \frac{\sum_{v \in S} d_v^+}{n} = \frac{|F|}{n}$$

be the average out-degree in D. We have

$$d^+ \leq \frac{4n-4}{n} \leq 4.$$

Using Jensen's inequality applied to the function $f(t) = 1/(t+1)$ on the interval $[0, \infty)$, and the first inequality in Theorem 3, we get

$$\beta(D) \geq \sum_{v \in S} \frac{1}{d_v^+ + 1} \geq \frac{n}{d^+ + 1} \geq \frac{n}{5}.$$

Let $S' \subset S$ be a set of disks of size at least $n/5$ and whose induced subgraph is acyclic in D. Move out far away the remaining set S'' of at most $4n/5$ disks, and note that the far away disks do not block any of the disks in S'. (More precisely, this happens in the new blocking graph on the set of disks obtained after these moves.) Perform a topological sort on the acyclic graph $D[S']$ induced by S', and fill the targets of these disks in that order using only target moves. Then fill the targets with the far away disks in any order. The number of moves is at most $n + 4n/5 = 9n/5$, as claimed. □

For any given $\varepsilon > 0$, a set $S' \subset S$ of size $\geq (\frac{1}{5} - \varepsilon)n$ whose induced subgraph is acyclic in D can be found with high probability in $O(n)$ expected time by repeating $O(1/\varepsilon)$ times the experiment in the proof of Theorem 3. The resulting number of moves is at most $n + (\frac{4}{5} + \varepsilon)n = (\frac{9}{5} + \varepsilon)n$. Assuming the intersection graph G is given, the total expected time of the algorithm is $O(n)$. Alternatively, a set $S' \subset S$ of size $\geq n/5$ whose induced subgraph is acyclic in D can be found in $O(n^2)$ expected time by repeating $O(n)$ times the experiment in the proof of Theorem 3, and resulting in at most $9n/5$ moves.

Clearly, the bipartite disk intersection graph G can be computed in $O(n^2)$ time. We present a faster algorithm in the next section.

3.1. *Computation of the Bipartite Disk Intersection Graph G*

The algorithm is based on the following property.

Lemma 4. *Let A, B, and C be three disks in the plane whose centers are a, b, and c and whose radii are r_a, r_b and r_c, respectively. Suppose that the disks B and C are disjoint and A is the smallest disk. If $x_a \leq x_b \leq x_c$ and the slopes of the lines ab and ac are in the range $[-\pi/6, \pi/6]$ then the disks A and C are disjoint.*

Proof. The conditions in the lemma give $|bc| \geq r_b + r_c$ and $r_a \leq \min\{r_b, r_c\}$. Without loss of generality we can assume that $y_b \geq y_c$, and that the slope of the line ab is $\pi/6$

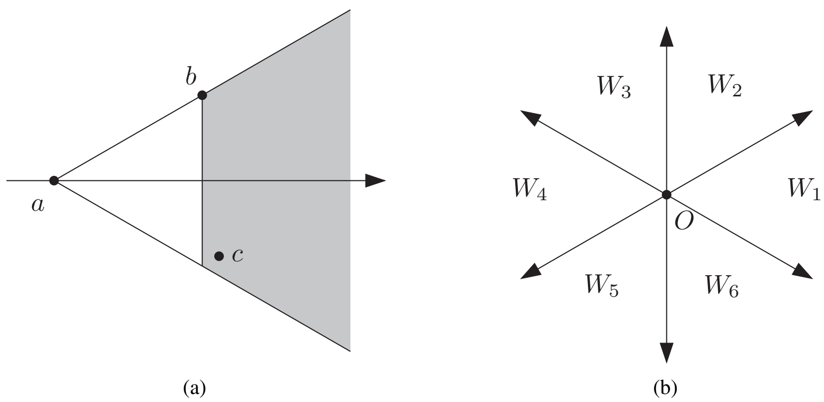

Fig. 4. (a) Three disks property. The shaded region is the locus of positions for c. (b) Six wedges.

since b can be moved vertically up otherwise, see Fig. 4. Then c is located in the shaded region bounded by the vertical line passing through b and the two lines with slopes $-\pi/6, \pi/6$ passing through a. Therefore, the inequality between the angles of Δabc gives

$$|ac| \geq |bc| \geq r_b + r_c \geq r_a + r_c,$$

and the lemma follows. □

Lemma 4 allows us to reduce the problem of computing the graph G to dynamic range searching as follows. For each of the six wedges shown in Fig. 4(b), we transform the axes into xy-axes and transform the coordinates of the points. For the ith wedge, $1 \leq i \leq 6$, we maintain two data structures $D_{i,S}$ and $D_{i,T}$ storing the centers of the disks from S and T, respectively, that have been processed so far. Each data structure allows us to find a point with the minimum x-coordinate as a key, in the range defined by a query point (the key of a center p in $D_{1,S}$ and $D_{1,T}$ is x_p). We use dynamic range search trees [14] for each $D_{i,X}$, $1 \leq i \leq 6$, $X \in \{S, T\}$.

Algorithm

1. Sort the disks in $S \cup T$ by radius. Initialize $D_{1,S}, \ldots, D_{6,S} = \emptyset$, and $D_{1,T}, \ldots, D_{6,T} = \emptyset$.
2. For each disk $\delta \in S \cup T$ in decreasing order of radius do the following steps:
   - 2.1. Let v be the vertex of G corresponding to δ and let a be the center of δ. Let $X \in \{S, T\}$ be the set containing δ and let $Y = \{S, T\} - X$.
   - 2.2. For each data structure $D_{i,Y}$, $1 \leq i \leq 6$, find the minimum key in the range corresponding to a. Let $c_1, \ldots, c_6$ be the centers corresponding to the found keys. Test the disks centered at $c_1, \ldots, c_6$ and select ones that intersect δ. Add the corresponding edges to G.
   - 2.3. Insert a into $D_{i,X}$ for each $1 \leq i \leq 6$.

The time for computing the smallest key and for updating the range search tree is $O(\log^2 n)$ [14]. The overall running time is $O(n \log^2 n)$. This can be further improved to $O(n \log n \log \log n)$ by using the technique of fractional cascading [10], [12].

4. Pseudodisks

In this section we extend the results in Theorem 4 to systems of pseudodisks. Let each object (planar shape) be described by a simple closed curve in the plane. A pair ω_1, ω_2 of shapes is a pair of *pseudodisks* if it satisfies the *pseudodisk property*: the sets $\omega_1 \backslash \omega_2$ and $\omega_2 \backslash \omega_1$ are connected [4]. Equivalently, the boundaries of any two shapes have at most two proper intersections. We say that start and target configurations S and T form a system (collection) of pseudodisks if $S \cup T$ does; obviously such intersections can occur only between an element of S and an element of T. Note that a pair S, T forming such a system depends primarily on the relative position of S with respect to T rather than on the objects themselves: for example, when reconfiguring a set of congruent horizontal thin rectangles (or segments) into vertical ones "overlapping" them, we do not have such a system, but we do if target positions are horizontal as well.

As in Section 3, we consider the bipartite intersection graph $G = (S, T, E)$, where edges correspond to start-object–target-object intersections. The result in Lemma 3 extends to systems of pseudodisks, however, its proof does not rely on the triangle inequality, but on the pseudodisk property instead.

Lemma 5. *Let S and T be the start and target configurations for a set of objects in the plane such that S and T form a collection of pseudodisks. Then the corresponding intersection graph $G = (S, T, E)$ is planar.*

Proof. First ignore start objects completely contained in targets and targets completely contained in start objects. Then for each intersecting pair (s, t) modify the boundary of one of them, say s (by a continuous motion), so that s and t become touching and have disjoint interiors. This can be done because no other start or target object intersects $\mathring{s} \cap \mathring{t}$, and by the pseudodisk property. Then the initial intersection graph becomes the *contact graph* of a family of simple closed curves (the new boundaries of start and target objects) with disjoint interiors: two curves are connected in the graph if they touch, i.e., their boundaries have at least a common point. It is known that the contact graph of such a family of curves is planar (see, e.g., [11] and [13]).

For objects completely contained in targets and targets completely contained in objects, add corresponding edges to the contact graph. The resulting graph is still planar (since the other one was) and is exactly the initial intersection graph. □

Remarks. For circular disks the following graph is planar: connect two vertices if there is a region covered by the corresponding two disks and by nothing else. (This region may not be adjacent to the boundary of the union of all disks.) This provides another proof of Lemma 3.

Table 1. Bounds on the number of moves in the sliding and lifting models.

Model	Disk type	Lower bound	Upper bound
Sliding	Congruent	$16n/15 - o(n)$	$3n/2 + o(n)$
	Arbitrary	$2n - o(n)$	$2n - 1$
Lifting	Congruent	$n + \Omega(n^{1/2})$	$n + O(n^{2/3})$
	Arbitrary	$\lfloor 5n/3 \rfloor$	$9n/5$

Note that the above lemma holds for an arbitrary bipartite pseudodisk intersection graph, not necessarily with objects (pseudodisks) congruent to one another in the two parts.

Finally, using the same argument as in Section 3, we obtain

Theorem 5. *Given a pair of start and target configurations S and T, each with n objects, such that S and T form a system of pseudodisks, $9n/5$ moves always suffice for transforming the start configuration into the target configuration.*

In particular, the above result holds for systems of homothetic copies of a convex object.

5. Concluding Remarks

The following related model—called the *sliding model*—has been considered in [3]. In one move a disk slides in the plane without intersecting any other disk, so that its center moves along an arbitrary (open) continuous curve. Clearly, any lower bound (on the number of moves) for lifting is also valid for sliding, and any upper bound (on the number of moves) for sliding is also valid for lifting. Another remark is that for lifting, those objects whose target position coincides with their start position can be safely ignored, while for sliding this is not true. It is quite interesting to compare the bounds on the number of moves for the two models, sliding and lifting, see Table 1 (the bounds for sliding are taken from [3]).

Acknowledgments

We are grateful to Ferran Hurtado for bringing this type of reconfiguration problems to our attention, and to János Pach for his valuable suggestions made on an earlier version of our manuscript.

References

1. M. Abellanas, S. Bereg, F. Hurtado, A. G. Olaverri, D. Rappaport, and J. Tejel, Moving coins, *Computational Geometry: Theory and Applications*, to appear.

2. N. Alon and J. Spencer, *The Probabilistic Method*, second edition, Wiley, New York, 2000.
3. S. Bereg, A. Dumitrescu, and J. Pach, Sliding disks in the plane, *Proceeding of Japan Conference on Discrete and Computational Geometry* 2004 (J. Akiyama, M. Kano, and X. Tan, editors), pp. 37–47, LNCS 3742, Springer, Berlin, 2005.
4. M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf, *Computational Geometry*, second edition, Springer, Berlin, 2000.
5. L. Danzer, B. Grünbaum, and V. Klee, Helly's theorem and its relatives, in *Convexity*, volume 7 of Proceeding of Symposia in Pure Mathematics, pp. 101–180, American Mathematical Society, Providence, RI, 1963.
6. E. Demaine, M. Demaine, and H. Verrill, Coin-moving puzzles, in *More Games of No Chance* (R. J. Nowakowski, editor), pp. 405–431, Cambridge University Press, Cambridge, 2002.
7. A. Dumitrescu and J. Pach, Pushing squares around, *Graphs and Combinatorics*, to appear. Also in *Proceedings of the* 20*th Annual Symposium on Computational Geometry* (SOCG '04), New York, June 2004, pp. 116–123.
8. M. Gardner, Penny puzzles, in *Mathematical Carnival*, pp. 12–26, Knopf, Washington, DC, 1975.
9. S. Jadhav and A. Mukhopadhyay, Computing a centerpoint of a finite planar set of points in linear time, *Discrete & Computational Geometry*, **12** (1994), 291–312.
10. S. Kapoor and M. Smid, New techniques for exact and approximate dynamic closest-point problems, *SIAM Journal on Computing*, **25** (1996), 775–796.
11. K. Kedem, R. Livne, J. Pach, and M. Sharir, On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles, *Discrete & Computational Geometry*, **1** (1986), 59–71.
12. K. Mehlhorn and S. Näher, Dynamic fractional cascading, *Algorithmica*, **5** (1990), 215–241.
13. J. Pach, Decomposition of multiple packing and covering, 2. *Kolloquium über Diskrete Geometrie*, Salzburg, 1980, pp. 169–178.
14. M. Smid, Closest-point problems in computational geometry, in *Handbook of Computational Geometry* (J.-R. Sack and J. Urrutia, editors), pp. 877–935, North-Holland, Amsterdam, 2000.

Received July 11, 2005, *and in revised form December* 20, 2005. *Online publication March* 20, 2006.