The Union of Moving Polygonal Pseudodiscs – Combinatorial Bounds and Applications

Mark de Berg* Hazel Everett† Leonidas J. Guibas‡

Abstract

Let $\mathcal{P}$ be a set of polygonal pseudodiscs in the plane with n edges in total translating with fixed velocities in fixed directions. We prove that the maximum number of combinatorial changes in the union of $\mathcal{P}$ is $\Theta(n^2\alpha(n))$. In general, if the pseudodiscs move along curved trajectories, then the maximum number of changes in the union is $\Theta(n\lambda_{s+2}(n))$, where s is the maximum number of times any triple of polygon edges meet in a common point.

We apply this result in two different settings. First, we prove that the complexity of the free space of a constant-complexity polygon translating amidst convex polyhedral obstacles with n edges in total is $O(n^2\alpha(n))$. Second, we show that the complexity of the space of lines missing a set of n convex homothetic polytopes of constant complexity in 3-space is $O(n^2\lambda_4(n))$. Both bounds are almost tight in the worst case.

1 Introduction

Let $\mathcal{P}$ be a set of polygons in the plane with n edges in total. Each polygon translates with a fixed velocity in a fixed direction. Our goal is to bound the number of changes in the combinatorial structure of the union of the polygons. For arbitrary convex polygons this problem is easy: the number of changes is $O(n^3)$, and this bound is tight in the worst case—see Section 2. If we put some restrictions on the polygons then the problem becomes more challenging. In this paper we study the case where the set $\mathcal{P}$ is a collection of pseudodiscs: at any time t, the boundaries of any two polygons intersect in at most two points, or in one segment. In Section 2 we prove that in this case the maximum number of changes in the union is only $\Theta(n\lambda_3(n))$, where $\lambda_q(n)$ is the maximum length of an (n,q)-Davenport-Schinzel sequence. The function $\lambda_q(n)$ is roughly linear in n for any constant q [1]; for instance, $\lambda_3(n) = \Theta(n\alpha(n))$, where $\alpha(n)$ denotes the extremely slowly growing functional inverse of the Ackermann function.

Our proof technique is robust with respect to the type of motions: if the polygons move along curved trajectories such that any vertex of one polygon crosses an edge of another

*Department of Computer Science, Utrecht University, P.O.Box 80.089, 3508 TB Utrecht, the Netherlands. Supported by the Dutch Organization for Scientific Research (N.W.O.) and by ESPRIT Basic Research Action No. 7141 (project ALCOM II: *Algorithms and Complexity*)

†Département d'Informatique, Université du Québec à Montréal, Case postale 8888, Succursale Centre-Ville, Montréal, Québec, Canada, H3C 3P8. Supported by NSERC and FCAR.

‡Department of Computer Science, Stanford University, Stanford, CA 94305. Supported by NSF/ARPA grant IRI-9306544, NSF grant CCR-9215219, and by the Stanford Integrated Manufacturing Association.

polygon a constant number of times, and any triple of edges meet in a common point at most s times, then we can show that the number of changes in the union is still roughly quadratic, namely $O(n\lambda_{s+2}(n))$. The polygons are even allowed to deform during the motions, as long as they remain a collection of polygonal pseudodiscs at any time.

The generality of our result makes it easy to apply in various settings. Suppose for example that we want to bound the complexity of the free space of a polygonal convex robot R of constant complexity moving amidst a collection of polygonal convex obstacles in the plane with n edges in total. If the robot is allowed to translate only, then the free space can be described as the complement of the union of a set of pseudodiscs—namely the Minkowski sums of the obstacles with the image of R under central reflection—which has $O(n)$ complexity [6]. When the robot is allowed to rotate also, we need three parameters to describe a placement: x- and y-coordinate of a reference point, and its orientation θ. Hence, the configuration space is three-dimensional. Notice that every cross-section of constant θ of the configuration space consists of a set of polygonal pseudodiscs. So if we sweep a plane h in θ-direction through configuration space, we get a set of moving pseudodiscs. The number of features of the free space corresponds exactly to the number of changes in the union of the pseudodiscs in the cross-section during the sweep. Thus our result applies, provided we can bound the number of times three edges of the pseudodiscs meet during the sweep. This way we can obtain the same bound on the complexity of the free space of a translating and rotating robot in the plane as Leven and Sharir [7]. In fact, our proof technique for bounding the changes in the union of moving pseudodiscs is very similar to their technique for bounding the free space complexity in the above-mentioned problem. The main difference is that our description is more general, which makes it directly applicable in many problems. We illustrate this with two examples, which are described in detail in Section 3.

The first example is another motion planning problem: we want to bound the complexity of the free space of a convex polygon translating amidst convex polyhedral obstacles in 3-space with n edges in total. Halperin and Yap [5] study this problem for the case where the robot is a triangle. By adapting the proof technique of Leven and Sharir they show an $O(n^2\alpha(n))$ bound on a certain type of features of the free space; they do not bound the total number of features. Our result on moving pseudodiscs is immediately applicable in this setting, yielding an $O(n^2\alpha(n))$ bound on the total free space complexity. It is not necessary here that the robot is a triangle; it can be any constant complexity convex polygon. Our technique does not work, however, for polyhedral robots. For this more general case, Aronov and Sharir [2] proved an $O(n^2\log^2 n)$ bound.[1] (Although our motion planning application is a special case of the problem considered by Aronov and Sharir, their results do not carry over to our setting of moving and deforming pseudodiscs in the plane.)

The second example of the applicability of our moving pseudodiscs result deals with a seemingly very different problem, namely the interaction between lines and polyhedra in 3-space. In many problems this interaction plays a fundamental role—the ray shooting problem is an example. It is therefore essential to understand the combinatorial and algorithmic issues involved in the interaction [4]. A natural way to classify lines with respect to a given set of polyhedra is to distinguish between lines that intersect one or more of the polyhedra, and lines that miss all polyhedra. For an arbitrary set of polyhedra the maximum complexity of the

[1] Aronov and Sharir do not assume that the polyhedra have constant complexity. For k Minkowski sums with a total of n edges, their bound is $O(nk\log^2 k)$.

space of missing lines is $\Theta(n^4)$, but in special cases better bounds are sometimes possible. For instance, Pellegrini [8] proved that the space of lines missing a starshaped polyhedron (such as a polyhedral terrain) with n edges is $O(n^3\beta(n))$, where $\beta(n)$ is a slowly growing function. (To be precise: $\beta(n) = 2^{c\sqrt{\log n}}$ for some constant c.) Using our result on moving pseudodiscs, we prove that the space of lines missing a set of n constant complexity convex homothets in 3-space, such as axis-aligned cubes, is also roughly cubic, namely $O(n^2\lambda_4(n)) = O(n^3\alpha(n))$.

2 The combinatorics of moving pseudodiscs

Let $\mathcal{P} = \{P_1, P_2, \ldots, P_k\}$ be a set of convex polygons in the plane with n edges in total. Each polygon P_i is moving along a trajectory ϑ_i. We denote the polygon P_i at time t by $P_i(t)$. We consider the case where the polygons are pseudodiscs. More precisely, we require that the set of polygons $P_i(t)$ forms a set of pseudodiscs at any time t. (A set of polygons is called a set of pseudodiscs if the boundaries of any pair of polygons intersect in at most two points, or in a connected set. The Minkowski sums of a set of disjoint obstacles with a fixed robot in the plane is an example of a set of pseudodiscs [6].) The union of the polygons $P_i(t)$ changes continuously as t increases. At certain times, however, the combinatorial structure of the union changes. The goal of this section is to prove the following theorem on the number of combinatorial changes.

Theorem 2.1 *Let $\mathcal{P}$ be a set of moving polygonal pseudodiscs in the plane with n edges in total. Suppose that the trajectories satisfy the following properties:*

- *Any vertex of one polygon crosses any edge of another polygon a constant number of times.*
- *Any triple of edges from three distinct polygons are concurrent (that is, meet in a common point) at most s times.*

Then the maximum number of changes in the union of the pseudodiscs is $\Theta(n\lambda_{s+2}(n))$.

If the pseudodiscs translate with fixed velocities into fixed directions, then any triple of edges meets at most once, and the maximum number of changes in the union will be $\Theta(n\lambda_3(n)) = \Theta(n^2\alpha(n))$.

In fact, our proof applies in a slightly more general setting than that of Theorem 2.1, which is summarized in the corollary below. Because the setting of Theorem 2.1 is more intuitive we shall describe the proof in this setting. It holds, however, almost verbatim in the setting of Corollary 2.2.

Corollary 2.2 *Let $\mathcal{P}$ be a set of moving and deforming objects in the plane that satisfies the following conditions:*

- *At any time, $\mathcal{P}$ forms a collection of convex polygonal pseudodiscs. (This implies that the edges of the polygons may deform but should remain straight line segments.)*
- *Edges may appear or disappear on the boundary of the objects, but the total number of edges ever appearing on the boundary of the objects is n in total.*

- *Any vertex of one polygon crosses any edge of another polygon a constant number of times.*
- *Any triple of edges from three distinct polygons are concurrent (that is, meet in a common point) at most s times.*

Then the maximum number of changes in the union of $\mathcal{P}$ *is* $\Theta(n\lambda_{s+2}(n))$.

The upper bound

We now prove the upper bound in Theorem 2.1. To simplify the presentation, we make the following general position assumption: at no time t are there four edges going through the same point. This is allowed because a degenerate situation can never lead to a maximum number of union changes. Now every combinatorial change in the union of the pseudodiscs corresponds to an event of one the following two types:

VE-event: A vertex of some polygon P_i crosses an edge of some other polygon P_j.

EEE-event: An edge of some polygon P_i crosses the intersection point of two edges of distinct polygons P_j, P_k.

(In fact, VE-events can be seen as special cases of EEE-events, where two of the three edges come from the same polygon.) We are only interested in the *external events*, that is, the events taking place on the union boundary. Note that if an edge becomes incident with another edge parallel to it, there is always a VE-event involved. Hence, by bounding the number of VE-events we also bound the number of such parallel-edge events.

The first property of the trajectories of the pseudodiscs, namely that a vertex crosses an edge a constant number of times, immediately implies a bound on the number of VE-events:

Observation 2.3 *The number of VE-events is* $O(n^2)$.

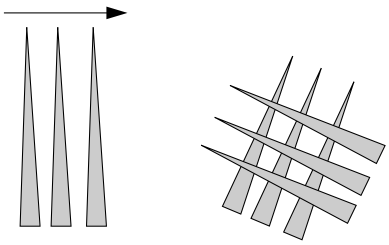

Figure 1: A cubic number of EEE-events.

The EEE-events form the difficult case. Potentially every triple of edges may give rise to a EEE-event, leading to a cubic number of EEE-events. If the polygons are arbitrary, there can even be a cubic number of *external* EEE-events. Figure 1 gives an example of this: a group of $n/3$ triangle moves to the right, and passes a group of $2n/3$ stationary triangles that form a grid. If the polygons are pseudodiscs, however, it is not possible to construct a grid-like union. In fact, the union of a collection of pseudodiscs has linear complexity [6]. In the remainder of this section we prove that the number of external EEE-events for a set of moving pseudodiscs is roughly quadratic, which implies a similar bound on the number of changes in their union.

Consider a fixed edge e of a polygon P_i. Without loss of generality, we may assume that e is stationary. (Think of an observer standing on e; this observer sees the other polygons pass by, and sees e as stationary.)

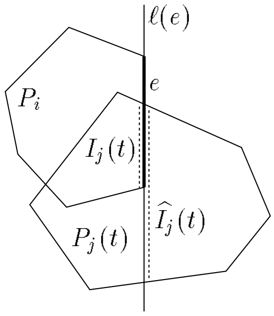

Figure 2: The intervals on edge e.

For $j \neq i$, we define $I_j(t) := e \cap P_j(t)$, that is, $I_j(t)$ denotes the intersection of polygon $P_j(t)$ with e at time t. We also define $\widehat{I}_j(t) := \ell(e) \cap P_j(t)$, where $\ell(e)$ denotes the line through e. Figure 2 illustrates these definitions. As the polygons move, the intervals $I_j(t)$ and $\widehat{I}_j(t)$ change. An external EEE-event involving the edge e corresponds to a change in the union of the intervals $I_j(t)$. Thus, there is an external EEE-event involving e at time t^* if the following situation arises: there are two intervals $I_j(t^*)$ and $I_k(t^*)$ sharing an endpoint p that lies in the interior of e and is not contained in any interval $I_l(t^*)$. We say that e *witnesses* this event if both $\widehat{I}_j(t^*)$ and $\widehat{I}_k(t^*)$ contain an endpoint of e. (This notion corresponds to the notion of one obstacle 'bounding' another obstacle in the paper by Leven and Sharir [7].)

Lemma 2.4 *An edge e witnesses $O(\lambda_{s+2}(n))$ EEE-events.*

Proof: Consider the *edge-time diagram of e*, which is defined as follows. The vertical axis in the diagram represents the line $\ell(e)$, and the horizontal axis represents time. For each P_j we draw the intervals $I_j(t)$ in the diagram, for $t \geqslant 0$. Thus we draw for each polygon the set $I_j := \bigcup_{t \geqslant 0} I_j(t)$. Figure 3 gives an example of an edge-time diagram for the case of pseudodiscs translating with fixed velocities in fixed directions. In that case the sets I_j are convex polygons, but this is not true in general. The number of (curved) edges on the boundary of I_j equals the number of edges of P_j, so overall there are $O(n)$ edges in the edge-time diagram. As observed earlier, an external EEE-event involving e corresponds to a change

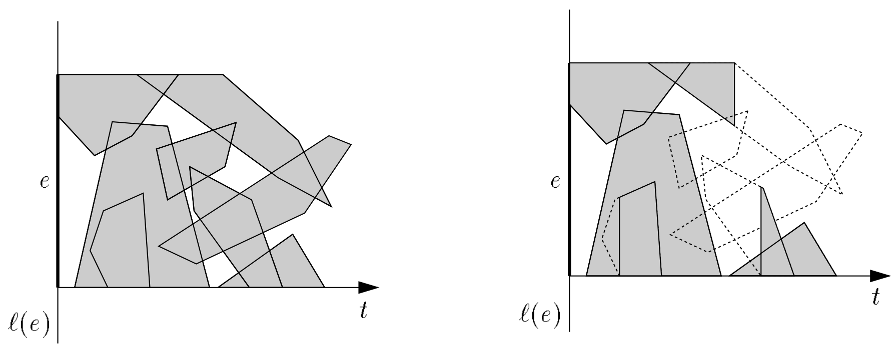

Figure 3: The edge-time diagram, and its modified version.

in the union of the intervals $I_j(t)$. Such a change in turn corresponds to a vertex of the union of the sets I_j in the edge-time diagram of e. However, not all such vertices correspond to EEE-events witnessed by e. Therefore we modify the diagram by drawing only the intervals $I_j(t)$ witnessed by e—see Figure 3. A EEE-event witnessed by e must be a vertex in this modified diagram. (The reverse in not necessarily true: a vertex in the modified diagram can lie inside the part of some I_j that was deleted in the modification. Such a vertex corresponds to a EEE-event that is not external. Since we are deriving an upper bound on the number of changes in the union boundary, we are allowed to count these events.) A set I_j in the modified diagram is 'grounded' to either the horizontal line through one endpoint of e or to the horizontal line through the other endpoint, that is, it is a histogram whose base lies on one of these lines. Consider the histograms whose bases lie on the lower of the two lines. A vertex of their union is a vertex of the upper envelope of the curved segments that form the boundary of these histograms. The trajectories of the pseudodiscs have the property that any triple of edges meets at most s times in a common point. Hence, two curved segments in the edge-time diagram intersect at most s times, so the upper envelope has complexity $O(\lambda_{s+2}(n))$ [1]. Similarly, the vertices of the union of the histograms whose bases are on the higher of the two lines are on a lower envelope, so there are only $O(\lambda_{s+2}(n))$ of them. The remaining vertices in the modified edge-time diagram are intersections between these two envelopes, of which there are $O(\lambda_{s+2}(n))$ as well. □

Lemma 2.4 does not count all the external events involving e, only the ones that are witnessed by e. Indeed, it can be shown that the total number of external events involving e can be as large as $\Theta(n\lambda_{s+2}(n))$. However, a EEE-event involves three edges. We might hope that any external EEE-event is witnessed by at least one of these edges. Unfortunately this is not always true. But it is still possible to prove that the total number of EEE-events is $O(n\lambda_{s+2}(n))$—see the next lemma—thus finishing the proof of the upper bound in Theorem 2.1.

Lemma 2.5 *The total number of EEE-events is $O(n\lambda_{s+2}(n))$.*

Proof: The key observation in the proof is the following. Let e and e' be intersecting edges of pseudodiscs P and P', respectively. Then e has an endpoint inside P', or e' has an endpoint inside P (or both), because the boundaries of P and P' intersects in at most two points. This implies that if there is a EEE-event involving edge e_1 of P_1, edge e_2 of P_2, and edge e_3 of P_3 then at least one of the following two cases occurs:

- there is an edge e_i with an endpoint inside both polygons P_j, for $i, j \in \{1, 2, 3\}$ and $j \neq i$, or
- e_1 has an endpoint inside P_2, e_2 has an endpoint inside P_3, and e_3 has an endpoint inside P_1 (or e_1 has an endpoint inside P_3, e_3 has an endpoint inside P_2, and e_2 has an endpoint inside P_1).

In the first case e_i witnesses the EEE-event. According to Lemma 2.4 there are only $O(\lambda_{s+2}(n))$ such events for e_i, so the total number of such events is $O(n\lambda_{s+2}(n))$.

In the second case the event may not be witnessed by any of the three edges. However, we shall prove that the total number of such unwitnessed events is linear in the number of witnessed events. Let t denote the time of such an event, involving edges e_1, e_2, and e_3 of polygons P_1, P_2, and P_3. The characterization above tells us that at time t the following holds (possibly after renumbering the polygons): polygon P_2 is on an (upper or lower) envelope in the modified edge-time diagram of e_1, polygon P_3 is on an envelope in the modified edge-time diagram of e_2, and polygon P_1 is on an envelope in the modified edge-time diagram of e_3. Let $t' < t$ be the last time before time t at which any of the three envelopes changed. In other words, t' corresponds to the last breakpoint before time t on one of the three envelopes. We charge the unwitnessed EEE-event at time t to the witnessed event at t', and we claim that any witnessed event can be charged at most twice this way. Indeed, consider a breakpoint on an envelope in the modified edge-time diagram of some edge e. After the breakpoint, the envelope is defined by some edge e'; the breakpoint can only be charged by a EEE-event involving e, e', and some third edge e''. There are only two candidates for this third edge, because it has to be on the upper or lower envelope in the modified edge-time diagram of e' at the time of the breakpoint. This proves the claim and, hence, the lemma. □

The lower bound

We now describe an example of a set $\mathcal{P}$ of n moving pseudodiscs in the xy-plane (actually, they will be squares) whose trajectories have the properties mentioned in Theorem 2.1 and whose union changes $\Theta(n\lambda_{s+2}(n))$ times.

The first $\lfloor n/2 \rfloor$ pseudodiscs are the squares $P_i := [2i : 2i+1] \times [0 : 1]$, for $0 \leqslant i \leqslant \lfloor n/2 \rfloor - 1$. These squares are all stationary.

To construct the second half of the set of pseudodiscs, we first go to an auxiliary plane, namely the ty-plane. The y-axis in this plane corresponds to the y-axis in the xy-plane, and the t-axis represents time. Let $A = \{a_1, \ldots, a_{\lceil n/2 \rceil}\}$ be a set of $\lceil n/2 \rceil$ t-monotone arcs in the ty-plane such that any pair of arcs intersects at most s times and whose upper envelope has $\Theta(\lambda_{s+2}(n))$ complexity. (Such a set can be constructed as follows: Take a set of n continuous functions, each pair of which intersects at most $s + 2$ times, and whose upper envelope has $\Theta(\lambda_{s+2}(n))$ complexity [3]. For each pair of functions, consider the first intersection point. Before this intersection point, one of the two functions cannot contribute to the upper envelope; delete this part, thereby reducing the number of intersections of the pair by one. Similarly, delete a part of one of the functions after the last intersection. After this has been done for every pair of functions, we have a collection of n arcs with at most s intersections per pair that has the same upper envelope as the original collection of functions.) Scale the set A such that all arcs lie strictly between the two horizontal lines $y = 0$ and $y = 1$. The idea is to create for each arc $a_j \in A$ a square Q_j with the following property: if (t^*, y^*) is a point on a_j, then at time t^* the top edge of Q_j has y-coordinate y^* and it intersects each of the stationary squares P_i. Figure 4 illustrates this. This requirement implies that the vertical velocity of Q_j at time t^* is determined by the tangent of the arc a_j at t-coordinate t^*. At times t^* when a_j is not defined, (that is, the line $t = t^*$ does not intersect a_j) Q_j should not intersect any of the squares P_i; to achieve this we park the square somewhere far enough from the scene. The parking lot should be such that squares can move in and out without intersecting any of the other parked squares. The set $\mathcal{P}$ consists of the stationary squares P_i and the squares Q_j.

Suppose for a moment that it is possible to construct the set $\mathcal{P}$. We claim that then there are $\Theta(n\lambda_{s+2}(n))$ changes in the union of $\mathcal{P}$. Indeed, consider a vertex (t^*, y^*) of the envelope of the set A. Let a_j be the arc that is on the envelope at time $t^* - \epsilon$, and let a_k be the arc

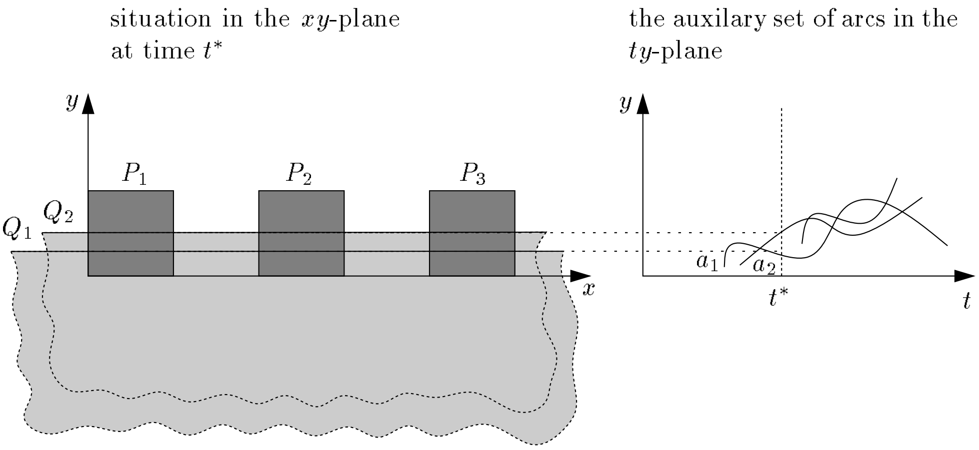

Figure 4: The lower bound construction.

that is on the envelope at time $t^* + \epsilon$, where ϵ is an infinitesimally small positive constant. In the xy-plane we then have the following situation. At time $t^* - \epsilon$ the top edge of Q_j intersects all stationary squares P_i, and it is the highest edge intersecting these squares. Thus the top edge of Q_j contributes a linear number of vertices to the union. At time $t^* + \epsilon$, this is true for Q_k. Hence, at time t^*, when the top edge of Q_k overtakes the top edge of Q_j, there are $\Theta(n)$ changes in the union of $\mathcal{P}$. Since the upper envelope of S has $\Theta(\lambda_{s+2}(n))$ vertices, the total number of changes in the union is $\Theta(n\lambda_{s+2}(n))$.

There is one problem, however: it is not possible to construct the squares Q_j such that they have exactly the properties stated above. In particular, the properties imply that each Q_j would appear instantaneously at the time corresponding to the t-coordinate of the left endpoint of a_j: just before this time it should not intersect any of the stationary squares P_i, and just after this time it should intersect all of them. This problem is easy to overcome. If a square Q_j must appear at time t^*, we simply stop the motion of the other squares temporarily and move Q_j from its parking spot to the place where it must appear. This motion is such that no horizontal edge of Q_j crosses a horizontal edge of another square and such that Q_j approaches the squares P_j from the left. When the square should disappear, we again stop the other squares and move the square back to its parking spot by first moving Q_j to the right and then back to its parking spot. Figure 5 illustrates how a square is moved from its parking spot to the place where it must appear. This way all the events corresponding to vertices on the envelope of A are realized. Moreover, any vertex of one square crosses an edge of another square at most a constant number of times. We claim that any triple of edges meets at most s times in a common point. To see this, observe that one of the edges of the triple must be a vertical edge belonging to a stationary square. If one of the other two edges is vertical as well, there is only one event for this triple of edges, due to the parking regulations. If both other edges are horizontal, the event corresponds to an intersection between two arcs in A, so there are at most s events for the triple.

We have constructed a set of n moving squares whose union changes $\Theta(n\lambda_{s+2}(n))$ times. This finishes the proof of Theorem 2.1.

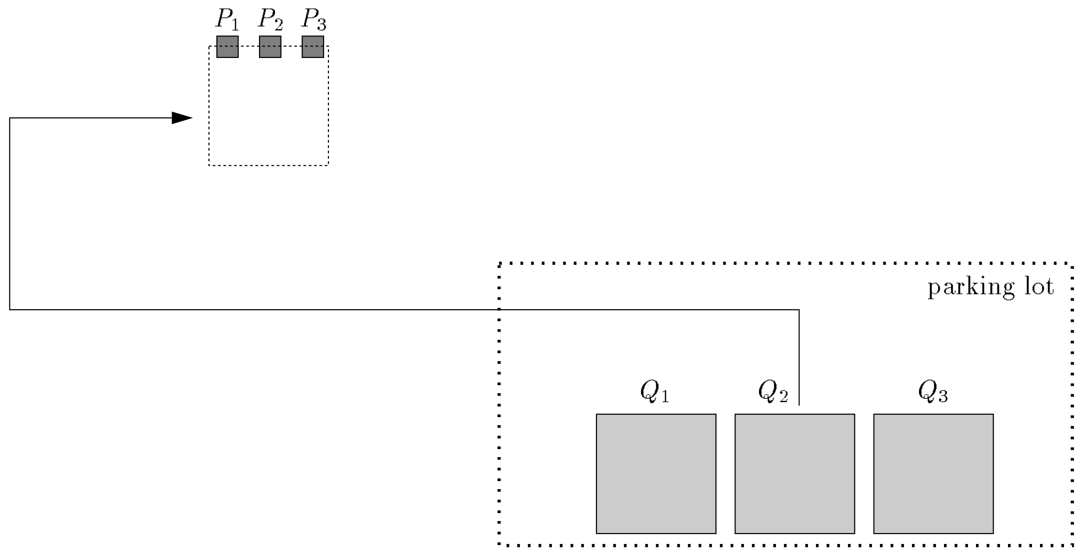

Figure 5: Moving a square out of the parking lot.

Remark 2.6 The lower bound construction for $s = 1$ as described above does not yield a set of pseudodiscs that move with constant velocity along a straight line. However, by choosing the sizes and velocities of the big squares carefully, it is possible to get $\Theta(n^2\alpha(n))$ changes in this case as well.

3 Applications

In this section we give two applications of the result we derived in the previous section.

3.1 A polygon translating amidst polyhedra in 3-space

Let $\mathcal{S}$ be a set of convex and disjoint polytopes in 3-space with n edges in total, and let R be a convex polygon of constant complexity. The polygon R is a robot that can translate (but not rotate) in 3-space, and the polyhedra in $\mathcal{S}$ are obstacles, which the robot is not allowed to intersect. Because R is not allowed to rotate, a placement of R can be described by the three coordinates of some reference point r. The space of all such placements is called the *configuration space* of the robot, and the space of all placements where the robot does not intersect any of the obstacles is called the *free space*. The free space is the complement of the union of the Minkowski sums of the obstacles with $-R$, the image of the robot under central reflection. These Minkowski sums are also called *configuration space obstacles*. Our goal is to prove the following bound on the complexity of the free space.

Theorem 3.1 *Let $\mathcal{S}$ be a set of convex disjoint polytopes in 3-space with n vertices in total, and let R be a robot that is a convex polygon of constant complexity. Then the free space of R under translational motions is $O(n^2\alpha(n))$ and $\Omega(n^2)$.*

The upper bound

Assume without loss of generality that R is parallel to the xy-plane. Consider a cross-section of the configuration space with a plane $h : z = c$, for some constant c. This cross-section describes the placements of R where its reference point lies in h. Because R is parallel to the the xy-plane, the cross-section of the configuration space on h is only influenced by the cross-sections of the obstacles with h. More precisely, the cross-section of h with any of the configuration space obstacles equals the Minkowski sum of the cross-section of the obstacle with $-R$. Hence, the cross-sections of the configuration space obstacles form a collection of polygonal pseudodiscs. So if we consider the cross-sections on the family of planes $h : z = c$, where c ranges from $-\infty$ to $+\infty$, we get a collection of moving and deforming pseudodiscs. We are thus in the position to apply Corollary 2.2. It remains to bound the three parameters in Corollary 2.2: the overall number of edges appearing on the pseudodisc boundaries, the number of times a vertex can cross an edge, and the number of times a triple of edges can meet.

Let's take a closer look at the pseudodiscs in a cross-section of the configuration space with a plane $h : z = c$. An edge of such a pseudodisc is the cross-section of a face of a configuration space obstacle with h, and a vertex of such a pseudodisc is the cross-section of an edge of a configuration space obstacle with h. Hence, the total number of edges that appear on the boundary of any of the pseudodiscs during the sweep equals the total number of faces of the configuration space obstacles, which is $O(n)$. Furthermore, a triple intersection between edges of the pseudodiscs corresponds to a point where the three corresponding faces of configuration space obstacles meet. Because the faces are planar, there is only such point. (We may assume general position.) Similarly, a vertex of a pseudodisc crosses an edge of another pseudodisc at most once during the sweep, because an edge of a configuration space obstacle intersects a face of another configuration space obstacle in at most one point. Hence, we may apply Corollary 2.2 with $s = 1$, and we get the upper bound stated in Theorem 3.1.

The lower bound

The obstacles in the lower bound construction are long and skinny tetrahedra placed as follows. Construct a 2D-grid of triangles in the xz-plane, similar to the grid depicted in Figure 1. Move half of these triangles slightly away from the xz-plane and inflate them a little to obtain a collection of disjoint tetrahedra. The collection of tetrahedra looks like a 2D-grid when viewed from $y = \infty$. The robot R is a skinny triangle parallel to the xy-plane, which is very long in y-direction. Every hole in the grid gives rise to a number of free placements where R touches two obstacle edges, which proves that the complexity of the free space is $\Omega(n^2)$.

3.2 Lines missing convex homothets in space

Let $\mathcal{S}$ be a set of n convex homothetic polytopes of constant complexity in 3-space. (A set of objects is called homothetic if they are identical up to translation and scaling. A set of cubes is an example of a set of homothetic polytopes.) We do not assume that the polytopes are disjoint. We call a line a *free line* if it does not intersect the interior of any of the polytopes from $\mathcal{S}$. The goal of this section is to prove the bound on the combinatorial complexity of $\mathcal{L}(\mathcal{S})$, the set of all free lines, stated in the following theorem.

Theorem 3.2 *Let $\mathcal{S}$ be a set of n convex homothetic polytopes of constant complexity in 3-space. The maximum complexity of the set $\mathcal{L}(\mathcal{S})$ of free lines is $O(n^2\lambda_4(n))$ and $\Omega(n^3)$.*

The upper bound

The complexity of $\mathcal{L}(\mathcal{S})$ is determined, up to a cubic factor, by the number of lines touching four edges of four distinct polytopes and missing all other polytopes. (For disjoint polytopes this is easy to see; for intersecting polytopes it requires a little more thought, but it is still true.) In general, the latter number can be $\Theta(n^4)$: Take a set of long and skinny tetrahedra that form a grid when viewed from $x = \infty$, copy this set and translate the new grid some distance into the x-direction. For every pair of holes, one from the first grid and one from the copied one, there is a unique combination of four edges that can be touched by a line through the two holes. If the polytopes are homothetic a grid-like construction cannot be made, and one would expect that the complexity of $\mathcal{L}(\mathcal{S})$ is less than $\Theta(n^4)$. In this section we prove that this is indeed true, and that the complexity is roughly cubic.

For a free line ℓ touching four edges, we can order the four edges from left to right along ℓ. We charge ℓ to the first of these edges. We shall prove that each edge gets charged $O(n\lambda_4(n))$ free lines in this manner.

Fix an edge e of one of the polytopes. Assume without loss of generality that e is contained in the plane $h_0 : x = 0$. From now on, we only consider the part of each polytope that lies to the right of h_0. Let ℓ be a free line touching four edges, of which e is the first one. Let p be the point where ℓ touches e, and project the polytopes onto the plane $h_1 : x = 1$, with p as the center of projection. Let x be the point where ℓ intersects h_1. The point x is the common intersection of three edges of three distinct projected polytopes. Because ℓ is free, x lies on the boundary of the union of the projected polytopes. Hence, the number of times we charge a free line to e equals the number of points $p \in e$ for which three edges of projected polytopes meet in a common point that is on the boundary of the union of the projected polytopes. (More precisely, for each such point p we should count the number of triples of concurrent edges.) The idea of the proof is to move a point p along e, and see how often a triple of concurrent edges arises as the projected polygons move (and change shape) on the projection plane h_1. Clearly, the number of such events is bounded by the number of combinatorial changes in the union of the projected polytopes. To use the results of the previous section we must prove that the projections are pseudodiscs. For parallel projections this is clear; in that case the projections are even homothets. That the projected polytopes are also pseudodiscs for perspective projections is proved next.

Let q_1 and q_2 be the two endpoints of e, and let $p(t) := (1-t)q_1 + tq_2$. The point $p(t)$ moves from q_1 to q_2 as t increases from 0 to 1. Let $\overline{P}(t)$ denote the projection of the polytope P onto the plane h_1, with $p(t)$ as the center of projection.

Lemma 3.3 *The set $\{\overline{P}(t) : P \in \mathcal{S}\}$ is a set of pseudodiscs for all $0 \leqslant t \leqslant 1$.*

Proof: Let P and P' be two polytopes from $\mathcal{P}$. We assume that no clipping was necessary for P and P', that is, that they were entirely to the right of the plane h_0. The proof can easily be adapted to the case where one or both of the polytopes lies partly to the left of h_0. Because P and P' are homothets, there is a point $q \in \mathbb{R}^3$—the *center of similarity* of P and P'—such that P' can be obtained by scaling P with respect to q.

Consider the projections $\overline{P}(t)$ and $\overline{P'}(t)$ for an arbitrary t with $0 \leqslant t \leqslant 1$. Let $\overline{q}(t)$ denote the projection of q onto h_1, with $p(t)$ as the center of projection. Because q is the center of similarity of P and P', there are two lines through $\overline{q}(t)$ that touch both $\overline{P}(t)$ and $\overline{P'}(t)$, as illustrated in Figure 6. The points where these two lines touch $\overline{P}(t)$ split $\partial\overline{P}(t)$ into two

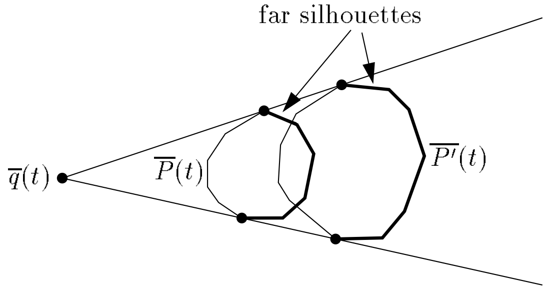

Figure 6: Illustration for the proof of Lemma 3.3.

pieces. We call the piece closest to $\overline{q}(t)$ the *near silhouette* of $\overline{P}(t)$, and we call the farthest piece the *far silhouette* of $\overline{P}(t)$. The near and far silhouette of $\overline{P'}(t)$ are defined analogously. The far silhouette of one of the polygons can intersect the near silhouette of the other polygon at most twice. To prove that $\overline{P}(t)$ and $\overline{P'}(t)$ are pseudodiscs it thus remains to show that the two near silhouettes, and the two far silhouettes, do not intersect each other.

To see that this is true we go back to three dimensions. Recall that q was a point such that P' can be obtained by scaling P with respect to q. Assume without loss of generality that P' is the bigger of the two polytopes. Consider the set of rays starting at q that are tangent to P (and, hence, to P'). The points of tangency on P form a curve on the boundary of P, which partitions the boundary of P into two pieces. The piece closest to q is called the *near half*, and the piece farthest from q is called the *far half* of the boundary. Note that the near and far silhouettes of $\overline{P}(t)$ are projections of pieces of the near half and far half of the boundary of P, respectively. Similarly the points of tangency on P' partitions its boundary into a near and a far half. If we take the two far halves of P and P' and consider all the segments joining q to them, we obtain two other homothetic convex solids F and F'. (In fact, F is the convex hull of $P \cup \{q\}$, and F' is the convex hull of $P' \cup \{q\}$.) The solid F is entirely contained in F'. Hence, from any point of view the silhouette of F has to be inside the silhouette of F', which implies that the far silhouettes do not intersect in the projection.

In a similar way it can be shown that the near silhouettes of $\overline{P}(t)$ and $\overline{P'}(t)$ do not intersect. □

The projections $\overline{P}(t)$ are pseudodiscs, which we need to apply the results of the previous section. Notice that the pseudodiscs change their shape during the motion, but this is of no importance; the only properties of the polygons that we need—besides that they are pseudodiscs at all times—are that a vertex crosses an edge a constant number of times, and that a triple of edges meets at most s times in a common point, for some parameter s—see Corollary 2.2.

The first property is satisfied because for a given vertex v of some polytope, and a given edge e' of another polytope, there is at most one line through e, v, and e'. Furthermore there are only two lines through any four-tuple of edges. Hence, we may apply Theorem 2.1 with $s = 2$ to bound the number of free lines touching a fixed edge e and three other edges. Summing over all edges e, we conclude that the total number of free lines touching four distinct edges is $O(n^2\lambda_4(n))$. This finishes the proof of the upper bound of Theorem 3.2.

The lower bound

We describe an example of a set $\mathcal{S}$ of n cubes in 3-space such that $\mathcal{L}(\mathcal{S})$ has $\Theta(n^3)$ complexity, thus finishing the proof of Theorem 3.2.

Figure 7 shows a top view of the construction. There are two sets of roughly $n/3$ small

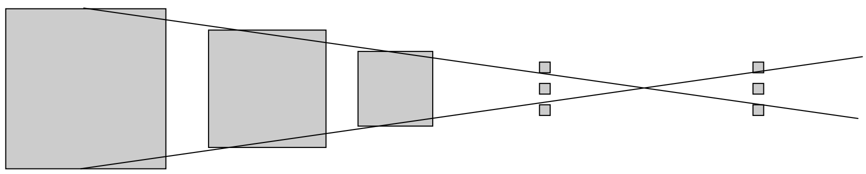

Figure 7: Top view of the lower bound construction.

cubes, whose top faces are in the plane $z = 0$. They are arranged in such a way that they generate a quadratic number of free lines touching two vertical edges, one from the left set and one from the right set. In the figure only the two extreme lines are shown. Then there is a third set, consisting of big cubes; the big cubes have their top faces in the plane $z = 0$ as well. These cubes can be placed in such a way that each of the $\Theta(n^2)$ lines generated by the small cubes touches all the top faces of the big cubes. Finally, we translate each of the big cubes slightly in the positive z-direction, such that the right edges of the top faces of these cubes are on a convex surface. In this way any two adjacent big cubes can be touched with a line that misses the other big cubes. If the translations are small enough, then each of the $\Theta(n^2)$ lines spawned by the small cubes can touch all adjacent pairs of big cubes, thus giving rise to $\Theta(n^3)$ distinct triples of edges that can be touched by a free line.

4 Concluding remarks

We have shown that the maximum number of changes in the union of n translating pseudodiscs is $\Theta(n^2\lambda_3(n))$, and we generalized this result to pseudodiscs moving (and deforming) along curved trajectories. We applied this result to prove bounds on the complexity of the free space of a constant complexity convex polygon translating amidst convex polyhedra in 3-space, and on the complexity of the space of lines missing a set of n convex homothets in 3-space.

We see two major challenges left.

The first is to extend our results to moving discs in the plane. This is equivalent to bounding the complexity of the union of a set of cylinder-like objects in 3-space. (The objects are not real cylinders, as the cross-section of such an object with any horizontal plane is a disc, instead of the cross-section with a plane perpendicular to its axis.) Thus the problem is closely related to bounding the complexity of the Voronoi diagram of a set of lines in 3-space.

The second challenge is to extend the result to fat polygons. Such an extension could be useful to prove a bound on the complexity of the space of lines missing a set of fat tetrahedra in 3-space, in the spirit of Section 3.2. We believe that this would be an important step in understanding the behavior of fat objects in 3-space.

Acknowledgements

We thank Godfried Toussaint for inviting us to the 1994 Barbados Workshop on Computational Geometry, where part of the research was performed. We also thank Otfried Schwarzkopf for his help in the proof of Lemma 3.3.

References

[1] P. K. Agarwal, M. Sharir, and P. Shor. Sharp upper and lower bounds on the length of general Davenport-Schinzel sequences. *J. Combin. Theory Ser. A*, 52:228–274, 1989.

[2] B. Aronov and M. Sharir. On translational motion planning in 3-space. In *Proc. 10th Annu. ACM Sympos. Comput. Geom.*, pages 21–30, 1994.

[3] M. J. Atallah. Dynamic computational geometry. In *Proc. 24th Annu. IEEE Sympos. Found. Comput. Sci.*, pages 92–99, 1983.

[4] B. Chazelle, H. Edelsbrunner, L. J. Guibas, and M. Sharir. Lines in space: combinatorics, algorithms, and applications. In *Proc. 21st Annu. ACM Sympos. Theory Comput.*, pages 382–393, 1989.

[5] D. Halperin and C.-K. Yap. Complexity of translating a box in polyhedral 3-space. In *Proc. 9th Annu. ACM Sympos. Comput. Geom.*, pages 29–37, 1993.

[6] K. Kedem, R. Livne, J. Pach, and M. Sharir. On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles. *Discrete Comput. Geom.*, 1:59–71, 1986.

[7] D. Leven and M. Sharir. On the number of critical free contacts of a convex polygonal object moving in two-dimensional polygonal space. *Discrete Comput. Geom.*, 2:255–270, 1987.

[8] M. Pellegrini. On lines missing polyhedral sets in 3-space. In *Proc. 9th Annu. ACM Sympos. Comput. Geom.*, pages 19–28, 1993.