

On the complexity of the free space of a translating box in $\mathbb{R}^3$

Gabriel Nivasch*

Abstract

Consider a convex polyhedral robot B that can translate (without rotating) amidst a finite set of non-moving polyhedral obstacles in $\mathbb{R}^3$. The *free space* $\mathcal{F}$ of B is the set of all positions in which B is disjoint from the interior of every obstacle.

Aronov and Sharir (1997) derived an upper bound of $O(n^2 \log n)$ for the combinatorial complexity of $\mathcal{F}$, where n is the total number of vertices of the obstacles, and the complexity of B is assumed constant.

Halperin and Yap (1993) showed that, if B is either a box or a “flat” convex polygon, then a tighter bound of $O(n^2 \alpha(n))$ holds. Here $\alpha(n)$ is the inverse Ackermann function.

In this paper we prove that if B is a box, then the complexity of $\mathcal{F}$ is $O(n^2)$. Furthermore, if B is a convex polygon whose edges come in parallel pairs, then the complexity of $\mathcal{F}$ is $O(n^2)$ as well. These results settle the question of the asymptotical worst-case complexity of $\mathcal{F}$ for a box, as well as for all convex polygons.

Keywords: motion planning, computational geometry, Minkowski sum, lower envelope

1 Introduction

One of the most basic problems studied in algorithmic motion planning is that of moving a rigid “robot” B among fixed obstacles. There are many possible instances of the problem, depending on the dimension of the ambient space ($\mathbb{R}^2$, $\mathbb{R}^3$, etc.), the shape of the robot, and the type of motion allowed (e.g. whether the robot is allowed to rotate or just to translate).

Usually we are given an initial placement of the robot, as well as a desired final placement, and the objective is to find a continuous motion that takes the robot from the initial placement to the final one, while avoiding the obstacles at all times (or determine that no such motion exists).

Each placement of the robot can be parametrized by a k -tuple of real numbers, where k is the number of degrees of freedom of the robot. The *configuration space* is thus a k -dimensional space, each point of which corresponds to a placement of the robot. The *free space* of the robot is the subset of the configuration space that corresponds to all placements that are *free* (or *legal*), in the sense that the robot does not intersect any obstacle. In the case where both the robot and the obstacles are polyhedra in $\mathbb{R}^3$ and the robot can only translate, the configuration space is also three-dimensional, and the free space has polyhedral boundaries, consisting of vertices, edges, and faces.

A basic parameter in the analysis of many motion planning algorithms is the *combinatorial complexity* of the free space, meaning the total number of vertices, edges and faces in the boundary of the free space. Hence, a natural question is to determine the worst-case complexity of the free space in different scenarios. See the survey chapter by Halperin, Salzman, and Sharir [2] for more background on algorithmic motion planning.

*Department of Computer Science, School of Computer Science, Ariel University, Ariel, Israel.
gabrieln@ariel.ac.il




Figure 1: Configuration that achieves $\Omega(n^2\alpha(n))$ triple contacts with a triangular robot. The robot is parallel to the xz -plane. There is a configuration of $n/2$ obstacle edges parallel to the xy -plane that form $\Omega(n\alpha(n))$ lower-envelope intersections, in which the edges are given slightly different z -coordinates (a). There are another $n/2$ obstacle edges parallel to the y -axis with different z -coordinates (b).

In this paper we consider the case in which the robot is either a “flat” convex polygon or a convex polyhedron that can only translate, and the obstacles are polyhedra in $\mathbb{R}^3$. We denote the total number of vertices of the obstacles by n . Then the total number of edges and faces of the obstacles is $\Theta(n)$ as well. The number of vertices of the robot is assumed constant.

As will be explained below in more detail, the complexity of $\mathcal{F}$ is asymptotically determined by the number of triple contacts between the robot and obstacles, i.e. the number of free placements in which three distinct robot features (vertices, edges, or faces) simultaneously make contact with three distinct obstacle features.

Hence, a trivial upper bound for the complexity of the free space is $O(n^3)$. Aronov and Sharir [1] improved the upper bound to $O(n^2 \log n)$. Halperin and Yap [3, 4] studied the cases where the robot is a “flat” convex polygon and a three-dimensional box. For both these cases they derived an upper bound of $O(n^2\alpha(n))$, where $\alpha(n)$ is the very slow-growing inverse Ackermann function. The inverse-Ackermann factor arises from consideration of upper and lower envelopes of segments.

Halperin and Yap [4] also proved that if the robot is a rectangle and the obstacles are n lines, then the complexity of $\mathcal{F}$ is $O(n^2)$.

Regarding lower bounds for the complexity of $\mathcal{F}$, a lower bound of $\Omega(n^2)$ is easy to obtain for any robot. For some types of robots (e.g. a triangle) one can obtain a lower bound of $\Omega(n^2\alpha(n))$ (Aronov and Sharir [1]). See Figure 1.

For the two-dimensional case, in which B is a convex polygon free to translate in $\mathbb{R}^2$ and the obstacles are polygons with a total of n vertices, the complexity of $\mathcal{F}$ is $O(n)$ (Kedem, Livne, Pach, and Sharir [6]).

Note that, by affine transformations, the cases where B is a cube or a box or a parallelepiped are all equivalent.




Figure 2: A fully-parallel polygon.

1.1 Our results

In this paper we prove the following:

Theorem 1. *Let B be a box-shaped robot that is free to translate in $\mathbb{R}^3$ amidst polyhedral obstacles that have a total of n vertices. Then the complexity of the free space of B is $O(n^2)$.*

We also study polygonal robots in more depth. We call a convex polygon *fully parallel* if its edges come in parallel pairs (not necessarily of the same length). See Figure 2. It is easy to see that the lower bound of $\Omega(n^2\alpha(n))$ shown in Figure 1 can be achieved with any polygonal robot that is not fully parallel, as well as with any polyhedral robot that can be projected to form a non-fully-parallel polygon.

We prove the following:

Theorem 2. *Let B be a fully-parallel convex polygon that is free to translate in $\mathbb{R}^3$ amidst polyhedral obstacles that have a total of n vertices. Then the complexity of the free space of B is $O(n^2)$ (where the hidden constant depends on B).*

Hence, the asymptotical worst-case complexity of $\mathcal{F}$ is now fully determined if B is a box (or parallelepiped), or any convex polygon.

2 Preliminaries

Let $B \subset \mathbb{R}^3$ be a polyhedral robot that can translate among a set $\mathcal{C} = \{C_1, \dots, C_k\}$ of pairwise-disjoint polyhedral obstacles in $\mathbb{R}^3$. A placement of B in space can be specified by a point v in *configuration space* $\mathbb{R}^3$. At that placement, the robot occupies the points $B + v = \{b + v : b \in B\}$. Such a placement is *free* (or *legal*) if $B + v$ is disjoint from the interior of every obstacle in $\mathcal{C}$. The set of all points in configuration space that correspond to free placements of the robot is called the *free space* of the robot, which we denote by $\mathcal{F}$.

An observation dating back to Lozano-Pérez and Wesley [7] is that the robot can be replaced by a point if the obstacles are appropriately expanded. Specifically, the free space is given by

$$\mathcal{F} = \mathbb{R}^3 \setminus \bigcup_{C \in \mathcal{C}} (\text{int}(C) \oplus (-B))$$

where $X \oplus Y = \{x + y : x \in X, y \in Y\}$ denotes the *Minkowski sum* of two sets of points.

We will assume for simplicity that the obstacles are in general position with respect to the robot, meaning, no obstacle edge is parallel to a robot face, no obstacle face is parallel to a robot edge, etc. It can be shown that such degeneracies can only decrease the complexity of $\mathcal{F}$. (See [1] and references cited there for more details on the general position assumption in the motion planning setting.)




Figure 3: A placement of B (here a cube) making three contacts with obstacles, in this case a vertex contact, an edge contact, and a face contact. This placement of B corresponds to a vertex of the free space.

The robot is in *contact* with an obstacle C in a certain placement if the robot intersects the boundary of C but not its interior in that placement. A *contact specification* is a pair $O = (f, g)$ where f is a feature (vertex, edge, or face) of the robot, and g is a feature of an obstacle, such that $\dim f + \dim g \leq 2$. A (not necessarily free) placement of the robot is said to *realize* the contact specification O if at that placement f intersects g .

We call the contact specification $O = (f, g)$ *generic*, *singly degenerate*, or *doubly degenerate*, according to whether if $\dim f + \dim g$ equals 2, 1, or 0, respectively. If O is generic, then we call it a *vertex contact*, an *edge contact*, or a *face contact*, according to whether the robot feature f is a vertex, an edge, or a face, respectively.

Due to the general position assumption, each face of $\mathcal{F}$ arises from a connected set of robot placements that realize a fixed generic contact specification.

Similarly, each edge of $\mathcal{F}$ corresponds to a connected set of robot placements that simultaneously realize two generic contact specifications. Some of these edges realize a singly-degenerate contact specification.

Finally, each vertex of $\mathcal{F}$ arises from a robot placement simultaneously realizing three generic contact specifications. See Figure 3 for an example. Some vertices of $\mathcal{F}$ correspond to the robot simultaneously realizing a singly-degenerate contact specification and a generic contact specification, while other vertices of $\mathcal{F}$ correspond to the robot realizing a doubly-degenerate contact specification.

In order to asymptotically bound the complexity of $\mathcal{F}$, it is enough to bound its number of vertices, since then the same asymptotic bound will also hold for the number of edges and faces. Furthermore, vertices of $\mathcal{F}$ that involve degenerate contact specifications are easily bounded by $O(n^2)$, so we only need to consider vertices of $\mathcal{F}$ that arise from three generic contact specifications.

3 Upper and lower envelopes

In this section we present some results, some of them new, on upper and lower envelopes, which we will need in subsequent sections.

Let S be a finite collection of graphs of possibly intersecting, possibly partial, functions $\mathbb{R} \rightarrow \mathbb{R}$. The *lower envelope* of S is the pointwise minimum of these function graphs, i.e. the parts of the function graphs that are visible from below at $y = -\infty$. Similarly, the *upper envelope* of S are the parts of the function graphs that are visible from above at $y = +\infty$.




Figure 4: The lower envelope of a set of segments, decomposed into concave chains.

The case that concerns us in this paper is the one in which S is a collection of n nonvertical line segments. In this case, the lower envelope of S consists of at most $2n - 1$ concave chains, where the endpoints of the concave chains are the segment endpoints that are visible from below (see Figure 4). Similarly, the upper envelope of S consists of at most $2n - 1$ convex chains.

Hart and Sharir [5] proved that, if S is a collection of nonvertical line segments, then the combinatorial complexity of each envelope of S is at most $O(n\alpha(n))$, where $\alpha(n)$ is the inverse-Ackermann function. Furthermore, this bound is worst-case tight (Wiernik and Sharir [8]).

Halperin and Yap (Lemma 2.3 in [4]) showed that if S_1 and S_2 are two collections of nonvertical line segments in the plane, with $S_1 \cup S_2$ in general position and $|S_1 \cup S_2| = n$, then each envelope (upper or lower) of S_1 intersects each envelope (upper or lower) of S_2 in at most $O(n\alpha(n))$ points. We refine this result:

Lemma 3. *Let S_1 and S_2 be two collections of nonvertical line segments in the plane, with $S_1 \cup S_2$ in general position and $|S_1 \cup S_2| = n$. Then:*

1. *The lower envelopes of S_1 and S_2 (and, similarly, the upper envelopes of S_1 and S_2) intersect one another in at most $O(n\alpha(n))$ points.*
2. *The lower envelope of S_1 and the upper envelope of S_2 (and vice versa) intersect one another in at most $O(n)$ points.*
3. *If no two segments in S_1 intersect, then each envelope of S_1 intersects each envelope of S_2 in at most $O(n)$ points.*
4. *If every segment in S_1 has larger slope than every segment in S_2 , then each envelope of S_1 intersects each envelope of S_2 in at most $O(n)$ points.*

Furthermore, all these bounds are worst-case tight.

Proof. Claim (1) follows from the Hart–Sharir result applied on $S_1 \cup S_2$. Worst-case tightness follows by taking a configuration S of n segments that realizes $\Theta(n\alpha(n))$ lower-envelope intersections, and considering a random partition of S into two parts. Each lower-envelope intersection has a probability of $1/2$ of its two segments falling into different parts. By linearity of expectation, the expected number of lower-envelope intersections between segments of different parts is $\Theta(n\alpha(n)/2) = \Theta(n\alpha(n))$, and there must be a partition that realizes at least this value.

For Claims (2) and (3), partition the plane into $2n + 1$ vertical slabs by tracing vertical lines through the endpoints of the segments. Within each slab, the lower envelope of each S_i is part of a single concave chain, and the upper envelope of each S_i is part of a single convex chain. Then, Claim (2) follows since a concave chain and convex chain can intersect in at most two points. And




Figure 5: A “flat” triangular robot B in $\mathbb{R}^3$ making three vertex contacts.

Claim (3) follows since in this case, each chain of S_1 consists of a single segment, so it intersects a chain of S_2 in at most two points.

For the cases of Claim (4) not yet proven, consider the lower envelope of $S_1 \cup S_2$. Due to the constraints on the slopes, each concave chain in the envelope contains at most one intersection between a segment of S_1 and a segment of S_2 . The case of the upper envelope is symmetric. $\square$

4 Triple vertex contacts with a triangular robot

Following Halperin and Yap [4], we first consider the following auxiliary problem: Suppose the robot B is a “flat” triangle that can translate amidst polyhedral obstacles in $\mathbb{R}^3$. We would like to bound the number of vertices of the free space $\mathcal{F}$ that correspond to triple vertex contacts, i.e. the number of free placements in which the three vertices of B make contact with three different obstacle faces (see Figure 5). Halperin and Yap derived an upper bound of $O(n^2\alpha(n))$ for the number of such contacts. We improve this bound to $O(n^2)$ by making a minor modification to their argument.

Lemma 4. *Let B be a triangular robot that is free to translate in $\mathbb{R}^3$ amidst polyhedral obstacles that have a total of n vertices. Then the number of triple vertex contacts that B can make with obstacles is $O(n^2)$.*

(Recall that if we do not restrict our attention to vertex contacts then $\Omega(n^2\alpha(n))$ triple contacts are possible, as shown in Figure 1.)

We now proceed to prove Lemma 4. Suppose without loss of generality that the triangle B is “horizontal”, i.e. parallel to the xy -plane. We can assume all obstacle faces are triangles. By the general-position assumption, the vertices of each triangular obstacle face have different z -coordinates. Let us split each face into two *subtriangles* by a horizontal line passing through the vertex with middle z -coordinate. Thus, each subtriangle has two non-horizontal sides and one horizontal side.

Let us see what happens to the intersection $f(z_0)$ between a subtriangle f and a horizontal plane $z = z_0$, as we vary z_0 . For every z_0 in the z -range of f , $f(z_0)$ is a line segment. As we increase z_0 at a constant rate, the line segment moves parallel to itself at constant speed, and its two endpoints also move at constant speed.

4.1 Contact specifications and covering

In our case, the relevant contact specifications are those of the form $O = (p, f)$ where p is a vertex of B and f is a subtriangle.




Figure 6: At this value $z = z'$, the contact specification $O_2 = (p_2, f_2)$ covers the contact specification $O_1 = (p_1, f_1)$ at the right, since q is to the right of $f_1(z')$ and the ray r from b with direction $\overrightarrow{p_1 p_2}$ intersects $f_2(z')$. As we slide the robot B along $f_1(z')$, it goes from not colliding with $f_2(z')$ to colliding with it.

Consider two contact specifications $O_1 = (p_1, f_1)$, $O_2 = (p_2, f_2)$, where $p_1 \neq p_2$, and where the z -ranges of f_1, f_2 overlap. Let $[z_0, z_1]$ be the intersection of the z -ranges of f_1, f_2 . Given $z' \in [z_0, z_1]$, we would like to define whether O_2 covers O_1 at $z = z'$. Let $\ell_1 = \ell_1(z')$ and $\ell_2 = \ell_2(z')$ be the lines containing the segments $f_1(z')$ and $f_2(z')$. Let q be the point of intersection of ℓ_1 and ℓ_2 , and suppose that $f_1(z')$ does not contain q . Let a, b be the endpoints of $f_1(z')$, with a left of b (i.e. with smaller x -coordinate). Let r be the ray with direction $\overrightarrow{p_1 p_2}$ emerging from the endpoint among a, b that is closer to q . If r intersects $f_2(z')$ then we say that O_2 covers O_1 . If q is to the right of $f_1(z')$ then we say that O_2 covers O_1 at the right; otherwise we say that O_2 covers O_1 at the left. See Figure 6.

The significance of O_2 covering O_1 is as follows. Suppose we slide the robot with its vertex p_1 moving along the segment $f_1(z')$ from one endpoint of $f_1(z')$ to the other. It might happen that during this motion, the edge $p_1 p_2$ starts colliding with the segment $f_2(z')$ and then stops colliding with it. See Figure 7. However, if O_2 covers O_1 then this is not possible: If the edge $p_1 p_2$ starts colliding with the segment $f_2(z')$, then the collision will continue until $p_1 p_2$ reaches the other endpoint of $f_1(z')$.

Since the endpoints of the segments $f_1(z'), f_2(z')$ move linearly with respect to z' , the range of values of z' at which O_2 covers O_1 forms a continuous subinterval of $[z_0, z_1]$.

Consider a placement of the robot realizing three vertex contact specifications $O_1 = (p_1, f_1)$, $O_2 = (p_2, f_2)$, $O_3 = (p_3, f_3)$. For each pair of indices $i, j \in \{1, 2, 3\}$, $i \neq j$, either O_i covers O_j or O_j covers O_i . Hence, one of the two following possibilities occurs: Either there is one contact specification that is covered by the other two, or the contact specifications are *circular* in the sense that O_i covers O_j , O_j covers O_k , and O_k covers O_i . See Figure 8.

4.2 Parametric plane of a contact specification

Given a contact specification $O_1 = (p_1, f_1)$, we define the *contact parametric plane* P_{O_1} in which we parametrize by a pair of real numbers (z', s) every robot placement in which the robot vertex p_1 makes contact with $f_1(z')$ (or with its supporting line). The number s represents the distance




Figure 7: Here $O_2 = (p_2, f_2)$ does not cover $O_1 = (p_1, f_1)$, and as we slide the robot B along $f_1(z')$, it collides with $f_2(z')$ but then stops colliding with it.




Figure 8: Robot placements realizing three vertex contact specifications O_1, O_2, O_3 . In the left and center figures, O_1 is covered by O_2 and O_3 . In the right figure, O_1 is covered by O_2 , which is covered by O_3 , which is covered by O_1 .




Figure 9: Three robot placements realizing a contact specification O_1 with the same z -value (left); and their representation in the parametric plane P_{O_1} (right). The shaded area within Δ_{O_1} represents robot placements that are illegal due to collision with f_2 .

between the left endpoint of $f_1(z')$ and the position of p_1 along $f_1(z')$. Thus, $P_{O_1} \cong \mathbb{R}^2$, and the robot placements that realize the contact O_1 correspond to a triangular region $\Delta_{O_1} \subset P_{O_1}$. We imagine that the z -axis in P_{O_1} is horizontal and the s -axis is vertical.

We can mark each point of Δ_{O_1} as *legal* or *illegal* depending on whether the corresponding robot placement intersects the interior of another obstacle or not.

Given another contact specification $O_2 = (p_2, f_2)$ with $p_2 \neq p_1$, the set of robot placements that realize both O_1 and O_2 corresponds either to the empty set or to a line segment within Δ_{O_1} . However, we plot in P_{O_1} only contact specifications that cover O_1 , and only for the z -ranges at which they cover it. Given O_2 that covers O_1 at a certain z -range, let σ_{O_2} be the segment in Δ_{O_1} that corresponds to all robot placements that realize the double contact O_1, O_2 in this z -range. If O_2 covers O_1 at the left, then all the points vertically below σ_{O_2} in the triangle Δ_{O_1} are illegal, and if O_2 covers O_1 at the right, then all the points vertically above σ_{O_2} in Δ_{O_1} are illegal. See Figure 9.

4.3 Families of non-intersecting segments in the parametric planes

Given a contact specification $O_1 = (p_1, f_1)$ and another robot vertex p_2 , let L_{O_1, p_2} (resp. R_{O_1, p_2}) be the set of all subtriangles f such that $O_2 = (p_2, f)$ covers O_1 at the left (resp. right) for some range of z . For each $f \in L_{O_1, p_2}$ (resp. R_{O_1, p_2}) let $\sigma(f)$ be the corresponding line segment in the parametric plane P_{O_1} . Let $\mathcal{L}_{O_1, p_2} = \{\sigma(f) : f \in L_{O_1, p_2}\}$ and $\mathcal{R}_{O_1, p_2} = \{\sigma(f) : f \in R_{O_1, p_2}\}$ be the sets of these line segments.

Here is the crucial observation that is missing in [4]: The segments in $\mathcal{L}_{O_1, p_2}$ are pairwise non-intersecting, since if $\sigma(f), \sigma(f') \in \mathcal{L}_{O_1, p_2}$ intersected, then the subtriangles f, f' themselves would intersect. Similarly, the segments in $\mathcal{R}_{O_1, p_2}$ are pairwise non-intersecting.

4.4 Counting triple vertex contacts

We are now ready to bound the number of robot placements making three simultaneous vertex contacts.

Consider first triple contacts in which one contact specification $O_1 = (p_1, f_1)$ is covered by the other two, $O_2 = (p_2, f_2)$, $O_3 = (p_3, f_3)$. In the parametric plane P_{O_1} , such a triple contact corresponds to the intersection point q of a segment in either $\mathcal{L}_{O_1, p_2}$ or $\mathcal{R}_{O_1, p_2}$ with a segment in either $\mathcal{L}_{O_1, p_3}$ or $\mathcal{R}_{O_1, p_3}$. Furthermore, in order for q to be legal, it is necessary for q to lie in the upper (resp. lower) envelope of $\mathcal{L}_{O_1, p_2}$ (resp. $\mathcal{R}_{O_1, p_2}$), as well as in the upper (resp. lower) envelope of $\mathcal{L}_{O_1, p_3}$ (resp. $\mathcal{R}_{O_1, p_3}$). Hence, Lemma 3 (Case 3) bounds the number of these intersection points q , for fixed O_1 , by $O(n)$. Since there are $O(n)$ choices of O_1 , the total bound for this type of triple-intersection contacts is $O(n^2)$.

Now consider circular triple contacts. Consider a robot placement at $z = z'$ realizing the triple contact $O_1 = (p_1, f_1)$, $O_2 = (p_2, f_2)$, $O_3 = (p_3, f_3)$ in which O_1 is covered by O_2 , O_2 is covered by O_3 , and O_3 is covered by O_1 . Say that all the coverings are at the right (the other cases are similar). This robot placement corresponds to a point q_1 in the segment $\sigma(f_2) \in \mathcal{R}_{O_1, p_2}$, which is the lowest segment of $\mathcal{R}_{O_1, p_2}$ at $z = z'$. The robot placement also corresponds to a point q_2 in the segment $\sigma(f_3) \in \mathcal{R}_{O_2, p_3}$, which is the lowest segment of $\mathcal{R}_{O_2, p_3}$. Finally, the robot placement also corresponds to a point q_3 in the segment $\sigma(f_1) \in \mathcal{R}_{O_3, p_1}$, which is the lowest segment of $\mathcal{R}_{O_3, p_1}$. All three points q_1, q_2, q_3 have z -coordinate equal to z' .

Consider the lower envelope of $\mathcal{R}_{O_1, p_2}$. This lower envelope is piecewise linear, with breakpoints at discrete values of z . Since the segments in $\mathcal{R}_{O_1, p_2}$ are pairwise nonintersecting, each breakpoint is an endpoint of a segment of $\mathcal{R}_{O_1, p_2}$. Let z_1 be the largest z -coordinate of a breakpoint that is smaller than z' . Define z_2, z_3 similarly, by looking at $\mathcal{R}_{O_2, p_3}, \mathcal{R}_{O_3, p_1}$, respectively. Let $z'' = \max\{z_1, z_2, z_3\}$.

Hence, in the z -range $[z'', z']$ the lowest segment of each $\mathcal{R}_{O_1, p_2}, \mathcal{R}_{O_2, p_3}, \mathcal{R}_{O_3, p_1}$ is the aforementioned $\sigma(f_2), \sigma(f_3), \sigma(f_1)$, respectively. We charge the triple contact O_1, O_2, O_3 to the breakpoint z'' .

In this charging scheme, each breakpoint in a lower envelope of a parametric plane is charged at most once, because we can reconstruct the triple contact given the breakpoint: Given z'' which is a breakpoint in $\mathcal{R}_{(p_1, f_1), p_2}$, say, we take the segment $\sigma(f_2)$ that is lowest in $\mathcal{R}_{(p_1, f_1), p_2}$ just after $z = z''$. Once we know f_2 , we know we need to look at $\mathcal{R}_{(p_2, f_2), p_3}$. Then we take the segment $\sigma(f_3)$ that is lowest in $\mathcal{R}_{(p_2, f_2), p_3}$ just after $z = z''$. Once we know f_3 , we know we need to look at $\mathcal{R}_{(p_3, f_3), p_1}$. Then we take the segment $\sigma(f_4)$ that is lowest in $\mathcal{R}_{(p_3, f_3), p_1}$ just after $z = z''$. If $f_4 = f_1$ then we have successfully reconstructed the triple contact-specification $(p_1, f_1), (p_2, f_2), (p_3, f_3)$ (which might or might not be realizable as a triple contact). If $f_4 \neq f_1$ then the breakpoint z'' we started with does not receive any charge.

Since each breakpoint is a segment endpoint and there are $O(n^2)$ segments in all the families $\mathcal{R}_{O, p}$ altogether, there are at most $O(n^2)$ circular triple contacts in which all coverings take place at the right. The other possibilities, in which some of the coverings take place at the left, are treated similarly.

This concludes the proof of Lemma 4.

Remark. Halperin and Yap [4] defined each $\mathcal{L}_{(p_1, f_1), p_2} \cup \mathcal{L}_{(p_1, f_1), p_3}$ as a single set, instead of as two separate sets. Therefore, the segments in their sets can intersect, and that is why they got the $\alpha(n)$ factor in their bound. Our approach of defining $\mathcal{L}_{(p_1, f_1), p_2}$ and $\mathcal{L}_{(p_1, f_1), p_3}$ as two separate sets (and similarly defining $\mathcal{R}_{(p_1, f_1), p_2}$ and $\mathcal{R}_{(p_1, f_1), p_3}$ as two separate sets) allows us to apply Lemma 3 (Case 3) and get a bound free of the $\alpha(n)$ factor.

Following Halperin and Yap, we conclude the following:

Corollary 5. *Let B be a polyhedral robot of constant complexity that is free to translate in $\mathbb{R}^3$ amidst polyhedral obstacles that have a total of n vertices. Then the number of triple vertex contacts that B can make with obstacles is $O(n^2)$.*

Proof. Every triple contact made by three vertices of B is also a triple contact made by a triangular robot B' spanned by those three vertices (though the opposite is not necessarily the case). $\square$

5 The case of a cubical robot

In this section we prove Theorem 1, regarding the case in which B is a cube (or a box or a parallelepiped). Assume for concreteness that $B = [0, 1]^3$ is an axis-parallel cube of side 1.

Let us go over all types of triple contacts $O_1 = (f_1, g_1)$, $O_2 = (f_2, g_2)$, $O_3 = (f_3, g_3)$.

Lemma 4 already addressed the case of three vertex contacts. Hence, we can assume without loss of generality that O_1 is not a vertex contact.

Due to the general position assumption, it is impossible to have two face contacts involving the same face or opposite faces of B .

Lemma 6. *Suppose f_1 and f_2 are the same edge or parallel edges of B , or that f_1 and f_2 are nonparallel faces of B , or that f_1 is a face of B and f_2 is an edge of f_1 or of the face opposite f_1 . Then there are at most $O(n^2)$ triple contacts involving f_1 and f_2 .*

Proof. In all these cases, once we fix obstacle features g_1, g_2 for the first two contacts, the line of movement of B that preserves these two contacts is axis-parallel. Suppose we move B in one direction along this line until it makes a third contact with some obstacle feature g_3 . Further movement in the same direction will make B intersect g_3 until B moves distance at least 1. But after B moves such a distance, it will no longer realize O_1 nor O_2 . Hence, for fixed g_1, g_2 , there are at most two possibilities for g_3 , one for each direction of movement. Since there are $O(n^2)$ choices for g_1, g_2 , the claim follows. (This argument is given in Lemma 2.1 of [4].) $\square$

In order to handle the remaining cases, we examine parametric planes of face and edge contacts.

5.1 Parametric plane of a face contact

Let $O_1 = (f_1, g_1)$ be a face contact. The set of all placements of B that realize O_1 can be parametrized by a pair of real numbers between 0 and 1. For example, if the face f_1 is parallel to the xy -plane, then we can represent every placement of B that realizes (f_1, g_1) by two real numbers $s, t \in [0, 1]$, where s denotes the distance from the vertex g_1 to the higher- x , y -parallel edge of f_1 , and t denotes the distance from g to the higher- y , x -parallel edge of f_1 . Real values of s and/or t outside of the range $[0, 1]$ give rise to placements in which g_1 lies in the containing plane of f_1 , though not in f_1 itself.

In the parametric plane $P_{O_1} \cong \mathbb{R}^2$, the unit square $\Gamma = [0, 1]^2 \subset P_{O_1}$ corresponds to the placements in which the contact O_1 actually takes place. Each point of P_{O_1} can be marked *legal* or *illegal* according to whether the corresponding robot placement intersects the interior of some obstacle or not.

As before, given another contact specification $O_2 = (f_2, g_2)$, the set of placements that realize both O_1 and O_2 corresponds either to a line segment $\sigma_{O_1, O_2} \subset P_{O_1}$ or to the empty set.

Whether O_2 is an edge contact with f_2 perpendicular to f_1 , or a vertex contact, the region of Γ that lies behind σ_{O_1, O_2} in two perpendicular axis-parallel directions is illegal. See Figure 10. Furthermore, if $O_2 = (f_2, g_2)$, $O'_2 = (f_2, g'_2)$ are two vertex contacts involving the same vertex f_2 of B , then the segments σ_{O_1, O_2} , σ_{O_1, O'_2} cannot intersect, because then the obstacle faces g_2, g'_2 would intersect.




Figure 10: Parametric plane of a face contact for a cube robot. The contact $O_1 = (f_1, g_1)$ defining the parametric plane is shown on the left, and two different options for the second contact $O_2 = (f_2, g_2)$, and the corresponding segments σ_{O_1, O_2} in the parametric plane, are shown on the right.




Figure 11: Parametric plane of an edge contact for a cube robot. The dotted lines in the parametric planes represent robot movements that preserve O_1 and are "horizontal" (parallel to the xy -plane).




Figure 12: If f_2 is one of the bolded edges, then a robot movement that preserves O_1 and O_2 will descend (decrease in z -coordinate) as it increases in x -coordinate. If f_2 is one of the dotted edges, then a robot movement that preserves O_1 and O_2 will ascend as it increases in x -coordinate.

5.2 Parametric plane of an edge contact

Now let $O_1 = (f_1, g_1)$ be an edge contact. Assume for concreteness that f_1 is the “vertical” edge of B that satisfies $x = y = 0$. In order for the contact O_1 to be possible, the obstacle segment g_1 must decrease in y -coordinate as it increases in x -coordinate.

We can represent each placement of B in which f_1 (or its extension) intersects g_1 (or its extension) by a pair (s, t) , where s is the distance between the lower- x endpoint of g_1 and the point of contact p , and t is the distance between p and the higher- z endpoint of f_1 .

In the resulting parametric plane $P_{O_1} \cong \mathbb{R}^2$, the placements in which the contact O_1 actually takes place correspond to a rectangle $\Gamma = [0, s_{\max}] \times [0, 1] \subset P_{O_1}$, where $s_{\max}$ is the length of g_1 . A movement of B that preserves the contact O_1 and is “horizontal” (preserves z -coordinate) corresponds in P_{O_1} to a line of a certain slope $m = m_{g_1}$.

Each point of P_{O_1} can be marked legal or illegal according to whether the corresponding placement of B is legal or illegal. Consider a second contact specification $O_2 = (f_2, g_2)$, where f_2 is either a vertex of B or an edge of B not parallel to f_1 . Then the set of points corresponding to placements that realize O_2 is either empty or a line segment $\sigma_{O_1, O_2} \subset P_{O_1}$. Furthermore, the region of Γ that is either vertically above or vertically below the segment σ_{O_1, O_2} (depending on whether f_2 lies on the lower- z or the higher- z side of B) is illegal. See Figure 11.

If $O_2 = (f_2, g_2)$, $O'_2 = (f_2, g'_2)$ are vertex contacts involving the same vertex f_2 of B , then the segments σ_{O_1, O_2} , σ_{O_1, O'_2} cannot intersect, because then the obstacle faces g_2 , g'_2 would intersect.

Out of the eight edges f_2 of B that are not parallel to f_1 , four of them yield segments σ_{O_1, O_2} that have slope larger than m for all choices of obstacle edge g_2 . The other four edges f_2 of B yield segments σ_{O_1, O_2} that have slope smaller than m for all choices of obstacle edge g_2 . See Figure 12.

5.3 The case analysis

Let us now go over the possibilities not handled by Lemmas 4 and 6.

Suppose at least one contact is a face contact. Say f_1 is a face of B . Then f_2, f_3 cannot both be edges of B , because then either they would be parallel, or else one of them would belong to f_1 or to the face opposite f_1 . Thus, one of f_2, f_3 , say f_2 , must be a vertex of B , while f_3 must be either another vertex of B or an edge perpendicular to f_1 . For each such choice of f_1, f_2, f_3 , there are $O(n)$ choices of g_1 . Given a choice of g_1 , consider the parametric plane P_{O_1} . The segments σ_{O_1, O_2} , over all choices of g_2 , form a family of segments that do not intersect one another. Hence,




Figure 13: When the robot is a fully-parallel polygon, in the parametric plane of an edge contact O_1 , the parallelogram-shaped region of height d above or below the segment σ_{O_1, O_2} is certainly illegal. Hence, we can divide the rectangle into horizontal strips and proceed as before.

we proceed as in Section 4.4, invoking Lemma 3 (Case 3). There are $O(n)$ triple contacts for each choice of g_1 , and hence a total of $O(n^2)$ contacts involving our choice of f_1, f_2, f_3 .

Now suppose no contact is a face contact. At least one contact, say O_1 , must be an edge contact. Hence, let f_1 be an edge of B . Each of f_2, f_3 must be an edge or a vertex of B , and no two of f_1, f_2, f_3 can be parallel edges. If one of f_2, f_3 is a vertex of B , then Lemma 3 (Case 3) applies again, and we again get a bound of $O(n^2)$.

Now suppose f_1, f_2, f_3 are three mutually nonparallel edges of B . For each f_i , let $h_{i,1}, h_{i,2}$ be the two faces of B that are perpendicular to f_i . If at least one f_i satisfies the property that the other two edges f_j, f_k belong to different faces $h_{i,1}, h_{i,2}$, then in the parametric plane P_{O_i} we are looking at intersections between an upper envelope and a lower envelope, so Lemma 3 (Case 2) applies, and we get once again a bound of $O(n^2)$.

The only remaining case is where f_1, f_2, f_3 are three edges incident to the same vertex of B . Fix an obstacle edge g_1 . In the parametric plane P_{O_1} , every segment σ_{O_1, O_2} has slope smaller than some slope $m = m_{g_1}$, while every segment σ_{O_1, O_3} has slope larger than m . Hence, Lemma 3 (Case 4) applies, and we again get a bound of $O(n)$ triple contacts for our choice of g_1 , or $O(n^2)$ contacts in total.

This case analysis concludes the proof of Theorem 1.

6 The case of a fully-parallel polygonal robot

Theorem 2, regarding a fully-parallel polygonal robot, follows by a few small modifications to the proof of Theorem 1.

Consider the case where two edge contacts $O_1 = (f_1, g_1), O_2 = (f_2, g_2)$ involve robot edges f_1, f_2 that are either equal or parallel. Let d and D , with $d < D$, be lengths of f_1 and f_2 . Then we argue as in the proof of Lemma 6, except that now there are at most $1 + \lceil D/d \rceil$ possibilities for g_3 , given g_1, g_2 . Since this number is a constant, we again get an overall bound of $O(n^2)$ for this case.

Next, consider an edge contact $O_1 = (f_1, g_1)$. Let f'_1 be the edge of B that is parallel to f_1 , and let d and D be the smaller and the larger length among f_1, f'_1 , respectively. The edges f_1, f'_1

partition the rest of the boundary of B into two parts, which we call $\pi_1(f_1)$ and $\pi_2(f_1)$. Consider a second contact $O_2 = (f_2, g_2)$ where f_2 is either a vertex of B or an edge of B different from f_1 and f'_1 . Consider the segment $\sigma_{O_1, O_2} \in S_{O_1, f_2}$ in the parametric plane P_{O_1} . Then there is a parallelogram-shaped region of height d , either vertically above or vertically below σ_{O_1, O_2} (depending on whether f_2 belongs to $\pi_1(f_1)$ or to $\pi_2(f_1)$), that is certainly illegal. See Figure 13. Hence, we can partition the parametric rectangle Γ into $\lceil D/d \rceil$ horizontal strips of width at most d . This operation might cut some segments in S_{O_1, f_2} into subsegments. Within each strip we are interested only in the lower or the upper envelope of the subsegments in that strip.

If one of f_2, f_3 is a vertex of B , then Lemma 3 (Case 3) applies.

Finally, suppose f_1, f_2, f_3 are three mutually nonparallel edges of B . Then at least one of the edges f_i satisfies the property that the other two edges f_j, f_k belong to opposite boundary parts $\pi_1(f_i), \pi_2(f_i)$. In the corresponding parametric plane P_{O_i} , the segments σ_{O_i, O_j} will have illegal regions above the segments, while the segments σ_{O_i, O_k} will have illegal regions below the segments (or vice versa). Hence, Lemma 3 (Case 2) applies.

With these modifications, Theorem 2 is proven.

7 Discussion and open problems

The main open problem is to close the gap between $O(n^2 \log n)$ and $\Omega(n^2 \alpha(n))$ for the general case of a convex polyhedral robot.

Another open problem is to obtain more refined bounds that depend on the complexity of the robot, as well as on the number of obstacles, and not just on the total number of obstacle vertices, under the assumption that the obstacles are convex.

Suppose there are k convex obstacles with a total of n vertices. For the general case where the robot is an r -vertex polyhedron, Aronov and Sharir [1] proved an upper bound of $O(rnk \log k)$. It is easy to achieve a lower bound of $\Omega(nk)$ for any robot. For a flat triangular robot, the construction of Figure 1 can be modified to achieve a lower bound of $\Omega(nk\alpha(k))$ (see the details in [1]). Halperin (personal communication) derived an upper bound of $O(nk)$ for the case of a segment-shaped robot. These are the currently best known bounds, as far as we know.

Acknowledgements. Thanks to Danny Halperin for suggesting me to look into this problem and for useful conversations.

References

- [1] Boris Aronov and Micha Sharir. On translational motion planning of a convex polyhedron in 3-space. *SIAM Journal on Computing*, 26(6):1785–1803, 1997. doi:10.1137/S0097539794266602.
- [2] Dan Halperin, Oren Salzman, and Micha Sharir. Algorithmic motion planning. In Csaba D. Toth, Joseph O’Rourke, and Jacob E. Goodman, editors, *Handbook of discrete and computational geometry*, pages 1311–1342. Chapman and Hall/CRC, 2017.
- [3] Dan Halperin and Chee-Keng Yap. Combinatorial complexity of translating a box in polyhedral 3-space. In *Proceedings of the ninth annual symposium on Computational geometry (SoCG 1993)*, pages 29–37. ACM, 1993. doi:10.1145/160985.160992.

- [4] Dan Halperin and Chee-Keng Yap. Combinatorial complexity of translating a box in polyhedral 3-space. *Computational Geometry*, 9(3):181–196, 1998. doi:10.1016/S0925-7721(97)00030-8.
- [5] Sergiu Hart and Micha Sharir. Nonlinearity of Davenport–Schinzel sequences and of generalized path compression schemes. *Combinatorica*, 6(2):151–177, 1986. doi:10.1007/BF02579170.
- [6] Klara Kedem, Ron Livne, János Pach, and Micha Sharir. On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles. *Discrete & Computational Geometry*, 1(1):59–71, 1986. doi:10.1007/BF02187683.
- [7] Tomás Lozano-Pérez and Michael A Wesley. An algorithm for planning collision-free paths among polyhedral obstacles. *Communications of the ACM*, 22(10):560–570, 1979. doi:10.1145/359156.359164.
- [8] Ady Wiernik and Micha Sharir. Planar realizations of nonlinear Davenport-Schinzel sequences by segments. *Discrete & Computational Geometry*, 3(1):15–47, 1988. doi:10.1007/BF02187894.