



Throwing a Sofa Through the Window

Thesis submitted in partial fulfillment of the requirements for the M.Sc.
degree in the
Blavatnik School of Computer Science, Tel-Aviv University

by

Itay Yehuda

This work has been carried out under the supervision of
Prof. Dan Halperin and Prof. Micha Sharir

August 2020

Acknowledgments

I deeply thank Pankaj Agarwal and Boris Aronov for useful interactions concerning this work. In particular, Boris has suggested the proof of the key topological property in the analysis in Chapter 4, and Pankaj has been instrumental in discussions concerning the range searching problem in Chapter 2. I am also grateful to Lior Hadassi for many chats and specifically the interaction involving the topological property presented in Chapter 4, and to Eytan Tirosh for the help with the alternative proof of Lemma 2.1.

Finally, I want to thank very deeply Micha Sharir and Dan Halperin, for the huge help with the whole process; from interacting with each and every problem, to teaching me how to type my solutions legibly. Without their assistance this thesis would not be as good as it is.

Abstract

We study several variants of the problem of moving a convex polytope in three dimensions through a flat rectangular (and sometimes more general) window. Specifically, we study variants in which the motion is restricted to translations only, discuss situations in which such a motion can be reduced to sliding (translation in a fixed direction) and present efficient algorithms for those variants. We then discuss the case of a window that is unbounded (has two infinite edges) and show that in this case, rotations are not necessary for passing the polytope through the window, an observation that leads to an efficient algorithm for this variant too. Then we study the importance of rotations by an example of a polytope that cannot pass through a certain window by translations only, but it can do so when rotations are allowed. We study also more general convex windows, and obtain some special properties of polytopes that can pass such a convex window. We then study the case of a circular window, and show that, for the regular tetrahedron K , there are two thresholds $1 > \delta_1 > \delta_2$ such that (i) K can slide through the window W if its diameter d is ≥ 1 , (ii) K cannot slide through W but can pass through it by a purely translational motion when $\delta_1 \leq d < 1$, (iii) K cannot pass through W by a purely translational motion but can do it when rotations are allowed when $\delta_2 \leq d < \delta_1$, and (iv) K cannot pass through W at all when $d < \delta_2$. This divides this motion planning problem into three sub-classes, with different capabilities: one dimensional translation (“sliding”), purely translational motion, and unrestricted motion (with all six degrees of freedom).

Contents

1	Introduction	1
2	Translation in a Fixed Direction	5
2.1	The Existence of an Orthogonal Sliding Motion	5
2.2	Finding a Sliding Motion	9
2.3	Implicitly Constructing All the Silhouettes	14
2.4	Improved Algorithm	16
3	Unbounded Windows	21
4	From Passing Through an Arbitrary Convex Window to Sliding Through a Gate	26
5	Purely Translational Motions	33
6	Rotations Are Needed	35
7	The Case of a Circular Window	38
7.1	Purely Translational Motion	38
7.2	General Motion	43
8	Conclusion	46

1

Introduction

Let K be a convex polytope (a ‘sofa’) in \mathbb{R}^3 with n edges, and let W be a rectangular window, placed, say, in the xy -plane in the axis-parallel position $[0, a] \times [0, b]$, where a and b are the respective width and height of W . We assume that the complement of W in the xy -plane is a solid wall that K must avoid. The problem is to determine whether K can be moved, in a collision-free manner, from any position that is fully contained in the upper halfspace $z > 0$, through W , to any position that is fully contained in the lower halfspace $z < 0$, and, if so, to plan such a motion (see Figure 1.1).

A continuous motion of a rigid body in three dimensions has six degrees of freedom, three of translation and three of rotation, and in the general form of the problem, we allow all six degrees. We will mainly study simpler versions where only restricted types of motion are allowed, such as purely translational motion (that has only three degrees of freedom), a translational motion in a fixed direction, that we refer to as *sliding* (one degree of freedom), or a translational motion combined with rotations around the vertical axis only (four degrees of freedom), etc. Some of our main results show that, in certain favorable situations, the existence of a general collision-free motion of K through W implies the existence of a restricted motion. This allows us to solve the problem in a significantly more efficient manner.

In terms of the *free configuration space* \mathcal{F} of K , all the placements of K that are fully contained in the upper halfspace are free, and form a connected subset \mathcal{F}^+ of \mathcal{F} . Similarly, all the placements of K that are fully contained in the lower halfspace are free, and form a connected subset \mathcal{F}^- of \mathcal{F} . Our problem, in general, is to determine whether both \mathcal{F}^+ and \mathcal{F}^- are contained in the same connected component of \mathcal{F} . This interpretation applies to the general setup, with six degrees of freedom, as well as to any other subclass of motion, with fewer degrees of freedom.

Motion planning is an intensively studied problem in computational geometry and robotics.

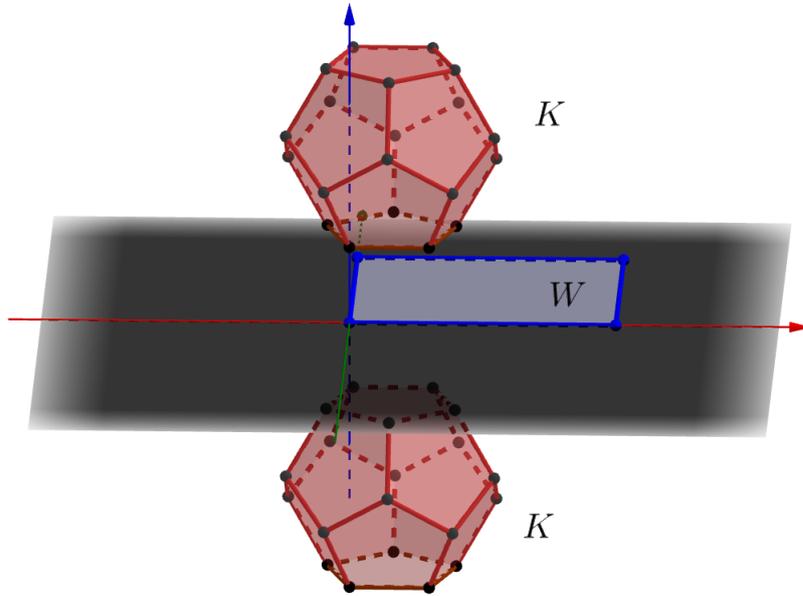


Figure 1.1: The general problem studied in this thesis: moving a convex polytope from an initial configuration above the xy -plane, to a target configuration below the xy -plane through a window on the xy -plane.

There is a systematic and general way to describe the free space \mathcal{F} using *constraint surfaces*, namely surfaces describing all the configurations where one feature on the boundary of the moving object (K in our case) touches a feature on the boundary of the work space space (W in our case); see, e.g., [7, 15, 21]. These surfaces partition the configuration space into cells such that each cell is either fully contained in \mathcal{F} or fully contained in the forbidden portion of the configuration space. This representation is based on the arrangement [14] of constraint surfaces, and is induced in our setting by $O(n)$ surfaces, since the combinatorial complexity (number of vertices, edges, and facets) of K is $O(n)$ and of the window is $O(1)$. By standard arguments in the study of arrangements of surfaces, the complexity of \mathcal{F} is bounded by $O(n^d)$, where d is the number of degrees of freedom [15]. In order to exploit this representation, we need to be able to construct it and then transform it into a graph (in the graph-theory sense) on which we will search for a solution motion. To this end we need further machinery, and we typically use *vertical decomposition*, a refinement of arrangements of surfaces [5]. Such constructions are easily implementable for motion planning with two degrees of freedom [10], but become practically complex for problems with three or more degrees of freedom. This has led researchers in robotics to develop alternative methods as we describe next.

There is a large suite of practical solutions to the motion-planning problem based on sampling-based techniques [6, Chapter 7],[13], the best known of which are PRM [19] and RRT [20], which have dozens of variants. While extremely successful in solving practical problems, they trade-off the completeness of the arrangement approach with efficiency. In particular the sampling-based techniques fail miserably, when the setting is tight [23, 24],

which is exactly the situation in the problems that we study in this thesis. If the polytope K is small relative to the window, then the problems become trivial. Therefore, our study is of instances where K is roughly the size of the window (the size comparisons are made precise below), and in such cases sampling-based techniques are inapplicable.

Toussaint [27] collected a variety of tight-setting motion planning problems under the title *movable separability of sets*. These problems are interesting both from a pure research perspective (see, e.g., [26] for an intriguing problem and its solution), but also from an applied perspective, since motion in tight settings often arises in manufacturing processes such as *assembly planning* [12] or *casting and molding* [4]. It is in Toussaint’s review that we encountered the problem of throwing a polytope through a window. Although Toussaint’s paper was published 35 years ago, to the best of our knowledge there has not been progress on this specific problem up till this thesis. We remark that the word *sofa* in the title of the thesis is borrowed from a classical problem of this flavor, in a two-dimensional setting, the *moving sofa* problem, or just the sofa problem (see, e.g., [8, 11]), which is to find the shape of largest area that can be moved through a corner in an L-shaped corridor whose legs have width 1 (see Figure 1.2).

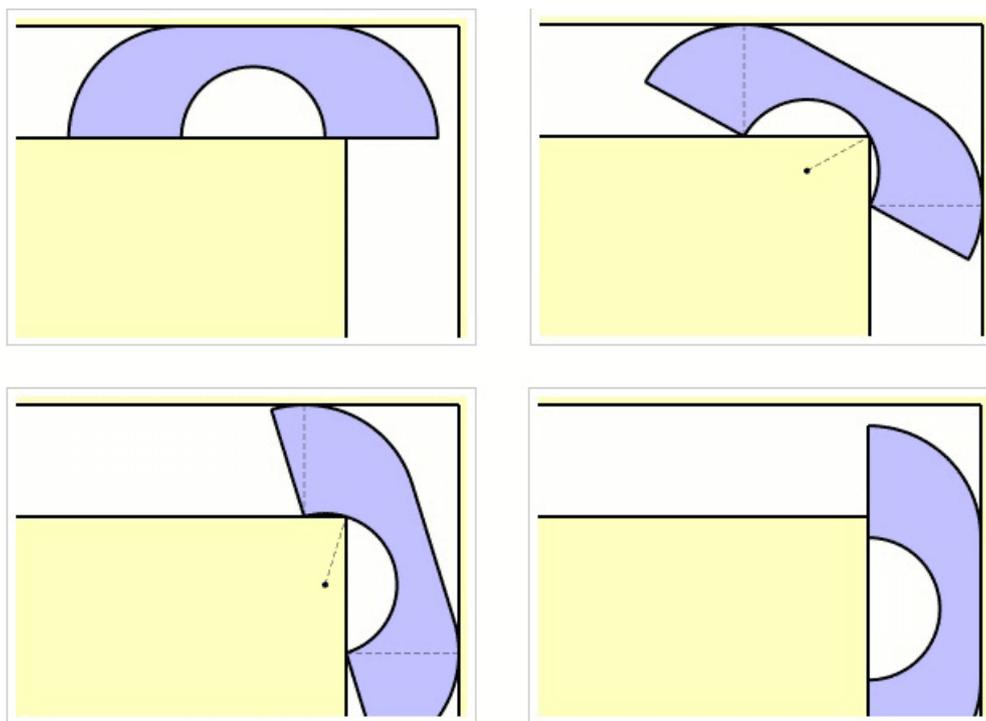


Figure 1.2: Moving a ‘sofa’ around the corner of an L-shaped corridor. Figure taken from <https://img.microsiervos.com/images2017/problema-del-sofa-Hammersley.png>

In general, \mathcal{F} is six-dimensional, so constructing an exact representation of \mathcal{F} , as is needed here, is both complicated and expensive. We do not address this setup in most of the thesis, but mainly focus on variants of the problem with fewer degrees of freedom, of the kinds mentioned above (General motions are considered in Chapters 6 and 7, and they are also

discussed in several other instances of the problem). In addition to allowing us to obtain considerably faster algorithms, these versions have nontrivial geometric properties, which are interesting in their own right, and which also allow us, in certain favorable situations, to reduce more general problems to simpler ones, with fewer degrees of freedom, thereby achieving faster algorithmic solutions.

Our results. We first consider, in Chapter 2, sliding motions (translations in a fixed direction) of K . We characterize the situations in which such a sliding motion exists, present efficient algorithms for finding such a motion when one exists, and report that there is no such a motion otherwise. We next consider in Chapter 3 unrestricted motion for the case where the window is unbounded in one direction (has two infinite edges)—we refer to such an unbounded window as a *gate*. We show that the existence of such a general collision-free motion through a gate implies the existence of a collision-free sliding motion through that gate, which is very easy to find when it exists. In Chapter 4 we present the connection between the existence of a general collision-free motion through any convex window W and sliding through a gate whose width depends on W . Specifically, we show that if the polytope K can pass through some convex window of diameter d , then, for any fixed direction \vec{v} , we can slide K through a gate of width d in the direction \vec{v} . We then consider in Chapter 5 purely translational motion of K through a rectangular window, and prove that the existence of such a purely-translational collision-free motion implies the existence of a collision-free sliding motion. In Chapter 6, we observe the importance of rotations. We give an example of a convex polytope that can move through a squared window by a collision-free motion that includes rotation, and show that in this case a purely translational motion does not exist. In Chapter 7 we consider the case where the window is circular. We show that for some polytopes, the various restricted families of motions are non-equivalent: For a sufficiently large window, there is a sliding motion. When the window is too small for sliding, it might be still large enough for a purely-translational motion. And if the window is too small for a purely-translational motion, it is possible that it is large enough for a general motion with six-degrees of freedom.

2

Translation in a Fixed Direction

In this chapter we address the case in which the movement is purely translational in a single fixed direction. Such a motion, to which we refer as a *sliding motion*, has only one degree of freedom. In the most restricted version (which is very easy to solve), we are given a fixed orientation of K and a fixed initial placement, and also the direction of motion. In this chapter we study a more general setting, in which we seek values for these parameters—orientation, initial placement, and direction of motion, for which such a sliding motion of K through W is possible (or determine that no such motion is possible).

In Section 2.1 we observe that if a sliding motion for K exists, then K can also slide in a direction orthogonal to the plane of the window. This is a special case of a more general and known result; we give a simpler proof of this result in our setting. We then describe, in Section 2.2, an algorithm to compute the orientation of K that will allow for sliding along this orthogonal direction. In Section 2.3, we address a sub-problem that arises in the algorithm, and we give an efficient algorithm to compute an implicit representation of all possible silhouettes of K . Finally, in Section 2.4 we present a more efficient algorithm to solve the orientation-determination problem, using intricate batch range-searching data structures.

2.1 The Existence of an Orthogonal Sliding Motion

For the most general version of the sliding motion, in which none of the parameters (orientation, initial placement, and direction of motion) is prespecified, we use the following key lemma:

Lemma 2.1. *If K can slide through W from some starting position in some direction, then*

K can slide through W, possibly from some other starting position and another orientation, by translating it in the z-direction.

Proof. Let K_0 be the starting position of K and \vec{v} be the direction of motion through W , for which the resulting sliding motion is collision-free. Form the infinite prism $\Pi_0 := \bigcup_{\lambda \in \mathbb{R}} (K_0 + \lambda \vec{v})$ that K_0 spans in direction \vec{v} . The premise of the lemma implies that the intersection of Π_0 with the xy -plane is contained in W .

Let W_0 be the orthogonal projection of W onto some plane orthogonal to \vec{v} . Note that W_0 is a parallelogram, and that, by construction, K_0 can pass through W_0 when translated in direction \vec{v} . By an old result, reviewed and proved by Debrunner and Mani-Levitska [9], it follows that, when mapped rigidly into the xy -plane, W_0 (the ‘shadow’ of W in direction \vec{v}) can be placed fully within W (see Figure 2.1).

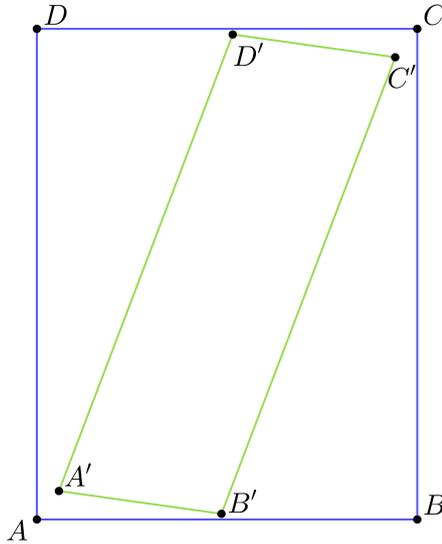


Figure 2.1: The projection of W (green) can be located in a congruent copy of W (blue).

Now rotate and translate \mathbb{R}^3 so that \vec{v} becomes the (negative) z -direction, and the image of W_0 is fully contained in (the former, untransformed copy of) W . Then the image of K under this transformation can be moved vertically down through W , in a collision-free manner, as asserted. \square

The proof of Debrunner and Mani-Levitska [9] is rather involved, and applies to an arbitrary planar convex shape (showing that it contains its projection in any direction). For the sake of completeness, we provide a simple alternative proof for the case of a rectangle.

Lemma 2.2. *Let W be a rectangle on some plane h . Let W_0 be the projection of W on the xy -plane. Then the xy -plane contains a congruent copy of W that contains W_0 .*

Proof. Denote the xy -projection by π . Let l be the intersection line of h and the xy -plane, and let α be the dihedral angle between these planes. Let p be an arbitrary point on h , and

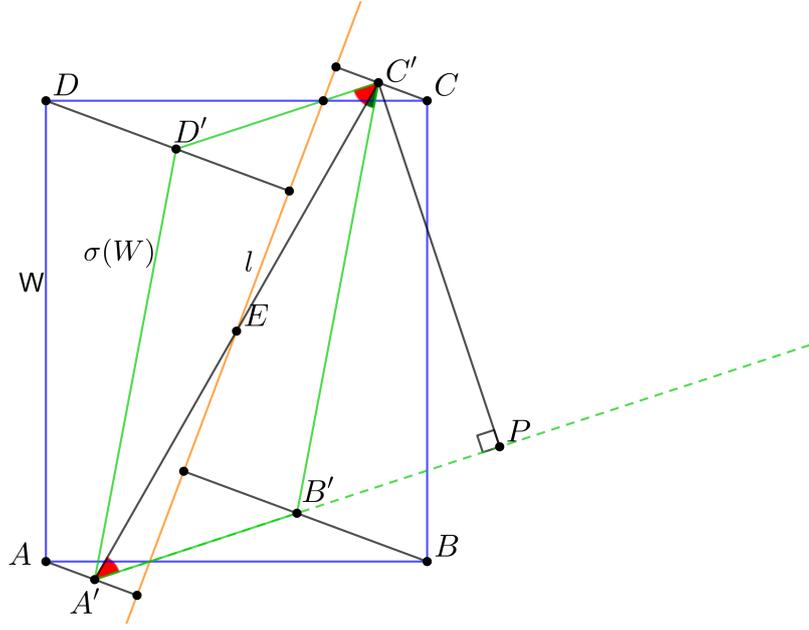


Figure 2.2: The window W (blue), the line l (orange) and the image $\sigma(W)$ (green).

let d be the distance from p to l . Then $\pi(p)$ lies at distance $d \cos \alpha$ from l (with the same nearest point on l). Informally, π moves every point in h closer to l by a factor of $\cos \alpha$. Then, instead of projecting h to the xy plane, we apply on h this linear transformation that moves every point closer to l by a factor of $\cos \alpha$. Denote this transformation by σ . This implies that every line segment in h is transformed to a shorter segment or of the same length—no line segment increases its length.

Let $W = ABCD$, and let $A' = \sigma(A), B' = \sigma(B), C' = \sigma(C), D' = \sigma(D)$. Let E denote the center of W (see Figure 2.2). Note that translating W on h keeps $\sigma(W)$ the same up to translation, so we may assume that l passes through E without loss of generality.

We use the following lemma:

Lemma 2.3. *Assume without loss of generality that B and C lie on one side of l , and that A and D lie on the other side (otherwise rename the vertices as $BCDA$), and that l intersects the ray \overrightarrow{BC} , namely the ray starting at B and passing through C (otherwise rename the vertices as $DCBA$). Then $\sphericalangle A'C'B' \leq \sphericalangle ACB$.*

Proof. Denote by T the intersection point of the lines BC and $B'C'$ (see Figure 2.3). As $\sigma(BC) = B'C'$ the line l must pass through T since it is the only point of BC that stays at the same location when applying σ . We then have:

$$\sphericalangle A'C'B' = \sphericalangle EC'B' = \sphericalangle TEC' + \sphericalangle C'TE \leq \sphericalangle TEC + \sphericalangle CTE = \sphericalangle ECB = \sphericalangle ACB.$$

□

Continuing with the proof of Lemma 2.2. there are two cases to consider:

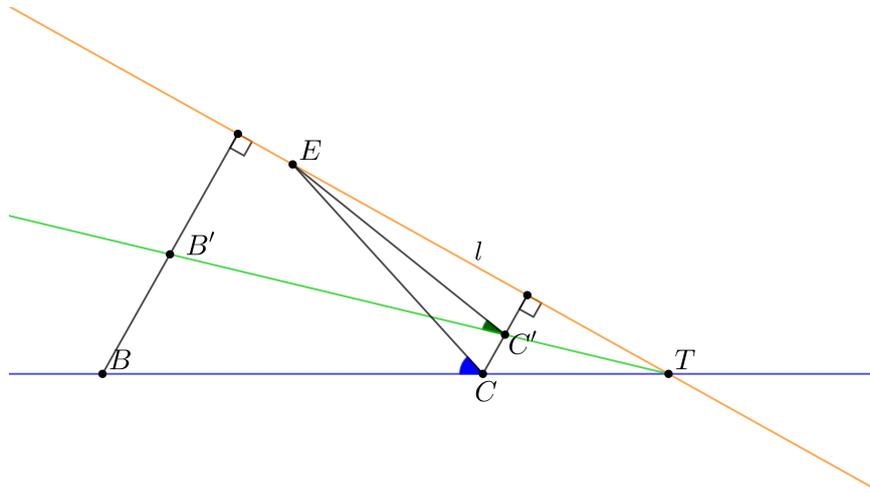


Figure 2.3: The side BC of W (blue), the side $B'C'$ (green), and the line l (orange).

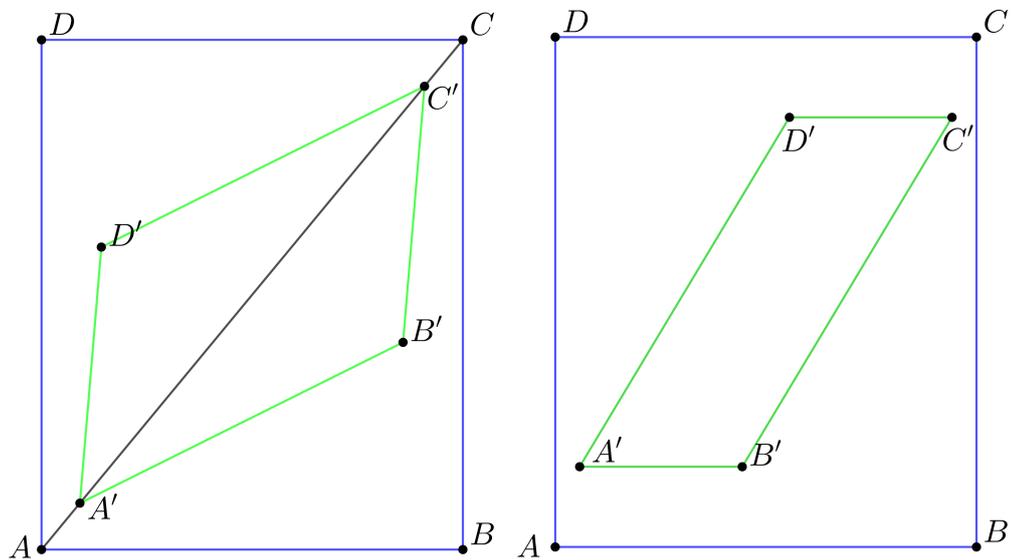


Figure 2.4: Placing $\sigma(W)$ in a congruent copy of W . Left: Placing the diagonal $A'C'$ on the diagonal AC . Right: Placing the side $A'B'$ parallel to the side AB .

For any pair of points P and Q we denote by \overline{PQ} the line through P and Q . If $\sphericalangle C'A'B' \leq \sphericalangle CAB$, then, since no line segment increases its length by applying σ , we have $A'C' \leq AC$. Denote by m the line \overline{AC} . Place $A'C'$ on m , such that the points A, A', C', C appear on m in this order and $A'C'$ is centered at E (see Figure 2.4, left). Note that the angle that \overline{AB} forms with m is greater than the angle that $\overline{A'B'}$ forms with l (by assumption), and that the angle that \overline{BC} forms with m is greater than the angle that $\overline{B'C'}$ forms with m (by Lemma 2.3). Hence B' is inside the triangle ABC . By symmetry, D is inside the triangle CDA , and therefore we successfully placed $\sigma(W)$ inside W .

if $\sphericalangle C'A'B' \geq \sphericalangle CAB$, draw from C' a line perpendicular to $\overline{A'B'}$ and denote the intersection by P (see Figure 2.2). We place $\sigma(W)$ inside W so that $A'B'$ is parallel to AB (see Figure 2.4, right). To do so, we need to prove that $C'P \leq CB$ and that $A'P \leq AB$. Indeed, we have:

$$CB \geq C'B' \geq C'P,$$

$$AB = CB \cdot \cot \sphericalangle CAB \geq C'P \cdot \cot \sphericalangle C'A'B' = A'P.$$

Therefore we successfully placed $\sigma(W)$ inside W . \square

2.2 Finding a Sliding Motion

Lemma 2.1 implies that K can slide through W if and only if there exists an orthogonal projection of K onto some plane $h_{\vec{v}}$, so that the image of K can be placed within a congruent copy of W . We specify such a projection by its direction, represented as a point on the unit sphere \mathbb{S}^2 . This fixes the *projected silhouette* of K , which is the boundary of the convex polygon obtained by the projection of K in direction \vec{v} , up to a possible rigid motion within the image plane $h_{\vec{v}}$. The silhouette itself is the cyclic sequence of vertices and edges of K , whose projections form the projected silhouette.¹ The silhouette and its projection do not change combinatorially, that is, when represented as a cyclic sequence of vertices and edges of K (or of their projections), as long as \vec{v} is not parallel to any face of K . The locus of directions that are parallel to a face f of K is the great circle γ_f of \mathbb{S}^2 that is parallel to f . We draw these $O(n)$ great circles on \mathbb{S}^2 , one circle for each face of K , and form their arrangement \mathcal{A}_0 (see [14]). (This arrangement is also known as the *aspect graph* of K ; see [22].) It consists of $O(n^2)$ faces, and, for all directions \vec{v} within the same face of \mathcal{A}_0 , the silhouette and its projection are fixed combinatorially, but the actual spatial positions of the projected vertices depend on the direction \vec{v} , and the projected silhouette can also rotate arbitrarily within the image plane $h_{\vec{v}}$. (Note that in this discussion we completely ignore translations of K , as they are irrelevant for the analysis and its conclusions.)

We assign some canonical coordinate frame to $h_{\vec{v}}$, and refer, for simplicity, to its axes as the x - and y -axes (they depend on \vec{v}). For example, excluding $O(1)$ problematic directions, which can be handled separately, and easily, we can take the x -axis within $h_{\vec{v}}$ to be the intersection of $h_{\vec{v}}$ with the xz -plane, and take the y -axis to be in the orthogonal direction within

¹The silhouette is indeed such a cycle of vertices and edges of ∂K for generic directions \vec{v} . When \vec{v} is parallel to a face f of K , the entire f is part of the silhouette.

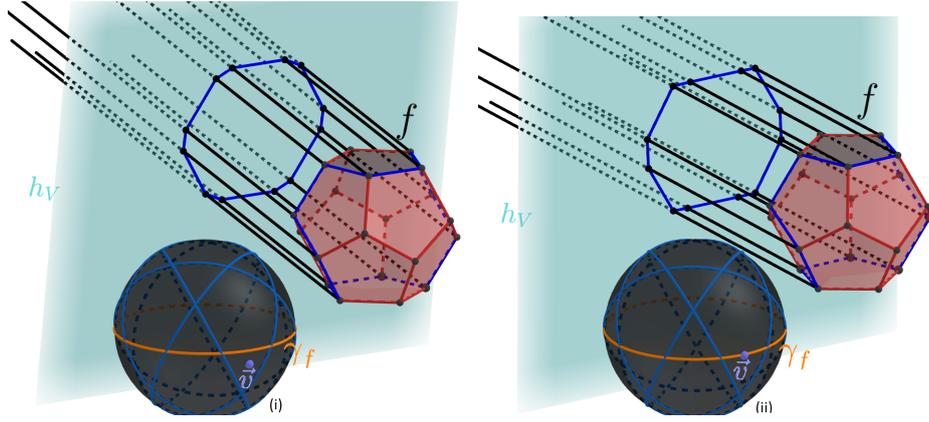


Figure 2.5: The silhouette and its projection: (i) The case of a generic \vec{v} . (ii) The case where \vec{v} is parallel to a face f of K (\vec{v} is on the great circle γ_f).

$h_{\vec{v}}$, oriented in the direction that has a positive y -component. The actual spatial location of the projected silhouette (up to translation, which we ignore) of K can be parameterized by (\vec{v}, θ) , where θ is the rotation of the projected silhouette within the image plane $h_{\vec{v}}$. We refer to (\vec{v}, θ) as the *view* of K . See Figure 2.5.

As we vary \vec{v} and θ , we want to keep track of the leftmost and rightmost vertices of the projected silhouette (in the x -direction), and of the topmost and bottommost vertices (in the y -direction, all with respect to the coordinate frame within $h_{\vec{v}}$). We succeed when we find a projection (in direction \vec{v}), followed by a rotation (by θ), for which the x -difference between the rightmost and leftmost vertices is at most a and the y -difference between the topmost and bottommost vertices is at most b . We reiterate that this is indeed the property that we need: It takes place in a slanted plane $h_{\vec{v}}$ with respect to an artificial coordinate frame within that plane, but using a suitable rotation of $h_{\vec{v}}$ we can make it horizontal and its coordinate frame parallel to the standard xy -frame. A subsequent suitable translation then brings the projected silhouette to within W , as desired.

Fix a face φ of \mathcal{A}_0 , and let w_1, w_2, \dots, w_m denote the cyclic sequence of the vertices of the projected silhouette, say in counterclockwise order, for views in φ . If the current leftmost vertex is some w_j , then it remains leftmost as long as neither of the two adjacent edges $w_{j-1}w_j$ and w_jw_{j+1} becomes y -vertical. (Recall that ‘leftmost’ and ‘ y -vertical’ are with respect to the artificial frame within $h_{\vec{v}}$.) The views (\vec{v}, θ) at which an edge e of K , say, $w_{j-1}w_j$ is y -vertical comprise a two-dimensional surface ρ_e in the three-dimensional space $V = \mathbb{S}^2 \times \mathbb{S}^1$ of views (\vec{v}, θ) . See Figure 2.6.

The discussion so far has been for views that have a combinatorially fixed silhouette. However, to make the algorithm for finding a sliding motion more efficient, we consider all possible silhouettes ‘at once’, using the following approach. After forming the aspect-graph arrangement \mathcal{A}_0 , as defined above, we replace each great circle γ_f on \mathbb{S}^2 by the cylindrical surface $\gamma_f^* = \gamma_f \times \mathbb{S}^1$, and collect these surfaces into a set Γ , of cardinality $O(n)$. Then, for each edge e of K (regardless of whether it is a silhouette edge or not), we form the surface ρ_e , as just defined, and collect these surfaces into a set Σ , of cardinality n . We now form

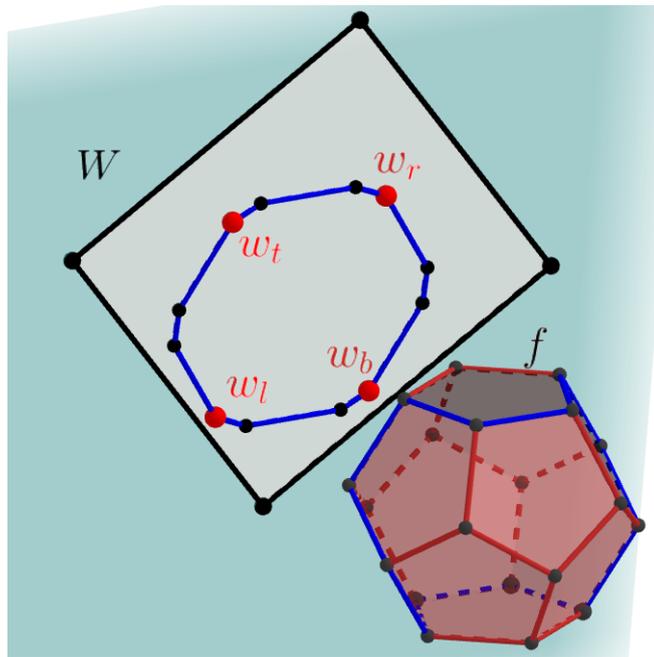


Figure 2.6: A view of K . To simplify the visualization, we rotate the containing window W rather than the projected silhouette. The leftmost, rightmost, topmost and bottommost vertices are highlighted.

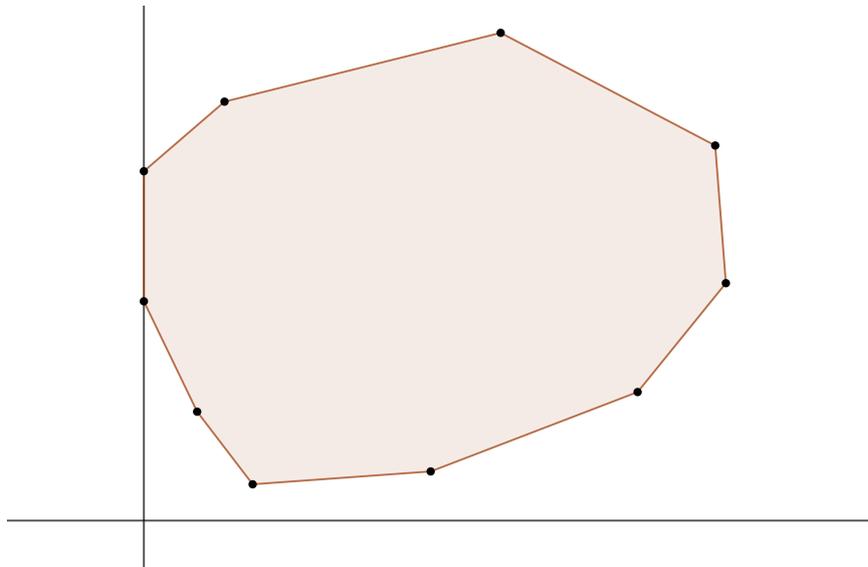


Figure 2.7: A discrete change of the leftmost vertex of the projected silhouette.

the three-dimensional arrangement $\mathcal{A} = \mathcal{A}(\Gamma \cup \Sigma)$ (note that all the surfaces of $\Gamma \cup \Sigma$ are two-dimensional). As is easily verified, for each three-dimensional cell τ of \mathcal{A} , the projected silhouette of K , and its four leftmost, rightmost, topmost and bottommost vertices (we refer to them collectively as the *extreme* vertices of the projected silhouette) are fixed for all views in τ . Since $|\Gamma \cup \Sigma| = O(n)$, the complexity of \mathcal{A} is $O(n^3)$.

To obtain a representation that is easy to process further, we construct the *vertical decomposition* of \mathcal{A} , which we denote as $\text{VD}(\mathcal{A})$. It is a decomposition of the three-dimensional cells of \mathcal{A} into a total of nearly cubic number of prism-like subcells (that we simply call prisms). See Sharir and Agarwal [25, Section 8.3] for more details. A sharp bound on its complexity (i.e., the number of prisms) is $O(n^2 \lambda_s(n))$, for some constant s (a sharp estimation of the value of s is not given in this thesis), where $\lambda_s(n)$ is the maximum length of a Davenport–Schinzel sequence of order s on n symbols; see [25]. The vertical decomposition can be constructed in time $O((n^2 \lambda_s(n) \log n))$ [3].

We now iterate over all prisms of $\text{VD}(\mathcal{A})$. For each prism τ , we retrieve the four extreme vertices of the projected silhouette, which are fixed for all views in τ , and check whether there is a view in τ for which these vertices, and thus all of the projected silhouette, fit into W (after suitable rotation and translation of W , as discussed above). To do so, denote these leftmost, rightmost, topmost and bottommost vertices as w_l , w_r , w_t and w_b , respectively. The x -coordinates x_{w_l} , x_{w_r} of w_l and w_r , and the y -coordinates y_{w_t} , y_{w_b} of w_t and w_b (within $h_{\vec{v}}$) are functions of (\vec{v}, θ) . We need to determine whether the region

$$S = S(w_l, w_r, w_t, w_b) := \{(\vec{v}, \theta) \in \mathbb{S}^2 \times \mathbb{S}^1 \mid x_{w_r}(\vec{v}, \theta) - x_{w_l}(\vec{v}, \theta) \leq a, y_{w_t}(\vec{v}, \theta) - y_{w_b}(\vec{v}, \theta) \leq b\},$$

which is exactly the region of views (\vec{v}, θ) at which W contains a (rotated and translated) copy of the projected silhouette with these four specific vertices as the extreme vertices of the projection, has a nonempty intersection with τ . Since S and τ are semialgebraic regions of constant complexity, this test can be performed, in a suitable (and standard) model of real algebraic computation, in constant time [10]. Summing over all prisms τ , the overall cost of these tests is proportional to the complexity of $\text{VD}(\mathcal{A})$, namely it is $O(n^2 \lambda_s(n))$.

To complete the description of the algorithm, we now consider the task of computing the four extreme vertices w_l , w_r , w_b and w_t of the silhouette, or, more precisely, the four (fixed) vertices of K that project to them, for each cell c of \mathcal{A} . As an easy by-product of the construction of $\text{VD}(\mathcal{A})$, each of its prisms can be associated with the cell of \mathcal{A} containing it, so the four extreme vertices will also be available for each prism of $\text{VD}(\mathcal{A})$.

By the nature of the surfaces forming \mathcal{A} , the projection of each cell c of \mathcal{A} onto \mathbb{S}^2 is fully contained in a single cell $\rho = \rho(c)$ of the two-dimensional aspect-graph arrangement \mathcal{A}_0 . For each such cell ρ , the discrete nature of the silhouette, as a cyclic sequence of vertices (and edges) of K , is fixed for every $\vec{v} \in \rho$ and for any $\theta \in \mathbb{S}^1$. Although we can do it faster, we simply iterate over the $O(n^2)$ cells of \mathcal{A}_0 , and for each cell ρ , compute the silhouette in $O(n)$ time, in brute force (by picking an arbitrary point \vec{v} in ρ , and by examining each edge of K for being part of the silhouette in direction \vec{v}). The overall cost of this step is thus $O(n^3)$.

Consider now a cell c of \mathcal{A} , and let $\rho = \rho(c)$ be the cell of \mathcal{A}_0 that contains the \mathbb{S}^2 -projection of c . Let (u_1, u_2, \dots, u_m) denote the cyclic counterclockwise sequence of vertices of K that forms the silhouette for directions in ρ , and let w_i denote the \mathbb{S}^2 -projection of

u_i , for $i = 1, \dots, m$. Since the vertices of K inducing w_l, w_r, w_b and w_t are fixed over c , it suffices to compute them for a fixed arbitrary view in c . We thus fix such a view (\vec{v}, θ) , and proceed as follows.

For each i , define the “derivative” of the silhouette at w_i to be the pair of vectors

$$(\mathbf{w}_i^-, \mathbf{w}_i^+) = (\overrightarrow{w_{i-1}w_i}, \overrightarrow{w_iw_{i+1}}),$$

where the vectors are represented in the coordinate frame induced by (\vec{v}, θ) in a plane orthogonal to \vec{v} , and where addition and subtraction of indices is modulo m . The extreme vertices w_l, w_r, w_b, w_t partition the silhouette into (at most) four subsequences: S_1 , between w_r and w_t , S_2 , between w_t and w_l , S_3 , between w_l and w_b , and S_4 , between w_b and w_r (see Figure 2.8), so that, for $w_i \in S_1$ (resp., S_2, S_3, S_4) both vectors $\mathbf{w}_i^-, \mathbf{w}_i^+$ lie in the second (resp., third, fourth, first) quadrant. For w_r (resp., w_t, w_l, w_b), the vectors lie, respectively, in the first and second (resp., second and third, third and fourth, fourth and first) quadrants.²

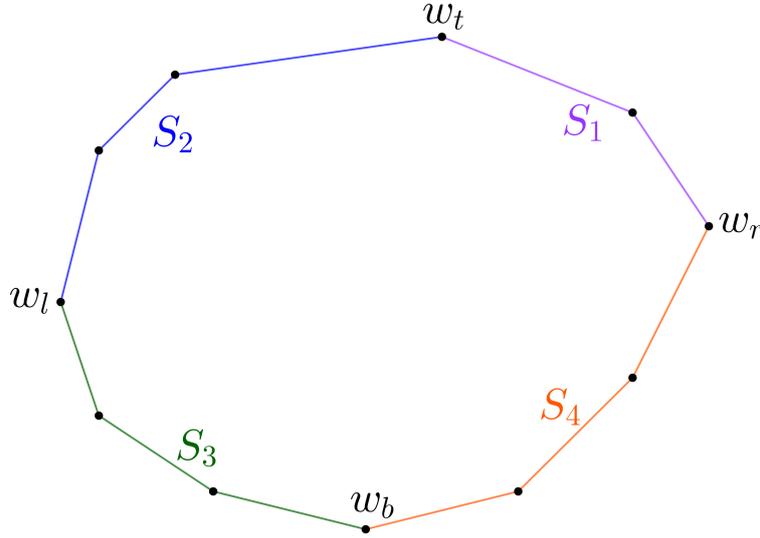


Figure 2.8: The sub-silhouettes S_1 (purple), S_2 (blue), S_3 (green), S_4 (orange). The extreme vertices w_r, w_t, w_l, w_b delimit these sub-silhouettes and are highlighted.

Using these observations, we find the four extreme vertices using binary search, as follows. We break the silhouette sequence into two linear subsequences at w_1 and $w_{m/2}$, and find the extreme vertices in each subsequence. Consider the subsequence $(w_1, w_2, \dots, w_{m/2})$. We compute the derivatives at w_1 and at $w_{m/2}$, and thereby identify the two respective sub-silhouettes that contain these vertices. Suppose for specificity that w_1 lies in S_1 and $w_{m/2}$ lies in S_3 . Then we know that our subsequence contains (only) w_t and w_l , and we can find each of them by a straightforward binary search, using the derivatives to guide the search.

²We gloss here over the easy special cases of degeneracy, in which the extreme vertices are not all distinct. In such cases some of the sub-silhouettes S_1, \dots, S_4 might be empty, and the rules for identifying the extreme vertices need to be adjusted.

We apply similar procedures in each of the other cases, and for the second subsequence $(w_{m/2}, w_{m/2+1}, \dots, w_1)$.

In conclusion, it takes $O(\log n)$ time to find the extreme vertices for each cell of \mathcal{A} , and thus also for each prism of $\text{VD}(\mathcal{A})$, for a total running time of $O(n^2 \lambda_s(n) \log n)$.

2.3 Implicitly Constructing All the Silhouettes

We do not have a faster approach to the full problem of finding a view for which the projection of K can fit inside W . Nevertheless, we present here an improved solution, which runs in $O(n^2 \log n)$ time, for the subproblem of (implicitly)³ constructing the silhouette for every face of the aspect-graph arrangement \mathcal{A}_0 . That is, our goal is to process the arrangement \mathcal{A}_0 of the great circles γ_f on \mathbb{S}^2 , and implicitly store with each of its faces φ the (fixed) circular list Λ_φ of the vertices of the projected silhouette. A key step towards this goal is to store, for each edge g of \mathcal{A}_0 , that separates two adjacent faces ζ, ζ' of \mathcal{A}_0 , the change between the silhouettes stored at ζ and at ζ' .

To understand this last issue, we note that g is a portion of some great circle γ_f , for some face f of K , so that f is visible for directions \vec{v} on one side of γ_f , say ζ , and invisible for directions \vec{v} on the other side, that is ζ' . For these nearby directions, when f is visible, one connected portion $\text{vis}(f)$ of ∂f is part of the silhouette, and when f is invisible, the complementary portion $\text{invis}(f)$ of ∂f becomes part of the silhouette. These two portions are fixed for each edge of γ_f , like g , but are not fixed throughout γ_f . Specifically, let g' be the next edge of \mathcal{A}_0 along γ_f . The common endpoint \vec{v} of g and g' is an intersection of γ_f with another great circle $\gamma_{f'}$. If f and f' are not adjacent along ∂K , the two portions $\text{vis}(f)$ and $\text{invis}(f)$ for g' are the same as for g , and no special treatment is needed. If f and f' share an edge e of K then, by construction, \vec{v} is the direction of e . Then, as is easily checked, e belongs to $\text{vis}(f)$ at g and to $\text{invis}(f)$ at g' or the other way around. (Note that, at \vec{v} itself, e might fail to belong to either silhouette.) See Figure 2.9 for an illustration.

This suggests that we trace the edges of \mathcal{A}_0 along each great circle γ_f in order, maintain the two portions $\text{vis}(f)$ and $\text{invis}(f)$ of ∂f for each edge g of γ_f , by a pair of pointers into the circular list of edges of ∂f (which we assume to be available from the DCEL representation of ∂K), and update these pointers by moving one of them past e , where e is the edge of f whose direction separates g from the next edge on γ_f (if such an e exists at all). The overall number of changes of these pointers, over all faces f of K , is thus proportional to the complexity of K , i.e., it is $O(n)$. However, since most of the $O(n^2)$ vertices of \mathcal{A}_0 do not entail any change of pointers, the overall cost of this preprocessing step is $O(n^2)$, which is the number of edges of \mathcal{A}_0 .

Having computed this data, over all edges of \mathcal{A}_0 , we now return to the main goal, of (implicitly) constructing the silhouette for each face of \mathcal{A}_0 . As a matter of fact, it will be more convenient to construct the (degenerate) silhouettes over each edge of \mathcal{A}_0 . We will

³An explicit construction of the silhouettes is doomed to require $\Theta(n^3)$ time in the worst case, since \mathcal{A}_0 has $\Theta(n^2)$ cells and each silhouette might be of linear size.

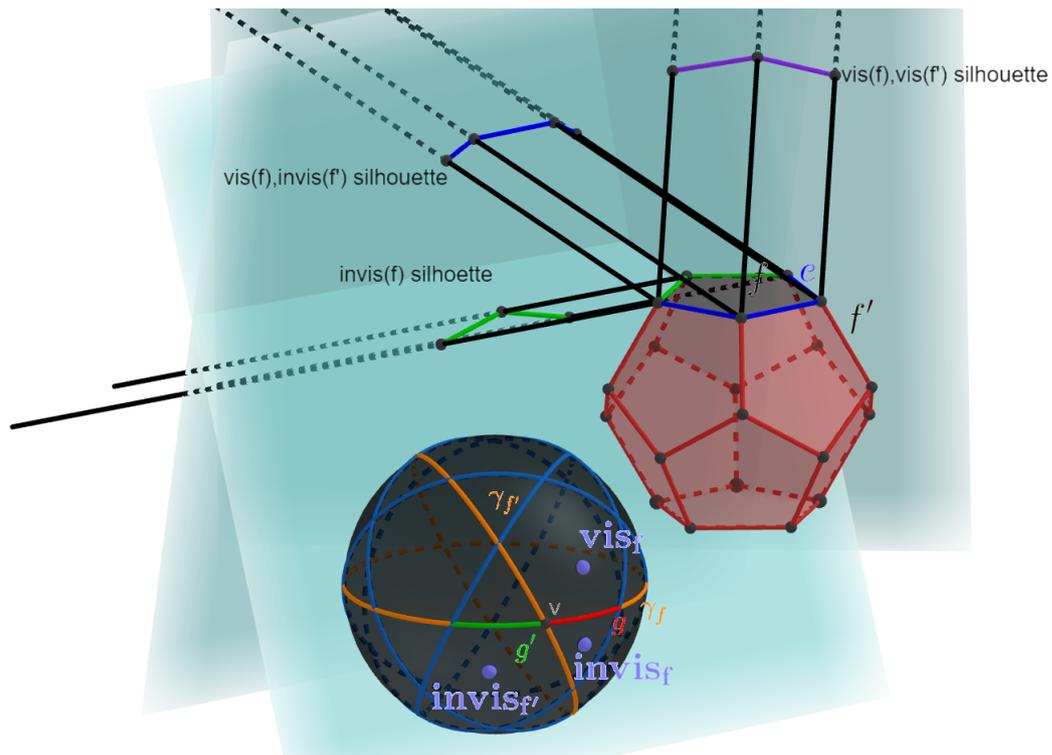


Figure 2.9: f is the top face and f' is the right face adjacent to f . When f is invisible from the direction \vec{v} (the antipodal to the point marked as vis_f on the unit sphere), $invis(f)$ is a part of the projected silhouette (green). When f is visible (the antipodal to the point marked as $invis_f$ on the unit sphere), the complementary portion $vis(f)$ is a part of the projected silhouette (blue). When \vec{v} is the direction marked as $invis_{f'}$ we are at the portion g' of γ_f (green arc) instead of g (red arc) and therefore the portion $vis(f)$ (purple) changes: e belongs to $vis(f)$ at g and to $invis(f)$ at g' .

later comment on how to extend the algorithm so that it can also retrieve the silhouette for each two-dimensional face of \mathcal{A}_0 .

We fix a face f of K , consider the corresponding great circle γ_f on \mathbb{S}^2 , and construct the silhouette over each edge of γ_f . We enumerate these edges as g_1, \dots, g_m , where $m = O(n)$ is the number of faces of K minus 1. (The sequence is actually cyclic but we consider it as linear, starting at some arbitrary edge g_1 .) We compute the silhouette over g_1 in $O(n)$ time, by brute force. We prepare a persistent search tree structure that stores all the versions of the silhouette over the edges g_i . To obtain the silhouette at g_i from that at g_{i-1} , we consider the great circle $\gamma_{f'}$ that separates g_{i-1} and g_i , and note that the silhouette changes by replacing the edges in $\text{vis}(f')$ by those in $\text{invis}(f')$, or vice versa, and we update our persistent search structure accordingly.

The overall number of changes, as we trace the entire γ_f , is proportional to the sum of the number of edges on the other faces of K (we encounter each other face exactly twice). Since each edge of K is counted in this sum twice, the overall number of updates over γ_f is $O(n)$, and the total cost of these updates is $O(n \log n)$. Summed over all great circles γ_f , this gives a total cost of $O(n^2 \log n)$ time.

The actual output of the procedure is slightly more involved, to facilitate the handling of the silhouettes over the two-dimensional faces of \mathcal{A}_0 . Note that the entire face f is always part of the silhouette over all edges along γ_f , but our output will store instead, for each edge g of γ_f , the two portions $\text{vis}(f)$ and $\text{invis}(f)$ of ∂f that replace f on the silhouette on the two sides of g (which have already been computed). This step is also done implicitly, by storing only the indices of the edges of f that delimit the portions $\text{vis}(f)$ and $\text{invis}(f)$. The overall cost of this supplementary maintenance is $O(n^2)$ time.

We can now access the silhouette in a given direction \vec{v} as follows. We locate \vec{v} in \mathcal{A}_0 . If it lies on an edge g , we access the silhouette stored at g using the persistent search structure associated with the great circle γ_f containing g . If \vec{v} lies in a face φ of \mathcal{A}_0 , we take an (arbitrary) edge g of φ , lying on some γ_f , identify whether f is visible or invisible for directions in φ , and access the version of the silhouette at g that uses, respectively, $\text{vis}(f)$ or $\text{invis}(f)$ as a subsequence of the silhouette.

If our goal is to report the silhouette in direction \vec{v} , we can do it in $O(\log n + k_{\vec{v}})$ time, where $k_{\vec{v}}$ is the number of edges of the silhouette. While being a nontrivial result, it is not significant in the worst case, since we can easily compute the silhouette at \vec{v} in brute force, in $O(n)$ time. A more significant performance improvement is obtained for tasks that involve searching in the silhouette, like finding the vertex of the silhouette that is extreme in some query direction, which can be accomplished in $O(\log n)$ time.

2.4 Improved Algorithm

We next present an improved, albeit more involved algorithm that solves the problem of finding an orientation of K that will allow for a sliding motion through W if one exists, in time $O(n^{8/3+\varepsilon})$, for any $\varepsilon > 0$. The problem of finding a direction \vec{v} in which we can slide K

through W is equivalent to the problem of finding a placement of W on some plane h such that the projected silhouette of K on h is contained in W , which in turn is equivalent to verifying that all the vertices of K are projected into that placement of W .

An equivalent way of checking for the latter characterization is to look for two unit vectors x and y (which will be the directions of the axes of W in the desired placement; note that h is spanned by x and y) that satisfy:

- (i) x and y are perpendicular to each other.
- (ii) For every segment e connecting two vertices of K we have $\langle x, e \rangle \leq a$.
- (iii) For every segment e connecting two vertices of K we have $\langle y, e \rangle \leq b$.

(Note that since we go over all unordered pairs of vertices of K in (ii), (iii), we actually require that $|\langle x, e \rangle| \leq a$ and $|\langle y, e \rangle| \leq b$ for each such segment e .) Every inequality in (ii) defines a halfspace that has to contain x . We intersect those $O(n^2)$ halfspaces, to obtain a convex polytope Q of complexity $O(n^2)$, and intersect Q with the unit sphere \mathbb{S}^2 to obtain the admissible region A of the vectors x that satisfy (ii), in $O(n^2 \log n)$. We apply the same procedure for y using the suitable collection of halfspaces in (iii), and obtain the admissible region B for the vectors y that satisfy (iii), also in $O(n^2 \log n)$. To satisfy also (i), we need to check whether there exist an orthogonal pair of vectors $x \in A, y \in B$. We use the following lemma:

Lemma 2.4. *Let S_A denote the set of all vertices of A , and let T_A denote the set of the points that are closest locally to the north pole of \mathbb{S}^2 along each circular arc of the boundary of A . (By choosing a generic direction for the north pole of \mathbb{S}^2 we may assume that T_A is finite and $|S_A \cup T_A| = O(n^2)$.) Define similarly the sets S_B, T_B . If there exist an orthogonal pair $(x, y) \in A \times B$ then there exist such an orthogonal pair so that either $x \in S_A \cup T_A$ or $y \in S_B \cup T_B$.*

Proof. We refer to an orthogonal pair in $A \times B$ as a *good pair*. Let (x, y) be a good pair such that x is as close to the boundary of A as possible. If there are multiple pairs with this property, pick the one in which x is the closest to the north pole. If there are still multiple pairs, pick an arbitrary pair among them. By continuity and the compactness of A and B , it is easy to show that such a “minimal pair” exist.

Several cases can arise:

1. x or y is one of the desired vertices. In this case we are done.
2. Both x and y lie in the interiors of A and B , respectively. In this case they can be moved slightly together in any direction, while maintaining their mutual orthogonality. In particular, x can get closer to the boundary of A so (x, y) is not the minimal pair.
3. x is on the boundary of A , and y is in the interior of B . Since we are not in Case 1, x lies in the relative interior of an edge of ∂A and is not the point on that edge that is

closest to the north pole. Then we have two available directions to move (x, y) slightly such that x remains on the same edge. One of these directions brings x to a point closer to the north pole, so (x, y) is not the minimal pair.

4. y is on the boundary of B (as in Case 3 we may assume that y lies in the relative interior of an edge of ∂B). In this case we fix x and move y along the great circle C_x of points perpendicular to x . Recall that y is on an edge of B , which is a circular arc γ . Every halfspace of the intersection contains the origin, so B is contained in the bigger portion C^+ (bigger than a hemisphere) of \mathbb{S}^2 that is bounded by the circle C containing γ . Since C_x is a great circle, it is bigger than C , so when moving y along C_x in at least one of the two possible directions, y enters C^+ (this is always true, regardless of the size of C_x , when the circles cross one another at y ; the fact that C_x is larger is needed when they are tangent at y), so it enters the interior of B . Now we are in one of the cases 2, 3 that we have already settled.

Having covered all possible cases, this completes the proof of the lemma. \square

We iterate over the points of $S_A \cup T_A$. For each such point v let C_v be the great circle of vectors perpendicular to v , and let \mathcal{C} denote the collection of these $O(n^2)$ great circles. We face the problem of determining whether any great circle in \mathcal{C} crosses B . This is the same as determining whether any great circle in \mathcal{C} crosses an arc of ∂B . This is a variant of the batched range searching paradigm, and we present next a detailed solution for this case. We apply a fully symmetric procedure to the collection of great circles orthogonal to the points of $S_B \cup T_B$ and to A . If we find a valid intersection it gives us a valid orthogonal pair. Otherwise, such a pair does not exist.

Detecting an intersection between the great circles of \mathcal{C} and the boundary arcs of B . We apply a central projection (from the center of \mathbb{S}^2) onto some plane, say a horizontal plane h lying below \mathbb{S}^2 (with a generic choice of the coordinate frame, we may assume that none of the points in $S_A \cup T_A \cup S_B \cup T_B$ are on the great circle that is parallel to h). This is a bijection of the open lower hemisphere onto h , in which (the lower portions of) great circles are mapped to lines, and (the lower portions of) circular arcs are mapped to arcs of conic sections (ellipses, parabolas, hyperbolas, or straight lines). This transforms the problem into a batched range searching problem, in which we have a set L of $M = O(n^2)$ lines (which arise from the great circles orthogonal to the points of $S_A \cup T_A$) and a set E of $N = O(n^2)$ pairwise disjoint arcs of conic sections (which are the projections of the arcs forming the boundary of B), and the goal is to determine whether any line in L crosses any arc in E . We note that the halfspaces from which we obtain B come in pairs that are symmetric to each other about the origin, so restricting the problem to the lower hemisphere incurs no loss of generality. We also note that there might be situations in which one of the great circles is fully contained in B , but these cases are easy to detect, e.g., by picking an arbitrary point on each great circle and checking whether it belongs to B , using a suitable point-location data structure on B .

To simplify the presentation, we assume that the arcs of E are elliptic arcs; handling the cases of parabolic or hyperbolic arcs is done in essentially the same manner.

Orient all the lines of L from left to right. We may assume that all the arcs in E are x -monotone (otherwise we break each arc that is not x -monotone at its leftmost and rightmost points, into at most three x -monotone subarcs). We orient all these (sub)arcs also from left to right. We also treat separately *convex* arcs, namely arcs for which the tangent directions turn counterclockwise as we traverse them from left to right, and *concave* arcs, for which the tangent directions turn clockwise. The treatments of these two subfamilies are fully symmetric, so we only consider the case of convex arcs.

A line ℓ intersects a convex x -monotone arc γ of some ellipse e , both oriented as above, if and only if one of the following conditions holds.

- (i) The two endpoints of γ lie on different sides of ℓ . See Figure 2.10(i).
- (ii) The two endpoints of γ lie to the left of ℓ and ℓ intersects e . For this to happen, γ must have a tangent that is parallel to ℓ . That is, the slope of ℓ must lie between the slopes of the tangents to γ at its endpoints. When all these conditions hold, it suffices to require that ℓ lies to the left of the right tangent to e with the same slope of ℓ . See Figure 2.10(ii,iii).

To test for intersections of type (i), we use a two-level data structure, where each level is a standard tree-like range searching structure for points and halfplanes (see [1, 2]). The first level collects the arcs that have one endpoint to the right of ℓ , and the second level tests whether any of these arcs has its other endpoint to the left of ℓ . Using the standard machinery for point-halfplane range searching (see, e.g., [1, Theorem 6.1], and also [2]), this takes time $O(M^{2/3}N^{2/3}\text{polylog}(M+N)) = O(n^{8/3}\text{polylog}(M+N))$.

To test for intersections of type (ii), we use a four-level data structure, where, as before, the first two levels are standard range searching structures for points and halfplanes, so that the first level collects the arcs that have their left endpoint to the left of ℓ , and the second level collects, from among the arcs in the output of the first level, those arcs that have their right endpoint also to the left of ℓ . The third level is a one-dimensional segment tree on the interval ranges of the slopes of the tangents to the arcs, and it collects those arcs whose tangent-slope range contains the slope of ℓ . Finally, the fourth level tests whether any of the arcs is such that its tangent that is parallel to ℓ passes to the right of ℓ .

To implement the fourth level, we note that the lines that are tangent to the ellipse e and have slope a can be written as $y = ax + \varphi_e^-(a)$ and $y = ax + \varphi_e^+(a)$, with $\varphi_e^-(a) < \varphi_e^+(a)$, where $\varphi_e^-(a)$ and $\varphi_e^+(a)$ are algebraic functions of constant degree that depend on e . If ℓ has the equation $y = ax + b$ then we need to test whether there exists an ellipse e such that $b > \varphi_e^-(a)$. We thus compute the lower envelope of the functions φ_e^- in time nearly linear in the number of arcs, and then, given a line $y = ax + b$, we test whether the point (a, b) lies above the envelope, in logarithmic time.

It is easy to see that in this case too, the overall cost is $O(n^{8/3}\text{polylog}(n))$. In conclusion, we have shown:

Theorem 2.5. *Given K and W as above, we can determine whether K can slide through W in a collision-free manner, and, if so, find such a sliding motion, in time $O(n^{8/3}\text{polylog}(n))$.*

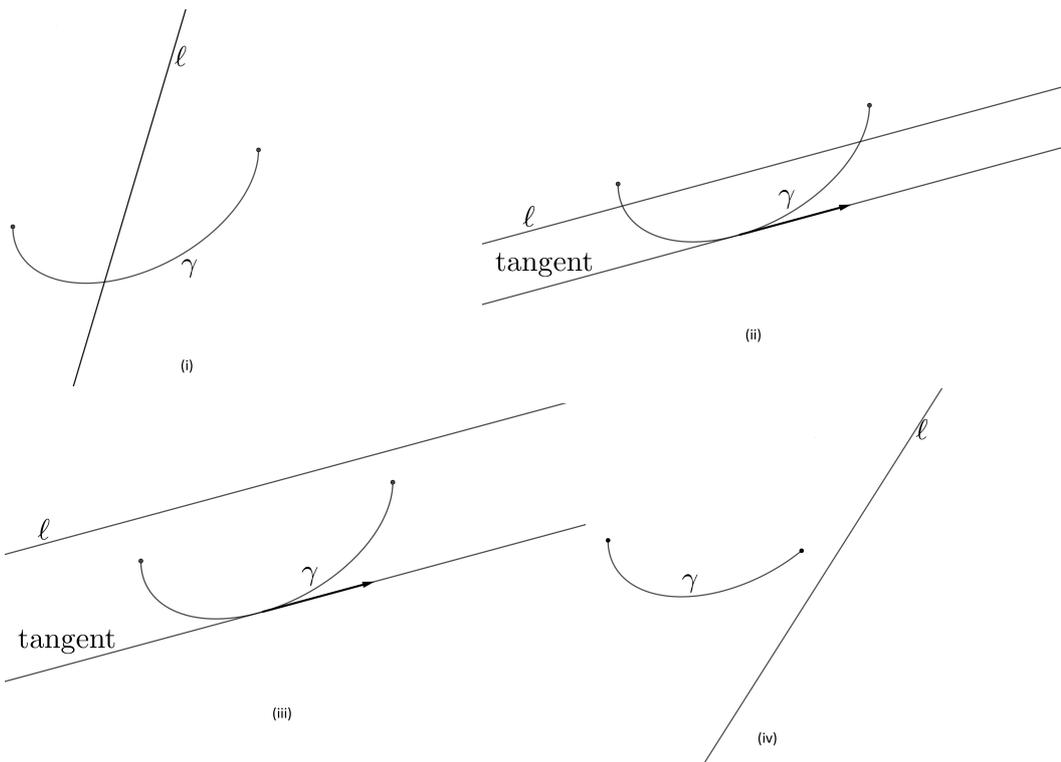


Figure 2.10: A line l intersecting a convex x -monotone elliptic arc γ : (i) The two endpoints of γ lie on different sides of l . (ii) The two endpoints lie to the left of l and l lies to the left of the parallel tangent to the arc. (iii) The two endpoints lie to the right of l (and then there is no intersection). (iv) The two endpoints lie to the left of l but γ has no tangent parallel to l (and then there is no intersection).

We are not aware of any published result that solves the specific problem at hand, of determining whether any great circle in \mathcal{C} crosses B , with comparable running time. A different solution, with a similar performance bound, was suggested to us by Pankaj Agarwal, and we thank him deeply for the useful interaction concerning this problem.

3

Unbounded Windows

In this chapter we consider the variant of the general problem in which W is an infinite slab in the xy -plane, bounded by, say, two vertical lines $x = 0$ and $x = a$. We refer to such a window as a *gate*. The problem is to determine whether K can be moved through W , by an arbitrary collision-free motion.

We claim that the problem can be reduced to the setup studied in the previous chapter, of translation in a single direction (sliding).

We first establish this claim for the case where K is a smooth compact strongly convex body, and then use a compactness argument to extend the result to convex polytopes. We note that strong convexity is not needed for the analysis in this chapter, but only in subsequent chapters. To avoid duplication, we use the same scheme for approximating a convex polytope, both here and in Chapter 4.

Let K be an arbitrary compact convex body in \mathbb{R}^3 . Let h denote the xz -plane, and let $g := h \cap W$, which is the segment $0 \leq x \leq a$, $z = 0$ within h . The two complementary rays to g within the x -axis form the only obstacles within h . Let π denote the orthogonal projection of 3-space onto h .

Assume that K can be moved through W by an arbitrary collision-free rigid motion, which we represent as a continuous map on $[0, 1]$ (a ‘time interval’), where, for each $t \in [0, 1]$, $K(t)$ denotes the placement of K at time t during the motion. For each $t \in [0, 1]$, $\partial\pi(K(t))$ is the projection of the silhouette of $K(t)$ on h . It is a time-varying convex region within h , whose shape is not rigidly fixed. For a convex polytope K , the projected silhouette $\partial\pi(K(t))$ is a time-varying convex polygon.

We have the following property, whose easy proof is omitted.

Lemma 3.1. *The motion $t \mapsto K(t)$ is collision-free, and moves K through W from a*

placement $K(0)$ in the upper halfspace to a placement $K(1)$ in the lower halfspace, if and only if the map $t \mapsto \pi(K(t))$ is collision-free within h , and moves the (time-varying) projection $\pi(K(t))$ through g from the placement $\pi(K(0))$ in the upper halfplane $z > 0$ to the placement $\pi(K(1))$ in the lower halfplane $z < 0$.

We note that in Lemma 3.1 the body K is not required to be smooth, but this requirement is needed for the proof of the following theorem.

Theorem 3.2. *Let K be a smooth compact convex body that can be moved, by a collision-free rigid motion, through W from a placement in the upper halfspace $z > 0$ to a placement in the lower halfspace $z < 0$. Then there exists a sliding collision-free motion of K through W .*

Proof. Let K be as in the theorem, and let $t \mapsto K(t)$ be a collision-free rigid motion that takes K through W , as in the theorem statement. For each t , $\pi(K(t))$ is also smooth (as a planar convex region). Put $\kappa(t) := \pi(K(t)) \cap g$, which is a subsegment of g (by assumption, and by Lemma 3.1, the intersection of $\pi(K(t))$ with the x -axis is always fully contained in g). $\kappa(t)$ is empty at the beginning and at the end of the motion, namely during some prefix interval and some suffix interval of $[0, 1]$ (if the motion is ‘crazy’ enough, $\kappa(t)$ might also be empty during some other inner intervals of $[0, 1]$). Nevertheless, since $\pi(K(t))$ crosses g from side to side, there must exist at least one closed maximal connected interval $I = [t_1, t_2]$ within $[0, 1]$ such that $\kappa(t) \neq \emptyset$ for all $t \in I$, and such that $\kappa(t_1)$ and $\kappa(t_2)$ are singletons, so that $\kappa(t_1)$ (resp., $\kappa(t_2)$) is the z -lowest (resp., z -highest) point of $\pi(K(t_1))$ (resp., of $\pi(K(t_2))$). See Figure 3.1 for an illustration.

Denote, for $t \in I$, the left and right endpoints of $\kappa(t)$ by $\kappa^-(t)$ and $\kappa^+(t)$, respectively, and let $\tau^-(t)$ (resp., $\tau^+(t)$) denote the tangent to $\pi(K(t))$ at $\kappa^-(t)$ (resp., at $\kappa^+(t)$), where we orient both tangents so that $\pi(K(t))$ lies to their right.

Since $\pi(K(t))$ is smooth, the two tangents are well defined and unique. Moreover, since the motion of $K(t)$ is continuous, so is the ‘motion’ of $\pi(K(t))$, and this is easily seen to imply that the directions $\mu^-(t)$ of $\tau^-(t)$, and $\mu^+(t)$ of $\tau^+(t)$ are also continuous functions of t .

Consider the map $\varphi(t)$ that maps $t \in I$ to the counterclockwise angle between $\mu^-(t)$ and $\mu^+(t)$. The map is undefined at t_1 and at t_2 , but we assume that it is defined everywhere in the interior of I (as would be the typical situation—see the comment made earlier). $\varphi(t)$ is clearly a continuous function. For t slightly larger than t_1 , $\varphi(t)$ has a small positive value, and for t slightly smaller than t_2 , $\varphi(t)$ is close to 2π . It follows, by continuity, that there exists $t_0 \in I$ for which $\varphi(t_0) = \pi$, that is, the two tangents at $\kappa^-(t_0)$ and at $\kappa^+(t_0)$ are parallel to each other. This means that $\pi(K(t_0))$ is contained in the slab σ , within h , bounded by the two tangent lines. This in turn implies that $K(t_0)$ is contained in the three-dimensional slab S which is the Cartesian product of σ and the y -axis. Moreover, the intersection of S with the xy -plane is a y -vertical slab that is contained in W (see Figure 3.2 for illustration). This in turn means that, if we fix the orientation of K to be that of $K(t_0)$, we can slide K within S through W (note that there are infinitely many ways to do so, each with its own y -component of the sliding direction). This completes the proof. \square

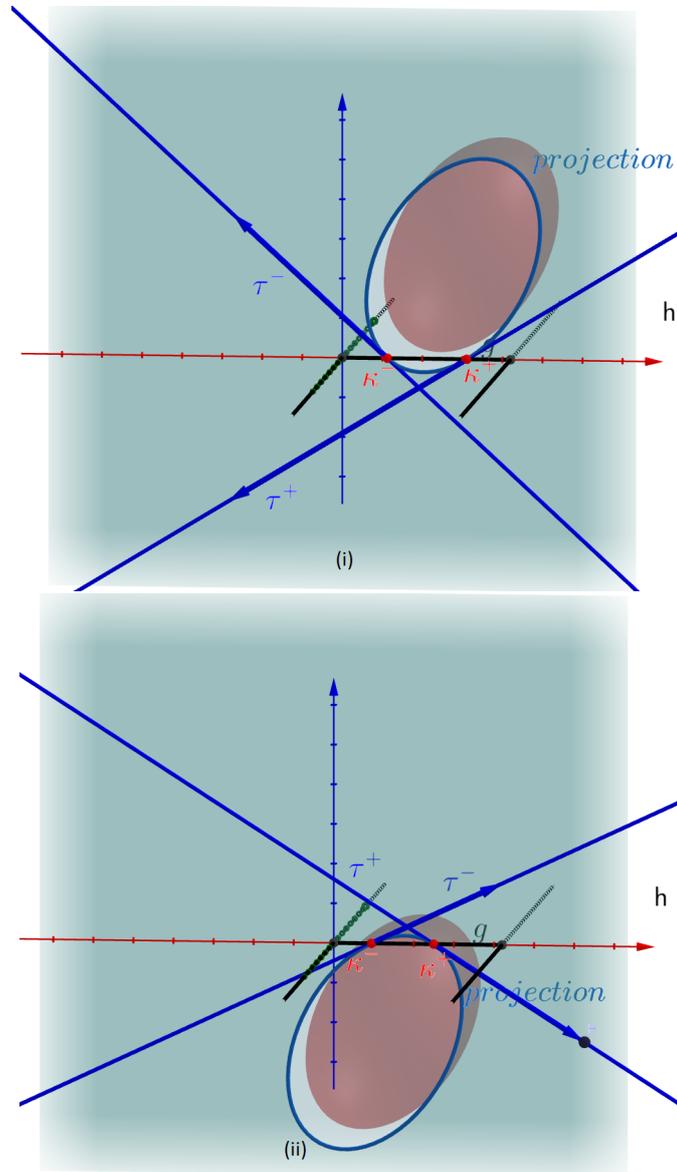


Figure 3.1: Moving the projection of K through g . Top: At the beginning of the crossing of g , the tangents $\tau^-(t)$ and $\tau^+(t)$ 'open up' (with respect to their sides that contain $K(t)$). Bottom: At the end of the crossing, they 'open down'.

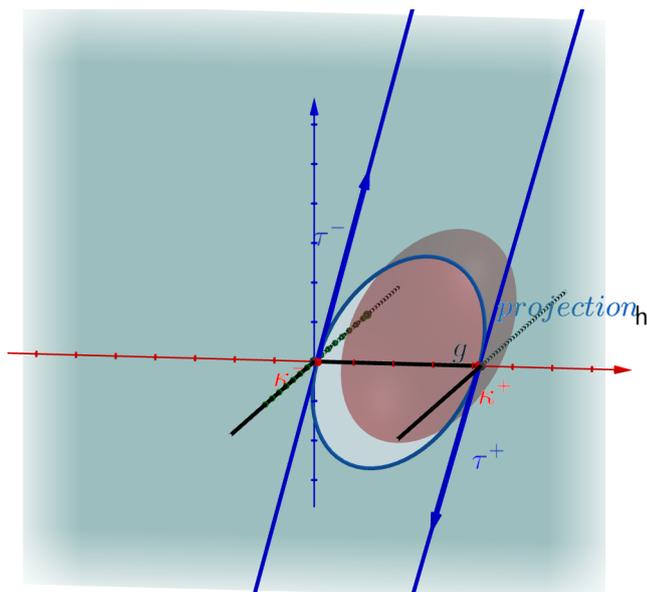


Figure 3.2: The critical instance t_0 where the tangents at $\kappa^-(t_0)$ and at $\kappa^+(t_0)$ become parallel.

To extend Theorem 3.2 to the case where K is a polytope, we use the following approximation scheme. Let D be some ball fully contained in K , with center c and radius ρ . For each $\delta > 0$, let L_δ be the Minkowski sum of K and a ball centered at the origin with radius δ , and define a map f_δ on \mathbb{S}^2 , so that, for each $\vec{v} \in \mathbb{S}^2$, $f_\delta(\vec{v}) = (1 - \delta)g(\vec{v}) + \delta\rho$, where $g(\vec{v})$ is the distance from c to ∂L_δ in direction \vec{v} . Define K_δ to be

$$K_\delta = \{c + tf_\delta(\vec{v})\vec{v} \mid \vec{v} \in \mathbb{S}^2, t \in [0, 1]\},$$

scaled down by a factor of $1 + \delta$. See Figure 3.3 for an illustration.

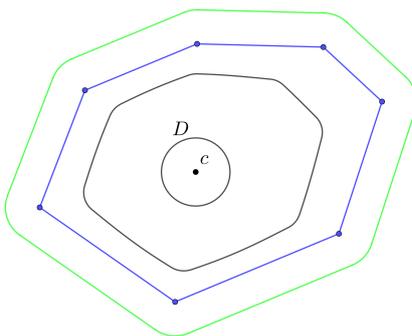


Figure 3.3: A two-dimensional illustration of the scheme for approximating a convex polytope by smooth strongly convex bodies: the convex polygon (blue), Minkowski sum with a small circle (green) and K_δ (black).

It is easily seen that K_δ is a smooth compact strongly convex object that is contained in K , and that $K_\delta \rightarrow K$ as $\delta \rightarrow 0$, in the sense that the Hausdorff distance between K and

K_δ tends to zero. Clearly, if K can be moved through W (by an arbitrary collision-free rigid motion), then so can K_δ .

For each $\delta > 0$, apply Theorem 3.2 to K_δ , to obtain a direction \vec{v}_δ and a rotation θ_δ orthogonal to \vec{v}_δ so that there is a sliding collision-free motion of K_δ in direction \vec{v}_δ from its view $(\vec{v}_\delta, \theta_\delta)$ through W . By compactness of \mathbb{S}^2 , there exists a sequence $\delta_i \downarrow 0$ such that \vec{v}_{δ_i} converges to some direction \vec{v} in \mathbb{S}^2 , and θ_{δ_i} converges to some rotation θ . By continuity, it follows that there exists a sliding collision-free motion of K through W in direction \vec{v} from its view (\vec{v}, θ) . That is, we have obtained the following result.

Theorem 3.3. *Let K be a convex polytope that can be moved by some collision-free rigid motion through a gate W . Then there exists a sliding collision-free motion of K through W .*

We can therefore apply the machinery of Theorem 2.5, and conclude that we can determine whether K can be moved through W by a collision-free motion in time $O(n^{8/3}\text{polylog}(n))$.

4

From Passing Through an Arbitrary Convex Window to Sliding Through a Gate

In this chapter we prove a similar yet different property of a convex polytope passing through some window. Here the window W is an arbitrary compact planar convex shape, not necessarily rectangular.

Theorem 4.1. *Let W be an arbitrary compact convex region in the xy -plane. Let K be a convex polytope that can be moved by some collision-free motion through W , and let d be the diameter of W (the maximum distance between any pair of points in W). Let h be an arbitrary plane, and let K_h be the projection of K on h . Then K_h can be rigidly placed between two parallel lines at distance d . That is, for any fixed direction \vec{v} , K can slide in direction \vec{v} through a gate of width d , in a plane perpendicular to \vec{v} .*

Before giving the proof, here is an interesting corollary of the theorem.

Corollary 4.2. *If K can be moved through a rectangular window W of dimensions $a \times b$ by some collision-free motion, then K can slide through a rectangle of dimensions $\min(a, b) \times \sqrt{a^2 + b^2}$.*

Proof. Assume without loss of generality that $a < b$. Since K can be moved by a collision-free motion through a rectangle of dimensions $a \times b$ it can be moved through a gate of width a , which we also assume to lie in the xy -plane. Therefore, by Theorem 3.3, K can slide through that gate, and Lemma 2.1 then implies that there exists a placement K_0 of K from which it can slide through the gate in the negative z -direction. Now project K_0 on the yz -plane and use Theorem 4.1 to place the projection between two parallel lines l_1, l_2 at distance $\sqrt{a^2 + b^2}$ (which is the diameter of W in this case), which, by rotating K_0 around the x -axis,

we can assume to be perpendicular to the xy -plane. All this implies that the xy -projection of (a rotated copy of) K_0 is contained in a rectangle of the asserted dimensions. Indeed, the projection of K_0 is contained in a y -vertical gate of width a , and this property remains true after we rotate K_0 around the x -axis. Therefore K can slide through such a rectangle in the negative z -direction (see Figure 4.1). \square

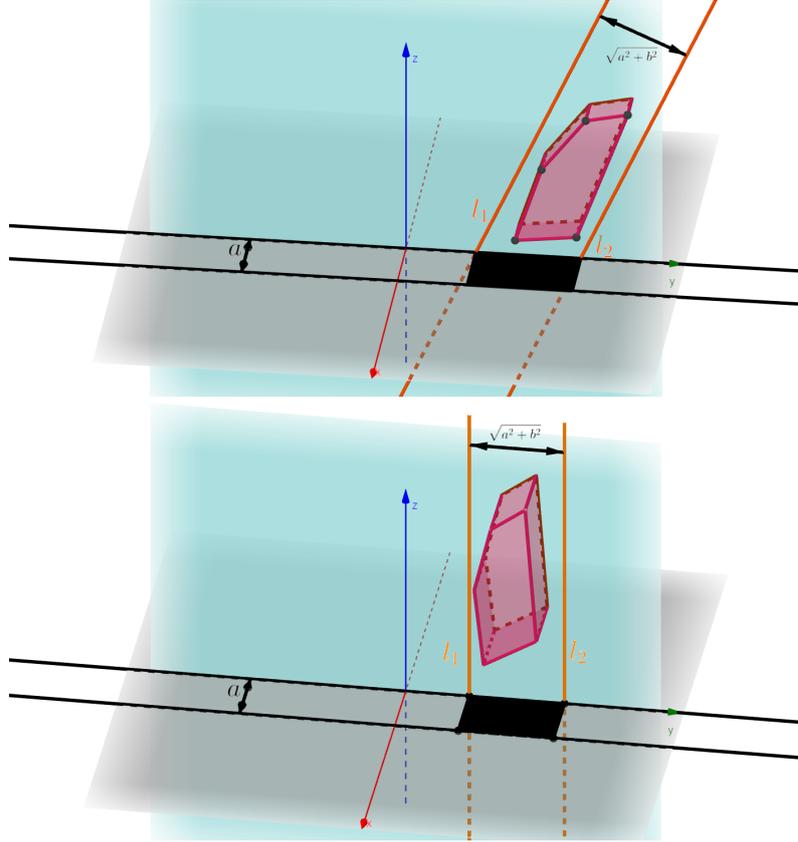


Figure 4.1: K can slide through a rectangle of dimensions $a \times \sqrt{a^2 + b^2}$. (top) The projection of K on the xy -plane is contained in a gate of width a (the black lines), and the projection of K on the yz -plane is contained between l_1, l_2 (orange). (bottom) After rotating K around the x -axis, l_1, l_2 are perpendicular to the xy -plane and the projection of K on the xy -plane is contained in such a rectangle.

We move now to prove the theorem.

Proof of Theorem 4.1. As before, we will first prove the theorem for smooth strongly convex compact bodies, and then extend the result to polytopes the same way as we did in the previous chapter. Consider the motion of K , now assumed to be smooth, strongly convex, and compact, in the normal setup, where W remains stationary in the xy -plane and K moves, during the time interval $[0, 1]$. We assume that at $t = 0$ (resp., at $t = 1$), K lies fully above (resp., below) the xy -plane.

Fix some direction \vec{v} , and let $C = C(\vec{v})$ denote the silhouette of K when viewed in direction \vec{v} . Let h be some plane orthogonal to \vec{v} , and let π_h denote the orthogonal projection

onto h . Parameterize a point $u \in C$ by the orientation θ of the tangent at $\pi_h(u)$ to $K_h := \pi_h(K)$ which is well defined since K is smooth, and let γ_h be the inverse of π_h ; that is, $\gamma_h(\theta)$ is the unique point $u \in C$ such that $\pi_h(u) = \theta$. Since K is assumed to be strongly convex, K_h is also strongly convex, and γ is a well-defined and continuous function on \mathbb{S}^1 . We extend γ to a bivariate function $\gamma^* : \mathbb{S}^1 \times [0, 1] \mapsto \mathbb{R}^3$, so that $\gamma^*(\theta, t)$ is the position (in the ambient 3-space) of $\gamma(\theta)$ at time t during the motion of K .

Let $\delta : \mathbb{S}^1 \times [0, 1] \mapsto \mathbb{R}$ be the function $\delta(\theta, t) = z(\gamma^*(\theta, t))$, namely, the z -coordinate of the corresponding point $\gamma(\theta)$ of C at time t . Note that at time $t = 0$ (resp., at time $t = 1$), δ is positive (resp., negative) at each θ , since K lies fully above (resp., below) the xy -plane at that time. Put $M := \max_{\theta \in \mathbb{S}^1} \delta(\theta, 0)$ and $m := \min_{\theta \in \mathbb{S}^1} \delta(\theta, 1)$. By our assumptions, $M > 0$ and $m < 0$.

The functions $\delta_0(\theta) = \delta(\theta, 0)$ and $\delta_1(\theta) = \delta(\theta, 1)$ are defined and continuous on \mathbb{S}^1 , and we extend each of them to the closed unit disk \mathcal{B}^1 bounded by \mathbb{S}^1 , in polar coordinates, which, for technical reasons, we write in reverse order as (θ, r) , by

$$\begin{aligned}\delta_0^*(\theta, r) &= r\delta_0(\theta) + (1-r)M \\ \delta_1^*(\theta, r) &= r\delta_1(\theta) + (1-r)m.\end{aligned}$$

It is easily checked that these extensions are well defined and continuous over \mathcal{B}^1 . Moreover, $\delta_0^*(\theta, r) > 0$ and $\delta_1^*(\theta, r) < 0$ for every θ .

We now take our function δ , which is so far defined on the side surface S of the cylinder $\mathbb{S}^1 \times [0, 1]$, and extend it to the entire boundary $S^* := S \cup B_0 \cup B_1$ of the cylinder, so that δ coincides with δ_0^* on the base B_0 of the cylinder at $t = 0$, and with δ_1^* on the base B_1 at $t = 1$. Clearly, the extended δ is well defined and continuous over S^* .

To simplify the forthcoming analysis, we identify S^* with the unit sphere \mathbb{S}^2 , which we parameterize by (θ, z) , where $\theta \in \mathbb{S}^1$ is the horizontal orientation of the point on \mathbb{S}^2 and z is its z -coordinate (so θ is not well defined at the north and south poles of \mathbb{S}^2). We use the simple homeomorphism f that maps a point $(\theta, t) \in S$ to $(\theta, t - 1/2) \in \mathbb{S}^2$, maps a point $(\theta, r) \in B_0$ to $(\theta, -1 + r/2) \in \mathbb{S}^2$, and maps a point $(\theta, r) \in B_1$ to $(\theta, 1 - r/2) \in \mathbb{S}^2$. See Figure 4.2 for an illustration. In what follows, we will mostly use \mathbb{S}^2 to represent S^* , except for a few technical observations.

Define a function G from \mathbb{S}^2 to \mathbb{R}^2 by

$$G(\theta, z) = (\delta(\theta, z), \delta(\theta + \pi, z)), \quad \text{for } (\theta, z) \in \mathbb{S}^2.$$

Our goal is to show that $G(\mathbb{S}^2)$ contains the origin. Note that, by construction, $G(f(B_0))$ is fully contained in the positive quadrant $Q_1 := \{(x, y) \mid x, y > 0\}$, and $G(f(B_1))$ is fully contained in the negative quadrant $Q_3 := \{(x, y) \mid x, y < 0\}$. Thus, if $G(\mathbb{S}^2)$ contains the origin then so does $G(f(S))$. Once this property is established, it provides us with a pair (θ, z) such that $\delta(\theta, z) = \delta(\theta + \pi, z) = 0$, which means that there are two antipodal points $u, v \in C$ that pass through W simultaneously. Therefore their distance must be at most the diameter of W , and hence also the distance between the parallel tangent planes through them, which is a slab parallel to \vec{v} of width at most d that contains K , as asserted.

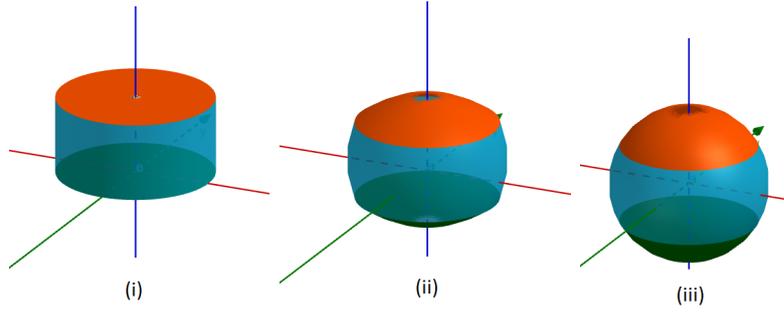


Figure 4.2: Identifying S^* with the unit sphere \mathbb{S}^2 . B_0 is shown in green, B_1 in orange, and S in light blue. In (i) S^* is depicted, in (ii) an intermediate snapshot of the deformation is shown, for visual convenience, and in (iii) the final unit ball is shown, divided into the three parts that correspond to B_0 , S and B_1 .

Assume to the contrary that $G(\mathbb{S}^2)$ does not contain the origin. Then we can normalize G to the function

$$H(\theta, z) := \frac{G(\theta, z)}{\|G(\theta, z)\|}, \quad \text{for } (\theta, z) \in \mathbb{S}^2,$$

which maps \mathbb{S}^2 continuously to the unit circle \mathbb{S}^1 . The function G , and thus also the function H , are symmetric with respect to the line $y = x$ in \mathbb{R}^2 , meaning that

$$\begin{aligned} G(\theta + \pi, z) &= \Sigma(G(\theta, z)), & \text{for } (\theta, z) \in \mathbb{S}^2, & \text{ and thus also} \\ H(\theta + \pi, z) &= \Sigma(H(\theta, z)), & \text{for } (\theta, z) \in \mathbb{S}^2, & \end{aligned}$$

where Σ is the reflection about $y = x$, that is, $\Sigma(x, y) = (y, x)$.

We now use the property that the real line is a *covering space* of \mathbb{S}^1 , in the specific (and easily verified) sense that the continuous map $p : \mathbb{R} \mapsto \mathbb{S}^1$, given by $p(x) = e^{2\pi ix}$, for $x \in \mathbb{R}$, is surjective, and, for each $\zeta \in \mathbb{S}^1$, there exists an open neighborhood U of ζ such that $p^{-1}(U)$ is the disjoint union of open sets in \mathbb{R} , each of which is mapped homeomorphically to U by p . The map p is called the *covering map*.

A well known property of covering spaces is the *lifting property* (reviewed, e.g., in [17]; see also [16]), a special case of which asserts, in the specific context used here, that, if φ is any continuous map from \mathbb{S}^2 to \mathbb{S}^1 then φ can be *lifted* to a map $\psi : \mathbb{S}^2 \mapsto \mathbb{R}$, so that $p \circ \psi = \varphi$. (Technically, this property holds when the domain of φ (and ψ), which is \mathbb{S}^2 in our case, is path connected, locally path connected, and simply connected, conditions that are trivially satisfied by \mathbb{S}^2 . Hence the lifting ψ does indeed exist.)

Applying the lifting property to the function H , we get a continuous mapping $T : \mathbb{S}^2 \mapsto \mathbb{R}$, such that $p \circ T = H$, so we have the property that

$$p(T(\theta + \pi, z)) = \Sigma(p(T(\theta, z))), \quad \text{for } (\theta, z) \in \mathbb{S}^2.$$

As is easily checked, we have $\Sigma(e^{iy}) = e^{i(\pi/2-y)}$, and therefore, for a point $x \in \mathbb{R}$, we have

$$\Sigma(p(x)) = \Sigma(e^{2\pi ix}) = e^{\pi i/2 - 2\pi ix} = p(1/4 - x), \quad \text{so}$$

$$p(T(\theta + \pi, z)) = p(1/4 - T(\theta, z)), \quad \text{for } (\theta, z) \in \mathbb{S}^2.$$

This in turn implies, by the definition of p , that

$$T(\theta + \pi, z) = 1/4 + k_{\theta, z} - T(\theta, z),$$

for some integer $k_{\theta, z}$. However, since T is continuous, there must be a single integer k such that $k_{\theta, z} \equiv k$ for all θ and z . That is, we have

$$T(\theta + \pi, z) + T(\theta, z) = 1/4 + k, \quad \text{for all } (\theta, z) \in \mathbb{S}^2. \quad (4.1)$$

By an easy application of the mean-value theorem (which is also a special case of the Borsuk-Ulam theorem in dimension 1), there exist θ_0 and θ_1 such that, recalling that the value $z = -1/2$ (resp., $z = 1/2$) corresponds to points on the lower (resp., upper) circle bounding S ,

$$\begin{aligned} T(\theta_0 + \pi, -1/2) &= T(\theta_0, -1/2) \\ T(\theta_1 + \pi, 1/2) &= T(\theta_1, 1/2). \end{aligned}$$

Substituting in (4.1), we get

$$T(\theta_0, -1/2) = T(\theta_1, 1/2) = 1/8 + k/2.$$

However, by construction, $H(\theta_0, -1/2)$ lies in the first quadrant Q_1 , and $H(\theta_1, 1/2)$ lies in the third quadrant Q_3 . Hence we have $T(\theta_0, -1/2) \in (0, 1/4) + \mathbb{Z}$ and $T(\theta_1, 1/2) \in (1/2, 3/4) + \mathbb{Z}$, but $1/8 + k/2$ can belong to only one of these sets (depending on whether k is even or odd). This contradiction shows that $G(\mathbb{S}^2)$, and thus also $G(f(S))$, contains the origin, as asserted.

So far the proof was for a smooth strongly convex compact bodies. The extension to the case of a convex polytope K is done exactly as in the proof of (that is, the argument preceding) Theorem 3.3. \square

A second proof. We provide an alternative proof of Theorem 4.1, and we are grateful to Boris Aronov for providing to us its main ingredients.

We use the same notations as in the previous proof. Similar to the first proof, the following, slightly more generally stated proposition is the main technical tool that we need.

Proposition 4.3. *Let $G: S \rightarrow \mathbb{R}^2$ be a continuous map, interpreted as the homotopy of the closed curve $\delta_0: \mathbb{S}^1 \rightarrow Q_1 \subset \mathbb{R}^2$, given by $\theta \mapsto G(\theta, 0)$, to the closed curve $\delta_1: \mathbb{S}^1 \rightarrow Q_3 \subset \mathbb{R}^2$, given by $\theta \mapsto G(\theta, 1)$. In addition, suppose that G is symmetric, in the sense that $G(\theta + \pi, t) = \Sigma(G(\theta, t))$, for all $\theta \in \mathbb{S}^1$ and $t \in [0, 1]$. Then there exist $\theta \in \mathbb{S}^1$, $t \in \mathbb{R}$ that satisfy $G(\theta, t) = O$, that is, G cannot miss the origin.*

Proof. Clearly, if $G(\theta, t) = O$ then we also have $G(\theta + \pi, t) = O$. Hence it suffices to show that there exists (θ, t) in $D := [0, \pi] \times [0, 1]$ (half the side surface of the cylinder) such that $G(\theta, t) = O$. Let Π be the image of D under G .

Consider the curve $\gamma_0: [0, 1] \rightarrow S$ defined by $t \mapsto (0, t)$, and its image Γ_0 under G , i.e., $\Gamma_0(t) = G(0, t) \in \mathbb{R}^2$. Let γ_1 and Γ_1 be defined similarly by $\gamma_1(t) := (\pi, t)$, so that $\Gamma_1(t) = G(\pi, t)$. Let γ'_1 and Γ'_1 be the reverses of γ_1 and Γ_1 , respectively — the same curves traversed in reverse direction.

Additionally, let $\zeta_0, \zeta_1: [0, \pi] \rightarrow S$ be the “half-circles” defined by $\theta \mapsto (\theta, 0)$ and $\theta \mapsto (\theta, 1)$, respectively, and $Z_i := G \circ \zeta_i$, for $i = 0, 1$. Let ζ'_0 and Z'_0 be the reverses of ζ_0 and Z_0 , respectively. See Figure 4.3 for an illustration.

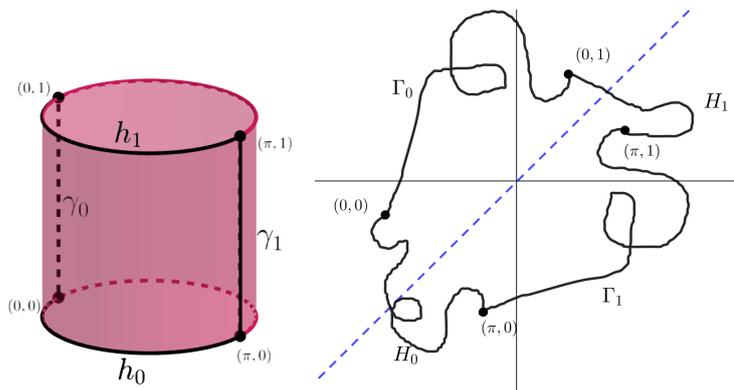


Figure 4.3: The boundary of half the cylinder is mapped to a closed loop with a nonzero winding number around the origin. Note that γ_0 and γ_1 are symmetric about the axis of the cylinder, and therefore $\Gamma_0 = \Sigma(\Gamma_1)$.

Concatenating γ_0 , ζ_1 , γ'_1 , and ζ'_0 in this order, we obtain a closed loop ℓ in S , which is the boundary of D , and its corresponding image $L := G(\ell)$ in the plane. By construction, ℓ bounds the topological disk $[0, \pi] \times [0, 1]$ in S , and L is a closed loop in \mathbb{R}^2 . We prove below that, if $O \notin L$, then L has a non-zero winding number around O . Given this property, we claim that O must lie in $\Pi = G([0, \pi] \times [0, 1])$. Indeed, if $O \notin \Pi$ then Π is contained in the punctured plane at the origin. Since L has a nonzero winding number around O , it is homotopic, within Π , to a curve obtained by looping around the origin a nonzero number of times. This curve is not homotopy-trivial—it is not homotopic to a point (within Π). On the other hand, ℓ is clearly homotopy-trivial within D , and therefore $L = G(\ell)$ is homotopic to a single point within $G(D) = \Pi$, a contradiction that establishes the proposition. \square

To complete the proof, we thus show:

Claim 4.4. *In the notation of the above proof, if L misses O , then the winding number of L around O is non-zero.*

Proof. Let $\arg(x, y)$ be the clockwise angle that the vector (x, y) makes with the positive x -axis and let, for a section λ of L , $\Delta\lambda$ be the integral of the change in $\arg \lambda(t)$ as t traces out λ from start to finish.

We will compute the winding number of L around the origin by breaking L into sections λ , computing the angle change $\Delta\lambda$ for each section, and adding up the numbers.

Let $\alpha := \arg G(0, 0) \in (0, \pi/2)$. Then by Σ -symmetry $\arg G(\pi, 0) = \pi/2 - \alpha$. Similarly, put $\beta := \arg G(0, 1) \in (\pi, 3\pi/2)$, so that $\arg G(\pi, 1) = 5\pi/2 - \beta \in (\pi, 3\pi/2)$. Since $Z_0 \subset Q_1$ (so Z_0 cannot wind around O), $\Delta Z_0 = \arg G(\pi, 0) - \arg G(0, 0) = (\pi/2 - \alpha) - \alpha = \pi/2 - 2\alpha$ and $\Delta Z'_0 = -\Delta Z_0 = 2\alpha - \pi/2$. Similarly, since $Z_1 \subset Q_3$, $\Delta Z_1 = \arg G(\pi, 1) - \arg G(0, 1) = (5\pi/2 - \beta) - \beta = 5\pi/2 - 2\beta$.

Γ_0 connects $G(0, 0)$ to $G(0, 1)$, so $\Delta\Gamma_0 = \arg G(0, 1) - \arg G(0, 0) + 2\pi k = \beta - \alpha + 2\pi k$, for some integer k , over which we have no control as we do not know how many times Γ_0 winds around the origin (we use here the assumption that Γ_0 avoids the origin). Because of Σ -symmetry, we must have $\Delta\Gamma_1 = -\Delta\Gamma_0$ and therefore $\Delta\Gamma'_1 = -\Delta\Gamma_1 = \Delta\Gamma_0$.

To summarize, the total change of the angle around L is equal to

$$\begin{aligned} \Delta\Gamma_0 + \Delta Z_1 + \Delta\Gamma'_1 + \Delta Z'_0 &= 2\Delta\Gamma_0 + \Delta Z_1 + \Delta Z'_0 \\ &= 2(\beta - \alpha + 2\pi k) + (5\pi/2 - 2\beta) + (2\alpha - \pi/2) \\ &= 2\pi(2k + 1). \end{aligned}$$

In particular, the total angle is not zero, no matter what the value of the integer k is, thereby completing the proof.

The remainder of the argument, namely that Proposition 4.3 implies the theorem, and the extension to the case of convex polytopes, is done exactly as in the first proof, thereby completing this second proof of the theorem. \square

5

Purely Translational Motions

In this chapter we solve another special case of the problem, in which arbitrary translations are allowed, with all three degrees of freedom, but rotations are not allowed. This case is solved by an easy combination of what we have proved so far. The following theorem is, in a sense, a strengthening of Lemma 2.1.

Theorem 5.1. *If K can be moved through a rectangular window W by some purely translational collision-free motion, then K can be moved through W , possibly from some other (translated and rotated) starting position, by sliding in the z -direction.*

Proof. Again, we first carry out the proof for the case where K is a smooth compact strongly convex body in three dimensions, and then extend the proof to the case where K is a convex polytope. We recall the proof of Theorem 3.2 and follow the notations used there. When K translates through W , its projection on the xz -plane is a fixed convex region that translates through the interval g on the x -axis, which is the x -projection of W . By the argument in the proof of Theorem 3.2, there is a time t during the motion at which the tangents τ^- and τ^+ become parallel, and form, when extended in the y -direction, a (possibly slanted) slab S that is orthogonal to the xz -plane, and that contains the placement of K at time t , so that the intersection of S with the xy -plane is a y -vertical strip of width at most a , whose x -projection is contained in that of W . Applying the same argument to the yz -plane (swapping the x - and y -directions), we get another time t' at which K is contained in another slab S' , orthogonal to the yz -plane, whose intersection with the xy -plane is an x -horizontal strip of width at most b , whose y -projection is contained in that of W (see Figure 5.1).

Hence, the intersection $\tau = S \cap S'$ is a (slanted) prism, whose cross-section with the xy -plane is a rectangle contained in W . Moreover, as is easily verified, τ contains some translated copy K_0 of K . Hence, K can slide through W from its placement K_0 in the unbounded direction of τ . By Lemma 2.1, K can also slide through W in the z -direction,

from a different, possibly rotated, initial placement.

The case where K is a convex polytope can be handled by the same limiting argument given in the proof of Theorem 3.3. \square

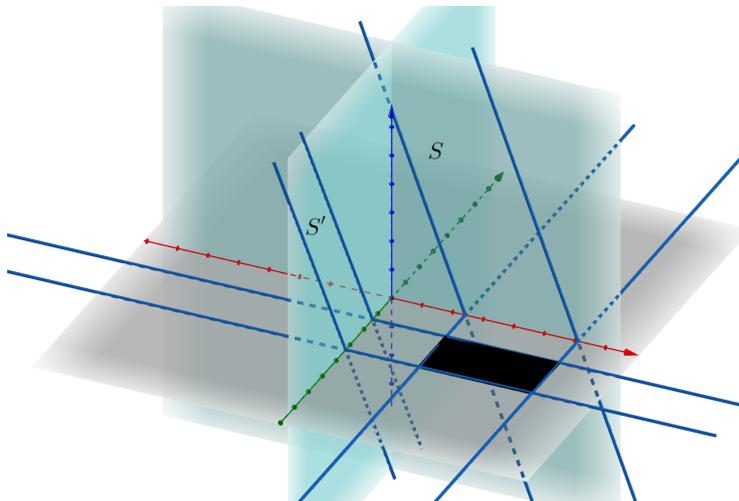


Figure 5.1: Some translated copy of K is contained within a slab S perpendicular to the xz -plane, and some other translated copy of K is contained in a slab S' perpendicular to the yz -plane, so that the x -projection of the intersection of S with the xy -plane is contained in that of W , and the y -projection of the intersection of S and the xy -plane is contained in that of W . We can thus translate K to a placement contained in the prism $S \cap S'$, from which K can slide through the intersection of $S \cap S'$ and the xy -plane (the black rectangle, which is contained in W).

This leads to an efficient algorithm, with running time $O(n^{8/3} \text{polylog}(n))$, for the problem considered in this chapter, following the algorithm given in Chapter 2.

6

Rotations Are Needed

So far we have considered sliding motions of a convex polytope through a window, and showed that in several cases it suffices to consider only such motions. However, this is not the case in general. We show in this and the following chapter that in general rotations are needed to obtain a collision-free motion of the polytope through the window.

Lemma 6.1. *let W be a squared window with side length $\sqrt{5}$. Let $A = (0, 0, 0)$, $B = (1, 3, 0)$, $C = (1, 0, h)$, $D = (0, 3, h)$ be four points, where $h \gg 1$ is a sufficiently large parameter. Let K be the tetrahedron $ABCD$ (see Figure 6.1). Then:*

1. *K cannot pass through W by any purely translational collision-free motion (for sufficiently large $h \gg 1$).*
2. *K can pass through W by a collision-free motion with only two degrees of freedom: translating in the z -direction combined with horizontal rotation (for any value of $h > 0$).*

Proof.

1. Assume to the contrary that there exists a purely translational motion of K through W . By Theorem 5.1, there exists some placement K_0 of K from which K can slide through W in the negative z -direction. Let $\pi(K_0)$ denote the vertical projection of K_0 onto the xy -plane. By the theorem, $\pi(K_0)$ can be rigidly placed inside W . Recall now that h is very large, which implies that, when transforming K to K_0 , the z -vertical direction turns by only a very small angle, for otherwise $\pi(K_0)$ would be very long and would not fit into such a square. More formally, for every $\varepsilon > 0$ there exists h_0 such that for every $h > h_0$ the angle by which the z -axis turns from K to K_0 is at most ε . As ε decreases to zero, the lengths of the projections of the segments AB, CD grow to $\sqrt{10}$, which is their original length, and the angle between them converges to some $0 < \phi < \frac{\pi}{2}$ (the exact angle is the angle obtained

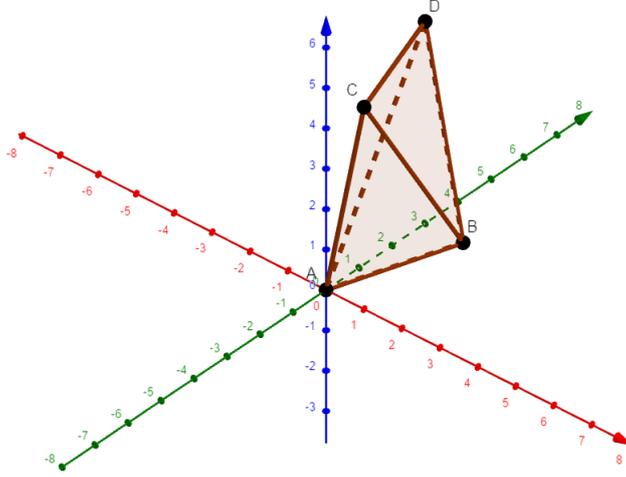


Figure 6.1: The tetrahedron $K = ABCD$.

when the z -axis remains the same, which is then $\phi = 2 \sin^{-1} \frac{1}{\sqrt{10}}$). Therefore, the projection $\pi(K_0)$ is the convex hull of two segments of length sufficiently close to $\sqrt{10}$, which is the diagonal of W , where the angle between them is sufficiently far from $0, \pi/2$. Hence $\pi(K_0)$ cannot be placed inside a square with side length $\sqrt{5}$. This contradiction establishes the first part of the theorem.

2. We move W instead of K , allowing it only to translate in the z -direction (so it always remains horizontal), and simultaneously rotate around its center (so the motion of W has only two degrees of freedom). More concretely, the center of W moves up along the line $x = 1/2, y = 3/2$. We parameterize the motion by a parameter $c \in [0, 1]$, so that at time c , W lies on the plane $z = ch$ and its center is at $(1/2, 3/2, ch)$. See Figure 6.2(left) for a schematic top view of K .

The cross section K_c of K at time c is shown (in green) in Figure 6.2(right). It is a quadrilateral $PQRS$, with $P = (c, 0)$, $Q = (1, 3(1 - c))$, $R = (1 - c, 3)$ and $S = (0, 3c)$. We place W around K_c so that PR lies at the middle of one diagonal of W (so W keeps rotating to align with this rotating segment). It is clear that the motion of W is continuous, and it remains to show that K_c always lies in (the placement at height ch , with the aligned diagonals, of) W .

It suffices to show that, at any time c during the motion, ΔPRS is contained in the isosceles right triangle with hypotenuse PR (this triangle is half of W , and the argument for the complementary half and for ΔPQR is fully symmetric). For this, it suffices to show that

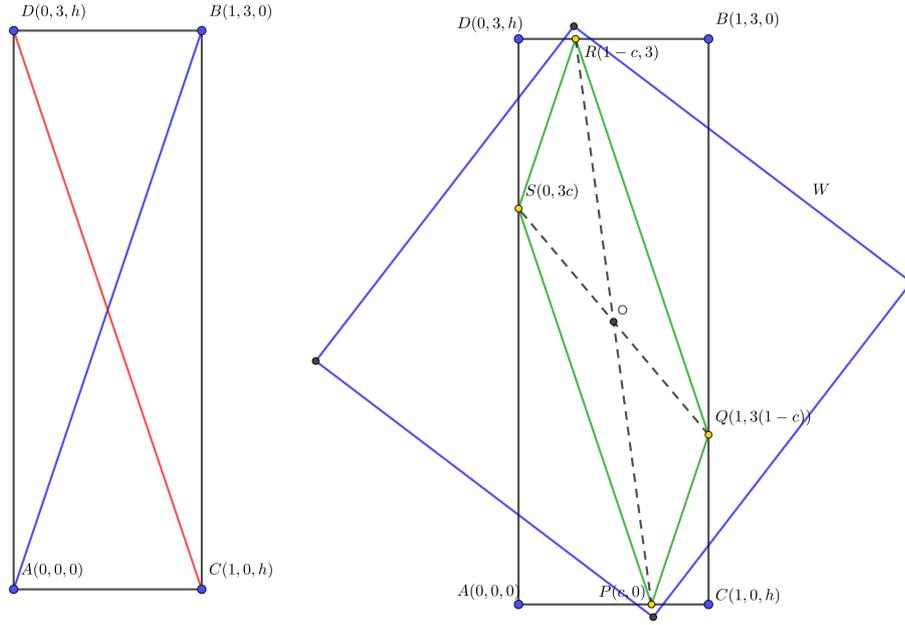


Figure 6.2: Placing the cross section of K inside W . Left: A schematic top view of K . Right: The cross section of K at time c (green), and a copy of W that contains it.

each of the angles $\sphericalangle SPR$, $\sphericalangle SRP$ is smaller than $\pi/4$. Note that the edges of $PQRS$ have fixed slopes, namely 3 and -3 , as they are parallel to the xy -projections of AB and CD . This implies that $\tan \sphericalangle SPQ = \tan \sphericalangle SRQ = \frac{3}{4} < 1$, so $\sphericalangle SPQ = \sphericalangle SRQ < \pi/4$. We have thus shown that K can move through W by (the dual version of) this motion, of translation in the z -direction combined with horizontal rotation. \square

7

The Case of a Circular Window

In this chapter we study the case where W is a circular window, of some diameter d . There are (at least) three possible types of motion of K through W : sliding, purely translational motion, and general motion, with all six degrees of freedom. In this chapter we show that these types are not equivalent. Our two main results are: (a) There are situations in which K can pass through W by a purely-translational collision-free motion but K cannot slide through W . (b) There are situations in which K can pass through W by a general collision-free motion but K cannot pass through W by a purely-translational motion.

We prove both results for the case where K is a regular tetrahedron of side length 1. We do so by showing the existence of two threshold parameters $\delta_2 < \delta_1 < 1$, so that, denoting by $\text{diam}(W)$ the diameter of W , (i) K can slide through W if $\text{diam}(W) \geq 1$, (ii) K cannot slide through W , but can pass through W by a purely translational motion, if $\delta_1 \leq \text{diam}(W) < 1$, (iii) K cannot pass through W by a purely translational motion, but can pass through W by a general motion, if $\delta_2 \leq \text{diam}(W) < \delta_1$, and (iv) K cannot pass through W at all if $\text{diam}(W) < \delta_2$.

Approximate values for these thresholds, obtained numerically, are $\delta_1 \approx 0.901388$ and $\delta_2 \approx 0.895611$.

7.1 Purely Translational Motion

Let K be a regular tetrahedron of side length 1. Assume first that $\text{diam}(W) \geq 1$. Then it is easy to show that K can slide through W . This is because K can be enclosed in a cylinder of diameter 1, such as a cylinder (of diameter 1) whose axis is orthogonal to two opposite edges of K . The interesting case is therefore when $\text{diam}(W) < 1$. We then have:

Theorem 7.1. *Let W be a circular window of diameter d .*

- (1) *K cannot slide through W , at any fixed orientation, if $d < 1$.*
- (2) *There exists a threshold $\delta_1 \approx 0.901388$, such that K can pass through W by a collision-free purely translational motion, if $d \geq \delta_1$.*

Proof.

(1) First, by the argument in the proof of Lemma 2.1, the projection of W on any plane can be rigidly placed inside W (for a circular window, the argument is actually simpler, since the projection is an ellipse whose major axis equals to the diameter of W), so, arguing as in the proof of Lemma 2.1, any polytope that can slide through W can also slide in the direction of the z -axis. It therefore suffices to show that K cannot be contained in a cylinder of diameter smaller than 1.

A proof of this fact can be found in [18]. For the sake of completeness, we reproduce here the proof. We first note that for any four vectors $\vec{v}_1, \vec{v}_2, \vec{v}_3, \vec{v}_4$ the following identity holds:

$$\frac{1}{2} \sum_{i=1}^4 \sum_{j=1}^4 |\vec{v}_i - \vec{v}_j|^2 = 3 \sum_{i=1}^4 |\vec{v}_i|^2 - 2 \sum_{1 \leq i < j \leq 4} \langle \vec{v}_i, \vec{v}_j \rangle = 4 \sum_{i=1}^4 |\vec{v}_i|^2 - \left| \sum_{i=1}^4 \vec{v}_i \right|^2.$$

Let K be the tetrahedron whose vertices are:

$$\begin{aligned} \vec{v}_1 &= \frac{1}{\sqrt{24}}(0, 0, 3), & \vec{v}_2 &= \frac{1}{\sqrt{24}}(\sqrt{8}, 0, -1), \\ \vec{v}_3 &= \frac{1}{\sqrt{24}}(-\sqrt{2}, \sqrt{6}, -1), & \vec{v}_4 &= \frac{1}{\sqrt{24}}(-\sqrt{2}, -\sqrt{6}, -1). \end{aligned}$$

It is indeed a regular tetrahedron of side length 1:

$$\begin{aligned} |\vec{v}_1 - \vec{v}_2|^2 &= \frac{1}{24}(8 + 16) = 1, & |\vec{v}_1 - \vec{v}_3|^2 &= \frac{1}{24}(2 + 6 + 16) = 1, \\ |\vec{v}_1 - \vec{v}_4|^2 &= \frac{1}{24}(2 + 6 + 16) = 1, & |\vec{v}_2 - \vec{v}_3|^2 &= \frac{1}{24}(18 + 6) = 1, \\ |\vec{v}_2 - \vec{v}_4|^2 &= \frac{1}{24}(18 + 6) = 1, & |\vec{v}_3 - \vec{v}_4|^2 &= \frac{1}{24}(24) = 1. \end{aligned}$$

Represent vectors in our 3-dimensional space as 3×1 column vectors. By some more algebra, we obtain

$$\begin{aligned} \sum_{i=1}^4 \vec{v}_i \vec{v}_i^T &= \\ \frac{1}{24} &\left[\begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 9 \end{pmatrix} + \begin{pmatrix} 8 & 0 & -\sqrt{8} \\ 0 & 0 & 0 \\ -\sqrt{8} & 0 & 1 \end{pmatrix} \right] + \\ &+ \frac{1}{24} \left[\begin{pmatrix} 2 & -\sqrt{12} & \sqrt{2} \\ -\sqrt{12} & 6 & -\sqrt{6} \\ \sqrt{2} & -\sqrt{6} & 1 \end{pmatrix} + \begin{pmatrix} 2 & \sqrt{12} & \sqrt{2} \\ \sqrt{12} & 6 & \sqrt{6} \\ \sqrt{2} & \sqrt{6} & 1 \end{pmatrix} \right] = \end{aligned}$$

$$\frac{1}{24} \begin{pmatrix} 12 & 0 & 0 \\ 0 & 12 & 0 \\ 0 & 0 & 12 \end{pmatrix} = \frac{1}{2} I_3.$$

Therefore, for any unit vector \vec{n} the following equation is satisfied:

$$\sum_{i=1}^4 \langle \vec{v}_i, \vec{n} \rangle^2 = \sum_{i=1}^4 |\vec{v}_i^T \vec{n}|^2 = \sum_{i=1}^4 \text{Tr}((\vec{v}_i \vec{v}_i^T) \cdot (\vec{n} \vec{n}^T)) = \frac{1}{2} \text{Tr}(\vec{n} \vec{n}^T) = \frac{1}{2}.$$

Note that $\sum_{i=1}^4 v_i = 0$, and hence:

$$\sum_{1 \leq i < j \leq 4} \langle \vec{n}, \vec{v}_i - \vec{v}_j \rangle^2 = \frac{1}{2} \sum_{i=1}^4 \sum_{j=1}^4 \langle \vec{n}, \vec{v}_i - \vec{v}_j \rangle^2 = 4 \sum_{i=1}^4 \langle \vec{n}, \vec{v}_i \rangle^2 = 2.$$

Assume that the smallest cylinder that contains K has diameter d . Let h be a plane perpendicular to the axis of the cylinder, let \vec{n}_1, \vec{n}_2 be two orthogonal unit vectors in h , let \vec{u}_i be the projection of \vec{v}_i on h , for $1 \leq i \leq 4$, and put $\vec{l}_{ij} = \vec{u}_i - \vec{u}_j$. It is easy to see that

$$|\vec{l}_{ij}|^2 = |\vec{u}_i - \vec{u}_j|^2 = \langle \vec{n}_1, \vec{u}_i - \vec{u}_j \rangle^2 + \langle \vec{n}_2, \vec{u}_i - \vec{u}_j \rangle^2 = \langle \vec{n}_1, \vec{v}_i - \vec{v}_j \rangle^2 + \langle \vec{n}_2, \vec{v}_i - \vec{v}_j \rangle^2.$$

We thus have $\sum_{1 \leq i < j \leq 4} |\vec{l}_{ij}|^2 = 4$. Consider the coordinate system in h whose axes are parallel to \vec{n}_1 and \vec{n}_2 , and whose origin is at the center of the intersection circle of h and the cylinder. In this coordinate system we have $|\vec{u}_i| \leq \frac{d}{2}$ for each i . Note that \vec{l}_{ij} remains the same and that $\sum_{i=1}^4 \vec{u}_i = 0$, as the projection of $\sum_{i=1}^4 \vec{v}_i = 0$, and we thus obtain:

$$\sum_{1 \leq i < j \leq 4} |\vec{l}_{ij}|^2 = \sum_{1 \leq i < j \leq 4} |\vec{u}_i - \vec{u}_j|^2 = 4 \sum_{i=1}^4 |\vec{u}_i|^2 \leq 16 \left(\frac{d}{2}\right)^2 = 4d^2.$$

Finally we get that $4 = \sum_{1 \leq i < j \leq 4} |\vec{l}_{ij}|^2 \leq 4d^2$, so $d \geq 1$, but in our case the diameter of W is strictly smaller than 1. We therefore conclude that K cannot slide through W , thus establishing Part (1).

(2) We now prove that although K cannot slide through W , it can pass through W by a purely translational movement, provided that $d = \text{diam}(W)$ is at least some threshold $\delta_1 < 1$, whose concrete value will be analyzed below. This holds for many orientations of K (but not for all orientations); this set of admissible orientations keeps shrinking as d approaches δ_1 .

Assume for now that the orientation of K is fixed. We claim that K can move through W at this fixed orientation, as above, if and only if every horizontal cross section of K can be enclosed in a disc of diameter d ; that is, the smallest enclosing disc of each cross section has diameter at most d . We refer to this property as the *small diameter property*. The ‘only if’ part is obvious. We briefly explain the ‘if’ part. Let $K(z)$ be the cross section of K at

height h . For every $x \in \partial K(z)$ let c_x be a horizontal circle of diameter d centered at x . That is, all the points within the plane of the cross section whose distance from x is at most $\frac{d}{2}$. Clearly, the intersection $R(z) = \bigcap_{x \in \partial K(z)} c_x$ denotes the set of all available positions for the center of W within that plane, such that it contains the cross section $K(z)$. $K(z)$ is a continuous function of z in the Hausdorff metric of sets, and hence so is $R(z)$. This is easily seen to imply that we can choose the position of the center of W for every cross section in a way that is continuous in z .

Assume without loss of generality that the initial placement of K is with its lowest vertex at $z = 0$, and let h denote the z -coordinate of the highest vertex. As above, denote by $K(z)$ the cross section of K at height z , for $z \in [0, h]$. Assume without loss of generality that all four vertices have distinct z -coordinates, and that the order of increasing z -coordinates of the vertices is A, B, C, D ; that is, $z_A < z_B < z_C < z_D$.

We claim that the small diameter property holds if and only if it holds for $K(z_B)$ and $K(z_C)$. Indeed, observing that these two cross sections are triangles, assume without loss of generality that the radius ρ of the smallest enclosing disc D_B of $K(z_B)$ is larger than or equal to that of $K(z_C)$. Enclose $K(z_C)$ by some disc D_C of radius ρ , and let E be the convex hull of $D_B \cup D_C$, which is a possibly slanted elliptic cylinder, each of whose horizontal cross sections is a congruent copy of the disc D_B . Since K has no vertices in the open slab $z_B < z < z_C$, it follows that the portion of K within the closed slab $z_B \leq z \leq z_C$ is the convex hull of $K(z_B) \cup K(z_C)$, and is therefore fully contained in E . Hence, for every $z_B < z < z_C$, $K(z)$ is contained in a disc of radius ρ . The cases of the slabs $z_A < z < z_B$ and $z_C < z < z_D$ are argued in the same manner. This establishes our claim.

In other words, we want to find orientations of K for which the (triangular) horizontal cross sections at the two middle vertices of K (in the z -direction) have smallest enclosing discs of diameters smaller than 1.

Denote the cross section $K(z_B)$ through B by BUV , where U is the point $AC \cap K(z_B)$ and V is the point $AD \cap K(z_B)$. Put $x = |AU|$ and $y = |AV|$, so $0 \leq x, y \leq 1$. Similarly, we write the triangular cross section $K(z_C)$ through C as CST , where S is the point $AD \cap K(z_C)$ and T is the point $BD \cap K(z_C)$, and put $z = |SD|$ and $w = |TD|$, so again $0 \leq z, w \leq 1$. See Figure 7.1 for an illustration. Note that we must have $x > y$ and $w > z$, for otherwise A and D would not have been the two z -extreme vertices of K .

The requirement that these two cross sections be parallel imposes the following relations between x, y, z , and w .

$$\begin{aligned} z &= \frac{x - y}{x} \\ w &= \frac{x - y}{x(1 - y)}. \end{aligned} \tag{7.1}$$

Indeed, since the two cross sections are parallel, they intersect any plane (not parallel to them) at parallel lines. In particular, we have $UV \parallel CS$ and $TS \parallel BV$, so the triangles AUV and ACS are similar, and so are the triangles DST and DVB . The first similarity

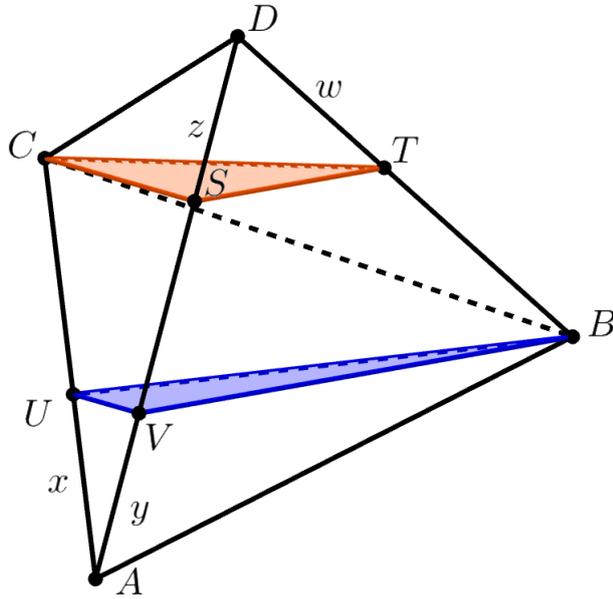


Figure 7.1: The horizontal cross sections of a regular tetrahedron through its two middle vertices.

implies that

$$x = \frac{AU}{AC} = \frac{AV}{AS} = \frac{y}{AS},$$

so $AS = y/x$, and then

$$z = AD - AS = 1 - AS = \frac{x - y}{x}.$$

The second similarity implies that

$$w = \frac{DT}{DB} = \frac{DS}{DV} = \frac{z}{1 - y} = \frac{x - y}{x(1 - y)},$$

thus establishing (7.1).

Note that, once we enforce $x > y$, the second inequality $w > z$ trivially holds.

The goal is then to search for orientations of K and for suitable choices of x and y (and thus of z and w too) for which the two cross sections have smallest enclosing discs of diameters smaller than 1. This is done as follows.

For a triangle Δ of side lengths a, b, c , the circumradius $r(\Delta)$ of Δ is given by the formula

$$r(\Delta) = \frac{abc}{4 \cdot \text{Area}(\Delta)}.$$

The area can be expressed by Heron's formula as

$$\text{Area}(\Delta)^2 = \pi(\pi - a)(\pi - b)(\pi - c),$$

where $\pi = (a + b + c)/2$ is half the perimeter. That is, we have

$$\begin{aligned}\text{Area}(\Delta)^2 &= \frac{1}{16}(a + b + c)(b + c - a)(a + c - b)(a + b - c) \\ &= \frac{1}{16}((a + b)^2 - c^2)(c^2 - (a - b)^2) \\ &= \frac{1}{16}(2a^2b^2 + 2a^2c^2 + 2b^2c^2 - a^4 - b^4 - c^4).\end{aligned}$$

Therefore,

$$r^2(\Delta) = \frac{a^2b^2c^2}{2a^2b^2 + 2a^2c^2 + 2b^2c^2 - a^4 - b^4 - c^4}. \quad (7.2)$$

Assume that the triangles BUV and CST are both acute, so their smallest enclosing discs coincide with their circumscribing discs. Apply this formula to each of the triangles BUV and CST . An easy application of the Law of Cosines yields

$$\begin{aligned}|BU|^2 &= 1 - x + x^2 \\ |BV|^2 &= 1 - y + y^2 \\ |UV|^2 &= x^2 - xy + y^2 \\ |CS|^2 &= 1 - z + z^2 \\ |CT|^2 &= 1 - w + w^2 \\ |ST|^2 &= z^2 - zw + w^2.\end{aligned}$$

Substituting these values in (7.2), once with $a^2 = |BU|^2$, $b^2 = |BV|^2$, $c^2 = |UV|^2$, and once with $a^2 = |CS|^2$, $b^2 = |CT|^2$, $c^2 = |ST|^2$, we get the values of the circumradii of the two triangles. If any of these triangles is obtuse, the radius of its smallest enclosing disc is half the longest edge.

The goal is, as said above, to find values of the parameters x, y that minimize the larger of these two radii (note that the choice of x and y determines the orientation of K , up to rotation about the z -axis because they determine a slice of K (namely, BUV) that has to be horizontal). By numerically testing a dense grid of values for x, y and running methods for finding the minimum of a function (computing the radius of the smallest enclosing disc using (7.2) for acute triangles, and half the longest edge for obtuse triangles), the optimizing parameters turned out to be $x \approx 0.43400$ and $y \approx 0.30265$, and the larger of the two diameters was ≈ 0.901388 . Setting δ_1 to this value completes the proof of Part (2). \square

7.2 General Motion

Theorem 7.2. *Let W be a circular window of diameter d . Then there exists a threshold $\delta_2 \approx 0.895611$, such that K can pass through W by a collision-free motion if and only if $d \geq \delta_2$.*

In other words, for diameters $\delta_2 \leq d < \delta_1$, the only way to move K through W is via a motion that also involves rotations, and for diameters $d < \delta_2$, no motion of K through W is possible.

Proof. We first construct the desired motion for $d \geq \delta_2$, which consists of five steps—sliding, rotation, sliding, rotation, and a final sliding. We use the setup and notations introduced in the analysis of the preceding section, and depicted in Figure 7.1. As earlier, it is more convenient to consider K as fixed, and W as moving around K .

Assume that the lowest vertex A lies on the xy -plane and inside W (see Figure 7.2 (i)). Start by sliding W up, possibly in a slanted direction, ensuring that it keeps containing the cross section of K with the plane supporting W , until W comes to contain B : See Figure 7.2 (ii). We want to choose the initial orientation of K so that the smallest enclosing disc of the horizontal (triangular) cross section of K through B , namely the triangle BUV , is of diameter at most d . As already noted, the orientation of K is determined by x and y , up to a possible rotation around the z -axis, as they determine the vertical direction of K (the one orthogonal to the triangle BUV).

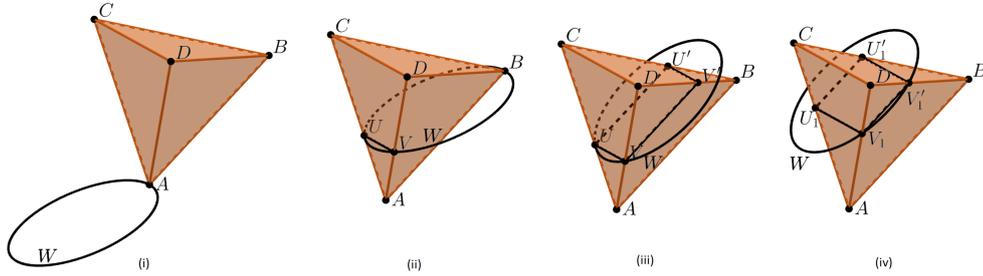


Figure 7.2: Moving W around K . (i) The initial configuration. (ii) W contains the triangle BUV . (iii) W contains the rectangle $UVV'U'$. (iv) W contains the symmetric rectangle, with edge lengths swapped. The remainder of the motion is a fully symmetric reversal of the first two steps.

We ran our numerical approximation scheme, and the smallest diameter of the smallest enclosing disc of BUV that we obtained was 0.895611, attained at $x = y = 0.391113$, and we take this value as our approximation of δ_2 . Note, incidentally, that this choice of parameters implies that the edge CD of K is horizontal. It also implies that $|UV| = x = y = 0.391113$.

We now rotate W about the line UV , in the direction that keeps A and B on one side of it. The cross section of K by the rotating plane is an isosceles trapezoid, and we keep rotating the plane until it becomes a rectangle $UVV'U'$. As is easily checked, we have $|UU'| = |VV'| = 1 - x = 0.608887$, and the diameter of the smallest enclosing disc of $UVV'U'$, which is its diagonal, is ≈ 0.72368 , much smaller than δ_2 . An easy adaptation of an argument used earlier shows that, during this rotation of W about UV , every cross section is contained in the corresponding rotated copy of the disc of diameter δ_2 whose bounding circle passes through U and V . See Figure 7.2(iii).

We then slide W in the direction perpendicular to $UVV'U'$. During this sliding the cross section of K remains rectangular, so that UV keeps increasing and UU' keeps decreasing, while the sum of their lengths remains 1. We stop when we reach a ‘symmetric’ rectangle where the side parallel to UV (resp., UU') is of length $1 - x$ (resp., x). See Figure 7.2(iv).

The situation that we have reached is fully symmetric to the one after the first two steps,

and we can now complete the motion by a symmetric reversal of the first two steps.

To complete the proof, for the case where $\text{diam}(W) < \delta_2$, we observe that in this case W cannot pass through any vertex of K , because then, by definition of δ_2 , the smallest enclosing disc of any cross section through any vertex would have diameter larger than $\text{diam}(W)$. \square

8

Conclusion

In this thesis we have studied a variety of problems concerning collision-free motion of a convex polytope through a planar window, under several kinds of motion — sliding (translation in a fixed direction), purely translational motion, and general motion. We have presented several properties and characterizations of such motions, and obtained efficient algorithms for several special cases.

There are many open problems and directions for further research. One such direction is to derive efficient algorithms for the most general problem with all six degrees of freedom. In general one expects a solution that runs in roughly $O(n^6)$ time, but the special structure of the problem suggests that faster solutions should be possible. We are currently pursuing this direction, by reducing the number of degrees of freedom to four, by forcing two edges of the polytope to touch two edges of the window. We believe that this should lead to an algorithm with running time close to $O(n^4)$.

Moreover, the free configuration space of the polytope has a combinatorial complexity of $O(n^4)$, because each edge of the window can touch at most one element of the polytope (a face, a vertex or an edge). This suggests that if we form the arrangement of these $O(n)$ ‘contact surfaces’ in the sixth-dimensional configuration space, its complexity should be only $O(n^4)$. Transforming this observation into a comparably efficient algorithm is more involved, and we are currently studying this approach.

Furthermore, we believe that this bound is tight, in the sense that for every positive integer m there is a polytope with complexity $n > m$, such that its free configuration space has combinatorial complexity of $\Theta(n^4)$, which will show that an $O(n^4)$ algorithm for the general problem might be optimal in some cases, if we compute the entire free configurations space.

In addition, in Chapter 6 we presented an example in which a rotation is needed to pass

the polytope through the window. However, in this construction we only used a rotation about the line perpendicular to the plane that contains the window. This suggests the conjecture that maybe every convex polytope that can pass through a rectangular window can also pass through it by a motion consisting of arbitrary translations, and rotations only about the line perpendicular to the plane that contains the window. The results of Chapter 7 show that for circular window this claim is false, but the status of the conjecture is still open for a rectangular window.

Bibliography

- [1] P. K. Agarwal, Simplex range searching, in *Journey Through Discrete Mathematics* (M. Loeb, J. Nešetřil and R. Thomas, eds.), Springer, Heidelberg 2017, pp. 1–30.
- [2] P. K. Agarwal and J. Erickson, Geometric range searching and its relatives, in *Advances in Discrete and Computational Geometry*, Contemp. Math. 223 (B. Chazelle, J. E. Goodman, and R. Pollack, eds.), AMS Press, Providence, RI, 1999, pp. 1–56.
- [3] M. de Berg, L. J. Guibas and D. Halperin, Vertical decompositions for triangles in 3-space, *Discrete Comput. Geom.* 15(1) (1996), 35–61.
- [4] P. Bose and D. Halperin and S. Shamaï, On the separation of a polyhedron from its single-part mold, *13th IEEE Conference on Automation Science and Engineering*, (2017), 61–66.
- [5] B. Chazelle, H. Edelsbrunner, L. J. Guibas and M. Sharir, A singly exponential stratification scheme for real semi-algebraic varieties and its applications, *Theor. Comput. Sci.* 84(1) (1991), 77–105.
- [6] H. Choset, K. M. Lynch, S. Hutchinson, G. Kantor, W. Burgard, L. E. Kavraki and S. Thrun, *Principles of Robot Motion: Theory, Algorithms, and Implementation*, MIT Press, 2005.
- [7] J.-C. Latombe, *Robot Motion Planning*, The Kluwer international series in engineering and computer science, Kluwer, volume 124, 1991.
- [8] H. T. Croft, K. J. Falconer and R. K. Guy, *Unsolved Problems in Geometry*, Problem Books in Mathematics, Springer Verlag, Heidelberg, 1991.
- [9] H. E. Debrunner and P. Mani-Levitska, Can you cover your shadows? *Discrete Comput. Geom.* 1 (1986), 45–58.
- [10] E. Fogel, D. Halperin and R. Wein, *CGAL Arrangements and their Applications - A Step-by-Step Guide*, Geometry and Computing, 7, Springer, 2012.
- [11] P. Gibbs, A computational study of sofas and cars, *Computer Science* 2 (2014), 1–5.
- [12] D. Halperin, J.-C. Latombe and R. H. Wilson, A general framework for assembly planning: The motion space approach, *Algorithmica* 3-4 (2000), 577–601.

- [13] D. Halperin, L. Kavraki and K. Solovey, Robotics, Chapter 51 in *Handbook of Discrete and Computational Geometry*, Chapman & Hall/CRC, 3rd edition, 2018, 1343–1376.
- [14] D. Halperin and M. Sharir, Arrangements, Chapter 28 in *Handbook of Discrete and Computational Geometry*, Chapman & Hall/CRC, 3rd edition, 2018, 723–762.
- [15] D. Halperin, O. Salzman and M. Sharir, Algorithmic Motion Planning, Chapter 50 in *Handbook of Discrete and Computational Geometry*, Chapman & Hall/CRC, 3rd edition, 2018, 1311–1342.
- [16] A. Hatcher, *Algebraic Topology*, Cambridge University Press, Cambridge, 2002, pages 61–62.
- [17] https://en.wikipedia.org/wiki/Covering_space
- [18] <https://math.stackexchange.com/questions/1364880/smallest-cylinder-into-which-a-regular-tetrahedron-can-fit>
- [19] L. E. Kavraki, P. Šestka, J.-C. Latombe and M. H. Overmars, Probabilistic roadmaps for path planning in high dimensional configuration spaces, *IEEE Transactions on Robotics* 12(4) (1996), 566–580.
- [20] J. J. Kuffner and S. M. LaValle, RRT-Connect: An efficient approach to single-query path planning, *IEEE International Conference on Robotics and Automation (ICRA)* (2000), 995–1001.
- [21] S. M. LaValle, *Planning Algorithms*, Cambridge University Press, 2006
- [22] W. H. Plantinga and C. R. Dyer, Visibility, occlusion, and the aspect graph, *Int. J. Computer Vision* 5 (1990), 137–160.
- [23] O. Salzman, M. Hemmer and D. Halperin, On the power of manifold samples in exploring configuration spaces and the dimensionality of narrow passages, *IEEE Trans Autom. Sci. Eng.* 12(2) (2015), 529–538.
- [24] O. Salzman, M. Hemmer, B. Raveh and D. Halperin, Motion planning via manifold samples, *Algorithmica*, 67(4) (2013), 547–565.
- [25] M. Sharir and P. K. Agarwal, *Davenport–Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.
- [26] J. Snoeyink and J. Stolfi, Objects that cannot be taken apart with two hands, *Discrete Comput. Geom.* 12 (1994), 367–384.
- [27] G. Toussaint, Movable separability of Sets, *Comput. Geom.* 2 (1985) 335–375.

תקציר

התיזה עוסקת בגרסה תלת-מימדית של בעיית הספה (the sofa problem). בעיית הספה (או בעיית הזזת הספה) היא בעיה עתיקה במתמטיקה, שהתפרסמה רשמית באמצע שנות השישים של המאה הקודמת. הבעיה עוסקת בתכונות הצורה המקסימלית שיכולה לעבור במסדרון ברוחב נתון עם פנייה של 90 מעלות. הבעיה שאנו דנים בה היא חקר תכונות הפאוניס הקמורים שיכולים לעבור דרך חור מלבני ("חלון") על קיר אינסופי.

אנחנו חוקרים את הבעיה בשני כיוונים: תכונות הכרחיות של פאוניס קמורים שיכולים לעבור דרך החלון, ואלגוריתמים יעילים לתכנון תנועה כני"ל במידה והיא קיימת.

אנחנו עוסקים בהרבה וריאנטים של הבעיה, כגון:

- כשהתנועה מוגבלת רק להזזה ובכיוון יחיד ("החלקה"),
- כשהתנועה מוגבלת רק להזזה (בלי סיבובים),
- כשזוג צלעות של החלון הן אינסופיות ("שער"),
- כשהחלון מעגלי.

הבעיה הוצגה לראשונה במאמר של Toussaint משנת 1985 ללא פיתרון, ולא מצאנו מאמר נוסף שדן בבעיה ספציפית זו.

התוצאות המוצגות בתיזה כוללות:

- בפרק 2 אנו דנים בהחלקה דרך חלון מלבני, ומראים שקיום החלקה של הפאון דרך החלון מבלי שהם יחתכו גוררת קיום של החלקה (אולי אחרת), בכיוון שמאונך למישור החלון. זה מאפשר לנו להתאים לכל החלקה מעין "צורה קנונית", ואנו מציגים אלגוריתם יעיל לקבוע האם קיימת החלקה שכזאת, ואם כן למצוא אותה.
- בפרק 3 אנו חוקרים את המקרה של תנועה כללית (הזזות וסיבובים) דרך שער (פס החסום על ידי שתי צלעות מקבילות לא חסומות). אנחנו מראים שקיום תנועה של הפאון דרך השער גוררת קיום החלקה שלו דרך השער, ומכאן קל לקבל אלגוריתם יעיל לתכנון תנועה שכזאת.
- בפרק 4 אנחנו עומדים על הקשר שבין חלון כללי שהפאון יכול לעבור דרכו, לעובי של שער שהחלון יכול לעבור דרכו. אנחנו מראים שאם הפאון יכול לעבור בחלון עם קוטר d (על ידי תנועה כללית כלשהי), אז לכל כיוון \vec{v} אפשר לסובב את הפאון סביב \vec{v} בלבד כך שיוכל להחליק דרך שער בעובי d שמונח על מישור שמאונך ל- \vec{v} .

- בפרק 5 אנו דנים בתנועה של פאון דרך חלון מלבני, המורכבת מהזזות בלבד. אנחנו מראים שקיום תנועה שכזאת גוררת קיום של החלקה דרך החלון, מה שמאפשר שימוש באלגוריתם שהצגנו בפרק 2.
- בפרק 6 אנו מראים שסיבובים לפעמים הכרחיים להעביר פאון דרך חלון. קונקרטי, אנו מציגים פאון שניתן להעביר דרך חלון מלבני (למעשה ריבועי) בתנועה המורכבת מסיבובים והזזות, אך לא ניתן להעבירו בחלון בתנועה המורכבת מהזזות בלבד.
- בפרק 7 אנו מציגים את ההבדלים בין שלושת סוגי התנועה—החלקה, תנועה של הזזות בלבד, ותנועה כללית—עבור חלון מעגלי. אנו מראים שהארבעון המשוכלל (בעל אורך צלע 1) יכול להחליק דרך חלון שקוטרו לפחות 1, אינו יכול להחליק אך יכול לעבור בתנועה של הזזות בלבד דרך חלון שקוטרו קטן מ-1 אך גדול מ-0.901388, אינו יכול לעבור בהזזות בלבד אך יכול לעבור בתנועה הכוללת סיבובים אם הקוטר בין 0.895611 ל-0.901388, ואינו יכול לעבור כלל אם הקוטר קטן יותר.



The Raymond and Beverly Sackler Faculty of Exact Sciences
Tel Aviv University

הפקולטה למדעים מדויקים
ע"ש ריימונד וברלי סאקלר
אוניברסיטת תל אביב

זריקת ספה דרך חלון

חיבור זה מוגש כחלק ממילוי הדרישות לקבלת
התואר "מוסמך למדעים" (M.Sc.) בבית הספר למדעי המחשב ע"ש בלוטניק
באוניברסיטת תל-אביב
ע"י

איתי יהודה

העבודה הוכנה באוניברסיטת תל-אביב
בהדרכתם של פרופסור דן הלפרין ופרופסור מיכה שריר
אב התש"פ