TEL-AVIV UNIVERSITY
RAYMOND AND BEVERLY SACKLER
FACULTY OF EXACT SCIENCES
SCHOOL OF COMPUTER SCIENCE

Minkowski Sums of Simple Polygons

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

by

Eduard Oks

The research work for this thesis has been carried out at Tel-Aviv University under the supervision of Prof. Micha Sharir

June 2005

Abstract

Let P be a simple polygon with m edges, which is the disjoint union of k simple polygons, all monotone in a common direction u, and let Q be another simple polygon with n edges, which is the disjoint union of ℓ simple polygons, all monotone in a common direction v. We show that the combinatorial complexity of the Minkowski sum $P \oplus Q$ is $O(k\ell mn\alpha(\min\{m,n\}))$. Some structural properties of the case $k = \ell = 1$ have been (implicitly) studied in [25]. We rederive these properties using a different proof, apply them to obtain the above complexity bound for $k = \ell = 1$, obtain several additional properties of the sum for this special case, and then use them to derive the general bound.

Acknowledgments

I deeply thank Prof. Micha Sharir for his guidance and his tremendous help in preparing this document. I also thank him for introducing me to the field of computational geometry. I also would like to thank Prof. János Pach for his help and contributions.

Contents

Chapter 1

Introduction

Given two sets $P, Q \subset \mathbb{R}^2$, their *Minkowski sum*, denoted by $P \oplus Q$, is defined as the set

$$P \oplus Q := \{p + q \mid p \in P, q \in Q\}.$$

This is a fundamental construct that arises in a wide range of applications, including robot motion planning [5, 19], assembly planning [11], and computer-aided design and manufacturing (CAD/CAM)[7].

Consider for example robot motion planning. This application involves placements and translational motion of an object in the presence of another object, which acts as a stationary obstacle. Let Q denote the obstacle, and P the robot, which is only allowed to translate. Assuming, without loss of generality, that the origin o lies in P, and denoting by $-P$ the reflection of P with respect to o, it follows by definition that $K = (-P) \oplus Q$ is the set of all vectors v such that translating P by v makes it intersect Q. In the study of motion planning, this sum is called a *configuration space obstacle, or $\mathcal{C}$-obstacle*. Hence the complement K^c of K is a representation of the space of all free placements of P (namely, placements disjoint from Q). This observation makes Minkowski sums a central tool in the analysis of translational motion planning (see, e.g., [5, 23]), and we will also use this interpretation in our analysis.

1.1 The combinational complexity of Minkowski Sums

Motivated by these applications, there has been much work on obtaining sharp bounds on the size of the Minkowski sums of two sets in two and three dimensions, and on developing fast algorithms for computing Minkowski sums. Throughout the thesis we only consider Minkowski sums of two polygons P, Q in the plane. It is well known that if both P and Q are convex polygons, with m and n vertices respectively, then $P \oplus Q$ is a convex polygon with at most $m + n$ vertices, and it can be computed in $O(m + n)$ time [5]

However, if only P is convex and Q is non-convex then $P \oplus Q$ has $\Theta(mn)$ vertices in the worst case (see Figure 1.1) [5, 18]. The proof relies on special properties of a set of *pseudodiscs*. We say that a pair of planar objects o_1 and o_2 are a pair of pseudodiscs if each

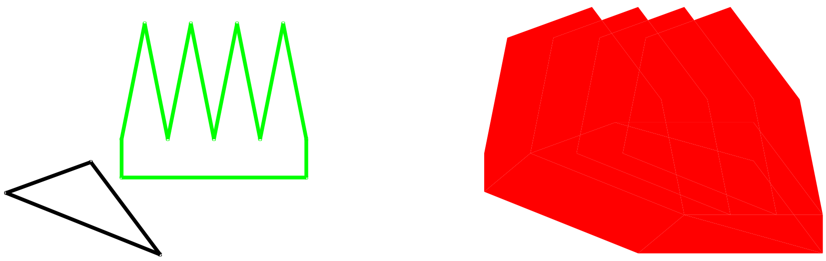

Figure 1.1: P is a convex polygon with m vertices and Q is a non-convex polygon with n vertices. The complexity of $P \oplus Q$ is $\Theta(mn)$. (Computed with Eyal Flato Minkowski Sum package [9].)

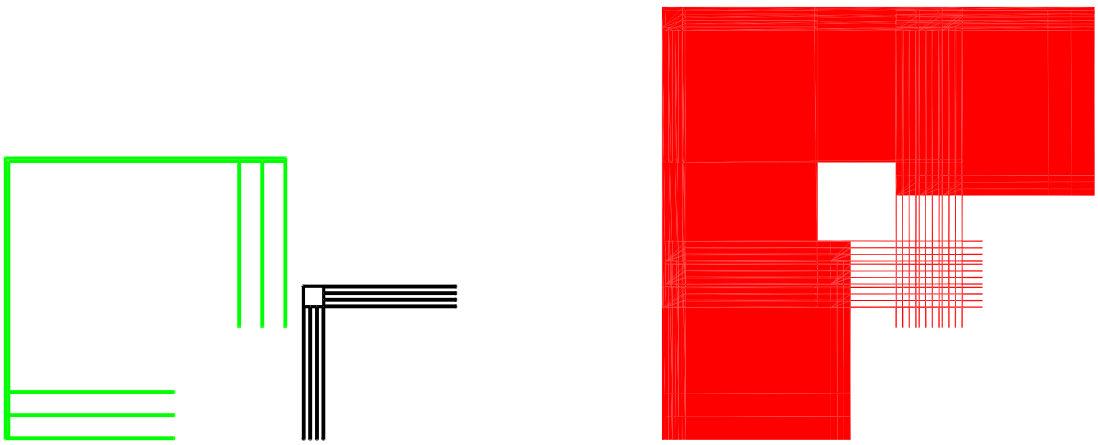

Figure 1.2: P and Q are non-convex polygons with m and n vertices respectively. The complexity of $P \oplus Q$ is $\Theta(m^2n^2)$. This is the worst case scenario. (Computed with Eyal Flato Minkowski Sum package [9].)

of the sets $o_1 \setminus o_2$ and $o_2 \setminus o_1$ is connected. Kedem et al. [18] proved that the number of vertices of the boundary of the union of a collection of n pseudodiscs is at most $6n - 12$, where a vertex is an intersection point of two boundaries. The time needed to construct $P \oplus Q$ in that case is $O(nm \log(nm))$.

Finally, if P and Q are non-convex polygons, then $P \oplus Q$ is a portion of the arrangement of $O(mn)$ segments, where each segment is the Minkowski sum of a vertex of P and an edge of Q, or vice-versa. Therefore the size of $P \oplus Q$ is $O(m^2n^2)$ and this bound is tight in the worst case. $P \oplus Q$ can be computed within that time [5, 16] (see Figure 1.2).

1.2 Related Work

1.2.1 Special cases of Minkowski sums

In the previous section we presented the well known combinatorial bounds on the complexity of the Minkowski sum of two polygonal sets. In motion planning applications, one is often interested in computing only a single connected component of the complement of $P \oplus Q$ [20]. Har-Peled et al. [14] showed that the complexity of a single face of the complement of $P \oplus Q$ is $\Theta(mn\alpha(n))$ in the worst case, where m and n are the number of vertices of P and Q respectively (without loss of generality $n < m$), and $\alpha(\cdot)$ is the functional inverse of Ackermann's function [24].

The special case where P is a simple polygon and Q is a line segment has been recently analyzed in [21], it was shown that in that case $P \oplus Q$ has at most $2n - 1$ edges, and this bound is tight in the worst case.

Ramkumar [22] presents a different approach to construct the outer face of the Minkowski sum. Existing methods rely on general algorithms for computing a single face in an arrangement of k line segments, which takes $O(k(\log k)(\alpha(k)))$ time. Instead, his algorithm exploits a new insight into the relationship between convolutions and Minkowski sums and, though asymptotically slower, has practical advantages for realistic polygon data. His method consists of traversing each cycle of the convolution, detecting self-intersections, and snipping off the loops thus created. In order to detect self-intersections, the algorithm adapts the geodesic triangulation ray-shooting data structure to answer ray-shooting queries on a dynamic polygonal line of size n, in $O(\log^2 n)$ amortized time. The algorithm constructs the outer face of the Minkowski sum of two simple polygons of size m and n, respectively, in time $O((k + (m+n)\sqrt{l}) \log^2(m+n))$ where k is the size of the convolution (k can be $O(mn)$ in the worst case) and l is the number of cycles in the convolution.

De Berg and van der Stappen [6] report on results concerning the relation between the fatness of the Minkowski sum of two sets and the fatness of these sets. The fatness of an object is determined by the emptiest ball centered inside the object and not fully containing it in its interior. Using this measure, they show that the fatness of $A \oplus B$ is at least as large as $\min(fatness(A); fatness(B))$, when A and B are connected closed and bounded sets in $\mathbb{R}^d$.

Monotone Polygons

Another special case where P is monotone, was studied by A.H. Barrera [2], who showed that the Minkowski sum of a monotone polygon P with n edges and a convex polygon Q with m edges can be calculated in $O(nm)$ time and space. For two monotone polygons P and Q (monotone in the same direction), an $O(nm \log(nm))$ time algorithm is presented, which is almost tight in the worst case, because it is shown that the number of edges in the sum is $\Theta(nm\alpha(min(n, m)))$ in the worst case (For example see Figure 1.3). In case when P is monotone and Q is simple, an $O((nm + k) \log(nm))$ time algorithm is given, where k in

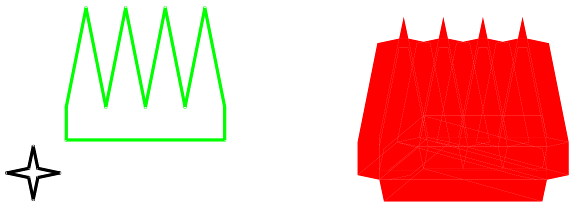

Figure 1.3: P and Q are x-monotone, non-convex polygons with m and n vertices respectively. The complexity of $P \oplus Q$ is $\Theta(nm\alpha(min(n,m))$ and the resulted polygon is x-monotone. (Computed with Eyal Flato Minkowski Sum package [9].)

the worst case can be $\Theta(n^2m)$, and the number of edges in the sum is $\Theta(n^2m)$. Barrera also proved that computing the Minkowski sum of two polygons is at least as hard as *sorting* $X + Y$ [3]. The best known time bound for solving this sorting problem is $O(n^2 \log(n))$, and it is an open problem whether this can be improved.

Discrete Approximations

Hartquist et al. [15] suggest a computing strategy for applications that use offsets, sweeps and Minkowski operations based on the ray-representation method. This method involves clipping a given input to a grid of rays and applying the mathematical definitions and operators (such as Minkowski sums) on the resulting discrete set. The authors aim to solve motion planning, process-modeling and visualization problems, and they present a hardware design for those applications.

Kavraki [17] uses the Fast Fourier Transform (FFT) algorithm on the bitmaps of a robot and obstacles to find the corresponding configuration-space obstacles for the robot translating among the obstacles. This method approximates the configuration space obstacles. The method is inherently parallel and can benefit from existing experience and special hardware for computing the FFT.

1.2.2 Applications

The translational robot motion planning problem planning is a convenient case study for Minkowski sum algorithms, and we therefore detail it and use it as an example in the rest of this thesis. There are many more applications in which the Minkowski sum operation is a useful tool. Some examples are listed here.

Polygon containment

Given two polygons P and Q in the plane, we wish to determine whether P can be contained (by translation, or by other geometric transformation) inside Q. This problem is known as the polygon containment problem [4]. If we only allow a translation (namely, the orientation of P is fixed), the problem can be solved as follows: Consider the complement of Q as an obstacle for the robot P, and try to place P such that it does not penetrate this obstacle. Practically, let B be the bounding box of Q and let $Q^c = B \setminus Q$ be the complement of Q. The free placements for P inside Q can be found by computing $Q^c \oplus (-P)$.

Geographic Information Systems; Cartographic generalization

Geographic Information Systems (GIS) are increasingly being studied in computational geometry. There are some problems in GIS that are closely related to our work. One of them was first posed by Marc van Kreveld [26]: Given two simple polygons, P and Q with m and n edges respectively, find the minimum length translation of one polygon relative to the other that will make the two polygons interior disjoint. The solution is based on the fact that P and Q are disjoint *if and only if* $P \oplus (-Q)$ does not contain the reference point of Q. Assuming that in their original placement P and Q intersect and that the reference point of Q is at the origin o (we can assume this without loss of generality), the shortest translation of Q relative to P that separates them is given by the point on the boundary of $P \oplus (-Q)$ which is closest to the origin o (this point can lie on the boundary of some hole of the Minkowski sum). The running time of this algorithm is $\Theta(n^2m^2)$ in the worst case, which is dominated by the complexity of the Minkowski sum algorithm.

Robust and efficient construction of planar Minkowski sums in practice

Flato and Halperin at [8, 10] present several different approaches to calculate the Minkowski sum of two simple polygons using the CGAL software library and its *planar map* package. The algorithms decompose each of the polygons P, Q into convex sub-pieces, form the Minkowski sums of the separate pieces, and then construct the union of these sub-sums. The algorithms differ in the implementation of the union step – calculation of the union of the Minkowski sub-sums.
In a subsequent work [1, 8] Agarwal, Flato and Halperin continue this research, by analyzing different decomposition methods, the first step of the Minkowski sum algorithm, such as *triangulation, convex decomposition with/without Steiner points, approximations and heuristics.* The emphasis of their work was to study the effect of the decomposition method on the efficiency of the overall process.
It is shown in these works that: (i) Triangulations are too costly. (ii) What constitutes a good decomposition for one of the input polygons depends on the other input polygon. Consequently, they develop a procedure for simultaneously decomposing the two polygons such that a "mixed" objective function is minimized. (iii) There are optimal decomposition algorithms that significantly expedite the Minkowski sum computation, but the decomposition itself is expensive to compute. In such cases, simple heuristics that approximate the optimal decomposition seem to perform very well in practice. They examined several criteria that affect the running time of the Minkowski sum algorithm. The most effective optimization is

minimizing the number of convex sub-polygons. Thus, triangulations which are widely used in the theoretical literature are not practical for the Minkowski sum algorithms.

1.3 Thesis Outline

The thesis presents a general technique for analyzing the complexity of the Minkowski sum of two simple polygons, using their partition into monotone pieces.

Specifically, let P be a simple polygon with m edges, which is the disjoint union of k simple polygons, all monotone in a common direction u, and let Q be another simple polygon with n edges, which is the disjoint union of ℓ simple polygons, all monotone in a common direction v. In Chapter 2 we show that the combinatorial complexity of the Minkowski sum $P \oplus Q$ is $O(k\ell mn\alpha(\min\{m,n\}))$. Some structural properties of the case $k = \ell = 1$ have been (implicitly) studied by Toussaint and ElGindy [25]. We re-derive these properties using a different proof and apply them to obtain the above complexity bound for $k = \ell = 1$. We obtain additional properties of the sum for this special case. Specifically, we show that the boundary of the Minkowski sum is the concatenation of two u-monotone and two v-monotone connected polygonal chains which are pairwise openly disjoint, and that the number of pockets in $P \oplus Q$ is $O(m + n)$. At the end we use all these properties to derive the complexity of the sum for the general case. (A pocket is a maximal sub-chain of the boundary that is monotone in one of the directions u, v and has unique local minimum or maximum in the other direction.) The bound is worst-case tight for $k = \ell = 1$, as follows from the construction of Barrera [2], and is almost tight in the general case $k = \Theta(m)$, $\ell = \Theta(n)$.

We give concluding remarks and suggest for further research in Section 2.6.

Chapter 2

Minkowski Sums of Monotone Polygons

In this chapter we study the coplexity of the Minkowski sum of two simple polygons, and express it in terms of the number of monotone sub-polygons into which each of them can be decomposed.

A simple polygon P is said to be *monotone* in direction u (also referred to as u-*monotone*) if every line orthogonal to u intersects P in a connected (possibly empty) interval. We can decompose any simple polygon P with m edges into simple sub-polygons, all monotone in some specified direction u, by drawing a vertical segment through each vertex of P which is a locally u-extremal point of ∂P, and by extending that segment *inside* P till it hits ∂P again. These segments decompose P into $O(m)$ pairwise openly disjoint u-monotone simple polygons, and this bound is tight in the worst case.

Let P be a u-monotone simple polygon with m edges, and let Q be a v-monotone simple polygon with n edges, for two (possibly different) directions u, v. We show (Theorem 2.2.1) that the complexity of $P \oplus Q$ in this case is only $O(mn\alpha(\min\{m, n\}))$, which is tight in the worst case. (The upper bound was obtained by Barrera [2] for the special case $u = v$. He also showed that the lower bound can be attained in this case.)

The proof relies on the following separation property, due to Toussaint and ElGindy [25]: Given *disjoint* monotone polygons P and Q as above, we can translate P to infinity, without colliding with Q, in at least one of the four directions $u \pm \pi/2$, $v \pm \pi/2$. This property implies that $P \oplus Q$ is *simply connected*, from which the complexity bound follows using known bounds on the complexity of a single face in an arrangement of line segments; see, e.g., [24].

We provide, in Theorem 2.1.1, an alternative proof of the result of [25], and then use it to obtain the asserted complexity bound. Moreover, we derive several additional structural properties of the sum $P \oplus Q$ of two monotone simple polygons. For example, we show that its boundary is the concatenation of four connected portions, two of which are u-monotone and two v-monotone. We also show that the number of *pockets* along $\partial(P \oplus Q)$ is only $O(m+n)$. This notion will be defined and analyzed in Section 2.4. This is roughly equivalent to asserting that the number of points on $\partial(P \oplus Q)$ that are locally x-extremal or y-extremal is $O(m+n)$.

We next use all these properties to prove the main result of the paper, which asserts that, if P is a simple polygon with m edges which is the disjoint union of k simple u-monotone sub-polygons, and Q is a simple polygon with n edges which is the disjoint union of ℓ simple v-monotone sub-polygons, for any (possibly distinct) directions u, v, then the complexity of $P \oplus Q$ is $O(k\ell mn\alpha(\min\{m,n\}))$. This (almost) properly interpolates between the two extreme cases $k = \ell = 1$ (where the bound is worst case tight), and $k = \Theta(m)$, $\ell = \Theta(n)$ (where we get an extra $\alpha(\cdot)$ factor).

2.1 Separating Two Monotone Chains

Theorem 2.1.1 (slightly reformulated) was already proven in [25]. We present here a different proof, using a functional representation of monotone polygonal paths.[1]

Theorem 2.1.1 *Let $f(x) : [a,b] \mapsto \mathbb{R}$, $g(y) : [c,d] \mapsto \mathbb{R}$ be (graphs of) continuous real functions defined on the above intervals of the x- and y-axes, respectively, that do not intersect each other. Then $f(x)$ can be translated to infinity along at least one of the four axis directions without colliding with $g(y)$.*

We say that a point p of the plane is *directly to the right* of another point q if the half-line starting at q and pointing to the positive x-direction passes through p. The notions of being directly to the left, directly above, and directly below, are defined in an analogous manner.

Lemma 2.1.2 *Suppose that $g(y)$ has a point directly to the right of the right endpoint of $f(x)$. Then $g(y)$ has no point directly to the left of any point of $f(x)$.*

Proof: It is enough to show, by symmetry, that this holds for every $y \geq f(b)$. If for every $y \geq f(b)$ we have $g(y) > b$, we are done. Otherwise, set $x_0 := g(f(b))$, $y_0 := f(b)$, $x_1 := b$. If there exist $y \geq f(b)$ with $g(y) \leq b$, then denote by y_1 the infimum of all such y. Then, by continuity, $g(y_1) = x_1$, and the statement (that $g(y)$ has no point directly to the left of any point of $f(x)$) holds on the interval $[y_0, y_1]$. If $f(x)$ remains under y_1 in every point to the left of x_0, then the statement holds on the whole interval $[y_0, d]$. Otherwise, there is a largest x_2 where (proceeding from right to left) $f(x)$ first attains y_1. Similarly, now $g(y)$ either remains to the right of x_2 all the way to the end $(d, g(d))$, or there exists y_2 where $g(y)$ first reaches x_2. See Figure 2.1.

This alternating construction terminates in finitely many steps, for otherwise we would obtain a bounded sequence (x_0, y_0), (x_1, y_0), (x_1, y_1), (x_2, y_1), (x_2, y_2), …, monotone in both coordinates, and its limit would be a common point of $f(x)$ and $g(y)$. It is easily seen that the termination of the process implies the statement of the lemma over the interval $[y_0, d]$, and a symmetric argument implies it for $[c, y_0]$. □

[1] We are grateful to János Pach for suggesting this proof, which has simplified our earlier analysis.

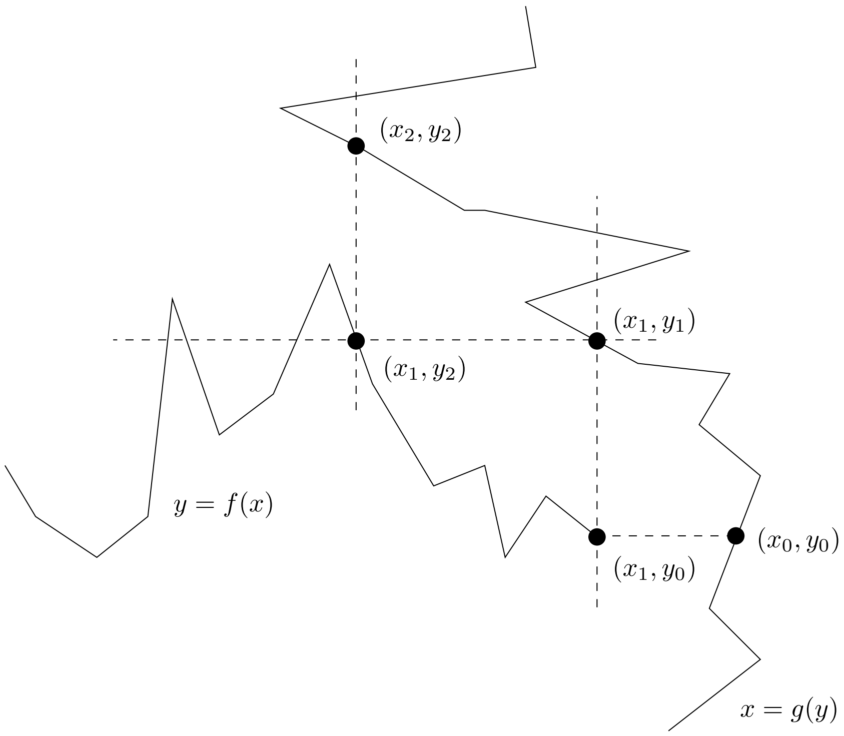

Figure 2.1: The 'staircase' of critical values $x_0, y_0, x_1, y_1, \ldots$

Lemma 2.1.3 *If any point of $g(y)$ is directly to the right of any point of $f(x)$, then the largest x_0 such that $f(x_0)$ is directly to the left of a point of $g(y)$ has the property that $g(y)$ has no point directly to the left of any point of $f([a, x_0])$.*

Proof: Put $x_0 = \sup\{x \in [a, b] \mid f(x) \in [c, d]$ and $g(f(x)) > x\}$, and apply Lemma 2.1.2 to g and to f restricted to $[a, x_0]$. □

Proof of Theorem 2.1.1: We can assume that $g(y)$ has a point below $f(x)$ and a point above $f(x)$, otherwise we are done. By Lemma 2.1.3 (and by symmetry), there is a smallest y^- such that $(g(y^-), y^-)$ is above $f(x)$, and then all points of $g(y)$ above y^- are not below $f(x)$, and there is a largest y^+ such that $(g(y^+), y^+)$ is below $f(x)$, and then the points of $g(y)$ below y^+ are not above $f(x)$. Obviously, we have $y^- > y^+$.

By definition of y^-, y^+, the points of $g(y)$ on the interval (y^+, y^-) are neither above nor below $f(x)$, so these $g(y)$ values do not belong to $[a, b]$. By Bolzano's theorem, they must all be smaller than a or all bigger than b. Assume that all $g(y) > b$ on this interval. Then, by the continuity of $g(y)$, we have $g(y^-) = b = g(y^+)$. So $f(b) \in (y^+, y^-)$, hence $b < g(f(b))$. Thus, we can apply Lemma 2.1.2 to conclude that no point of $g(y)$ is directly to the left of any point of $f(x)$, so $f(x)$ can be translated to the left. □

Remarks: (1) Theorem 2.1.1 also holds when the graphs $f(x), g(y)$ touch each other without crossing at any finite number of points. We omit details of this extension.

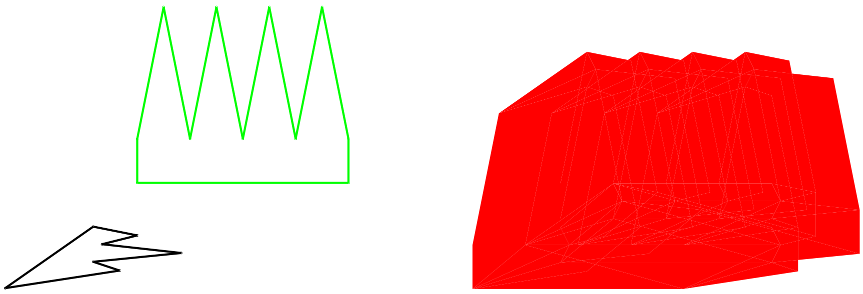

Figure 2.2: P and Q are non-convex polygons with m and n vertices respectively, P is x-monotone and Q is y-monotone. The complexity of $P \oplus Q$ is $\Theta(nm\alpha(min(n,m))$ and the boundary can be presented as the concatenation of two x-monotone and two y-monotone connected polygonal chains. (Computed with Eyal Flato Minkowski Sum package [9].)

(2) Theorem 2.1.1 also holds when we replace $f(x), g(y)$ by any pair of bounded connected arcs, each monotone in some direction, and these directions are not required to be orthgonal to each other (as is the case in the theorem). If u, v are the directions of monotonicity, the theorem asserts that one arc can be translated to infinity in one of the four directions $u \pm \pi/2$, $v \pm \pi/2$, without meeting the other arc. Indeed, apply a 'shearing' affine transformation which maps the direction $u + \pi/2$ to the positive y-direction, and maps $v - \pi/2$ to the positive x-direction. This transforms the scenario into the one studied above, and an application of Theorem 2.1.1 in the new scenario, combined with the inverse shearing transformation, establishes the asserted property.

(3) Theorem 2.1.1 also holds when we replace $f(x), g(y)$ by any pair of simple polygons, monotone in the x- and y-directions, or, as in (2), in any two directions. This extension follows easily from the preceding analysis, and is the one proved in [25] (using a different approach).

2.2 Minkowski Sum of Two Monotone Polygons

Theorem 2.2.1 *Let P and Q be two simple monotone polygons in two (possibly different) directions, having m and n edges, respectively. Then the complexity of the Minkowski sum $P \oplus Q$ is $O(mn\alpha(\min\{m,n\}))$, (see Figure 2.2 for example).*

Proof: Suppose that P is monotone in direction u and that Q is monotone in direction v. Arguing as in Remark (2) of the preceding section, we may assume that u is the x-direction and that v is the y-direction. Let $\widetilde{P} = -P$ denote the reflection of P about the origin. Let t be a vector in the plane such that $t \notin P \oplus Q$. Then, by definition, $\widetilde{P}_t = \widetilde{P} + t$ is disjoint from Q. By Theorem 2.1.1 and Remark (3) following it, we can translate $\widetilde{P}_t$ in one of the four coordinate directions all the way to infinity, so that it does not intersect Q during the motion. This implies that there is a ray ρ in one of the four axis directions that emanates from t and is disjoint from $P \oplus Q$. This in turn implies that the complement of $P \oplus Q$ has no bounded components ('holes' of $P \oplus Q$), and thus $P \oplus Q$ is simply connected.

In other words, the boundary of $P \oplus Q$ is connected, and coincides with the boundary of the unbounded face of its complement. Let Σ denote the set of all line segments of the form $e + v$, where e is an edge of P and v is a vertex of Q, or e is an edge of Q and v is a vertex of P. Σ consists of $2mn$ segments, and any point on $\partial(P \oplus Q)$ must be contained in one of these segments. As is well known (see, e.g., [24]), the complexity of any single face in an arrangement of $2mn$ segments is $O(mn\alpha(mn))$. To obtain the slightly improved asserted bound, assume, without loss of generality, that $m \leq n$. Note that Σ can be represented as the union of $2m$ subsets, each consisting of the sums of all edges of Q with a fixed vertex of P, or of the sums of all vertices of Q with a fixed edge of P. Each subset consists of pairwise (openly) disjoint segments, so the complexity of the sub arrangement that they form is $O(n)$. We then apply the *Combination Lemma* of Har-Peled [13], which implies that the complexity of a single face in the overlay of $2m$ arrangements, each of complexity $O(n)$, is $O(mn\alpha(m))$. See also [?] for an alternative proof. □

2.3 The Boundary of the Sum of Two Monotone Polygons

In what follows, we assume that P and Q are monotone in the x- and y-directions, respectively. As noted above, this involves no loss of generality.

Theorem 2.3.1 *Let P and Q be two simple polygons monotone in the x- and the y-directions, respectively. Then the boundary of $S = P \oplus Q$ is the concatenation of two x-monotone and two y-monotone connected polygonal chains, which are pairwise openly disjoint.*

See Figure 2.3.

We use in the proof the interpretation, already mentioned above, of $S = P \oplus Q$ as the space of all "forbidden" translations of $\widetilde{P} = (-P)$ at which it intersects Q, which we regard as stationary. The boundary of S is the set of all translations where $\widetilde{P}$ touches Q, but does not intersect its interior.

By Theorem 2.1.1, each point $v \in \partial S$ can be classified into one (or more) of the four following types:

Top, if $\widetilde{P}$ can be moved from v to infinity in the positive y-direction without penetrating into Q.

Bottom, if $\widetilde{P}$ can be moved from v to infinity in the negative y-direction without penetrating into Q.

Left, if $\widetilde{P}$ can be moved from v to infinity in the negative x-direction without penetrating into Q.

Right, if $\widetilde{P}$ can be moved from v to infinity in the positive x-direction without penetrating into Q.

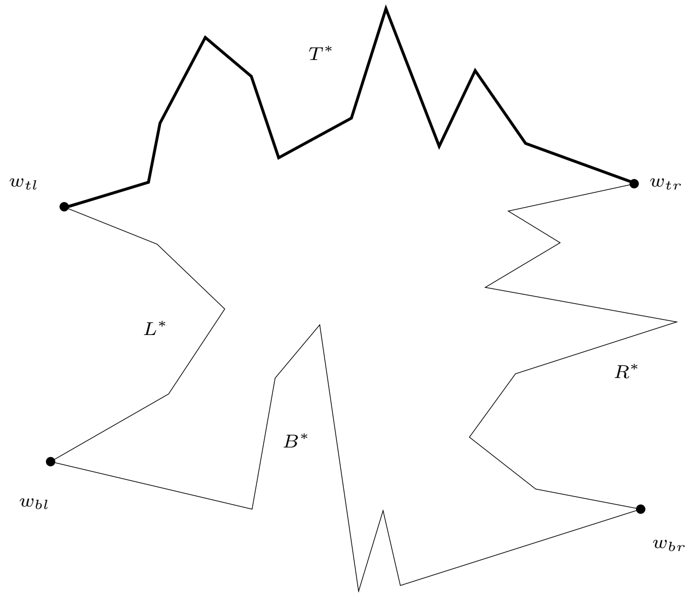

Figure 2.3: The Minkowski sum of two monotone polygons, and (one possible) partition of its boundary into x-monotone portions T^*, B^*, and y-monotone prtions L^*, R^*, delimited by the points $w_{tr}, w_{tl}, w_{bl}, w_{br}$. The portion T^* is highlighted.

We can therefore write ∂S as the union of four subsets T, B, L, R, where T (resp., B, L, R) consists of all top (resp., bottom, left, right) points on ∂S. By definition, all of these sets are closed. These sets are not necessarily disjoint, but the only points of ∂S that belong to $T \cap B$ are the leftmost and rightmost points of S. Similarly, only the topmost and bottommost points of S can belong to $L \cap R$. Any other pair of sets can have a more substantial intersection.

Proof of Theorem 2.3.1: Let w_t, w_b, w_l, w_r denote respectively the highest, lowest, leftmost, and rightmost points of ∂S. (We assume general position which makes these points unique). These four points partition ∂S into four connected portions, which we denote as the northeastern portion NE (lying clockwise from w_t to w_r), the southeastern portion SE (lying clockwise from w_r to w_b), the southwestern portion SW (lying clockwise from w_b to w_l), and the northwestern portion NW (lying clockwise from w_l to w_t). Note that the points w_t, w_b, w_l, w_r need not be distinct, although we always have $w_t \neq w_b$ and $w_l \neq w_r$. See Figure 2.4.

It is easily seen that

$$NE \subseteq T \cup R, \quad SE \subseteq B \cup R, \quad SW \subseteq B \cup L, \quad NW \subseteq T \cup L.$$

More precisely, except possibly for their endpoints, these chains satisfy

$$NE \cap (B \cup L) = \emptyset, \quad SE \cap (T \cup L) = \emptyset, \quad SW \cap (T \cup R) = \emptyset, \quad NW \cap (B \cup R) = \emptyset.$$

Lemma 2.3.2 *Let $u, v \in NE$ such that v lies clockwise to u. It is impossible that $u \notin T$ and $v \notin R$. Symmetric statements hold for SE, SW, and NW.*

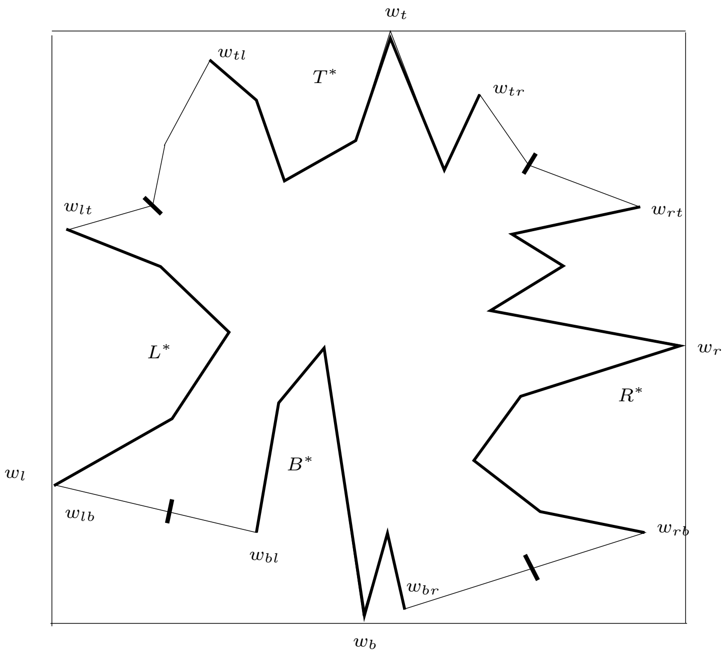

Figure 2.4: The partition of ∂S into the portions NE, SE, SW, and NW.

Proof: Suppose to the contrary that such a pair of points u, v exists. Clearly, we have $u \in R$ and $v \in T$. Hence the rightward-directed ray ρ_u emanating from u and the upward-directed ray ρ_v emanating from v are both openly disjoint from S. It is easily seen that these rays must cross each other. Indeed, it is impossible to draw the simple clockwise-directed connected polygonal chain NE, so that it starts at w_t, ends at w_r, lies below w_t and to the left of w_r, and passes first through u and then through v, so that the rays ρ_u and ρ_v are openly disjoint from NE and from each other. This is because such a drawing would yield a plane embedding of $K_{3,3}$, as is illustrated in Figure 2.5.

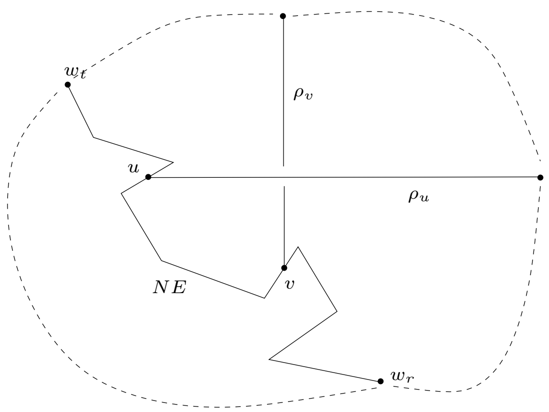

Figure 2.5: If ρ_u and ρ_v do not intersect, we obtain an impossible plane embedding of $K_{3,3}$.

Hence the two rays intersect, at some point z, as is illustrated in Figure 2.6. Let $\widetilde{P}_u = \widetilde{P} + u$, $\widetilde{P}_v = \widetilde{P} + v$, denote the placements of $\widetilde{P}$ with its reference point placed at u, v, respectively. Since $u, v \in \partial S$, $\widetilde{P}_u$ and $\widetilde{P}_v$ touch Q, but do not penetrate into it. Move $\widetilde{P}_u$ to the right until its reference point reaches z, and then move it down until the reference point reaches v. That is, the reference point traces the chain $J := \overline{uz} \cup \overline{zv}$, and the area swept by $\widetilde{P}$ during this motion is $P' := \widetilde{P} \oplus J$. By construction, P' and Q are openly disjoint. See Figure 2.7. By construction, J is both (weakly) x- and y-monotone. Thus P' is also x-monotone, since it is the Minkowski sum of two x-monotone polygons (see, e.g., [2]). Since P' and Q are openly disjoint, it follows from Theorem 2.1.1 (and the subsequent Remark (3)) that we can move P' to infinity along one of the four coordinate directions, without penetrating into Q. However, P' contains both $\widetilde{P}_u$ and $\widetilde{P}_v$, and thus both these polygons can be moved to infinity *in the same direction*, or, in other words, both u and v belong to the same subset of ∂S, which contradicts the facts that $u, v \notin B \cup L$, $u \notin T$ and $v \notin R$. The corresponding statements for SE, SW, and NW are proved in a fully symmetric manner. □

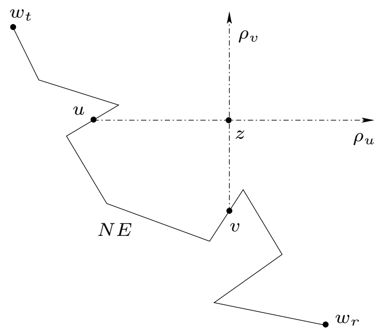

Figure 2.6: The configuration in Lemma 2.3.2.

Lemma 2.3.2 implies that we can partition NE into two openly disjoint connected subchains, T_{NE} and R_{NE}, with a common endpoint w, such that T_{NE} connects w_t to w and is contained in T, and R_{NE} connects w to w_r and is contained in R. Symmetrically, we obtain similar partitions $SE := R_{SE} \cup B_{SE}$, $SW := B_{SW} \cup L_{SW}$, and $NW := L_{NW} \cup T_{NW}$. The point w need not be unique. For example, if NE is monotone in both the x- and y-directions, *any* point along it can serve as the delimiter w. See Figure 2.4 where the delimiters w are highlighted; for NE, any point between w_{tr} and w_{rt} can serve as a delimiter, and similarly for the other three chains. We refer to the loci of the delimiters w as the *buffer zones* of NE, SE, SW, and NW. Set $T^* := T_{NW} \cup T_{NE} \subseteq T$, $R^* := R_{NE} \cup R_{SE} \subseteq R$, $B^* := B_{SE} \cup B_{SW} \subseteq B$, and $L^* := L_{SW} \cup L_{NW} \subseteq L$. Each of these four sets is connected, and they constitute the desired partition of ∂S, as asserted in Theorem 2.3.1. □

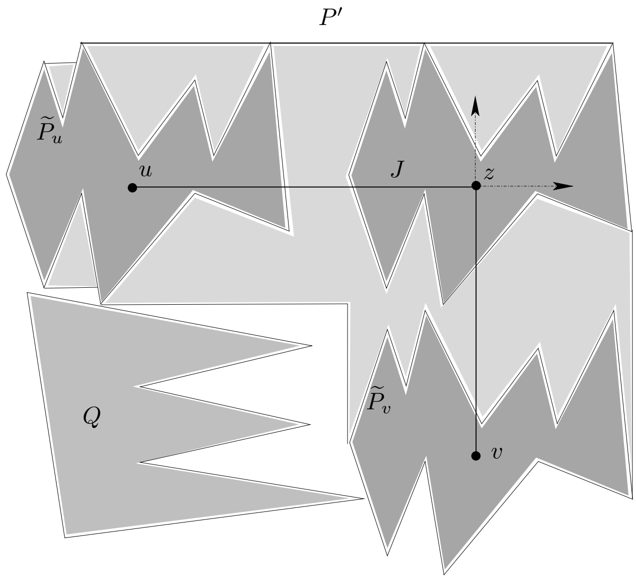

Figure 2.7: The swept polygon P', obtained as the area swept by $\widetilde{P}$ as it translates from u to v via z, is disjoint from Q.

2.4 Pockets in the Minkowski Sum of Monotone Polygons

We next bound the number of *pockets* in S. A *top pocket* is a maximal connected portion γ of T^* which is the concatenation of two connected portions α and β, such that (when proceeding in clockwise direction) α is monotone decreasing in y and monotone increasing in x, and β is monotone increasing in both x and y. Consequently, the common endpoint of α and β is locally y-minimal in T^*, and the two other endpoints of α, β are locally y-maximal. *Bottom pockets* in B^*, *left pockets* in L^*, and *right pockets* in R^* are defined in an analogous manner. See Figure 2.8. Note that the pockets are pairwise openly disjoint, and that their union is ∂S minus the four buffer zones of NE, SE, SW, and NW.

Theorem 2.4.1 *The number of pockets in $P \oplus Q$ is $O(m+n)$.*

Proof: Let $\gamma = \alpha \| \beta$ be a top pocket with lowest vertex v, incident to both α and β. In the interpretation of placing $\widetilde{P}$ around Q, v is a translation of $\widetilde{P}$ into a free placement at which one of the following situations arise:

Single contact: A vertex p of $\widetilde{P}$ that is locally y-maximal on (the bottom portion of) $\partial\widetilde{P}$ touches the unique y-maximal vertex of Q. See Figure 2.9(a).

Double contact: Two points p, p' of the bottom portion of $\partial\widetilde{P}$ touch two corresponding points q, q' of ∂Q that are locally top boundary points. Moreover, if we move $\widetilde{P}$ slightly to the left then it penetrates into Q in the vicinity of one of these contacts, and if we

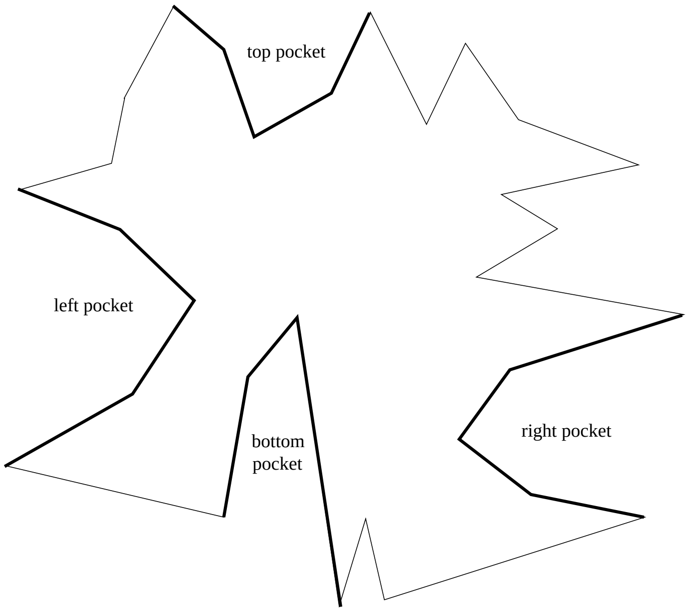

Figure 2.8: Pockets of ∂S.

move $\widetilde{P}$ slightly to the right then it penetrates into Q in the vicinity of the other contact. (It is easily checked that these penetrations cannot both occur in the vicinity of the same contact, because this would contradict the monotonicity of either P or Q.) See Figure 2.9(b,c).

A top pocket whose lowest point is generated by a single contact can be uniquely charged to the corresponding vertex p of $\widetilde{P}$, for a total of $O(m)$ such pockets. The same bound holds for bottom pockets of this kind, and the number of left and right pockets of this kind is $O(n)$. (The constants of proportionality in these bounds are smaller than 1.)

Consider next a top pocket whose lowest point v is generated by a double contact of two points p, p' on the bottom boundary of $\widetilde{P}$ with two corresponding points q, q' of ∂Q that are locally top boundary points. Assume, without loss of generality, that p lies to the left of p'. Since v is the lowest point of a pocket, we cannot move $\widetilde{P}$ to the left or to the right without immediately penetrating into Q. As noted above, one of the following two cases must arise:

Case A: As we move $\widetilde{P}$ slightly to the left, p penetrates into Q, and as we move it slightly to the right, p' penetrates into Q. See Figure 2.9(b).

Case B: As we move $\widetilde{P}$ slightly to the left, p' penetrates into Q, and as we move it slightly to the right, p penetrates into Q. See Figure 2.9(c).

We first claim that case A is impossible. Indeed, consider the portion of the bottom boundary of $\widetilde{P}$ at the placement v as the graph of a continuous function $y = f(x)$. Let $\bar{q}$ (resp., $\bar{q}'$) be a point in the interior of Q which coincides with p (resp., with p') as we move $\widetilde{P}$ slightly to the left (resp., to the right). Connect $\bar{q}$ and $\bar{q}'$ by a y-monotone polygonal chain within Q, which we regard as the graph of a continuous function $x = g(y)$. By construction,

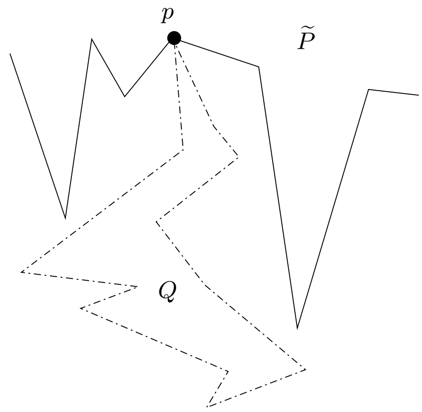

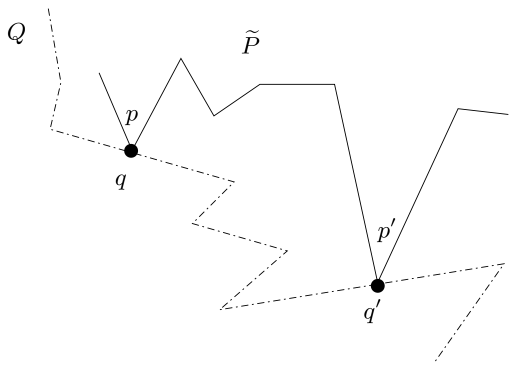

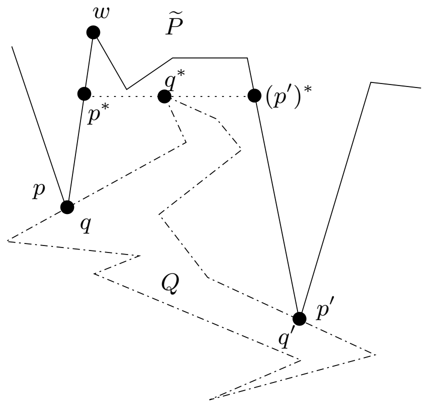

Figure 2.9: Three kinds of placements that correspond to the lowest point of a top pocket. (a) A single contact. (b) An impossible double contact. (c) A double contact, and the (pocket) vertex w of $\widetilde{P}$ being charged.

$g(y)$ has a point directly to the right of p' (namely, $\bar{q}'$), and a point directly to the left of p (namely, $\bar{q}$), which contradicts Lemma 2.1.2, thus showing that case A is impossible.

In case B, the local penetrations of $\widetilde{P}$ into Q in the vicinities of the two pairs of coincident points $p = q$ and $p' = q'$ (at the placement v of $\widetilde{P}$) imply that the following property holds: Separate $\widetilde{P}$ and Q locally near $p = q$ by a line ℓ, and locally near $p' = q'$ by a line ℓ'. Then ℓ has a positive slope, $\widetilde{P}$ lies locally to its left and above it, and Q lies locally to its right and below it. Symmetrically, ℓ' has a negative slope, $\widetilde{P}$ lies locally to its right and above it, and Q lies locally to its left and below it. This implies, arguing as in Section 2.1, that the top vertex q^* of Q lies above p and p' and below the portion γ of the graph $y = f(x)$ that connects p and p' (as defined in case A). Let p^* (resp., $(p')^*$) denote the closest point on γ that lies directly to the left (resp., to the right) of q^*; the preceding analysis implies that both points exist. It follows that the global maximum of $y = f(x)$ between p^* and $(p')^*$ must occur at an interior point w, which is the highest point of a bottom pocket of $\widetilde{P}$

(alternatively, the lowest point of a top pocket of P). See Figure 2.9(c).

We charge the top pocket of v to w, and claim that *this charging is unique.* Indeed, suppose to the contrary that another top pocket is also charged to w. Let v_1 denote its lowest point, and let the corresponding double contact be determined by points p_1, p_1' of $\widetilde{P}$ with corresponding points q_1, q_1' of Q.

It is more convenient for this stage of the analysis to regard $\widetilde{P}$ as the stationary and Q as the translating polygon. We thus have two placements of Q, which for simplicity we denote as Q and Q_1. The stationary bottom boundary of $\widetilde{P}$ contains the five points p, p_1, w, p', p_1', so that p and p_1 lie to the left of w, p' and p_1' lie to the right of w, and w lies above the four other points. The polygon Q touches $\widetilde{P}$ at the two points p, p', which coincide with the respective points $q, q' \in Q$, and the polygon Q_1 touches $\widetilde{P}$ at the two points p_1, p_1', which coincide with the respective points $q_1, q_1' \in Q_1$. Let t be the translation vector that satisfies $Q_1 = Q + t$. Since the situation is symmetric in Q and Q_1 we may assume, withot loss of generality, that t has a positive x-component, and consider two cases, depending on the sign of the y-component of t. (The general position assumption allows us to assume that t is not horizontal.)

Suppose first that t has a positive y-component. Refer to Figure 2.10. Since case B applies to both Q and Q_1, there exists a point $\tilde{q} \in Q$ directly to the left and arbitrarily close to p'. Hence, $\tilde{q}_1 = \tilde{q} + t \in Q_1$ lies to the right and above p'. This implies that the highest point q_1^* of Q_1 lies above and to the right of p'. Indeed, it clearly lies above p'. Suppose to the contrary that it lied to the left of p'. By the preceding arguments, q_1^* lies below the graph of $y = f(x)$ and has a point on that graph directly to its right. By continuity, the closest such point must lie to the left of p'. This, combined with the fact that $\tilde{q}_1 = \tilde{q} + t \in Q_1$ lies to the right and above p', is easily seen to contradict Lemma 2.1.2, thus showing that q_1^* lies to the right of p'.

It follows that the two closest points $p_1^*, (p_1')^*$ on $y = f(x)$ that lie directly to the left and to the right of q_1^* both lie to the right of p'. But then the pocket associated with Q_1 should have charged a vertex of $\widetilde{P}$ that lies between p_1^* and $(p_1')^*$, and thus lies to the right of p', so it could not have charged w.

The case where t has a negative y-component is argued in a fully symmetric fashion. Using the point p_1 instead of p', we show that the highest point q^* of Q is 'trapped' in a pocket of $\widetilde{P}$ that lies fully to the right of p_1, and thus the pocket associated with the placement Q cannot charge w. This shows that w can be charged at most once, as asserted.

To sum up, we have shown that the top pockets of S that are generated by double contacts can be uniquely charged to top pockets of P. Symmetrically, bottom pockets of S of this kind can be uniquely charged to bottom pockets of P. Yet another symmetric argument, in which the roles of P and Q are interchanged, shows that left (resp., right) double-contact pockets of S can be uniquely charged to left (resp., right) pockets of Q. Adding the bound on the number of single-contact pockets, we conclude that the total number of pockets of S is $O(m + n)$. □

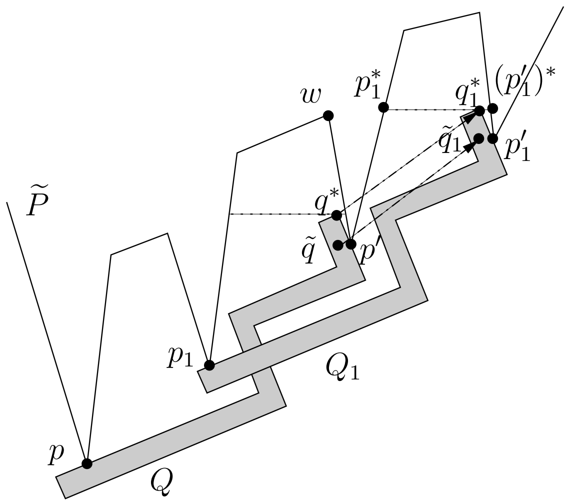

Figure 2.10: Illustrating the proof that two top pockets of S of type B cannot both charge the same pocket of P.

2.5 Minkowski Sum of Non-monotone Simple Polygons

The preceding machinery allows us to derive the main result of this paper:

Theorem 2.5.1 *Let P be a simple polygon with m edges, which can be decomposed into k simple subpolygons, all monotone in the x-direction, and let Q be a simple polygon with n edges, which can be decomposed into ℓ simple subpolygons, all monotone in the y-direction. Then the complexity of $P \oplus Q$ is $O(k\ell mn\alpha(\min\{m,n\}))$. The same holds if the x- and y-directions are replaced by two arbitrary directions.*

Proof: Let $P_1, \ldots, P_k$ be the k subpolygons in the decomposition of P, and let $Q_1, \ldots, Q_\ell$ be the ℓ subpolygons in the decomposition of Q. Let m_i denote the number of edges of P_i, for $i = 1, \ldots, k$, and let n_i denote the number of edges of Q_i, for $i = 1, \ldots, \ell$. We have $\sum_{i=1}^{k} m_i = O(m)$ and $\sum_{i=1}^{\ell} n_i = O(n)$.

Put $S_{ij} := P_i \oplus Q_j$, for $i = 1, \ldots, k$ and $j = 1, \ldots, \ell$. Clearly, $S = P \oplus Q = \bigcup_{i,j} S_{ij}$. By Theorem 2.2.1, the complexity of S_{ij} is $O(m_i n_j \alpha(\min\{m_i, n_j\}))$. Hence, the total number of edges of all the sums S_{ij} is

$$O\left(\left(\sum_{i=1}^{k}\sum_{j=1}^{\ell} m_i n_j\right)\alpha(\min\{m,n\})\right) = O(mn\alpha(\min\{m,n\})).$$

For each i and j, let T_{ij}^*, B_{ij}^*, L_{ij}^* and R_{ij}^* denote the four connected portions of ∂S_{ij}, as provided by Theorem 2.3.1. Let X denote the collection of all the chains T_{ij}^* and B_{ij}^*, and let Y denote the collection of all the chains L_{ij}^* and R_{ij}^*. X is a set of $2k\ell$ x-monotone polygonal chains. The number of intersections of any pair of such chains is proportional to the number of their edges, which is easily seen to imply that the complexity of the arrangement $\mathcal{A}(X)$ is $O(k\ell mn\alpha(\min\{m,n\}))$. Similarly, the complexity of $\mathcal{A}(Y)$ is also $O(k\ell mn\alpha(\min\{m,n\}))$.

The complement of S is the union of some faces of the arrangement $\mathcal{A}(X \cup Y)$. Let H denote the collection of these faces. H contains one (the unique) unbounded face, and the rest are bounded faces ('holes' of S). By the Combination Lemma for planar arrangements (see [24]), the overall complexity of all the faces of H (that is, the complexity of S) is proportional to the complexity of $\mathcal{A}(X)$ plus the complexity of $\mathcal{A}(Y)$ plus $|H|$. Hence, Theorem 2.5.1 is an immediate consequence of the following lemma.

Lemma 2.5.2 *The number of holes of $P \oplus Q$ is $O(k\ell mn\alpha(\min\{m,n\}))$.*

Proof: Let f be a bounded hole in H. If ∂f contains a vertex of either $\mathcal{A}(X)$ or $\mathcal{A}(Y)$, we charge f to that vertex, and thus conclude that the number of such holes is $O(k\ell mn\alpha(\min\{m,n\}))$. Otherwise, f is a convex polygon, whose boundary consists of a sequence of edges, alternating between edges of $\mathcal{A}(X)$ and edges of $\mathcal{A}(Y)$. Clearly, f has an even number of edges.

Let v be the lowest vertex of f, and suppose that it is incident to an edge e of $\mathcal{A}(X)$ and to an edge e' of $\mathcal{A}(Y)$. Suppose, without loss of generality, that e bounds f to the left of v and that e' bounds f to the right of v; see Figure 2.11. In this case e is (a portion of) an edge of some T^*_{ij} and e' is (a portion of) an edge of some $L^*_{i'j'}$. Clearly, f is a portion of a face f_0 of the arrangement $\mathcal{A}(T^*_{ij} \cup L^*_{i'j'})$, and v is a local y-minimum of f_0. (Note that the case $i = i'$, $j = j'$ is impossible, because T^*_{ij} cannot meet L^*_{ij} in such a way.)

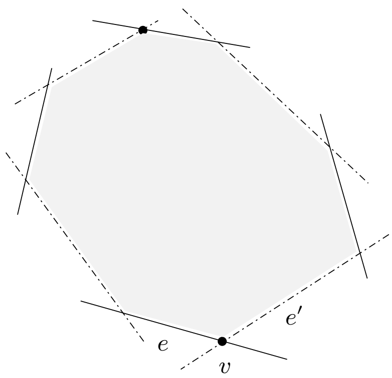

Figure 2.11: A convex hole in H.

A simple application of Morse theory to f_0 shows that, if f_0 is not y-monotone, then the number of local y-minima of f_0 is proportional to the number of points of ∂f_0 which are local y-extrema of the complement of f_0 (i.e., reflex locally y-extremal vertices of ∂f_0). See, e.g., [12, Lemma 2.4] for a similar argument. Any such point u is a local y-extremal vertex of either T^*_{ij} or $L^*_{i'j'}$. The latter chain has only two such vertices, and the number of such vertices on the former chain T^*_{ij} is 1 plus the number of top pockets of S_{ij}. Hence, this number is $O(m_i)$, by Theorem 2.4.1. We repeat this argument to all the faces of $\mathcal{A}(T^*_{ij} \cup L^*_{i'j'})$ which are not y-monotone, to all other combinations of sub-boundaries of S_{ij} and $S_{i'j'}$, and to all combinations of i, j, i', j', to conclude that the overall number of holes f that satisfy all the above conditions is

$$O\left(\sum_{i,j,i',j'}(m_i+m_{i'})\right)=O(mk\ell^2)=O(k\ell mn).$$

Suppose then that f_0 is y-monotone. Then f_0 has a unique y-minimal point (namely, v). If ∂f_0 contains a vertex of either S_{ij} or of $S_{i'j'}$ then we charge v (uniquely) to such a vertex. Summing over all such faces f_0 and over all i,j,i',j', we conclude that the number of holes f that fall into this subcase is

$$O\left(\sum_{i,j,i',j'}(m_in_j+m_{i'}n_{j'})\alpha(\min\{m,n\})\right)=O(k\ell mn\alpha(\min\{m,n\})).$$

We may thus assume that f_0 is convex and bounded, and that the edges of its boundary alternate between (portions of) edges of T^*_{ij} and (portions of) edges of $L^*_{i'j'}$. No edge of the top boundary of f_0 can belong to T^*_{ij}, and thus the top boundary of f_0 consists of a *single edge* of $L^*_{i'j'}$. Similarly, the left boundary of f_0 consists of a single edge of T^*_{ij}. We distinguish between two cases:

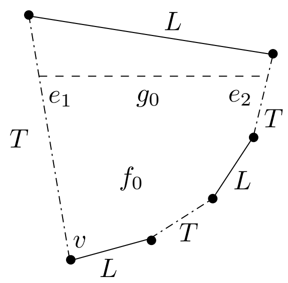

Figure 2.12: A convex hole f_0 of $T^*_{ij}\cup L^*_{i'j'}$ with more than four edges, and the pair (e_1,e_2) of edges to which f_0 is charged.

Case A: f_0 has more than four edges. See Figure 2.12. Note that the unique left edge e_1 of f_0 spans the entire y-range of the hole. We can therefore connect e_1 to the highest right edge e_2 of f_0 (excluding the top edge of f_0) by a horizontal segment g_0, as shown in Figure 2.12. Note that e_2 is also (a portion of) an edge of T^*_{ij}. Let $\bar{e}_1$ and $\bar{e}_2$ denote the edges of T^*_{ij} that contain, respectively, e_1 and e_2. We draw in the plane a graph G whose vertices are the edges of T^*_{ij}, which are drawn as they are. The edges of G are all the pairs $(\bar{e}_1,\bar{e}_2)$ obtained from the faces f_0 that fall into this subcase, and are drawn in the plane as the above horizontal connecting segments g_0. Clearly, G is planar.

We claim that G is simple. Indeed, suppose to the contrary that there exist two faces f_0,f_1 that cause the same pair of edges $\bar{e}_1,\bar{e}_2$ of T^*_{ij} to be connected by two respective horizontal segments g_0,g_1. Suppose, without loss of generailty, that g_0 lies higher than g_1, and refer to Figure 2.13. Note that f_0 has at least one additional right edge e_3 that is contained in an edge $\bar{e}_3$ of T^*_{ij}. The x-monotonicity of T^*_{ij} implies that the entire edge $\bar{e}_3$ must be contained

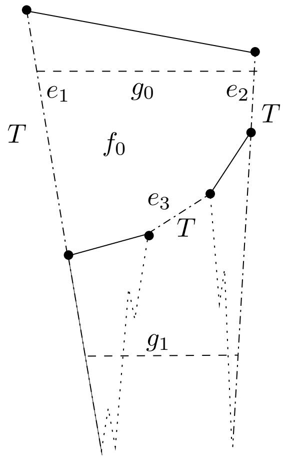

Figure 2.13: Illustrating the proof that G is simple.

in the upper wedge defined by the lines containing $\bar{e}_1$ and $\bar{e}_2$. Therefore, $\bar{e}_3$ appears along T^*_{ij} between $\bar{e}_1$ and $\bar{e}_2$. Moreover, $\bar{e}_3$ lies above g_1, and the right endpoint of $\bar{e}_1$ lies below g_1. It follows that the portion of T^*_{ij} between $\bar{e}_1$ and $\bar{e}_3$ must cross g_1, which is impossible. This contradiction shows that G is simple. Hence the number of edges of G, and thus the number of faces f_0 of the type under consideration, is proportional to the number of edges of T^*_{ij}. Summing this bound over all i, j, i', j', we conclude that the overall number of holes f that give rise to faces f_0 of type (A) is $O(k\ell mn\alpha(\min\{m, n\}))$.

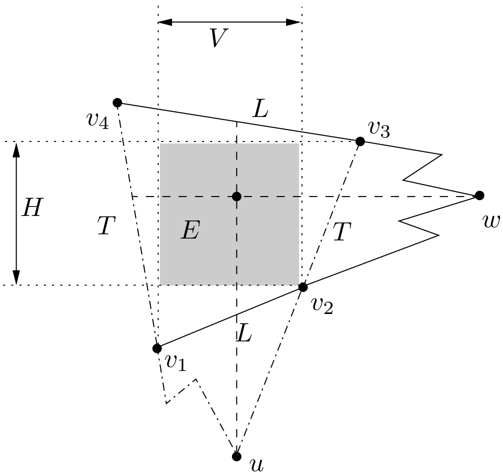

Figure 2.14: A convex quadrangular hole f_0 of $T^*_{ij} \cup L^*_{i'j'}$, and the pair of pockets to which f_0 is charged.

Case B: f_0 is a quadrilateral. See Figure 2.14. Let v_1, v_2, v_3, v_4 denote the vertices of f_0 in counterclockwise order, starting from the bottom vertex $v = v_1$. Let V denote the vertical strip between v_1 and v_2, and let H denote the horizontal strip between v_2 and v_3. It follows that the rectangle $E = V \cap H$ is fully contained in f_0. See Figure 2.14. Since the left edge of

f_0 has negative slope, and the right edge (which belongs to T^*_{ij}) has positive slope, it follows that the portion of T^*_{ij} between these edges must contain at least one pocket, and that its lowest vertex u lies in V. Similarly, the portion of $L^*_{i'j'}$ between the bottom and the top edges of f_0 must contain at least one pocket, and its rightmost vertex w lies in H. Hence, the vertical line through u and the horizontal line through w meet inside E, and thus inside f_0. We can therefore charge f_0 to the pair of pockets of T^*_{ij} and of $L^*_{i'j'}$ that are associated with u and w, respectively, and conclude that any such pair of pockets is charged at most once. Hence, by Theorem 2.4.1, the number of faces f_0 of type (B) is $O(m_i n_{j'})$. Summing this bound over all i, j, i', j', we conclude that the overall number of holes f that give rise to faces f_0 of type (B) is $O(k\ell mn)$.

Thus the number of holes of $P \oplus Q$ is $O(k\ell mn\alpha(\min\{m, n\}))$. This completes the proof of Lemma 2.5.2, and thus also the proof of Theorem 2.5.1. $\square$

Remarks: (1) As already remarked, the bound in Theorem 2.5.1 is slightly suboptimal when $k = \Theta(m)$ and $\ell = \Theta(n)$, which is the case of arbitrary simple polygons P and Q. In this case the worst case tight bound is $O(m^2n^2) = O(k\ell mn)$ [5]. It would be interesting to finetune our analysis so as to make our bound equal to this bound in the general case.

(2) The subpolygons $P_1, \ldots, P_k$ in the decomposition of P need not be pairwise openly disjoint, and the theorem continues to hold provided that the overall number of their edges is still $O(m)$. A similar extension applies to the decomposition of Q.

2.6 Conclusion

We have presented a general technique for analyzing the complexity of the Minkowski sum of two simple polygons, expressed in terms of their partition into monotone pieces. We have derived some interesting properties of the Minkowski sum of two polygons, each monotone in a different direction. Our bound nicely interpolates between the special case of two monotone polygons and the general case of two arbitrary simple polygons. It is worst case tight in the former case, and nearly so in the latter case.

Bibliography

[1] P.K. Agarwal, E. Flato and D. Halperin, Polygon decomposition for efficient construction of Minkowski sums, *Comput. Geom. Theory Appl.* 21:39–61, 2002.

[2] A. H. Barrera, Computing the Minkowski sum of monotone polygons, *IEICE Trans. Inf. & Syst.*, E80-D (2):218–222, 1997.

[3] A. H. Barrera, Finding an $O(n^2 \log n)$ algorithm is sometimes hard, *In Proc. 8th Canad. Conf. Comput. Geom.*, pages 289–294. Carleton University Press, Ottawa, Canada, 1996.

[4] B. Chazelle, The polygon containment problem, In F. P. Preparata, editor, *Computational Geometry*, volume 1 of *Adv. Comput. Res.*, pages 1–33. JAI Press, London, England, 1983.

[5] M. de Berg, M. van Kreveld, M. Overmars and O. Schwarzkopf, *Computational Geometry, Algorithms and Applications*, 2nd Edition, Springer Verlag, Berlin-Heidelberg, 2000.

[6] M. de Berg and A. van der Stappen, On the fatness of minkowski sums, *Technical Report UUCS199939*, Dept. of Computer Science, Utrecht University, 1999.

[7] G. Elber and M.S. Kim, editors, *Special Issue of Computer Aided Design: Offsets, Sweeps and Minkowski Sums*, 31, 1999.

[8] E. Flato, *Robust and efficient construction of planar Minkowski sums*, M.Sc. Dissertation, School of Computer Science, Tel Aviv University, 2000.

[9] E. Flato homepage. http://www.math.tau.ac.il/∼flato.

[10] E. Flato and D. Halperin, Robust and efficient construction of planar Minkowski sums, in *Abstracts 16th European Workshop on Computational Geometry*, Eilat, 2000, pages 85–88.

[11] D. Halperin, J.C. Latombe, and R. H. Wilson, A general framework for assembly planning: The motion space approach, *Algorithmica*, 26:577–601, 2000.

[12] D. Halperin and M. Sharir, Almost tight upper bounds for the single cell and zone problems in three dimensions, *Discrete Comput. Geom.* 14:385–410, 1995.

[13] S. Har-Peled, Multicolor combination lemmas, *Comput. Geom. Theory Appl.* 12:155–176, 1999.

[14] S. Har-Peled, T. M. Chan, B. Aronov, D. Halperin, and J. Snoeyink, The complexity of a single face of a Minkowski sum, *In Proc. 7th Canad. Conf. Comput. Geom.*, pages 91–96, 1995.

[15] E. Hartquist, J. Menon, K. Suresh, H. Voelcker, and J. Zagajac, A computing strategy for applications involving offsets, sweeps, and Minkowski operations, *Comput. Aided Design*, 31(4):175–183, 1999.

[16] A. Kaul, M. A. O'Connor, and V. Srinivasan, Computing Minkowski sums of regular polygons, *In Proc. 3rd Canad. Conf. Comput. Geom.*, pages 74–77, 1991.

[17] L. E. Kavraki, Computation of configuration space obstacles using the Fast Fourier Transform, *IEEE Trans. Robot. Autom.*, 11:408–413, 1995.

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

[19] J.C. Latombe, Robot Motion Planning, *Kluwer Academic Publishers*, Boston, 1991.

[20] R. Pollack, M. Sharir, and S. Sifrony, Separating two simple polygons by a sequence of translations, *Discrete Comput. Geom.*, 3:123–136, 1988.

[21] G. Pustylnik and M. Sharir, The Minkowski sum of a simple polygon and a segment, *Inform. Process. Lett.* 85:179–184, 2003.

[22] G.D. Ramkumar, An algorithm to compute the Minkowski sum outer face of two simple polygons, *Proc. 12th Annu. ACM Sympos. Comput. Geom.*, 234-241, 1996.

[23] M. Sharir, Efficient algorithms for planning purely translational collision-free motion in two and three dimensions, *Proc. IEEE Symp. on Robotics and Automation* (1987), 1326–1331.

[24] M. Sharir and P.K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications,* Cambridge University Press, Cambridge-New York-Melbourne, 1995.

[25] G.T. Toussaint and H.A. ElGindy, Separation of two monotone polygons in linear time, *Robotica* 2:215–220, 1984.

[26] M. van Kreveld, Twelve computational geometry problems from cartographic generalization, Manuscript. Presented at the Dagstuhl meeting on Computational Geometry, March, 1999.