

Sweeping Minimum Perimeter Enclosing Parallelograms: Optimal Crumb Cleanup

Yonit Bousany*

Mary Leah Karker†

Joseph O’Rourke‡

Leona Sparaco§

Abstract

We examine the problem of pushing all the points of a planar region into one point using parallel sweeps of an infinite line, minimizing the sum of the lengths of the sweep vectors. We characterize the optimal 2-sweeps of triangles, and provide a linear-time algorithm for convex polygons.

1 Introduction

A fascinating problem was introduced by Dumitrescu and Jiang in [DJ09]:¹ What is the optimal way to sweep points in the plane all to one target point? In their model, the sweeper is an infinite line that moves orthogonal and parallel to itself, pushing points along the sweep direction. The cost of a sweep is the distance swept (with no cost for how many points are pushed at once, as in the “earth mover’s distance”). The cost of multiple sweeps is the sum of each individual sweep’s cost, with no cost assessed for repositioning in preparation for the next sweep. The problem is to sweep all the points into one point, selected by the algorithm.

The sweep line may be viewed as a “table crumb” or “crumb scraper,” the instrument waiters use in restaurants to clean the table cloth between courses.

In [DJ09], several variants of the problem are considered, most involving sweeping a discrete set of points. Here we explore another natural variant. In our model the crumbs form a continuous simply connected region of the plane, for example, a simple polygon. We use a different sweep model, where the infinite sweep line L may translate by any vector v , not only by vectors orthogonal to L as in [DJ09]. Each point swept moves by v , with the sweep cost $|v|$. This model is appropriate if the crumbs have sufficient friction to avoid sliding along the sweeper. Our goal is to find the optimal sweep strategy and its cost for a given shape.

Our contributions are as follows:

1. We show that sometimes more than two sweeps are

needed to optimally sweep nonconvex shapes, but we conjecture that two sweeps always suffice for convex shapes.

2. Restricting attention to two sweeps, the best two sweeps are determined by the minimum perimeter enclosing parallelogram.
3. We characterize the optimal two sweeps for triangles.
4. We provide a linear-time algorithm for finding the minimum perimeter enclosing parallelogram (and so the optimal two sweeps) for any convex n -gon.

Concerning this last point, there has been considerable work on both minimal area and minimal perimeter enclosures, all using some variant of Toussaint’s “rotating calipers” algorithm for finding the minimum area enclosing rectangle [Tou83]. Two aspects of these algorithms are crucial for efficiency: determining if one or more edges of the enclosure must be flush with the enclosed hull, and an “interspersing” lemma that avoids backtracking in the search. In general, establishing these conditions has proven more difficult for minimum perimeter than for minimal area. Only relatively recently has a linear-time algorithm for finding a minimum perimeter triangle appeared [BM03], and it is quite intricate. It was followed by a polynomial-time algorithm for minimum perimeter k -gons in [MP08].

It appears that explicit attention has only been paid to enclosing parallelograms when motivated by an application. Image compression drove the algorithm in [STV⁺95] for finding a minimum area enclosing parallelogram. Here two adjacent edges are flush and interspersing is straightforward. Our whimsical “application” led to our investigation of minimal perimeter parallelograms. That one edge must be flush was established in [MP08]. Our contribution here is establishing an interspersing lemma (Lemma 10), allowing a rotating calipers algorithm to achieve linear time.

2 Multiple Sweeps

Lemma 1 (1-Sweep) *One sweep suffices if and only if all the points swept are collinear, in which case the cost is the length of the shortest segment containing all the points.*

*ybousany@smith.edu.

†mkarker@wesleyan.edu.

‡Dept. of Computer Science, Smith College, Northampton, MA 01063, USA. orourke@cs.smith.edu.

§lsparaco@uvm.edu.

¹They credit Paweł Żyliński as the originator.

Proof. All points swept move parallel to the sweep, and so can only end up at one target point if they lie along a line. \square

Lemma 2 (2-Sweeps) *The optimal 2-sweep strategy for any shape S is to sweep parallel to the edges of the minimum perimeter parallelogram enclosing S , with cost half the perimeter.*

Proof. By Lemma 1, the first sweep must push S to a line segment. Two sweeps then determine a parallelogram, and the cost under our model is the sum of the lengths of two consecutive edges. \square




Figure 1: Two ways to sweep an equilateral triangle.

Example. Figure 1 shows two ways to sweep a unit equilateral triangle. In (a) two sweeps are used: down the $\sqrt{3}/2$ altitude, and then left across the base. In (b) four sweeps are used, the first two using $1/(2\sqrt{3})$ and $1/\sqrt{3}$ to produce a $\frac{1}{3}$ by $\frac{2}{3}$ parallelogram, which is then reduced to a point with two more sweeps. The total cost for both strategies is $1 + \sqrt{3}/2$. Indeed, we have found many distinct strategies for this shape (including infinite series of sweeps) that all result in the same cost, which we conjecture to be optimal. An argument in [DJ09] gives a lower bound of $\sqrt{3}$, smaller by ≈ 0.13 .




Figure 2: Three sweeps can be better than two.

Nonconvex Shapes. Figure 2 shows a simple example where three sweeps are superior to two sweeps. The 3-sweeps illustrated cost $\sqrt{2} + 2 \approx 3.41$ but the 2-sweep cost (using Lemmas 2 and 4 below) is $\sqrt{5} + 1/2 \approx 3.73$. Henceforth we study two sweeps:

Conjecture 3 *The optimal cost of sweeping any convex shape S can be achieved by two sweeps.*

3 One Flush, Two Flush

Restricting attention to two sweeps reduces the problem via Lemma 2 to finding a minimum perimeter enclosing parallelogram. Lemma 1 in [MP08] implies this:

Lemma 4 (1-Flush) *A minimum perimeter parallelogram enclosing a polygon Q has one edge flush with the convex hull of Q .*

They prove a somewhat different (and more general) result, but when specialized to parallelograms, the above is implied by their proof.

We now further restrict attention to triangles. Obviously for triangles, either just one edge is flush, or two edges are flush. We characterize precisely the conditions separating these two cases:

Theorem 5 *Let T be a triangle, and P its minimum perimeter enclosing parallelogram. If T is obtuse, P is flush with the two edges forming the obtuse angle. If T is acute, P is a rectangle flush with the shortest edge.*

Normalize the triangle T so that its longest side has length 1; let the other two sides have lengths $a \leq b \leq 1$. We can then view the space of all triangles as points (a, b) . Figure 3 illustrates the import of the theorem. Although the lower diagonal half of this diagram suffices to represent all triangles (a, b) , we find it intuitively helpful to display the full square. We must have




Figure 3: Theorem 5 in the (a, b) space.

$a + b \geq 1$ to satisfy the triangle inequality. Let θ be the angle at the ab -vertex. All points along the circular arc have $\theta = 90^\circ$, those below are obtuse, those above, acute. The triple point $(a, b) = (\frac{\sqrt{2}}{2}, \frac{\sqrt{2}}{2})$ represents the $45^\circ-45^\circ-90^\circ$ triangle where 1-flush against a is also 1-flush against b and so is 2-flush against a and b .

Proof. Observe that $a + b \leq a + 1 \leq b + 1$. Therefore 2-flush against the two shorter sides always beats the other possible 2-flush cases.

We analyze the competing cases illustrated in Figures 4 and 5.




Figure 4: θ Acute: 1-flush against the shortest side wins.

θ Acute. That the 1-flush rectangles illustrated in Figure 4 are the appropriate cases to consider is not justified until Corollary 9. Denote by h_a the height of the parallelogram that is flush against the shortest side a .

Claim 1: 1-flush against shortest side a beats all 2-flush cases (bottom row of Figure 4).

Because we have a right triangle, $h_a < b$. Thus $h_a + a < a + b$. So 1-flush against a beats 2-flush against a and b and hence beats all 2-flush cases.

Claim 2: 1-flush against shortest side a beats 1-flush against longest side.

Let h_1 be the height of the parallelogram flush against the longest side. Because we have a right triangle, $h_a < 1$. Therefore :

$$\begin{aligned} h_a(1-a) &< (1-a) \\ h_a - h_a a &< 1-a \\ h_a + a &< 1 + h_a + a \end{aligned}$$

Computing the triangle area two ways yields $\frac{1}{2} \cdot 1 \cdot h_1 = \frac{1}{2} \cdot h_a a$, and therefore $h_1 = h_a a$. Thus $h_a + a < 1 + h_1$.

Claim 3: 1-flush against shortest side a beats 1-flush against median-length side b .

Define h_b as the height of the parallelogram flush against side b . Since $a < b$, it follows that

$$\begin{aligned} a(1 - h_a/b) &< b(1 - h_a/b) \\ a - ah_a/b &< b - bh_a/b < b + ah_a/b \end{aligned}$$

Expressing the area of the triangle two ways yields $\frac{1}{2}ah_a = \frac{1}{2}bh_b$, and hence $ah_a/b = h_b$. Thus $a + h_a < b + h_b$.

Therefore 1-flush against the shortest side a wins in all cases.




Figure 5: θ Obtuse: 2-flush against the shorter sides wins.

θ Obtuse. Lemma 7 below shows that the only 1-flush case that is relevant is 1-flush against the longest side, because 1-flush against the shorter sides reduces to 2-flush against the shorter sides (in the notation of that lemma, $\Delta y = 1$).

Claim: 2-flush against shorter sides beats 1-flush against the longest side. See Figure 5. The 2-flush sweeps cost $a + b$, and the 1-flush sweeps cost $1 + h_1$, so we aim to prove

$$1 + h_1 \geq a + b.$$

Expressing the area of the triangle in two ways, $\frac{1}{2}1 \cdot h_1 = \frac{1}{2}ab \sin \theta$, leads to $h_1 = ab \sin \theta$. Substituting and squaring, our new goal is

$$(1 + ab \sin \theta)^2 \geq (a + b)^2.$$

Expanding yields

$$1 + 2ab \sin \theta + (ab \sin \theta)^2 - (a^2 + b^2 + 2ab) \geq 0.$$

The Law of Cosines implies that

$$1^2 = a^2 + b^2 - 2ab \cos \theta.$$

Substituting for 1 in our inequality gives

$$a^2 + b^2 - 2ab \cos \theta + 2ab \sin \theta + (ab \sin \theta)^2 - a^2 - b^2 - 2ab \geq 0,$$

which simplifies to

$$2ab(\sin \theta - \cos \theta - 1) + (ab \sin \theta)^2 \geq 0.$$

Now, because $(\sin \theta - \cos \theta) \geq 1$ throughout the obtuse range $\theta \in [\pi/2, \pi]$, the first term is nonnegative. And because $\sin \theta \geq 0$, the second term is as well, and we have established the inequality.²

²The proof of this case is due to Anastasia Kurdia, significantly simpler than our original proof. A referee also suggested another simple proof.

In terms of the viewpoint of Figure 3, the transition between obtuse and acute θ is the quarter-circle arc. Let Γ be the locus of (a, b) -points representing triangles where the 2-flush and 1-flush instances in Figure 5 have identical perimeter. Γ is an algebraically complicated curve, shown in Figure 6. The inequality just established implies that Γ falls outside the acute/obtuse transition throughout its length. For example, 1-flush




Figure 6: Points on $\Gamma = b(a)$ represent acute θ .

for isosceles triangles does not start to beat 2-flush until $\theta \leq 2 \cos^{-1} \frac{4}{5} \approx 74^\circ$. The reason Γ does not appear in Figure 3 is that the 1-flush cases against the shortest side dominate. □

Corollary 6 *The minimum perimeter parallelogram enclosing a random triangle is twice as likely to have two sides flush than it is to have just one side flush.*

Proof. In the model that selects the three triangle vertices uniformly and randomly on a circle, it is well known that the probability of an obtuse angle is $\frac{3}{4}$. □

4 Algorithm for Convex Polygons

With the 1-flush lemma (Lemma 4) in hand, an outline of an algorithm to find the best parallelogram P enclosing a convex n -gon C is easy: for each edge e of C selected as being flush with P , find the optimal orientation of the other side of P . We will show how this best second side orientation can be found in $O(n)$ time, then in $O(\log n)$ time, and finally in $O(1)$ amortized time, leading to successively faster algorithms, culminating in a linear-time algorithm.

First we consider a situation with the 1-flush side of P fixed as horizontal, and the free side of P pivoting on two points, ignoring the actual shape of C between the pivots, as illustrated in Figure 7(a). The minimum




Figure 7: The minimum perimeter parallelogram is achieved with $\theta = \cos^{-1}(0.2) \approx 78^\circ$. Here $(\Delta x, \Delta y) = (0.2, 1.6)$.

perimeter parallelogram is determined solely as a function of the vertical separation of the pivots. Let the pivots be separated by $(\Delta x, \Delta y)$ when the height of the parallelogram is normalized to 1, as in the figure. The relevant range for θ is $(\theta_{\min}, \theta_{\max})$ where

$$\begin{aligned} \theta_{\min} &= \max\{0, \tan^{-1}(\Delta y/\Delta x)\} \\ \theta_{\max} &= \min\{\pi, \pi + \tan^{-1}(\Delta y/\Delta x)\}. \end{aligned}$$

These limits serve to keep the polygon (which must include the segment connecting the pivots) inside a positively oriented parallelogram. For example, in the figure, the range is $(7^\circ, 180^\circ)$.

Lemma 7 *For $\theta \in (\theta_{\min}, \theta_{\max})$, the perimeter $\rho(\theta)$ of P has a unique minimum at the value $\theta^* = \cos^{-1}(\Delta y)$.*

Proof. Letting the heights of the two pivots be y_1 and $y_1 + \Delta y$, $\rho(\theta)$ can be expressed as a function of $\{y_1, \Delta y, \Delta x, \theta\}$. The derivative of $\rho(\theta)$ solves explicitly to the claimed expression, independent of y_1 and Δx . □

The lemma is illustrated in Figure 7(b). Our proof of this lemma is via explicit computation of a complex expression, but the succinctness of the solution suggests there may be a simple geometric proof.

Now imagine rotating “calipers” around the convex polygon C , with angle varying from $\theta = 0$ with respect to the x -axis and flush edge, to $\theta = 180^\circ$. The left pivot walks down the left side of C while the right pivot walks up the right side of C . For each θ , the pair of pivots determine a Δy value (right minus left vertical height), which determines the optimal angle θ^* via Lemma 7. But for many pivots, the line at this angle is not tangent to C , but rather penetrates C . Figure 8 illustrates this on the left side. When the θ^* value penetrates C at either pivot, it means that the caliper tangents determining those pivots do not determine a minimum perimeter parallelogram: we are on the downslope in Figure 7(b) and the perimeter can be decreased by moving to another pivot. When the θ^* lines are also tangent to C at both pivots, we are at the unique local minimum of the perimeter. A second circumstance that determines the




Figure 8: Tangent lines at angle θ determine pivots, and pivots determine θ^* values. Here it is only the last $\theta^* = 67^\circ$ that is also tangent to the determining pivot.

minimum perimeter is when θ^* penetrates C below at one pivot and penetrates C above at the next pivot, in which case the minimum is determined by the angle of the edge between.

Now, Δy is a monotonically increasing step function of the caliper angle θ , as shown in Figure 9(a), and so $\cos^{-1}(\Delta y)$ is a monotonically decreasing step function (b). We seek an angle θ^* that satisfies $\theta^* = \cos^{-1}[\Delta y(\theta^*)]$, for such a θ^* is both a minimum by Lemma 7 and tangent to the pivots. In Figure 9(b) this is depicted




Figure 9: Graphs of Δy and $\cos^{-1}(\Delta y)$ for Figure 8.

by intersecting the $\cos^{-1}(\theta)$ staircase with the diagonal line $\theta^* = \theta$. When that diagonal pierces the staircase on a horizontal step (as illustrated), then the minimum is achieved rocking on both pivots; when the diagonal pierces a vertical step, the minimum is achieved with the corresponding edge flush.

This analysis establishes two facts:

Lemma 8 *With one edge of P fixed flush with an edge of C , the minimum perimeter parallelogram is unique, and can be found in $O(\log n)$ time by binary search for the intersection point with the arccos staircase illustrated in Figure 9(b).*

We will not justify this further because we only use the uniqueness. A consequence of the preceding two lemmas is this:

Corollary 9 *If a convex polygon C has one edge e such that all other edges of C project orthogonally into e , then the minimum perimeter enclosing parallelogram with one side flush with e is the enclosing rectangle with base e .*

Proof. When the pivot points are the two endpoints of e —which they can be because of the projection assumption—then $\Delta y = 0$ and by Lemma 7, $\theta^* = 90^\circ$. By Lemma 8 this is the unique minimum. \square

Optimal Pivots. The following lemma shows that the search for P based on edge e of C chosen to satisfy the 1-Flush lemma need not start anew for the next edge e' of C :

Lemma 10 *Let p and q be the left and right side pivots for the minimum perimeter parallelogram P with base edge flush with edge e of C . Then the left and right pivots p' and q' for the optimal parallelogram P' flush with the next edge e' counterclockwise of e on C , are either the same as p and q or counterclockwise of p and q .*

Proof. The value of Δy decreases to $\Delta y' < \Delta y$ for the same pivots p and q with respect to the new base line containing e' . With $\Delta y'$ smaller, $\theta^{*'} = \cos^{-1} \Delta y'$ is larger; see Figure 9. Not only is $\theta^{*'} > \theta^*$, it is larger with respect to a base line that has rotated counterclockwise from the base line from which θ^* is measured. Consequently, the side of P' realizing $\theta^{*'}$ is rotated counterclockwise of the line of P realizing θ . \square

This leads to a linear-time algorithm. First find the optimal P for any start edge e_0 of C as the one flush, in $O(n)$ time. Then, for each successive counterclockwise edge e of C , fix that as the one flush, and search for the optimal pivots counterclockwise of the previous pivots. Lemma 10 guarantees this will find the best parallelogram. Examples are shown in Figure 10.




Figure 10: Minimum perimeter parallelogram in red. (a) 1-flush minimum. (b) 2-flush minimum.

5 Open Problems

The main open problem is to prove or disprove Conjecture 3. One small step toward that would be to prove that three sweeps are never needed for triangles. Because the first sweep of a 3-sweep sequence reduces a triangle to a quadrilateral, this could be settled with a

characterization of the optimal 2-sweeps of quadrilaterals, in a form similar to that of Theorem 5.

Although we cannot think of any but contrived applications, it would be interesting to find either a minimum surface area enclosing parallelepiped in 3D, or one that minimizes the 1-skeleton length. The minimum volume parallelepiped has application to identifying important surfaces of a protein, and an $O(n^6)$ algorithm is developed in [VW04].

Acknowledgements. We thank Anastasia Kurdia and the referees for insightful comments.

References

- [BM03] Binay Bhattacharya and Asish Mukhopadhyay. On the minimum perimeter triangle enclosing a convex polygon. In *Lecture Notes Comput. Sci.*, volume 2866, pages 84–96. Springer-Verlag, 2003.
- [DJ09] Adrian Dumitrescu and Minghui Jiang. Sweeping points. *Algorithmica*, 2009. To appear, 2010.
- [MP08] Joseph S.B. Mitchell and Valentin Polishchuk. Minimum-perimeter enclosures. *Information Processing Letters*, 107(3–4):120–124, 2008.
- [STV⁺95] Christian Schwarz, Jürgen Teich, Alek Vainshtein, Emo Welzl, and Brian L. Evans. Minimal enclosing parallelogram with application. In *Proc. 11th Annu. ACM Sympos. Comput. Geom.*, pages C34–C35, 1995.
- [Tou83] Godfried T. Toussaint. Solving geometric problems with the rotating calipers. In *Proc. IEEE MELECON '83*, pages A10.02/1–4, 1983.
- [VW04] Frédéric Vivien and Nicolas Wicker. Minimal enclosing parallelepiped in 3D. *Comput. Geom. Theory Appl.*, 29:177–190, 2004.