International Journal of Computational Geometry & Applications

A MONOTONIC CONVOLUTION FOR MINKOWSKI SUMS

Victor Milenkovic
Department of Computer Science, University of Miami
Coral Gables, FL 33124-4245, USA
vjm@cs.miami.edu

Elisha Sacks
Computer Science Department, Purdue University
West Lafayette, IN 47907-2066, USA
eps@cs.purdue.edu

We present a monotonic convolution for planar regions A and B bounded by line and circular arc segments. The Minkowski sum equals the union of the cells with positive crossing numbers in the arrangement of the convolution, as is the case for the kinetic convolution. The monotonic crossing number is bounded by the kinetic crossing number, and also by the maximum number of intersecting pairs of monotone boundary chains, which is typically much smaller. We give a Minkowski sum algorithm based on the monotonic convolution. The running time is $O(s + n\alpha(n)\log(n) + m^2)$, versus $O(s + n^2)$ for the kinetic algorithm, with s the input size and with n and m the number of segments in the kinetic and monotonic convolutions. For inputs with a bounded number of turning points and inflection points, the running time is $O(s\alpha(s)\log s)$, versus $\Omega(s^2)$ for the kinetic algorithm. The monotonic convolution is 37% smaller than the kinetic convolution and its arrangement is 62% smaller based on 21 test pairs.

Keywords: Minkowski sum; convolution; kinetic framework.

1. Introduction

Minkowski sums of planar regions are an important computational geometry concept with diverse applications, including packing, path planning, and assembly. The Minkowski sum of A and B is $A \oplus B = \{a + b \,|\, a \in A, b \in B\}$. The kinetic convolution[1] is the standard construction method. The input is the oriented curve segments that form the region boundaries. The boundaries are convolved into a curve that partitions the plane into cells with crossing numbers. The cells with positive crossing numbers comprise the Minkowski sum. Fig. 1a–c illustrates the algorithm for A a convex polygon and B a hole bounded by a convex polygon.

The kinetic algorithm is suboptimal in that the portion of the convolution in the interior of $A \oplus B$ is generated, partitioned, and then discarded. The convolution can be pruned with a curvature test. The resulting convex convolution[2] consists of the Minkowski sum boundary along with some interior segments (Fig. 1d). Its cells do not have crossing numbers, so an alternate $A \oplus B$ membership test is used.

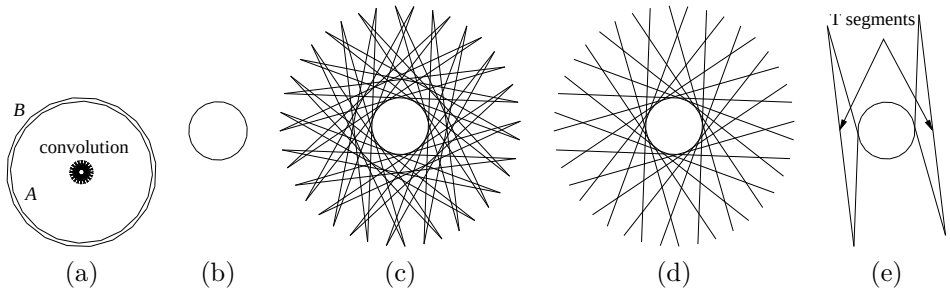

Fig. 1. (a) Part A in hole B and their kinetic convolution (actual size). (b) Minkowski sum x10. (c) Kinetic convolution x10. (d) Convex convolution x10. (e) Monotonic convolution x10.

We present a monotonic convolution whose cells have crossing numbers that determine membership in $A \oplus B$ (Fig. 1e, Sec. 4–5). Unlike the kinetic convolution, the monotonic convolution depends on the coordinate frame. The region boundaries are partitioned into x-monotone chains that meet at turning points. The monotonic convolution is obtained by forming convex convolutions for all pairs of A/B chains, computing their envelopes, and adding curve segments at concave turning points (the T segments in Fig. 1e). The monotonic crossing number is bounded by the kinetic crossing number, and also by the maximum number of intersecting pairs of monotone boundary chains, which is typically much smaller. In our example, the monotonic crossing number is 1 at infinity, 0 in the Minkowski sum, and 2 inside the triangular cells. In contrast, the kinetic crossing number (Fig. 1c) is 1 at infinity, rises to $\Omega(s)$ for s-sided A and B, then falls to 0 inside the Minkowski sum.

We give a Minkowski sum algorithm (Sec. 6) based on the monotonic convolution. The input regions are bounded by s line and circular arc segments. The running time is $O(s + n\alpha(n)\log(n) + m^2)$, versus $O(s + n^2)$ for the kinetic algorithm, with n and m the number of segments in the kinetic and monotonic convolutions, and with $\alpha(n)$ the inverse Ackerman function. For inputs with a bounded number of turning points and inflection points, the running time is $O(s\alpha(s)\log s)$, versus $\Omega(s^2)$ for the kinetic algorithm. This is the case in our example. The monotonic convolution is 37% smaller than the kinetic convolution and its arrangement is 62% smaller based on 21 test pairs of industrial part shapes (Sec. 7).

2. Definitions

A planar region has an outer boundary and zero or more inner boundaries (Fig. 2a). A boundary is a simple loop of line and arc segments. The endpoints of segment a are designated tail(a) and head(a), so that the region interior is to the left when a is traversed from tail to head. Arcs are split at vertical turning points, so every segment lies on its upper or lower semicircle. Arc a has a center, center(a), and a signed radius, radius(a), which is positive/negative on the upper/lower semicircle.

The rightward (outward) normal vector of a line segment, a, is $(y, -x)$ with

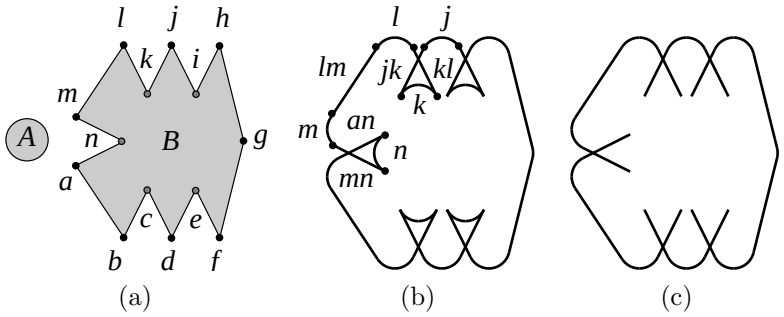

Fig. 2. (a) Planar regions; (b) Kinetic convolution; (c) Convex convolution.

$(x, y) = \text{head}(a) - \text{tail}(a)$. The rightward normal vector of an arc segment at p is $(p - \text{center}(a))/\text{radius}(a)$ when $\text{head}(a)_x < \text{tail}(a)_x$ and is negated otherwise. The rightward normal angles at $\text{tail}(a)$ and $\text{head}(a)$, called $\alpha(a)$ and $\beta(a)$, are in $[-\pi, 0]/[0, \pi]$ for a lower/upper semicircle. An arc is convex when $\alpha(a) < \beta(a)$ and is concave otherwise. The angle interval of a is $[\alpha(a), \beta(a)]$ when a is convex and is $[\beta(a), \alpha(a)]$ otherwise. The point on a with angle θ is called $\text{point}(a, \theta)$.

The inputs to the kinetic convolution[1] are polygonal tracings. Imagine a wheelchair driving along the boundary of a polygon. When it reaches a vertex, it executes a pure turn to orient itself along the next edge. The turns eliminate boundary orientation discontinuities. Since we include arcs in boundaries, we represent a pure turn as a zero radius arc: a set of pairs $\langle p, \theta \rangle$ with p a vertex and with $\theta \in [\alpha, \beta]$ a rightward normal angle. Between each pair of consecutive segments, a and b, we insert a smoothing arc, e, with $\text{center}(e) = \text{tail}(e) = \text{head}(e) = \text{head}(a)$, $\text{radius}(e) = 0$, $\alpha(e) = \beta(a)$, and $\beta(e) = \alpha(b)$. In Fig. 2a, convex/concave smoothing arcs are labeled with black/gray circles. Smoothing arcs are split at 0 and π like input arcs. They are treated identically to input arcs in our algorithms.

3. Prior work

With turns replaced by smoothing arcs, the kinetic convolution, $A \otimes_k B$, is the sum of all pairs of boundary points with equal rightward normal angles. Boundary segments $a \in A$ and $b \in B$ contribute a segment, e, to the sum when their angle intervals overlap. In Fig. 2b, selected sum segments are labeled with their B segments, and the A segment is the upper/lower semicircle when the e normal vector points up/down. If a or b is a line segment with angle α, e is a line segment with $\alpha(e) = \beta(e) = \alpha$. Let $t_a = \text{tail}(a)$ and $h_a = \text{head}(a)$ for a line segment, $t_a = h_a = \text{point}(a, \alpha)$ for an arc, and likewise for b. The endpoints of e are $\text{tail}(e) = h_a + h_b$ and $\text{head}(e) = t_a + t_b$ when a or b is a concave arc, and the opposite otherwise. If a and b are arcs whose shared angle interval is $[\alpha, \beta]$, e is an arc with $\alpha(e) = \alpha$ and $\beta(e) = \beta$ when a and b are both convex/concave, and the angles reversed otherwise; $\text{center}(e) = \text{center}(a) + \text{center}(b)$,

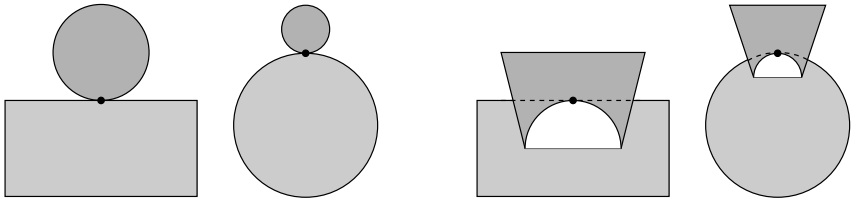

Fig. 3. Curvature test.

radius(e) = radius(a) + radius(b), tail(e) = point$(a, \alpha(e))$ + point$(b, \alpha(e))$, and head(e) = point$(a, \beta(e))$ + point$(b, \beta(e))$.

The Minkowski sum of two bounded, simply connected regions is obtained by constructing the convolution arrangement and assigning the cells crossing numbers. The unbounded cell is assigned crossing number 0. The other cells are traversed in breadth-first order where two cells are neighbors if they share an edge. Suppose cell b is visited from cell a with crossing number c. Cell b is assigned crossing number $c + 1$ when the ab edge is crossed from right to left, as determined by its tail and head, and is assigned $c - 1$ otherwise. In Fig. 2b, the large inner cell has crossing number 1 and the five small cells have crossing number 2. If t lies in a cell with crossing number K, $A \cap (-B + t)$ has K connected components. We call K the kinetic intersection number of A and $-B + t$. The boundary of $A \oplus B$ consists of the segments that separate cells of zero and positive crossing number.

Another Minkowski sum algorithm[3] decomposes A and B into convex pieces, sums the pieces, and combines the results. The advantage is that convex regions are simple to sum. The algorithm is less efficient than the kinetic one because the union of the piece sum boundaries is a superset of $A \otimes_k B$. Also, concave boundary segments can only be approximated, typically by many line segments.

A large portion of the kinetic convolution can lie in the Minkowski sum interior, as Fig. 1a–c illustrates for a convex polygon in a convex polygonal hole. The convolution can be pruned by excluding convexly incompatible sums.

The kinetic convolution contains points that a curvature test excludes from the Minkowski sum boundary. Boundary segments $a \in A$ and $b \in B$ are convexly incompatible when the sum of their curvatures is negative. The curvature is 0 for a line segment, $1/r$ for a convex arc of radius r, and $-1/r$ for a concave arc. Smoothing arcs have infinite curvature. If points $p \in a$ and $q \in b$ have equal rightward normals, A and $-B + p + q$ have a point of tangency and intersect in the neighborhood of this point, so $p + q$ is in the Minkowski sum interior (Fig. 3).

The convex convolution, $A \otimes_c B$, is the set of sums of convexly compatible pairs of boundary points[2] (Fig. 2c). We have $A \otimes_c B \subseteq A \otimes_k B$, yet $A \otimes_c B$ is a superset of the boundary of $A \oplus B$. The cells of the arrangement do not have crossing numbers. The Minkowski sum is obtained by selecting a point t in each cell and testing

whether A intersects $-B + t$, which takes $O(s \log s)$ time with s the input size.

4. Monotonic intersection number

We define the monotonic intersection number of regions A and B and prove that it is positive when A and B intersect, is zero otherwise, and is bounded by the kinetic intersection number. The next section shows how to construct a monotonic convolution curve whose crossing number at t is the monotonic intersection number of A and $-B + t$. The Minkowski sum consists of the cells with positive crossing numbers, since these are the ones in which A and $-B + t$ intersect.

We make the simplifying assumption that neither A nor B has a vertical line segment in its boundary. In Sec. 6, we show how to handle vertical line segments in the input through a form of symbolic perturbation.

A segment is upper/lower when its tail is right/left of its head, which means that the region interior is below/above it. An upper/lower chain is a maximal sequence of upper/lower segments on a region boundary. The chains are partially ordered in y. Pairs of upper and lower chains meet at turning points. A turning point is convex/concave when the lower chain is below/above the upper chain. In Fig. 2a, the upper chains are *ghijklm* and *na,* the lower chains are *abcdefg* and *mn,* the convex turning points are *a,g,* and *m,* and the concave turning point is *n.*

Let O be the number of pairs consisting of an upper chain of A that overlaps in x a lower chain of B. (We show that O also equals the number of overlaps between lower chains of A and upper chains of B.) If an upper chain of A overlaps in x and is always less in y than a lower chain of B, we say the pair of chains is facing. Similarly, an upper chain of B can face a lower chain of A. Let F be the number of pairs of facing chains. Let T be the number of concave turning points of A that lie interior to B plus the number of concave turning points of B that lie interior to A. The monotonic intersection number is $M = O - F - T$.

We define the pair intersection number, P, as the number of pairs of connected components of A and B that intersect. Trivially, $P \leq K$ with K the kinetic intersection number because K is the number of connected components in $A \cap B$.

Theorem 1. *If A and B are bounded and simply connected, then $P \leq M \leq K$.*

Proof. Sweep A and B with a vertical sweep line, L, from left to right. As the sweep line moves, maintain the value of P, M, and K for $H \cap A$ and $H \cap B$, where H is the half-plane to the left of L. We have $P = M = K = 0$ when $H \cap A = H \cap B = \emptyset$. The values can change at turning points of A or B and at the intersection of a chain of A with a chain of B. We show that $P \leq M \leq K$ after each of the nine possible events (Fig. 4). When the sweep ends, $H \cap A = A$ and $H \cap B = B$.

The proof covers the generic case in which events have distinct x values and crossings are transverse. The general case can be handled by symbolic perturbation.
1. Free left convex turning point If L sweeps through a left convex turning point $a \in A$ such that $a \notin B$ (free), $H \cap A$ has a new upper and a new lower chain

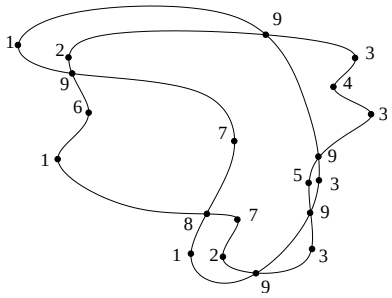

Fig. 4. Intersection number events.

with a as their common left endpoint. $H \cap A$ has a new connected component, but it does not intersect $H \cap B$, so P and K stay the same. If $L \cap B$ has k components, L intersects k upper chains and k lower chains of B. Hence, O increases by k whether we are counting upper chains of A paired with lower chains of B or vice versa. T remains the same. $L \cap B$ has k_- components below a and k_+ components above a with $k_- + k_+ = k$. The lower chain of A starting at a faces k_- upper chains of B and the upper chain of A starting at a faces k_+ lower chains of B. Thus, F increases by $k_- + k_+ = k$ and $M = O - F - T$ is unchanged.

2. Embedded left convex turning point If $a \in A \cap B$, the new connected component of $H \cap A$ intersects a single component of $H \cap B$ in a single region, so P and K increase by 1. One component of $L \cap B$ contains a. The lower and upper chain of A at a face k_- and k_+ upper and lower chains of B where this time $k_- + k_+ = k - 1$. F increases by $k - 1$ and $M = O - F - T$ increases by 1.

3. Right convex turning point If $a \in A$ is a right convex turning point, sweeping past a does not change P, M, or K.

4. Free left concave turning point If $a \in A$ is a left concave turning point with $a \notin B$, M remains the same, as in the convex case. No new component or intersection region is created, so P and K stay the same.

5. Embedded left concave turning point If $a \in A \cap B$, O increases by k, the number of components of $L \cap B$, as in the convex case, and F increases by $k - 1$. This time T increases by 1, so M stays the same. No new component or intersection region is created, so P and K stay the same.

6. Free right concave turning point If $a \in A$ is a right concave turning point and $a \notin B$, then O, F, T, and hence M stay the same. Since $a \notin A \cap B$, K stays the same. P does not increase, since the number of A components does not increase.

7. Embedded right concave turning point If $a \in A \cap B$, T increases by 1 and M decreases by 1. We show that two $H \cap A$ components merge at a, which implies that P decreases by 1. The two corresponding components of $H \cap A \cap B$ merge, so K decreases by 1.

Two vertical segments of $L \cap A$ merge at a. Suppose they are connected by a curve in $H \cap A$. Once the two segments merge, the curve can be extended to the right of a to form a loop around a. Such a loop cannot exist because A is simply connected, so the curve cannot exist, which means that the vertical segments are in different components of $H \cap A$ before a.

8. Intersection of facing chains If an upper chain of A intersects a lower chain of B from below at their leftmost intersection p, their subchains in $H \cap A$ and $H \cap B$ go from facing to non-facing as L sweeps through p. F decreases by 1, so M increases by 1. A new region of $H \cap A \cap B$ appears, so K increases by 1. If the component of A immediately below p did not intersect the component of B immediately above p, P also increases by 1; otherwise, it stays the same.

9. Intersection of other chains If p is not the leftmost intersection, the subchains in $H \cap A$ and $H \cap B$ were not facing, so F and M stay the same. K increases by 1. The component of A below p and the component of B above p already intersected, so P stays the same. If an upper chain crosses a lower chain from above, two upper chains cross, or two lower chains cross, no numbers change. □

The theorem tells us that M is positive when A and B intersect and is zero otherwise. If they intersect, $P > 0$, so $M > 0$. If not, $P = K = 0$, so $M = 0$. The theorem does not extend to multiply connected regions. A counterexample appears in Fig. 5a: region A intersects disk B, so $K = P = 1$, but $M = 2 - 0 - 2 = 0$. We extend the theorem to $\overline{A}$ (the complement of A) and B bounded and simply connected. The general case reduces to these two cases, as shown in Sec. 6. We define $M = C + O - F - T$ where C is the number of upper chains (or lower chains) of B. Connected region A is wider than B if its projection onto the x-axis is longer.

Corollary 1. *If $\overline{A}$ and B are bounded and simply connected and $\overline{A}$ is wider than B, then $P \leq M \leq K$ for A and B, and $P = M = K = 1$ if $B \subseteq A$.*

Proof. Suppose the projection of B is to the right of the left endpoint of the projection of $\overline{A}$. Replace A by its intersection, S, with a square that contains $\overline{A}$ and B, and has the leftmost point of $\overline{A}$ on its left side (Fig. 5b). P stays the same because a component of B that intersects A also intersects S. K stays the same because $A \cap B = S \cap B$. O increases by C because the bottom of the square overlaps C upper chains of B. F and T stay the same because there are no new facing chains or embedded concave corners. By Thm. 1, $P \leq C + O - F - T \leq K$. For $B \subseteq A$, $P = K = 1$ because A and B each have one component and $A \cap B = B$. Since $P \leq M \leq K$, $M = 1$ also. □

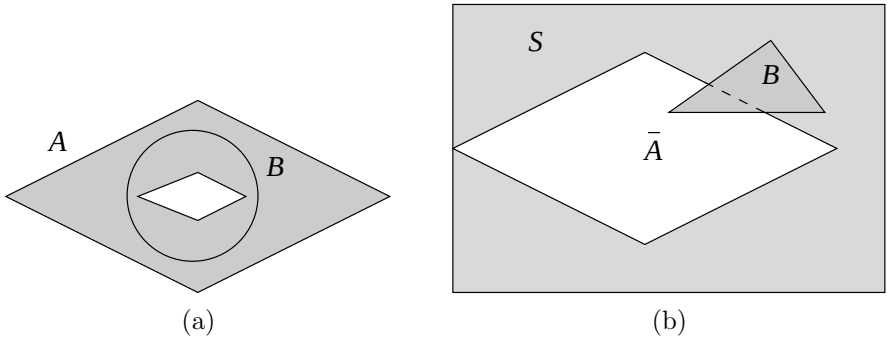

Fig. 5. (a) Counterexample to Thm. 1 for multiply connected A; (b) S for $\overline{A}$ and B.

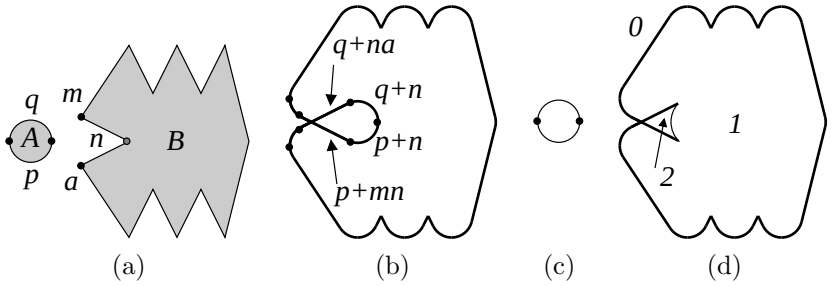

Fig. 6. (a) Planar regions; (b) $O-F$ segments; (c) T segments; (d) Monotonic convolution.

5. Monotonic convolution

If we replace B by $-B+t$, M becomes a function of the point t. The plane partitions into open regions where $M(t)$ is constant. Define the M segments to be the multi-set of oriented segments on these region boundaries such that each segment is oriented with larger M on the left and its multiplicity is the increase in M between the left and right regions. We define the monotonic convolution to be the M segments. The M segments have the property that for all t_1, t_2 not on any segment, the crossing number (Sec. 3) from t_1 to t_2 is $M(t_2) - M(t_1)$. The M segments plus the value of M in one cell represent the function $M(t)$.

To construct the monotonic convolution, we construct the $O-F$ segments and the T segments, analogously defined. The union of the $O-F$ segments and the $-T$ segments (T segments with orientation reversed) is the monotonic convolution. When constructing the union of multi-sets, the overlapping portion of segments with opposite orientation is canceled. Suppose segments ab and cd are on the same line or circle with overlapping domains $a_x < d_x < b_x < c_x$ and with opposite orientations. Their union is the ad portion of ab and the cb portion of cd. Fig. 6 illustrates monotonic convolution construction. The union of the $O-F$ segments $q+n$ and $p+n$ with the $-T$ circle is the concave side of the $M=2$ region.

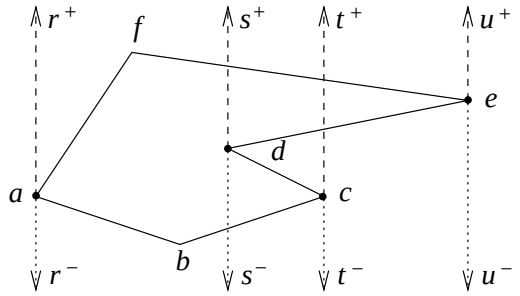

Fig. 7. Decomposition of region $abcdef$ into Π shapes $u^- efar^-$, $t^- cds^-$, $r^+ abct^+$, and $s^+ deu^+$.

5.1. $O - F$ segments

Define the Π shape of an upper chain, U, to be $\Pi(U) = U \oplus Y_-$ where $Y_- = \{(0, y) \,|\, y \leq 0\}$ is the negative y-axis. Define the Π shape of a lower chain, L, to be $\Pi(L) = L \oplus Y_+$ where $Y_+ = \{(0, y) \,|\, y \geq 0\}$. We chose the Π notation because the region looks like a pedestal. The Π shape of an upper/lower chain is bounded by the chain and by the downward/upward vertical rays at its endpoints. It is smoothed as described in Sec. 2. Figure 7 shows the four Π shapes for the region $abcdef$: $\Pi(efa)$ has sides r^- and u^- and $\Pi(de)$ has sides s^+ and u^+.

Minkowski sums commute with Π shapes. For upper chains U_1 and U_2, $\Pi(U_1) \oplus \Pi(U_2) = U_1 \oplus Y_- \oplus U_2 \oplus Y_- = (U_1 \oplus U_2) \oplus (Y_- \oplus Y_-) = (U_1 \oplus U_2) \oplus Y_- = \Pi(U_1 \oplus U_2)$. Since $\Pi(U_1) \oplus \Pi(U_2)$ is a Π shape, it consists of all the cells of the convex convolution $\Pi(U_1) \otimes_c \Pi(U_2)$, except for the uppermost cell. Define the crust to be the upper envelope of the convolution oriented right to left. For lower chains L_1 and L_2, similar arguments show that $\Pi(L_1) \oplus \Pi(L_2) = \Pi(L_1 \oplus L_2)$ and that the Minkowski sum is bounded from below by the lower envelope of $\Pi(L_1) \otimes_c \Pi(L_2)$. Define the crust to be the lower envelope oriented left to right.

Theorem 2. *The $O - F$ segments of regions A and B are the union of crusts of pairs of upper/lower chains from A and B.*

Proof. We discuss right chain endpoints (Fig. 8). A lower chain, c, and an upper chain, d, of A meet at p. A lower chain, e, and an upper chain, f, of B meet at q. The $c + e$ crust meets the $d + f$ crust at $p + q$. The analysis of left endpoints is the same, but the orientations of the O segments and vertical F segments are reversed.

The O segment is the vertical line $t_x = p_x + q_x$ with upward orientation, since d overlaps $-f + t$ in x to the left of this line. Chains d and $-f + t$ are facing if d is below $-f + t$ along every vertical line, which is equivalent to $\Pi(d) \cap \Pi(-f + t) = \emptyset$, or $t \notin \Pi(d) \oplus \Pi(f)$, since $\Pi(-f + t) = -\Pi(f) + t$. Each crust segment is a $-F$ segment because $-F$ is larger below than above. The upward vertical ray from $p + q$, oriented downward, is a $-F$ segment because the chains do not overlap on the right and are facing on the left. Chains d and $-f + t$ do not contribute the

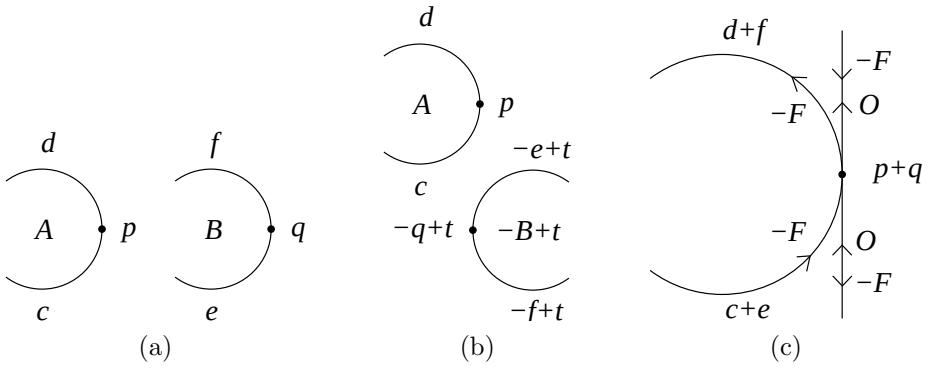

Fig. 8. (a) A and B; (b) A and $-B+t$; (c) O and $-F$ segments.

downward vertical ray from $p+q$ because the chains do not overlap on the left and are not facing on the right. Chains $-e+t$ and c contribute the $c+e$ crust and the downward ray from $p+q$ to the set of $-F$ segments.

The $O-F$ segments are just the crusts because the O and F verticals cancel. Consider the verticals at $t_x = p_x + q_x$ (Fig. 8). When t crosses the vertical, O increases by 1 because d overlaps $-f+t$ in x. F also increases by 1: d faces $-f+t$ when $t_y > p_y + q_y$ and $-e+t$ faces c otherwise. Hence, $O-F$ is unchanged. □

5.2. *T segments*

The value of T for A and $-B+t$ is the number of concave turning points $p_1 \in A$ such that $p_1 \in -B+t$ plus the number of concave turning points $p_2 \in B$ such that $-p_2 + t \in A$. We have $t \in p_1 + B$ and $t \in A + p_2$ respectively. The boundaries of these regions are the T segments.

Each concave turning point becomes a convex turning point of an upper and lower Π shape. These convex Π turning points generate convex convolution segments that can be $O-F$ segments. All such $O-F$ segments cancel with the corresponding T segments. The monotonic convolution is a subset of the kinetic convolution when all the convolution segments generated by turning points appear as $O-F$ segments, as in Fig. 6. In Fig. 9, concave turning point p on upper chain pq generates convolution segment abc of which ab is an $O-F$ segment, but bc is not because it is below the envelope. The T segment is $abcd$ of which ab cancels with the $O-F$ segment. Segment bc is in the monotonic convolution, but is not in the kinetic convolution, which contains cd only.

6. Minkowski sum algorithm

The Minkowski sum of bounded, simply connected regions A and B is computed as follows: form the $O-F$ segments, form the $-T$ segments, form their union to obtain $A \otimes_m B$, construct its arrangement, compute crossing numbers, and return

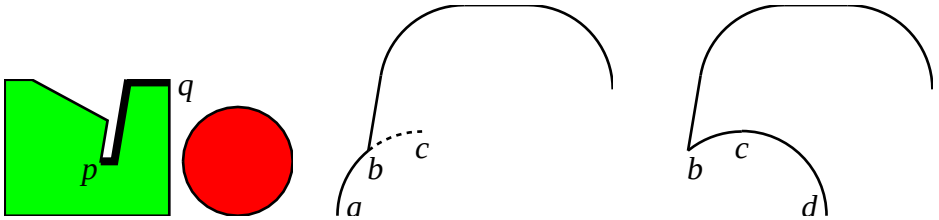

Fig. 9. Partial cancellation: (a) regions; (b) $O - F$ segments for pq; (c) M segments.

the edges that separate cells with zero and positive crossing numbers.

The $O - F$ segments are formed from all pairs of upper or lower chains: (1) smooth the chains at the endpoints, (2) find all segment pairs whose angle intervals overlap; (3) sum the convexly compatible pairs; (4) compute the upper or lower envelope. Step 1 takes $\mathrm{O}(s)$ time for two regions with s boundary segments. Step 2 takes $\mathrm{O}(s+n)$ for two regions with n sum pairs.[4] Step 3 takes $\mathrm{O}(n)$. Step 4 takes $\mathrm{O}(n \log n)$ for line segments and $\mathrm{O}(n\alpha(n) \log n)$ for arcs with $\alpha(n)$ the inverse Ackerman function.[5]

We symbolically perturb vertical input segments by labeling them as upper or lower. A vertical segment that joins two upper/lower segments is labeled upper/lower. Otherwise, a vertical segment with a rightward/leftward normal is labeled upper/lower. Two upper/lower chains are summed as before with vertical segments ignored. The endpoints of the sum are determined by summing corresponding endpoints of the chains, including endpoints of vertical segments. Vertical segments are added to bridge vertical gaps between the sum endpoints and the envelope and within the envelope. These modifications take time linear in the input plus output size.

The $-T$ segments are generated. A turning point is defined to be the junction of an upper and lower segment, including verticals. If p_2 is a concave turning point of B, generate a copy of the boundary of A translated by p_2 with orientations reversed. For segment e_1 of A, generate e_1' with $\mathrm{tail}(e_1') = \mathrm{head}(e_1)+p_2$, $\mathrm{head}(e_1') = \mathrm{tail}(e_1)+p_2$, and for an arc $\mathrm{center}(e_1') = \mathrm{center}(e_1)+p_2$ and $\mathrm{radius}(e_1') = \mathrm{radius}(e_1)$. Handle a concave turning point p_1 of A analogously. The running time is $\mathrm{O}(sc)$ with $c = \mathrm{O}(s)$ the number of concave turning points.

Segments are canceled. Gather all $O - F$ and $-T$ segments with the same line or circle. Sort the segment endpoints by x and cancel backwards intervals against forward intervals. Generate new segments with the same curves using the vertices whose x-coordinates bound uncanceled intervals. The running time is $\mathrm{O}(n + sc)$ using hashing, since there are $\mathrm{O}(n)$ $O - F$ segments and $\mathrm{O}(sc)$ $-T$ segments.

The arrangement of $A \otimes_m B$ is computed in $\mathrm{O}(m^2)$ time with $m = \mathrm{O}(n+sc)$ the number of convolution segments. An output sensitive bound is no better because the arrangement size can be $\Omega(m^2)$ when the Minkowski sum size is constant. When

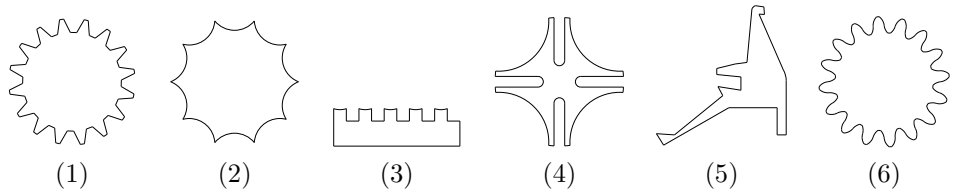

Fig. 10. Test profiles.

the regions have a bounded number of turning points and inflection points, m is $O(s)$. The arrangement has $O(s)$ size because it consists of a bounded number of monotone chains, which intersect $O(s)$ times. The kinetic convolution contains $\Omega(s)$ chains. Fig. 1 shows that the size of its arrangement is $\Omega(s^2)$. Assign crossing numbers and compute $A \oplus B$ as described in Sec. 5.

If $\overline{A}$ and B are bounded and simply connected, let a and b be the leftmost and rightmost vertices of $\overline{A}$; define c and d for B likewise. If $b_x - a_x \leq d_x - c_x$, then $A \oplus B$ is the entire plane. Otherwise, proceed as before, but assign crossing number 1 to the unbounded cell in accordance with Cor. 1. If A and B are multiply connected, construct the sum of the outer loops and the sum of each outer loop with each inner loop of the other region. $A \oplus B$ is the intersection of the resulting sums. For general A and B, the union of the sums of all pairs of connected components is $A \oplus B$. The running time is linear in the arrangement size in all cases.

7. Experimental results

We compare the monotonic convolution to the kinetic convolution on six industrial part profiles (Fig. 10). The number of boundary segments without smoothing arcs is $s = 64, 10, 23, 24, 22, 96$. We tested all 21 pairs of profiles. The tests were performed with robust floating point algorithms that will be reported separately.

Figs. 11 and 12 show two typical pairs and Table 1 summarizes the results for all 21 pairs. The average number of kinetic, convex, and monotonic convolution segments is 821, 486, and 516; the average number of arrangement edges is 2342, 1469, and 902; and the percentage of edges in the Minkowski sum is 6, 10, and 16. The monotonic convolution is 37% smaller than the kinetic convolution and only 6% larger than the convex convolution. Its arrangement is 61% smaller than the kinetic arrangement and 37% smaller than the convex arrangement because envelopes are more compact than convolution segments. The Minkowski sum computation time is 42% smaller with the monotonic convolution than with the kinetic convolution. We conclude that the monotonic convolution provides crossing numbers like the kinetic convolution, while matching the simpler convex convolution in complexity and beating its arrangement complexity and computation time.

The monotonic convolution is not a subset of the kinetic convolution when convex convolution segments involving concave turning points do not become $O-F$ segments, and hence the corresponding $-T$ segments are not canceled (Figs. 9, 11,

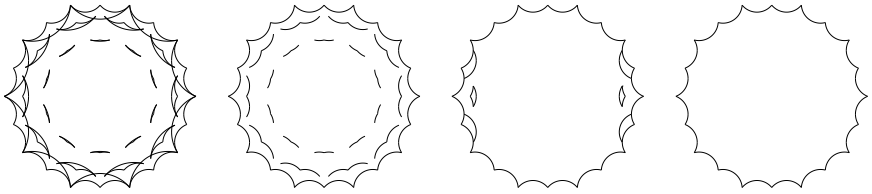

Fig. 11. Pair 2/2: kinetic (a), convex (b), and monotonic (c) convolutions; (d) Minkowski sum.

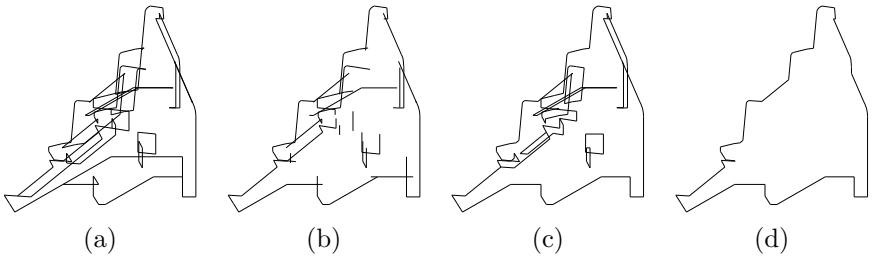

Fig. 12. Pair 5/5: kinetic (a), convex (b), and monotonic (c) convolutions; (d) Minkowski sum.

and 12). Although cancellation helps in some cases, such as in Figs. 1 and 6, it might hurt on average. We can define an alternate convolution consisting of the $O - F$ segments alone. Its crossing number at t is $M'(t) = O(t) - F(t)$ by Thm. 2. If A or $\overline{A}$ is bounded and simply connected and so is B, this crossing number is positive when A intersects $-B + t$ and is zero otherwise. If they intersect, $M > 0$, which implies $M' > 0$ because $T \geq 0$. If not, $M = 0$ by Thm 1 and $T = 0$ by definition, so $M' = 0$. We plan to investigate the theoretical and practical value of this convolution.

Acknowledgments

Research supported by NSF grants IIS-0082339, CCF-0306214, and CCF-0304955.

References

1. L. Guibas, L. Ramshaw, and J. Stolfi. A Kinetic Framework for Computational Geometry. In *Proceedings of the 24th IEEE Symposium on Foundations of Computer Science*, pages 100–111, 1983.
2. Kaul A., M. A. O'Connor, and V. Srinivasan. Computing minkowski sums of regular polygons. In *Proceedings of the Third Canadian Conference on Computational Geometry*, pages 74–77, 1991.
3. Pankaj K. Agarwal, Eyal Flato, and Dan Halperin. Polygon decomposition for efficient construction of minkowski sums. In *European Symposium on Algorithms*, pages 20–31, 2000.

Table 1. Comparison of kinetic and monotonic convolutions: n, c, and m segments in the kinetic, convex, and monotonic convolutions; e_n, e_c, and e_m edges in their arrangements; and m_s segments in the Minkowski sum.

pair	n	c	m	e_n	e_c	e_m	m_s
1/1	832	448	610	1920	1126	754	163
1/2	768	424	472	2093	1292	753	164
1/3	689	445	404	2268	1520	679	134
1/4	990	560	638	4000	2260	1265	193
1/5	561	343	410	2658	1243	1018	168
1/6	2778	1476	1928	6657	4243	3015	279
2/2	108	72	44	194	134	46	28
2/3	244	166	141	888	525	196	85
2/4	363	257	212	1093	758	377	105
2/5	272	169	167	806	489	227	80
2/6	1234	774	676	2166	1892	985	232
3/3	134	73	26	237	161	43	43
3/4	511	357	225	1500	948	261	93
3/5	286	185	117	953	619	137	99
3/6	1167	741	623	3166	2168	1007	138
4/4	389	271	196	822	649	247	89
4/5	353	229	241	1041	661	403	89
4/6	1614	968	1040	4709	3064	2097	298
5/5	124	84	110	240	195	157	52
5/6	986	583	702	4111	1981	2212	212
6/6	2328	1188	1601	4725	3293	1945	236

4. L. J. Guibas and R. Seidel. Computing convolutions by reciprocal search. *Discrete Comput. Geom.*, 2:175–193, 1987.
5. Micha Sharir and Pankaj K Agarwal. *Davenport-Schinzel sequences and their geometric applications*. Cambridge University Press, New York, 1995.