Rectangle-of-influence triangulations

Therese Biedl* Anna Lubiw* Saeed Mehrabi* Sander Verdonschot†

1 Background

The concept of *rectangle-of-influence (RI) drawings* is old and well-studied in the area of graph drawing. A graph has such a drawing if we can assign points to its vertices such that for every edge (v, w) the *supporting rectangle* (i.e., the minimal closed axis-aligned rectangle $R(v, w)$ containing v and w) contains no other points. In the original setup, the graph had to have an edge for every pair of points with an empty supporting rectangle (the *strong* model, see e.g. [8]). Later papers focus on *weak* RI-drawings, where for every edge the supporting rectangle must be empty, but not all such edges must exist. Of particular interest are planar weak RI-drawings of planar graphs, since these can always be deformed to reside on an $n \times n$-integer grid. See e.g. [9, 1].

Our Results: In this paper, we take two computational geometry problems—how to triangulate a point set and how to flip between two geometric triangulations—and apply them in the setting of rectangle-of-influence drawings. In particular, we show that any point set can be triangulated (with some exceptions near the convex hull edges, which we show to be necessary) such that the resulting planar straight-line graph (PSLG) is an RI-drawing. Next, we turn to the problem of *flipping* among geometric triangulations, i.e., converting one triangulation into another through the operation of flipping the diagonal of one quadrangle. We show that any RI-triangulation can be converted into any other RI-triangulation by $O(n^2)$ such flipping-operations (and $\Omega(n^2)$ flips are required for some RI-triangulations). Moreover, all intermediate triangulations are also RI-triangulations. The main idea is that the L^∞-Delaunay-triangulation (defined below) is an RI-triangulation; it hence suffices to find one flip-operation that gets the RI-triangulation "closer" to the L^∞-Delaunay-triangulation in some sense. We also study how to flip from any triangulation to an RI-triangulation while getting "closer".

Existing literature: Triangulating point sets and polygons is one of the standard problems in computational geometry. Any set of n points can be triangulated in $O(n \log n)$ time (e.g. by computing the Delaunay triangulation), and the interior of a polygon can be triangulated in $O(n)$ time [4]. The Delaunay triangulation has been generalized to other "unit discs". In particular, for any convex compact shape C, the C-Delaunay-triangulation is a triangulation such that for every edge (v, w) there exists a homothet of C that contains v and w and no other points [5]. Aurenhammer and Paulini [2] studied a number of their properties. If C is a unit square (i.e., the unit-circle in the L^∞-metric), then this triangulation is called the *L^∞-Delaunay-triangulation*. The L^∞-Delaunay-triangulation can be computed in $O(n \log n)$ time [5]. Observe that the L^∞-Delaunay-triangulation is an RI-triangulation, but an RI-triangulation need not be a C-Delaunay-triangulation for any C because the supporting rectangles are not necessarily homothets of each other or expandable into such.

Flipping among triangulations is also a well-studied problem; see [3] for an overview of many variants and existing results. It is very easy to see that any (geometric) triangulation T_1 can be flipped into any other triangulation T_2 via the intermediary of the Delaunay triangulation T_D: We can always find a flip that gets us closer to the Delaunay triangulation (in the sense that some angle-sum increases), so keep flipping from T_1 until we reach T_D. Also compute the flips from T_2 to T_D, and reversing these flips and combining the two flip-sequences then gives the result. For C-Delaunay-triangulations, a similar result holds: we can always flip to get "closer" to the C-Delaunay-triangulation [2].

Notation: Let P be a set of n points that we assume to be in *general position* in the sense that no two points are on a horizontal or vertical line, and no 4 points are on a square. For any two points $p, q \in P$, define the *supporting rectangle* $R(p, q)$ to be the minimal axis-aligned rectangle containing p and q. Define a *supporting square* $S(p, q)$ to be a minimal axis-aligned square containing p and q; $S(p, q)$ is not unique. For any two points $p, q \in P$, we call a supporting rectangle/square of (p, q) *empty* if it contains no points of P other than p and q.

An edge (p, q) between points of P is called an *RI-edge* (*L^∞-edge*) if $R(p, q)$ is empty (resp., some supporting square $S(p, q)$ is empty). Note that an L^∞-edge is an RI-edge. An *RI-polygon* is a polygon for which every edge is an RI-edge. A planar straight-line graph (PSLG) is a *triangulation* if every interior face

*David R. Cheriton School of Computer Science, University of Waterloo, Waterloo, ON N2L 3G1, Canada, {alubiw,biedl, s2mehrabi}@uwaterloo.ca. TB and AL supported by NSERC.

†Computational Geometry Lab, School of Computer Science, Carleton University, Ottawa, Ontario, Canada. sander@cg.scs.carleton.ca

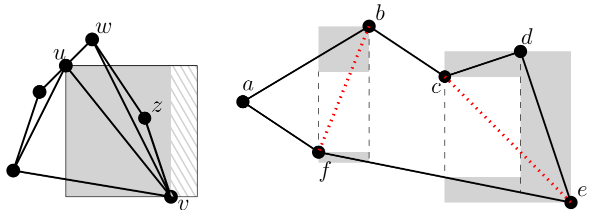

Figure 1: (Left) A triangulation. (v, w) is not locally RI. (u, v) is locally L^∞ (therefore locally RI), but not globally L^∞. (Right) A polygon (solid) with its trapezoidation (thin dashed). Edge (b, f) would be added with the first method, edge (c, e) with the second.

of the PSLG is a triangle. An *RI-triangulation* (*L^∞-triangulation*) is a triangulation for which every edge is an RI-edge (L^∞-edge). We sometimes use "unrestricted" triangulation for a triangulation that need not be an RI-triangulation. A triangulation is *maximal* if it contains as many edges as possible while staying within the additional requirements that we impose. Thus, a *maximal RI-triangulation* is an RI-triangulation with as many edges as possible while having only RI-edges, and similarly for maximal L^∞-triangulations.

For an edge (u, v) in a triangulation, vertex w is *facing* (u, v) if there exists an interior face $\{u, v, w\}$. Every interior edge has exactly two vertices facing it. We say that (u, v) is *locally RI* (resp. *locally L^∞*) if $R(u, v)$ (resp.. some supporting square $S(u, v)$) contains none of the vertices facing (u, v). Sometimes we say that an RI-edge (L^∞-edge) is *globally RI* (*globally L^∞*).

2 RI-triangulating an RI-polygon

We first show that any RI-polygon P can be *RI-triangulated*, i.e., made into a triangulation by adding only RI-edges in its interior (presuming no extra points are inside P). To do so, first find a trapezoidation of P, i.e., extend vertical subdivision lines from all vertices. This can be done in linear time [4].

We add RI-edges in two ways. First, check whether there is any trapezoid that has vertical subdivision lines on both its left and right sides, and for which the two vertices v, w that caused these lines are on opposite sides (top/bottom) of the trapezoid. If so, add the diagonal (v, w). This is an RI-edge since $R(v, w)$ is contained within the supporting rectangles of the top and bottom edge of T and the interior of P, all of which are empty. See edge (b, f) in Fig. 1.

Secondly, if no such trapezoid exists, then any remaining face is *x-monotone*, i.e., it consists of two x-monotone chains from left to right. Since the first method does not apply, one of the two chains is a single edge. Consider one such piece with (say) the bottom chain a single edge (v, w). All vertices in the top chain are outside $R(v, w)$ and hence have larger x-coordinate than v, w. Let u be a local maximum in the top chain, and connect the neighbors of u. The new edge is an RI-edge, because its supporting rectangle is contained in the union of the ones of the edges incident to u as well as $R(v, w)$ and the interior of P. See edge (c, e) in Fig. 1 (with $v{=}f, w{=}e, u{=}d$).

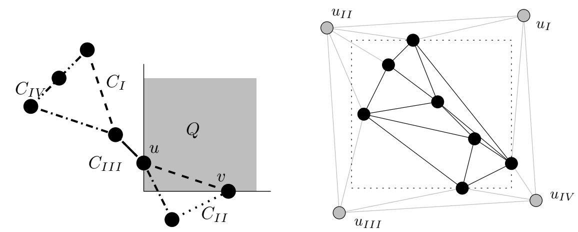

Figure 2: The maximal-hull of a set of points, and how to add corner-points and edges to them.

So we add RI-edges until all interior faces are triangles. This takes linear time (once the trapezoidation is found), since finding the local minimum/maximum for the second rule can be done in $O(1)$ amortized time.

Theorem 1 *Every RI-polygon can be triangulated using only RI-edges in linear time.*

3 Outer face considerations

The remaining sections deal with triangulations of point sets, rather than polygons. Here there arise some complications at the outer face. For any maximal (unrestricted) triangulation of P, the outer face consists of the convex hull $CH(P)$. If some edge of $CH(P)$ is not an RI-edge, then it obviously cannot be in an RI-triangulation. We begin by characterizing the outer face of any RI-triangulation.

The maxima-hull: We need some definitions that are closely related to the rectilinear convex hull (see e.g. [10]) and the maxima of a set of vectors (see e.g. [7]). Define a *first quadrant* of a point p to be the set $\{(x, y) : x \geq x(p), y \geq y(p)\}$. Define the *first-quadrant-chain* C_I to be all those points p in P for which the first quadrant relative to p contains no other points of P; we sort these points by increasing x-coordinate (hence by decreasing y-coordinate), and connect them with straight-line segments in this order. Similarly define three other chains C_{II}, C_{III}, C_{IV} using the three other types of quadrants. Note that C_I and C_{II} share one endpoint (the one with maximum x-coordinate), and similarly for the other chains, so we can combine the four polygonal chains into one closed polygonal chain that we call the *maxima-hull* $MH(P)$. Note that some chains may have edges in common, but no two edges cross, so $MH(P)$ may self-overlap but it has a well-defined interior region. See Fig. 2.

In the appendix, we prove the following.

Lemma 2 *For any point set P, the maxima-hull consists of L^∞-edges (hence RI-edges).*

Lemma 3 *For any point set P, any RI-edge (u, v) is within the region bounded by $MH(P)$.*

Combining them gives:

Corollary 1 *An RI-triangulation is maximal if and only if its outer face consists of the maxima-hull.*

Adding corner-points: It will be cumbersome to deal directly with edges that are on the convex hull but not RI-edges. To simplify our life, we add points as follows. For any point set P, let P^+ be the set obtained by adding four *corner-points* $u_I, u_{II}, u_{III}, u_{IV}$ that form an axis-aligned rectangle (we rotate it slightly to be in general position) and are outside the bounding box of P, with u_i in the ith quadrant relative to all points of P. See Fig. 2. Notice that the convex hull of P^+ consists of L^∞-edges.

Lemma 4 *Let T^+ be a maximal RI-triangulation of P^+, and let $T := T^+ - \{u_I, u_{II}, u_{III}, u_{IV}\}$. Then T is a maximal RI-triangulation of T.*

Proof. Clearly any edge of T is an RI-edge, so we only need to argue maximality. Assume (p, u_I) is an edge in T^+ for some $p \neq u_{II}, u_{III}, u_{IV}$. Then $R(p, u_I)$ contains no other point, and is to the right and/or above u_{II}, u_{III}, u_{IV}. So expanding $R(p, u_I)$ into the first quadrant of p does not add points of P, hence p is on the maxima-hull. So removing $\{u_I, u_{II}, u_{III}, u_{IV}\}$ from T^+ leaves an RI-triangulation where the outer face is the maxima-hull. By Corollary 1 this is maximal. □

4 RI-triangulating a point set

In this section, we study how to find a maximal RI-triangulation of a given point set. It is obvious that this exists (for example the L^∞-Delaunay triangulation will do), but our algorithm is especially simple.

Theorem 5 *A maximal RI-triangulation of a point set P can be computed in $O(n \log n)$ time, or $O(n)$ time if P is sorted by x-coordinate.*

Proof. As before, add corner-points $u_I, u_{II}, u_{III}, u_{IV}$ to obtain point set P^+. Sort the points by x-coordinates, and add an edge between any two consecutive points; these are clearly RI-edges. Also add the cycle $u_I, u_{II}, u_{III}, u_{IV}$; these are also RI-edges. Now we have a PSLG whose faces consist of RI-polygons. Triangulate each polygon with Theorem 1. We obtain an RI-triangulation T^+ of P^+, and it is clearly maximal since all convex-hull edges are in it. By Lemma 4, deleting the four corner-points gives the result. □

5 Flipping and RI-triangulations

In this section, we investigate flipping while maintaining RI-triangulations or (if we start with an unrestricted one) getting closer to an RI-triangulation. The natural "intermediary" here is the L^∞-Delaunay triangulation, which is an RI-triangulation. So we show that any triangulation can be flipped to an RI-triangulation while getting "closer", and then that any RI-triangulation can be flipped to the L^∞-Delaunay triangulation while remaining an RI-triangulation throughout.

5.1 Flipping to an RI-triangulation

Let P^+ be a point set for which any convex hull edge is an RI-edge (we will argue later how to remove this assumption). Let T be an arbitrary triangulation of P^+. A *bad triangle* $\{u, v, w\}$ in T is a face $\{u, v, w\}$ such that $v \in R(u, w)$. After possible rotation, assume that u is in the 2nd quadrant and w is in the 4th quadrant relative to v, with edge (u, w) routed above v. Define the *special region* of bad triangle $\{u, v, w\}$ to be all points p above (u, w) with $x(u) \leq x(p) \leq x(v)$ (including u) and all points p to the right of (u, w) with $y(v) \leq x(p) \leq y(w)$ (including w). See Fig. 3(a). The definition is symmetric for the other three possible rotations of a bad triangle. Now define for any triangulation T the potential function $\Phi(T)$ to be the sum, over all bad triangles $\{u, v, w\}$, of the number of points in the special region of $\{u, v, w\}$.

Lemma 6 *Let T be any triangulation of P^+ with $\Phi(T) > 0$. Then there exists an edge in T that we can flip so that the resulting triangulation T' satisfies $\Phi(T') < \Phi(T)$.*

Proof. Since $\Phi(T) > 0$, it must have at least one bad triangle, and hence edges that are not locally RI. Of all those edges, let (u, w) be the one that maximizes the L_1-distance between its endpoints, thus $|x(u) - x(w)| + |y(u) - y(w)|$ is maximum among all edges that are not locally RI. Since convex hull edges are RI-edges, we know that (u, w) is an interior edge and has two facing vertices. Let v be a vertex facing (u, w) that is in $R(u, w)$, and assume again that (after possible rotation) the bad triangle $\{u, v, w\}$ has u (or w) in the 2nd (4th) quadrant of v with (u, w) above v.

Let z be the other vertex facing (u, w). We claim that $x(z) > x(u)$. We know that z is above the line through (u, w) since $\{z, u, w\}$ is a face. If we had $x(z) < x(u)$, then $u \in R(z, w)$, and so (z, w) is not locally RI but its L_1-distance is longer than the one of (u, w), contradicting our choice of edge (u, w).

Therefore we know $x(z) > x(u)$, and symmetrically one argues $y(z) > y(w)$. Notice that therefore quadrilateral $\{u, v, w, z\}$ is convex and so edge (u, w) is flippable. We claim that regardless of the position of z, doing this

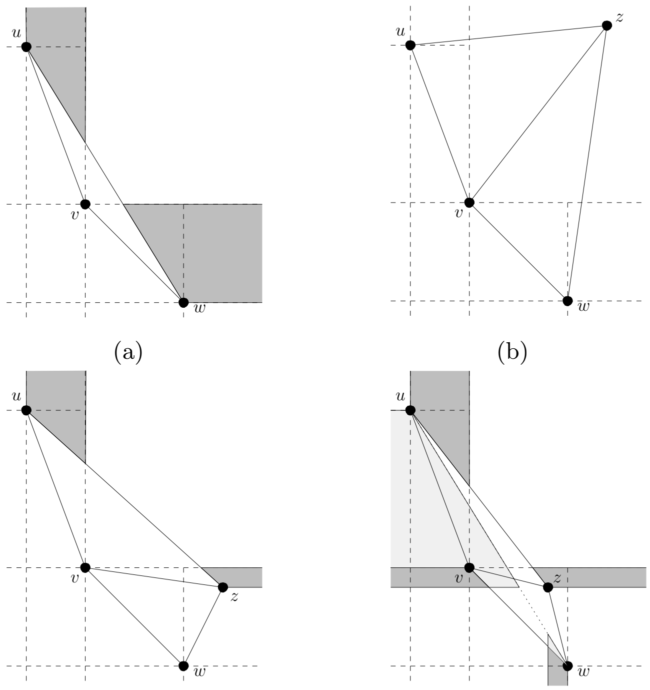

Figure 3: (a) The special region (shaded) of bad triangle $\{u, v, w\}$. (b-d) Possible positions of the other facing vertex z. Light gray regions were in the special region of the bad triangle $\{u, z, w\}$.

flip improves the potential function. To show this, we distinguish where z is located

- $x(z) > x(v), y(z) > y(v)$ (see Fig. 3(b)): In this case, neither of the two new triangles $\{u, z, v\}$ and $\{v, z, w\}$ is bad. Since triangle $\{u, v, w\}$ used to be bad, Φ decreases by at least 2.
- $x(z) > x(w), y(z) < y(v)$ (see Fig. 3(c)): In this case, the new triangle $\{u, z, v\}$ is bad, but $\{v, z, w\}$ is not. The special region of triangle $\{u, v, z\}$ is a strict subset of the one for triangle $\{u, v, w\}$, and in particular, excludes w. Hence Φ decreases.
- $z \in R(v, w)$ (see Fig. 3(d)): In this case, both new triangles $\{u, v, z\}$ and $\{v, w, z\}$ are bad. But both triangles $\{u, v, w\}$ and $\{u, z, w\}$ that were removed were bad, and the special regions of the new triangles are strict subsets of the special regions of the old triangles that, in particular, contain u and w only once, instead of twice. So again Φ decreases.

The cases for $x(z) < x(v), y(z) > y(u)$ and for $z \in R(u, v)$ are symmetric. □

Note in particular that if $\Phi(T) = 0$, then it has no bad triangles (because any bad triangle has at least two points in its special region); therefore it has no edge that is not locally RI.

Lemma 7 *Let T be a triangulation such that all edges are locally RI and all outer face edges are globally RI. Then all edges are globally RI.*

Proof. Suppose that some edges of T are not globally RI, hence contain points inside their supporting rectangles. Let $e = (u, v)$ be the edge with the closest such point, say z. Since e is an interior edge by assumption, it has two facing vertices; let w be the vertex facing e that is on the same side of the supporting line of e as z is. Suppose that z and w lie right of e, and u is the left endpoint of e; see Fig. 1. Observe that w must lie either above $R(u, v)$ (within the same x-range) or to the right of $R(u, v)$ (within the same y-rage), else some edge of $\{u, v, w\}$ would not be locally RI or $\{u, v, w\}$ would not be a face. Assume w lies strictly above $R(u, v)$, within the same x-range. Then (v, w) is also not globally RI, since z lies in $R(v, w)$. But z is closer to (v, w) than to e, contradicting our choice of e. □

Putting the two lemmas together, we can hence flip from any triangulation to an RI-triangulation while steadily improving the potential-function. Initially there are $O(n)$ bad triangles, each of which defines a special region containing at most n points, so $\Phi(T) \in O(n^2)$. Each flip improves the function, and we are done when it is 0, which means that the number of flips is $O(n^2)$. We summarize:

Lemma 8 *Any maximal triangulation of P^+ can be converted into an RI-triangulation of P^+ using $O(n^2)$ flips.*

We argue in the appendix that this bound is sometimes tight.

Lemma 9 *There are triangulations that cannot be converted into an RI-triangulation with $o(n^2)$ flips.*

Proof. Consider a set of points spread evenly over two opposing convex chains, such that the only possible RI-edges between the two chains connect point i on the first chain to point i and $i + 1$ on the opposing chain (see Fig. 4). The edges connecting consecutive points on each chain are not intersected by any other possible edge, which implies that these edges are present in every triangulation and can never be flipped. Thus, the region between the two chains is independent of the outside regions. Further, by construction, this center region has a unique RI-triangulation (Fig. 4, right). Now consider a triangulation that includes the long diagonal connecting one end of the first convex chain to the opposite end of the other chain (Fig. 4, left). Transforming this triangulation into the unique RI-triangulation requires $\Omega(n^2)$ flips, by an argument analogous to the one given by Hurtado et al. [6, Theorem 3.8] □

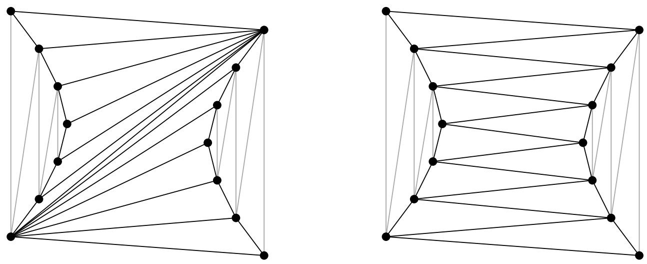

Figure 4: Turning the triangulation on the left into an RI-triangulation requires $\Omega(n^2)$ flips for the region between the chains to become the only possible RI-triangulation shown on the right.

Arbitrary point sets: It remains to argue how to handle point sets P where not all convex-hull edges are RI-edges. Assume we are given a maximal triangulation T of P, and we would like to flip that to a maximal RI-triangulation T_R. Add corner-points $u_I, u_{II}, u_{III}, u_{IV}$ as before to obtain point set P^+ and connect them in a cycle. Connect u_i (for $i \in \{I, II, III, IV\}$) to all points p on the convex hull of P for which quadrant i contains only p and u_i. This gives a triangulation T^+ of P^+ because the convex hull is the outer face T. See Fig. 2. The convex hull of P^+ consists of RI-edges, so there exists a sequence σ of flips that turns T^+ into a maximal RI-triangulation T_R^+.

Lemma 10 *Throughout flip-sequence σ, all edges incident to corner-points are RI-edges. Therefore any edge being flipped is not incident to a corner-point.*

Proof. The first claim implies the second because flipped edges were not locally RI. We prove the first claim for u_I only. Apart from the edges to u_{II} and u_{IV}, vertex u_I has an edge only to a point p for which the first quadrant is empty, so $R(u_I, p)$ is empty and the claim holds. We now show that any time we flip an edge (u, w) such that the new edge is (v, u_I) for some u, v, w, this new edge is an RI edge. Since u_I is outside the bounding box of P, it is outside $R(u, w)$. Since (u, w) was not locally RI (by our choice of edges to flip), therefore $v \in R(u, w)$. In the naming of Fig. 3, we have $u_I = z$ and the case of Fig. 3(b) applies since u_I is in the first quadrant of v. We already knew that in this case (v, z) is locally RI. But since $z = u_I$, edge (v, z) is globally RI: $R(v, u_I)$ lies within the union of $R(u, u_I)$ and $R(w, u_I)$ (both empty, because these edges are incident to u_I and hence RI-edges), and the interior of triangles $\{u, v, w\}$ and $\{u, u_I, w\}$ (which are faces). □

We now create a sequence of flips and edge deletions for T that leads to a maximal RI-triangulation by mirroring the flip-sequence of T^+. We maintain the claim that at any time triangulation T equals T^+ with the corner-points removed. Clearly this holds initially. Say the next flip for T^+ was to flip (u, w) to (v, z). We know $u, w \in P$. If $v, z \in P$, then (by induction) the 4-cycle u, v, w, z that exists in T^+ also exists in T, and so we can do the exact same flip in T and the claim holds. Else, one of v, z is a corner-point. Do an edge-deletion in T, i.e., remove edge (u, w) without adding a new one. The claim still holds since one end of (v, z) is not in P.

We end with a triangulation T_R of P that equals T_R^+ with the corner-points removed. By Lemma 4, this is a maximal RI-triangulation.

Theorem 11 *We can convert any maximal triangulation into an RI-triangulation by doing $O(n^2)$ flips and $O(n)$ edge-deletions.*

5.2 Flipping between RI-triangulations

As explained earlier, to flip between maximal RI-triangulations while maintaining an RI-triangulation, it suffices to show that every maximal RI-triangulation can be flipped into the L^∞-Delaunay-triangulation. To prove this, we use again a potential-function argument, but with a different function. We need the following:

Lemma 12 *[2] Let T be a maximal triangulation where all edges are locally L^∞. Then all edges are globally L^∞.*

Our potential function depends on having fixed, for every edge (u, v) of the current triangulation, a particular supporting square $S(u, v)$, and counting the number of points of P in it. We define $\Psi(T)$ to be the sum, over all edges (u, v) of the number of points in $S(u, v)$. When we flip, we are free to choose a supporting square for the new edge.

Lemma 13 *Let T be a maximal RI-triangulation that is not the L^∞-Delaunay triangulation. There exists an edge e such that flipping e results in an RI-triangulation T' and we can assign a supporting square to e such that $\Psi(T') < \Psi(T)$.*

Proof. By Lemma 12 T has some edge (u, w) that is not locally L^∞. Up to symmetry, we may assume that the height Y of $R(u, w)$ is no smaller than its width. All supporting squares of (u, w) are then in the horizontal strip H of height Y between u and w.

Since T is a maximal RI-triangulation, its outer face is the maxima-hull and consists of L^∞-edges. So (u, w) is an interior edge. Let v and z the two vertices facing (u, w). We claim that both v and z must be in strip H. To see this, recall that we have an RI-triangulation, and hence rectangle $R(u, w)$ is empty. This rectangle bisects strip H. So if, say, v is not in H, then at most one facing vertex is in H, which means that to one side of $R(u, w)$ in H there is no facing vertex of (u, v). We could hence pick a supporting square S' of (u, w) that consists of

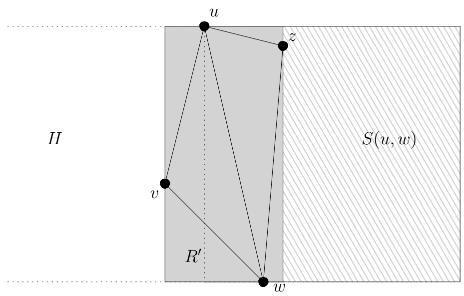

Figure 5: Finding a supporting square for (v, x).

$R(u, w)$ extended to that side. Then S' contains neither v nor z, and (u, w) would be locally L^∞, a contradiction.

So both v and z are in strip H. By planarity they must be on opposite sides of $R(u, w)$, say v is to the left and z is to the right. Quadrangle $\{u, v, w, z\}$ hence is drawn convex and edge (u, v) is flippable. Define R' be the minimum rectangle containing u, v, w, z, and notice that it is contained in the union of the supporting rectangles of the edges $(u, v), (v, w), (u, w), (u, z), (z, w)$. Since we had an RI-triangulation, therefore R' contains no points other than these 4. Also u, v, w, z are all on the boundary of R'. Therefore $R(v, z)$ is empty and the new edge (v, z) is an RI-edge.

Now we explain how to find a supporting square for (v, z). Let X be the width of R', and notice that $X < Y$, since $X = Y$ would imply four points on a square and $X > Y$ would mean that some square within R' supports (u, w) and does not contain v, z, contradicting that (u, w) is not locally L^∞. Now consider the union of R' and the square $S(u, w)$ that was used as supporting square for (u, w). Since (u, w) was not locally L^∞, at least one of $\{v, z\}$ is in $S(u, w)$, hence $S(u, w) \cup R'$ contains at most one more point than $S(u, w)$. Let R'' be the rectangle obtained by shrinking $S(u, w) \cup R'$ to height $Y - \varepsilon$ in such a way that $u, w \notin R''$. We choose ε so small that v and z remain in R'' and such that $Y - \varepsilon > X$. So R'' contains at least one fewer points than $S(u, w)$. Finally shrink R'' in width until it is a square; we can do this and retain v and z in it, since R'' is taller than R' is wide. Using the resulting square for $S(v, z)$ decreases Ψ as desired. □

Theorem 14 *Any maximal RI-triangulation T can be converted into any other maximal RI-triangulation T' by doing $O(n^2)$ flips, and all intermediate triangulations are maximal RI-triangulations.*

Proof. As before it suffices to argue this if T' is the L^∞-Delaunay triangulation. Compute an arbitrary set of supporting squares for T. Initially there are $O(n)$ edges in the triangulation, each of which has at most n points in its supporting square, so $\Psi(T) \in O(n^2)$. Applying the above flip means that with $O(n^2)$ flips we get to the L^∞-Delaunay-triangulation while maintaining an RI-triangulation. □

This bound is tight. Fig. 6 shows two RI-triangulations of a point set that forms two convex chains. Hurtado et al. [6, Theorem 3.8] showed that their flip distance is $\Omega(n^2)$ even without the restriction of using only RI-triangulations.

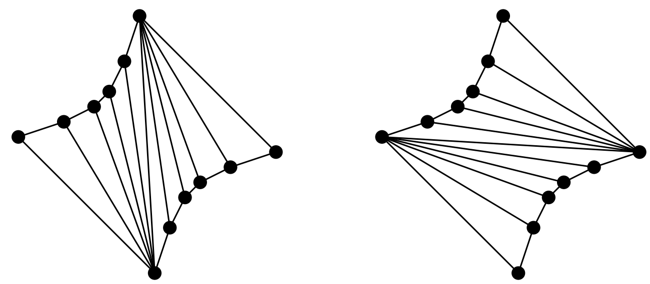

Figure 6: A pair of RI-triangulations such that $\Omega(n^2)$ flips are required to transform one into the other.

References

[1] S. Alamdari and T. Biedl. Open rectangle-of-influence drawings of non-triangulated planar graphs. In *Proc. GD'12*, pages 102–113, 2013.

[2] F. Aurenhammer and G. Paulini. On shape Delaunay tessellations. *Inf. Process. Lett.*, 114(10):535–541, 2014.

[3] P. Bose and F. Hurtado. Flips in planar graphs. *Comput. Geom.*, 42(1):60–80, 2009.

[4] B. Chazelle. Triangulating a simple polygon in linear time. *Discrete Comput. Geom.*, 6:485–524, 1991.

[5] R. Drysdale. A practical algorithm for computing the Delaunay triangulation for convex distance functions. In *Proc. SODA '90*, pages 159–168, 1990.

[6] F. Hurtado, M. Noy, and J. Urrutia. Flipping edges in triangulations. *Discrete Comput. Geom.*, 22(3):333–346, 1999.

[7] H. T. Kung, F. Luccio, and F. P. Preparata. On finding the maxima of a set of vectors. *J. ACM*, 22(4):469–476, 1975.

[8] G. Liotta, A. Lubiw, H. Meijer, and S. Whitesides. The rectangle of influence drawability problem. *Comput. Geom.*, 10(1):1–22, 1998.

[9] K. Miura, T. Matsuno, and T. Nishizeki. Open rectangle-of-influence drawings of inner triangulated plane graphs. *Discrete Comput. Geom.*, 41(4):643–670, 2009.

[10] T. Ottmann, E. Soisalon-Soininen, and D. Wood. On the definition and computation of rectlinear convex hulls. *Inf. Sci.*, 33:157–171, 1984.

A Proof of Lemma 2

We first need a helper-result.

Lemma 15 *Let (u, v) be any edge on the first-quadrant chain C_I, say $x(u) < x(v)$. Let Q be the first quadrant of point $(x(u), y(v))$, i.e., it has u and v on its boundary. Then no point (other than u, v) is in Q.*

Proof. See Figure 2. Assume to the contrary that Q contained points other than u, v, and let w be the one that maximizes the x-coordinate. Since the first quadrants relative to u and v are empty, w must be in $R(u, v)$. By general position we have $x(u) < x(w) < x(v)$ and $y(v) < y(w) < y(u)$. The first quadrant of w cannot contain any other point, for all such points would be to the right of w (contradicting the choice of w). So point w should have been in C_I, contradicting that u, v were consecutive in C_I. □

To complete the proof of Lemma 2, note that for any edge (u, v) on C_I, rectangle $R(u, v)$ is part of this first-quadrant, so it is empty, and we can expand it into a square supporting u, v that is empty. So any edge on C_I is an L^∞-edge. Similar arguments hold for the other three quadrant-chains, which proves Lemma 2.

B Proof of Lemma 3

Proof. We aim to show that any RI-edge (u, v) is on or below $C_I \cup C_{II}$. Similar proofs in the other three directions show that (u, v) is either on or enclosed by the maxima-hull as desired.

Up to renaming and symmetry, we may assume that u is left of v and higher than v. Consider the maximal vertical strip S containing (u, v). This strip must intersect $C_I \cup C_{II}$, since the rightmost point of P is in C_I and the leftmost of P is in C_{II}. Let C be the part of $C_I \cup C_{II}$ within S, say its ends are p_u (on a vertical line with u) and p_v (on a vertical line with v). Applying Lemma 15 (or the equivalent for second quadrants) to the edge containing p_u shows that the vertical ray upward from p_u contains no other points of P. So either $p_u = u$, or p_u is above u. Similarly $p_v = v$ or p_v is above v.

In what follows, we use the term "vertex of C" for a point on C that is also a point of P, while "point of C" refers to an arbitrary point that belongs to C. If C has no vertices, then it is a single line segment $\overline{p_u p_v}$, and by the above (u, v) is below that. So assume that C has vertices.

Since C (as part of $C_I \cup C_{II}$) consists of an increasing chain followed a decreasing chain, its minima (with respect to y-coordinate) appear at the ends. Since p_v is above v, therefore all vertices of C have y-coordinate at least $y(v)$. Since none of these vertices are inside $R(u, v)$ (recall that (u, v) is an RI-edge), they in fact must have y-coordinate at least $y(u)$, hence be above $R(u, v)$. In consequence the only edge of C that can intersect $R(u, v)$ is the edge incident to p_v, but this edge is also above (u, v) since p_v is above v. So all of C is above (u, v) as desired. □

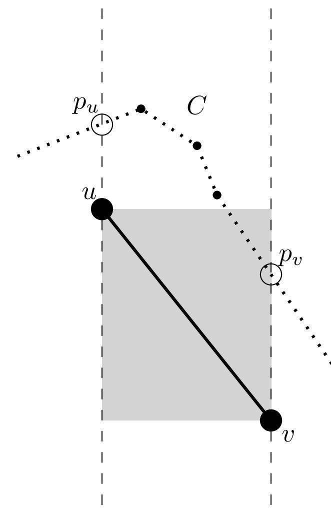

Figure 7: For the proof of Lemma 3.