Triangulating with High Connectivity

Tamal Krishna Dey*
Michael B. Dillencourt†
Subir K. Ghosh‡
Jason M. Cahill†

Abstract

We consider the problem of triangulating a given point set, using straight-line edges, so that the resulting graph is "highly connected." Since the resulting graph is planar, it can be at most 5-connected. Under the nondegeneracy assumption that no three points are collinear, we characterize the point sets with three vertices on the convex hull that admit 4-connected triangulations. More generally, we characterize the planar point sets that admit triangulations having neither chords nor complex (i.e., nonfacial) triangles. We also show that any planar point set can be augmented with at most 2 extra points to admit a 4-connected triangulation. All our proofs are constructive, and the resulting triangulations can be constructed in $O(n \log n)$ time. We conclude by stating several open problems. In particular, it is open whether a polynomial-time algorithm exists for determining whether a point set with no degeneracy restrictions and no restrictions on the number of extreme points admits a 4- or 5-connected triangulation.

Keywords. Discrete geometry, triangulations, k-connectedness.

May 20, 1994
Revised December 7, 1995

*Dept. of Computer Science and Engineering, I.I.T. Kharagpur, 721302, India
†Department of Information and Computer Science, University of California, Irvine, California 92717-3425, USA
‡Computer Science Group, Tata Institute for Fundamental Research, Bombay 400005, India

1 Introduction

We consider the problem of obtaining a planar network of maximum connectivity when the vertex locations are specified and only straight-line edges are allowed. Given a finite planar point set S and an integer k, we say that S is *k-connectible* if there exists a k-connected planar graph with straight-line edges having vertex-set S.

Planar point sets that are 1-, 2-, and 3- connectible are easily characterized, and no planar point set is k-connectible for $k > 5$. Our main result is a characterization of the conditions under which a planar point set in general position[1] and having exactly 3 extreme vertices is 4-connectible (Theorem 3.2). So, in particular, this theorem characterizes those point sets with three extreme vertices that admit Schlegel diagrams [7] of 4-connected 3-polyhedra.

Theorem 3.2 is actually a consequence of a more general theorem, Theorem 3.1, which characterizes the conditions under which a planar point set in general position can form the vertices of a triangulation having neither chords nor complex (i.e., nonfacial) triangles. One consequence of these results is that any planar point set in general position becomes 4-connectible if we are allowed to add 2 additional (Steiner) points (Theorem 3.7). All our proofs are constructive, and the graphs can be constructed in $O(n \log n)$ time. In Section 5, we present examples showing why the restrictions in our Theorems are necessary, and we state several open problems. In particular, it remains open whether there is a polynomial algorithm for determining whether a point set is 5-connectible, and whether there is a polynomial algorithm for determining whether a point set with no restriction on the number of extreme points is 4-connectible.

We know of no previous work on the problem of determining whether a given set is k-connectible. There have been many other examples of problems where the input is a set of geometrical data and the desired output is an object that can be described in combinatorial or graph-theoretical (non-metric) terms. Specific examples of such problems include the joint triangulation problem (determining whether two labeled point sets admit triangulations such that a labeled edge in one is a labeled edge in the other) [11], finding a minimum set of lines that cover a given finite planar point set [9], finding a tree with low stabbing number that spans a given planar point set [1], and determining whether a planar point set admits a triangulation whose dual has a Hamiltonian cycle [2].

In a sense, the problem addressed in this paper can be viewed as the inverse of the problem of drawing a planar graph, which has been the subject of considerable attention [5]. In particular, our results complement recent work, motivated by floorplanning problems in VLSI circuit design, concerning layout of triangulations having no complex triangles. A floorplan in VLSI circuit design is essentially a dissection of a rectangle into a finite number of non-overlapping sub-rectangles. It is known that a triangulated planar graph has a rectangular dual which is a floorplan only if it does not have a complex triangle [13].

Triangulations without complex triangles have been previously studied from a purely graph-theoretical perspective by one of the authors [6]. A classical theorem of graph theory asserts that any triangulation having neither complex triangles nor chords is Hamiltonian [12].

[1] By general position, we mean that no three points are collinear. Most of the terms used in this section are defined in Section 2.

2 Preliminaries

A *polygon* is a closed, bounded subset of the plane whose boundary is a simple cycle consisting of the union of a finite collection of line segments. The endpoints of the line segments are the *vertices* of the polygon; $|P|$ denotes the number of vertices of polygon P. We say that a point x is inside, on, or outside polygon P according to whether x is (respectively) in the interior of, on the boundary of, or in the complement of P.

Let S be a finite set of planar points. A *triangulation* of S is a planar graph T with vertex-set S such that all edges are line segments, the boundary of the outer face is the boundary of the convex hull, and all faces of T with the possible exception of the exterior face are bounded by triangles. A *chord* of a triangulation T is an edge connecting two nonconsecutive vertices on the boundary, and a *complex triangle* is a triangle that does not form the boundary of a face; see Figure 2.1. A triangulation is said to be *noncomplex* if it has neither chords nor complex triangles.

A graph is *k-connected* if it is connected and it remains connected whenever $k-1$ vertices and their attached edges are removed. A planar point set S is *k-connectible* if there exists a k-connected planar graph with vertex set S such that all edges are line segments. Since adding edges to a graph cannot decrease the connectivity, S is k-connectible if and only if there is a k-connected triangulation with vertex set S.

The following characterizations of k-connected triangulations are immediate consequences of results established in [8].

Lemma 2.1 A triangulation is 3-connected if and only if it does not have a chord.

Lemma 2.2 A triangulation T is 4-connected if and only if

(A1) T does not have a chord.

(A2) T does not have a complex triangle.

(A3) No interior vertex is connected to two or more non-consecutive vertices on the boundary of T.

Lemma 2.3 If the boundary of the outer face of a triangulation T is a triangle, then T is 4-connected if and only if it has no complex triangle.

The following lemma contains the elementary facts about connectibility mentioned in the introduction.

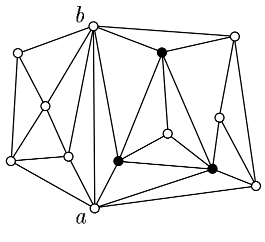

Figure 2.1: Chords and complex triangles: ab is a chord, and the three dark vertices form the boundary of a complex triangle.

Lemma 2.4 Let S be a planar point set.

(E1) S is always 1-connectible.

(E2) S is 2-connectible if and only if the points of S do not all lie on a single line.

(E3) S is 3-connectible provided S is 2-connectible and S does not consist of the vertices of a convex polygon.

(E4) S cannot be k-connectible for $k > 5$.

PROOF. Statements (E1) and (E2) are obvious; note that our general-position assumption implies that S is always 2-connectible. (E3) follows from Lemma 2.1, since we can choose an interior point $v \in S$ (i.e., a point not on the boundary of the convex hull of S), connect all points on the convex hull boundary to v, and then add each remaining interior point of S to the triangulation in such a way that no chords are introduced. (E4) follows from Euler's formula. □

We conclude this section with an easy but technical graph-theoretical lemma that will be useful in the proof of our main theorem. We use the graph-theoretic terminology of [3]. In particular, $K_{m,n}$ denotes the complete bipartite graph on a vertex set with bipartition (X, Y), where $|X| = m$ and $|Y| = n$. If graph G has vertex-set V and $V' \in V$, the subgraph of G *induced* by V' has vertex-set V' and edge-set consisting of all edges with both endpoints in V'.

Lemma 2.5 Let G be a subgraph of $K_{3,3}$, with bipartition B and C. Suppose that every vertex in B has degree at least 2, and that G has no induced $K_{2,2}$ subgraph. Then G is $K_{3,3}$ minus a perfect matching (or, equivalently, G is a cycle).

PROOF. Let $B = \{b_1, b_2, b_3\}$ and $C = \{c_1, c_2, c_3\}$. Note that the total degree of the vertices of C must be at least 6, and no vertex can have degree greater than 3. If some vertex of C had degree 3, the presence of an induced $K_{2,2}$ subgraph would follow immediately. So every vertex of C must have degree exactly 2. Hence G has exactly six edges, which in turn implies that every vertex of B must have degree exactly 2.

Assume, without loss of generality, that the two neighbors of b_1 are c_2 and c_3, and that c_2 is also joined to b_3. The absence of an induced $K_{2,2}$ implies that b_3 cannot be joined to c_3. Hence b_3 must be joined to c_1, c_3 must be joined to b_2, and then b_2 must be joined to c_1. Hence G consists of the cycle $b_1c_2b_3c_1b_2c_3b_1$. □

3 Characterizing point sets admitting noncomplex triangulations

We assume throughout this section that our point sets satisfy the general-position assumption introduced in Section 1: no three points are collinear. In Section 5, we give an example showing why this assumption is important. We also assume that all point sets have at least four points.

The planar point set S shown in Figure 3.1 does not admit a noncomplex triangulation. Indeed, in any triangulation of S, vertex x must be connected to every other vertex, as are consecutive vertices around the convex set $S - \{x\}$. Any triangulation of S must also contain a chord of the

Figure 3.1: A planar point set that does not admit a noncomplex triangulation.

convex hull of $S - \{x\}$. This chord and the two edges joining its endpoints to x form a complex triangle.

Theorem 3.1, which we establish below, states that the example of Figure 3.1 is essentially the only 3-connectible planar point set that does not admit a noncomplex triangulation. One consequence of this theorem is Theorem 3.2, which characterizes those planar point sets with triangular convex hull boundaries that are 4-connectible. This result, in turn, allows us to conclude that any planar point set becomes 4-connectible with the addition of at most two Steiner points. The proof of Theorem 3.1 is constructive, and leads to an $O(n \log n)$ algorithm for constructing a noncomplex triangulation if one exists.

The following definition captures the salient properties of the example of Figure 3.1. A planar point set is *anomalous* if it contains a point x such that the following properties hold:

(B1) S has exactly three extreme vertices, one of which is x.

(B2) The set $S - \{x\}$ consists of the vertices of a convex polygon, P.

Theorem 3.1 If S is a planar point set in general position, then S admits a noncomplex triangulation if and only if (1) it is not anomalous, and (2) it is not the set of vertices of a convex polygon.

Theorem 3.2 If S is a planar point set in general position, with exactly three points on the convex hull boundary, then S is 4-connectible if and only if it is not anomalous.

Theorem 3.2 is an immediate consequence of Theorem 3.1 and Lemma 2.3. The necessity of conditions (1) and (2) in Theorem 3.1 follows from Lemma 2.1 and the preceding discussion of Figure 3.1. To establish sufficiency of these conditions, we let S be a finite planar point set that satisfies conditions (1) and (2), and we show how to construct a noncomplex triangulation of S. Throughout this section, if A is a polygon whose vertices are a subset of S, we will use the phrase "a triangulation of the region bounded by A" to mean a triangulation whose vertices are the points of S lying on or inside A and having boundary A.

Our construction relies heavily on the *convex layer* structure [4]. The outermost convex layer of S is the boundary of the convex hull, and each subsequent convex layer is defined, recursively, to be the boundary of the convex hull of the set obtained by removing the vertices of all previously defined convex layers from S. We let k be the number of convex layers, with layer 1 the outermost layer and layer k the innermost layer. All layers except the innermost layer consist of boundaries of convex polygons. The innermost layer may consist of either a single point, a line segment, or the boundary of a convex polygon. If $k = 1$, S forms the vertices of a convex polygon, so we may assume that $k \geq 2$. We refer to the region between two consecutive layers in the convex layer structure as *interlayer regions*. If $k > 2$, we refer to layers other than the innermost or outermost layers as *intermediate* layers.

Our construction also requires adding edges between two consecutive layers in the convex layer structure to create a triangulation of the interlayer region. The edges that have one endpoint on each of the two layers are called *cross edges*.

Consider the following strategy: compute the convex layers of S, triangulate the innermost layer (if it is a convex polygon), and add cross edges to each interlayer region to produce a triangulation. If we attempt to use this strategy to produce a noncomplex triangulation, it can fail in one of three ways. First, if the innermost layer is a convex polygon with four or more vertices, then chords are introduced when it is triangulated. If the two endpoints of the chord are connected to a common vertex at layer $k-1$ during the triangulation of the region between layer k and $k-1$, a complex triangle results. Second, a triangulation of an interlayer region may produce a complex triangle consisting of two cross edges and one edge from the inner or outer layer. (Both of the preceding conditions arise if we attempt to apply this strategy to Figure 3.1.) Third, if any intermediate convex layer is a triangle, the strategy clearly fails to produce a noncomplex triangulation. With appropriate modification, however, these difficulties can be overcome for non-anomalous point sets. In essence, the first two difficulties are addressed by carefully choosing the cross edges, and the third difficulty is addressed by "borrowing" a point from the next outer layer.

3.1 Two key lemmas

The following two lemmas are used to extend noncomplex triangulations from one layer to the next layer. Given a triangulation let pqr and pqs be two interior triangular faces incident on edge pq. The edge pq is called *flippable* if the quadrilateral $prqs$ can be triangulated with the triangles prs and qrs (i.e., if we can replace edge pq with rs). Thus an edge is flippable if and only if it is not a convex hull edge and the union of the two faces incident on it is a convex quadrilateral. If pq is flippable, the operation of replacing it with the opposite diagonal in the quadrilateral (i.e., rs) is referred to as *flipping* edge pq.

Lemma 3.3 Let pqx be a complex triangle in a triangulation T, and suppose that edge pq is flippable. Then flipping pq cannot introduce any new complex triangle, with the possible exception of triangle rsx. This exception can occur only if rx and sx are edges of T (i.e., if p, q, r, and s all have x as a neighbor in T).

Proof. (See Figure 3.2(a) and (b)). Let T be a triangulation containing a complex triangle pqx, and let pqr and pqs be the two triangular faces of T incident on pq. Assume that at least one of the segments rx and sx is not an edge of T. Let T' be the triangulation obtained from T by flipping edge pq, and suppose that T' contains a complex triangle that was not a triangle in T. This complex triangle must be incident on the new edge rs. Any complex triangle of T' incident on rs must have an edge piercing either segment px or qx. But this is impossible, since px and qx are edges of T' (since they were edges of T and flipping edge pq did not affect them). $\square$

Lemma 3.4 Let A and B be two consecutive convex layers, with B lying completely inside A, such that $|B| \geq 4$. Suppose that the region bounded by A is triangulated so that:

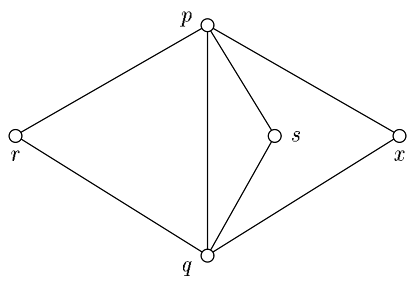

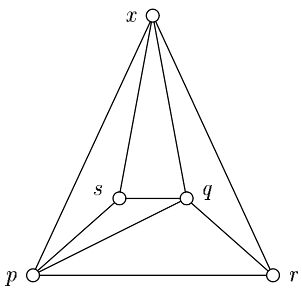

Figure 3.2: (a) Proof of Lemma 3.3. (b) Illustration of the exceptional condition in the statement of Lemma 3.3.

(Z1) The region between A and B is triangulated with cross edges.

(Z2) The region bounded by B is triangulated without any complex triangle, and there is no complex triangle incident on a chord of B.

Then it is possible to construct a noncomplex triangulation of the region bounded by A.

Notice that our lemma allows the possibility that A is a triangle. We remark that in applications of Lemma 3.4, condition (Z2) will be enforced in one of two ways. If B is the innermost layer, the chords of B will be carefully chosen so that (Z2) holds. If B is not the innermost layer, then the triangulation of B will be a noncomplex triangulation, so (Z2) will hold by definition.

PROOF. Let T be a triangulation of the region bounded by A satisfying properties (Z1) and (Z2). Assume T has a complex triangle pqr; otherwise there is nothing to prove. It follows from (Z1) and (Z2) that the polygon B must lie entirely on or inside this complex triangle; moreover, after suitable relabeling, the boundary of the triangle must be defined by either (i) a vertex p of B and an edge qr of A, or (ii) a vertex p of A and an edge qr of B.

In case (i) let x and w be the two vertices adjacent to p on B, labeled so that edges pq and pr are incident on triangles pqx and prw respectively. Let t be the neighbor of q on the boundary of A that is not r. We first have to dispose of a special subcase, illustrated in Figure 3.3. Suppose that A is a triangle tqr, and that one of the vertices of A (without loss of generality, assume r) is adjacent to all boundary vertices of B. There are then two complex triangles: pqr and pxr. Let v be the neighbor of x (other than p) on the boundary of B. Notice that since B is not a triangle, vertices v, p, w, and x are all distinct. Notice also that hypothesis (Z2) implies that B has no chords. We proceed as follows (see Figure 3.3(a)):

(F1) If possible, flip edge pr for tw. This flip eliminates both triangles pqr and pxr. By Lemma 3.3, this flip creates no new complex triangles.

(F2) If (F1) is impossible, flip edge rx for qv if possible. This flip creates no new complex triangle (by Lemma 3.3). Triangle prx is eliminated. Triangle pqr still exists, but we are no longer in the special subcase, so the argument of the next paragraph can be applied. (Note that we have not modified either polygon A or B, we have simply flipped one cross edge for another.)

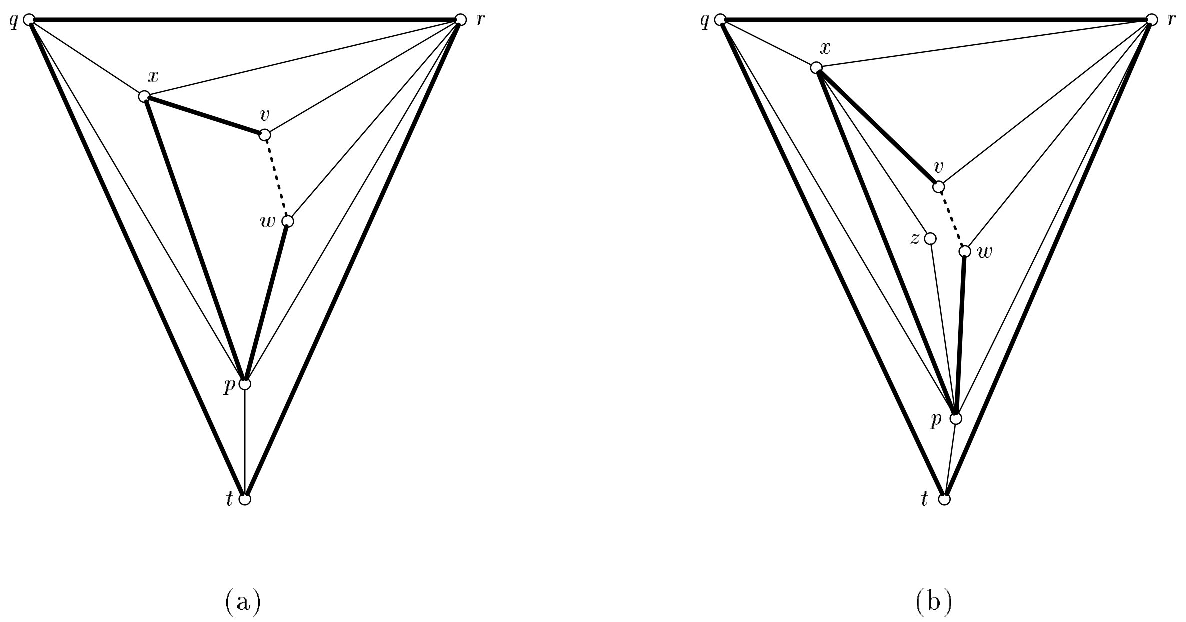

Figure 3.3: Special subcase in Proof of Lemma 3.4. (a) If possible, flip edge pr or edge rx. (b) If neither of these flips is possible, flip px, and then pq.

(F3) If neither (F1) nor (F2) is possible, let z be the vertex inside B such that pzx is a facial triangle (see Figure 3.3(b)). Flip px for qz, and then flip pq for tz. It follows from the impossibility of (F1) and (F2) that both these flips are possible. This sequence eliminates both triangles pqr and pxr. Moreover, neither flip creates a new complex triangle (by two applications of Lemma 3.3).

Having disposed of the special subcase, we now consider case (i) assuming that either A is not a triangle, or A is a triangle but no vertex of A is connected to every boundary vertex of B (Figure 3.4). Let H be the open halfplane supported by px that contains q, and let H' be the open halfplane supported by pw that contains r. We claim that at least one of pq and pr is flippable. If this claim were false, then there would be no vertex of A other than q, r in $H \cup H'$. Since B is contained entirely inside A, this is impossible, and the claim holds. Hence we can always flip one of pq and pr to destroy triangle pqr.

It remains to show that this flip does not introduce a new complex triangle. Without loss of generality, assume that edge pq is being flipped to edge tx. If A is not a triangle (Figure 3.4(a)), then r and t cannot be neighbors because A has no chords. If A is a triangle but no vertex of A is connected to every boundary vertex of B (Figure 3.4(a)), then r is not connected to x. Either assumption implies that p, q, t, and x cannot all have r as a common neighbor, so Lemma 3.3 implies that the flip does not introduce any new complex triangles.

In case (ii), observe that every vertex of B is connected to p by a cross edge, so (Z2) implies that B has no chords. Let x and w be the two vertices on B adjacent to q and r on B respectively (Figure 3.5). We claim that some edge of pqr is flippable. Let H and H' be the two open halfplanes supported by qx and rw respectively that contain p. If any vertex of A other than p lies in $H \cup H'$ (Figure 3.5(a)), at least one of the edges of pq and pr is flippable. If edge pq is flipped for edge tx, r and x cannot be neighbors (because rx would be a chord of B), so no new complex triangle is

(a) (b)

Figure 3.4: Case (i) in Proof of Lemma 3.4. (a) A is not a triangle. (b) A is a triangle, but not all vertices of B have a common neighbor on A.

(a) (b)

Figure 3.5: Case (ii) in Proof of Lemma 3.4. (a) Edge pq is flippable. (b) Neither edge pq nor edge pr is flippable.

introduced by Lemma 3.3. A symmetric argument applies if edge pr is flipped.

If neither pq nor pr is flippable (Figure 3.5(b)), all vertices of A except p lie in the region $H_c \cap H'_c$, where H_c and H'_c are the open halfplanes supported by qx and rw respectively that do not contain p. This implies that edge qr is flippable. The edge replacing qr connects an interior point of B (which cannot have p as a neighbor) with a point on A. So we can once again apply Lemma 3.3 to argue that no new complex triangle is introduced. This completes the proof of Lemma 3.4. $\sqcup\!\sqcap$

3.2 Proof of the main results

Now we proceed to prove the main result (Theorem 3.1). The plan of the proof is as follows. We first show that *any* planar point set S in general position with $k \geq 3$ has a noncomplex triangulation. We then show that if a planar point set S in general position with $k = 2$ fails to have a noncomplex triangulation, it must be anomalous.

Assume $k \geq 3$. For $j = 1, \ldots, k-2$, define a *good level-j triangulation* to be a triangulation for which the following properties hold:

1. The boundary of the outer face consists of layer j.
2. All points of S inside the boundary are vertices of the triangulation.
3. The triangulation is noncomplex.

The proof of the theorem for $k \geq 3$ follows by induction from the following two lemmas, since a noncomplex triangulation is simply a good level-1 triangulation:

Lemma 3.5 *S has a good level-$(k-2)$ triangulation.*

Lemma 3.6 *Given a good level-j triangulation, for $2 \leq j \leq k-2$, we can construct a good level-$(j-1)$ triangulation.*

Proof of Lemma 3.5: Let A be layer $k-2$, B layer $k-1$, and C layer k. We distinguish six cases, which we group as follows. B may be either a polygon having more than three vertices (Part I) or a triangle (Part II). Within each part, C may consist of a single vertex ($|C| = 1$), a line segment ($|C| = 2$), or a convex polygon ($|C| > 2$).

Part I: B is not a triangle: We first obtain a noncomplex triangulation of the region bounded by B, using the appropriate case (1, 2, or 3) below. It then follows immediately from Lemma 3.4 that the region bounded by A has a noncomplex triangulation (i.e., that S has a good level-$(k-2)$ triangulation).

Case 1: B is not a triangle, $|C| = 1$: Stellate B from the single vertex inside it.

Case 2: B is not a triangle, $|C| = 2$: Let $C = \{c_1, c_2\}$. Because B is not a triangle, it has two nonadjacent vertices p and q on opposite sides of the line c_1c_2. Triangulate the region between B and the segment c_1c_2 by first adding the four edges connecting each of c_1 and c_2 with each of p and q, and then adding a single edge connecting each of the remaining vertices of B to either c_1 or c_2 to obtain a triangulation. Because B is a convex polygon, there is only one way to perform the second step so that no edges cross.

Case 3: B is not a triangle, $|C| \geq 3$: C is the boundary of a convex polygon, inside B. First, compute some arbitrary triangulation of the region between B and C using cross edges. Let R be the set of vertices of C that are connected to two or more vertices of B. The set R necessarily contains at least two vertices of C. There are three subcases. R may consist of a pair of consecutive vertices of C (Subcase 3a). Otherwise, either C is not a triangle and R contains at least one pair of nonconsecutive vertices of C (Subcase 3b), or C is a triangle and $R = C$ (Subcase 3c).

Subcase 3a: R contains only a single pair of consecutive vertices of C, c_1 and c_2. This subcase is illustrated in Figure 3.6(a). All vertices of C are joined to a single vertex x of B. Assume that c_2 is the counterclockwise neighbor of c_1 on C; let c_1' be the clockwise neighbor of c_1, c_2' the counterclockwise neighbor of c_2. If C is a triangle, then $c_1' = c_2'$, but that does not affect the following argument. Now let x_1 and x_2 be, respectively, the counterclockwise and clockwise neighbors of x about the boundary of B. If the segment x_1c_1' does not intersect the interior of polygon C, then we can obtain a noncomplex triangulation of the region bounded by B by first flipping edge xc_1 and then triangulating the interior of C by joining c_1 to every vertex of C. If the segment x_2c_2' does not intersect the interior of polygon C, a similar construction works. Suppose neither of these last two conditions holds; then segments x_1c_1' and x_2c_2' both intersect segment c_1c_2. Let x' be the counterclockwise neighbor of x_1

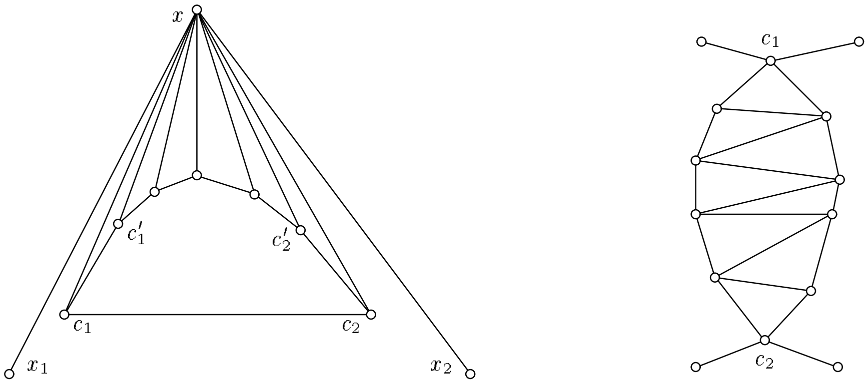

(a) (b)

Figure 3.6: Obtaining a good triangulation when the first two layers are convex polygons. (a) Subcase 3a: c_1 and c_2 are consecutive, and are the only two vertices of the inner polygon that have two neighbors on the outer polygon. (b) Subcase 3b: c_1 and c_2 are two nonconsecutive vertices of the inner polygon that have two neighbors on the outer polygon.

about B. Since B is not a triangle, $x' \neq x_2$. Since B is convex, a noncomplex triangulation of the region bounded by B can be obtained by first deleting edge c_1c_2 and an appropriate set of cross edges and then connecting x' to every vertex of C. The "appropriate set of cross edges" consists of all edges that interfere with adding the edges between x' and C (i.e., edge x_1c_2, if it exists, and any edges connecting x_1 with vertices on the counterclockwise arc of B from x' and x_2).

Subcase 3b: C is not a triangle, and R contains at least one pair of nonconsecutive vertices of C. Let c_1 and c_2 be a pair of nonconsecutive vertices of C in R (see Figure 3.6(b)). Removing these two vertices from the boundary of C creates two nonempty arcs, C' and C'', so no vertex of B can be joined to a vertex of C' and a vertex of C''. Hence if C is triangulated using only edges with one endpoint on C' and the other on C'' (in other words, triangulated so that c_1 and c_2 are ears), then no such edge can participate in a complex triangle. This produces a triangulation T of the region bounded by B, involving layers B and C, that satisfies all conditions of Lemma 3.4. Hence, there is a noncomplex triangulation of the region bounded by B.

Subcase 3c: C is a triangle $c_1c_2c_3$, and all vertices of C belong to R. In this case, there cannot be a complex triangle. Indeed, if there were a complex triangle consisting of two vertices from C and one vertex of B, the third vertex of C would have only one neighbor on B and hence could not be in R. So suppose the only remaining possibility happens, namely a complex triangle consisting of two vertices from B (say x and y) and one vertex (say c_1) of C. Since the triangulation was obtained using cross edges, xy cannot be a chord of B, so it must be an edge of B. The subgraph induced by x, y, and the vertices of C necessarily contains the three edges of C, the edge xy, and cross edges c_1x and c_1y. Since a planar graph on five vertices cannot be a complete graph, at least one of the edges c_2x, c_2y, c_3x, c_3y must be missing, so either c_2 or c_3 fails to be in R. Hence this complex triangle cannot exist, by contraposition.

Part II: B is a triangle: If B is a triangle, we "borrow" an appropriate vertex p from A, and use this vertex to "augment" B to a quadrilateral, B', whose vertices are the borrowed vertex and the three vertices of B. We show that there is always a way to choose p (and hence B') so that there is a noncomplex triangulation with boundary B'. The process works in two stages.

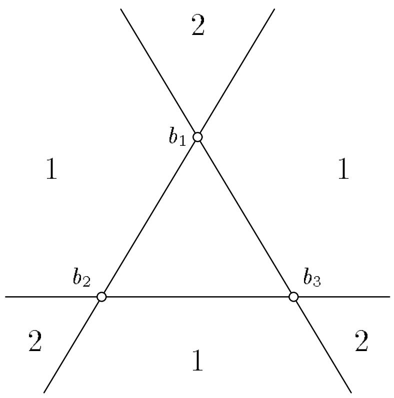

Figure 3.7: The arrangement of the lines supporting triangle B in Part II of the proof of Lemma 3.5.

1. **The augmentation stage:** We choose the vertex p and the augmented quadrilateral B', and construct a triangulation of the region bounded by B', in a case-specific manner. The details depend on whether C is a single point, a line segment, or a polygon (cases 4, 5, and 6 below). The resulting triangulation of the region bounded by B' has neither complex triangles nor chords.

2. **The completion stage:** After the augmentation stage is complete, we complete the triangulation of the region bounded by A. We do this by adding cross edges between A and B'. This second stage is the same for all three cases. The polygon B' is not necessarily convex, and A and B' have exactly one point in common (namely, the "borrowed" vertex p).

Before describing the process, we introduce some terminology. Let $B = b_1b_2b_3$. Consider the arrangement of the three lines that support the three edges of B, shown in Figure 3.7. This arrangement has seven planar regions: the triangular region B, and six other regions exterior to the triangle. We call the three exterior regions bounded by three lines *type-1* regions, and the three exterior regions bounded by two lines *type-2* regions, or *cones*.

We describe the completion stage first. Suppose that A contains a vertex in a type-1 region. We will see below that if this happens, the augmenting vertex p chosen by the augmentation stage is in a type-1 region. Hence the augmented quadrilateral B' is convex, as illustrated in Figure 3.8(a). It is then a simple matter to triangulate the region between A and B' using cross edges that do not involve p, which guarantees that the resulting triangulation of the region bounded by A does not have complex triangles or chords.

Now suppose that all vertices of A are in type-2 regions with respect to $B = b_1b_2b_3$. The augmented quadrilateral B' constructed by the augmentation stage will be non-convex only at a single vertex. We may assume, after suitable relabeling of the b_i's, that the boundary of B' is the cycle $pb_1b_2b_3$, and that the non-convex vertex of B' is b_1 (see Figure 3.8(b)). Since there are no vertices of A in type-1 regions, each type-2 region (in particular, the cone supported at b_2) must contain at least one vertex of A. To produce a noncomplex triangulation of A, first *temporarily* add the edge pb_2, and let B'' be the triangle pb_2b_3. Next, triangulate the region between A and B'', using only cross edges that are not incident on p. The resulting triangulation of the region bounded by A has a single complex triangle, namely pb_2b_3. We claim that edge pb_2 is flippable. Indeed, if pb_2 were not flippable then the halfplane supported by the segment b_1b_2 and containing p would contain

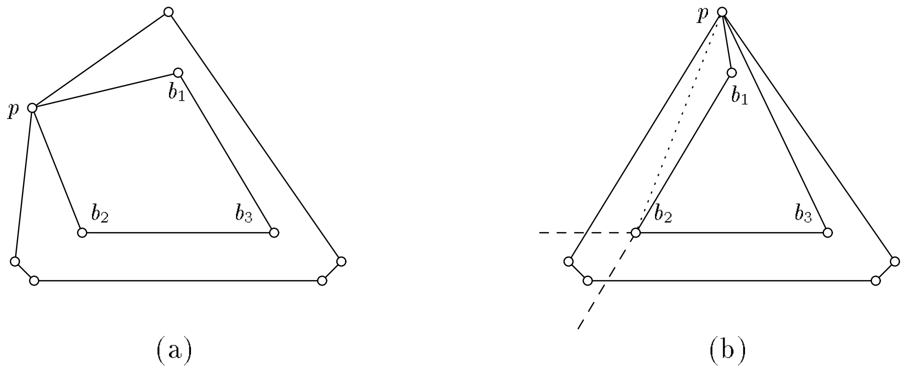

Figure 3.8: Obtaining a noncomplex triangulation of the region bounded by A from a noncomplex triangulation of the region bounded by B'. (a) When possible, p is chosen from a type-1 region. (b) The case where p cannot be chosen from a type-1 region.

no vertex of A other than p. But, as noted above, the cone supported at b_2, which lies within this halfplane, contains a vertex of A. Hence edge pb_2 is flippable. Since the polygon B' has no chords, it does not contain the edge b_1b_3. So by Lemma 3.3, the triangulation obtained by flipping edge pb_2 is noncomplex.

We now describe the augmentation stage. In each of the following three cases, we choose a point p on the boundary of A, which we use to augment B to a quadrilateral B', and we produce a noncomplex triangulation of the area bounded by B'. Moreover, if A has any type-1 vertices, then p is always chosen to be a type-1 vertex.

Case 4: B is a triangle, $|C| = 1$: Let c be the unique point inside B. If any vertices of A are in type-1 regions, let p be one such vertex. Otherwise, let p be any vertex of A. Without loss of generality, let b_1b_3 be the edge of B crossed by the segment pc. Let B' be the polygon with boundary $pb_1b_2b_3$, and stellate B' from c.

Case 5: B is a triangle, $|C| = 2$: Let $C = \{c_1, c_2\}$. Let $\{b_1, b_2, b_3\}$ be the vertices of B. If there is a point of A in a type-1 region, let p be such a point, and use p to augment B to a convex quadrilateral, B'. Let x and y be two non-consecutive vertices of B' that lie on opposite sides of the line c_1c_2. Join each of x and y to each of c_1 and c_2, and connect c_1 and c_2 to the two remaining vertices of B to obtain a noncomplex triangulation of the region bounded by B'.

Now suppose there is no point of A in a type-1 region. Relabel the vertices of B if necessary so that b_1 is separated from b_2 and b_3 by the line c_1c_2. The points b_2, c_1, c_2, and b_3 form the vertices of a convex quadrilateral; let $b_2c_1c_2b_3$ be the consecutive vertices of that convex quadrilateral (this may require swapping the labels of c_1 and c_2). Let p be any vertex of A in the cone at b_1. Assume that pc_2 crosses the edge b_1b_3 (otherwise, swap the labels on c_1 and c_2, and swap the labels on b_2 and b_3). Augment B to the quadrilateral $B' = pb_1b_2b_3$, and delete the edge b_1b_3. Add the edges pc_2, b_1c_1, b_1c_2, b_2c_1, b_3c_1, and b_3c_2. This completes the noncomplex triangulation of the region bounded by the quadrilateral B'.

Case 6: B is a triangle, $|C| \geq 3$:

Subcase 6a: C is a triangle: Define a vertex of C to be *restricted* if there is only one vertex of B to

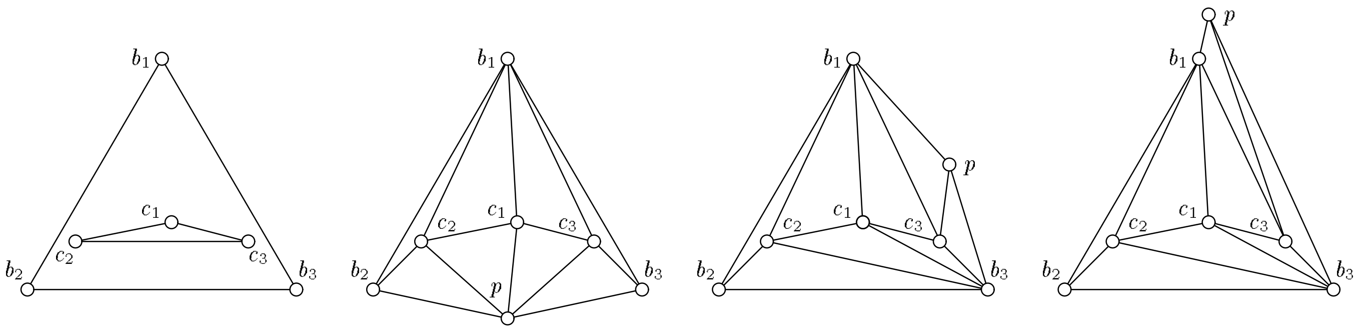

Figure 3.9: Obtaining a noncomplex triangulation when the first two layers are triangles and the inner layer has a restricted vertex. (a) Vertex c_1 is a restricted vertex. (b) Vertex p in the type-1 region separated from triangle B by the line b_2b_3. (c) Vertex p in the type-1 region separated from triangle B by the line b_1b_3. (d) The case where p cannot be chosen from a type-1 region.

which it can be connected by a line segment that does not enter the interior of C, such as vertex c_1 in Figure 3.9(a). It is easy to see that C cannot possibly have more than one restricted vertex.

If c_1 is a restricted vertex of C, let b_1 be the only vertex to which c_1 can be connected by a line segment not entering the interior of C. Assume that the remaining four vertices of B and C are labeled so that $b_2c_2c_3b_3$ forms a convex quadrilateral. If there is a vertex p in a type-1 region, then p can be used to augment B to B' so that a noncomplex triangulation of the region inside B' is obtained. Indeed, there are essentially only two possibilities, illustrated in Figures 3.9(b) and (c). If there is no vertex in a type-1 region, choose p in the type-2 region consisting of the cone supported at b_1, as shown in Figure 3.9(d). Assume that the segment pc_1 crosses the edge b_1b_3; if not, swap labels b_2 and b_3, and swap labels c_2 and c_3. Augment B to the quadrilateral $pb_1b_2b_3$, delete edges c_2c_3 and b_1b_3, and add edges b_1c_1, b_1c_2, b_1c_3, b_2c_2, b_3c_1, b_3c_2, b_3c_3, and pc_3.

Now suppose that C does not have any restricted vertices. We claim that there is a triangulation of the region between B and C such that each vertex of C has exactly two neighbors on C. Once such a triangulation is obtained, it is straightforward to verify that for any choice of augmenting vertex, a noncomplex triangulation of the region inside B' can be obtained by flipping a single edge of triangle B.

It remains to prove the claim. Define the (B, C)-visibility graph to be the bipartite graph whose vertices are the vertices of triangles B and C, with b_i and c_j connected by an edge iff the line segment from b_i to c_j does not enter the interior of C. If b_ic_j is an edge of this graph, we say that b_i "sees" c_j and c_j "sees" b_i.

Now suppose that there are two vertices of B (say b_2 and b_3) and two vertices of C (say c_2 and c_3) such that b_2 sees c_3, b_3 sees c_2, and the two segments b_2c_3 and b_3c_2 cross. If c_1 is on the same side of the line c_2c_3 as b_2 and b_3 (Figure 3.10(a)), then we can obtain the required triangulation by adding the edges b_2c_2, c_2b_1, b_1c_3, c_3b_3, b_3c_1, and c_1b_2. If c_1 is on the other side of this line, then since c_1 is not restricted, it must lie in one of the regions marked X, Y, or Z in Figure 3.10(b). It is straightforward to verify that for any one these regions, the required triangulation exists.

Finally, suppose that there does not exist a pair of edges in the (B, C) visibility graph corre-

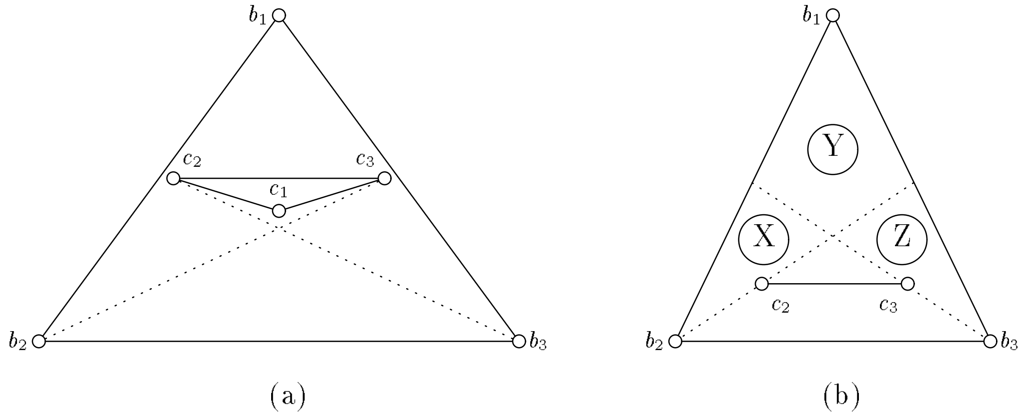

Figure 3.10: Obtaining a noncomplex triangulation when the first two layers are triangles and the inner layer has no restricted vertex.

sponding to a pair of crossing segments. Then the (B, C)-visibility graph contains no induced $K_{2,2}$ subgraph. Since each vertex of B sees at least two vertices of C, it follows from Lemma 2.5 that the (B, C) visibility graph consists of a cycle, which implies the claim.

Subcase 6b: C **is not a triangle**: If there is a vertex p of A in a type-1 region, use p to augment B to a convex quadrilateral B' and apply the argument of Case 3.

Suppose there is no vertex p of A in any type-1 region. Compute a triangulation of the region between B and C, using cross edges. If there are at least three vertices of C that are incident on two cross edges, then any type-2 vertex can be used as an augmenting vertex, in conjunction with a retriangulation of C, similar to the construction of Subcase 3b of Part I. Otherwise, a single vertex of B, which we assume to be b_1, is connected to all vertices of C. Choose p in the type-2 region supported by the vertex b_1 of B. Let c_1, c_2, c_1', and c_2' be defined as in Case 3. The situation is as depicted in Figure 3.11. The set of edges connecting p to vertices of C all lie outside C. Since C is not a triangle, there are at least four such segments. Hence at least one of the two edges b_1b_2 and b_1b_3

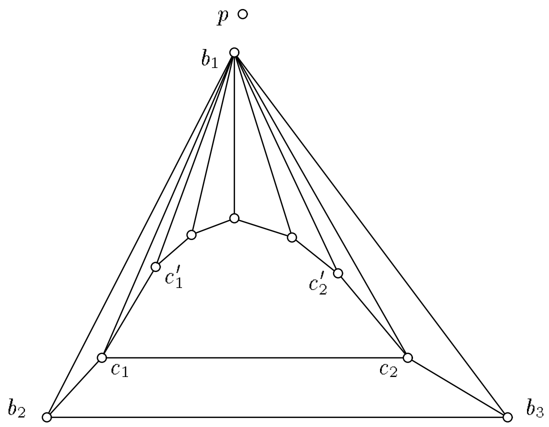

Figure 3.11: Subcase 6b: Construction for the case when all vertices of C are connected to a single vertex of triangle B.

is intersected by at least two such edges. Without loss of generality, assume edge b_1b_3 is intersected by edges pc_2 and pc'_2. Augment B to the quadrilateral $B' = pb_1b_2b_3$. Delete edge b_1b_3, and add edges pc_2 and pc'_2. Ensure that b_3 and c_1 are connected; this edge may already be there, but if not, flip edge b_2c_2 for b_3c_1. Since C is not a triangle, c_1 and c'_2 cannot be adjacent on C. Triangulate the interior of polygon C by connecting all vertices to c_2, analogous to Subcase 3a. This construction produces a non-convex triangulation of the region inside the augmented quadrilateral B'.

□

Proof of Lemma 3.6: Let A be layer $j-1$, B layer j. If B is not a triangle, then triangulate the region between A and B with cross edges and apply Lemma 3.4. If B is a triangle, augment B to B' by borrowing an appropriate vertex p from A, and then produce a noncomplex triangulation of B'. The borrowing proceeds as in Part II of the proof of Lemma 3.5. It is straightforward to verify that for any choice of p, at least one edge of B can be flipped for an edge with one endpoint on p such that the resulting triangulation of B' has no complex triangles. □

Proof of Theorem 3.1: If $k > 2$, the theorem follows by induction, using Lemma 3.5 as the base and Lemma 3.6 for the inductive step.

To complete the proof of Theorem 3.1, it suffices to address the case $k = 2$. If the outer layer has 4 or more points, the existence of a noncomplex triangulation follows immediately from the constructions in Part I of Lemma 3.5. So suppose the outer layer has 3 points. If there is exactly one point inside the outer layer, the unique possible triangulation is noncomplex. If there are exactly two points inside the outer layer, the configuration is anomalous.

Now suppose the inner layer consists of a convex polygon and the configuration is not anomalous. Triangulate the region between the outer layer and the inner layer using cross edges. If no point on the outer layer is connected to all points on the inner layer, then there are three different points on the inner layer with two neighbors on the outer layer. So either the inner layer is a triangle (and we get the graph of the octahedron), or there are two nonconsecutive points on the inner layer with two neighbors on the outer layer (and we can proceed as in Subcase 3b of Lemma 3.5.)

Finally, suppose the outer layer is a triangle, and the triangulation with cross edges connects a single vertex of the outer layer, say x, to all points of the inner layer. The condition (B1) holds. Since the configuration is not anomalous, (B2) cannot hold. It follows that some vertices of the inner layer currently connected only to x could also be joined to another vertex of the outer layer, say y. By making this change (which also requires deleting some other edges from x and adding some edges from y), we obtain a configuration in which there are three different points on the inner layer with two neighbors on the outer layer. This situation was dealt with in the preceding paragraph. This completes the proof of Theorem 3.1. □

We conclude this section by noting that we can always make a point set S in general position admit a 4-connected triangulation if we allow extra (Steiner) points. Indeed, let p be any point on the boundary of the convex hull of S. Add two points q and r such that all points of $S - \{p\}$ are inside the triangle pqr, choosing the points carefully so that (B2) does not hold. The resulting set $S' = S \cup \{q, r\}$ is not anomalous and hence, by Theorem 3.2, is 4-connected. We have shown:

Theorem 3.7 A planar point set S in general position can be augmented using at most two extra points so that it admits a 4-connected triangulation.

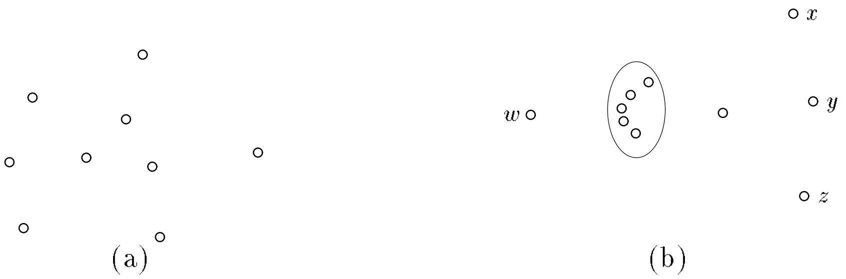

Figure 5.1: Two point sets that are not 4-connectible.

4 Algorithms

It is straightforward to show that constructing a good level-$(j-1)$ triangulation from a good level-j triangulation takes $O(n_j + n_{j-1})$ time where n_j and n_{j-1} are the number of vertices in layers j and $j-1$, respectively. Hence, if S admits a noncomplex triangulation, such a triangulation can be constructed in $O(n)$ time once we have the convex layers of S. Convex layers of S can be constructed in $O(n \log n)$ time using the algorithm of [4]. Also, we can check whether a point set is anomalous in $O(n \log n)$ time by directly checking condition (B1) and, if necessary, (B2). So we have the following theorem.

Theorem 4.1 Given a planar point set S in general position, in $O(n \log n)$ time we can either construct a non-complex triangulation of S if it admits one, or report that no such triangulation exists.

Since the complexity of computing a triangulation of a planar point set is $\Omega(n \log n)$ [10], this algorithm is optimal to within a constant factor.

5 Remarks and open problems

In this paper we have characterized the point sets that admit a noncomplex triangulation. This solves the question of 4-connectibility for a point set with three extreme vertices. However, it does not solve the 4-connectibility problem in general, and this problem remains open.

The two point sets in Figure 5.1 illustrate two ways that a planar point set can fail to be 4-connectible. The set shown in Figure 5.1(a) fails to be 4-connectible because there are fewer interior points than convex hull edges. Any triangulation of this point set must either have a chord, violating condition (A1) of Lemma 2.2, or two triangles having distinct convex hull edges as their bases but sharing an interior point as their common apex. In the latter case, condition (A3) of Lemma 2.2 is violated. Figure 5.1(b) also fails to be 4-connectible, even though it has more interior points than convex hull edges. To see this, note that if a 4-connected triangulation of this set exists, then one of the circled points (call it p) would have to be connected to y; otherwise x and z would have a common interior neighbor, violating (A3). But then p is connected to both w and y, so (A3) fails anyway.

We have not addressed the condition of 5-connectibility. It follows from the results in [8] that a triangulation is 5-connected if it satisfies conditions (A1)-(A3), has no complex (i.e., nonfacial)

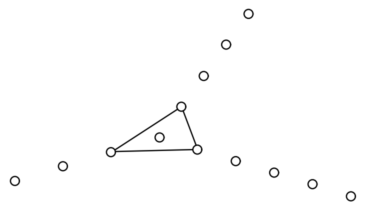

Figure 5.2: A configuration not in general position that does not admit a non-complex triangulation. The complex triangle shown must be present in any triangulation of this point set.

quadrilateral, and has no interior edge connected to two or more nonconsecutive boundary vertices. A simpler problem than general 5-connectibility might be characterizing those planar point sets that admit triangulations without complex quadrilaterals.

Finally, we briefly discuss the general position assumption made in this paper, namely that no three points are collinear. Figure 5.2 illustrates a point set that does not admit a non-complex triangulation. The points lie on three lines. The complex triangle shown in the figure must be present in any triangulation of this point set. The authors conjecture that this is essentially the only non-anomalous point set not admitting a noncomplex triangulation. More precisely, we conjecture that any such point set must consist of points along three rays with a common origin point, and must include the origin point.

Acknowledgment

The authors gratefully acknowledge stimulating discussions with David Eppstein. They also thank Henk Meijer for pointing out an error in an earlier version of the paper.

References

[1] P. K. Agarwal. Partitioning arrangements of lines: II. Applications. *Discrete & Computational Geometry*, 5(6):533–573, 1990.

[2] E. M. Arkin, M. Held, J. S. B. Mitchell, and S. S. Skiena. Hamiltonian triangulations for fast rendering. In J. van Leeuwen, editor, *Algorithms—ESA '94: Second Annual European Symposium*, pages 36–47, Utrecht, The Netherlands, September 1994. Springer-Verlag Lecture Notes in Computer Science 855.

[3] J. A. Bondy and U. S. R. Murty. *Graph Theory with Applications*. North-Holland, New York, NY, 1976.

[4] B. Chazelle. On the convex layers of a planar set. *IEEE Transactions on Information Theory*, 31(4):509–517, July 1985.

[5] G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis. Algorithms for drawing graphs: An annotated bibliography. *Computational Geometry: Theory and Applications*, 4(5):235–282, October 1994.

[6] M. B. Dillencourt. Hamiltonian cycles in planar triangulations with no separating triangles. *Journal of Graph Theory*, 14(1):31–49, 1990.

[7] B. Grünbaum. *Convex Polytopes*. Wiley Interscience, New York, NY, 1967.

[8] J-P. Laumond. Connectivity of planar triangulations. *Information Processing Letters*, 34(2):87–96, March 1990.

[9] N. Megiddo and A. Tamir. On the complexity of locating linear facilities in the plane. *Operations Research Letters*, 1(5):194–197, November 1982.

[10] F. P. Preparata and M. I. Shamos. *Computational Geometry: An Introduction*. Springer-Verlag, New York, NY, 1985.

[11] A. Saalfeld. Joint triangulations and triangulation maps. In *Proceedings of the Third ACM Symposium on Computational Geometry*, pages 195–204, Waterloo, Ontario, June 1987.

[12] H. Whitney. A theorem on graphs. *Annals of Mathematics*, 32:378–390, 1931.

[13] K. Yeap and M. Sarrafzadeh. Floor-planning by graph dualization: 2-concave rectilinear modules. *SIAM Journal on Computing*, 22(3):500–526, June 1993.