HAPPY ENDINGS FOR FLIP GRAPHS *

David Eppstein [†]

ABSTRACT. We show that the triangulations of a finite point set form a flip graph that can be embedded isometrically into a hypercube, if and only if the point set has no empty convex pentagon. Point sets with no empty pentagon include intersections of lattices with convex sets, points on two lines, and several other infinite families. As a consequence, flip distance in such point sets can be computed efficiently.

1 Introduction

Triangulations, partitions of the convex hull of a point set into triangles having the given points as vertices and meeting edge-to-edge, form a fundamental object of study in computational geometry, of importance due to applications including computer graphics and finite element mesh generation. A point set may have many different triangulations, and there are many known algorithms for finding high-quality triangulations using various optimization criteria [3]. Flips, flip distance, and the flip graph (all defined below) form important tools in organizing the set of possible triangulations, and in performing local optimization heuristics for finding good triangulations.

A *flip* in a triangulation is the replacement of two triangles that share an edge by the other two triangles that cover the same quadrilateral. Flipping has long been important in the study of triangulations, both topologically since the discovery that any two triangulations of the same topological surface can be connected to each other by flips [33] and computationally from the fact that repeated flips can be used to compute Delaunay triangulations [26, 29]. However, despite this long study, much about flipping remains mysterious. In particular, it is not known how to compute efficiently the *flip distance* between triangulations, the minimum number of flips needed to transform one triangulation into another. Even for convex point sets, it remains unknown whether flip distances may be computed in polynomial time, although in this case tight linear bounds are known on how large the flip distance can be as a function of the number of points [20, 30]. For nonconvex points, flip distances may grow superlinearly compared to the number of points [5, 17, 21] and it is unknown how to approximate this number efficiently and accurately. The higher dimensional analogue of flip distance is not, in general, well-defined: there exist pairs of triangulations that cannot be converted into each other by flips [28].

*This work was supported in part by NSF grant 0830403 and by the Office of Naval Research under grant N00014-08-1-1015.

[†]Computer Science Department, Donald Bren School of Information and Computer Sciences, University of California, Irvine, eppstein@uci.edu

Figure 1: A high-level view of the flip graph of a 3×3 grid of points. The 64 triangulations of this point set are grouped so that each flip within a group adds or removes an edge between adjacent grid points and the flips between two groups involve longer edges. Within each group, the flips are shown directly as edges connecting pairs of triangulations, while the number of lines between each pair of groups indicates the number of flip graph edges connecting pairs of triangulations in those two groups.

The set of triangulations of a planar point set, and the flip distances between any two triangulations, form a *metric space*: flip distances are symmetric, nonnegative, and satisfy the triangle inequality. The metric space of triangulations and flip distances can be represented combinatorially as an undirected graph, the *flip graph*, that has a vertex per triangulation and an edge per flip between two triangulations (for examples, see Figures 1 and 15). The flip distance between any two triangulations is equal to the length of the shortest path between the corresponding vertices in this graph. However, except for very small numbers of points it is not feasible to compute flip distances by constructing the flip graph explicitly and applying a graph shortest path algorithm: the number of triangulations

Figure 2: A triangulation of a point set in convex position (with a designated root edge of the convex hull shown dashed) and its dual binary tree (with a dashed vertical line segment connecting to the root vertex).

may be exponential in the number of points of the input point set (in fact, for points in general position, the number of triangles is always exponential, due to the existence of linearly many independently flippable edges [21]) so the flip graph is generally too large to construct.

There are also strong connections between flip graphs, rotation in binary trees, and polyhedral combinatorics. If a set of n points is in *convex position* (its points form the vertices of a convex polygon, with one polygon edge chosen arbitrarily as root), then its triangulations are in one-to-one correspondence, by a form of planar duality, to the binary trees with $n - 1$ leaves [30]; see Figure 2 for an example. According to this correspondence, a flip of a triangulation corresponds to a binary tree rotation, a standard operation in the theory of data structures that interchanges the heights of two adjacent tree nodes while preserving the left-to-right traversal ordering of the tree nodes [30]. From this point of view, the flip graphs of convex point sets have been shown to be skeletons of polytopes, called *associahedra* or *Stasheff polytopes* (Figure 15) [31, 32]. Associahedra are examples of *secondary polytopes* [4, 16], which can be defined for arbitrary point sets without restricting the points to be in convex position. However the vertices of the secondary polytope describe only the *regular triangulations*, triangulations that can be formed by orthogonal projection of a higher dimensional convex polyhedron. For some sets of points that are not in convex position, there exist non-regular triangulations. Therefore, in such cases, the vertices and edges of the secondary polytope do not represent the entire flip graph.

In this paper we study flip distances and flip graphs of highly nonconvex point sets, those in which no five points of the set form the vertices of an empty pentagon. Rabinowitz [27] has called this property of having no empty pentagon the “pentagon property”. It is known (a relative of the famous Happy Ending problem of Erdős and Szekeres [15]) that sets of ten or more points in general position must always have an empty pentagon [18], but our results apply to point sets not necessarily in general position; for instance, the intersection of the integer lattice with any bounded convex subset of the plane forms a set with no

empty pentagon. We show that the flip graph of any point set with no empty pentagon is a *partial cube*, a graph that can be embedded isometrically into a hypercube [8]. Conversely any point set with an empty pentagon cannot have a flip graph that is a partial cube. Via this characterization, flip distances for triangulations of point sets without empty pentagons may be calculated easily and efficiently.

Emo Welzl, in an invited talk at Graph Drawing 2006 [34],¹ claimed that, for any $n \times n$ grid of points in the plane, the flip graph of the grid can be represented (not necessarily isometrically) as an induced subgraph of a hypercube. This is a special case of our result, and was the original inspiration for the investigations reported herein.

Subsequent to the preliminary conference publication of this paper [11], Abel *et al.* have also investigated point sets with the empty pentagon property: as they show, sufficiently large point sets with the empty pentagon property have arbitrarily large subsets of collinear points [1]. The analogous result for points with no empty quadrilateral is an immediate consequence of our Theorem 1.

2 New results

We prove the following results:

- A finite point set has exactly one triangulation if and only if it has no empty quadrilateral. Moreover, the finite point sets having this property are exactly the vertex sets of planar graphs with dilation one; in previous work we completely classified these sets.
- The flip graph of a finite point set forms a partial cube if and only if the point set has no empty pentagon. Moreover, the point sets having this property are exactly those for which another associated geometric graph, the *quadrilateral graph*, is a forest.
- If a set of n points has no empty pentagon, it has $O(n^2)$ empty quadrilaterals. There exist point sets for which this bound is tight.
- We can find an empty pentagon in a set of n points, if one exists, or if there is no empty pentagon construct the quadrilateral graph, in time $O(n^2)$.
- If a set of n points has no empty pentagon, then we may compute the flip distance between any two triangulations of the point set in time $O(n^2)$.
- We provide a polynomial-time-computable lower bound on the flip distance between triangulations in any finite point set. When the point set has no empty pentagon, or consists solely of the vertices of a pentagon, this bound is exactly equal to the flip distance. However this bound is inaccurate for any point set containing an empty hexagon. It would be of interest to characterize more precisely the point sets on which this bound is tight.

¹See also Kaibel and Ziegler [22].

Figure 3: A point set and its quadrilateral graph.

In addition we provide several new examples of finite point sets that do not have an empty pentagon. They include: any intersection of a lattice with a bounded convex set, any finite set of points drawn from the union of two lines, any finite set of points drawn from the union of three rays with coincident apexes and no concave angle, several other infinite families of examples, and several sporadic examples.

3 Geometric definitions

For any point set P , let $\text{CH } P$ denote the (closed) convex hull of P . If $\text{CH}(P \setminus \{a\}) \neq \text{CH } P$ we say that a is a *vertex* of the hull. Let $|P|$ denote the cardinality of P . We say that P contains an *empty k -gon* if there exists a subset S of P with $|S| = k$ such that $P \cap \text{CH } S = S$ and such that each point of S is a vertex of $\text{CH } S$. That is, S forms the vertices of a convex polygon which does not include any other points of P . In particular we say that P contains an *empty quadrilateral* when there exists a set S with these properties and with $k = |S| = 4$, and we say that P contains an *empty pentagon* when there exists a set S with these properties and with $k = |S| = 5$. Note that three collinear points may never belong to an empty k -gon, because the middle point of the three belongs to the convex hull but is not a vertex.

We define a *diagonal* of P to be a line segment ab such that $P \cap \text{CH}\{a, b\} = \{a, b\}$; that is, in the notation above, it is an empty 2-gon. Equivalently, a diagonal is an edge in the *visibility graph* of P . A *triangulation* of P is any graph having as its vertices all the points of P and as its edges a maximal set of diagonals that do not cross each other. Removing the edges and vertices of the triangulation from the plane partitions the remaining points of the plane into connected components that we call *faces*. For a maximal set of diagonals, there is one unbounded face (the set-theoretic complement of the convex hull of P) and all remaining faces must be triangles; see, e.g., [3, Lemma 1].

In addition, we define the following geometric graphs for a point set P :

- Let $\text{DT } P$ denote the Delaunay triangulation of P , in which an edge is present if and

only if there exists a closed circle containing no other points of P . We will only consider Delaunay triangulations of point sets in which no four points are cocircular, so the Delaunay triangulation is uniquely defined.

- We define the *quadrilateral graph* $QG\ P$ to be a graph having a vertex for each diagonal of P , and an edge connecting diagonals ab and cd whenever line segments ab and cd cross and $\{a, b, c, d\}$ form the vertices of an empty quadrilateral (Figure 3).
- Let $FG\ P$ denote the *flip graph* of P , a graph with a vertex for each triangulation of P and an edge connecting two triangulations that differ by a single *flip*: the removal of one diagonal and its replacement by another diagonal. Observe that the two diagonals involved in a flip must necessarily be connected by an edge in the quadrilateral graph.

The *flip distance* between any two triangulations is defined to be the minimum number of flips needed to transform one triangulation into the other, or equivalently the unweighted distance between the corresponding vertices in the flip graph.

A *Delaunay flip* of a diagonal ac is a diagonal bd such that bd crosses ac and such that bd belongs to $DT\{a, b, c, d\}$. Any triangulation that is not Delaunay can be transformed into the Delaunay triangulation by replacing diagonals according to a sequence of Delaunay flips [26, 29]. In any point set, each diagonal ac that is not itself part of the Delaunay triangulation has at least one Delaunay flip; for instance, if b and d are chosen to be the points on either side of line ac such that angles abc and adc are as large as possible, then replacing ac by bd is a Delaunay flip.

4 Partial cubes

An *isometric embedding* of a graph G into another graph H is a function f from the vertices of G to the vertices of H that preserves distances: if u and v are any two vertices in G , then $d_G(u, v) = d_H(f(u), f(v))$. If G is isometrically embedded into H , it must form an induced subgraph of H , but not every induced subgraph is isometrically embedded. A *partial cube* is a graph that can be embedded isometrically into a hypercube. Another way of stating this is that, if G is a partial cube, we can label each vertex of G with a bitvector, in such a way that unweighted distance in G equals Hamming distance of labels. Thus, a partial cube structure on a graph enables us to look up distances between vertices easily; see e.g. Chepoi *et al.* [7] for data structures using partial cube embeddings to compute distances in certain graph families. There are many applications and examples of partial cubes [13], and a mathematical characterization of these graphs due to Djoković [8] allows them to be recognized efficiently [12].

Although much is known about partial cubes, for the purposes of this paper we need only some simple observations about them:

- Any partial cube is bipartite. This follows from the facts that a partial cube is isomorphic to an induced subgraph of a hypercube and that every hypercube is bipartite. Alternatively, one can derive a bipartition from the parity of the number of nonzero bits in the bitvectors labeling the vertices of a partial cube.

Figure 4: A planar graph on six vertices with dilation one.

- Any tree is a partial cube. To find a partial cube labelling of a tree, choose a root, label each vertex v by a vector of bits, one bit per edge, and set the bit value for edge e in v 's label to 1 if e belongs to the path from v to the root and 0 otherwise.
- The Cartesian product of partial cubes is another partial cube. The labels for vertices in the product can be formed by concatenating the labels from the factors of the product.
- If G is a partial cube, and G' is embedded isometrically into G , then G' is also a partial cube. The labels for the vertices of G' can be copied from the labels of their images in G .

5 Quadrilaterals, dilation, and unique triangulation

As a warmup to our main result (Theorem 2) describing several equivalent characterizations for point sets with no empty pentagon, we prove a similar equivalence of characterizations for point sets with no empty quadrilateral.

If G is a graph embedded without crossings as a straight line drawing in the plane, and we define the length of each edge in G to be the Euclidean distance between its end-points, then the *dilation* of a pair of vertices u, v in G is the ratio between their distance in G and their Euclidean distance. The dilation of G is the maximum dilation of any two of its vertices. A graph has dilation one (Figure 4) if and only if, between any two vertices, there exists a path connecting them that lies entirely along the line segment between them.

Theorem 1. *The following conditions are equivalent for a finite point set P :*

1. P has no empty quadrilateral.
2. $QG P$ is an independent set.
3. P has only one triangulation.
4. P is the set of vertices of a non-crossing embedded graph with dilation one.

Proof. It is simplest to prove that the complements of these conditions are equivalent.

($\bar{1} \Rightarrow \bar{2}$): If P has an empty quadrilateral, the two diagonals of this quadrilateral form an edge in the quadrilateral graph.

($\bar{1} \Rightarrow \bar{3}$): If P has an empty quadrilateral, there exists a triangulation T that includes the boundary edges of the quadrilateral and either one of its two diagonals (because any non-crossing set of line segments can be extended to a triangulation). But flipping the diagonal of this quadrilateral in T produces a second triangulation; therefore, P has more than one triangulation.

($\bar{1} \Rightarrow \bar{4}$): Suppose that P has an empty quadrilateral. Then any dilation-one graph G on P would have to include line segments connecting opposite vertices of this quadrilateral; these line segments cross within the quadrilateral, contradicting the requirement that G be embedded without crossings. Therefore, no such dilation-one graph exists.

($\bar{2} \Rightarrow \bar{1}$): If $QG\ P$ has an edge, its endpoints correspond to two crossing diagonals in P , and the four endpoints of these diagonals form an empty quadrilateral.

($\bar{3} \Rightarrow \bar{1}$): If P has more than one triangulation, then perturb the point set by an affine transformation if necessary so that the Delaunay triangulation is unique; this perturbation does not affect the existence or nonexistence of empty quadrilaterals. After this perturbation, there exists a triangulation T that is not Delaunay, because there are multiple triangulations but only one Delaunay triangulation. Then T has a Delaunay flip, and the endpoints of the flipped edges form an empty quadrilateral.

($\bar{4} \Rightarrow \bar{3}$): Suppose P is not the vertex set of a dilation-one graph, and let T be any triangulation of P . Since no graph on P has dilation one, this is in particular true of T ; that is, there exists at least one pair of vertices in P such that no path of T lies along the line segment connecting these two vertices. Among pairs of vertices with this property, let ab be the closest such pair. Then there can be no point of P on line segment ab , for if such a point c existed then ac or bc would be a closer pair with the same property. Thus, ab must be a diagonal of P that does not belong to T . There exists a triangulation T' containing segment ab (since any non-crossing set of line segments can be extended to a triangulation); T and T' are different triangulations since one of them contains ab and the other does not. $\square$

In previous work [10,9] we classified all point sets that can be vertex sets of dilation-one planar graphs. They are:

- n collinear points (point sets with no empty triangle).
- $n - 1$ points on a line, and one point off the line.
- $n - 2$ points on a line, and two points on opposite sides of the line, such that the $n - 2$ points and the two points have disjoint convex hulls.
- $n - 2$ points on a line, and two points on opposite sides of the line, such that the $n - 2$ points and the two points have convex hulls that intersect at one of the $n - 2$ points on the line.

Figure 5: Lemma 1: two empty quadrilaterals sharing two edges lead to an empty pentagon.

- A set of six points of a type shown in Figure 4.

Therefore, these are exactly the point sets that have no empty quadrilateral, and that have only one triangulation.

6 Pentagons and partial cubes

Any finite point set can be perturbed by a suitably chosen affine transformation in such a way that no four perturbed points are cocircular, without changing the orientation of any triple of points. In particular, such transformations preserve the set of triangulations of the point set as well as its flip graph. For this reason, we can assume when necessary throughout this section that the point set under consideration has no four cocircular points. With this assumption, we can orient each edge $ab\text{-}cd$ of QGP from ab to cd whenever $DT\{a, b, c, d\}$ has cd as its diagonal. In the next sequence of lemmas we show that the assumption of no empty pentagon causes QGP to have a particularly simple structure: it is a forest in which the edges in each tree are oriented towards a single vertex representing an edge of $DT P$.

Lemma 1. *If a finite point set P contains two empty quadrilaterals that share two edges, then P contains an empty pentagon.*

Proof. The two shared edges must have a vertex in common, for otherwise it would not be possible to form two different quadrilaterals with those edges. Therefore, let the two quadrilaterals have vertices $\{a, b, c, d\}$ and $\{a, b, c, e\}$. By the assumption that both quadrilaterals are empty, neither d nor e can be contained in the other quadrilateral. But the graph formed by the union of the edges of the two quadrilaterals is $K_{2,3}$, whose only planar drawings have one vertex inside the quadrilateral formed by the edges disjoint from it. Therefore, the six edges of the two quadrilaterals do not form a planar drawing, and there

Figure 6: Lemma 2: an empty quadrilateral restricts the location of the apex of an empty triangle to two wedges.

exists a pair of edges, one from each quadrilateral, that cross each other. Without loss of generality let these two crossing edges be cd and ae , as shown in Figure 5.

Note that the set $S = (P \cap \text{CH}\{c, d, e\}) \setminus \{c, d\}$ is nonempty, because it contains e . Among all points in S choose f to be one that minimizes angle cdf ; if two or more points have the same minimum angle, choose f to be the one closest to d . Then $\text{CH}\{a, b, c, d, f\} = \text{CH}\{a, b, c, d\} \cup \text{CH}\{c, d, f\}$. The quadrilateral $\text{CH}\{a, b, c, d\}$ is empty by assumption, and the triangle $\text{CH}\{c, d, f\}$ is empty because any point within it would contradict the choice of f as minimizing angle cdf or as being closest to d with the given angle. Therefore $\text{CH}\{a, b, c, d, f\}$ is empty, and forms an empty pentagon for P . $\square$

Lemma 2. *Suppose that a finite point set P contains no empty pentagon, that a, b, c, d are the vertices of an empty quadrilateral in clockwise order, and that $\{a, c, e\}$ are the vertices of an empty triangle. Let W_1 be the closed wedge bounded by lines ab and ad that is (among the four wedges bounded by these two lines) the one opposite quadrilateral $abcd$, and symmetrically let W_2 be the closed wedge bounded by lines bc and cd that is (among the four wedges bounded by these two lines) the one opposite quadrilateral $abcd$. Then $e \in W_1 \cup W_2$.*

Proof. For any other placement of e , either $abce$ or $aecd$ would be an empty quadrilateral and we could apply Lemma 1 to find an empty pentagon, contradicting the assumption that P contains no empty pentagon. $\square$

An example of the lemma, showing the two wedges in which e must be contained, is illustrated in Figure 6.

Lemma 3. *Let finite point set P contain no empty pentagon and have no four cocircular points, and let ac be a diagonal of P that does not belong to DTP . Then ac has a unique Delaunay flip bd .*

Figure 7: Lemma 3: in a set with no empty pentagons, if bd is a Delaunay flip for ac , then there can be no other Delaunay flip ef .

Proof. Let bd be any Delaunay flip for ac . Diagonal bd belongs to $DT\{a, b, c, d\}$ if and only if there exists a circle O having a and c on its boundary and b and d in its interior, if and only if the sum of angles abc and adc is at least π . If $ac - ef$ is any edge in the quadrilateral graph of P , with $ef \neq bd$, then by Lemma 2 we must have e and f belonging to the two wedges W_1 and W_2 . Necessarily one of e and f must belong to each wedge, for otherwise ef would not cross ac contradicting the assumption that $ac - ef$ belongs to the quadrilateral graph.

Wedge W_1 is bounded by a ray with apex a along line ab and on the opposite side of a from b ; this line crosses O twice, once at a and once on the other side of b from a . Therefore, the boundary ray of W_1 is entirely outside O . By a symmetric argument, the other three boundary rays of W_1 and W_2 are also entirely outside O . Therefore W_1 and W_2 themselves are exterior to O (Figure 7). Since O is an empty circle for ac among the set of four sites $\{a, c, e, f\}$, ac must be an edge of $DT\{a, c, e, f\}$ and ef cannot be a Delaunay flip for ac . $\square$

Lemma 4. *If a finite point set P contains no empty pentagon, QGP is a forest.*

Proof. We perturb P by an affine transformation if necessary so that no four points are cocircular; this perturbation does not change QGP . For each edge $ac-bd$ of QGP , either bd is a Delaunay flip of ac or vice versa. We orient each edge from ac to bd if bd is the Delaunay flip of ac , and the other orientation otherwise. Repeated Delaunay flipping must always lead to the Delaunay triangulation [26, 29], so this orientation is acyclic and (by Lemma 3) has at most one outgoing edge per vertex. Therefore each connected component of QGP is a tree and QGP is a forest. $\square$

Lemma 5. *If finite point set P contains no empty pentagon, and T is any triangulation of P , then T contains exactly one edge from each tree in QGP .*

Figure 8: Lemma 6: if abc is a non-Delaunay empty triangle, then one of its vertices must be an endpoint of the Delaunay flip for the opposite edge.

Proof. We perturb P by an affine transformation if necessary so that no four points are cocircular; this perturbation does not affect the statement of the lemma. According to the orientation from the proof of Lemma 4, each tree in $QG P$ has a unique Delaunay edge (the edge at the root of the tree), so the lemma holds when T is the Delaunay triangulation. For arbitrary T the result follows by induction on the number of flips needed to reach the Delaunay triangulation: each Delaunay flip reduces this number by one and preserves the number of edges per tree in $QG P$. $\square$

Lemma 6. *Let finite point set P have no empty pentagon and no four cocircular points, and let abc be an empty triangle that is not part of $DT P$. Then at least one vertex of abc is an endpoint of the Delaunay flip of the opposite edge.*

Proof. If c is an endpoint for the Delaunay flip for ab , the result holds; otherwise, suppose that the Delaunay flip for ab is de . Swap the names of points a, b or of points d, e if necessary so that ac crosses the boundary of quadrilateral $adbe$ on edge be (Figure 8). Let S be the set $(P \cap CH\{a, c, e\}) \setminus \{a, c\}$; S is non-empty, because it contains e . Let f be the point of S that minimizes angle acf ; if multiple points have the same minimizing angle, choose f to be the one closest to c . Then $abcf$ is an empty quadrilateral, for any point inside it would form a sharper angle than f or be closer to c with the same angle, contradicting the choice of f .

Angle afc is at least equal to angle aec , which exceeds angle aeb . Moreover, angle abc exceeds angle adb , for otherwise c would not belong to the wedge for empty quadrilateral $adbe$ described in Lemma 2. Therefore, $afc + abc > aeb + adb > \pi$ (where the second inequality $aeb + adb > \pi$ characterizes the fact that de is a Delaunay flip for ab); since $afc + abc > \pi$, bf is a Delaunay flip for ac . Therefore, b is an endpoint for the Delaunay flip of the opposite edge ac . $\square$

Define the *flip number* $\#ab$ of a diagonal ab in a point set P with no four cocircular points to be the minimum number of flips needed to reach a triangulation containing ab from $\text{DT } P$. If ab is a Delaunay triangulation edge, define $\#ab = 0$.

Lemma 7. *Let P be a finite point set with no empty pentagon and no four cocircular points. Suppose that abc and acd are empty triangles, intersecting only on non-Delaunay edge ac , such that $\#ac \geq \max\{\#ab, \#bc, \#cd, \#ad\}$. Then bd is the Delaunay flip of ac .*

Proof. If ab and bc are both Delaunay edges, then any sequence of Delaunay flips that produces the Delaunay triangulation from some triangulation containing ac must eventually flip ac , without ever flipping ab and bc ; therefore b is an endpoint of the Delaunay flip for ac . If on the other hand ab is not a Delaunay edge, then b must be an endpoint of the Delaunay flip for ac , for otherwise by Lemma 6 ac would be an edge of the convex hull of the four vertices forming the Delaunay flip for ab or bc , so any flip sequence producing that edge from the Delaunay triangulation would have to contain ac before its final step, contradicting the assumption that $\#ac \geq \max\{\#ab, \#bc\}$. Thus regardless of which edges are Delaunay, b is an endpoint of the Delaunay flip for ac . By a symmetric argument d is also an endpoint of the Delaunay flip; therefore the flip is bd . $\square$

Theorem 2. *The following three conditions are equivalent for a finite point set P :*

1. P has no empty pentagon.
2. $\text{QG } P$ is a forest.
3. $\text{FG } P$ is a partial cube.

Proof. $(2 \Rightarrow 1)$: If $\text{QG } P$ is a forest, P must have no empty pentagon, because the diagonals of an empty pentagon would form a cycle of five vertices in $\text{QG } P$, but a forest has no cycles.

$(1 \Rightarrow 2)$: Lemma 4.

$(3 \Rightarrow 1)$: If $\text{FG } P$ is a partial cube, P must have no empty pentagon, because the diagonals of an empty pentagon can be flipped one for another to form a cycle of five vertices in $\text{FG } P$, but every partial cube is bipartite.

$(1 \Rightarrow 3)$: In the remainder of the proof we show that, if P has no empty pentagon, $\text{FG } P$ is a partial cube. Let Π be the partial cube formed as the Cartesian product of the trees in the forest $\text{QG } P$. A vertex in Π can be identified by selecting one vertex for each tree in $\text{QG } P$. By Lemma 5, any triangulation T has one diagonal corresponding to a vertex in each tree of $\text{QG } P$, so T corresponds to a vertex in Π . We show below that this identification of triangulations with vertices in Π is isometric, by showing both upper and lower bounds relating flip distance to distance in Π .

We begin by showing that flip distance is at least equal to distance in Π . If two triangulations T_1 and T_2 differ by a flip, then the vertices in $\text{QG } P$ representing their sets of diagonals differ only in a single tree of $\text{QG } P$, and their images of T_1 and T_2 in Π are adjacent. More generally, let T_1 and T_2 be any two triangulations of P , at flip distance δ . In any sequence of δ flips changing T_1 to T_2 , each flip corresponds in Π to an adjacency

between vertices, so (by the triangle inequality for distances in Π) the distance between the images of T_1 and T_2 in Π is at most δ .

To complete the proof, we show that flip distance is at most equal to distance in Π , by showing that any two triangulations can be connected by a short sequence of flips. Suppose T_1 and T_2 have images in Π that are δ apart. We perturb the input by an affine transformation if necessary so that no four points are cocircular, allowing us to define Delaunay flips and flip numbers; let ab be the edge in the symmetric difference of T_1 and T_2 maximizing $\#ab$. The two triangles based on this edge (in whichever of T_1 or T_2 contains the edge) meet the preconditions of Lemma 7, so there exists a Delaunay flip of ab in the triangulation T_i containing ab . This Delaunay flip must change T_i to a triangulation in which the representative vertex in QGP is closer than ab to the representative vertex in the other triangulation, so the flipped triangulation has distance $\delta - 1$ from the other triangulation. By repeating this choice of flipping the edge with the maximum flip number, we can find a sequence of δ flips that transforms T_1 to T_2 , so the flip distance is at most equal to δ .

Since the flip distance and the distance in Π are both upper bounded by the other, they are equal. Since we have embedded FGP isometrically in a partial cube Π , FGP is itself a partial cube. $\square$

For a set P of n points with no empty pentagon, QGP has a number of edges less than its number of vertices, as it is a forest. But there are $O(n^2)$ vertices in QGP , because each corresponds to a pair of points in P . Thus, in such a point set, there can be at most $O(n^2)$ empty quadrilaterals. Placing $n/2$ points on each of two parallel lines forms a set of n points with no empty pentagon and $(n/2 - 1)^2$ empty quadrilaterals, showing that this quadratic bound is tight.

7 Examples

We describe briefly some examples of finite point sets without empty pentagons, showing that they exhibit great variety despite the need for some collinearities. We do not have a complete classification of such sets, but they appear to allow significantly greater complexity than the point sets without empty quadrilaterals classified in earlier work and described earlier in this paper.

- Any point set with no empty quadrilateral clearly also has no empty pentagon. Thus, all the types of point sets with no empty quadrilateral, as described in Section 5, also form examples of point sets with no empty pentagon.
- If P is a point set with no empty pentagons, and C is any convex set in the plane, then $P \cap C$ again has no empty pentagons. In addition, any projective transformation of P that preserves the orientation of all its triangles results in a point set with no empty pentagons.
- The union of two lines has no convex pentagons at all, let alone empty convex pentagons (Figure 9, far left). Therefore, one can form finite sets of points with no convex

Figure 9: Continuous point sets with no pentagons: (a) two lines, (b) three rays, (c) two opposite wedges, (d) a wedge and a line segment, and (e) a quadrilateral and a line segment.

pentagons by taking any finite subset of the union of two lines.

- Similarly, if three rays share a common apex and the three wedges bounded by these rays each form angles less than π (Figure 9, near left), the points from these three rays do not form any convex pentagons, and one can form finite point sets with no convex pentagon by taking any finite subset of these three rays.
- Let W_1 and W_2 be two wedges, with $\bar{W}_1$ and $\bar{W}_2$ the complementary wedges bounded by the other two rays on the same lines. Suppose further that the apex of W_1 is contained in $\bar{W}_2$ and the apex of W_2 is contained in $\bar{W}_1$. Let V_1 be a subset of the boundary of W_1 that is entirely contained in $\bar{W}_2$, and let V_2 be a subset of the boundary of W_2 that is entirely contained in $\bar{W}_1$ (Figure 9, center). Then $V_1 \cup V_2$, and any of its finite subsets, do not contain the vertices of any convex pentagon.
- Let V be the boundary of a wedge W , and S be a line segment contained in the complementary wedge $\bar{W}$, such that the line containing S is disjoint from W (Figure 9, near right). Then again $V \cup S$ and its finite subsets do not form any convex pentagons.
- The vertices of a convex quadrilateral, and any subset of a line segment that connects two opposite sides of the quadrilateral, do not form any convex pentagon (Figure 9, far right).
- Any lattice in the plane (for instance the lattice of points with integer coordinates) has no empty pentagons [2, 24, 25]. This follows immediately from the observation that there are only four possible combinations of the parity of the x -coordinate and the parity of the y -coordinate of a grid point. Therefore, among any five grid points, some two have the same parity for both coordinates, and their midpoint is also a grid point, causing the pentagon formed by the five points to be non-empty. The stronger assertion that any lattice pentagon $abcde$ has another lattice point within the small pentagon $S(a, b, c, d, e)$ formed by its inner diagonals (Figure 10) has supposedly appeared as a problem in the Russian mathematics olympiad, and can be solved by a simple case analysis:

Figure 10: Any lattice pentagon has another lattice point in it, and more specifically in the smaller closed pentagon formed by its inner diagonals.

- Suppose we have five points a, b, c, d, e forming (in clockwise order) a convex pentagon in a grid, with m the midpoint of some two points with the same parity. If m lies within $S(a, b, c, d, e)$, as shown in the left and center of Figure 10, we are done.
- If m lies on one of the internal diagonals of pentagon $abcde$, but does not belong to the inner pentagon $S(a, b, c, d, e)$, we may assume without loss of generality that m is the midpoint of segment ac . Then one of the two pentagons $abmde$ or $mbcde$ is convex (depending on the relative position of the inner pentagon and m on segment ac) and has an inner pentagon that is a subset of $S(a, b, c, d, e)$. By induction on the number of grid points contained in the given pentagon, this new smaller pentagon has a grid point in its inner pentagon, which must therefore also belong to $S(a, b, c, d, e)$.
- In the remaining case, m is a midpoint of one of the outer edges of the pentagon, as shown in the center and right of Figure 10; without loss of generality m lies on edge ab . Now, $S(m, b, c, d, e)$ is not in general a subset of $S(a, b, c, d, e)$, but it is a subset of $S(a, b, c, d, e) \cup T$ where T is the triangle formed by lines ac , bd , and be . As before, by induction, $S(m, b, c, d, e)$ contains a lattice point f . If f is in $S(a, b, c, d, e)$, we are done. If f is not in $S(a, b, c, d, e)$, it is in T , and $afcde$ is a convex pentagon such that $S(a, f, c, d, e)$ is a subset of $S(a, b, c, d, e)$; by induction a third time, $S(a, f, c, d, e)$ contains a lattice point, which is also a lattice point in $S(a, b, c, d, e)$.

Thus, intersecting the lattice with a bounded convex set results in a finite set P without empty pentagons. An example of a flip graph for a point set formed from a 3×3 square subset of a lattice can be seen in Figure 1.

- One may also intersect the lattice with a bounded convex set K and then remove any subset of points that lie on the boundary of the convex hull of the intersection, without changing the property of having no empty pentagons. The resulting reduced set P has no empty pentagon. To see this, suppose that E is any empty convex polygon in P . We consider the following cases:
 - If none of the removed boundary points lies on the boundary of E , then the grid points in E are the same as those in the intersection of the grid with K , which has no empty convex pentagon as above. Therefore, E has at most four sides.

- If E has two consecutive sides ab and bc from which boundary points have been removed, let p be the removed point closest to b on side ab and let q be the removed point closest to c on side ac , and consider the grid point $r = p + q - b$ (the fourth point of a parallelogram with vertices p , b , and q). Because distance ab is an integer multiple of distance pb , a is at least twice as far from b as p is, and similarly c is at least twice as far from b as q is. Therefore, r must lie within the convex hull of abc , so the only way for E to be empty is for abc to be a triangle with r on its boundary edge ac . , we may find an affine transformation of the plane such that that the lattice from which P is formed is the integer lattice, with these two consecutive sides of E on the coordinate axes of the plane and E entirely within the positive quadrant. That is, the two sides of E are the line segments from the origin to the points $(x, 0)$ and $(0, y)$ for some integers x and y , both greater than one. But then E must contain the point $(1, 1)$; for E to be an empty polygon with respect to P , the point $(1, 1)$ must be one of the removed points, $x = y = 2$, and E must be a triangle.
- In the remaining case E has a single side bc from which a boundary point has been removed, let ab be the other side of E with vertex b , let p be the removed point closest to b on segment bc , and let $q = a + p - b$ be the fourth point of a parallelogram with vertices a , b , and p . This parallelogram cannot contain any grid points other than its vertices, because if it did then by point-reflection symmetry the half-parallelogram abp would contain a grid point, which would lie in the interior of E and therefore could not be one of the removed points; this contradicts the assumption that E is empty. Thus, $abpq$ is a period parallelogram of the integer lattice, and we may perform an affine transformation (without changing the emptiness of E) to bring a to the point $(0, 1)$, b to $(0, 0)$, p to $(1, 0)$, and q to $(1, 1)$. Then E is entirely contained within the positive quadrant, and all vertices of E must lie on the two lines $y = 0$ and $y = 1$, because for any other point r in the positive quadrant the convex hull of $abcr$ would contain q in its interior. But any convex polygon whose vertices lie on two lines can only be a triangle or a quadrilateral.

Therefore, as claimed, the sets formed in this way also have the empty pentagon property.

- The point sets in Figure 11, derived from nested regular polygons, do not have any empty pentagons. In each case an additional level of nesting would lead to an empty pentagon.
- Let P be a finite subset of the union of two lines, and let Q be a finite subset of a single line, such that any empty triangle of P contains a point from Q . Then $P \cup Q$ contains no empty pentagons, for any pentagon must have three noncollinear vertices from P and some point from Q would cause it to be nonempty (Figure 12). Similar constructions are possible in which P is a subset of one of the other pentagon-free configurations in Figure 9, and Q is an appropriate subset of a line segment separating the two components of the configuration.

Figure 11: Point sets with no empty pentagon derived from nested regular polygons.

Figure 12: A finite subset P of the union of two lines (the dark circles) and a collinear point set Q (light circles) that includes a point in every empty triangle of P .

- Kára et al. [23] observe that the infinite set of points with Cartesian coordinates $(12i, 6)$, $(3i, 3)$, $(4i, 2)$, or $(6i, 0)$ for integer values of i form a set with no empty pentagon; the same is true for the finite sets formed by intersecting this infinite set with any bounded convex set.

8 Finding pentagons and flips

Our classification of point sets with no empty convex quadrilateral allows us to determine whether a point set has an empty quadrilateral, in linear time. For larger constant values of k , an algorithm of Eppstein *et al.* [14] allows the existence of an empty convex k -gon to be determined in time $O(n^3)$. We now describe a simple technique that solves the problem of testing for the existence of an empty pentagon more quickly, in $O(n^2)$ time.

Lemma 8. *Let q be the unique point with minimum y -coordinate in a point set P , and suppose that the remaining points are given in sorted order according to the slopes of the lines they form with q . Then we can find an empty pentagon in P with q as one of its vertices, if such a pentagon exists, in linear time.*

Proof. If two points in P form the same slope with q , remove all but the point nearest q ; this does not change the existence of an empty pentagon with q as vertex, because none of the removed points can be in an empty polygon with q . Connect the remaining points by a polygonal chain in the sorted order, and connect the ends of this chain to q ; this forms a simple polygon each point of which is visible from q . As long as this polygon contains a concave vertex other than q , simplify the polygon by removing such a vertex and adding a diagonal connecting its two neighbors; this simplification process produces a partition of the initial polygon into a collection of triangles, disjoint from q , and a *pseudotriangle* (a simple polygon with three non-reflex vertices) of which q is a vertex (Figure 13). The planar dual graph of this partition is a tree, with the pseudotriangle containing q as root, and each triangle of this partition is visible from q through the triangle edge connecting it to its parent in this dual tree.

We assert that P has an empty pentagon, with q as vertex, if and only if some two adjacent triangles in this partition (neither having q as vertex) form a convex quadrilateral. In one direction, it is clear that any such quadrilateral forms with q a convex pentagon. In the other direction, suppose that there is no convex quadrilateral of this type, and consider any four points a, b, c, d in $P \setminus \{q\}$. As the diagonals added in constructing our partition form an outerplanar graph, some two of these points, say a and b , must not be connected by a diagonal of the partition, because outerplanar graphs cannot contain a complete graph K_4 on four vertices. Let A be the region containing a that is closest to q in the partition, and similarly let B be the region containing B that is closest to q . Let U denote the path in the tree dual to the partition connecting A to B , corresponding to a collection of regions in our partition (shown shaded in the figure), and connect a to b by a path u of diagonals along the part of the boundary of U farthest from q . Then, by the condition that no two adjacent triangles in the partition be convex, u follows a concave angle at each of its vertices, and u together with edges qa and qb forms a pseudotriangle. As $\text{CH}\{q, a, b\}$ is nonempty, a, b, c, d , and q cannot form part of an empty pentagon. Therefore, there can be no empty pentagon with q as boundary.

The partition into triangles, and the test for the existence of two adjacent triangles forming a convex quadrilateral, are easily performed in linear time. $\square$

To perform this computation efficiently, we need for each point q the radial sorted order of the points above it. The algorithm below finds these sorted orders efficiently by applying projective duality to transform the problem into one of constructing a line arrangement [6]; it is by now a standard component in the computational geometry repertoire, but we include it for completeness.

Lemma 9. *Suppose we are given a set of points, no two of which have the same y -coordinate. Then we may determine, for each point q in the set, the radial sorted order around q of the points with greater y -coordinate, in a total time of $O(n^2)$.*

Figure 13: A set of points, connected in radial order as viewed from its bottom vertex q to form a polygon, and partitioned into triangles by a maximal set of diagonals disjoint from q . No two adjacent triangles form a convex quadrilateral, from which it follows that any two vertices a and b can be connected by a concave chain of diagonals. The shaded region corresponds to a path connecting a to b in the dual graph of the partition.

Proof. We assume without loss of generality (by perturbing the points if necessary) that no two input points determine a vertical line. If the perturbation consists of a small enough rotation of the plane, this will not change the radial sorted ordering; here “small enough” means that the y -coordinate ordering does not change. To find a suitable rotation angle, find a line ℓ determined by some two input points that is neither horizontal nor vertical, but that forms an angle that is as sharp as possible with one of the coordinate axes; this may easily be done in $O(n^2)$ time. Then, a rotation of the input set by half the angle formed by ℓ with a coordinate axis will preserve the y -ordering of the points and eliminate all vertical lines determined by pairs of points.

We consider two planes, which we call the primal and the dual. The primal plane is given by coordinates x, y , and contains (non-vertical) lines defined by equations of the form $y = mx + b$, for pairs of numbers m, b . Analogously, the dual plane is given by coordinates X, Y , and contains (non-vertical) lines defined by equations of the form $Y = MX + B$, for pairs of numbers M, B . To each of our input points (x, y) we associate the dual line determined by the numbers $M = -x, B = y$, and to each line $y = mx + b$ in the primal plane we associate the dual point $X = m, Y = b$. Then, with this association, $y = mx + b$ if and only if $Y = MX + B$; that is, an input point lies on some primal line if and only if the line dual to the input point contains the point dual to the primal line. Thus, the dual points where any two dual lines cross correspond to the primal lines determined by two input points, and the sorted order of these dual points by their X -coordinates (that is, the order in which they appear along the dual lines) is the same as the slope ordering of the lines determined by pairs of points in the primal plane.

Thus, by forming the set of lines in the dual plane corresponding to the input points in the primal plane, and applying an algorithm for constructing line arrangements [6], we may find in total time $O(n^2)$, for each point q , the sorted order of all the other points by

the slopes of the lines they form with q . We may then remove from this sorted order the points below q , and move the subsequence of lines with negative slope to the start of the order, to form the radial ordering of the points above q . $\square$

Theorem 3. *We may test whether any set of n points contains an empty pentagon, in time $O(n^2)$.*

Proof. We rotate the point set if necessary, as in the proof of Lemma 9, so that no two points have the same y -coordinate. We then apply Lemma 8 to each point q in the point set, and to the set of points having larger y -coordinates than each q . Sorting all of the input points radially around each point q may be performed in time $O(n^2)$ by Lemma 9, after which the time to process each point q is linear. $\square$

Theorem 4. *If a set P of n points has no empty pentagon, we can construct QG P in time $O(n^2)$.*

Proof. As in the algorithm for finding an empty pentagon, we rotate the points so that no two have the same y -coordinate, sorting the points above each point q radially around q , and triangulate the polygon formed by connecting the points in this radial order. This triangulation must be unique, for otherwise it would be possible to flip some two of its triangles, but then (as in the proof of Lemma 8) there would be an empty pentagon in P .

A similar argument to that in Lemma 8 shows that, in an input set with no empty pentagon, a set of four points forms the endpoints of the diagonals in a flip, with q having the minimum y -coordinate of the four points, if and only if the other three points form one of the triangles in the triangulated polygon for q . In more detail, let a , b , and c be any three points in radial order abc as viewed from q . In one direction, suppose some two of these three points are not visible to each other within the polygon. then $abcq$ cannot be an empty convex quadrilateral and does not represent a flip in QG; in this case, also, abc cannot be a triangle in the triangulation.

In the other direction, suppose that all three points a , b , and c are visible to each other. Then diagonal ac must pass within the polygon, below point b , so $abcq$ is empty and convex and represents a flip in QG. In this case, by the visibility of its three vertices, abc is a triangle in some triangulation of the polygon, and by the uniqueness of the triangulation it is a triangle in the constructed triangulation.

Thus, in total time $O(n^2)$ we may find the flips forming all of the edges in QG. Combining these edges to form an adjacency list representation of the graph QG itself is straightforward within the same time bound. $\square$

9 Computing and estimating flip distance

From the partial cube embedding in Theorem 2 it follows that the flip distance between two triangulations T and T' in a point set P with no empty pentagon may be computed by finding the edges e from T and e' from T' that correspond to vertices in each component

Figure 14: A hexagon (left), its quadrilateral graph (center), and the complete bipartite graph of distances in the quadrilateral graph from $\{ac, ae, ce\}$ to $\{bd, bf, df\}$ (right).

of $QG P$, computing the distance between the corresponding pair of vertices in $QG P$, and adding the numbers obtained in this fashion for each component of $QG P$. We observe that this distance computation can be implemented as an algorithm that runs in time $O(n^2)$ without explicitly constructing $QG P$: flip both triangulations to the Delaunay triangulation, find the flips that are used in one but not both of these two flip sequences, and return the number of these flips.

We now describe a technique for estimating flip distances heuristically in any point set.

Theorem 5. *Let T and T' be any two triangulations of the same point set P . Form a complete bipartite graph with the diagonals of T on one side and the diagonals of T' on the other, and label the edge between any two diagonals by the distance between those diagonals in $QG P$. Let M be the minimum weight of a perfect matching in this complete bipartite graph. Then M can be computed in polynomial time, and provides an underestimate of the true flip distance between T and T' .*

Proof. For any sequence of flips that converts one triangulation T to another T' , we may partition the flips into a collection of paths in $QG P$, in which each path starts at a vertex of $QG P$ representing a diagonal of T and ends at a vertex representing a diagonal of T' : at the start of the sequence of flips form a length-0 path for each diagonal in T , and whenever flipping a diagonal xy to another diagonal uv , concatenate the quadrilateral graph edge $xy-uv$ onto the end of the path that previously ended in xy . At each stage of this procedure the current set of paths each start at a vertex representing a diagonal of the current triangulation, so when this procedure finishes processing all flips it will have formed the desired set of paths.

Thus, the total number of flips in the flip sequence must be equal to the total lengths of these paths. The edges in the matching correspond to a set of paths in $QG P$ connecting the same endpoints, with minimum total length among all possible such sets of paths; therefore, the total weight of the matching is an underestimate of the length of any

Figure 15: The flip graph of a hexagon. Note that the top and bottom triangulations are at distance four from each other.

flip sequence. □

For instance, consider the hexagon shown on the left of Figure 14, with QGP shown in the center of the same figure. There are two triangulations of the hexagon that use none of its long diagonals: one, T , using diagonals $\{ac, ae, ce\}$, and the other, T' , using diagonals $\{bd, bf, df\}$. On the right of Figure 14 is a complete bipartite graph with the diagonals $\{ac, ae, ce\}$ on one side of the bipartition and the diagonals $\{bd, bf, df\}$ on the other side. It is easily seen that the minimum weight matching in this complete bipartite graph has total weight three. Thus, we can conclude that any flip sequence transforming T to T' requires at least three flips. However, the actual flip distance between T and T' in this example is four, as can be seen from the flip graph of the hexagon (Figure 15).

This estimate may be useful in A^* -algorithm based heuristic search for the true flip distance between two triangulations, as A^* depends on the existence of an estimate that is guaranteed to be an underestimate [19]. Additionally, there are some point sets P for which this method is guaranteed to produce the flip distance correctly rather than yielding an inaccurate estimate. For instance, if P has no empty pentagons, then a pair of diagonals from T and T' will have finite distance in QGP if and only if they belong to the same tree in QGP , so the minimum weight perfect matching will match each diagonal of T with the unique diagonal of T' in the same tree of QGP . Thus, in this case, our estimate ends up calculating the flip distance exactly. If P is itself a pentagon, also, this estimate can be seen by a simple case analysis to calculate the flip distance exactly. However,

whenever P contains an empty hexagon, there exist triangulations T and T' for which M is a strict underestimate of the flip distance, because we can embed the two triangulations with diagonals $\{ac, ae, ce\}$ and $\{bd, bf, df\}$ from the example above into the empty hexagon of P . Thus, the nonexistence of an empty hexagon is a necessary condition for our estimate to be exact. It would be of interest to determine whether it is also a sufficient condition.

References

- [1] Z. Abel, B. Ballinger, P. Bose, S. Collette, V. Dujmović, F. Hurtado, S. D. Kominers, S. Langerman, A. Pór, and D. R. Wood. Every large point set contains many collinear points or an empty pentagon, [arXiv:0904.0262](https://arxiv.org/abs/0904.0262). Preprint, 2009.
- [2] J. R. Arkininstall. Minimal requirements for Minkowski's theorem in the plane, I. *Bull. Austral. Math. Soc.* 22:259–274, 1980, [doi:10.1017/S0004972700006560](https://doi.org/10.1017/S0004972700006560).
- [3] M. W. Bern and D. Eppstein. Mesh generation and optimal triangulation. *Computing in Euclidean Geometry*, 2nd edition, pp. 47–123. World Scientific, Lecture Notes Series on Computing, 1995, <http://www.ics.uci.edu/~eppstein/pubs/BerEpp-CEG-95.pdf>.
- [4] L. J. Billera, P. Filliman, and B. Sturmfels. Constructions and complexity of secondary polytopes. *Advances in Math.* 83(2):155–179, 1990, [doi:10.1016/0001-8708\(90\)90077-Z](https://doi.org/10.1016/0001-8708(90)90077-Z).
- [5] P. Bose and F. Hurtado. Flips in planar graphs. *Computational Geometry: Theory & Applications* 42(1):60–80, 2009, [doi:10.1016/j.comgeo.2008.04.001](https://doi.org/10.1016/j.comgeo.2008.04.001), http://www.scs.carleton.ca/research/tech_reports/2006/download/TR-06-09.pdf.
- [6] B. Chazelle, L. J. Guibas, and D. T. Lee. The power of geometric duality. *BIT* 25(1):76–90, 1985, [doi:10.1007/BF01934990](https://doi.org/10.1007/BF01934990).
- [7] V. Chepoi, F. Dragan, and Y. Vaxès. Center and diameter problems in plane triangulations and quadrangulations. *Proc. 13th ACM-SIAM Symp. Discrete Algorithms*, pp. 346–355, January 2002, <http://www.cs.kent.edu/~dragan/centdiam5.ps>.
- [8] D. Z. Djoković. Distance preserving subgraphs of hypercubes. *J. Combinatorial Th., Ser. B* 14:263–267, 1973, [doi:10.1016/0095-8956\(73\)90010-5](https://doi.org/10.1016/0095-8956(73)90010-5).
- [9] V. Dujmović, D. Eppstein, M. Suderman, and D. Wood. Drawings of planar graphs with few slopes and segments. *Computational Geometry: Theory & Applications* 38:194–212, 2007, [doi:10.1016/j.comgeo.2006.09.002](https://doi.org/10.1016/j.comgeo.2006.09.002), [arXiv:math.CO/0606450](https://arxiv.org/abs/math.CO/0606450).
- [10] D. Eppstein. Dilation-free planar graphs. Web page <http://www.ics.uci.edu/~eppstein/junkyard/dilation-free/>, 1997.
- [11] D. Eppstein. Happy endings for flip graphs. *Proc. 23rd Symp. Computational Geometry*, pp. 92–101. ACM, 2007, [arXiv:cs.CG/0610092](https://arxiv.org/abs/cs.CG/0610092).

- [12] D. Eppstein. Recognizing partial cubes in quadratic time. *Proc. 19th Symp. Discrete Algorithms*, pp. 1258–1266. SIAM, January 2008, [arXiv:0705.1025](https://arxiv.org/abs/0705.1025).
- [13] D. Eppstein, J.-C. Falmagne, and S. Ovchinnikov. *Media Theory*. Springer-Verlag, 2007.
- [14] D. Eppstein, M. Overmars, G. Rote, and G. Woeginger. Finding minimum area k -gons. *Discrete Comput. Geom.* 7(1):45–58, 1992, [doi:10.1007/BF02187823](https://doi.org/10.1007/BF02187823), <http://www.ics.uci.edu/~eppstein/pubs/EppOveRot-DCG-92.pdf>.
- [15] P. Erdős and G. Szekeres. A combinatorial problem in geometry. *Compositio Math* 2:463–470, 1935, http://archive.numdam.org/article/CM_1935__2__463_0.pdf.
- [16] I. M. Gel’fand, M. M. Kapranov, and A. V. Zelevinsky. Newton polyhedra of principal A -determinants. *Soviet Math. Dokl.* 308:20–23, 1989.
- [17] S. Hanke, T. Ottman, and S. Schuierer. The edge-flipping distance of triangulations. *J. Universal Comput. Sci.* 2(8):570–579, 1996, [doi:10.3217/jucs-002-08-0570](https://doi.org/10.3217/jucs-002-08-0570).
- [18] H. Harborth. Konvexe Fünfecke in ebenen Punktmengen. *Elem. Math.* 33(5):116–118, 1978.
- [19] P. E. Hart, N. J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Trans. Systems Science and Cybernetics* SSC4(2):100–107, 1968, [doi:10.1109/TSSC.1968.300136](https://doi.org/10.1109/TSSC.1968.300136).
- [20] F. Hurtado and M. Noy. Graph of triangulations of a convex polygon and tree of triangulations. *Computational Geometry: Theory & Applications* 13(3):179–188, 1999, [doi:10.1016/S0925-7721\(99\)00016-4](https://doi.org/10.1016/S0925-7721(99)00016-4).
- [21] F. Hurtado, M. Noy, and J. Urrutia. Flipping edges in triangulations. *Discrete Comput. Geom.* 22(3):333–346, 1999, [doi:10.1007/PL00009464](https://doi.org/10.1007/PL00009464), http://www.math.unam.mx/~urrutia/online_papers/Flipping.pdf.
- [22] V. Kaibel and G. M. Ziegler. Counting lattice triangulations. *Surveys in Combinatorics 2003*, 307 edition, pp. 277–308. Cambridge University Press, Lond. Math. Soc. Lect. Note Ser., 2003, [arXiv:math.CO/0211268](https://arxiv.org/abs/math.CO/0211268).
- [23] J. Kára, A. Pór, and D. R. Wood. On the chromatic number of the visibility graph of a set of points in the plane. *Discrete Comput. Geom.* 34(3):497–506, 2005, [doi:10.1007/s00454-005-1177-z](https://doi.org/10.1007/s00454-005-1177-z).
- [24] L. F. Klosinski, G. L. Alexanderson, and L. C. Larson. The fifty-first William Lowell Putnam mathematical competition. *Amer. Math. Monthly* 98(8):719–727, 1991, <http://www.jstor.org/stable/2324423>.
- [25] L. C. Larson and G. Gilbert. Problems. *Math. Mag.* 66(5):338–345, 1993, <http://www.jstor.org/stable/2690520>.
- [26] C. L. Lawson. Generation of a triangular grid with applications to contour plotting. Memo 299, Jet Propulsion Laboratory, 1972.

- [27] S. Rabinowitz. Consequences of the Pentagon Property. *Geombinatorics* 14:208–220, 2005, <http://www.mathpropress.com/stan/bibliography/consequences.pdf>.
- [28] F. Santos. A point set whose space of triangulations is disconnected. *J. Amer. Math. Soc.* 13(3):611–637, 2000, [doi:10.1090/S0894-0347-00-00330-1](https://doi.org/10.1090/S0894-0347-00-00330-1).
- [29] R. Sibson. Locally equiangular triangulations. *The Computer J.* 21(3):243–245, 1973, [doi:10.1093/comjnl/21.3.243](https://doi.org/10.1093/comjnl/21.3.243).
- [30] D. D. Sleator, R. E. Tarjan, and W. P. Thurston. Rotation distance, triangulations and hyperbolic geometry. *J. Amer. Math. Soc.* 1(3):647–681, 1988, [doi:10.2307/1990951](https://doi.org/10.2307/1990951).
- [31] J. D. Stasheff. Homotopy associativity of H -spaces, I. *Trans. Amer. Math. Soc.* 108:275–292, 1963, [doi:10.2307/1993608](https://doi.org/10.2307/1993608).
- [32] D. Tamari. The algebra of bracketings and their enumeration. *Nieuw Arch. Wisk. (3)* 10:131–146, 1962.
- [33] K. Wagner. Bemerkungen zum Vierfarbenproblem. *J. Deutsch. Math.-Verein.* 46:26–32, 1936.
- [34] E. Welzl. The number of triangulations on planar point sets. *Graph Drawing, 14th Int. Symp., GD 2006*, 4372 edition, pp. 1–5. Springer-Verlag, Lecture Notes in Computer Science, 2007, [doi:10.1007/978-3-540-70904-6_1](https://doi.org/10.1007/978-3-540-70904-6_1), http://www.inf.ethz.ch/personal/emo/PublFiles/Triangulations_LNCS4372%_GD06.pdf.