Plane geometric graph augmentation: a generic perspective*

Ferran Hurtado[†] Csaba D. Tóth[‡]

August 18, 2011

Abstract

Graph augmentation problems are motivated by network design, and have been studied extensively in optimization. We consider augmentation problems over plane geometric graphs, that is, graphs given with a crossing-free straight-line embedding in the plane. The geometric constraints on the possible new edges render some of the simplest augmentation problems intractable, and in many cases only extremal results are known. We survey recent results, highlight common trends, and gather numerous conjectures and open problems.

1 Introduction

Let $G = (V, E)$ be a graph. We say that a second graph $G' = (V, E \cup E')$ obtained by adding a set E' of edges to G is an *(edge) augmentation* of G. The goal of this operation is to ensure that the augmented graph G' has some desired property. Usually, one would like to achieve the goal at a minimum cost, which is typically measured by the *number* of new edges, although weighted versions are also possible. In this survey, we consider edge augmentation only, but we note that in general one could augment a graph with both new *vertices* and edges, or even *subdivide* an edge (by replacing an edge with a path).

A *geometric graph* $G = (V, E)$ is a graph drawn in the plane such that the vertex set V is a set of points in the plane, and the set of edges E consists of line segments with endpoints in V, whose relative interiors are disjoint. Two edges of a geometric graph *cross* if they have an intersection point lying in the relative interior of both edges. We consider *crossing-free* (or *noncrossing*) geometric graphs, where no two edges cross. The terms *plane geometric graph*, *(crossing-free) segment configuration*, and *planar straight-line graph* (for short PSLG) will also be used here as synonyms for crossing-free geometric graphs. Rather than using only one of these terms in this panoramic paper, we use them all interchangeably, as otherwise a reader who follows the references may be confused by the diversity of the terminology in the literature.

*Research by Ferran Hurtado has been partially supported by projects MEC MTM2009-07242, and Gen. Catalunya DGR 2009SGR1140. Research by Csaba Tóth has been partially supported by NSERC grant RGPIN 35586.

†Departament de Matemàtica Aplicada II, Universitat Politècnica de Catalunya (UPC), Barcelona, Spain. *ferran.hurtado@upc.edu*

‡Department of Mathematics and Statistics, University of Calgary, Calgary, AB, Canada. Research by Tóth has been conducted at Tufts University, Medford, MA, USA. *cdtoth@ucalgary.ca*

Compatibility and visibility. Two crossing-free geometric graphs, $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$, are *compatible* if their union $(V_1 \cup V_2, E_1 \cup E_2)$ is also a crossing-free geometric graph. In this survey we focus on augmentation problems in which a geometric graph $G = (V, E)$ is augmented to a graph compatible with G. For example, G could be a set of disjoint segments and we may want to add new segments among the endpoints in order to obtain a crossing-free spanning tree with certain desirable properties. Alternatively, we may be given an arbitrary noncrossing geometric graph G and we might want to add the minimum number of edges to increase its vertex or edge-connectivity. In a third example, we may consider whether from any given plane spanning tree G we can construct an augmentation G' containing a Hamiltonian cycle or, when V is even, a perfect matching (Fig. 1).

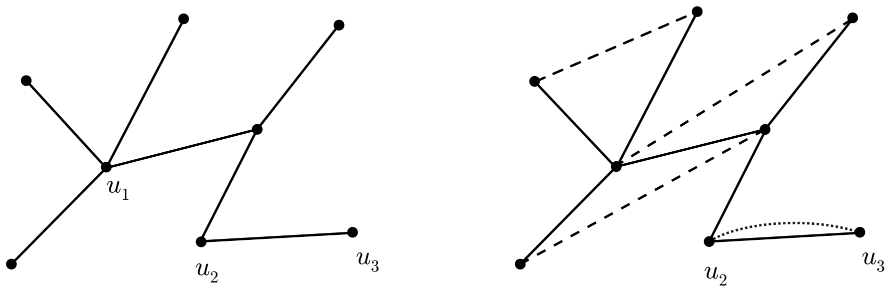

Figure 1: The tree on the left does not contain any perfect matching. Augmenting the tree with the three dashed edges on the right results in a graph that contains a perfect matching. For example, the three new edges together with the original edge u_2u_3 form a perfect matching.

The possible new edges that may be added to a geometric graph $G = (V, E)$ can be interpreted in the context of *visibility* problems, which emerged in the late eighties and early nineties. We say that two vertices $p, q \in V$ *see* each other (*i.e.*, they are mutually *visible*) if the segment pq does not cross any edge in E and its relative interior does not contain any vertex in V. The edges in E together with all visibility edges form a geometric graph, called the *segment endpoint visibility graph* for the segment configuration E. Notice that a segment between two vertices in V belongs to this graph if and only if it is compatible with E. We refer the readers to the surveys [11,50,87,115] for properties and related results stated inside the visibility framework.

For a set of points S in general position (that is, no three points on a line), we denote by $K(S)$ the complete geometric graph on vertex set S, that is, where every pair of points is joined by an edge.

A unifying framework. The concepts of augmentation and compatibility allow to describe, under a unifying framework, several problems that have a common flavor and have attracted substantial research, yet make use of different terminology or aim at different goals. A unified view helps to identify and understand common methods and common difficulties.

Our survey is not intended to be exhaustive, since there are many different variants of the augmentation problems under this framework. Rather, we have selected a sample of representative problems that provide the reader a global comprehension of the field. Along the way, we also present several open problems and unsettled conjectures.

In the augmentation problems we consider, we are typically given a planar straight-line graph G and a property $\mathcal{P}$; and our goal is to find an augmentation G' with property $\mathcal{P}$. When the augmentation is not necessarily feasible, one may seek efficient decision algorithms and combinatorial characterizations. If the augmentation is feasible, then the goal is to find the minimum number of new edges required. This includes possible combinatorial characterization of the minimum number, efficient algorithms to construct a minimum augmentation, approximation algorithms, as well as bounds on the extremal number of edges required over certain classes of graphs.

2 An introductory example: augmenting a matching

A set of disjoint line segments in the plane is, in effect, a crossing-free straight-line drawing of a perfect matching, where the segment endpoints are the vertices. This is a basic scenario for augmentation problems, in which we can neatly see geometry and graph theory interplay. In this section we review some fundamental results about this particular case.

2.1 From perfect matchings to Hamiltonian cycles and paths

A *polygonization* of a given point set S is a simple polygon whose vertex set is exactly S, in other words, a crossing-free spanning cycle in $K(S)$. It is easy to see that every set of $n \geq 3$ points in general position admits a polygonization (*e.g.*, every minimum Euclidean TSP is crossing-free). Consider now a set M of n disjoint line segments, let us denote by S_M the set of their endpoints, and assume that S_M is in general position. If the matching M can be augmented with $|M|$ edges to a noncrossing cycle P (*i.e.*, a simple polygon), then every other edge of P belongs to M. Such a polygon P is also called an *alternating polygonization* of M. It is easy to see that M does not always admit an alternating polygonization (see Fig. 2). Rappaport [92] proved that it is NP-complete to decide whether a plane geometric graph can be augmented to a spanning cycle; and it is also conjectured to be NP-complete to decide whether a noncrossing matching admits an alternating polygonization. Studying this problem had been suggested by Toussaint around 1985. Rappaport, Imai and Toussaint [93] proved that the decision problem is polynomially solvable in some special cases, for example when the segments in M are *convexly independent*, that is, each of them has at least one endpoint in the convex hull $\mathrm{CH}(S_M)$.

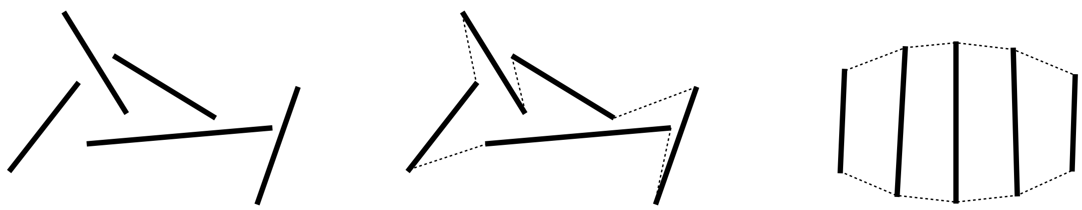

Figure 2: A set of five segments (left), which admits an alternating polygonization (middle). This is not possible for the five segments on the right of the figure because their endpoints are in convex position, and the only simple polygonization they admit is the boundary of their convex hull, missing the three central segments.

Let us remark that the noncrossing geometric constraint is the core of the difficulty of this problem. If we disregard crossings, we are simply in the combinatorial scenario, and any perfect matching M in the complete graph K_{2n} can be augmented in $2^n(n-1)!$ ways to a Hamiltonian cycle in which every other edge belongs to M. For embedded geometric graphs, we have seen an example of segments not admitting any alternating polygonization (Fig. 2). When they exist, it is possible that there is only one, for example, if M consists of every other edge of a convex $2n$-gon. But there can be exponentially many alternating polygonizations, as in the configuration in Fig. 3. Nevertheless, their number is bounded above by c^n for some constant c (see Section 3), in contrast to more than $n!$ different alternating cycles among abstract graphs.

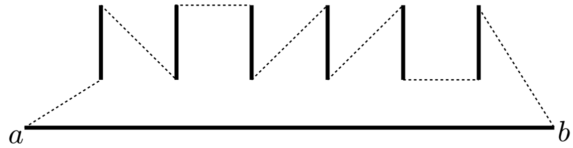

Figure 3: For this set of $n = 7$ segments we can easily construct 2^{n-1} different alternating polygonizations. Imagine that we travel from a to b, traversing the $n-1$ vertical segments in our trajectory. At every departure we can choose whether we arrive to the top or to the bottom end point of the next vertical segment.

We have seen that a noncrossing matching does not necessarily admit a Hamiltonian cycle. Mirzaian [81] conjectured that every noncrossing matching with noncollinear vertices can be augmented to a Hamiltonian plane geometric graph (that is, the augmented graph contains a Hamiltonian cycle, but this cycle does not have to contain M). Several years later, Hoffmann and Tóth [59] confirmed this conjecture in the affirmative. They construct a noncrossing cycle P incrementally, starting from the convex hull $\mathrm{CH}(S_M)$. The polygon P is then successively extended to pass through more segment endpoints, while it remains compatible with M. In a first phase of their algorithm, P is incrementally extended to include the second endpoint of every segment that already has an endpoint along P (similarly to Mirzaian's technique to handle convexly independent segments [81]); and simultaneously they construct a tiling of $\mathrm{CH}(S_M)$ such that every nonempty tile is adjacent to a unique edge of P. At the end of the first phase, all tiles are convex, and every remaining edge in $M \setminus P$ lies in the interior of some tile. In the second phase, a Hamiltonian polygon is computed, by induction, in each tile that contains noncollinear segments of M. The small polygons in the tiles, as well as collinear segments in some other tiles, are fused to P by modifying the common edge of the tile with P (see Fig. 4 for an example).

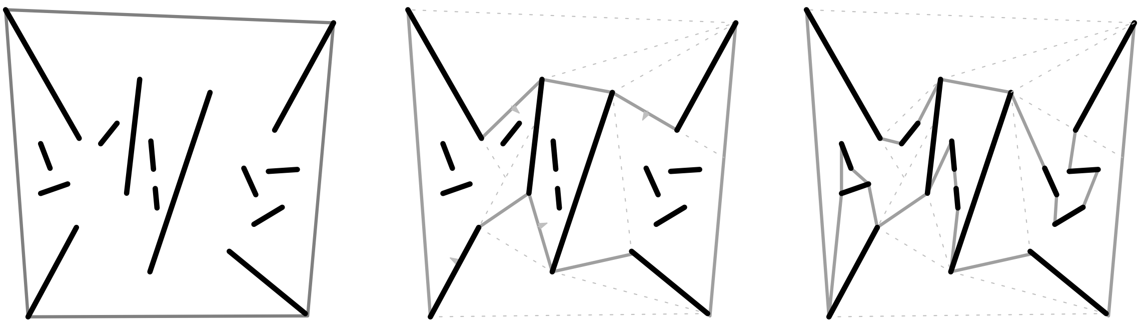

Figure 4: A set M of disjoint line segment, where P is initially the boundary of the convex hull $\mathrm{CH}(S_M)$ (left). Polygon P is extended until it contains either both endpoints or neither endpoint of each segment, every segment in $M \setminus P$ lies in a convex tile adjacent to a unique edge of P, which are indicated by small arrow heads (middle). In each convex tile containing noncollinear segments, there is a compatible spanning cycle by induction, which can be fused to P via the corresponding edge of P (right).

Mirzaian [81] also made a slightly stronger conjecture: that every noncrossing matching M

with noncollinear vertices would admit a *circumscribing polygon*, that is, a polygonization of S_M in which every segment of M would be either an edge or an internal diagonal of the polygon (external diagonals are excluded). He proved this property for convexly independent segments. However, Urabe and Watanabe found a counterexample to this conjecture [114]. O'Rourke and Rippel found a new family of segments admitting a circumscribing polygonization, namely sets of unit segments such that the line containing each segment misses all the other segments [88].

Pach and Rivera proved that every set M of n segments has a subset $M' \subseteq M$ of size roughly $\sqrt[3]{n}$, such that M' admits a circumscribing polygon [89]. Let us note, though, that this polygon may not be compatible with M, it may cross some edges in $M \setminus M'$. If we, instead, insist on alternating polygons that are compatible with all segments in M, then the best result one can prove is that there are always two segments in M that can be augmented to a simple quadrilateral compatible with M [62]. The bound of at most two is tight, as shown by the example in Figure 2, right.

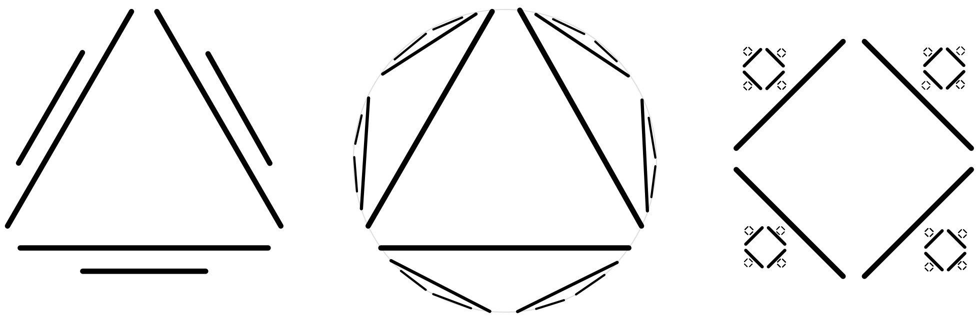

Figure 5: The set of segments that does not admit a spanning alternating path (left). This idea can be extended to a tree-like construction (middle) in which the endpoints of the segments are in convex position, in such a way that every segment "hides" two subtrees having the same structure. Only a subset of logarithmic size can be included in any alternating simple path. A similar construction (right) with orthogonal segments arranged in a 4-ary tree.

Around 1992, Urrutia asked what is the maximum length of a compatible alternating path for a perfect matching M. It is a significant relaxation to require an *open* polygonal path rather than a closed polygon. For example, the segments in the right of Figure 2, in fact, can be augmented to an alternating Hamiltonian path. However such a path is not always possible (Figure 5, left), and it is not difficult to construct sets of n segments in which any noncrossing alternating path compatible with M has $O(\log n)$ length (Figure 5, middle). Hoffmann and Tóth [58] proved that for every set of n disjoint segments, there is an alternating path of length $\Omega(\log n)$ compatible with M, hence showing that the upper bound is tight up to a multiplicative constant. For n disjoint orthogonal line segments, Tóth [111] proved an asymptotically tight bound of $\Theta(\sqrt{n})$ (Figure 5, right).

2.2 From perfect matchings to encompassing trees

We have seen in the previous section that not every perfect matching M can be augmented to a Hamiltonian cycle. It is not very difficult to prove, though, that there is always an *encompassing tree*, that is, a compatible spanning tree, that contains M as a subgraph: one can construct an encompassing tree from any triangulation of M (Fig. 6). Nevertheless, a naïve algorithm might yield some high degree node in the tree. Is it always possible to find an encompassing tree of bounded degree?

Bose and Toussaint [20] proved in 1992 that every n disjoint segments admit an encompassing tree with maximum degree at most seven, and gave an $O(n \log n)$ algorithm for its construction. Shortly afterwards, they improved on this result, showing that every set of disjoint segments can be encompassed by a spanning tree of maximum degree three (a binary tree), which can be constructed in optimal $\Theta(n \log n)$ time [18]. The bound on the degree is the best possible, because an encompassing tree with maximum degree two would be an alternating Hamiltonian path, which we know is not always possible. Later Souvaine and Tóth [106] obtained a generalization: every (disconnected) PSLG on n vertices can be augmented into a connected PSLG such that the degree of each vertex increases by at most two, and the augmentation can be computed in $O(n \log n)$ time.

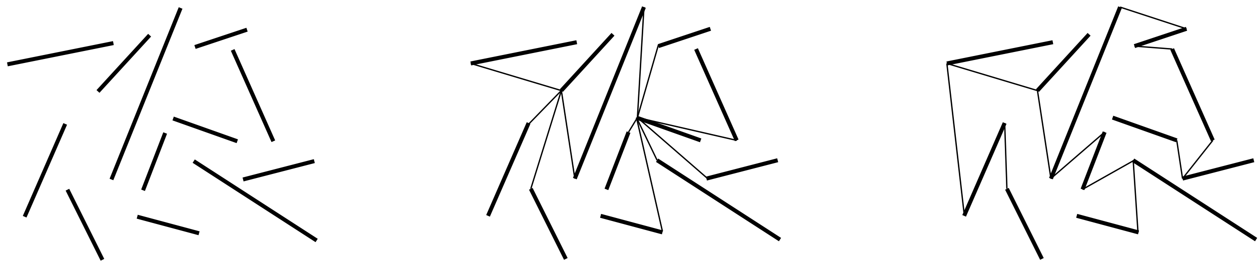

Figure 6: The set of twelve segments (left), which can be encompassed by a tree in which the maximum degree of a node is eight (middle) but also by a binary tree (right).

A colored version of this problem was posed in 2004 in the conference version of [66]: given a set of disjoint segments with a proper vertex coloring (two endpoints of each segments are colored differently), is it always possible to find a plane encompassing tree such that the new edges also respect the coloring? Hurtado *et al.* [66] gave an affirmative answer. Hoffmann and Tóth [57] later proved that there is always a *binary* encompassing tree respecting the initial coloring of the segment endpoints. This result was also generalized for arbitrary segment configurations [60]: every vertex-colored PSLG without singleton components can be augmented into a connected PSLG such that the degree of each vertex increases by at most two and the new edges respect the coloring. The exclusion of singleton components is necessary, since it is possible that a singleton is visible only from vertices of the same color (Figure 7, left).

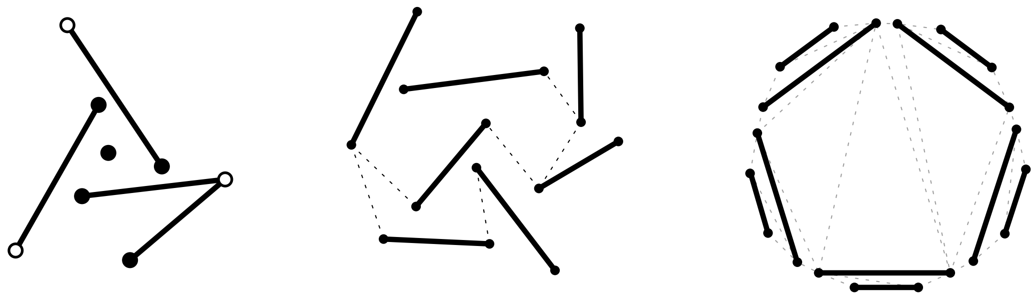

Figure 7: A vertex-colored graph where a singleton vertex is only visible from vertices of the same color (left). Disjoint line segments can be augmented to a pointed binary spanning tree (middle). A set of 10 segments, where for every pseudo-triangulation there is a vertex of degree at least 7 (right).

Another variation arose in the context of *pseudo-triangulations*, a topic that spurred substantial attention in the first decade of the 21st century (see the survey [95]) and that was introduced by

Streinu [107]. A *pseudo-triangle* is a simple polygon having exactly three convex vertices, and a *pseudo-triangulation* of a set S of n points is a noncrossing geometric graph, whose bounded faces decompose the convex hull $\mathrm{CH}(S)$ into pseudo-triangle faces. While the number of triangles and edges in a triangulation of S depends on the number of vertices on $\mathrm{CH}(S)$, it is always possible to decompose $\mathrm{CH}(S)$ into exactly $n-2$ pseudo-triangles, using exactly $2n-3$ edges, if S is in general position, no matter how many points of S lie in the interior of the convex hull. These decompositions minimize the number of pseudo-triangles (notice that a triangulation is also a pseudo-triangulation) and have the property that every vertex $p \in S$ is *pointed*: one of the faces incident to p has a reflex angle at p. The pointedness property is related to the rigidity of graphs [108], which is one of the reasons for the interest in pseudo-triangulations. Motivated by this framework, Hoffmann *et al.* [56] proved that every set of disjoint segments has an encompassing tree in which every vertex is pointed and has a degree at most three (Figure 7, middle). Such a pointed binary tree can be augmented to a minimal pseudo-triangulation, and building on that they also proved that every set of disjoint line segments in the plane has an encompassing minimal pseudo-triangulation whose maximum vertex degree is bounded by seven, and this bound cannot be improved (Figure 7, right).

2.3 Looking for a second matching

A variant close in spirit to the original problem on alternating polygonizations of a matching was posed by Aichholzer *et al.* [2]. Every alternating polygonization has an even number of edges and is the union of two disjoint perfect matchings, M and M'. On the other hand, if M and M' are two disjoint and compatible perfect matchings of a point set, then their union is a set of simple polygons (cycles in $K(S_M)$), each of which has an even number of edges. Let us observe that the matching in the right part of Figure 2 has an odd number of edges and does not admit any compatible disjoint perfect matching. Several other constructions were given in [2] for an odd number of disjoint segments without any compatible disjoint perfect matchings. The authors conjectured that for every noncrossing perfect matching M with an even number of edges in general position there is a disjoint compatible perfect matching (*Disjoint Compatible Matching Conjecture*).

The positive results in [2] towards a solution were only partial. They confirmed the conjecture in some special cases (namely, for convexly independent or orthogonal segments), and they also proved that there is always a set of alternating polygons, all of them compatible with M, encompassing at least $4/5$ of the segments in M. On the other hand, the techniques used for these results led the authors to pose stronger conjectures, involving the *convex subdivision* of the free space around the line segments. In fact, this geometric tool has been used in solving nearly all the problems described in this section. One easy way to construct a convex subdivision for the segments in M is the following (refer to Figure 8). For each endpoint q of a segment $pq = s \in M$, extend s along the ray $\overrightarrow{pq}$ beyond q until it hits another segment, a previous extension, or (if it is not blocked) to infinity. Different orders in which the extensions are drawn may yield different subdivisions, but in all cases the plane is subdivided into $n+1$ convex cells, where $n = |M|$. Clearly, not all convex subdivisions can be obtained in this way, *e.g.* the minimum number of cells in a convex subdivision may be far fewer than $n+1$. In the *dual multigraph* D associated with a convex subdivision, the vertex set $V(D)$ is the set of convex cells, and every segment endpoint p corresponds to an edge in $E(D)$ between two cells incident to p (double edges are possible if a segment in M lies on the common boundary between two cells).

The additional conjectures in [2] were stated in the terms of the dual multigraph D associated with a "suitable" convex subdivision. Orienting an edge of D towards a node $v \in V(D)$ can be

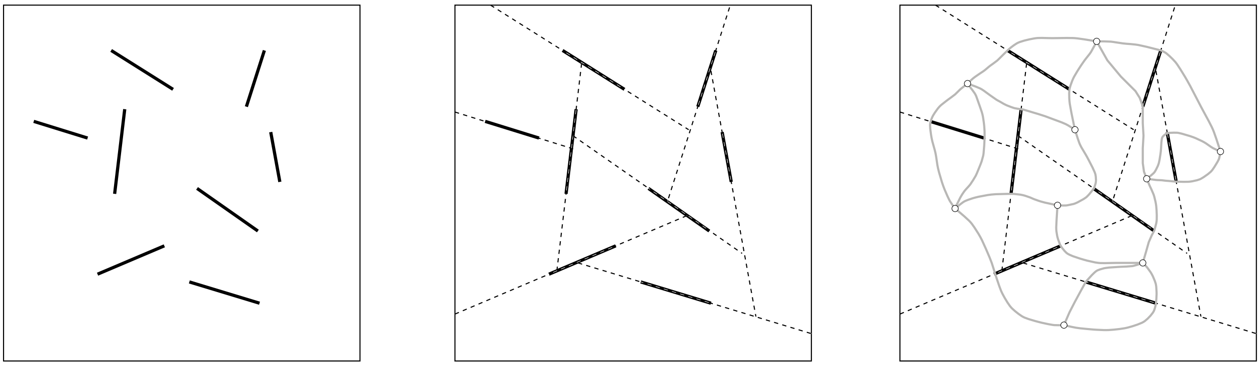

Figure 8: The segments on the left part can be extended to produce a convex subdivision of the plane, as shown in the middle. The dual multigraph of this subdivision is shown on the right.

seen as assigning a segment endpoint to the cell corresponding to v. In an *even orientation* of an undirected (multi-)graph, edges are oriented in such a way that every node has an even in-degree. In particular, an even orientation of D means that an even number of segment endpoints are assigned to each convex cell of the subdivision. Now, it is not difficult to see that in this case, one can match the endpoints assigned to each cell using disjoint segments lying in the interior of the cell (and hence compatible with M), with the only exceptional case that a cell is assigned to exactly two segments endpoints, which are the two endpoints of the same segment of M. We can encode this information in the dual graph: two adjacent edges of D are said to be in *conflict* if they correspond to the two endpoints of a segment in M; and an even orientation of D is *conflict-free* if no two conflicting edges are oriented into a node of in-degree 2. Hence, if M admits a convex subdivision where the dual graph has a conflict-free even orientation, this immediately implies the existence of a disjoint compatible matching M'.

A solution to the Disjoint Compatible Matching Conjecture has been claimed very recently by Ishaque *et al.* [68], and follows the approach described in the preceding paragraph. For every set of $n \geq 2$ disjoint line segments in the plane in general position, they construct a convex subdivision such that the associated dual graph D contains two edge-disjoint spanning trees. The subdivision is obtained by an iterative process, in which the number of cells may drop below $n+1$, but the number of edges remains $2n$, since they are in bijection with the segment endpoints. They also show that every multigraph that has an even number of edges and contains two edge-disjoint spanning trees must have a conflict-free even orientation. Based on the argument in the previous paragraph, this already implies the Disjoint Compatible Matching Conjecture.

None of the stronger conjectures formulated by Aichholzer *et al.* [2] have been confirmed yet, but some of them have been refuted. For example, if we insist that the convex subdivision must be constructed by extending the segments successively beyond their endpoints (as in Fig. 8), then the dual graph does not always contain two edge-disjoint spanning trees [7]. In the proof of Ishaque *et al.* [68], it was essential to work with a broader class of convex subdivisions, which may have fewer than $n+1$ convex cells.

3 An extreme case: augmenting empty graphs

A peculiar yet very interesting augmentation problem arises when the input graph has no edges: we are only given a set S of points in the plane and some graph property $\mathcal{P}$. The most restrictive graph property requires the augmentation to be isomorphic with a given planar graph $G = (V, E)$, with $|S| = |V| = n$ vertices. In other words, the problem is to *embed G with straight-line edges on top of S* without crossings.

Notice that this problem is not always feasible, and that the specific configuration of the points in S makes a world of difference. For example, no graph with at least $2n-2$ edges can be embedded on top of a set of n points in convex position, because this set admits at most $2n-3$ pairwise non-crossing geometric edges, which would constitute a triangulation of its convex hull. On the other hand, even if we fix a set of n points with a triangular convex hull (and so any triangulation has $3n-6$ edges), this point set may not accommodate all n-vertex planar graphs. An example is given in Figure 9: many maximal planar graphs with n vertices contain no vertex of degree 3 or less, or no vertex of degree $n-1$, and none of them can be realized on top of the set shown in Figure 9.

Figure 9: A set S of $n = 9$ points with a triangular convex hull (left). A maximal plane graph on top of S will be a triangulation T with $3n-6$ edges. However, the edges shown on the right must appear in every triangulation of S, and they only differ in the edges that triangulate the gray region Q. Therefore T will have at least two vertices of degree 3, and at least one with degree $n-1$.

3.1 Plane drawings of specific graphs and classes of graphs

Given an n-vertex planar graph G and an n-element point set S, it is NP-complete to decide whether G admits a straight-line embedding on top of S. Cabello [23] proved that this is NP-complete problem even if G is restricted to be 2-connected and 2-outerplanar.

However, there are finite point sets, perhaps much larger than n, that accommodate all n-vertex planar graphs. A point set S is called *n-universal* if every planar graph with n nodes admits a straight-line embedding on top of a subset of S. For instance, the union of vertex sets of arbitrary straight-line embeddings of *all* n-vertex planar graphs is n-universal, although this point set is quite large, with up to $e^{\Theta(n)}$ points. It is a longstanding open problem to find the smallest size of an n-universal point set for all $n \geq 0$ (see Problem 45 in [29]). It is known that an $(n-1) \times (n-1)$ section of the integer lattice is n-universal [44, 101]. However, no n-universal set of points with a subquadratic number of elements is known. Chrobak and Karloff [27] proved that every n-universal set must have at least $1.098n$ points, and Kurowski [79] improved the lower bound to $1.235n$.

It is also remarkable that for certain graph classes the configuration of the points—provided they are in general position—is no obstacle at all. For example, Gritzmann *et al.* [55] proved that every outerplanar graph of n vertices admits a crossing-free plane embedding on top of any set of n points in general position. There are also efficient algorithms for the construction of the embedding [15,76]. In the case of a tree $T = (V, E)$, it is even possible to select a node $v \in V$ and a point p in the given point set S, and realize T on top of S with the additional constraint that v is mapped to p. This was proved by Ikebe *et al.* [67] and an efficient associated algorithm is described in [19].

There are many variants of similar embeddability problems. Brass *et al.* [21], for example, consider the *geometric simultaneous embedding* of two planar graphs in which two graphs should admit a plane realization on top of the same vertex set, possibly with the additional requirement that they are compatible; see [43] for a list of results on this subject. We do not pursue here this extension or the many variants that appear in the rich literature of the *Graph Drawing* area of problems. We refer the interested reader to Chapter 9 in [22] for geometric graphs, and to [31,109], Chapter 21 in [99], and Chapter 52 in [52], for the more generic framework.

3.2 The number of embeddings

Some graphs admit many straight-line embeddings on top of a point set S in general position. For instance, the cycle C_n has several embeddings on an n-element point set, unless the points are in convex position. It was the subject of intense research to find upper and lower bounds on the maximum possible number of embeddings. Here we briefly survey two cases and provide references for other graph classes.

Let S be a set of n points in the plane in general position, labeled by integers from 1 through n. There is an obvious combinatorial upper bound on the number of polygonizations given by the number of cyclic permutations, $(n-1)!$, which is roughly n^n, neglecting exponential and polynomial terms. However, it is clear that most of the polygons generated in this way would have crossings. Notice that the solution to the Euclidean traveling salesman problem, that is, the spanning cycle of minimal total length, is necessarily crossing-free, *i.e.*, a polygonization, because any two crossing edges can be replaced by two edges of their convex hull, resulting in a strictly shorter spanning cycle. This motivates why this counting problem, going back to Newborn and Moser [85], has been intensely investigated.

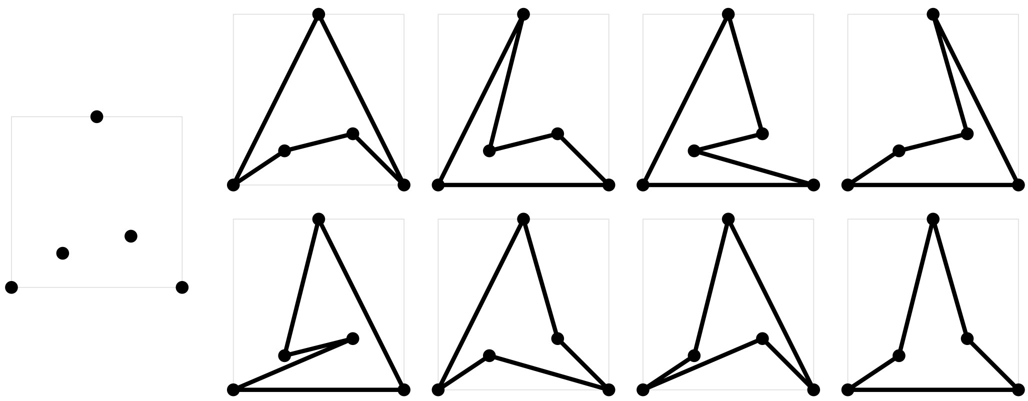

Figure 10: A set S of five points in the plane, and all eight straight-line embeddings of C_5 on top of S.

A major step in estimating the number straight-line embeddings was achieved by Ajtai, Chvátal, Newborn, and Szemerédi [6]. They proved that every set of n points in the plane admits at most c^n crossing-free geometric graphs, with a constant $c = 10^{13}$. It is worth mentioning that the *crossing lemma*, a corner stone in geometric graph theory, was proved in this paper as a lemma for this result[1]. The upper bound for the number of polygonizations on an n-element point set has been improved substantially in a series of papers. One of the latest upper bounds is $O(68.66^n)$ by Dumitrescu *et al.* [36], which has recently been further improved to $O(54.55^n)$ by Sharir *et al.* [103]. The best current lower bound comes from the so-called *double chain* configuration (see Figure 11), which is known to have more than 4.64^n different polygonizations [47]. If we denote by $p(n)$ the maximum number of polygonizations over all sets of n points in general position, the currently best bounds can be summarized as $4.64^n \lesssim p(n) \lesssim 54.55^n$, where we omit polynomial factors and display the dominating exponential term only.

The bound c^n given by Ajtai *et al.* [6] also applies to triangulations. Counting triangulations is an old problem going back to Euler, who considered the case of points in convex position, counted by the *Catalan numbers* (see [102] for an historical account). Possibly David Avis was the first to ask about the maximal number of triangulations over generic point sets [118]. Denote by $\operatorname{tr}(S)$ the number of triangulations a point set S admits, and by $\operatorname{tr}(n) = \max_{|S|=n} \operatorname{tr}(S)$ the maximal value over all n-element point sets. Assume that we have a bound $\operatorname{tr}(n) \leq c_t^n$ for some constant c_t. Then, since every crossing-free geometric graph on n vertices can be augmented to a triangulation and a triangulation has at most $3n - 6$ edges, we infer that $2^{3n-6}c_t^n \approx 8^n c_t^n$ is an upper bound for the number of plane geometric graphs on top of any n-element point set.

Similarly, if for some class of n-vertex graphs $\mathcal{G}$ one can derive an upper bound of g^n on the number of graphs in $\mathcal{G}$ that are contained as a subgraph in a triangulation on n vertices, then $g^n c_t^n$ is an upper bound for the different realizations of graphs in $\mathcal{G}$ on any n-element point set. This explains why improving on the value of the constant c_t has been the subject of a long list of papers, see [118,105] for an account and references. Currently the best upper bound is $\operatorname{tr}(n) \leq 30^n$, obtained by Sharir and Sheffer [102], refining the method given by Sharir and Welzl [104], which combines the use of random triangulations with a charging scheme. On the opposite direction, the *double chain* point configuration admits 8^n triangulations [47]. This lower bound construction was widely believed to be best possible until Aichholzer *et al.* [5] introduced the *double zig-zag chain*, and proved that it admits $\Omega(8.48^n)$ triangulations (Figure 11). The double chain consists essentially of two "flat" convex polygons P and Q facing each other. In the double zig-zag chain, P and Q are the most basic type of *almost-convex polygons*, a class introduced by Hurtado and Noy [63]. Using other almost-convex polygons in a similar manner, Dumitrescu *et al.* [36] have very recently constructed n-element point sets that admit $\Omega(8.65^n)$ triangulations. While the gap between the upper and lower bounds keeps narrowing and there is no firm conjecture on what the right constant may be in the exponent, nonetheless one may say that the lower bounds have always been thought to be closer to the true value.

Other graphs, such as perfect matchings, spanning paths, pseudo-triangulations and many more have been studied from this perspective. Details on bounds and references can be found in [5] and [105].

[1] The crossing lemma was independently proved by Leighton [80], as well.

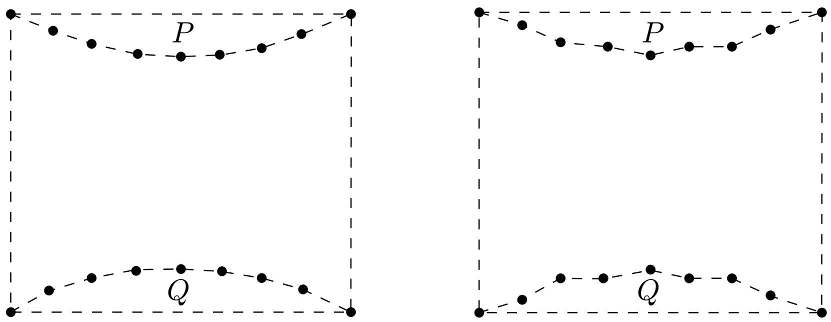

Figure 11: The double chain (left) consists of two convex chains that face each other. The line connecting any pair of points in the lower (resp. upper) chain leaves above (resp. below) all the points in the other chain. The dashed edges are present in every triangulation, so the three resulting regions are triangulated independently. The double zig-zag chain (right) is very similar, but reflex and convex vertices alternate, and the line defined by two consecutive vertices in the lower chain leaves exactly one point from the lower chain above; the situation is symmetric for the upper chain.

3.3 Spanning graphs with desirable properties

In many applications, we do not necessarily want to embed a specific graph on top of a point set S, but rather we would like to construct a spanning graph on S with certain properties. The properties may be purely graph theoretical (e.g., connectivity) or specific to the geometric realization (e.g., dependent on Euclidean lengths or angles). In this section, we consider some representative examples of each type.

Connectivity. Given a point set S in the plane in general position, and an integer $k \geq 0$, we could like to embed a k-connected (resp., k-edge-connected) graph on top of S. For $k = 1$, it is clear that every point set admits a spanning tree, which is the smallest connected graph on n vertices. As noted above, $n \geq 3$ points in general position also admit a spanning cycle, which is the smallest 2-connected graph on n vertices. However, n points in convex position do not admit 3-connected augmentation, since every maximal augmentation is a triangulation in which at least two vertices have degree 2. It is easy to see, though, that any other point set (*i.e.*, not in convex position) on $n \geq 4$ vertices admits a 3-connected graph: start with a star centered at a vertex in the interior of the convex hull, and complete it into a wheel graph, which is 3-connected. For 4- and 5-connectivity, the analogous question about feasibility is not so easy anymore.

To determine whether a point set admits a k-connected graph, it is enough to consider whether it admits a k-connected *triangulation*, since every planar straight-line graph can be triangulated and additional edges can only increase the connectivity. Dey *et al.* [30] proved that every set whose convex hull consists of 3 vertices admits a 4-connected triangulation, with the only exception of the point set in Figure 9, left. No characterization is known for generic point sets (with arbitrary convex hull) that admit 4- or 5-connected triangulations. Note that k-connectivity is infeasible for $k \geq 6$, since very planar graph has a vertex of degree at most 5.

If we can decide whether there exists a k-connected graph on top of a given point set S, the next question is to find one with the fewest possible edges. If we allow crossings, then every set of n vertices admits a 3-connected graph with $\lceil 3n/2 \rceil$ edges, which is best possible since the degree of every vertex must be at least 3. For constructing a *crossing-free* 3-connected graph on top of a point set S, the location of the points (or, rather, their order type) already matters. García *et al.* [45] proved that if we are given a set of n points, $h < n$ of which lie on the boundary of the convex hull, then it admits a 3-connected planar straight-line graph with $\max(\lceil 3n/2 \rceil, n + h - 1)$ edges, and this bound is the best possible for each point set. This implies, in particular, that a cubic (that is, absolute minimum size) 3-connected graph is possible if $n \geq 4$ is even and at most $n/2 + 1$ points lie on the convex hull (Figure 12, left). The proof in [45] is algorithmic, and 3-connected graphs of the above size can be computed in polynomial time.

Angle and slope conditions. A frequently used generalization of triangulations is a *convex decomposition*, which is a plane geometric graph where all bounded faces are convex, and the bounded faces jointly tile the convex hull of the vertex set. Every point set admits a convex decomposition (since every triangulation is a convex decomposition). While a triangulation on top of n points, h of which lie on the convex hull, has exactly $3n - h - 3$ edges, a convex decomposition may have as few as n edges (if the points are in convex position). Improving upon earlier results [84, 61], Sakai and Urrutia [100] showed recently that every set of n points, h of which are on the convex hull, admits a convex decomposition with $\frac{7}{3}n - h$ edges, which is the best possible for infinitely many (n, h) pairs (Figure 12, middle), but not known to be optimal in general. The currently best lower bound, $\frac{23}{11}n - 3$, in terms of n is due to García-Lopez and Nicolás [49].

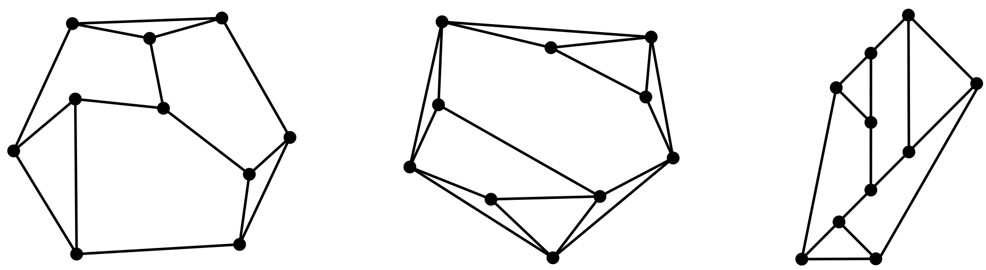

Figure 12: A set of $n = 10$ points, $h = 6$ of which lie on the boundary of the convex hull, admits a 3-connected cubic PSLG (left). A minimum convex decomposition for $n = 10$ points, $h = 5$, with $\lfloor \frac{7}{3}n - h \rfloor = 18$ edges (middle). A straight-line embedding of a 3-connected cubic graph where the 15 edges have only 6 different slopes (right).

It is well known [12] that among all triangulations of a point set, the Delaunay triangulation maximizes the minimum angle (intuitively, all triangles are as "fat" as possible). Aichholzer *et al.* [4] proved that every point set admits a triangulation in which each point is incident to a triangle that has an angle of at least $2\pi/3$ at that point, and the bound $2\pi/3$ is the best possible. There are similar results [4, 14, 35, 40] for angle-constrained spanning cycles, however all known results consider arbitrary geometric graphs (*with* possible crossings). It is an open problem to determine the minimum and maximum possible angles that a *crossing-free* spanning cycle, spanning path, or spanning tree can have, over all n-element point sets in general position.

In a problem closely related to the angle constraints on adjacent edges, Dujmović *et al.* [34] studied the crossing-free geometric graphs that can be embedded in the plane with few different edge *slopes*. They prove that every planar graph with n vertices has a straight-line embedding in the plane with at most $2n - 10$ different slopes, based on a canonical decomposition introduced by de Fraysseix [44]. They also construct triangulations that require at least $n + 2$ different slopes. It remains an open problem to determine the minimum number of slopes sufficient for the straight-line embedding of every planar graph with n vertices. It is clear that a vertex of degree Δ forces at least $\lceil \Delta/2 \rceil$ slopes in any straight-line drawing. Dujmović *et al.* [34] show that $\lceil \Delta/2 \rceil$ slopes are sufficient for the embedding of every tree of maximum degree Δ; and 6 slopes are sufficient for the embeding of every 3-connected cubic planar graph [73] (Figure 12, right). But in general, it is not known whether the "slope number" of a planar graph is a function of the maximum degree.

Spanners. The *stretch factor* of a geometric graph $G = (V, E)$ is the maximum ratio

$$\max_{u,v \in V} \frac{|\text{path}(u, v)|}{|uv|},$$

where $|\text{path}(u, v)|$ is the Euclidean length of the shortest path between u and v, and $|uv|$ is the Euclidean distance between points u and v. The ratio $|\text{path}(u, v)|/|uv|$ is also called the *detour* between points p and q. A geometric graph $G = (V, E)$ is a k-spanner if its stretch factor is at most k. For a set of n points, the only 1-spanner is the complete geometric graph, which has $\binom{n}{2}$ edges and many crossings whenever $n \geq 5$. Chew [26] proved that for every point set, there is a plane geometric graph with stretch factor at most 2, and it is the dual graph of the Voronoï diagram induced by equilateral triangles. Chew conjectured that the stretch factor of the standard Delaunay triangulation (induced by disks) is at most $\pi/2 \approx 1.5708$. Chew's conjecture has recently been refuted by Bose *et al.* [17]. The best current lower bound was found by a computer search: Xia and Zhang [121] presented a point set whose Delaunay triangulation has a stretch factor of 1.5907. The best current upper bound, 1.998, has recently been announced by Xia [120], improving the previous best upper bound of $4\sqrt{3}\pi/9 \approx 2.418$ by Keil and Gutwin [77]. It is not known what is the maximum stretch factor of a Delaunay triangulation, and, more interestingly, what is the minimum stretch factor that a plane geometric graph can have on top of any finite point set.

Giannopoulos *et al.* [51] and Gudmundsson and Smid [53] proved independently that, for a given set of n points and an integer K, it is NP-hard to find the plane geometric graph of minimum stretch factor and at most K edges. It is still possible, however, that one can efficiently compute a plane geometric graph with minimum stretch factor for a given point set if there is no limitation on the number of edges.

In the last 20 years, researchers have studied extensively crossing-free spanners that have several desirable properties simultaneously, such as bounded stretch factor, bounded degree, and small Euclidean length [16]. One recent result in this thread, by Kanj *et al.* [71], states that for every integer $k \geq 14$ and every n-element point set, there is a PSLG with maximum degree at most k and stretch factor at most $(1 + \frac{2\pi}{k\cos(\pi/k)}) \cdot 4\sqrt{3}\pi/9$, and such a PSLG can be computed in $O(n)$ time. For a detailed historical account, variants of the problem, and algorithmic results, we refer to a survey book by Narasimhan and Smid [83].

4 Generic plane augmentation problems

In the previous two sections, we have surveyed augmentation problems for perfect matchings and empty graphs. In the most general version of geometric graph augmentation, we are given an arbitrary planar straight-line graph $G = (V, E)$ and a graph property $\mathcal{P}$. Our task is to augment G into a PSLG with property $\mathcal{P}$ or report that no such augmentation is possible. Here $\mathcal{P}$ may be an abstract graph property or a property specific to geometric graphs. At any rate, we can consider only properties of planar graphs, since the output has to be planar. We consider two illustrative examples: One is connectivity augmentation, where property $\mathcal{P}$ is k-connectivity or k-edge-connectivity for some integer k, $1 \leq k \leq 5$. The other property is *detour at most* d, for some real $d \geq 1$, which is specific to geometric graphs.

4.1 Connectivity augmentation

The k-connectivity (resp., k-edge-connectivity) augmentation problem asks for the minimum number of edges that augment a given graph $G = (V, E)$ into a k-connected (resp., k-edge-connected) graph $G' = (V, E \cup E')$. For abstract graphs, both the vertex- [70] and edge-connectivity [42,82,117] versions have polynomial-time solutions for every integer k; and there are linear time solutions for $k = 2$ [38, 64, 65, 94]. In general, vertex connectivity is technically more difficult to handle. When k is part of the input, Végh [116] recently gave a polynomial time algorithm for the k-connectivity augmentation of a $(k-1)$-connected graph.

Kant and Bodlaender [74] considered connectivity augmentation problems over planar graphs, where both the input and output are required to be planar (*i.e.*, no minor isomorphic to K_5 or $K_{3,3}$). They showed that such a *planarity-preserving* edge-connectivity augmentation is NP-hard already for $k = 2$, and later Rutter and Wolff [98] showed this for vertex-connectivity as well. Fialko and Mutzel [41] and Gutwenger *et al.* [54] proposed constant-factor approximations. However, in a planarity preserving augmentation, the output graph may not be compatible with every straight-line embedding of the input graph, and so these results are not applicable when the input is a planar straight-line graph.

Connectivity augmentation over planar straight-line graphs is significantly more restrictive. For instance, a path on n vertices has a planarity preserving augmentation to a 2-connected Hamiltonian cycle by adding one edge (between the endpoints of the path). However if the path is embedded in the plane as a zig-zag chain on n points in convex position (Fig. 13), then 2-edge-connectivity augmentation requires $\lfloor n/2 \rfloor$ new edges, 2-connectivity requires $n - 2$ new edges, and the connectivity augmentation is not even feasible for $k \geq 3$. Rutter and Wolff [98] proved that the k-connectivity and k-edge-connectivity augmentation problems are NP-hard over plane geometric graphs for $k = 2, 3, 4$, and 5. As noted above, the problem is infeasible for $k \geq 6$ because every planar graph has a vertex of degree at most 5. Currently, there are no approximation algorithms. Research focused on determining whether augmentation is feasible, and on extremal bounds: what is the minimum number of new edges that are sufficient for the k-connectivity (resp., k-edge-connectivity) augmentation of *any* plane geometric graph on n vertices?

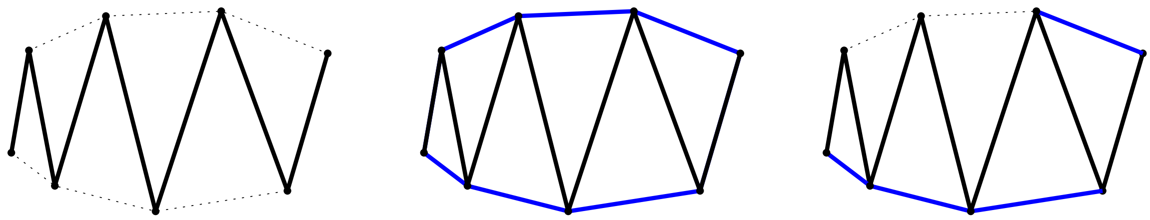

Figure 13: A zig-zag path on 8 vertices in convex position (left). 2-connectivity augmentation required 6 new edges (middle). 2-edge-connectivity augmentation required 4 new edges (right).

Vertex-connectivity. Abellanas *et al.* [1] were the first to obtain extremal bounds. They proved that a connected PSLG with b cut-vertices can be augmented to a 2-connected PSLG with b new edges. The case of the zig-zag path, with $n-2$ cut vertices, shows that this bound is the best possible. It follows that every connected PSLG on n vertices can be augmented to a 2-connected PSLG with at most $n-2$ new edges, and this bound is also the best possible. It is not known, however, what is the minimum number of edges sufficient for the 2-connectivity augmentation of any (possibly disconnected) PSLG with n vertices.

It is easy to see that every PSLG in general position can be augmented to 2-connectivity, since every triangulation is 2-connected. This is no longer true for 3-connectivity. Tóth and Valtr [113] proved that a PSLG $G=(V,E)$ can be augmented to 3-connectivity if and only if V is not in convex position and no edge in E is a proper diagonal of the convex hull $\mathrm{CH}(V)$. Similar combinatorial characterizations are not known for 4- or 5-connectivity.

Edge-connectivity. Every PSLG can be augmented to a 2-edge-connected PSLG. Tóth and Valtr [113] characterized the PSLGs that can be augmented to 3-edge-connectivity, they are called *3-edge-augmentable.* Specifically, a PSLG $G=(V,E)$ is 3-edge-augmentable if and only if V is not in convex position and there is no edge $e \in E$ such e is a proper diagonal of $\mathrm{CH}(V)$ and all vertices on one side of e lie on the boundary of $\mathrm{CH}(V)$. Similar combinatorial characterizations are not known for 4- or 5-connectivity.

Let us point out two easy but important tools that simplify the study of edge-connectivity augmentation. The first tool is the concept of *k-edge-connected* components in a graph $G=(V,E)$: it is a maximal subset $V_c \subset V$ such there are at least k edge-disjoint paths between any two vertices of V_c. By Menger's theorem, a graph G is k-edge-connected if and only if it has only one k-edge-connected component. The second tool is the use of multi-edges. Abellanas *et al.* [1] proved that if a PSLG $G=(V,E)$ can be augmented to 2-edge-connectivity with m new edges such that some new edges are duplications of existing edges in E, then the augmentation is also possible with m new edges and no double edges. An analogous result holds for 3-edge-connectivity augmentation if G is 3-edge-augmentable [8].

Target edge-connectivity	**1**	**2**	**3**
Arbitrary PSLG	$n-1$	$\geq \lfloor (4n-4)/3 \rfloor$	$2n-2$
Connected PSLG	0	$\lfloor (2n-2)/3 \rfloor$	$\geq \lfloor (4n-4)/3 \rfloor$
2-edge-connected PSLG	0	0	$n-2$

Table 1: The minimum number of new edges sufficient for raising the edge-connectivity of any PSLG on n vertices in general position to a target $k=1,2,3$. Tight bounds and lower bounds.

Extremal bounds are known for the minimum number of edges sufficient for increasing the edge-connectivity by one for any PSLG with n vertices (see Table 1). It is easy to see that every disconnected PSLG with $c \geq 2$ components can be augmented to a connected PSLG by adding exactly $c-1$ new edges. The maximum value of c is n, attained for empty graph. Abellanas *et al.* [1] conjectured that every connected PSLG on n vertices can be augmented to a 2-edge-connected PSLG with at most $\lfloor (2n-2)/3 \rfloor$ new edges, and the this was later confirmed by Tóth [112]. This upper bound is the best possible: the lower bound construction consists of a triangulation with m vertices, with a leaf added in each of the $2m-5$ bounded faces, and 3 pairwise invisible leaves in the outer face (Fig. 14, left). This PSLG has $n=m+(2m-5)+3=3m-2$ vertices, and each of the $2m-2$ leaves requires one new edge for 2-edge-connectivity. The proof of the upper bound is constructive: a polynomial-time augmentation algorithm adds new edges in each face of the input graph independently. The key tool is the *geodesic hull* of all 2-edge-connected components adjacent to a face (Fig. 14, middle). The boundary of the geodesic hull is a closed polygonal chain P that visits all 2-edge-connected components, and so one can add some edges of P into the input graph (possibly creating double edges), and merge them into a single 2-edge-connected component (Fig. 14, middle). At the end of the algorithm, double edges can be replaced with single edges by [1].

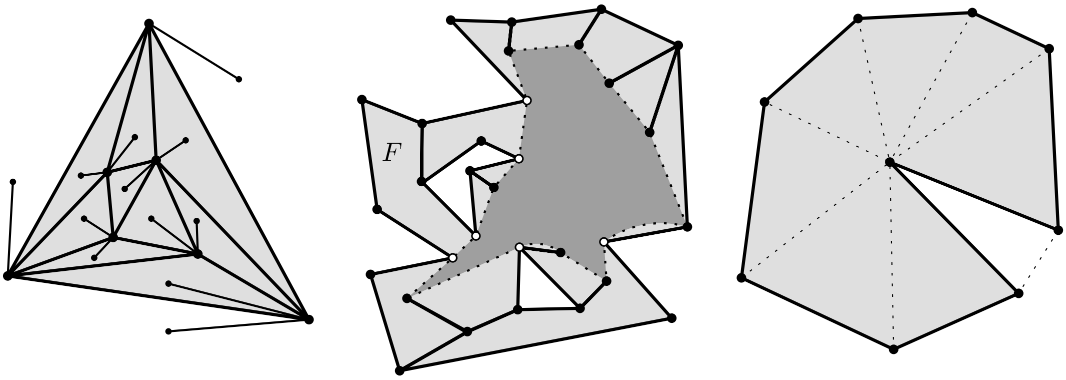

Figure 14: A PSLG on $n = 19$ vertices, including $\lfloor (2n-2)/3 \rfloor = 12$ leaves (left). The geodesic hull of one vertex of each 2-edge-connected component incident to a face F (middle). A 2-edge-connected graph on $n = 9$ vertices that requires $n - 2 = 7$ new edges to 3-edge-connectivity (right).

Tóth and Valtr [113] showed that $n - 2$ edges are always sufficient for raising the edge-connectivity of a 3-edge-augmentable PSLG on n vertices from 2 to 3. The upper bound is based on the simple observation that a new edge can decrease the number of 3-edge-connected components by one, and if there are n such components initially, then the first two new edges can create a K_4 minor, which is a 3-connected component with 4 vertices. This upper bound is the best possible: if $G = (V, E)$ is a Hamiltonian cycle with exactly one vertex in the interior of $\mathrm{CH}(V)$, then the only 3-edge-connected augmentation is the wheel graph with $n - 2$ new edges (Fig. 14, right).

For abstract graphs, one can solve k-connectivity-augmentation optimally in a multi-phase algorithm, where each phase increments the edge-connectivity by one [25]. This is not the case for geometric graphs, since edges added in one phase could become obstacles in subsequent phases. Al-Jubeh *et al.* [8] showed that any PSLG on n vertices in general position can be augmented to 3-edge-connectivity with at most $2n-2$ new edges, if the augmentation is feasible, and this bound is the best possible. No tight extremal bounds are known for raising the edge-connectivity from 0 to 2 or from 1 to 3 (see Table 1). Lower bound constructions, similar to one in the left part of Fig. 14, show that in the worst case $\lfloor (4n-4)/3 \rfloor$ new edges are needed for both problems.

Augmenting crossing-free straight-line trees. Abellanas *et al.* [1] proved that every path on n vertices in general position can be augmented to 2-edge-connectivity with $\lfloor n/2 \rfloor$ edges, and this bound is the best possible (Fig. 13). They conjectured that $\frac{n}{2} + O(1)$ edges are sufficient for augmenting any tree on n vertices to 2-edge-connectivity. However, Tóth [112] constructed a family of trees that require $\frac{17}{33}n - O(1)$ new edges. García and Tejel [48] improved the lower bound to $\frac{6}{11}n - O(1)$. Consider the tree in the left part of Fig. 15. It has 8 leaves, 5 of which lie in the interior of the convex hull and are pairwise invisible to each other. Identify one of the interior leaves A with an exterior leaf of a scaled copy of the original construction as in the right part of

Fig. 15(right). Then the number of vertices goes up by 11. The leaf A, which required one new edge, has been replaced by a subgraph with 7 leaves. Two of these leaves see each other, and the remaining 5 leaves are isolated. However, even if we join the two new mutually visible leaves to each other, the original edge A is still not contained in any circuit, so these 7 leaves require 7 new edges. By iterating this step k times, we obtain a tree with $13 + 11k$ vertices that require $7 + 6k$ new edges.

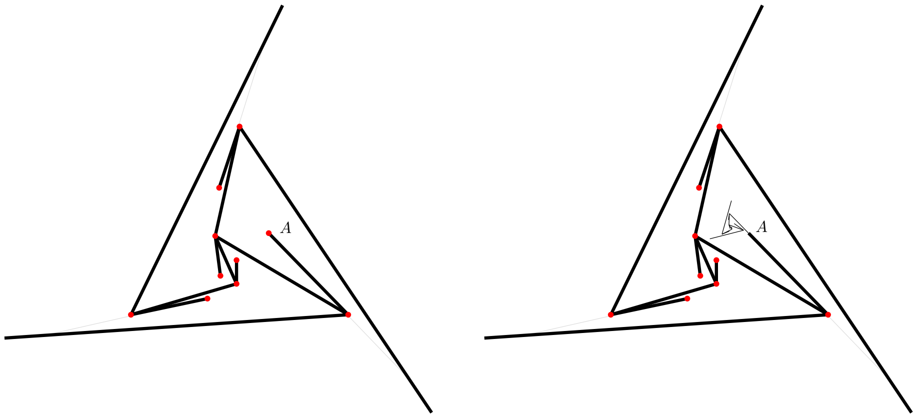

Figure 15: A tree with 13 vertices, including 8 leaves, which requires 7 new edges (left). Replacing A with a scaled copy of the original construction, the number of vertices increases by 11, and the number of leaves increases by 7, and it requires 6 more new edges (right).

Currently, the best upper bound for trees is the same as for any PSLG: every connected PSLG with n vertices can be augmented to a 2-edge-connected PSLG by adding at most $\lfloor (2n-2)/3 \rfloor$ new edges. We emphasize again that the difficulty is posed by requiring *straight-line* edges. If the new edges are allowed to be Jordan arcs, then the lower bound of $\lfloor n/2 \rfloor$ is the best possible [112].

4.2 Network optimization through augmentation

Motivated by wireless communication networks, Kranakis *et al.* [78] studied 2-edge-connectivity augmentation for noncrossing subgraphs of unit disk graphs. They prove that if $G = (V, E)$ is a connected PSLG in general position, it has b cut edges, and all edges have length at most 1, then G can be augmented to a 2-edge-connected PSLG with at most b new edges, each of which has length at most 3. The bound 3 on the length of the new edges cannot be improved (Fig. 16, left), but the bound on the number of new edges is not known to be optimal. Similarly, a noncrossing tree with n vertices in general position and edges of length at most 1 can be augmented to a 2-edge-connected PSLG with at most $\lfloor 5n/6 \rfloor$ new edges, each of length at most 3. It follows from the reduction of Rutter and Wolff [98] that finding the minimum number of new edges remains NP-complete under the restriction that the new edges have bounded length.

In typical applications in wireless communication, the edges of a graph or digraph are given implicitly. The input is a set S of points in the plane. Every assignment of angular domains with apices at S and radius r (which correspond to directional antennae of uniform range r stationed at points in S) define an geometric digraph $G = (V, E)$ such that $(u, v) \in E$ if and only if one of the angular domains with apex at u contains v and the distance $|uv|$ is at most r. A digraph defined in this way can be augmented by increasing the radius r. Note, however, that some edge pairs may inevitably cross in this model. Problems in the literature [24, 28, 33] ask for the minimum radius and minimum angle (alternatively, sum of angles) that guarantee strong connectivity.

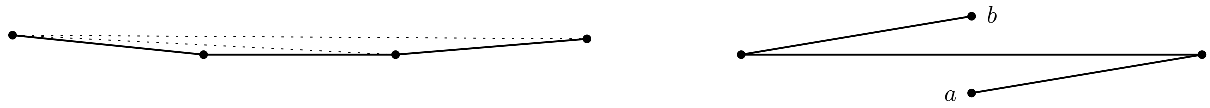

Figure 16: A path with unit length edges, that requires an edge of length $3 - \varepsilon$ for 2-edge-connectivity augmentation (left). A path between a and b where the stretch factor is large, but cannot be improved with additional edges (right).

Giannopoulos *et al.* [51] showed that it is NP-hard to decide whether a given PSLG $G = (V, E)$ can be augmented with at most K new edges such that its stretch factor drops below a given threshold $\lambda > 1$. Farshi *et al.* [39] considered the problem of finding a single new edge that would maximally decrease the stretch factor of a geometric graph, where both the input graph and the new edge may have crossings; they proposed efficient algorithms for both finding and approximating the optimal "shortcut." Wulff-Nilsen [119] generalized the problem to arbitrary metric spaces. Aronov *et al.* [10] studied the variant of this problem where the input graph contains a singleton (which must be connected to the remainder of the graph), but they also allow crossings in the output. No results are known for finding such "bottleneck" edges over crossing-free geometric graphs in the plane. Note that if crossings are not allowed, then the stretch factor of the augmentation cannot be decreased below any constant threshhold. The detour between two endpoints of a zig-zag path can be arbitrarily large (Fig. 16, right), but in some cases it cannot be improved by adding new edges.

In all applications mentioned so far, it was clear that the addition of new edges can only improve the network (new edges can neither decrease the connectivity nor increase the detour). We note here that augmentation may deteriorate the quality of a network. For example, there are simple instances of noncooperative network congestion games (*c.f.*, Braess's paradox) where the total latency at Nash equilibrium increases if we add a new edge. Roughgarten [96, 97] showed that this deterioration can be arbitrarily large, even in planar networks. Network design problems, under various capacity and demand constraints, go beyond the scope of this paper. We refer to a comprehensive survey by Tardos and Wexler [110]. Current techniques in this domain are unable to enforce the crossing-free condition, and it would be very interesting to study congestion games for embedded networks in a geometric setting.

5 Concluding remarks

The family of augmentation problems is quite large. As mentioned in the introduction, we have only discussed some selected representatives problems here. We conclude with briefly mentioning some problems that have not been included.

- We omitted the *weighted versions* of problems, in which the weight of an edge is its Euclidean length, and the weight of a PSLG is its total edge length. The complexity of only very few weighted optimization problems are known. Given a set of n points in the plane, the Euclidean minimum spanning tree (EMST) can be computed in $O(n \log n)$ time, as a minimum spanning tree in the Delaunay triangulation of the points. However, finding a minimum weight triangulation is NP-hard [86]. The minimum weight PSLGs from most graph classes (*e.g.*, spanning trees, spanning cycles) are often automatically noncrossing, but the crossing-free property has to be explicitly enforced for *maximum* weight PSLGs, which is impossible to model with standard optimization techniques. Approximation algorithms for some variants have been presented in [9, 37].

- We have not reviewed *combinatorial games*, in which two players construct disjoint compatible PSLGs on a given point set by incrementally augmenting the empty graph, following certain rules. Many graph creation games have been considered (*e.g.* maker-breaker or avoider-enforcer games) for noncrossing geometric graphs [3], but many open problems remain.
- We did not mention any problem that involves Steiner points (that is, augmenting a graph with both vertices and edges, and possibly subdividing existing edges). Patrignani [90] showed that it is NP-complete to decide whether a partial embedding of a graph can be completed to a straight-line embedding. This problem is, in fact, NP-complete already for trees [13].

References

[1] M. Abellanas, A. García, F. Hurtado, J. Tejel, and J. Urrutia, Augmenting the connectivity of geometric graphs. *Comput. Geom. Theory Appl.* **40** (3) (2008), 220-230.

[2] O. Aichholzer, S. Bereg, A. Dumitrescu, A. García, C. Huemer, F. Hurtado, M. Kano, A. Márquez, D. Rappaport, S. Smorodinsky, D. Souvaine, J. Urrutia, and D. Wood, Compatible geometric matchings. *Comput. Geom. Theory Appl.* **42** (2009), 617–626.

[3] O. Aichholzer, D. Bremner, E. D. Demaine, F. Hurtado, E. Kranakis, H. Krasser, S. Ramaswami, S. Sethia, and J. Urrutia, Games on triangulations. *Theor. Comp. Sci.* **343** (1-2) (2005), 42–71.

[4] O. Aichholzer, T. Hackl, M. Hoffmann, C. Huemer, A. Pór, F. Santos, B. Speckmann, and B. Vogtenhuber, Maximizing maximal angles for plane straight line graphs. In *Proc. 10th Workshop on Algorithms and Data Structures*, vol. 4619 of LNCS, 2007, pp. 458–469.

[5] O. Aichholzer, T. Hackl, C. Huemer, F. Hurtado, H. Krasser, and B. Vogtenhuber, On the number of plane geometric graphs. *Graphs and Combinatorics* **23** (1) (2007), 67–84.

[6] M. Ajtai, V. Chvátal, M. M. Newborn, and E. Szemerédi, Crossing-free subgraphs. *Annals Discrete Math.* **12** (1982), 9-12.

[7] M. Al-Jubeh, M. Hoffmann, M. Ishaque, D. L. Souvaine, and C. D. Tóth, Convex partitions with 2-edge connected dual graphs. *J. Combin. Opt.* (2010), in print.

[8] M. Al-Jubeh, M. Ishaque, K. Rédei, D. L. Souvaine, and C. D. Tóth, Tri-edge-connectivity augmentation for planar straight line graphs. *Algorithmica* (2011), to appear.

[9] N. Alon, S. Rajagopalan, and S. Suri, Long non-crossing configurations in the plane. *Fundam. Inform.* **22** (4) (1995), 385–394.

[10] B. Aronov, K. Buchin, M. Buchin, M. van Kreveld, M. Löffler, J. Luo, R. I. Silveira, and B. Speckmann, Connect the dot: computing feed-links with minimum dilation. In *Proc. Algorithms Data Structures Sympos. (WADS)*, vol. 5664 of LNCS, 2009, Springer, pp. 49-60.

[11] T. Asano, S. K. Ghosh and T. C. Shermer, Visibility in the plane, in: J.R. Sack, J. Urrutia (Eds.), *Handbook on Computational Geometry*, Elsevier, Amsterdam, 2000, pp. 829–876.

[12] F. Aurenhammer and Y. Xu, Optimal triangulations. In C. A. Floudas and P. M. Pardalos (eds.): *Encyclopedia of Optimization*, Vol. 4, Kluwer, 2001, pp. 160–166.

[13] A. Bagheri and M. Razzazi, Planar straight-line point-set embedding of trees with partial embeddings *Inf. Proc. Lett.* **110** (2010), 521–523.

[14] I. Bárány, A. Pór, and P. Valtr, Paths with no small angles. *SIAM J. Discrete Math.* **23** (4) (2009), 1655–1666.

[15] P. Bose, On embedding an outer-planar graph in a point set. *Comput. Geom. Theory Appl.* **23** (2002), 303-312.

[16] P. Bose, J. Gudmundsson, and M. Smid, Constructing plane spanners of bounded degree and low weight. *Algorithmica* **42** (2005), 249–264.

[17] P. Bose, L. Devroye, M. Löffler, J. Snoeyink, and V. Verma, The spanning ratio of the Delaunay triangulation is greater than $\pi/2$. *Comput. Geom. Theory Appl.* **44** (2) (2011), 121–127.

[18] P. Bose, M. E. Houle, and G. T. Toussaint, Every set of disjoint line segments admits a binary tree. *Discrete Comput. Geom.* **26** (3) (2001), 387–410.

[19] P. Bose, M. McAllister, and J. Snoeyink, Optimal algorithms to embed trees in a point set. *Journal of Graph Algorithms and Applications* **1** (2) (1997), 1-15.

[20] P. Bose and G. T. Toussaint, Growing a tree from its branches. *J. Algorithms* **19** (1995), 86–103.

[21] P. Brass, E. Cenek, C. Duncan, A. Efrat, C. Erten, D. Ismailescu, S. Kobourov, A. Lubiw, and J. Mitchell. On simultaneous planar graph embeddings. *Comput. Geom. Theory Appl.* **36** (2) (2007), 117–130.

[22] P. Brass, W. Moser and J. Pach, *Research Problems in Discrete Geometry.* Springer, 2005.

[23] S. Cabello, Planar embeddability of the vertices of a graph using a fixed point set is NP-hard. *J. Graph Algorithms and Applications* **10** (2006), 353–363.

[24] I. Caragiannis, C. Kaklamanis, E. Kranakis, D. Krizanc, and A. Wiese, Communication in wireless networks with directional antennas. In *Proc. 20th ACM Sympos. Parallel Algorithms and Architectures*, 2008, ACM Press, pp. 344–351.

[25] E. Cheng and T. Jordán, Succesive edge-connectivity augmentation problems. *Math. Program.* **84** (1999), 577–593.

[26] L. P. Chew, There are planar graphs almost as good as the complete graph. *J. Computer and System Sciences* **39** (2) (1989), 205–219.

[27] M. Chrobak and H. Karloff, A lower bound on the size of universal sets for planar graphs. *SIGACT News* **20** (1989), 83-86.

[28] M. Damian and R. Flatland, Spanning properties of graphs induced by directional antennas. In *Proc. 20th Fall Workshop on Computational Geometry*, Stony Brook, NY, 2010,

[29] E. D. Demaine, J. S. B. Mitchell, and J. O'Rourke (eds.), *The Open Problems Project.* `http://cs.smith.edu/~orourke/TOPP/welcome.html`

[30] T. K. Dey, M. B. Dillencourt, S. K. Ghosh, J. M. Cahill, Triangulating with high connectivity. *Comput. Geom. Theory Appl.* **8** (1997), 39-56.

[31] G. Di Battista, P. Eades, R. Tamassia, G. Tollis, and Ioannis, *Graph Drawing: Algorithms for the Visualization of Graphs.* Prentice Hall, 1998.

[32] D. P. Dobkin, S. J. Friedman, and K. J. Supowit, Delaunay graphs are almost as good as complete graphs. *Discrete Comput. Geom.* **5** (1990), 399–407.

[33] S. Dobrev, E. Kranakis, D. Krizanc, O. Morales, J. Opatrny, and L. Stacho, Strong connectivity in sensor networks with given number of directional antennae of bounded angle. In *Proc. 4th Conf. on Combin. Opt. Appl.*, vol. 6509 of LNCS, Springer, 2010, pp. 72-86.

[34] V. Dujmović, D. Eppstein, M. Suderman, and D. R. Wood, Drawings of planar graphs with few slopes and segments. *Comput. Geom. Theory Appl.* **38** (2007), 194-212.

[35] A. Dumitrescu, J. Pach, and G. Tóth, Drawing Hamiltonian cycles with no large angles. In *Proc. 17th Sympos. Graph Drawing (GD'09)*, vol. 5849 of LNCS, 2010, Springer, pp. 3–14.

[36] A. Dumitrescu, A. Schulz, A. Sheffer, and C. D. Tóth, Bounds on the maximum multiplicity of some common geometric graphs. In *Proc. Sympos. Theorerical Aspects of Comp. Sci. (STACS)*, 2011, LIPICS, Schloss Dagstuhl, vol. 5 of LIPICS, Schloss Dagstuhl, pp. 637-648.

[37] A. Dumitrescu and C. D. Tóth, Long non-crossing configurations in the plane. *Discrete Comput. Geom* **44** (4) (2010), 727–752.

[38] K. P. Eswaran, R. E. Tarjan, Augmentation problems. *SIAM J. Comput.* **5** (4) (1976), 653-665.

[39] M. Farshi, P. Giannopoulos, and J. Gudmundsson, Improving the stretch factor of a geometric network by edge augmentation. *SIAM J. Comput.* **38** (1) (2008), 226–240.

[40] S. P. Fekete and G. J. Woeginger, Angle-restricted tours in the plane. *Comput. Geom. Theory Appl.* **8** (4) (1997), 195–218.

[41] S. Fialko and T. P. Mutzel, A new approximation algorithm for the planar augmentation problem. In *Proc. 9th ACM-SIAM Symposium on Discrete Algorithms*, 1998, pp. 260–269.

[42] A. Frank, Augmenting graphs to meet edge-connectivity requirements. *SIAM J. Discrete Math.* **5** (1) (1992), 22-53.

[43] F. Frati, M. Kaufmann, and S. Kobourov, Constrained simultaneous and near-simultaneous embeddings. In *Proc. 15th Sympos. on Graph Drawing (GD'07)*, vol. 4875 of LNCS, 2008, Springer, pp. 268–279.

[44] H. de Fraysseix, J. Pach, and R. Pollack, How to draw a planar graph on a grid. *Combinatorica* **10** (1) (1990), 41-51.

[45] A. García, C. Huemer, F. Hurtado, J. Tejel, and P. Valtr, On triconnected and cubic plane graphs on given point sets. *Comput. Geom. Theory Appl.* **42** (9) (2009), 913–922.

[46] A. García, F. Hurtado, M. Noy, and J. Tejel, Augmenting the connectivity of outerplanar graphs, *Algorithmica* **56** (2) (2010), 160–179.

[47] A. García, M. Noy, and J. Tejel, Lower bounds on the number of crossing-free subgraphs of K_N. *Comput. Geom. Theory Appl.* **16** (2000), 211–221.

[48] A. García and J. Tejel, private communication.

[49] J. García-Lopez and M. Nicolás, Planar point sets with large minimum convex partitions. In *Abstracts of the 22nd European Workshop on Comput. Geom.*, Delphi, 2006, pp. 51-54.

[50] S. K. Ghosh, *Visibility algorithms in the plane.* Cambridge University Press, Cambridge, United Kingdom, 2007.

[51] P. Giannopoulos, R. Klein, Ch. Knauer, M. Kutz, and D. Marx, Computing geometric minimum-dilation graphs is NP-hard. *Int. J. Comput. Geom. Appl.* **20** (2) (2010), 147–173.

[52] J. E. Goodman, J. O'Rourke (eds.), *Handbook of Discrete and Computational Geometry.* Chapman & Hall and CRC Press, 2nd edition, 2004.

[53] J. Gudmundsson and M. Smid, On spanners of geometric graphs. *Intern. J. Foundations of Comp. Sci.* **20** (1) (2009), 135–149.

[54] C. Gutwenger, P. Mutzel, and B. Zey, On the hardness and approximability of planar biconnectivity augmentation. In *Proc. 15th Computing and Combin. Conf. (COCOON)*, vol. 5609 of LNCS, Springer, 2009, pp. 249–257.

[55] P. Gritzmann, B. Mohar, J. Pach, and R. Pollack, Embedding a planar triangulation with vertices at specified points. *Amer. Math. Monthly* **98** (2) (1991), 165-166.

[56] M. Hoffmann, B. Speckmann, and C. D. Tóth, Pointed binary encompassing trees: simple and optimal. *Comput. Geom. Theory Appl.* **43** (1) (2010), 35–41.

[57] M. Hoffmann and C. D. Tóth, Alternating paths through disjoint line segments. *Inf. Proc. Lett.* **87** (6) (2003), 287–294.

[58] M. Hoffmann and C. D. Tóth, Pointed and colored binary encompassing trees. In *Proc. 21st Sympos. Comput. Geom.*, 2005, ACM Press, pp. 81–90.

[59] M. Hoffmann and C. D. Tóth, Segment endpoint visibility graphs are Hamiltonian. *Comput. Geom. Theory Appl.* **26** (1) (2003), 47-68.

[60] M. Hoffmann and C. D. Tóth, Vertex-colored encompassing graphs, manuscript, 2010, `http://math.ucalgary.ca/~cdtoth/colored-encom.pdf`.

[61] K. Hosono, On convex decompositions of a planar point set, *Discrete Math.* **309** (2009), 1714-1717.

[62] M. E. Houle, private communication.

[63] F. Hurtado and M. Noy, Counting triangulations of almost-convex polygons. *Ars Combinatoria* **45** (1997), 169-179.

[64] T.-S. Hsu, Simpler and faster biconnectivity augmentation, *J. Algorithms* **45** (1) (2002), 55-71.

[65] T.-S. Hsu and V. Ramachandran, On finding a minimum augmentation to biconnect a graph, *SIAM J. Computing* **22** (5) (1993), 889-912.

[66] F. Hurtado, M. Kano, D. Rappaport, and C. D. Tóth, Encompassing colored crossing-free geometric graphs. *Comput. Geom. Theory Appl.* **39** (1) (2008), 14–23.

[67] Y. Ikebe, M. A. Perles, A. Tamura, and S. Tokunaga, The rooted tree embedding problem into points in the plane. *Discrete Comput. Geom.* **11** (1994), 51-63.

[68] M. Ishaque, D. L. Souvaine, and C. D. Tóth, Disjoint compatible geometric matchings. In *Proc. 27th Sympos. on Comput. Geom. (Paris, 2011)*, ACM Press, pp. 125-134.

[69] M. Ishaque, D. L. Souvaine, C. D. Tóth, and A. Winslow, Constrained tri-connected planar straight line graphs. In *Abstracts of the 20th Fall Workshop on Comput. Geom. (Stony Brook, NY, 2010).*

[70] B. Jackson and T. Jordán, Independence free graphs and vertex connectivity augmentation. *J. Combin. Theory, Ser. B* **94** (2005), 31–77.

[71] I. A. Kanj, L. Perkovic, and G. Xia, On spanners and lightweight spanners of geometric graphs. *SIAM J. Comput.* **39** (6) (2010), 2132–2161.

[72] G. Kant, Augmenting outerplanar graphs. *J. Algorithms* **21** (1996), 1–25.

[73] G. Kant, Hexagonal grid drawings. In *Proc. 18th Workshop in Graph-Theoretic Concepts in Computer Science (WG'92)*, vol. 657 of LNCS, 1993, Springer, pp. 263-276.

[74] G. Kant and H. L. Bodlaender, Planar graph augmentation problems. In *Proc. 2nd Workshop on Algorithms and Data Structures*, vol. 519 of LNCS, Springer, 1991, pp. 286-298.

[75] G. Kant and H. L. Bodlaender, Planar graph augmentation problems. Technical report RUU-91-25, Utrecht University, 1991.

[76] M. Kaufmann and R. Wiese, Embedding vertices at points: Few bends suffice for planar graphs. *Journal of Graph Algorithms and Applications* **6** (1) (2002), 115-129.

[77] J. M. Keil and C. A. Gutwin, Classes of graphs which approximate the complete Euclidean graph. *Discrete Comput. Geom.* **7** (1992), 13–28.

[78] E. Kranakis, D. Krizanc, O. Morales, and L. Stacho, Bounded length, 2-edge augmentation of geometric planar graphs. In *Proc. 4th Conf. on Combin. Opt. Appl.*, vol. 6508 of LNCS, Springer, 2010, pp. 385-397.

[79] M. Kurowski, A 1.235 lower bound on the number of points needed to draw all n-vertex planar graphs. *Inf. Proc. Lett.* **92** (2004), 95-98.

[80] T. Leighton, New lower bound techniques for VLSI. *Math. Systems Theory* **17** (1984), 47–70.

[81] A. Mirzaian, Hamiltonian triangulations and circumscribing polygons of disjoint line segments. *Comput. Geom. Theory Appl.* **2** (1)(1992), 15–30.

[82] H. Nagamochi and T. Ibaraki, Augmenting edge-connectivity over the entire range in $O(nm)$ time. *J. Algorithms* **30** (1999), 253–301.

[83] G. Narasimhan and M. Smid, *Geometric Spanner Networks.* Cambridge University Press, Cambridge, United Kingdom, 2007.

[84] V. Neumann-Lara, E. Rivera-Campo, and J. Urrutia, A note on convex decompositions of point sets in the plane, *Graphs and Combinatorics* **20** (2) (2004), 223-231.

[85] M. Newborn and W. O. J. Moser, Optimal crossing-free Hamiltonian circuit drawings of K_n. *J. Combin. Theory, Ser. B* **29** (1980), 13–26.

[86] W. Mulzer and G. Rote, Minimum-weight triangulation is NP-hard. *J. ACM* **55** (2) (2008), article 11.

[87] J. O'Rourke, Visibility. In J. E. Goodman and J. O'Rourke (eds.), *Handbook of Discrete and Computational Geometry.* CRC Press LLC, Boca Raton, FL, 2nd edition, 2004, pp. 643-665.

[88] J. O'Rourke and J. Rippel, Two segment classes with Hamiltonian visibility graphs. *Comput. Geom. Theory Appl.* **4** (1994), 209-218.

[89] J. Pach and E. Rivera, Two segment classes with Hamiltonian visibility graphs. *Comput. Geom. Theory Appl.* **10** (1998), 121-124.

[90] M. Patrignani, On extending a partial straight-line drawing. *International Journal of Foundations of Computer Science* **17** (5) (2006), 1061-1069.

[91] J. Plesník, Minimum block containing a given graph. *Arch. Math.* **27** (6) (1976), 668–672.

[92] D. Rappaport, Computing simple circuits from a set of line segments is NP-complete. *SIAM J. Comput.* **18** (6) (1989), 1128–1139.

[93] D. Rappaport, H. Imai, and G. T. Toussaint, Computing simple circuits from a set of line segments. *Discrete Comput. Geom.* **5** (3) (1990), 289-304.

[94] A. Rosenthal and A. Goldner, Smallest augmentations to biconnect a graph. *SIAM J. Computing* **6** (1977), 55-66.

[95] G. Rote, F. Santos, and I. Streinu, Pseudo-triangulations—a survey. In J. E. Goodman and J. Pach (eds.), *Surveys on Discrete and Computational Geometry: Twenty Years Later*, vol. 453 of Contemporary Mathematics, 2008, AMS, pp 343–410.

[96] T. Roughgarden, On the severity of Braess's Paradox: Designing networks for selfish users is hard. *J. Comput. Syst. Sci.* **72** (5) (2006), 922–953.

[97] T. Roughgarden, Routing games. Chapter 18 in: N. Nisan *et al.* (eds.) *Algorithmic Game Theory*, Cambridge University Press, 2007.

[98] I. Rutter and A. Wolff, Augmenting the connectivity of planar and geometric graphs. Technical report, Universität Karlsruhe, 2008. Number 2008–3.
`http://digbib.ubka.uni-karlsruhe.de/volltexte/1000007814`

[99] J. R. Sack, J. Urrutia (eds.), *Handbook of Computational Geometry.* North Holland, 2000.

[100] T. Sakai and J. Urrutia, Convex decompositions of point sets in the plane. In *Proc. 7th Japan Conference on Computational Geometry and Graphs*, JAIST, 2009.

[101] W. Schnyder, Embedding planar graphs on the grid. In *Proc. 1st ACM-SIAM Sympos. Discrete Algorithms*, 1990, pp. 138-148.

[102] M. Sharir and A. Sheffer, Counting triangulations of planar point sets. —em Electronic J. Combinat. **18** (2011), P70.

[103] M. Sharir, A. Sheffer, and E. Welzl, Counting plane graphs: perfect matchings, spanning cycles, and Kasteleyn's technique. manuscript, 2011.

[104] M. Sharir and E. Welzl, On the number of crossing-free matchings (cycles, and partitions). *SIAM J. Comput.* **36** (3) (2006), 695–720.

[105] A. Sheffer, Numbers of plane graphs,
`http://www.cs.tau.ac.il/~sheffera/counting/PlaneGraphs.html`

[106] D. L. Souvaine and C. D. Tóth, A vertex-face assignment for plane graphs. *Comput. Geom. Theory Appl.* **42** (5) (2009), 388-394.

[107] I. Streinu, A combinatorial approach to planar non-colliding robot arm motion planning. In *Proc. 41st Sympos. Foundations of Comp. Sci. (FoCS)*, IEEE, Los Alamitos, CA, 2000, pp. 443–453,

[108] I. Streinu, Pseudo-triangulations, rigidity and motion planning. *Discrete Comput. Geom.* **34** (2005), 587–635.

[109] R. Tamassia (ed.), *Handbook of Graph Drawing and Visualization*, CRC Press, in preparation.

[110] E. Tardos and T. Wexler, Inefficiency of equilibria in network formation games. Chapter 19 in: N. Nisan *et al.* (eds.), *Algorithmic Game Theory*, Cambridge University Press, 2007.

[111] C. D. Tóth, Alternating paths along axis-parallel segments. *Graphs and Combinatorics* **22** (4) (2006), 527–543.

[112] C. D. Tóth, Connectivity augmentation in planar straight line graphs. *Europ. J. of Combinatorics*, to appear. (Preliminary version in: *Proc. Intl. Conf. on Topological and Geometric Graph Theory*, Paris, 2008, pp. 51-54.)

[113] C. D. Tóth and P. Valtr, Augmenting the edge connectivity of planar straight line graphs to three. In *Proc. 13th Spanish Meeting on Comput. Geom.*, Zaragoza, 2009.

[114] M. Urabe and M. Watanabe, On a counterexample to a conjecture of Mirzaian. *Comput. Geom. Theory APpl.* **2** (1) (1992), 51-53.

[115] J. Urrutia, Art gallery and illumination problems. In J. R. Sack and J. Urrutia (eds.), *Handbook on Computational Geometry*, Elsevier, Amsterdam, 2000, pp. 973-1127.

[116] L. Végh, Augmenting undirected node-connectivity by one. In *Proc. 42nd Sympos. on Theory of Computing (STOC)*, ACM Press, 2010, pp. 563-572.

[117] T. Watanabe and A. Nakamura, Edge-connectivity augmentation problems. *J. Comput. System Sci.* **35** (1987), 96–144.

[118] E. Welzl, The number of triangulations on planar point sets. In *Proc. 14th International Symposium (GD'06)*, vol. 4372 of LNCS, Springer, 2007, pp. 1–4.

[119] C. Wulff-Nilsen, Computing the dilation of edge-augmented graphs in metric spaces. *Comput. Geom. Theory Appl.* **43** (2010), 68-72.

[120] G. Xia, Improved upper bound on the stretch factor of Delaunay triangulations, in *Proc. 27th Sympos. on Comput. Geom. (SoCG)*, ACM Press, 2011, pp. 264–273

[121] G. Xia and L. Zhang, Improved lower bound for the stretch factor of Delaunay triangulations. In *Proc. 20th Fall Workshop on Comput. Geom.*, Stony Brook, NY, 2010.