Counting Polygon Triangulations is Hard

David Eppstein
Computer Science Department, University of California, Irvine, USA
eppstein@uci.edu

Abstract

We prove that it is #P-complete to count the triangulations of a (non-simple) polygon.

1 Introduction

In 1979, Leslie Valiant published his proof that it is #P-complete to compute the permanent of a 0-1 matrix, or equivalently to count the perfect matchings of a bipartite graph [41]. This result was significant in two ways: It was surprising at the time that easy (polynomial time) existence problems could lead to intractable counting problems, and it opened the door to hardness proofs for many other counting problems. Beyond individual problems, several broad classifications of hard graph counting problems are now known. For instance, it is #P-complete to count k-colorings or determine the value of the Tutte polynomial at any value outside of a small finite set of exceptions [23], and all problems of counting homomorphisms to a fixed directed acyclic graph are either #P-complete or polynomial [18].

In computational geometry, it has been shown that counting the vertices or facets of high-dimensional convex polytopes is #P-complete [28], and that computing the expected total length of the minimum spanning tree of a stochastic subset of three-dimensional points is #P-hard [25]. Additionally, when testing existence of a geometric structure is hard [5, 9, 31, 29, 26] it is just hard to determine whether its count is nonzero. However, we know of no past hardness proofs for counting easy-to-construct two-dimensional structures. Although there has been significant research on counting non-crossing configurations in the plane, including matchings, simple polygons, spanning trees, triangulations, and pseudotriangulations, the complexity of these problems has remained undetermined. Research on these problems has instead focused on determining the number of configurations for special classes of point sets [20, 8, 24], bounding the number of configurations as a function of the number of points [2, 4, 39, 37, 17, 38, 1, 33, 3, 36, 21, 10], developing exponential- or subexponential-time counting algorithms [12, 7, 42, 5, 30, 13], or finding faster approximations [6, 27].

In this paper we bring the two worlds of #P-completeness and counting planar structures together by proving the following theorem:

▶ **Theorem 1.** *It is #P-complete to count the number of triangulations of a given polygon.*

Necessarily, the polygons of our hardness construction have holes, as it is straightforward to count the triangulations of a simple polygon in polynomial time by dynamic programming [19, 32, 16]. The polygons resulting from the construction can be drawn with all vertices on a grid of polynomial size. Our proof strategy is to develop a polynomial-time counting reduction from the problem of counting independent sets in 3-regular bipartite planar graphs (here the independent sets are not necessarily maximum nor maximal), which was proved
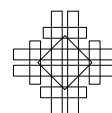
#P-complete by Xia, Zhang, and Zhao [43]. We reduce counting independent sets in 3-regular planar graphs to counting maximum non-crossing subsets of a special class of line segment arrangements, which we in turn reduce to counting triangulations.

2 Preliminaries

2.1 Polygons and triangulation

A *planar straight-line graph* consists of finitely many closed line segments in the Euclidean plane, disjoint except for shared endpoints. The endpoints of these segments can be interpreted as the vertices, and the segments as the edges, of an undirected graph drawn with straight edges and no crossings in the plane. The *faces* of the planar straight-line graph are the connected components of its complement (that is, maximal connected subsets of the plane that are disjoint from the segments of the graph). In any planar straight-line graph, exactly one unbounded face extends beyond the bounding box of the segments; all other faces are bounded. A segment of the graph forms a *side* of a face if the interior of the segment intersects the topological closure of the face. As with any graph, a planar straight-line graph is d-regular if each of its vertices is incident to exactly d line segments.

For the purposes of this paper, we define a *polygon P* to be a 2-regular planar straight-line graph in which there is a bounded face ϕ whose sides are all the segments of P. If ϕ exists, it is uniquely determined from P. We call ϕ the *interior* of P. The connected components of the graph are necessarily simple cycles of line segments, exactly one of which separates ϕ from the unbounded face. If there is more than one connected component of the graph, we call the other components *holes*, and if there are no holes we call P a *simple polygon*.

We define a *triangulation* of a polygon P to be a planar straight-line graph consisting of the edges of P and added segments interior to P, all of whose vertices are vertices of P, partitioning the interior of P into three-sided faces. As is well known, every polygon has a triangulation. A triangulation of an n-vertex polygon can be found in $O(n \log n)$ time, for instance by constrained Delaunay triangulation, and this can be improved to $O(n \log h)$ for polygons with h holes [11]. The known exponential or sub-exponential algorithms for counting triangulations of point sets [7, 30] can be adapted to count triangulations of polygons in the same time bounds.

2.2 Counting complexity

The complexity class #P is defined as the class of functional algorithmic problems for which the desired output counts the accepting paths of some nondeterministic polynomial-time Turing machine.

▶ **Lemma 2.** *Computing the number of triangulations of a polygon is in #P.*

Proof. The output is the number of accepting paths of a nondeterministic polynomial-time Turing machine that guesses the set of edges in the triangulation, and verifies that these edges form a triangulation of the input. ◀

#P-hardness and #P-completeness are defined using reductions, polynomial-time transformations from one problem X (typically already known to be hard) to another problem Y that we wish to prove hard. Three types of reduction are in common use for this purpose:

- *Turing reductions* consist of an algorithm for solving problem X in polynomial time given access to an oracle for solving problem Y.

- *Polynomial-time counting reductions* consist of two polynomial-time transformations: a transformation σ that transforms inputs to X into inputs to Y, and a second transformation τ that transforms outputs of Y back to outputs of X. The reduction is valid if, for every input χ to problem X, $\tau(Y(\sigma(\chi))) = X(\chi)$.
- *Parsimonious reductions* consist of a polynomial-time transformation from inputs of X to inputs of Y that preserve the exact solution value.

A problem Y is defined to be #P-hard for a given class of reductions if every problem X in #P has a reduction to Y. Y is #P-complete if, in addition, Y is itself in #P. Composing two reductions of the same type produces another reduction, so we will generally prove #P-hardness or #P-completeness by finding a single reduction from a known-hard problem and composing it with the reductions from everything else in #P to that known-hard problem.

The reductions that we construct in this work will be polynomial-time counting reductions. However, we rely on earlier work on #P-completeness of graph problems that uses the weaker notion of Turing reductions. Therefore, we will prove that our geometric problems are #P-complete under Turing reductions. If the graph-theoretic results are strengthened to use counting reductions (and in particular if counting maximum independent sets in regular planar graphs is #P-complete under counting reductions) then the same strengthening will apply as well to counting triangulations. For the remainder of this paper, however, whenever we refer to #P-hardness or #P-completeness, it will be under Turing reductions.

2.3 Counting independent sets

Xia, Zhang, and Zhao [43] proved that it is #P-complete to count the vertex covers in a connected 3-regular bipartite planar graph (subsets of vertices that touch all edges). A set of vertices is a vertex cover if and only if its complement is an independent set, so it immediately follows that it is also #P-complete to count independent sets in the same graphs (the counts are the same). It is also #P-complete to find the number of *maximum* independent sets in a planar bipartite graph of maximum degree three [40], but for our purposes the regularity of the graph is more important than the maximality of the independent set.
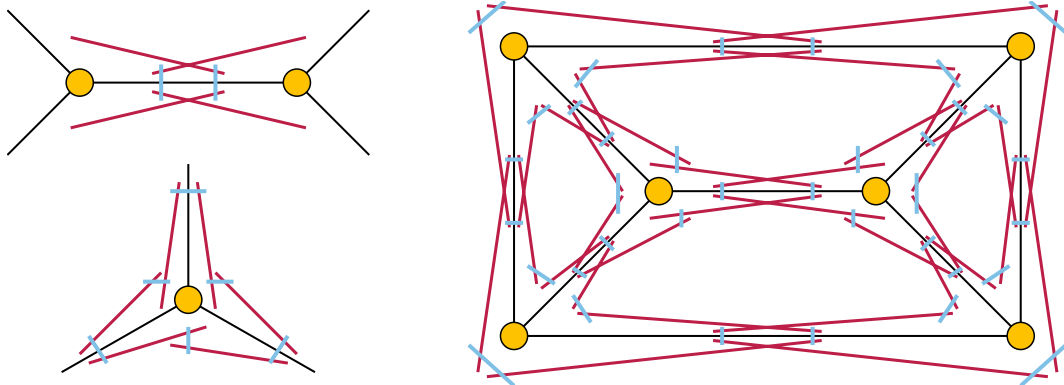
3 Red–blue arrangements

In this section we define and prove hard a counting problem that will be an intermediate step in our hardness proof for triangulations. It involves counting maximum non-crossing subsets of certain special line segment arrangements. Counting maximum non-crossing subsets of arbitrary line segment arrangements can easily be shown to be hard: It follows from the hardness of counting maximum independent sets in planar graphs [40] and from the proof of Scheinerman's conjecture that every planar graph can be represented as an intersection graph of line segments [34, 14]. So the significance of the reduction that we describe in this section is that it provides arrangements with a highly constrained form, a form that will be useful in our eventual reduction to counting triangulations.

▶ **Definition 3.** *We define a* red–blue arrangement *to be a collection of finitely many line segments in the plane (specified by the Cartesian coordinates of their endpoints) with the following properties:*
- *Each line segment is assigned a color, either red or blue.*
- *Each intersection point of two segments is a proper crossing point of exactly two segments.*
- *Each blue segment is crossed by exactly two other segments, both red.*
- *Each red segment is crossed by exactly three other segments, in the order blue–red–blue.*
- *The union of the segments forms a connected subset of the plane.*

Figure 1 Transformation from a 3-regular planar graph G to a red–blue arrangement. Left: the gadgets for edges and vertices of the graph. Right: an example of the whole arrangement for the graph of a triangular prism.

See the right side of Figure 1 for an example. We will be interested in the maximum non-crossing subsets of such an arrangement: sets of as many segments as possible, no two of which cross.

In a red–blue arrangement, consider the graph whose vertices represent line segments and whose edges represent bichromatic crossings (crossings between a red segment and a blue segment). Then this graph is a disjoint union of cycles, each of which has even length with vertices alternating between red and blue. We call the cycles of this graph *alternating cycles* of the arrangement.

▶ **Lemma 4.** *Every red–blue arrangement has equal numbers of red and blue segments. If there are n red and n blue segments, the maximum non-crossing subsets of the arrangement all have exactly n segments. Within each alternating cycle of the arrangement, a maximum non-crossing subset must use a monochromatic subset of the cycle (either all the red segments of the cycle, or all the blue segments of the cycle).*

Proof. The equal numbers of red and blue segments in the whole arrangement follow from the decomposition of the arrangement into alternating cycles and the equal numbers within each cycle. Within a single cycle, there are exactly two maximum non-crossing subsets, the subsets of red and of blue segments, each of which uses exactly half of the segments of the cycle. Therefore, no non-crossing subset of the whole arrangement can include more than n segments (half of the total), and a non-crossing subset that uses exactly n segments must be monochromatic within each alternating cycle. There exists at least one non-crossing subset of exactly n segments, namely the set of blue segments.   ◀

Given a connected 3-regular planar graph G, we will construct a red–blue arrangement A_G from it, as follows.

We begin by finding a straight-line drawing of G. We then replace each vertex v of G by a twelve-segment alternating cycle (Figure 1, lower left). Each edge of G incident to v will have two red segments on either side of it, within the two faces bounded by that edge, and these segments form the six red segments of the alternating cycle. Three blue segments, drawn near v within the three faces incident to v, connect pairs of red segments within each face. Another three blue segments cross the three edges incident to v (which are not themselves part of the arrangement) and connect the two red segments on either side of the edge.

In this way, each edge uv of G will have four red segments near it (two on each side from each of its two endpoints) and two blue segments crossing it (one from each endpoint). We arrange these segments so that, on each side of uv, the two red segments cross, and there are no other crossings except those in the alternating cycles (Figure 1, upper left). An example of the whole red–blue arrangement produced by these rules is shown in Figure 1, right.

By using an initial drawing of G within an $n \times n$ unit grid, where n is the number of vertices of G [35, 15], and then scaling the grid by a polynomial factor, we can create room to place each segment of the corresponding red–blue arrangement near its corresponding edge or vertex in the drawing of G, with integer coordinates for its endpoints. In this way, the entire construction can be performed in polynomial time, using segments whose endpoints have integer coordinates of polynomial magnitude. Because the coordinates are small integers, the angles between any two crossing segments will be at least inverse-polynomial. We omit the details.

▶ **Lemma 5.** *Let G be a connected 3-regular planar graph, and let A_G be constructed from G as above. Then the independent sets of G correspond one-for-one with the maximum non-crossing subsets of A_G, with an independent set having size k if and only if the corresponding non-crossing subset has exactly $6k$ red segments.*

Proof. Given an independent set Z in G, one can construct a non-crossing subset of A_G of size n, with $6|Z|$ red segments, by choosing all the red segments in the alternating cycles around members of Z, and all the blue segments in the remaining alternating cycles. Because Z includes no two adjacent vertices, the subsets chosen in this way will be non-crossing. Z can be recovered as the set of vertices from which we used red segments of alternating cycles, so the correspondence from independent sets to non-crossing subset constructed in this way is one-to-one.

By Lemma 4, these non-crossing subsets are maximum. By the same lemma, every maximum non-crossing subset consists of red segments in the alternating cycles of some vertices of G and blue segments in the remaining alternating cycles. No two adjacent vertices of G can have their red segments chosen as that would produce a crossing, so there are no maximum non-crossing subsets other than the ones coming from independent sets of G. ◀

This leads to the main result of this section:

▶ **Lemma 6.** *It is #P-complete to compute the number of maximum non-crossing subsets of a red–blue arrangement A.*

Proof. The problem is clearly in #P. By Lemma 5, the construction of the red–blue arrangement A_G from a 3-regular planar graph G is a parsimonious reduction from the known #P-complete problem of counting independent sets in 3-regular planar graphs. ◀

4 Reduction to triangulation

In this section we describe our reduction from line segment arrangements to polygons. The rough idea is to thicken each segment of the given red–blue arrangement to a rectangle with rounded ends, and form the union of these rectangles. In the part of the polygon formed from each segment, many more triangulations will use diagonals running end-to-end along the segment than triangulations that do not, causing most triangulations to correspond to sets of diagonals from a maximum non-crossing subset of the arrangement. With a careful choice of the shapes of the thickened segments in this construction, we can recover the number of maximum non-crossing subsets of the arrangement from the number of triangulations.
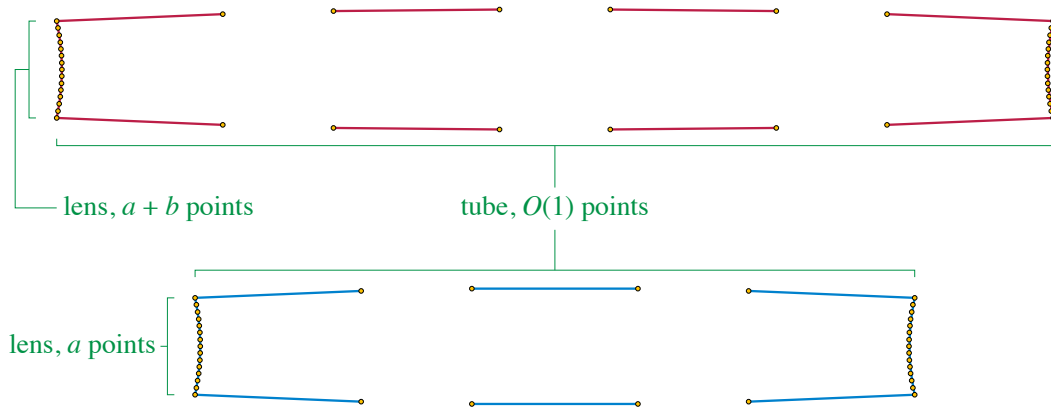
Figure 2 Polygonal gadgets for replacing the red segments (top) and blue segments (bottom) of a red–blue arrangement.

4.1   From segments to polygons

The gadgets that we use to replace the red and blue segments of a red–blue arrangement are shown (not to scale) in Figure 2. They resemble the cross-section of a telescope, and we have named their parts accordingly. Each gadget consists of $O(1)$ vertices spaced along two convex curves, which we call the tube. These two curves bend slightly outwards from each other, but both remain close to the center line of the gadget. Consecutive pairs of points are connected to each other along the tube, leaving either three gaps (for red segments) or two gaps (for blue segments) where other segments of the arrangement cross the given segment. The two vertices on each side of each gap are shared with the gadget for the crossing segment.

At the ends of the tubes, the two convex curves are connected to each other by two concave curves (the lenses of the gadget), each containing larger numbers of vertices that are connected consecutively to each other without gaps. Each blue lens has some number a of vertices, and each red lens has some larger number $a + b$ of vertices, where a and b are both positive integers to be determined later as a function of the number of segments in the arrangement. The key geometric properties of these gadgets are:

- Each gadget forms a collection of polygonal chains whose internal vertices (the vertices that are not the end-point of any of the chains) are not part of any other gadget and whose endpoints are part of exactly one other gadget. The union of all the gadgets of the arrangement forms a single polygon (with holes).
- Each vertex of the lens at one end of the gadget is visible within the polygon to each vertex of the lens at the other end of the gadget, and to each vertex of the tube of the gadget, but not to any vertices of its own lens that it is non-adjacent to.
- All of the vertices that are visible to a vertex of the lens of a gadget belong to the same gadget. (This is not an automatic property of the gadgets as drawn, but can be achieved by making the gadgets sufficiently narrow relative to the spacing of points along their tubes, depending on the angles at which different segments cross each other.)
- Each pair of points that are visible to each other are both part of one gadget. (Again, this is not an automatic property, but can be achieved by using narrow-enough gadgets.)

We denote by P_A the polygon obtained by replacing each segment of a red–blue arrangement A by a gadget. Figure 3 gives an example of a red–blue arrangement A and its corresponding polygon P_A, drawn approximately and schematically (not to scale).
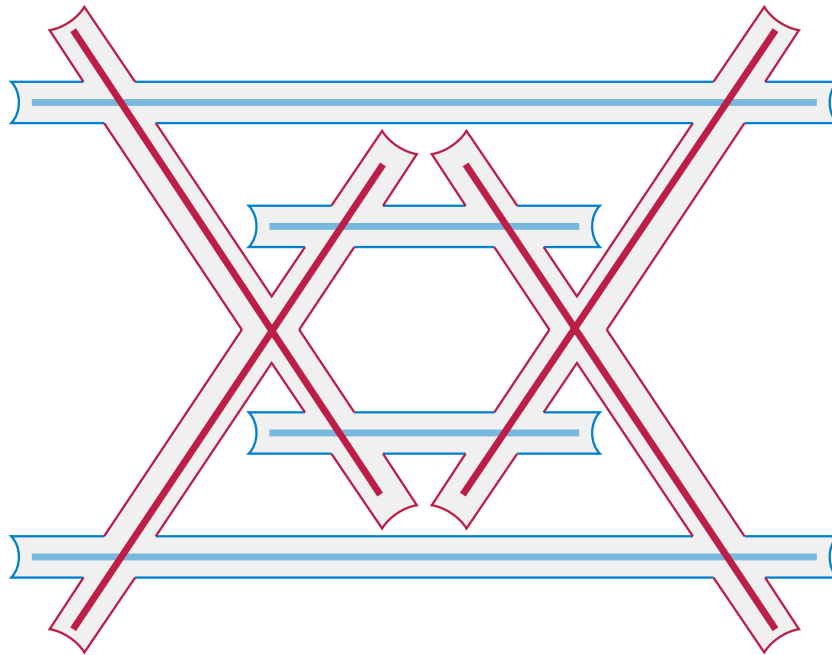
Figure 3 A red–blue arrangement A and the polygon P_A obtained from it by replacing its segments by gadgets (schematic view, not to scale).

As long as the segments of A have endpoints whose coordinates are integers of polynomial magnitude, and the parameters a and b are chosen to also be of polynomial magnitude, it will be possible to carry out the construction of P_A from A in polynomial time and for the coordinates of the resulting polygon to be integers of (larger) polynomial magnitude. We omit the details.

4.2 Triangulation within a gadget

For a triangulation T of a polygon P_A, and a gadget S of P_A corresponding to a segment s of A, we define the *local part* of T in S to be the set of triangles that have a side on one of the two lenses of S. We say that segment s is *active* in T if its local part includes at least one triangle that touches both lenses of S, and that s is *passive* otherwise.

▶ **Lemma 7.** *The active segments of a triangulation form a non-crossing set.*

Proof. If two segments cross, any lens-to-lens triangle within one segment crosses any lens-to-lens triangle of the other. But in a triangulation, no two triangles can cross each other.   ◀

▶ **Lemma 8.** *Let T be a triangulation of P_A, let s be an active segment of T, let S be the corresponding gadget, and let Π_S be the polygon formed by intersecting P_A with the convex hull of S. Then the intersection of T and Π_S is a triangulation of Π_S.*

Proof. By the construction of P_A, every edge of T must belong to a single gadget (as there are no other visibilities between pairs of vertices in P_A). By the same argument as in Lemma 7, no segment with its endpoints in a single other gadget than S can cross Π_S. Therefore, all edges of T that include an interior point of Π_S must connect two vertices of S, and lie within Π_S. It follows that all triangles of T that include an interior point of Π_S must also lie within Π_S, so that T restricts to a triangulation of Π_S.   ◀
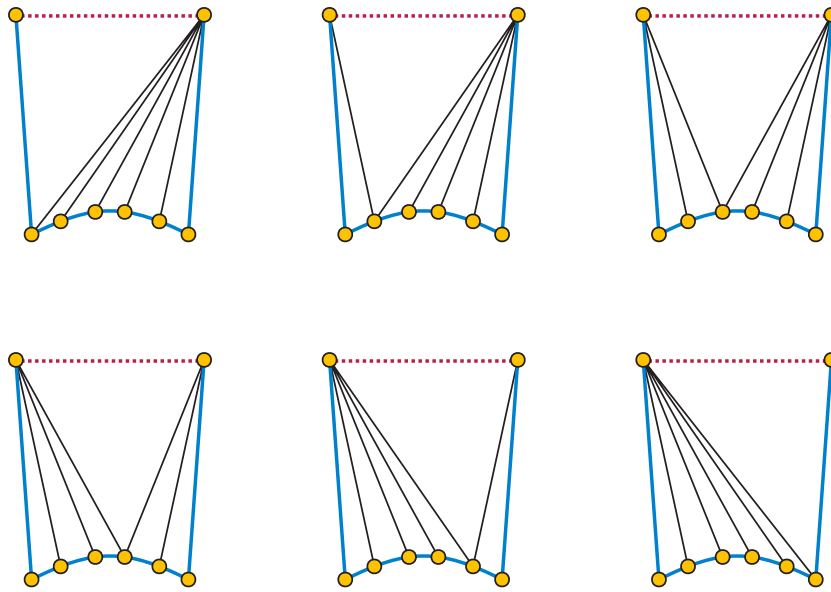
Figure 4 Illustration for Lemma 10: Closing off a lens by connecting the endpoints of its path (red dashed edge) forms a polygon with as many distinct triangulations as lens vertices.

▶ **Corollary 9.** *Let T be a triangulation of P_A, and let L be a lens in polygonal chain C_L of a gadget S. Suppose that some active segment of T separates C_L from the rest of S. Let Π_L be the polygon formed from C by adding one more edge connecting the two endpoints of C. Then T intersects Π_L in a triangulation of Π_L.*

Proof. By Lemma 8, T contains the edge connecting the endpoints of C. This edge separates Π_L from the rest of the polygon, so each triangle of T must lie either entirely within or entirely outside Π_L. But the triangles within Π_L cover Π_L (as the triangles of T cover all of P_A) so they form a triangulation of Π_L. ◀

These claims already allow us to produce a precise count of the triangulations whose active segments form a maximum non-crossing subset, as a function of the number of red segments in the subset, which we will do in the next section. However, we also need to bound the number of triangulations with smaller sets of active segments, and show that (with an appropriate choice of the parameters a and b) they form a negligible number of triangulations compared to the ones coming from maximum non-crossing subsets.

4.3   Counting within a gadget

It is straightforward to count triangulations of the polygon formed by a single lens:

▶ **Lemma 10.** *Let Π_L be the polygon formed from a lens by closing its path off by a single edge uv, as in Corollary 9. Then the number of distinct triangulations of Π_L equals the number of vertices of the lens (a for a blue lens, $a + b$ for a red lens).*

Proof. Any triangulation of Π_L is completely determined by the choice of apex of the triangle that has uv as one of its sides. All the remaining edges of the triangulation must each connect one vertex of the lens to the remaining visible endpoint of uv, the only vertex visible to that lens vertex (Figure 4). The number of choices for the apex equals the number of lens vertices. ◀
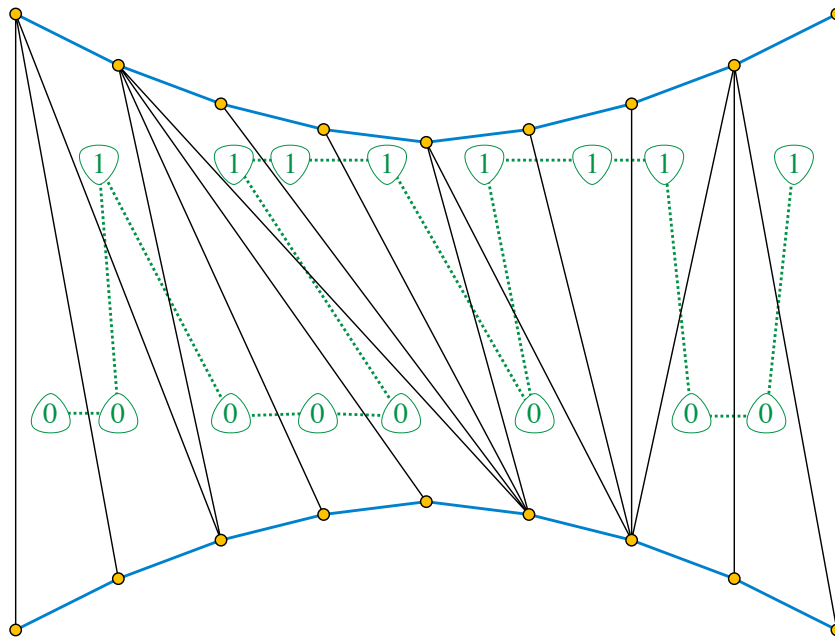
Figure 5 Encoding a triangulation of the convex hull of two lenses by a sequence of bits.

More generally, for any passive segment, we have:

▶ **Lemma 11.** *Let S be the gadget of segment s. Then the number of combinatorially distinct ways to choose the local part in S of a triangulation of P_A for which s is passive is polynomial in the number of lens vertices of S.*

Proof. Recall that the triangles of the local part each have one lens edge as one of their sides. Within each lens, the only choice is where to place the apex of each of these triangles. There are $O(1)$ choices of apex vertex, and within each lens the edges whose triangles have a given apex must form a contiguous subsequence. There are only polynomially many ways of partitioning the sequences of lens edges into a constant number of contiguous subsequences. ◀

A similar argument also shows:

▶ **Lemma 12.** *Let T be a triangulation of P_A and let S be a gadget of P_A. Then only $O(1)$ vertices of S can participate in triangles of T which are not part of the local part of T in S.*

Proof. S has $O(1)$ vertices outside of its two lenses, so it has $O(1)$ non-lens vertices that can participate in non-local triangles. A lens vertex can participate in a non-local triangle in one of two ways: either the triangle has one lens vertex and two non-lens vertices, or it has two vertices from opposite lenses and one non-lens vertex. There are $O(1)$ edges of T between non-lens vertices of S, each of which is part of two triangles in T, so there are $O(1)$ triangles in $S \cap T$ with only one lens vertex. Each non-lens vertex of S can participate in only one triangle whose other two vertices are on opposite lenses, so again there are $O(1)$ triangles of this type. Therefore, there are $O(1)$ lens vertices in non-local triangles. ◀

It is sufficiently messy to count the triangulations of active gadgets that we provide here only approximate bounds. However, the exact number of triangulations can be found in polynomial time using the algorithm for counting triangulations of simple polygons [19, 32, 16].
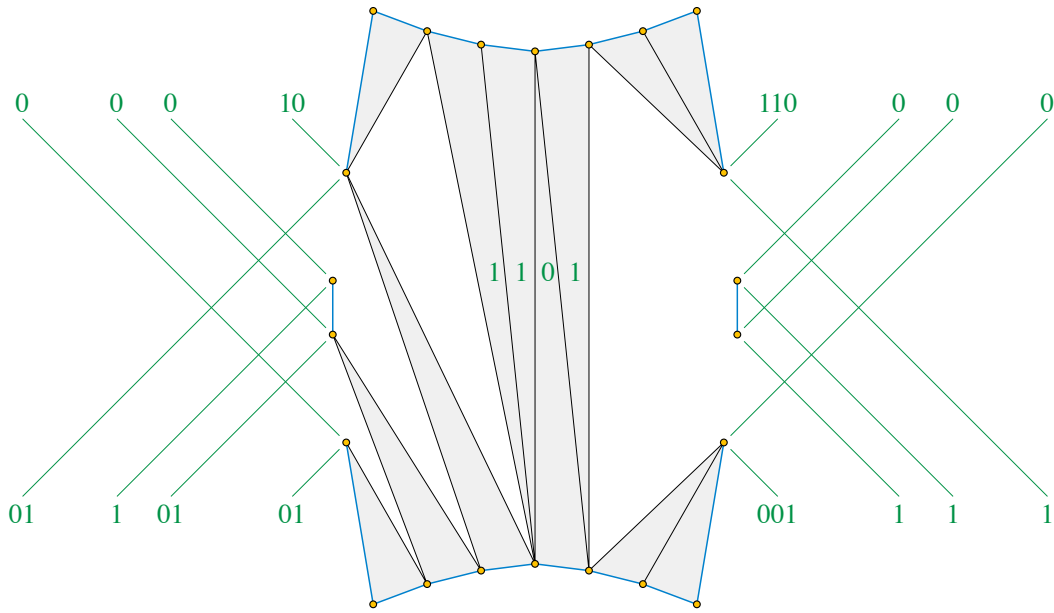
**Figure 6** Encoding the local part of a triangulation by a sequence of bits.

▶ **Lemma 13.** *Let Π_S be the polygon formed by intersecting polygon P_A with the convex hull of gadget S, as in Lemma 8. Let ℓ be the number of lens vertices of Π_S (2a for a blue gadget, $2a + 2b$ for a red gadget). Then the number of triangulations of Π_S is $\Theta(2^\ell/\sqrt{\ell})$.*

Proof. One can obtain $\Omega(2^\ell/\sqrt{\ell})$ triangulations by choosing a maximal set of non-crossing lens-to-lens diagonals in Π_S, which form the edges of a triangulation of the convex hull of the two lenses, and then choosing arbitrarily a triangulation of the remaining polygons of $O(1)$ vertices on either side of the convex hull of the two lenses. The triangles within any triangulation of the convex hull of the two lenses have a path as their dual graph, and if they are ordered along the path then the triangulation itself is determined by a sequence of bits (one bit per triangle) that denote whether each triangle includes an edge of one lens or of the other lens [22] (Figure 5). There are $\ell - 2$ bits in the sequence, exactly half of which must be zeros and half of which must be ones, so the number of these triangulations is

$$\binom{\ell - 2}{(\ell - 2)/2} = \Omega\left(\frac{2^\ell}{\sqrt{\ell}}\right).$$

In the other direction, we can produce an overestimate of the number of distinct local parts of triangulations of Π_S by a similar method of counting balanced binary strings of slightly greater length. As in Figure 5, consider one of the two lenses as being "upper" (labeled by 1's in the binary string) and the other as being "lower" (labeled by 0's). Similarly, we can describe the two sides of the tube of the gadget as being left or right. Suppose also that each side of a tube of the given gadget has t internal vertices (vertices that do not belong to either lens). We describe the local part of a triangulation by a sequence of bits, as follows (Figure 6):

- ▬ t blocks of 0-bits, each terminated by a single 1-bit, specifying how many edges of the lower lens form triangles whose apex is at each of the internal vertices of the left tube.
- ▬ t blocks of 1-bits, each terminated by a single 0-bit, specifying how many edges of the upper lens form triangles whose apex is at each of the internal vertices of the left tube.
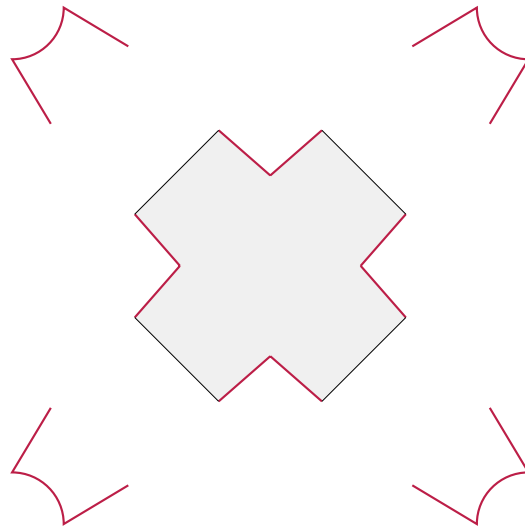
Figure 7 The twelve-sided polygon in the center of two crossing red gadgets.

- A sequence of bits as before describing the left-to-right sequence of triangles that have all three vertices on the two lenses, with a 0-bit for a triangle with a side on the lower lens and a 1-bit for a triangle with a side on the upper lens.
- t blocks of 0-bits, each terminated by a single 1-bit, specifying how many edges of the lower lens form triangles whose apex is at each of the internal vertices of the right tube.
- t blocks of 1-bits, each terminated by a single 0-bit, specifying how many edges of the upper lens form triangles whose apex is at each of the internal vertices of the right tube.

The resulting sequence of bits uniquely describes each distinct local part, has $2\ell - 2 + 4t$ bits, and has equal numbers of 0- and 1-bits. Not all such sequences of bits describe a valid local part (they may specify triangles connecting the lower and upper lenses to the tubes that cross each other) but this is non-problematic. There are $O(2^\ell/\sqrt{\ell})$ possible sequences of bits of this type, so there are $O(2^\ell/\sqrt{\ell})$ local parts of triangulations of Π_S. Each local part leaves remaining untriangulated regions on the left and right sides of Π_S with a bounded total number of vertices (as in the proof of Lemma 12) so it can be extended to a complete triangulation in $O(1)$ distinct ways. Therefore, the total number of triangulations of Π_S is $O(2^\ell/\sqrt{\ell})$. ◀

4.4 Global counting

The key properties of the number of triangulations of P_A (with an appropriate choice of a and b) that allow us to prove our counting reduction are that

- the number of triangulations coming from any particular maximum non-crossing subset of A can be determined only from n and the number of red segments in the subset, and
- non-crossing subsets that use fewer segments than the maximum, or that use the maximum total number of segments but fewer than the maximum possible number of red segments, contribute a negligible fraction of the total number of triangulations.

In this section we make these notions precise.

▶ **Definition 14.** *We define the following numerical parameters:*
- *α is the number of triangulations of Π_S for a blue gadget S, as bounded asymptotically as a function of parameter a by Lemma 13.*

- β is the number of triangulations of Π_S for a red gadget S, as bounded asymptotically as a function of parameters a and b by Lemma 13.
- γ is the number of triangulations of the twelve-sided polygon formed in the center of two crossing red gadgets by the four shared vertices of the two gadgets and their four neighbors in each gadget (Figure 7).
- $q_{n,r}$, for any $n \geq 0$ and $0 \leq r \leq n/2$, is $\alpha^{n-r} \beta^r \gamma^{(n-2r)/2} a^{2r}(a+b)^{2(n-r)} 2^{3r}$.

We do not derive an explicit formula for α and β. Therefore, it is important for us that they can be calculated in polynomial time by dynamic programming, as (for appropriate choices of a and b) they count triangulations of simple polygons of polynomial size.

▶ **Lemma 15.** *Let A be a red–blue arrangement with n red and n blue segments. Let Ξ be a maximum non-crossing subset of A, with $r = r(\Xi)$ red segments. Then the number of triangulations of P_A in which Ξ forms the active set is $q_{n,r}$.*

Proof. By Lemma 8, the triangulations with Ξ active can be partitioned into the polygons Π_S for the active segments of Ξ and the remaining sub-polygons left by the removal of these polygons, each of which can be triangulated independently and each of which contributes a term to the product in the definition of $q_{n,r}$. There are r active red segments and $n-r$ active blue segments, each of which contributes a factor of β or α (respectively) to the product by which $q_{n,r}$ was defined. There are $2r$ lenses of passive blue segments and $2(n-r)$ lenses of passive red segments, each of which contributes a term of a or $(a+b)$ to the product, respectively, by Lemma 10.

The r active red segments and r passive red segments that they cross leave $n-2r$ red segments that are passive but not crossed by another red segment. These $n-2r$ red segments form $(n-2r)/2$ twelve-sided polygons where their gadgets cross each other in pairs, with each pair contributing a factor of γ to the product.

The remaining factor of 2^{3r} in the product comes from the r passive blue gadgets, each of which has a central quadrilateral between the two red active segments that cross it, and from the r passive red gadgets that are crossed by an active red gadget, each of which has two central quadrilaterals between the consecutive pairs of the three active segments that cross it. These $3r$ quadrilaterals each can be triangulated in two ways. ◀

▶ **Corollary 16.** *Let $a \geq b$. Then $q_{n,r}/q_{n,r-1} = \Theta(2^{2b})$.*

Proof. Increasing the exponent of β by one and decreasing the exponent of α by one leads to the 2^{2b} change in the total. All of the other changes to the formula for $q_{n,r}$ are bounded either by a constant (independent of a and b) or by a power of $a/(a+b)$, which is at least $1/2$ by the assumption that $a \geq b$. ◀

▶ **Lemma 17.** *Suppose the parameters a and b are both upper-bounded by polynomials of n. Then*

$$\log_2 q_{n,r} = 2an + 2rb \pm O(n \log n),$$

where the constant in the O-notation depends on the bounds on a and b.

Proof. In the formula for $q_{n,r}$, each factor of α contributes $2a - O(\log n)$ to the logarithm, each factor of β contributes $2a + 2b - O(\log n)$, each factor of a or b contributes $O(\log n)$, and each factor of γ or 2 contributes $O(1)$. The result follows by adding these contributions according to their exponents. ◀

▶ **Lemma 18.** *Suppose the parameters a and b are both upper-bounded by polynomials of n. Let p denote the number of triangulations of P_A that do not have a maximum non-crossing subset of A as their active set. Then*

$$\log_2 p \leq 2a(n-1) + 2b(n-1) + O(n \log n).$$

Proof. There are 2^{2n} subsets of the $2n$ segments of A, and a smaller number of these subsets that can be a non-maximum active set of a triangulation. The choice of this subset adds $O(n)$ to the logarithm. For each non-maximum active set, there can be at most $n-1$ active segments, and the triangulations of the corresponding polygons adds at most $2a(n-1) + 2b(n-1) - O(n \log n)$ to the logarithm. The local parts of the passive gadgets can be chosen in a number of ways per gadget that is polynomial in a and b, adding another $O(n \log n)$ term to the logarithm. Once the active gadgets have been triangulated and the local parts of the passive gadgets have been chosen, the regions that remain to be triangulated have a total of $O(n)$ vertices, so the number of ways to triangulate them contributes another $O(n)$ to the logarithm. ◀

4.5 Completing the reduction

We have already described how to transform an arrangement A into a polygon P_A, modulo the choice of the parameters a and b. To complete the description of our reduction, we need to set a and b and we need to describe how to recover the number of maximum non-crossing subsets of A from the number N of triangulations of P_A.

Given a red–blue arrangement A with n red and n blue segments, set $b = 2n$. By Corollary 16, $q_{n,r-1}$ and $q_{n,r}$ differ by a factor of approximately 2^{4n}, much larger than the 2^{2n} bound on the total number of non-crossing sets of segments. Therefore, for any given r and for all sufficiently large n, the number of triangulations with a maximum non-crossing set of active segments that includes fewer than r red segments is strictly less than $q_{n,r}$. Next, set $a = 3n^2$. This is large enough that, comparing the bounds of Lemma 17 and Lemma 18, the number p of triangulations whose active set is non-maximum is strictly smaller than $q_{n,0}$, as the $O(n \log n)$ term and larger multiple of b in the formula for $\log_2 p$ are not large enough to make up for the smaller multiple of a in the same formula.

With these choices of a and b, the same reasoning also shows that (for all sufficiently large n) the total number of triangulations of any $P_{A'}$ for any arrangement A' with fewer red and blue segments than A is strictly smaller than $q_{n,0}$. Therefore, from the total number N of triangulations of P_A we can unambiguously determine the size of A.

With these considerations, we are ready to prove the correctness of our reduction:

Proof of Theorem 1. Counting triangulations of a given polygon P is in #P by Lemma 2, so we can complete the proof by describing a polynomial-time counting reduction from the number of maximum non-crossing subsets of a red–blue arrangement A (proved #P-hard in Lemma 6). To transform A into a polygon, let n be the number of red segments in A. Our reduction will work for all sufficiently large n; for the bounded values of n that are not sufficiently large, directly compute the number of maximum non-crossing subsets of A and construct a polygon with the same number of triangulations using the formula for numbers of triangulations of lens polygons of Lemma 10. Otherwise, choose $a = 4n^2$ and $b = 3n$ as above, and construct the polygon P_A. To transform the number N of triangulations of the given polygon into the number of maximum non-crossing subsets of A, first check whether this number is sufficiently small that it comes from our special-case construction for bounded values of n. If so, directly decode it to the number of maximum non-crossing subsets.

In the remaining case, recover n as the unique value that could have produced a polygon P_A with N triangulations. For each choice of r from n down to 0 (in decreasing order) compute the number of maximum non-crossing subsets with r red segments as $\lfloor N'/q_{n,r} \rfloor$ (where N' starts with the value N and is reduced as the algorithm progresses) and then replace N' by $N' \bmod q_{n,r}$ before proceeding to the next value of r. Sum the numbers of maximum-non-crossing subsets obtained for each value of r to obtain the total number of maximum non-crossing subsets.

When computing the number of non-crossing subsets for each value of r, the contributions from triangulations whose active segments include more red segments than r will already have been subtracted off, by induction. The contributions from triangulations whose active segments include fewer red segments than r, or from triangulations that do not have maximum non-crossing sets of active segments, will sum to less than a single multiple of the number of triangulations for each non-crossing set with the given number of red segments, as discussed above. Therefore, each number of non-crossing subsets is computed correctly.   ◀

5   Conclusions and open problems

We have shown that counting triangulations of polygons (with holes) is #P-complete under Turing reductions. It would be of interest to tighten this result to show completeness under counting reductions, or even under parsimonious reductions. Can this be done, either by strengthening the type of reduction used for the underlying graph problem that we reduce from (independent sets in regular planar graphs) or by finding a different reduction for triangulations that bypasses the Turing reductions used for this graph problem?

In a triangulation of a polygon with holes, every hole has a diagonal connecting its leftmost vertex to a vertex to the left of it in another boundary component. By testing all combinations of these left diagonals, and using dynamic programming to count triangulations of the simple polygon formed by cutting the input along one of these sets of diagonals (avoiding triangulations that use previously-tested diagonals) it is possible to count triangulations of an n-vertex polygon with h holes in time $O(n^{h+3})$. Is the dependence on h in the exponent of n necessary, or is there a fixed-parameter tractable algorithm for this problem?

More generally, there are many other counting problems in discrete geometry for which we neither know a polynomial time algorithm nor a hardness proof. For instance, we do not know the complexity of counting triangulations, planar graphs, non-crossing Hamiltonian cycles, non-crossing spanning trees, or non-crossing matchings of sets of n points in the plane. Are these problems hard?

—— **References** ——

1   Oswin Aichholzer, Victor Alvarez, Thomas Hackl, Alexander Pilz, Bettina Speckmann, and Birgit Vogtenhuber. An improved lower bound on the minimum number of triangulations. In Sándor Fekete and Anna Lubiw, editors, *Proceedings of the 32nd International Symposium on Computational Geometry (SoCG 2016)*, volume 51 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages A7:1–A7:16, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. `doi:10.4230/LIPIcs.SoCG.2016.7`.

2   Oswin Aichholzer, Thomas Hackl, Clemens Huemer, Ferran Hurtado, Hannes Krasser, and Birgit Vogtenhuber. On the number of plane geometric graphs. *Graphs and Combinatorics*, 23(suppl. 1):67–84, 2007. `doi:10.1007/s00373-007-0704-5`.

3   Oswin Aichholzer, Ferran Hurtado, and Marc Noy. A lower bound on the number of triangulations of planar point sets. *Comput. Geom. Theory and Applications*, 29(2):135–145, 2004. `doi:10.1016/j.comgeo.2004.02.003`.

4   Oswin Aichholzer, David Orden, Francisco Santos, and Bettina Speckmann. On the number of pseudo-triangulations of certain point sets. *J. Combin. Theory Ser. A*, 115(2):254–278, 2008. `doi:10.1016/j.jcta.2007.06.002`.

5   Victor Alvarez, Karl Bringmann, Radu Curticapean, and Saurabh Ray. Counting triangulations and other crossing-free structures via onion layers. *Discrete Comput. Geom.*, 53(4):675–690, 2015. `doi:10.1007/s00454-015-9672-3`.

6   Victor Alvarez, Karl Bringmann, Saurabh Ray, and Raimund Seidel. Counting triangulations and other crossing-free structures approximately. *Comput. Geom. Theory and Applications*, 48(5):386–397, 2015. `doi:10.1016/j.comgeo.2014.12.006`.

7   Victor Alvarez and Raimund Seidel. A simple aggregative algorithm for counting triangulations of planar point sets and related problems. In Timothy Chan and Rolf Klein, editors, *Proceedings of the 29th Annual Symposium on Computational Geometry (SoCG'13)*, pages 1–8, New York, 2013. ACM. `doi:10.1145/2462356.2462392`.

8   Emile E. Anclin. An upper bound for the number of planar lattice triangulations. *J. Combin. Theory Ser. A*, 103(2):383–386, 2003. `doi:10.1016/S0097-3165(03)00097-9`.

9   Takao Asano, Tetsuo Asano, and Hiroshi Imai. Partitioning a polygonal region into trapezoids. *J. ACM*, 33(2):290–312, 1986. `doi:10.1145/5383.5387`.

10   Andrei Asinowski and Günter Rote. Point sets with many non-crossing perfect matchings. *Comput. Geom. Theory and Applications*, 68:7–33, 2018. `doi:10.1016/j.comgeo.2017.05.006`.

11   Marshall W. Bern and David Eppstein. Mesh generation and optimal triangulation. In Ding-Zhu Du and Frank K. Hwang, editors, *Computing in Euclidean Geometry*, volume 4 of *Lecture Notes Series on Computing*, pages 47–123. World Scientific, 2nd edition, 1995.

12   Sergei Bespamyatnikh. An efficient algorithm for enumeration of triangulations. *Comput. Geom. Theory and Applications*, 23(3):271–279, 2002. `doi:10.1016/S0925-7721(02)00111-6`.

13   Hervé Brönnimann, Lutz Kettner, Michel Pocchiola, and Jack Snoeyink. Counting and enumerating pointed pseudotriangulations with the greedy flip algorithm. *SIAM J. Comput.*, 36(3):721–739, 2006. `doi:10.1137/050631008`.

14   Jérémie Chalopin and Daniel Gonçalves. Every planar graph is the intersection graph of segments in the plane: extended abstract. In Michael Mitzenmacher, editor, *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 631–638, 2009. `doi:10.1145/1536414.1536500`.

15   Hubert de Fraysseix, János Pach, and Richard Pollack. How to draw a planar graph on a grid. *Combinatorica*, 10:41–51, 1990. `doi:10.1007/BF02122694`.

16   Q. Ding, J. Qian, W. Tsang, and C. Wang. Randomly generating triangulations of a simple polygon. In Lusheng Wang, editor, *Computing and Combinatorics: 11th Annual International Conference, COCOON 2005, Kunming, China, August 16–19, 2005, Proceedings*, volume 3595 of *Lecture Notes in Computer Science*, pages 471–480. Springer, Berlin, 2005. `doi:10.1007/11533719_48`.

17   Adrian Dumitrescu, André Schulz, Adam Sheffer, and Csaba D. Tóth. Bounds on the maximum multiplicity of some common geometric graphs. *SIAM J. Discrete Math.*, 27(2):802–826, 2013. `doi:10.1137/110849407`.

18   Martin Dyer, Leslie Ann Goldberg, and Mike Paterson. On counting homomorphisms to directed acyclic graphs. *J. ACM*, 54(6):A27:1–A27:23, 2007. `doi:10.1145/1314690.1314691`.

19   Peter Epstein and Jörg-Rüdiger Sack. Generating triangulations at random. *ACM Trans. Model. Comput. Simul.*, 4(3):267–278, 1994. `doi:10.1145/189443.189446`.

20   Philippe Flajolet and Marc Noy. Analytic combinatorics of non-crossing configurations. *Discrete Math.*, 204(1-3):203–229, 1999. `doi:10.1016/S0012-365X(98)00372-0`.

21   Alfredo García, Marc Noy, and Javier Tejel. Lower bounds on the number of crossing-free subgraphs of K_N. *Comput. Geom. Theory and Applications*, 16(4):211–221, 2000. `doi:10.1016/S0925-7721(00)00010-9`.

22   Ferran Hurtado, Marc Noy, and Jorge Urrutia. Flipping edges in triangulations. *Discrete Comput. Geom.*, 22(3):333–346, 1999. `doi:10.1007/PL00009464`.

23   F. Jaeger, D. L. Vertigan, and D. J. A. Welsh. On the computational complexity of the Jones and Tutte polynomials. *Math. Proc. Cambridge Philos. Soc.*, 108(1):35–53, 1990. `doi:10.1017/S0305004100068936`.

24   Volker Kaibel and Günter M. Ziegler. Counting lattice triangulations. In C. D. Wensley, editor, *Surveys in Combinatorics 2003: Papers from the 19th British Combinatorial Conference held at the University of Wales, Bangor, June 29–July 4, 2003*, volume 307 of *London Math. Soc. Lecture Note Ser.*, pages 277–307. Cambridge Univ. Press, Cambridge, UK, 2003. `arXiv:math/0211268`.

25   Pegah Kamousi and Subhash Suri. Stochastic minimum spanning trees and related problems. In Philippe Flajolet and Daniel Panario, editors, *Proceedings of the Eighth Workshop on Analytic Algorithmics and Combinatorics, ANALCO 2011, San Francisco, California, USA, January 22, 2011*, pages 107–116. SIAM, 2011. `doi:10.1137/1.9781611973013.12`.

26   Goos Kant and Hans L. Bodlaender. Triangulating planar graphs while minimizing the maximum degree. *Inform. and Comput.*, 135(1):1–14, 1997. `doi:10.1006/inco.1997.2635`.

27   Marek Karpinski, Andrzej Lingas, and Dzmitry Sledneu. A QPTAS for the base of the number of crossing-free structures on a planar point set. *Theoretical Computer Science*, 711:56–65, 2018. `doi:10.1016/j.tcs.2017.11.003`.

28   Nathan Linial. Hard enumeration problems in geometry and combinatorics. *SIAM J. Algebraic Discrete Methods*, 7(2):331–335, 1986. `doi:10.1137/0607036`.

29   Anna Lubiw. Decomposing polygonal regions into convex quadrilaterals. In Joseph O'Rourke, editor, *Proceedings of the 1st Symposium on Computational Geometry, Baltimore, Maryland, USA, June 5-7, 1985*, pages 97–106, New York, 1985. ACM. `doi:10.1145/323233.323247`.

30   Dániel Marx and Tillmann Miltzow. Peeling and nibbling the cactus: subexponential-time algorithms for counting triangulations and related problems. In Sándor Fekete and Anna Lubiw, editors, *Proceedings of the 32nd International Symposium on Computational Geometry (SoCG 2016)*, volume 51 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages A52:1–A52:16, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. `doi:10.4230/LIPIcs.SoCG.2016.52`.

31   Alexander Pilz and Carlos Seara. Convex quadrangulations of bichromatic point sets. In *Proceedings of the 33rd European Workshop on Computational Geometry (EuroCG 2017)*, 2017. URL: `https://mat-web.upc.edu/people/carlos.seara/data/publications/internationalConferences/EuroCG-17_paper_23.pdf`.

32   Saurabh Ray and Raimund Seidel. A simple and less slow method for counting triangulations and for related problems. In *Proceedings of the 20th European Workshop on Computational Geometry (EuroCG 2004)*, 2004. URL: `https://hdl.handle.net/11441/55368`.

33   Francisco Santos and Raimund Seidel. A better upper bound on the number of triangulations of a planar point set. *J. Combin. Theory Ser. A*, 102(1):186–193, 2003. `doi:10.1016/S0097-3165(03)00002-5`.

34   Edward R. Scheinerman. *Intersection Classes and Multiple Intersection Parameters of Graphs*. PhD thesis, Princeton University, 1984.

35   Walter Schnyder. Embedding planar graphs on the grid. In David S. Johnson, editor, *Proceedings of the First Annual ACM–SIAM Symposium on Discrete Algorithms, SODA 1990, 22-24 January 1990, San Francisco, California, USA*, pages 138–148, 1990. URL: `https://dl.acm.org/citation.cfm?id=320176.320191`.

36   Raimund Seidel. On the number of triangulations of planar point sets. *Combinatorica*, 18(2):297–299, 1998. `doi:10.1007/PL00009823`.

37   Micha Sharir and Adam Sheffer. Counting triangulations of planar point sets. *Electron. J. Combin.*, 18(1):P70:1–P70:74, 2011. URL: `https://emis.ams.org/journals/EJC/Volume_18/PDF/v18i1p70.pdf`.

38   Micha Sharir, Adam Sheffer, and Emo Welzl. Counting plane graphs: perfect matchings, spanning cycles, and Kasteleyn's technique. *J. Combin. Theory Ser. A*, 120(4):777–794, 2013. `doi:10.1016/j.jcta.2013.01.002`.

39 Micha Sharir and Emo Welzl. On the number of crossing-free matchings, cycles, and partitions. *SIAM J. Comput.*, 36(3):695–720, 2006. `doi:10.1137/050636036`.

40 Salil P. Vadhan. The complexity of counting in sparse, regular, and planar graphs. *SIAM J. Comput.*, 31(2):398–427, 2001. `doi:10.1137/S0097539797321602`.

41 L. G. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8(2):189–201, 1979. `doi:10.1016/0304-3975(79)90044-6`.

42 Manuel Wettstein. Counting and enumerating crossing-free geometric graphs. *J. Comput. Geom.*, 8(1):47–77, 2017. URL: `https://jocg.org/index.php/jocg/article/view/280`.

43 Mingji Xia, Peng Zhang, and Wenbo Zhao. Computational complexity of counting problems on 3-regular planar graphs. *Theoretical Computer Science*, 384(1):111–125, 2007. `doi:10.1016/j.tcs.2007.05.023`.