

Counting Triangulations and other Crossing-Free Structures Approximately

Victor Alvarez* Karl Bringmann† Saurabh Ray‡ Raimund Seidel§

April 25, 2014

Abstract

We consider the problem of counting straight-edge triangulations of a given set P of n points in the plane. Until very recently it was not known whether the *exact* number of triangulations of P can be computed asymptotically faster than by enumerating all triangulations. We now know that the number of triangulations of P can be computed in $O^*(2^n)$ time [9], which is less than the lower bound of $\Omega(2.43^n)$ on the number of triangulations of *any* point set [30]. In this paper we address the question of whether one can approximately count triangulations in sub-exponential time. We present an algorithm with sub-exponential running time and sub-exponential approximation ratio, that is, denoting by Λ the output of our algorithm and by c^n the exact number of triangulations of P , for some positive constant c , we prove that $c^n \leq \Lambda \leq c^n \cdot 2^{o(n)}$. This is the first algorithm that in sub-exponential time computes a $(1 + o(1))$ -approximation of the base of the number of triangulations, more precisely, $c \leq \Lambda^{\frac{1}{c^n}} \leq (1 + o(1))c$. Our algorithm can be adapted to approximately count other crossing-free structures on P , keeping the quality of approximation and running time intact. In this paper we show how to do this for matchings and spanning trees.

1 Introduction

Let P be a set of n points on the plane. A crossing-free structure on P is a straight-line plane graph with vertex set P . Examples of crossing-free structures include triangulations, trees, matchings, and spanning cycles, also known as the polygonizations of P , among others. Let X denote certain type of crossing-free structures and let $\mathcal{F}_X(P)$ denote the set of all crossing-free structures on P of type X . One of the most intriguing

*Information Systems Group, Universität des Saarlandes, Germany, alvarez@cs.uni-saarland.de.

†Max Planck Institute for Informatics, Karl Bringmann is a recipient of the *Google Europe Fellowship in Randomized Algorithms*, and this research is supported in part by this Google Fellowship. kbringma@mpi-inf.mpg.de.

‡Ben Gurion University of the Negev, Israel saurabh@math.bgu.ac.il.

§Fachrichtung Informatik, Universität des Saarlandes, Germany, rseidel@cs.uni-saarland.de.

problems in Computational Geometry is the following: given P and a particular type of crossing-free structures X , how fast can we compute the cardinality of $\mathcal{F}_X(P)$?

Among all kinds of crossing-free structures on P , triangulations are perhaps the most studied ones, so let us first focus on them. How fast can we compute the number of triangulations of P ? For starters, although triangulations are the most studied, the problem of computing the number of *all* triangulations of P is, in general, neither known to be #P-hard, nor do we know of a polynomial time algorithm to even approximate that number. When the point set is in convex position, an easy recurrence relation can be derived showing that the number of triangulations spanned by n points in convex position is C_{n-2} , where C_k is the k -th Catalan number. F. Hurtado and M. Noy [21] were able to find further formulas to exactly compute the number of triangulations of “almost convex sets”. Unfortunately, the approach of finding exact formulas does not seem to take us much further.

For *any* given set of points P it is possible to enumerate all its triangulations in time proportional to its number of triangulations. This is because the *flip graph** of triangulations is connected, see [23, 33]. Thus any graph traversal algorithm like DFS or BFS can be used to enumerate the vertices of the flip graph of triangulations of P . One limitation of such traversal algorithms is, however, that the amount of memory used is proportional to the number of vertices in the graph — which is known to *always* be exponential, because the number of triangulations of P lies between $\Omega(2.43^n)$ [30] and $O(30^n)$ [29]. Using a general technique due to D. Avis and K. Fukuda, called Reverse Search [10], it is possible to enumerate triangulations while keeping the memory usage polynomial in n . This technique has been further improved in [11, 22]. However, it is important to observe that, since the number of triangulations is exponential in n , counting triangulations by enumeration takes exponential time, so a natural question is whether one can do significantly better.

At this point let us mention that it is not known whether the aforementioned bounds of $\Omega(2.43^n)$ and $O(30^n)$ for the number of triangulations are realizable, i.e., whether there are indeed set of points having that many triangulations. In terms of realizability, the best we know at this point is that there exists a set of n points having $O^*(3.47^n)^\dagger$ [3] triangulations, but there is also a set of points where this number grows to at least $\Omega(8.65^n)$ [15]. This leads to the interesting observation that the number of triangulations is not minimized when the set of points is in convex position; we already mentioned that the number of triangulations of set of n points in convex position is C_{n-2} , which asymptotically grows as $\Theta^*(4^n)$.

In [1] O. Aichholzer introduced the notion of the “path of a triangulation”, T-path from now on, that allows a divide-and-conquer approach to speed up counting. From

*The flip graph is a graph whose vertex set is the set of triangulations of P and where there is an edge between vertices of the flip graph if both corresponding triangulations can be transformed into each other by flipping exactly one of their edges.

[†]In the notation $\Omega^*(\cdot)$, $O^*(\cdot)$, and $\Theta^*(\cdot)$ we neglect polynomial terms and we just present the dominating exponential term.

empirical observations, this approach seems to count triangulations in time sub-linear in the number of triangulations, that is, apparently faster than enumeration. Unfortunately, no proof is known that this algorithm is *always* faster than enumeration, and no good analysis of its running time has been obtained yet. Subsequently, a simple algorithm based on dynamic programming was presented in [26]. This algorithm empirically appears to be substantially faster than Aichholzer’s algorithm. From the limited empirical data, the running time seems to be proportional to the square root of the number of triangulations. However, as with Aichholzer’s algorithm, the worst case running time of this algorithms seems difficult to analyze and no bound on the running time has been found.

Very recently two algorithms whose running times could actually be analyzed were presented. The first algorithm, presented in [8], combines Aichholzer’s idea of T-paths with a sweep-line algorithm and runs in time proportional to the largest number of T-paths found during execution (within a *poly*(n) factor). In [8] the number of T-paths is shown to be at most $O(9^n)$. It is important to observe that, for very particular and well-studied configurations of points, the number of T-paths can be shown to be significantly smaller than the number of triangulations. Thus, at least for those configurations, the algorithm of [8] counts triangulations faster than by using enumeration techniques. Unfortunately, this algorithm turned out to be very slow in practice, which is most probably due to the fact that the number of T-paths can still be very large — in [14] a configuration having at least $\Omega(4^n)$ T-paths was shown.

The second algorithm, presented in [6], uses a divide-and-conquer approach based on the onion layers of the given set of points. This algorithm was shown to have a running time of at most $O^*(3.1414^n)$ and is likely to have a running time sub-linear in the number of triangulations since it is widely believed that the number of triangulations spanned by *any* set of n points is at least $\Omega^*(3.47^n)^{\ddagger}$, see [28, 4, 3]. From the experimental point of view, the second algorithm turned out to be significantly faster, for certain configurations of points, than the one shown in [26]. These experiments can be found in [7]. Finally, in 2013, an algorithm with running time $O^*(2^n)$ was presented [9]. This last algorithm shows that enumeration algorithms for triangulations can indeed *always* be beaten, as all point sets with n points have at least $\Omega(2.43^n)$ triangulations. From an experimental point of view it was also shown to be significantly faster than all previous algorithms on a variety of inputs [9].

With respect to crossing-free structures other than triangulations the situation is very similar. In [22] a general framework for enumerating crossing-free structures (including spanning trees and perfect matchings) was presented. However, as for triangulations, the number of enumerated objects is again in general exponential. Table 1 gives the latest asymptotic bounds on the number of triangulations, perfect matchings, and spanning trees, all crossing-free, found on a set of n points P in the plane. The symbols \leq, \geq should be understood as upper and lower bound, respectively.

[‡]The exact term is $\sqrt{12}^{n-\Theta(\log n)}$ [3].

	Triangulations	Perfect matchings	Spanning trees
$\forall P \leq$	$O(30^n)$ [29]	$O(10.05^n)$ [31]	$O(141.07^n)$ [19]
$\forall P \geq$	$\Omega(2.43^n)$ [30]	$\Omega^*(2^n)$ [17]	$\Omega^*(6.75^n)$ [16]
$\exists P \leq$	$O^*(3.47^n)$ [3]	$O^*(2^n)$ [17]	$O^*(6.75^n)$ [16]
$\exists P \geq$	$\Omega(8.65^n)$ [15]	$\Omega^*(3^n)$ [17]	$\Omega(12.52^n)$ [20]

Table 1: Asymptotic bounds for various classes of crossing-free structures on P (a set of n points in the plane). The selected (gray) lower bounds are tight.

We remark that it has been proven that the number of perfect matchings, and the number of spanning trees, among others, is minimized when P is in convex position, see [2] and references therein. That is, the selected (gray) lower bounds in Table 1 are tight, i.e., there exists a set of n points P having at most that many perfect matchings [17], and spanning trees [16]. The interested reader is referred to [2] for a list of other classes of crossing-free structures on P whose cardinality is minimized when P is in convex position, and to [32, 15] for up-to-date lists of asymptotic bounds for other crossing-free structures.

Thus, the question arises again whether we can count these structures faster than by enumeration. In this respect, in [8] an algorithm was shown that counts pseudo-triangulations in time proportional to the largest number of pseudo-triangulation paths [5] that the algorithm encounters. It is, however, still not known whether there are always considerably less pseudo-triangulation paths than pseudo-triangulations. In [27] an algorithm was shown that counts *all* crossing-free structures of a given set of points P in time sub-linear in the number of counted objects, thus achieving the desired exponential speed-up in this case. New fast counting algorithms appeared very recently: In [34, 35] it was shown, among other results, that the number of *all* crossing-free structures can be computed in time $O^*(2.839^n)$, while perfect matchings and convex partitions can be counted in time $O^*(2^n)$. These algorithms generalize the idea presented in [9] and show that enumeration can, at least in these cases, *always* be beaten. It is, however, still open whether for spanning trees and spanning cycles, two of the most popular classes of crossing-free structures, the same can be said, see [34, 35] for partial results in this direction.

1.1 Our contribution

The $O^*(2^n)$ algorithm of [9] for counting triangulations *exactly* seems hard to beat at this point. Instead, we relax the goal to computing an *approximation* of the number of triangulations and pose the question of whether one can reduce the runtime in this setting. The answer presented in this paper is, to the best of our knowledge, the first result in this new line of research.

Note that, since for all sets of n points the number of triangulations is $\Omega(2.43^n)$ [30] and $O(30^n)$ [29], the quantity $\Theta(\sqrt{30 \times 2.43}^n)$ approximates the exact number of trian-

gulations within a factor of $O\left(\sqrt{30/2.43}^n\right)$. Thus, one can trade the exponential time of an exact algorithm for a polynomial time algorithm with exponential approximation ratio. In this paper we bridge the gap between these two solutions by presenting an algorithm with sub-exponential running time and sub-exponential approximation ratio.

Let $\mathcal{F}_X(P)$ denote the set of all crossing-free structures of type X on P , where X could mean triangulations, matchings, or spanning trees. The main result of this paper is the following:

Theorem 1. *Let P be a set of n points on the plane. Then a number Λ can be computed in time $2^{O(\sqrt{n}\log n)}$ such that $|\mathcal{F}_X(P)| \leq \Lambda \leq |\mathcal{F}_X(P)| \cdot 2^{O(n^{3/4}\sqrt{\log n})} = |\mathcal{F}_X(P)|^{1+o(1)}$.*

At the end of § 4 we mention a trade-off between the running-time and the approximation factor stated in Theorem 1.

While the approximation factor of Λ is rather big, the computed value is of the same order of magnitude as the correct value of $|\mathcal{F}_X(P)|$, that is, we compute a $(1 + o(1))$ -approximation of the base of the exponentially large value $|\mathcal{F}_X(P)|$. More precisely, if we denote $|\mathcal{F}_X(P)|$ by c^n , for some positive c that depends on P and X and lies between two positive constants, then we have $c \leq \Lambda^{\frac{1}{n}} \leq c^{1+o(1)} \leq (1 + o(1))c$. Also, this approximation can be computed in sub-exponential time, which, at least theoretically, is asymptotically faster than the worst-case running times of the algorithms presented in [8, 6, 9, 34, 35]. This is certainly very appealing.

The rest of the paper is divided as follows: We start in § 2 with basic preliminaries. For simplicity, we first show in § 3 the algorithm for counting triangulations approximately and in § 4 the corresponding proof of Theorem 1. We generalize the algorithm for counting matchings and spanning trees in § 5. We close the paper in § 6 with some remarks and conclusions.

2 Preliminaries

Our algorithm uses simple cycle separators as the main ingredient, originally presented in [25] by G. L. Miller, and improved in [13] by H. N. Djidjev and S. M. Venkatesan. The following theorem accounts for both results:

Theorem 2 (Separator Theorem). *Let T be a triangulation of a set of n points in the plane such that the unbounded face is a triangle. Then there exists a simple cycle C , made of edges of T , of size at most $\sqrt{4n}$, that separates the set A of vertices of T in its interior from the set B of vertices of T in its exterior, such that the number of elements of each one of A and B is at most $\frac{2n}{3}$.*

An example of a simple cycle separator can be seen in Figure 1(a). Observe, however, that the Separator Theorem does not imply that *every* triangulation of a set of points contains a *unique* simple cycle separator. One can easily come up with examples in which a triangulation contains more than one simple cycle separator. The important part here is that *every* triangulation contains *at least* one simple cycle separator.




Figure 1: Left — A simple cycle separator C , shown in black, of a triangulation T . The set A of vertices of T contained in the interior C_{int} of C , shown in dark gray, is separated by C from the set B of vertices of T in the exterior of C . Right — The outer boundary $\partial\mathcal{P}$ of a sub-problem \mathcal{P} is shown in black. The size of $\partial\mathcal{P}$ could be larger than three, and thus we cannot apply the Separator Theorem directly. We circumvent this by arbitrarily triangulating the outer face, adjacencies are shown as dashed edges. The interior of a simple cycle separator for this augmented configuration is shown in dark gray. This separator can be seen as a path separator connecting the vertices u_1, u_2 of $\partial\mathcal{P}$.

3 Counting triangulations approximately

The idea for an approximate counting algorithm is suggested by the Separator Theorem: (1) Enumerate the set \mathcal{S} of *all* possible simple cycle separators that we can find in the given set P ; the size of each separator in \mathcal{S} is at most $\sqrt{4n}$. (2) For each simple cycle separator $C \in \mathcal{S}$, recursively compute the number of triangulations of each of the parts A and B , specified by the Separator Theorem, that are delimited by C . We then multiply the number we obtain for the sub-problem $A \cup C$ by the number we obtain for sub-problem $B \cup C$, and we add these products over all cycle separators of \mathcal{S} .

Note that the sub-problem corresponding to the outer part B contains C as a hole. Thus, in general a sub-problem \mathcal{P} of our algorithm can be described by a pair $\mathcal{P} = (R, P)$, where R is a polygonal region in the plane with polygonal holes (at initialization R is the interior of the convex hull of the given point set) and $P \subset R$ is a set of points; for simplicity we assume that P contains the vertices of R . The task is to count the number of triangulations of P in R . We let $|\mathcal{P}| := |P|$.

Note that we cannot directly apply the Separator Theorem to such a general sub-problem, since (1) the outer face can contain more than three vertices and (2) we may have holes. To overcome these problems, we fix an arbitrary triangulation of the outer face as well as of all holes of the region R (in a combinatorial sense), see Figures 1(b) and 2(a). From now on we call any triangulation of P in R together with the fixed triangulations of the outer face and holes of R a *triangulation of \mathcal{P}* . Note that such a triangulation of \mathcal{P} satisfies all requirements of the Separator Theorem, so that a small

simple cycle separators exists. We remark that such a simple cycle separator can alternatively be thought of as paths connecting some holes and possibly the outer face of R , see Figure 2(a).

Hence, we can enumerate all simple cycle separators that may appear in any triangulation of \mathcal{P} , we postpone the details to Lemma 1. For each such separator C (with vertices $V(C)$ and interior C_{int}) we split up the sub-problem $\mathcal{P} = (R, P)$ into

$$\begin{aligned}\mathcal{P}_A &:= \mathcal{P}_A(C) := (R \cap C_{\text{int}}, V(C) \cup (P \cap C_{\text{int}})) \quad \text{and} \\ \mathcal{P}_B &:= \mathcal{P}_B(C) := (R \setminus C_{\text{int}}, V(C) \cup (P \setminus C_{\text{int}})),\end{aligned}$$

see Figure 2(b). Note that both regions $R \cap C_{\text{int}}$ and $R \setminus C_{\text{int}}$ are again polygonal regions with polygonal holes. Moreover, the vertices of C are part of both sub-problems.

The algorithm is given in pseudo-code in Algorithm 1 of page 7. With this algorithm we clearly (over-)count the triangulations of the given point set P . We show in § 4.1 that the overcounting is kept under control as long as we solve small recursive sub-problems exactly, see lines 8 and 9. We discuss the running time in § 4.2.

Algorithm 1 Counting triangulations approximately

Input: Set P of n points in the plane.

Output: Number Λ , computed in time $2^{o(n)}$, such that $c^n \leq \Lambda \leq c^{n+o(n)}$, where c^n is the exact number of triangulations of P .

```

1:  $\Gamma \leftarrow$  Stopping size dependent on  $n$  ▶ Global variable defining a stopping condition.
2:  $R \leftarrow \text{Conv}(P)$  ▶  $R$  is initially the interior of the convex hull of  $P$ .
3:  $\mathcal{P} \leftarrow (R, P)$  ▶  $\mathcal{P}$  is the initial sub-problem.
4:  $\text{CNT\_TRL\_APPROX}(\mathcal{P})$  ▶ Function call to compute the number of triangulations. See below.

5: function  $\text{CNT\_TRL\_APPROX}(\mathcal{P} = (R, P))$ 
6:    $\Lambda \leftarrow 0$ 
7:    $n \leftarrow |P|$ 
8:   if  $n \leq \Gamma$  then
9:      $\Lambda \leftarrow \text{CNT\_TRL\_EXACT}(\mathcal{P})$  ▶ Exact counting for small sub-problems.
10:  else
11:     $\mathcal{S} \leftarrow \text{ALL\_CYCLE\_SEPARATORS}(\mathcal{P})$  ▶ Enumerating the set  $\mathcal{S}$  of all cycle separators of  $\mathcal{P}$ .  
Each separator is of size at most  $\sqrt{4n}$ .
12:    for each  $C \in \mathcal{S}$  do
13:       $V(C) \leftarrow$  set of vertices of  $C$ 
14:       $C_{\text{int}} \leftarrow$  interior of  $C$ 
15:       $\mathcal{P}_A \leftarrow (R \cap C_{\text{int}}, V(C) \cup (P \cap C_{\text{int}}))$  ▶ The outer face of  $\mathcal{P}_A$  is  $C$ .
16:       $\mathcal{P}_B \leftarrow (R \setminus C_{\text{int}}, V(C) \cup (P \setminus C_{\text{int}}))$  ▶ In  $\mathcal{P}_B$  separator  $C$  becomes a hole.
17:       $\Lambda_A \leftarrow \text{CNT\_TRL\_APPROX}(\mathcal{P}_A)$  ▶ Recursive computation.
18:       $\Lambda_B \leftarrow \text{CNT\_TRL\_APPROX}(\mathcal{P}_B)$  ▶  $|\mathcal{P}_A|, |\mathcal{P}_B| \leq \frac{2n}{3} + \sqrt{4n}$ 
19:       $\Lambda \leftarrow \Lambda + (\Lambda_A \cdot \Lambda_B)$ 
20:    end for
21:  end if
22:  return  $\Lambda$ 
23: end function

```




Figure 2: Left — A sub-problem $\mathcal{P} = (R, P)$. Region R , shown in light gray, is in general a polygonal region with holes, shown in dark gray. The triangulation of each hole and of the outer face is arbitrary — the triangulation of the outer face is not shown for simplicity. The adjacencies of every simple separator C of \mathcal{P} must follow the adjacencies of the triangulated holes if the separator intersects them; separator C is shown with dashed lines. Right — The recursive sub-problems $\mathcal{P}_A, \mathcal{P}_B$ that stem from C are shown separately.

Our first lemma tells us how much time is spent enumerating the set \mathcal{S} of all simple cycle separators in terms of the problem size $n = |\mathcal{P}|$. This corresponds to line 11 of Algorithm 1.

Lemma 1. *Let \mathcal{P} be a sub-problem, $n = |\mathcal{P}|$, and consider the set \mathcal{S} of all simple cycle separators of size at most $\sqrt{4n}$ over all triangulation of \mathcal{P} . The set \mathcal{S} can be enumerated in time $2^{O(\sqrt{n} \log n)}$.*

Proof. Let $\mathcal{P} = (R, P)$. We know by the Separator Theorem and the discussion beneath that every triangulation contains at least one separator C . Moreover, the size of C is at most $\sqrt{4n}$. Thus, searching by brute-force will do the job. We simply enumerate all sequences of at most $\sqrt{4n}$ points of P and test each such sequence to determine whether it (1) forms a simple cycle, (2) is consistent with the triangulations of the outer face and holes of R , and (3) fulfills the dividing properties of the Separator Theorem.

Clearly, the number of tested sequences is $2^{O(\sqrt{n} \log n)}$. Verifying whether a sequence of points is indeed a separator can be done in polynomial time. Thus, the total time spent remains being $2^{O(\sqrt{n} \log n)}$. ■

4 Proof of Theorem 1

We can now proceed with the corresponding proof of Theorem 1. We first prove that the approximation ratio of the algorithm for counting triangulations is sub-exponential and then we prove that its running time is sub-exponential as well.

4.1 Quality of approximation

We analyze the *overcounting factor* $S(\mathcal{P}, \Gamma)$ defined as the value returned by our algorithm on the sub-problem \mathcal{P} (using stopping size Γ for exactly solving small sub-problems) divided by the total number of triangulations of \mathcal{P} . This determines the approximation quality of our algorithm. We already argued that every triangulation is counted at least once, so $S(\mathcal{P}, \Gamma) \geq 1$. In this section we prove an upper bound on $S(\mathcal{P}, \Gamma)$.

By the proof of Lemma 1 we also obtain that the number of simple cycle separators for a sub-problem \mathcal{P} of size n is at most $2^{O(\sqrt{n} \log n)}$. Since at *every* stage of the recursion of the counting algorithm *no* triangulation of P can contain more separators than the total number of simple cycle separators found at that stage, we can express the overcounting factor of the algorithm by the following recurrence,

$$S(\mathcal{P}, \Gamma) \leq \sum_C S(\mathcal{P}_A(C), \Gamma) \cdot S(\mathcal{P}_B(C), \Gamma) \leq 2^{O(\sqrt{n} \log n)} \cdot S(\mathcal{P}_A^*, \Gamma) \cdot S(\mathcal{P}_B^*, \Gamma),$$

where the summation is over all separators C available at the level of recursion, $\mathcal{P}_A(C)$ and $\mathcal{P}_B(C)$ are the sub-problems as explained before, and $\mathcal{P}_A^*, \mathcal{P}_B^*$ are the subproblems corresponding to a separator C^* that maximizes the term $S(\mathcal{P}_A(C), \Gamma) \cdot S(\mathcal{P}_B(C), \Gamma)$ over all C . Recall that whenever the current recursive sub-problem contains $\leq \Gamma$ points we stop the recursion and compute the number of triangulations of the sub-problem exactly, see lines 8 and 9 of Algorithm 1. Hence, we have $S(\mathcal{P}, \Gamma) = 1$ whenever $|\mathcal{P}| \leq \Gamma$. We can now write:

$$S'(\mathcal{P}, \Gamma) := \log(S(\mathcal{P}, \Gamma)) \leq O(\sqrt{n} \log n) + S'(\mathcal{P}_A^*, \Gamma) + S'(\mathcal{P}_B^*, \Gamma).$$

Our goal now is to prove the following lemma:

Lemma 2. *Let \mathcal{P} be a sub-problem of size n and assume $\Gamma = n^{\Omega(1)}$, $n > \Gamma$, and Γ is at least a sufficiently large constant. Then we have*

$$S'(\mathcal{P}, \Gamma) = O\left(\left(\frac{n}{\sqrt{\Gamma/3}} - \sqrt{n}\right) \log \Gamma\right).$$

Proof. We use induction over the size n of \mathcal{P} . We have

$$\begin{aligned} S'(\mathcal{P}, \Gamma) &\leq O(\sqrt{n} \log n) + S'(\mathcal{P}_A^*, \Gamma) + S'(\mathcal{P}_B^*, \Gamma) \\ &\leq O(\sqrt{n} \log n) + c \left(\frac{n_A}{\sqrt{\Gamma/3}} - \sqrt{n_A} + \frac{n_B}{\sqrt{\Gamma/3}} - \sqrt{n_B} \right) \log \Gamma, \end{aligned} \quad (1)$$

where n_A, n_B are the sizes of the sub-problems \mathcal{P}_A^* and \mathcal{P}_B^* of \mathcal{P} , respectively, and c is some sufficiently large positive constant. By the Separator Theorem, we can express $\alpha n \leq n_A \leq \alpha n + \sqrt{4n}$ and $\beta n \leq n_B \leq \beta n + \sqrt{4n}$, for some α and β with (1) $0 < \beta, \alpha \leq \frac{2}{3}$ and (2) $\alpha + \beta = 1$.

Now let us focus on the term $\frac{n_A}{\sqrt{\Gamma/3}} - \sqrt{n_A} + \frac{n_B}{\sqrt{\Gamma/3}} - \sqrt{n_B}$ of equation (1) above:

$$\begin{aligned}
\frac{n_A}{\sqrt{\Gamma/3}} - \sqrt{n_A} + \frac{n_B}{\sqrt{\Gamma/3}} - \sqrt{n_B} &= \frac{n_A + n_B}{\sqrt{\Gamma/3}} - \sqrt{n_A} - \sqrt{n_B} \\
&\leq \frac{\alpha n + \sqrt{4n} + \beta n + \sqrt{4n}}{\sqrt{\Gamma/3}} - \sqrt{n_A} - \sqrt{n_B} \\
&\leq \frac{n + 4\sqrt{n}}{\sqrt{\Gamma/3}} - \sqrt{\alpha n} - \sqrt{\beta n} \\
&= \frac{n + 4\sqrt{n}}{\sqrt{\Gamma/3}} - \sqrt{n}(\sqrt{\alpha} + \sqrt{\beta}) \\
&\leq \frac{n + 4\sqrt{n}}{\sqrt{\Gamma/3}} - \sqrt{n}(1 + \varepsilon)
\end{aligned}$$

The last inequality is obtained by minimizing $\sqrt{\alpha} + \sqrt{\beta}$. Since we mentioned before that $0 \leq \beta \leq \alpha \leq \frac{2}{3}$ and $\alpha + \beta = 1$, the minimum of $\sqrt{\alpha} + \sqrt{\beta}$ is attained at $(\alpha, \beta) = (\frac{2}{3}, \frac{1}{3})$, and is strictly larger than one, so we can choose $\varepsilon > 0$. Now, since Γ is sufficiently large, we have $\frac{4\sqrt{n}}{\sqrt{\Gamma/3}} \ll \varepsilon\sqrt{n}$, so $\frac{4\sqrt{n}}{\sqrt{\Gamma/3}} - \varepsilon\sqrt{n} \leq -\varepsilon'\sqrt{n}$, for some positive constant ε' . Thus we can continue as follows:

$$\frac{n_A}{\sqrt{\Gamma/3}} - \sqrt{n_A} + \frac{n_B}{\sqrt{\Gamma/3}} - \sqrt{n_B} \leq \frac{n + 4\sqrt{n}}{\sqrt{\Gamma/3}} - \sqrt{n}(1 + \varepsilon) \leq \frac{n}{\sqrt{\Gamma/3}} - (1 + \varepsilon')\sqrt{n}. \quad (2)$$

Combining equations (1) and (2) we obtain

$$\begin{aligned}
S'(\mathcal{P}, \Gamma) &\leq O(\sqrt{n} \log n) + c \left(\frac{n}{\sqrt{\Gamma/3}} - (1 + \varepsilon')\sqrt{n} \right) \log \Gamma \\
&\leq c \left(\frac{n}{\sqrt{\Gamma/3}} - \sqrt{n} \right) \log \Gamma + O(\sqrt{n} \log n) - c \cdot \varepsilon' \sqrt{n} \log \Gamma
\end{aligned}$$

If we choose Γ to be sufficiently large, say $\Gamma \geq n^\delta$, for some constant $\delta > 0$, then the negative term $-c \cdot \varepsilon' \sqrt{n} \log \Gamma$ is larger, for sufficiently large c , than the $O(\sqrt{n} \log n)$ term. Hence, we can conclude that $S'(\mathcal{P}, \Gamma) \leq O\left(\left(\frac{n}{\sqrt{\Gamma/3}} - \sqrt{n}\right) \log \Gamma\right)$, which proves the induction step.

It still remains to prove that the inductive claim holds for the boundary condition, so let \mathcal{Q} be a recursive sub-problem of size $\leq \Gamma$. As \mathcal{Q} stems from an application of the Separator Theorem, it is easy to see that $|\mathcal{Q}| \geq \Gamma/3$. Thus, we have $S'(\mathcal{Q}, \Gamma) = 0 \leq c \left(\frac{|\mathcal{Q}|}{\sqrt{\Gamma/3}} - \sqrt{|\mathcal{Q}|} \right) \log \Gamma$. Lemma 2 follows. \blacksquare

Now, let Λ be the number computed by our algorithm. Recall that $|\mathcal{F}_T(P)|$ is the exact number of triangulations of P . By setting $\Gamma = \sqrt{n} \log n$ we obtain an overcounting

factor of the algorithm of

$$S(\mathcal{P}, \Gamma) = 2^{S'(\mathcal{P}, \Gamma)} = 2^{O\left(\frac{n \log \Gamma}{\sqrt{\Gamma}}\right)} = 2^{O\left(n^{\frac{3}{4}} \sqrt{\log n}\right)}.$$

Hence $|\mathcal{F}_T(P)| \leq \Lambda \leq |\mathcal{F}_T(P)| \cdot 2^{O\left(n^{\frac{3}{4}} \sqrt{\log n}\right)} = |\mathcal{F}_T(P)|^{1+o(1)}$. This completes the qualitative part of Theorem 1. It remains to discuss the running time of the algorithm.

4.2 Running time

The running time of the algorithm can be expressed with the following recurrence:

$$T(n) \leq 2^{O(\sqrt{n} \log n)} \cdot T\left(\frac{2n}{3} + \sqrt{4n}\right).$$

Taking again $T'(n) := \log(T(n))$ yields $T'(n) := T'\left(\frac{2n}{3} + \sqrt{4n}\right) + O(\sqrt{n} \log n)$, which can then be solved using the well-known Akra-Bazzi Theorem for recurrences, see [24]. This yields $T'(n) = O(\sqrt{n} \log n)$. There is, however, one detail missing, the stopping condition Γ . In order to use the Akra-Bazzi Theorem we need a boundary condition of $T(n) = 1$ for $1 \leq n \leq n_0$ (for some constant n_0), but in the algorithm we stop the recursion whenever a sub-problem \mathcal{Q} is of size $\leq \Gamma$ (which is dependent on the size n of the original point set). At that point we compute the exact number of triangulations of \mathcal{Q} — say using the algorithm presented in [9]. This gives $T(|\mathcal{Q}|) = 2^{O(|\mathcal{Q}|)} = 2^{O(\Gamma)}$. Hence the exponent in the running time of the algorithm is bounded from above by the solution of $T'(n)$, as given by the Akra-Bazzi Theorem, plus $O(\Gamma)$, i.e., $T(n) = 2^{O(\sqrt{n} \log n + \Gamma)}$. If as before we assume that $\Gamma = \sqrt{n} \log n$ then we end up with $T(n) = 2^{O(\sqrt{n} \log n)} = 2^{o(n)}$, which concludes the proof of Theorem 1 for triangulations.

As a final remark observe that we could have used other values for Γ , rather than $\sqrt{n} \log n$, without violating any argument in the proofs. This yields a tradeoff with running time $2^{O(\Gamma)}$ and approximation ratio $2^{O\left(\frac{n \log \Gamma}{\sqrt{\Gamma}}\right)}$ for any $\sqrt{n} \log n \leq \Gamma \leq n$. Although the quality of the approximation improves with larger Γ , the running time increases. Since we see no way of not having overcounting with this algorithm, we regard $\Gamma = \sqrt{n} \log n$ as the most reasonable setting.

5 Extension to other crossing-free structures

In this section we show how to count crossing-free structures other than triangulations. The idea is to use the framework first developed in [6] that allows to exactly count crossing-free structures using the constrained Delaunay triangulation (CDT) Δ^S constrained to contain the crossing-free structure S . For completeness we will briefly describe the constrained Delaunay triangulation as explained in [6, 7].

The constrained Delaunay triangulation (CDT) Δ^S of P was first introduced in [12]. Formally, it is the triangulation T of P containing S such that no edge e in $T \setminus S$ is flippable in the following sense: Let Δ_1, Δ_2 be triangles of P sharing e . The edge e is flippable if and only if $\square = \Delta_1 \cup \Delta_2$ is convex, and replacing e with the other diagonal of \square increases the smallest angle of the triangulation of \square . One of the most important properties of constrained Delaunay triangulations is its uniqueness if no four points of P are cocircular. Thus, under standard non-degeneracy assumptions, *there is a unique CDT for any given set of mandatory edges*. For a good study on constrained Delaunay triangulations we suggest the book [18] by Ø. Hjelle and M. Dæhlen.

Without loss of generality we will assume that no four points of P are cocircular. This can always be achieved by perturbing P while keeping the order type of P .

Using the constrained Delaunay triangulation, algorithms to exactly count crossing-free matchings and spanning cycles (polygonizations) of P were shown in [6, 7]. The main idea is the following: Assume we want to count all elements $S \in \mathcal{F}_{\mathcal{C}}(P)$ of a certain class \mathcal{C} of crossing-free structures on P . Now, instead of counting all elements of $\mathcal{F}_{\mathcal{C}}(P)$ directly, we count the pairs (Δ^S, S) for all $S \in \mathcal{F}_{\mathcal{C}}(P)$, i.e., we count every element $S \in \mathcal{F}_{\mathcal{C}}(P)$ embedded in its CDT Δ^S . This yields the correct number, since Δ^S is unique w.r.t. S by our non-degeneracy assumption. As we can see the embedding of S in Δ^S as annotating every edge of the triangulation Δ^S by a bit specifying whether that edge belongs to S , we end up counting annotated triangulations. For this we can use similar ideas as in the case of counting triangulations: Having a set \mathcal{S} of separators for triangulations we enumerate all separators $C \in \mathcal{S}$ and use the divide-and-conquer approach to count the crossing-free structures S such that $C \subseteq \Delta^S$. For example, \mathcal{S} could be simple cycle separators, presented in § 2. It is important to observe that when independently solving the smaller sub-problems given by C , we have to have a way to locally verify whether $C \subseteq \Delta^S$, i.e., we have to make sure that choices in a sub-problem do not depend on choices in other sub-problems, as otherwise we might get an overcounting by much more than the number of separators. Recall that we want to keep overcounting under control while achieving sub-exponential runtime. In order to make the previous general idea clear we will adapt the way matchings are counted in [6, 7].

5.1 Counting matchings approximately

Let M be a matching of P , not necessarily perfect. Let Δ^M be the CDT of M . As stated before, we are working under the assumption that no four points of P are cocircular, thus Δ^M is unique w.r.t. M . Now, annotate Δ^M as follows:

- Each vertex v of Δ^M is annotated with a number m_v that indicates the vertex of M that v is matched to. If v is unmatched in M we set m_v to, say, 0.
- Each edge e of Δ^M is annotated with a bit b_e that indicates whether e belongs to M or not.

Let us denote by $\overline{\Delta}^M$ the CDT Δ^M annotated according to the previous rules. Let $C \in \mathcal{S}$ be a separator contained in Δ^M . Separator C inherits all the information, i.e., annotations, from $\overline{\Delta}^M$. The separator thus annotated will be denoted by $C_{\overline{\Delta}^M}$.

5.1.1 The algorithm

An annotated triangulation is said to be *legal* if and only if it is identical to $\overline{\Delta}^M$, for some matching M of P . Now observe that there is a one-to-one correspondence between matchings of P and legal annotated triangulations. Thus, instead of counting matchings directly we may count legal annotated triangulations. To do so we proceed essentially as in the algorithm presented in § 3: (1) For any sub-problem \mathcal{P} we enumerate *all* separators in \mathcal{S} . (2) We enumerate *all* annotations for each separator $C \in \mathcal{S}$. Each such annotated separator splits \mathcal{P} into smaller regions yielding (independent) sub-problems. In each recursive sub-problem we count legal annotated triangulations that are consistent with the annotated separator, i.e., for example, if two adjacent vertices of the separator have been annotated, and they agree to be matched to each other and the edge connecting them is annotated to be in the matching, then in future sub-problems other edges adjacent to those two vertices cannot be annotated to be in a matching as well. (3) We stop the recursion whenever we find a sub-problem of size at most Γ , just as in the algorithm shown in § 3, and compute the *exact* number of legal annotated triangulations that are consistent with the annotations of the boundary of the sub-problem.

It should be clear by now that the only sub-problems that will contribute to the final computed number of matchings are the ones for which the algorithm, in its whole run, could complete a full annotated constrained Delaunay triangulation without finding any violation of the annotations inherited by the separators that led to that triangulation. There are, however, two details we have not yet taken care of: (1) How do we verify that *each* edge that is not annotated to be in a matching M is not flippable in $\overline{\Delta}^M$? As it stands right now, for an annotated separator C dividing \mathcal{P} into sub-problems \mathcal{P}_A and \mathcal{P}_B , we run the algorithm on \mathcal{P}_A and \mathcal{P}_B independently, although an edge e of C could become flippable depending on the two triangles containing e that lie in different sub-problems. (2) How do we *exactly* count legal annotated triangulations that are consistent with the annotations of the boundary of a sub-problem of size at most Γ ? We will tackle each difficulty in turn.

5.1.2 Triangular cycle separators

To overcome difficulty (1) mentioned before, we make use of what we can call “fat” separators, which are very similar to a construction in [6, 7]. Let T be a triangulation of sub-problem \mathcal{P} and let $C \subset T$ be a cycle separator. Every edge of C is contained in two triangles of T (recall that we arbitrarily triangulated the outer face, so indeed every edge is contained in exactly two triangles). For every such edge choose those two triangles. We make C fat by considering the union of all those chosen triangles (as a

set of edges), see Figure 3. We call the resulting fat version of the cycle separator C a *triangular cycle separator* and denote it by C° .




Figure 3: Left — a triangulation and a simple cycle, shown in gray and black respectively. Right — the cycle is made “fat” by attaching the shown triangles to it.

Fat separators allow us to overcome the difficulty regarding edge-flippability as follows. Instead of using simple cycle separators, as we did in § 3, we use triangular cycle separators for counting matchings. Specifically, in the recursive splitting the whole separator C° is part of both sub-problems. Since for every edge e contained in a fat separator C° at least one of its incident triangles is fully contained in C° , we can check flippability of e in one of the sub-problems, namely the one containing the other triangle incident to e . Thus, flippability can be checked independently in each sub-problem.

5.1.3 Counting legal annotated triangulations exactly

Assume that at a certain recursion level we have finally reached a sub-problem \mathcal{Q} of size at most Γ . The boundary of this sub-problem has been (partially) annotated and now we have to exactly count *all* legal annotated triangulations that are consistent with these annotations. In order to achieve this we make use of enumeration algorithms for triangulations, see [10, 11, 22]. For each enumerated triangulation T of \mathcal{Q} we consider *each* subset $E' \subseteq E = E(T)$ of its set of edges E . For each such E' we check whether E' forms a valid matching, is consistent with the annotations of the boundary of \mathcal{Q} , and violates no flippability condition. We count all sets E' satisfying these conditions. It should be clear that using this brute-force approach we are exactly counting *all* legal annotated triangulations that are consistent with the annotations of \mathcal{Q} .

The number of triangulations of \mathcal{Q} can be expressed as c^Γ , for some positive constant c . The number of edges in each triangulation of \mathcal{Q} is $O(\Gamma)$. Thus, enumerating every subset E' of edges takes time $2^{O(\Gamma)}$. Verifying whether the chosen subset of edges forms a consistent matching and checking for flippability of the edges can easily be done within the same time bound. Hence, overall we can solve sub-problems of size Γ in time $2^{O(\Gamma)}$.

5.1.4 Quality of approximation and running time

Having overcome the two difficulties we mentioned previously, we are now left with the analysis of the approximation quality and runtime of the algorithm. Fortunately, essentially all the work has been done in § 4.

From Lemma 1 we obtain that the total number of simple cycle separators is $2^{O(\sqrt{n} \log n)}$. Now, we make each such simple cycle separator fat by attaching at most two triangles to each of its edges. It is not hard to see that the total number of fat simple cycle separators, or triangular cycle separators, is $2^{O(\sqrt{n} \log n)}$, just observe that for every edge of a simple cycle separator we choose two additional points that will complete the attached triangles to make it fat. Next consider the annotations of the triangular cycle separators. It is easy to verify that the number of different annotations for the vertices of a separator is at most $2^{O(\sqrt{n} \log n)}$ and the number of annotations for its edges is at most $2^{O(\sqrt{n})}$. Therefore, the total number of *annotated* triangular cycle separators is $2^{O(\sqrt{n} \log n)}$.

Now, observe that for the analysis in § 4 we essentially only used the following three facts. (1) The number of (annotated) separators of each sub-problem is at most $2^{O(\sqrt{n} \log n)}$. (2) A sub-problem of size Γ can be solved in time $2^{O(\Gamma)}$. (3) All further checks can be performed in polynomial time. All three facts still hold in the case of matchings. Hence, the analysis from § 4 goes through and we obtain the same asymptotic bounds for approximation quality and runtime. This finishes the analysis of the algorithm for counting matchings approximately.

5.2 Counting trees approximately

Having shown how to count matchings approximately using annotations similar to [6, 7], we now show an annotation scheme that allows to count (crossing-free) spanning trees of P approximately. Fixing an arbitrary vertex $p^* \in P$ we root every spanning tree F at p^* by orienting all of its edges towards p^* . This way, every point p in $P \setminus \{p^*\}$ has exactly one outgoing edge, and we call the other end of that edge the *parent* of p in F , denoted by $\text{par}(p)$ (for p^* we set $\text{par}(p^*) := \text{nil}$). Moreover, we denote by $d(p)$ the depth of p in F , i.e., its distance to p^* in F . Note that the resulting oriented spanning trees of P rooted at p^* are in one-to-one correspondence to (ordinary) spanning trees of P , so we may count the former. In the following we will write for short *spanning tree* for oriented spanning trees of P rooted at p^* . We annotate the CDT Δ^F of a spanning tree F of P as follows:

- Every vertex v of F is annotated by the pair $\langle \text{par}(v), d(v) \rangle$.
- Every edge e of Δ^F is annotated with a bit b_e that again represents whether e is part of F or not.

We remark that one can reconstruct the spanning tree already from the data $\text{par}(v)$ or the bits b_e . However, annotations by $\text{par}(v)$ and b_e that locally look like a spanning tree could still contain cycles (that can only be detected globally). The additional annotation by $d(v)$ is used to check acyclicity locally, as explained later.

As for matchings, the annotated CDT Δ^F will be denoted by $\overline{\Delta}^F$. Again there is a one-to-one correspondence between (oriented, rooted at p^*) spanning trees of P and legal annotated triangulations. Therefore, we will again count the latter, just as we did with matchings. In fact, the algorithm is the same as in 5.1.1, the only difference is annotation scheme, which encodes a different class of crossing-free structures.

Nevertheless, in the case of spanning trees two things that are less obvious: (1) How to exactly count all legal annotated triangulations of sub-problems of size at most Γ , and (2) the correctness of the annotation scheme: why do we only (over-)count spanning trees? Let us start with the former.

Assume that we arrive at a sub-problem \mathcal{Q} of size at most Γ . In \mathcal{Q} there are already some annotations present and we have to count all annotated triangulations that are consistent with those annotations. We will proceed as before: We enumerate each triangulation T of \mathcal{Q} and every subset $E' \subseteq E = E(T)$ of edges of T that have no annotation set yet. All edges of E' are chosen to be in the spanning tree. At this point we can check for flippability conflicts. If we find any conflict we abort and consider the next subset of edges, otherwise we continue. Since we are interested in (oriented, rooted at p^*) spanning trees, we enumerate *every* possible orientation of the edges in E' , observe that there are at most $2^{O(|E'|)} = 2^{O(\Gamma)}$ possible orientations we have to go through. This orientation sets the parent of every vertex of the edges of E' . Recall that we want to orient towards the root p^* . Now we test whether this orientation of the edges of E' is consistent with the annotations already present in \mathcal{Q} , e.g., a vertex v that was already annotated in \mathcal{Q} , because it belonged to some separator, must have its parent set. If this parent is in \mathcal{Q} , then v and its parent must be vertices of some edge of E' , otherwise we can abort this run. This in particular tests that the *unique* parent of every vertex of T is set correctly if this parent belongs to \mathcal{Q} as well.

Having chosen E' and having set orientations for its elements, note that we are left with some set of connected components F_1, \dots, F_k formed by the (oriented) edges of T that we have chosen to be in the tree, see Figure 4(a). Now, for each connected component F_i , $1 \leq i \leq k$, let v^i denote an arbitrary vertex that was found fully annotated when first entering \mathcal{Q} , any vertex of the boundary of \mathcal{Q} could be such a vertex since they are all fully annotated. Perform a Breadth-first search (BFS) of F_i starting at v^i as if the edges of F_i had no orientations. Let v be the current vertex the BFS is about to expand its neighborhood of, so in the beginning $v = v^i$, and let d_v be its depth. Set the depth of every vertex w of F_i that was found while expanding the neighborhood of v to $d_v + 1$ or $d_v - 1$ depending on the orientation of the edge $e = vw$, i.e., if e is oriented towards w , then w is the parent of v and thus we set its depth to $d_v - 1$, and if e is oriented towards v we set it to $d_v + 1$. If the depth of w is found to be already set when doing the BFS, then we just verify that the value the algorithm would write there




Figure 4: A sub-problem \mathcal{Q} of size at most Γ is shown in gray. The boundary of \mathcal{Q} is shown simplified, i.e., no triangular cycle separator is shown. The set of edges E' chosen to be in a tree is shown along with an orientation of its elements. For simplicity, the triangulation of \mathcal{Q} that the set E' is chosen from is not shown, so in particular we are assuming that there is no flippability conflict. In (a) there are three connected components F_1, F_2, F_3 , separated by dashed lines for clarity. The gray vertices on the boundary of \mathcal{Q} are fully annotated, so we can start a BFS for each component at any of those vertices. The labels r, s, t represent the annotated depth of those vertices. Starting with the vertex at depth t , every vertex of that component gets the shown depth assigned by the BFS. By the time we arrive at the vertices of depth $t+4$ on the boundary of \mathcal{Q} , since they are fully annotated, we compare $t+4$ with the depth therein annotated and verify for consistency. If they are inconsistent then we abort the run with this set of orientations to the edges of E' . In (b) there are four connected components but the single vertex in the middle will never be annotated since there is no edge in E' that has it as a vertex, so even if we found consistent annotations for the other components, this run is aborted.

and the value already there coincide. If those two numbers do not coincide, then we abort and discard this set of orientations given to the edges of E' , otherwise we continue with the BFS. If the whole connected component F_i was successfully traversed, then we manage to find a consistent set of annotations for it and we can continue with another connected component, if any. If after considering every connected component there is a vertex of T that has not been fully annotated, then we dismiss this run as well, since we were not able to find a consistent set of annotations for E' in \mathcal{Q} , see Figure 4(b), otherwise we declare this run a success. With this method we clearly exactly count *all* legal annotated triangulations that are consistent with the annotations we started with for sub-problem \mathcal{Q} .

It remains to show that, overall, *only* spanning trees of P are (over-)counted.

Lemma 3. *Let P be a set of n points. Then our annotation scheme encodes the spanning trees of P unequivocally and can be checked locally.*

Proof. Consider a legal annotated triangulation \overline{T} of P . The annotations b_e of its edges induce a set of edges F such that \overline{T} is the CDT Δ^F of F (by flippability constraints). Moreover, the edges in F are consistent with the annotations $\langle \text{par}(v), d(v) \rangle$ of the vertices

of \overline{T} , and the latter give F an orientation. Thus, every vertex except for the root p^* has a unique parent, so it has out-degree 1, and F consists of exactly $n - 1$ edges. It remains to show that F has no undirected cycle, which would imply that the annotations induce an (oriented, rooted at p^*) spanning tree of P . For this, first note that any undirected cycle in F would also be a directed cycle (according to the orientation of F given by the parent annotations), since all vertices have out-degree at most 1. Moreover, consistency of the annotated depths $d(v)$ implies that $d(\text{par}(v)) = d(v) - 1$ for all vertices v . Hence, we also cannot have any directed cycle in F , as otherwise for at least one edge of the cycle the depths cannot be consistent. This finishes the proof. ■

Finally, it remains to show what approximation factor and running time we obtain with this algorithm. For the former, the explanation given in § 5.1.4 follows verbatim. For the running time it follows almost verbatim, the only difference is that when choosing a subset of edges E' that will be part of a spanning tree, we have to orient them as well. This orientation comes with a $2^{|E'|} = 2^{O(\Gamma)}$ overhead. That is, instead of having a runtime of $T(\Gamma) = c^\Gamma \cdot 2^{O(\Gamma)} = 2^{O(\Gamma)}$, as we had when counting matchings in § 5.1.4, we now have a runtime of $2^{O(\Gamma)} \cdot T(\Gamma) = 2^{O(\Gamma)}$, which is asymptotically the same. The time required to perform the breath-first searches have no effect in asymptotic terms either. Thus, choosing again $\Gamma = \sqrt{n} \log n$ gives us an overall running time of $T(n) = 2^{O(\sqrt{n} \log n)}$. Hence, both algorithms have the same asymptotic behavior.

6 Conclusions

In this paper we have shown algorithms that given a set of points count triangulations, crossing-free matchings, and crossing-free spanning trees approximately. Both the approximation ratio and the running time are sub-exponential. We would like to remark that using the set of annotations for crossing-free spanning cycles shown in [6] we can modify the algorithms presented in this paper to also count crossing-free spanning cycles approximately, again with sub-exponential approximation ratio and running time. After showing the algorithms for matchings and spanning trees, this extension is straightforward. On the other hand, the annotations shown in this paper to count spanning trees were not known before. They are compatible with the algorithm of [6], and combining the two yields an algorithm to count crossing-free spanning trees *exactly* in time $n^{O(k)}$, where k is the number of onion layers of the given set of points. Specifically, as long as k is fixed we can exactly count spanning trees in polynomial time, which is an interesting result in itself.

Finally, we can express the cardinality of any of the classes of crossing-free structures considered in this paper as c^n , where c depends on the given set of points but lies between two positive constants. Although the approximation ratios shown in this paper are rather large, our algorithms compute a $(1 + o(1))$ -approximation of the base c , and it does so in sub-exponential time. No algorithm with this property was known before. However,

it remains an open problem to find an algorithm with sub-exponential approximation ratio and running time $2^{o(\sqrt{n} \log n)}$, e.g., polynomial.

Acknowledgement We would like to thank the anonymous referees for their valuable suggestions.

References

- [1] O. Aichholzer. The path of a triangulation. In *Proceedings of the 15th annual Symposium on Computational Geometry*, SoCG '99, pages 14–23, New York, NY, USA, 1999. ACM. One citation on page 2.
- [2] 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):67–84, 2007. One citation on page 4.
- [3] O. Aichholzer, F. Hurtado, and M. Noy. A lower bound on the number of triangulations of planar point sets. *Computational Geometry*, 29(2):135–145, 2004. 4 citations on pages 2, 3, and 4.
- [4] O. Aichholzer and H. Krasser. The point set order type data base: A collection of applications and results. In *Proceedings of the 13th Canadian Conference on Computational Geometry*, CCCG '01, pages 17–20, 2001. One citation on page 3.
- [5] O. Aichholzer, G. Rote, B. Speckmann, and I. Streinu. The zigzag path of a pseudo-triangulation. In *Algorithms and Data Structures*, WADS '03, pages 377–388. Springer Berlin Heidelberg, 2003. One citation on page 4.
- [6] V. Alvarez, K. Bringmann, R. Curticapean, and S. Ray. Counting crossing-free structures. In *Proceedings of the 28th annual Symposium on Computational Geometry*, SoCG '12, pages 61–68, New York, NY, USA, 2012. ACM. 10 citations on pages 3, 5, 11, 12, 13, 15, and 18.
- [7] V. Alvarez, K. Bringmann, R. Curticapean, and S. Ray. Counting triangulations and other crossing-free structures via onion layers. *Computing Research Repository (CoRR)*, 2013. Available at <http://arxiv.org/abs/1312.4628>. 6 citations on pages 3, 11, 12, 13, and 15.
- [8] V. Alvarez, K. Bringmann, and S. Ray. A simple sweep line algorithm for counting triangulations and pseudo-triangulations. *Computing Research Repository (CoRR)*, 2013. Available at <http://arxiv.org/abs/1312.3188>. 5 citations on pages 3, 4, and 5.
- [9] V. Alvarez and R. Seidel. A simple aggregative algorithm for counting triangulations of planar point sets and related problems. In *Proceedings of the 29th annual Symposium on Computational Geometry*, SoCG '13, pages 1–8, New York, NY, USA, 2013. ACM. 7 citations on pages 1, 3, 4, 5, and 11.
- [10] D. Avis and K. Fukuda. Reverse search for enumeration. *Discrete Applied Mathematics*, 65(1–3):21–46, 1996. First International Colloquium on Graphs and Optimization. 2 citations on pages 2 and 14.
- [11] S. Bespamyatnikh. An efficient algorithm for enumeration of triangulations. *Computational Geometry*, 23(3):271–279, 2002. 2 citations on pages 2 and 14.

- [12] L. P. Chew. Constrained Delaunay triangulations. In *Proceedings of the 3rd annual Symposium on Computational Geometry*, SoCG '87, pages 215–222, New York, NY, USA, 1987. ACM. One citation on page 12.
- [13] H. N. Djidjev and S. M. Venkatesan. Reduced constants for simple cycle graph separation. *Acta Informatica*, 34(3):231–243, 1997. One citation on page 5.
- [14] A. Dumitrescu, B. Gärtner, S. Pedroni, and E. Welzl. Enumerating triangulation paths. *Computational Geometry*, 20(1–2):3–12, 2001. Selected papers from the 12th Annual Canadian Conference on Computational Geometry. One citation on page 3.
- [15] A. Dumitrescu, A. Schulz, A. Sheffer, and C. D. Tóth. Bounds on the maximum multiplicity of some common geometric graphs. *SIAM Journal on Discrete Mathematics*, 27(2):802–826, 2013. 3 citations on pages 2 and 4.
- [16] P. Flajolet and M. Noy. Analytic combinatorics of non-crossing configurations. *Discrete Mathematics*, 204(1–3):203–229, 1999. Selected papers in honor of Henry W. Gould. One citation on page 4.
- [17] A. Garcia, M. Noy, and J. Tejel. Lower bounds on the number of crossing-free subgraphs of K_n . *Computational Geometry*, 16(4):211–221, 2000. One citation on page 4.
- [18] Ø. Hjelle and M. Dæhlen. *Triangulations and Applications (Mathematics and Visualization)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006. One citation on page 12.
- [19] M. Hoffmann, A. Schulz, M. Sharir, A. Sheffer, C. Tóth, and E. Welzl. Counting plane graphs: Flippability and its applications. In *Thirty Essays on Geometric Graph Theory*, pages 303–325. Springer New York, 2013. One citation on page 4.
- [20] C. Huemer and A. de Mier. Lower bounds on the maximum number of non-crossing acyclic graphs. *arXiv preprint 1310.5882*, 2013. Available at <http://arxiv.org/abs/1310.5882>. One citation on page 4.
- [21] F. Hurtado and M. Noy. Counting triangulations of almost-convex polygons. *Ars Combinatoria*, 45:169–180, 1997. One citation on page 2.
- [22] N. Katoh and S.-I. Tanigawa. Fast enumeration algorithms for non-crossing geometric graphs. *Discrete & Computational Geometry*, 42(3):443–468, 2009. 3 citations on pages 2, 3, and 14.
- [23] C. L. Lawson. Generation of a triangular grid with applications to contour plotting. *Jet Propul. Lab. Techn. Memo*, 299, 1972. One citation on page 2.
- [24] T. Leighton. Notes on better master theorems for divide-and-conquer recurrences. Technical report, Massachusetts Institute of Technology, 1996. Available at <http://courses.csail.mit.edu/6.046/spring04/handouts/akrabazzi.pdf>. Accessed April 4, 2014. One citation on page 11.
- [25] G. L. Miller. Finding small simple cycle separators for 2-connected planar graphs. *Journal of Computer and System Sciences*, 32(3):265–279, 1986. One citation on page 5.
- [26] S. Ray and R. Seidel. A simple and less slow method for counting triangulations and for related problems. In *Proceedings of the 20th European Workshop on Computational Geometry*, EWCG, 2004. One citation on page 3.
- [27] A. Razen and E. Welzl. Counting plane graphs with exponential speed-up. In *Rainbow of Computer Science*, pages 36–46. Springer Berlin Heidelberg, 2011. One citation on page 4.

- [28] F. Santos and R. Seidel. A better upper bound on the number of triangulations of a planar point set. *Journal of Combinatorial Theory, Series A*, 102(1):186–193, 2003. One citation on page 3.
- [29] M. Sharir and A. Sheffer. Counting triangulations of planar point sets. *Electr. J. Comb.*, 18(1):P70, 2011. 3 citations on pages 2 and 4.
- [30] M. Sharir, A. Sheffer, and E. Welzl. On degrees in random triangulations of point sets. *J. Comb. Theory Ser. A*, 118(7):1979–1999, Oct. 2011. 4 citations on pages 1, 2, and 4.
- [31] M. Sharir and E. Welzl. On the number of crossing-free matchings, cycles, and partitions. *SIAM Journal on Computing*, 36(3):695–720, 2006. One citation on page 4.
- [32] A. Sheffer. Numbers of plane graphs. <http://www.cs.tau.ac.il/~sheffera/counting/PlaneGraphs.html>. Accessed April 4, 2014. One citation on page 4.
- [33] R. Sibson. Locally equiangular triangulations. *The Computer Journal*, 21(3):243–245, 1978. One citation on page 2.
- [34] M. Wettstein. Algorithms for counting crossing-free configurations. Master’s thesis, Dept. of Computer Science, ETH Zürich, 2013. 3 citations on pages 4 and 5.
- [35] M. Wettstein. Counting and enumerating crossing-free geometric graphs. In *Proceedings of the 30th annual Symposium on Computational Geometry*, SoCG ’14, 2014. To appear. 3 citations on pages 4 and 5.