ON PROPERTIES OF RANDOM DISSECTIONS
AND TRIANGULATIONS*

NICLA BERNASCONI, KONSTANTINOS PANAGIOTOU†,
ANGELIKA STEGER

Received November 30, 2007

In this work we study properties of random graphs that are drawn uniformly at random from the class consisting of biconnected outerplanar graphs, or equivalently dissections of large convex polygons. We obtain very sharp concentration results for the number of vertices of any given degree, and for the number of induced copies of a given fixed graph. Our method gives similar results for random graphs from the class of triangulations of convex polygons.

1. Introduction & Results

Dissections and triangulations of a convex n-gon are well-studied objects. A dissection is a partition of the polygon into polygonal regions by means of non-crossing diagonals. Triangulations are a special case of dissections, where all regions are triangles. It is a simple and standard exercise in any combinatorics course to obtain that the number of triangulations t_n is equal to the $(n-2)$nd Catalan number, i.e. $t_{n+2} = \frac{1}{n+1}\binom{2n}{n}$. The number of dissections d_n, however, is a much harder object. While an explicit formula involving a sum over products of binomial coefficients belongs to the classical repertoire of advanced combinatorics, see e.g. [1, p. 74], an asymptotic formula was obtained only a few years ago by Flajolet and Noy [3], who showed that $d_n \sim cn^{-3/2}\rho_\mathcal{D}^{-n}$ where $\rho_\mathcal{D} := 3 - 2\sqrt{2} \doteq 0.1716$. Throughout the

Mathematics Subject Classification (2000): 05C80, 05C07

paper, with $\sim$ we will mean "asymptotically equal", in the sense that the ratio of both terms tends to 1 as n grows.

If we advance to the question of properties of a random dissection or triangulation (which is meant to denote a dissection/triangulation drawn uniformly at random from the class of all dissections respectively triangulations of a given convex n-gon), practically nothing seems to be known for dissections. For a random triangulation Gao and Wormald used methods from analytic combinatorics to determine in [6] the limiting distribution of the maximum vertex degree, and obtain quite precise bounds on the number of vertices of degree k in [7]. (Those papers and in addition [8] also study these and related questions for more general types of triangulations.)

The main reason why these questions are so difficult is that contrary to the standard Erdős and Rényi model of random graphs $G_{n,p}$, in random dissections and triangulations (and other graph classes with structural side constraints) the edges are not independent. Therefore most tools from classical random graph theory are not applicable in this context.

In this paper we show that recent progress in the construction of so-called Boltzmann samplers by Duchon, Flajolet, Louchard, and Schaeffer [2] and Fusy [5] can be used to reduce the study of degree sequences and subgraph counts to properties of sequences of independent and identically distributed random variables – to which we can apply standard Chernoff bounds to obtain extremely tight results.

Our results. Let $\mathcal{LD}_n$ denote the class of dissections of labeled convex n-gons, and let LD_n be a graph drawn uniformly at random from $\mathcal{LD}_n$. For a labeled dissection LD we shall denote by $\deg(k; \mathsf{LD})$ the number of vertices in LD with degree k. In our first theorem we determine the asymptotic value of the random variable $\deg(k; \mathsf{LD}_n)$ and provide very tight bounds for the tail probabilities. For brevity we write "$(1 \pm \varepsilon)X$" to denote the interval $((1-\varepsilon)X, (1+\varepsilon)X)$.

Theorem 1.1. *Let* $d_k := (k-1)p^2(1-p)^{k-2}$, *where* $p := 2 - \sqrt{2}$, *and let* $k_0 = k_0(n)$ *be the largest integer such that* $d_{k_0}n > (\log n)^3$. *There is a constant* $C > 0$ *such that for every* $k \le k_0$ *and every* $\frac{(\log n)^2}{\sqrt{d_k n}} < \varepsilon = \varepsilon(n) < 1$ *the following holds for sufficiently large* n:

$$\mathbb{P}\left[\deg(k; \mathsf{LD}_n) \in (1 \pm \varepsilon) \cdot d_k \cdot n\right] \ge 1 - e^{-C\varepsilon^2 \frac{d_k}{k} n}.$$

Furthermore, if $k \in [k_0 + 1, 10 \log n]$, *then*

$$\mathbb{P}\left[\deg(k; \mathsf{LD}_n) < (\log n)^4\right] \ge 1 - kn^{-\log n}.$$

For all remaining k *we have that* $\mathbb{P}[\deg(k; \mathsf{LD}_n) = 0] \to 1$.

From Theorem 1.1 it is easy to derive information about the *maximum vertex degree* $\Delta(\mathsf{LD}_n)$ of a random element from $\mathcal{LD}_n$.

Corollary 1.2. *Let* $p := 2 - \sqrt{2}$, *and set* $b := \frac{1}{1-p}$. *Then*

$$\mathbb{P}\big[\Delta(\mathsf{LD}_n) \notin \big(\log_b n - O(\log\log n), 10\log n\big)\big] = o(1).$$

With our method it seems not possible to improve this result, but we believe that the maximum degree for a random element of $\mathcal{LD}_n$ is given by the lower bound.

Conjecture 1.3. *Let* $p := 2 - \sqrt{2}$, *and set* $b := \frac{1}{1-p}$. *Then*

$$\mathbb{P}\big[\Delta(\mathsf{LD}_n) = \log_b n + \Theta(\log\log n)\big] \geq 1 - o(1).$$

Next we turn to subgraph counts. For an unlabeled dissection H we denote by $\mathsf{copy}\,(H; \mathsf{LD})$ the number of induced copies of H in LD.

Theorem 1.4. *Let* H *be an unlabeled dissection on* n_H *vertices, such that* $n_H = o(\log n)$. *There exist* r_H *(which we define in Section 3.3) and* $c_H := \frac{1}{2} r_H \cdot q^{n_H - 3}$, *where* $q := \frac{2-\sqrt{2}}{2}$, *such that for every* $\frac{\log n}{\sqrt{n8^{-n_H}}} < \varepsilon = \varepsilon(n) < 1$ *and* n *sufficiently large*

$$\mathbb{P}\big[\,\mathsf{copy}\,(H; \mathsf{LD}_n) \in (1 \pm \varepsilon) \cdot c_H \cdot n\,\big] \geq 1 - e^{-\left(\frac{1}{8}\right)^{n_H} \varepsilon^2 n}.$$

For triangulations we obtain similar results, except that some constants change. Let $\mathcal{LT}_n$ denote the class of triangulations of labeled convex n-gons and let LT_n be a graph drawn uniformly at random from $\mathcal{LT}_n$.

Theorem 1.5. *For a random triangulation* LT_n *Theorem 1.1 holds with* LD_n *replaced by* LT_n, *if we let* $p := 1/2$ *instead of* $p := 2 - \sqrt{2}$. *Similarly, if* H *denotes an unlabeled triangulation on* n_H *vertices, then Theorem 1.4 holds for* LT_n *if we let* $p := 1/2$, *and* $q := 1/2$.

Note that for $p = 1/2$ we have $d_k = (k-1)2^{-k}$. For triangulations these values for the expected number of vertices of degree k were already determined by Gao and Wormald in [7]. However, for small k our bounds on the deviation probabilities are much tighter. In particular, note that for constant k our tail bounds are of the form $e^{-\Theta(n)}$, which is comparable to what we are used to from classical random graph theory.

Techniques & Outline. Let $\mathcal{D}$ be the class of *edge-rooted, unlabeled* dissections of convex polygons. The root is an *oriented* edge, such that the outer face is on its right side. It is easy to see that every edge-rooted, unlabeled dissection with $n \geq 3$ vertices gives rise to precisely $\frac{(n-1)!}{2}$ distinct labeled dissections. Hence the degree sequence of a random labeled dissection equals the degree sequence of a random edge-rooted, unlabeled dissection; the same holds obviously for the subgraph count. Formally, we have the following statement.

Theorem 1.6. *For a random edge-rooted unlabeled dissection* D_n *Theorem 1.1 and Theorem 1.4 hold if we replace* LD_n *with* D_n.

The above discussion can easily be adapted for the case of triangulations, where $\mathcal{T}$ is the class of edge-rooted unlabeled triangulations – we omit the obvious details.

The greatest benefit of "switching" from $\mathcal{LD}$ to $\mathcal{D}$ (and similarly from $\mathcal{LT}$ to $\mathcal{T}$) is that the classes $\mathcal{D}$ and $\mathcal{T}$ allow a so-called *decomposition*, which is a *unique description* of the class in terms of general-purpose combinatorial constructions (see Section 3). These constructions appear frequently in modern systematic approaches to asymptotic enumeration and random sampling of several combinatorial structures. It is beyond the scope of this paper to survey these results, and we refer the reader to [4] and references therein for a detailed exposition.

One advantage of the knowledge of the decomposition is that it allows us to develop *mechanically* algorithms that sample objects from the graph class in question by using the framework of *Boltzmann samplers*. This framework was introduced by Duchon et al. in [2], and was extended by Fusy [5] to obtain an (expected) linear time sampler for planar graphs. Our main contribution here is to exploit samplers for dissections and triangulations to reduce in a very general way the problems of determining the degree sequence and counting small subgraphs to properties of *independent* random variables, see Section 3.

We shall now give a short review of the concept of Boltzmann samplers, tailored to our intended application. Let $\mathcal{G}$ be a class of unlabeled graphs, and let $\mathcal{G}_n$ the subset of graphs in $\mathcal{G}$ with n vertices. We will write $g_n := |\mathcal{G}_n|$. Let $G(z) = \sum_{n \geq 0} g_n z^n$ be the *ordinary generating function* of $\mathcal{G}$. In the Boltzmann model of parameter x, we assign to any object $\gamma \in \mathcal{G}$ the probability

$$(1.1) \qquad\qquad \mathbb{P}_x[\gamma] = \frac{x^{|\gamma|}}{G(x)},$$

if the expression above is well-defined. A *Boltzmann sampler* $\Gamma G(x)$ for $\mathcal{G}$ is an algorithm that generates graphs from $\mathcal{G}$ according to (1.1). In [2] several general procedures, which translate common combinatorial construction rules like union, set, etc. into Boltzmann samplers are given. Notice that the probability above only depends on the choice of x and on the size of γ, such that every object of the same size has the *same* probability of being generated. This means that if we condition on the output being of a certain size n, then the Boltzmann sampler is a *uniform* sampler of the class $\mathcal{G}_n$. The parameter x "tunes" the expected size of the output, and the larger we make it, the larger the expected size of a random object from $\mathcal{G}$ becomes.

The remainder of the paper is structured as follows. The next section presents a few facts about Boltzmann samplers that we are going to exploit later. In Section 3 we first introduce the concept of *pre*degree and *post*degree, which are central to our work, and then construct and analyze a specific Boltzmann sampler for dissections. This will then allow us to prove Theorem 1.1. In Sections 3.3 and 4 we sketch how to modify this proof in order to obtain Theorems 1.4 and 1.5.

2. Properties of the Boltzmann Model

In this section we state and prove some basic facts which we are going to apply several times. We start by observing that as long as the probability (1.1) is not too small for large $|\gamma|$, we can construct by *rejection* an efficient sampler that always outputs an object of the desired size with high probability. In fact, suppose that for a class of graphs $\mathcal{G}$ the corresponding ordinary generating function $G(z)$ becomes singular at $\rho_{\mathcal{G}}$, and that the value $G(\rho_{\mathcal{G}})$ is finite. Moreover, suppose that there exist constants $c, \alpha > 0$ such that $|\mathcal{G}_n| \sim cn^{-\alpha}\rho_{\mathcal{G}}^{-n}$. Then we obtain with (1.1) that the probability for the whole graph class $\mathcal{G}_n$ (that is, the sum of the probabilities of all the elements on n vertices) is $\mathbb{P}_x[\mathcal{G}_n] \sim \frac{1}{G(x)}cn^{-\alpha}\left(\frac{x}{\rho_{\mathcal{G}}}\right)^n$. Note that if we set $x < \rho_{\mathcal{G}}$, the probability of drawing a graph of size n is exponentially small, while if $x = \rho_{\mathcal{G}}$, then this probability is proportional to a polynomial. Hence, from now on we will concentrate on the case $x = \rho_G$.

Suppose that we are given an algorithm $\Gamma G(\rho_{\mathcal{G}})$ that generates graphs according to distribution (1.1) with $x = \rho_{\mathcal{G}}$. With this notation, consider the following simple algorithm.

$\widetilde{\Gamma}G(n):$ **for** $\left(i = 1, \ldots, n^{\alpha + 3/2}\right)$ **let** $\gamma_i \leftarrow \Gamma G(\rho_{\mathcal{G}})$
                 $I := \{j \mid 1 \leq j \leq n^{\alpha + 3/2} \text{ and } |\gamma_j| = n\}$
                 if $I = \emptyset$ **return** $\perp$
                 else return $\gamma_{\min\{j \mid j \in I\}}$

Proposition 2.1. *Let $\mathcal{G}$ be a class of graphs and $c, \alpha > 0$ constants such that $\mathbb{P}_{\rho_{\mathcal{G}}}[\mathcal{G}_n] \sim cn^{-\alpha}$. For sufficiently large n $\mathbb{P}[\widetilde{\Gamma}G(n) = \bot] \leq e^{-n}$. Furthermore, $\mathbb{P}[\widetilde{\Gamma}G(n) = \gamma \mid \widetilde{\Gamma}G(n) \neq \bot] = |\mathcal{G}_n|^{-1}$ for any $\gamma \in \mathcal{G}_n$.*

Proof. Since the calls to $\Gamma G(\rho_{\mathcal{G}})$ are independent we may estimate with much room to spare

$$\mathbb{P}[\widetilde{\Gamma}G(n) = \bot] \leq \left(1 - (1 - o(1))cn^{-\alpha}\right)^{n^{\alpha+3/2}} \leq \left(e^{-\frac{c}{2}n^{-\alpha}}\right)^{n^{\alpha+3/2}} \leq e^{-n}.$$

The second statement follows immediately from the fact that $\mathbb{P}[\Gamma G(\rho_{\mathcal{G}}) = \gamma_1] = \mathbb{P}[\Gamma G(\rho_{\mathcal{G}}) = \gamma_2]$ for all $\gamma_1, \gamma_2 \in \mathcal{G}_n$. ∎

In other words, $\widetilde{\Gamma}G(n)$ will succeed with very high probability, and is a uniform sampler for graphs from $\mathcal{G}_n$. The next proposition states that if we can prove that the probability that a random object generated by $\widetilde{\Gamma}G(n)$ has a property with probability at least p, then a graph drawn uniformly at random from $\mathcal{G}_n$ has that property with at least the same probability. The proof is a straightforward application of the proposition above.

Proposition 2.2. *Let $\mathcal{G}$ be a class of graphs and $\widetilde{\Gamma}G$ be its corresponding sampler as above. Furthermore, let $\mathcal{P} \subset \mathcal{G}$ and suppose that there is a $p > 0$ such that $\mathbb{P}[\widetilde{\Gamma}G(n) \in \mathcal{P}] \geq p$. Denote by G_n a graph drawn uniformly at random from $\mathcal{G}_n$. Then $\mathbb{P}[\mathsf{G}_n \in \mathcal{P}] \geq p$.*

Proof. Note that for every $\gamma \in \mathcal{G}_n$ we have $\mathbb{P}[\widetilde{\Gamma}G(n) = \gamma \mid \widetilde{\Gamma}G(n) \neq \bot] = |\mathcal{G}_n|^{-1}$. Hence

$$\begin{aligned} \mathbb{P}[\mathsf{G}_n \in \mathcal{P}] &= \mathbb{P}[\widetilde{\Gamma}G(n) \in \mathcal{P} \mid \widetilde{\Gamma}G(n) \neq \bot] \\ &\geq \mathbb{P}[\widetilde{\Gamma}G(n) \in \mathcal{P} \text{ and } \widetilde{\Gamma}G(n) \neq \bot] = \mathbb{P}[\widetilde{\Gamma}G(n) \in \mathcal{P}]. \quad ∎ \end{aligned}$$

3. Dissections of Convex Polygons

In the remainder of the paper we shall denote with slight abuse of notation by a "dissection" an edge-rooted, unlabeled dissection. Recall that the root of a dissection is an oriented edge, such that the outer face is on its right side. A dissection is then either a single edge, or an (ordered) sequence of $i \geq 2$ dissections along the face containing the root edge, where $i-1$ pairs of vertices are glued together, see Figure 1.

This yields that the ordinary generating function $D(z)$ enumerating dissections satisfies (see also [3] for a more general treatment)
(3.1)
$$D(z) = z^2 + \frac{D(z)^2}{z} + \frac{D(z)^3}{z^2} + \cdots + \frac{D(z)^i}{z^{i-1}} + \cdots = \frac{z}{4}\left(1 + z - \sqrt{z^2 - 6z + 1}\right).$$

Figure 1. Decomposition of rooted dissections of convex polygons.

It is easy to see that D becomes singular at $\rho_D := 3 - 2\sqrt{2} \doteq 0.1716$, and sophisticated tools from analytic combinatorics (see [3], Theorem 4) yield that there is a constant $c > 0$ such that $d_n := |\mathcal{D}_n| \sim cn^{-3/2}\rho_D^{-n}$.

Before we proceed we shall make a few important definitions. For an edge-rooted dissection D we will denote by the *root vertex* of D the tail of the root edge of D, and by *end vertex* the head of the root edge. Furthermore, we shall denote by $\mathsf{rdeg}\,(D)$ the degree of the root vertex of D.

For a dissection D, let us now consider its internal dual graph T_D (see Figure 2 for an illustration), in which we add a vertex for every face of D except for the outer one, and connect two vertices if and only if the corresponding faces share an edge in D. The dual graph is a rooted tree, where the root is the vertex corresponding to the face containing the root edge of D. For a vertex v of D, let $f_v \in V(T_D)$ be that vertex in the dual graph that corresponds to a face that contains v, and among all such vertices is that vertex which is closest to the root of T_D. We call f_v the *characteristic vertex* of v in T_D. Let $e_1(v)$ and $e_2(v)$ be the two edges on the outer face of D that are incident to v, where $e_1(v)$ is the first that is encountered when transversing the outer face of D in the direction of its root edge. Observe that the face corresponding to f_v partitions the edges incident to v in two parts, one containing $e_1(v)$, the other containing $e_2(v)$. We define the *predegree* $\mathsf{pred}\,(v; D)$ of v as the number of edges in the part containing $e_1(v)$, while the *postdegree* $\mathsf{postd}\,(v; D)$ of v is defined as the number of edges in


**Figure 2.** A dissection and its corresponding internal dual tree.

Figure 3. The first cycle of a dissection, with dissections "sitting" on its edges.

the part containing $e_2(v)$. We will simply write $\mathsf{pred}\,(v)$ and $\mathsf{postd}\,(v)$ if it is clear of which dissection we are talking about. Finally, let $\mathsf{pred}\,(k; D)$ be the number of vertices with predegree k in D, and define $\mathsf{postd}\,(k; D)$ similarly.

Notice that the sum of predegree and postdegree of a vertex is exactly equal to the degree of the vertex. As we will see in the following subsections, these two parameters will be much easier to handle than counting directly vertices with a fixed degree in random dissections.

Before we present the details of our analysis we shall collect some properties of the predegree and the postdegree of a vertex in a dissection. Recall that a dissection D is either a single oriented edge, or it consists of a cycle $C_D := \{\{v_1, v_2\}, \{v_2, v_3\}, \ldots, \{v_\ell, v_1\}\}$ of length ℓ, which is rooted at the edge (v_1, v_ℓ), and the edge $\{v_i, v_{i+1}\}$ is replaced by an edge-rooted dissection D_i – see Figure 3. Having this, we can easily prove the following lemma.

Lemma 3.1. *Let $D \in \mathcal{D}$ be a dissection, and let $v_1, \ldots, v_\ell$ and $D_1, \ldots, D_{\ell-1}$ be as defined above. Then the following holds for all $1 \le i < \ell$. If $v \in D_i$ and $v \notin \{v_i, v_{i+1}\}$, then*

$$\mathsf{pred}\,(v; D) = \mathsf{pred}\,(v; D_i) \quad \text{and} \quad \mathsf{postd}\,(v; D) = \mathsf{postd}\,(v; D_i).$$

Furthermore, if $v = v_i$, i.e. v is the root of a dissection D_i, then

$$\mathsf{pred}\,(v_i; D) = \deg\,(v_i; D_i) = \mathsf{rdeg}\,(D_i).$$

Proof. The statement is clearly true if D is a single edge. If D has at least 3 vertices, then the characteristic face of a vertex v in D_i such that $v \notin \{v_i, v_{i+1}\}$ is a face of D_i. Similarly, the edges $e_1(v)$ and $e_2(v)$ needed for the definition of predegree and postdegree are edges of the border of D_i. Hence, the predegree and postdegree of v depend only on D_i, which proves the first statement. For the second one, observe that the characteristic vertex f_v of v corresponds to the cycle C_D. Its predegree is therefore precisely the degree of the root vertex of D_i. ∎

Next we define an operation on the class $\mathcal{D}$. Given a dissection D, this operation first reflects D at an axis perpendicular to the root edge (such that we obtain the same dissection, but the root edge will now have the outer face on its left side, i.e., the resulting graph is not an element of $\mathcal{D}$), and then invert the direction of the root edge (so we obtain a graph in $\mathcal{D}$). We will call this operation *reflection-rotation*, and denote it by $\mathrm{rr}(D)$. The following lemma is an immediate consequence of the definition of rr – we state it without proof.

Lemma 3.2. *For every element of $\mathcal{D}$ and for every vertex $v \in D$*

i) $\mathsf{pred}\,(v;\,\mathrm{rr}(D)) = \mathsf{postd}\,(v;\,D)$,

ii) $\mathsf{postd}\,(v;\,\mathrm{rr}(D)) = \mathsf{pred}\,(v;\,D)$.

Furthermore, $\mathrm{rr}(\cdot)$ is a 1-to-1 mapping from $\mathcal{D}_n$ to $\mathcal{D}_n$, and $\mathrm{rr}(\mathrm{rr}(D)) = D$.

The remainder of this section is structured as follows. In Section 3.1 we will design a sampler for dissections of convex polygons, and prove some fundamental properties of it. Section 3.2 deals with the analysis of an execution of this sampler, which eventually yields the degree sequence of random dissections. Finally, in Section 3.3 we demonstrate how we can obtain tight estimates for the number of (small) subgraphs.

3.1. A Sampler for Dissections

According to the decomposition of the class of dissections (Figure 1) and the translation rules in [2], a Boltzmann sampler for $\mathcal{D}$ starts with a cycle of a certain length (given by an appropriate probability distribution), and then substitutes every edge distinct from the root edge with another randomly generated dissection. More formally, define the *cycle distribution* $\mathrm{Cyc}(x)$ with parameter x by

$$(3.2) \qquad c_\ell(x) := \mathbb{P}\,[\,\mathrm{Cyc}(x) = \ell\,] = \begin{cases} \frac{x^2}{D(x)} & \text{if } \ell = 2, \\ \left(\frac{D(x)}{x}\right)^{\ell-2} & \text{if } \ell > 2, \\ 0 & \text{otherwise.} \end{cases}$$

Then the sampler $\Gamma D(x)$ for $\mathcal{D}$ is given by the following algorithm.

$\Gamma D(x):$ $\ell \leftarrow \mathrm{Cyc}(x)$
 if $(\ell = 2)$ **return** a single edge
 else
  $\gamma \leftarrow \ell$-cycle $\{\{v_1, v_2\}, \ldots, \{v_{\ell-1}, v_\ell\}, \{v_\ell, v_1\}\}$
  for $(i = 1 \ldots \ell - 1)$
   $\gamma_i \leftarrow \Gamma D(x)$
   $\gamma \leftarrow$ identify (v_i, v_{i+1}) with the root of γ_i
  return γ, rooted at (v_1, v_ℓ)

The next lemma follows directly from the compilation rules in [2].

Lemma 3.3. $\mathbb{P}[\Gamma D(x) = D] = \frac{x^{|V(D)|}}{D(x)}$ *for $D \in \mathcal{D}$.*

With the above result we obtain straightforwardly an asymptotic estimate for the probability that $\Gamma D(x)$ outputs an object of a given size.

Lemma 3.4. *Let* $0 < x \leq \rho_{\mathcal{D}}$. *There is a constant* $C = C(x) > 0$ *such that*

$$\mathbb{P}\left[|\Gamma D(x)| = n\right] \sim Cn^{-3/2} \left(\frac{x}{\rho_{\mathcal{D}}}\right)^n.$$

Proof. By applying Lemma 3.3 and using the estimate $|\mathcal{D}_n| \sim cn^{-3/2}\rho_{\mathcal{D}}^{-n}$ we obtain

$$\mathbb{P}\left[|\Gamma D(x)| = n\right] = \frac{|\mathcal{D}_n| \cdot x^n}{D(x)} \sim \frac{cn^{-3/2}\rho_{\mathcal{D}}^{-n}x^n}{D(x)} = Cn^{-3/2} \left(\frac{x}{\rho_{\mathcal{D}}}\right)^n. \qquad \blacksquare$$

Observe that if we choose $x = \rho_{\mathcal{D}}$, then the lemma above states that the probability that $\Gamma D(x)$ outputs an object of size n is proportional to $n^{-3/2}$. Thus, Proposition 2.1 guarantees the existence of an exact-size sampler $\widetilde{\Gamma D}(n)$, which performs n^3 independent calls to $\Gamma D(\rho_{\mathcal{D}})$, and has probability of success at least $1 - e^{-n}$, i.e., it returns an empty graph $\perp$ with probability at most e^{-n}.

The next lemma summarizes some key properties of $\Gamma D(x)$ that we shall exploit later.

Lemma 3.5. *Let* D *be the output of* $\Gamma D(x)$, *and let* r *be the root vertex of* D. *Then there is a unique sequence of values* $\alpha_1, \ldots, \alpha_N$ *such that the following is true:*

i) $\alpha_N = 2$ *and* $\Gamma D(x)$ *outputs* D *if and only if the random values it used in its execution were precisely this sequence.*
ii) $N = |E(D)|$.
iii) $|V(D)| = \left|\{i \mid 1 \leq i \leq N \text{ and } \alpha_i = 2\}\right| + 1$.
iv) $\deg(r; D) = \mathsf{rdeg}(D) = \min\limits_{k=1,\ldots,N}\{k \mid \alpha_k = 2\}$.
v) *For every* $k \geq 1$ *the number of vertices different from the two endpoints of the root edge in* D *that have predegree* k *is equal to the quantity*

$$\left|\{1 \leq i \leq N - k \mid \alpha_i = 2, \text{ and } \alpha_{i+k} = 2, \text{ and } \alpha_{i+j} \neq 2 \text{ for } 1 \leq j < k\}\right|.$$

Proof. Before we prove the statements, let us make an important observation. Suppose that when $\Gamma D(x)$ finishes its execution, we draw the resulting dissection on the plane such that its root (i.e., the edge (v_1, v_ℓ) generated in the very first call of $\Gamma D(x)$) is as demonstrated in Figure 1, thus it is oriented from left to right. Then $\Gamma D(x)$ has the property that it first generated the cycle that contains the root edge (v_1, v_ℓ), and then generated the other parts of the dissection in *clockwise* order: it successively created $\ell - 1$ dissections $D_1, \ldots, D_{\ell-1}$, starting from the one which was attached to the

edge (v_1, v_2), then the one which was attached to (v_2, v_3), and so on, always proceeding in a clockwise manner – cf. Figure 3. Note that the same also applies for all recursive calls to $\Gamma D(x)$. With this observation we shall prove the statements *i)–v)* with induction over the structure of D.

It is straightforward to see that if D is a single edge, then *i)* is trivially true. On the other hand, if D is composed out of the dissections $D_1, \ldots, D_{\ell-1}$, then by the induction hypothesis there are $\ell - 1$ sequences $A_1, \ldots, A_{\ell-1}$ of values, such that $\Gamma D(x)$ outputs D_i if and only if the random values that it used in its execution were precisely A_i. Since $\Gamma D(x)$ can output D if and only if it first had generated a cycle of length ℓ, and then generated the dissections $D_1, \ldots, D_{\ell-1}$ in exactly this order, only the sequence of random values $(\ell, A_1, \ldots, A_{\ell-1})$ will cause $\Gamma D(x)$ to output D. Furthermore, the last used variable α_N must be equal to 2, since otherwise $\Gamma D(x)$ would have initiated at least one more recursive call (and hence α_N would not have been the end of the list). This proves *i)*.

For the remainder of the proof we fix the following notation. We will assume that $\Gamma D(x)$ outputs D_j if and only if it used the sequence of random values $A_j = (\alpha_{j,1}, \ldots, \alpha_{j,N_j})$, and we will write $(\alpha_1, \ldots, \alpha_N) \equiv (\ell, A_1, \ldots, A_{\ell-1})$, where $\ell \geq 3$.

To see *ii)*, note that the statement is trivial if D is just an edge. Moreover, if D has more than two vertices, then by applying the induction hypothesis we may assume that $|E(D_j)| = N_j$ for $j = 1, \ldots, \ell-1$. But then we can derive

$$|E(D)| = 1 + \sum_{j=1}^{\ell-1} |E(D_j)| = 1 + \sum_{j=1}^{\ell-1} N_j = |(\ell, A_1, \ldots, A_{\ell-1})| = N,$$

which is precisely *ii)*.

The third statement of the lemma is easily seen to be true for dissections with two vertices. If D is composed, then by the induction hypothesis we have for all $1 \leq j \leq \ell-1$ that

$$|V(D_j)| = \left|\{i \mid 1 \leq i \leq N_j \text{ and } \alpha_{j,i} = 2\}\right| + 1.$$

But then we may deduce that

$$|V(D)| = \sum_{j=1}^{\ell-1} |V(D_j)| - (\ell - 2) \quad \text{(as the vertices } v_2, \ldots, v_{\ell-1}$$
$$\text{are counted in the sum twice)}$$
$$= \sum_{j=1}^{\ell-1} \left[\left|\{i \mid 1 \leq i \leq N_j \text{ and } \alpha_{j,i} = 2\}\right| + 1 \right] - (\ell - 2)$$
$$= \left|\{i \mid 1 \leq i \leq N \text{ and } \alpha_i = 2\}\right| + 1 \quad \text{(as } \alpha_1 = \ell \geq 3\text{)}.$$

In order to see *iv)*, let us assume first that D is an oriented edge. Then it is easy to see that $\Gamma D(x)$ will output D if and only if the first random value evaluates to 2 (and no other calls to $\Gamma D(x)$ will be initiated). But then $\alpha_1 = 2$, and the statement is true as the degree of the root of D is precisely 1. On the other hand, if D is composed out of the dissections $D_1, \ldots, D_{\ell-1}$, then we have $\deg(r; D) = 1 + \deg(r; D_1)$. Now note that r is also the root vertex of D_1, as $\Gamma D(x)$ identifies the edge $(v_1, v_2) \equiv (r, v_2)$ with the root edge of D_1 when constructing D. Hence, we may apply the induction hypothesis, and we obtain

$$\deg(r; D) = 1 + \deg(r; D_1) = 1 + \min_{1 \le i \le N_1} \{i \mid \alpha_{1,i} = 2\}.$$

But the last quantity equals $\min_{1 \le i \le N} \{i \mid \alpha_i = 2\}$, as we have $\alpha_1 = \ell \ge 3$ and due to *i)* there is at least one element in the sequence A_1 that equals 2. This proves *iv)*.

Finally, to prove *v)*, we denote a vertex different from the endpoints of the root edge of a dissection as an *inner* vertex. Observe that *v)* is true for dissections which consist only of a single edge, as we then do not have any inner vertex. In every other case, due to the induction hypothesis we have that for all $1 \le j \le \ell - 1$, the number of inner vertices of D_j that have predegree k is

$$\left| \{1 \le i \le N_j - k \mid \alpha_{j,i} = 2, \alpha_{j,i+k} = 2, \text{ and } \alpha_{j,i+x} \ne 2 \text{ for } 1 \le x < k\} \right|.$$

In order to count the number of inner vertices with predegree k in γ we apply Lemma 3.1. Then we have that this number equals the number of inner vertices of the dissections $D_1, \ldots, D_{\ell-1}$ with that predegree, plus the number of vertices among $v_2, \ldots, v_{\ell-1}$, that have predegree k. But according to the second statement of Lemma 3.1, the latter number equals the number of dissections among $D_2, \ldots, D_{\ell-1}$, that have *root degree* equal to k.

Now let d_j be the minimal index in A_j, such that $\alpha_{j,d_j} = 2$. Due to *iv)* we know that $\mathsf{rdeg}(D_j) = d_j$, and hence the quantity $s_k := \left| \{2 \le j \le \ell - 1 \mid d_j = k\} \right|$ counts the number of root vertices of $D_2, \ldots, D_{\ell-1}$ that have degree k. Thus, from the previous discussion we may deduce that the number of inner vertices of D with predegree k is precisely

$$\sum_{j=1}^{\ell-1} \left| \{1 \le i \le N_j - k \mid \alpha_{j,i} = 2, \alpha_{j,i+k} = 2, \alpha_{j,i+x} \ne 2 \text{ for } 1 \le x < k\} \right| + s_k =$$

$$\sum_{j=1}^{\ell-1} \left| \{d_j \le i \le N_j - k \mid \alpha_{j,i} = 2, \alpha_{j,i+k} = 2, \alpha_{j,i+x} \ne 2 \text{ for } 1 \le x < k\} \right| + s_k.$$

To complete the proof, note that we know from *i)* that for all j we have $\alpha_{j,N_j} = 2$. The above quantity is thus equal to

$$\left|\left\{1 \leq i \leq N - k \mid \alpha_i = 2, \alpha_{i+k} = 2, \alpha_{i+x} \neq 2 \text{ for } 1 \leq x < k\right\}\right|,$$

as $\alpha_1 = \ell > 2$. ∎

Lemma 3.5 allows us to identify every dissection D with a sequence $(\alpha_1, \ldots, \alpha_N)$ that generates it. Moreover, the probability that $\Gamma D(x)$ outputs D is exactly equal to the probability that an (infinite) sequence $(\alpha_1, \alpha_2, \ldots)$, in which the α_i's are drawn independently according to (3.2), starts with $(\alpha_1, \ldots, \alpha_N)$. In order to study properties of a random dissection we will thus proceed in two steps:

– establish a correspondence between properties of a dissection and properties of such sequences,
– bound the probability that a random dissection has a specific property in terms of the probability that a sequence drawn according to (3.2) will have the corresponding property.

Note that the statements *ii)–v)* from Lemma 3.5 already contribute to the first part. We shall close this section with an additional lemma that will help us to further establish a correspondence between properties of dissections and properties of sequences of values. In order to formulate it we need some notation. From Lemma 3.5 we know that the α_i's that are equal to 2 play a special rôle: they determine the number of vertices in the dissection. In order to emphasize this we subdivide every sequence $(\alpha_1, \ldots, \alpha_N)$ into blocks as follows. A *block* of size ℓ is a subsequence $(\alpha_{i+1}, \ldots, \alpha_{i+\ell})$ such that $\alpha_i = \alpha_{i+\ell} = 2$ and $\alpha_{i+k} > 2$ for all $k = 1, \ldots, \ell - 1$. (For notational convenience we assume that $\alpha_0 \equiv 2$ so that the sequence starts with a block; the value α_0 is never used by the sampler.)

In our final lemma of this section we establish how predegree and postdegree of a given vertex in D change during the execution of the sampler. To formulate the lemma we fix some more notation. Observe that $\Gamma D(x)$ uses the values in the sequence $A = (\alpha_1, \alpha_2, \ldots)$ *sequentially* in order to generate a dissection. We denote by $\widehat{D}_i$ the part of the final dissection that $\Gamma D(x)$ has generated at the iteration when everything up to and including the $(i-1)$st block has been read. In particular, $\widehat{D}_1$ is just a single edge. Furthermore, denote by $u_1, \ldots, u_{|V(D)|}$ the vertices of D, where u_1 is the root vertex, and u_i is the neighbor of u_{i-1} in *clockwise* direction around the border of D. Finally, let $e_i = \{w_i, w_i'\}$ be the edge on the border of $\widehat{D}_i$ which is going to be identified with the root edge of the dissection that will be generated in the current (recursive) call to $\Gamma D(x)$, and let $e_1 \equiv \widehat{D}_1$.

Lemma 3.6. *With the notation above, we have that for all $1 \leq i < |V(D)|$ the following two properties hold:*

i) $w_i = u_i$, *and* $\mathsf{pred}\,(u_i; D) =$ *size of the i-th block read by $\Gamma D(x)$.*
ii) $\mathsf{postd}\,(u_i; D) = \mathsf{postd}\,(u_i; \widehat{D}_{i+j})$, *for all $0 \leq j \leq |V(D)| - i$.*

Proof. The proof of the lemma is by induction over the structure of D. If D is a single oriented edge the statement is obtained straightforwardly. Otherwise, D is a sequence of dissections $D_1, \ldots, D_{\ell-1}$ along the face containing the root edge, where the root vertex of D_i is identified with the end vertex of D_{i-1}, and the root vertex of D_1 is joined by an edge with the end vertex of $D_{\ell-1}$ (cf. Figure 3). Observe that there exists indices $1 = i_1 < i_2 < \cdots < i_{\ell-1}$ such that u_{i_j} is the root vertex of D_j.

Recall that the sampler will start constructing the dissection D_j after having completed the generation of D_{j-1}, as it serves the recursive calls in clockwise order around the root edge. In particular, therefore we have that $\widehat{D}_{i_j}$ is the dissection consisting of $D_1, \ldots, D_{j-1}$, and the cycle surrounding the root face.

With this in mind, we see that the claims of the lemma follow for vertices *not* contained in the root face by applying the induction hypothesis and Lemma 3.1. It thus remains to show the claims for vertices u_{i_j}. *i)* follows immediately from the observations above and Lemmas 3.1 and 3.5. To see *ii)* observe that

$$\mathsf{postd}\,(u_{i_j}; \widehat{D}_{i_j}) = \text{degree of the end vertex of } D_{j-1} \text{ (in } D_{j-1})$$
$$= \mathsf{postd}\,(u_{i_j}; D),$$

due to Lemma 3.1 and Lemma 3.2 and note that the postdegree will not change any more, as the generation of D_{j-1} was completed. ∎

3.2. The Degree Sequence of Random Dissections

In this section we shall prove Theorem 1.1. We start by quoting a simple version of the standard Chernoff bounds (see for example [9]).

Lemma 3.7 (Chernoff Inequalities). *Let $X \sim \mathrm{Bin}(n, p)$ and $\mu := \mathbb{E}\,[X] = np$. Then for every $0 < \varepsilon < 1$*

$$(3.3) \qquad\qquad \mathbb{P}\,[X \in (1 \pm \varepsilon)\mu] \geq 1 - 2e^{-\frac{\varepsilon^2}{3}\mu}.$$

First we shall prove a lemma that is similar in spirit to Theorem 1.1, but considers just the predegrees instead of the full degrees of the vertices. Observe that statement *v)* of Lemma 3.5 provides a close connection between the number of vertices of predegree k and the number of blocks of length k in the input sequence $(\alpha_1, \alpha_2, \ldots)$ of the sampler $\Gamma D(x)$. Observe also that the fact that the α_i's are sampled independently according to the distribution given in (3.2) implies that the length of a block is geometrically distributed. Now, by applying Proposition 2.1 we obtain that there is an algorithm that returns a dissection of size exactly n (or nothing), which exploits $\Gamma D(x)$. By keeping track of the probabilities, both observations together will allow us to make precise statements regarding the number of vertices with a given predegree in random dissections of size n.

Lemma 3.8. *Let $p_\ell := c_2(\rho_\mathcal{D})(1 - c_2(\rho_\mathcal{D}))^{\ell-1}$, where $c_2(\rho_\mathcal{D})$ is defined in (3.2), and let $\ell_0 = \ell_0(n)$ be the largest integer such that $p_{\ell_0} n > \frac{1}{9}(\log n)^2$. Furthermore, let d_k be defined as in Theorem 1.1. There is a constant $C > 0$ such that for every $\frac{(\log n)^2}{\sqrt{p_\ell n}} < \varepsilon = \varepsilon(n) < 1$ the following holds for sufficiently large n. If $\ell \le \ell_0$*

$$\mathbb{P}\left[\,\mathsf{pred}\,(\ell; \mathsf{D}_n) \in (1 \pm \varepsilon) \cdot p_\ell \cdot n\,\right] \ge 1 - e^{-C\varepsilon^2 p_\ell n}.$$

Furthermore, if $\ell \in [\ell_0 + 1, 5\log n]$, then

$$\mathbb{P}\left[\,\mathsf{pred}\,(\ell; \mathsf{D}_n) < (\log n)^2\,\right] \ge 1 - n^{-\log n}.$$

Finally, if $\ell > 5\log n$, then $\mathbb{P}[\mathsf{pred}\,(\ell; \mathsf{D}_n) = 0] \to 1$.

Proof. Let

(3.4) $$L = (\alpha_1, \alpha_2, \ldots)$$

be an (infinite) sequence of random variables which were drawn independently according to the distribution (3.2), where we set $x = \rho_D$. Furthermore, let $\Gamma D(L)$ be the sampler $\Gamma D(x)$ for the special case $x = \rho_\mathcal{D}$, with the following modification: $\Gamma D(L)$ proceeds in exactly the same way as $\Gamma D(\rho_\mathcal{D})$, but every time that $\Gamma D(\rho_\mathcal{D})$ would make a random choice (i.e., draw a random variable), $\Gamma D(L)$ reads the next unused entry from L. Due to Lemma 3.3, the algorithm $\Gamma D(L)$ is a Boltzmann sampler for $\mathcal{D}$, i.e. $\mathbb{P}[\Gamma D(L) = \gamma] = \frac{\rho_\mathcal{D}^{|\gamma|}}{D(\rho_\mathcal{D})}$, where the probability is taken over the random choices in L.

With this in mind, by applying Lemma 3.4 and Proposition 2.1 we obtain that there is a sampling algorithm $\widetilde{\Gamma} D(n)$ which performs n^3 independent calls to $\Gamma D(L)$ such that for each of these n^3 calls it uses a new list L_i as defined in (3.4).

In the sequel we define two events $\mathfrak{A}$ and $\mathfrak{B}$ such that their intersection implies the event "$\mathsf{pred}(\ell; \widetilde{\Gamma}D(n)) \in (1 \pm \varepsilon)p_\ell n$" if $\ell \leq \ell_0$. (We discuss at the end of the proof how we have to modify the event $\mathfrak{B}$ in order to obtain the statements for $\ell > \ell_0$.)

($\mathfrak{A}$) $\widetilde{\Gamma}D(n) \neq \perp$.

($\mathfrak{B}$) For $j \geq 1$, let $N_{i,j}$ be the position of the j-th occurrence of the value "2" in the list L_i, and for $j \geq 2$ let $\Delta_{i,j} = N_{i,j} - N_{i,j-1}$. Furthermore, let $X_{i,j}$ be the indicator variable for the event "$\Delta_{i,j} = \ell$". Then $X_i := \sum_{j=2}^{n-1} X_{i,j} \in (1 \pm \varepsilon/2)p_\ell n$ for all $1 \leq i \leq n^3$.

Now assume that $\mathfrak{A}$ and $\mathfrak{B}$ occur simultaneously. Clearly, due to $\mathfrak{A}$ there is an index i_0 such that $\Gamma D(L_{i_0})$ outputs a dissection with precisely n vertices. But then, due to Lemma 3.5 (statements $iii)$ and $v)$), the number of vertices in the output of $\Gamma D(L_{i_0})$ that are different from the endpoints of its root edge and that have predegree ℓ is equal to

$$\left| \left\{ 1 \leq i \leq N_{i_0, n-1} - \ell \mid \alpha_i = 2, \alpha_{i+\ell} = 2, \alpha_{i+x} \neq 2 \text{ for } 1 \leq x < \ell \right\} \right|.$$

But this quantity equals precisely the variable X_{i_0} defined in the event $\mathfrak{B}$. Hence, $\mathfrak{B}$ implies that the number of vertices with predegree ℓ in the output of $\widetilde{\Gamma}D(n)$ is in the interval

$$\left[(1 - \varepsilon/2)p_\ell n, (1 + \varepsilon/2)p_\ell n + 2 \right] \subseteq (1 \pm \varepsilon)p_\ell n,$$

whenever n is large enough.

In the sequel we will show that $\mathbb{P}[\mathfrak{A}] \geq 1 - e^{-n}$ and $\mathbb{P}[\mathfrak{B}] \geq 1 - e^{-C'\varepsilon^2 p_\ell n}$, for a suitably chosen constant $C' > 0$. Then the proof of the first part of the lemma follows easily:

$$\mathbb{P}\left[\mathsf{pred}(\ell; \mathsf{D}_n) \in (1 \pm \varepsilon) \cdot p_\ell \cdot n \right] = \mathbb{P}\left[\mathsf{pred}(\ell; \widetilde{\Gamma}D(n)) \in (1 \pm \varepsilon) \cdot p_\ell \cdot n \mid \mathfrak{A} \right]$$

$$\geq \mathbb{P}[\mathfrak{B} \mid \mathfrak{A}] = \frac{\mathbb{P}[\mathfrak{A} \text{ and } \mathfrak{B}]}{\mathbb{P}[\mathfrak{A}]} \geq 1 - e^{-n} - e^{-C'\varepsilon^2 p_\ell n} \geq 1 - e^{-C\varepsilon^2 p_\ell n},$$

for an appropriately chosen constant $C > 0$.

The fact that $\mathbb{P}[\mathfrak{A}] \geq 1 - e^{-n}$ follows immediately from Proposition 2.1. To obtain the lower bound for the probability of $\mathfrak{B}$ we first consider an arbitrary but fixed index $1 \leq i \leq n^3$. Observe that the fact that the α_i's are independent implies that the values $\Delta_{i,2}, \ldots, \Delta_{i,n-1}$ are independent. Furthermore note that $\mathbb{P}[\Delta_{i,j} = \ell] = c_2(\rho_\mathcal{D})(1 - c_2(\rho_\mathcal{D}))^{\ell-1} = p_\ell$, as this is the probability that there are $\ell - 1$ consecutive values in L_i that are greater than 2, followed by a "2". So X_i is distributed as $\mathrm{Bin}(n-2, p_\ell)$, and with Lemma 3.7 we obtain that there is a constant $C'' > 0$ such that

$$\mathbb{P}[X_i \notin (1 \pm \varepsilon/2)p_\ell n] \leq e^{-C''\varepsilon^2 p_\ell n}.$$

Let $\overline{\mathfrak{B}}$ denote the complement of event $\mathfrak{B}$. We estimate $\mathbb{P}[\overline{\mathfrak{B}}]$ by the union bound, and exploit the lower bound for ε:

$$\mathbb{P}[\overline{\mathfrak{B}}] = \mathbb{P}[\exists i\colon X_i \notin (1\pm\varepsilon/2)p_\ell n] \le n^3\cdot e^{-C''\varepsilon^2 p_\ell n} \le e^{-C'\varepsilon^2 p_\ell n}.$$

This concludes the proof for $\ell \le \ell_0$. Now let $\ell \in [\ell_0+1, 5\log n]$, and define the event $\mathfrak{B}'$ to denote the event $\mathfrak{B}$ with the only difference that we require $X_i < (\log n)^2 - 1$. As above, it is easy to argue that $\mathfrak{A}$ and $\mathfrak{B}'$ imply the event "$\mathsf{pred}(\ell; \widetilde{\varGamma}D(n)) < (\log n)^2$". (Here we used that the predegree of the endvertex of a dissection is always one. The minus one term in the definition of event $\mathfrak{B}'$ thus suffices to take care of both root vertices.) To complete the proof for such ℓ, recall that X_i is distributed like $\mathrm{Bin}(n-2, p_\ell)$ and note that due to $p_\ell n \le \frac{1}{9}(\log n)^2$ we may estimate with $\binom{a}{b} \le (\frac{ea}{b})^b$

$$\mathbb{P}[\overline{\mathfrak{B}'}] = \mathbb{P}\left[\exists i\colon X_i \ge (\log n)^2 - 1\right] \le n^3 \binom{n-2}{(\log n)^2-1} p_\ell^{(\log n)^2-1}$$

$$\le n^3 \left(\frac{enp_\ell}{(\log n)^2-1}\right)^{(\log n)^2-1} \le n^{-\log n},$$

whenever n is sufficiently large. If $\ell > 5\log n$ we observe that $c_2(\rho_\mathcal{D}) \doteq 0.58579$ and hence $p_\ell n = o(n^{-3})$. We define the event $\mathfrak{B}''$ as above, with the difference that we require $X_i = 0$ *and* $\mathsf{rdeg}(\varGamma D(L_i)) \ne \ell+1$. We easily see that $\mathfrak{A}$ and $\mathfrak{B}''$ imply the event "$\mathsf{pred}(\ell; \widetilde{\varGamma}D(n)) = 0$", since the condition $\mathsf{rdeg}(\varGamma D(L_i)) \ne \ell+1$ implies that the predegree of the root vertex of $\varGamma D(L_i)$ is different from ℓ. Finally, to bound $\mathbb{P}[\mathfrak{B}'']$ we observe that due to *iv)* of Lemma 3.5 we have

$$\mathbb{P}[\overline{\mathfrak{B}''}] = \mathbb{P}[\exists i\colon X_i > 0] + \mathbb{P}[\exists i\colon \mathsf{rdeg}(\varGamma D_{\le n}(L_i)) = \ell+1]$$

$$\le n^3(\mathbb{E}[X_i] + p_{\ell+1}) \overset{(\mathbb{E}[X_i]\le np_\ell)}{=} o(1). \qquad \blacksquare$$

Recall that in Section 3 we introduced a reflection-rotation operation rr that is a bijective mapping from $\mathcal{D}_n$ to $\mathcal{D}_n$, and maps a dissection D to a dissection $D' = \mathrm{rr}(D)$ such that $\mathsf{pred}(\ell; D) = \mathsf{postd}(\ell; D')$, cf. Lemma 3.2. As $\mathrm{rr}(\mathrm{rr}(D)) = D$ this implies in addition that the number of dissections in $\mathcal{D}_n$ with a specified number of vertices with postdegree ℓ is equal to the number of dissections with the same number of vertices with predegree ℓ.

Lemma 3.9. *Let $\mathsf{D_n}$ be drawn uniformly at random from $\mathcal{D}_n$. For all ℓ and k we have*

$$\mathbb{P}[\mathsf{pred}(\ell; \mathsf{D_n}) = k] = \mathbb{P}[\mathsf{postd}(\ell; \mathsf{D_n}) = k].$$

With the above facts at hand we can now prove the theorem on the degree distribution of a random dissection.

Proof of Theorem 1.1. We first consider the case $k \leq k_0$; at the end of the proof we will point out the modifications that have to be made to obtain the result for all other k as well. Let $1 \leq \ell \leq k-1$ and denote for a dissection D by $\mathsf{PRED}(\ell; D) := \{v \in V(D) \mid \mathsf{pred}(v; D) = \ell\}$ the set of vertices with predegree ℓ in D, and define similarly $\mathsf{POSTD}(\ell; D)$. Our proof strategy is to show that with probability at least $1 - e^{-\tilde{C}\varepsilon^2 \frac{d_k}{k} n}$, for some $\tilde{C} > 0$,

$$(3.5) \qquad \left| \mathsf{PRED}(k-\ell; \mathsf{D_n}) \cap \mathsf{POSTD}(\ell; \mathsf{D_n}) \right| \in (1 \pm \varepsilon) p_\ell p_{k-\ell} n,$$

where p_ℓ (as in Lemma 3.8) is defined as $c_2(\rho_\mathcal{D})(1 - c_2(\rho_\mathcal{D}))^{\ell-1}$, with $c_2(\rho) = \frac{\rho_\mathcal{D}^2}{D(\rho_\mathcal{D})} = 2 - \sqrt{2}$, namely $c_2(\rho_\mathcal{D})$ is the p in the statement of Theorem 1.1. Then

$$p_\ell p_{k-\ell} = c_2(\rho_\mathcal{D})(1 - c_2(\rho_\mathcal{D}))^{\ell-1} c_2(\rho_\mathcal{D})(1 - c_2(\rho_\mathcal{D}))^{k-\ell-1}$$
$$= c_2(\rho_\mathcal{D})^2 (1 - c_2(\rho_\mathcal{D}))^{k-2} = \frac{d_k}{k-1},$$

with d_k as defined in the statement of the theorem. Therefore, if we prove (3.5), and since the degree of a vertex in a dissection equals the sum of its predegree and postdegree, summing up the above expression for $1 \leq \ell \leq k-1$ yields for sufficiently large n the claim of the theorem.

From now on let $1 \leq \ell \leq k-1$ be fixed and let

$$(3.6) \qquad L^{=\ell} = (\alpha_1, \alpha_2, \dots) \quad \text{and} \quad L^{\neq \ell} = (\beta_1, \beta_2, \dots)$$

be two (infinite) sequences of variables drawn independently according to the distribution (3.2), where we set $x = \rho_D$. Furthermore, let $\Gamma D(L^{=\ell}, L^{\neq\ell})$ be the sampler $\Gamma D(x)$ for the special case $x = \rho_\mathcal{D}$, with the following modification: $\Gamma D(L^{=\ell}, L^{\neq\ell})$ proceeds in exactly the same way as $\Gamma D(\rho_\mathcal{D})$, but every time that $\Gamma D(\rho_\mathcal{D})$ would make a random choice (i.e., draw a random variable), $\Gamma D(L^{=\ell}, L^{\neq\ell})$ reads the next unused value from one of the lists $L^{=\ell}$ or $L^{\neq\ell}$. The choice of the list is made according to the following rules:

($R1$) We decide on the list from which to read the values once at the beginning and then *only after* a "2" that was read, i.e., after the end of a block was reached in the current list.

($R2$) In the beginning the sampler draws from $L^{=\ell}$ if $\ell = 1$, and otherwise from $L^{\neq\ell}$.

($R3$) After a "2" was read we proceed as follows. Let $\widehat{D}_i$ be the (partial) dissection generated by $\Gamma D(L^{=\ell}, L^{\neq\ell})$ after having processed in total precisely $i-1$ blocks (such that some of them were read from $L^{=\ell}$, and the remaining ones from $L^{\neq\ell}$). Moreover, let $e_i = \{w_i, w_i'\}$ denote the edge

on the border of $\widehat{D}_i$ which is going to be identified with the root edge of the dissection that will be generated in the current recursive call to the sampler, such that w_i is identified with the root vertex of that dissection. Then $\Gamma D(L^{=\ell}, L^{\neq \ell})$ will draw the variables of the next block from $L^{=\ell}$ if w_i has postdegree precisely ℓ in $\widehat{D}_i$, and otherwise from $L^{\neq \ell}$.

As $\Gamma D(L^{=\ell}, L^{\neq \ell})$ simply imitates the execution of $\Gamma D(\rho_\mathcal{D})$, it is due to Lemma 3.3 a Boltzmann sampler for $\mathcal{D}$, i.e.

$$(3.7) \qquad \mathbb{P}\big[\Gamma D(L^{=\ell}, L^{\neq \ell}) = \gamma\big] = \frac{\rho_\mathcal{D}^{|\gamma|}}{D(\rho_\mathcal{D})},$$

where the probability is taken over the random choices in $L^{=\ell}$ and $L^{\neq \ell}$.

With this in mind, by applying Lemma 3.4 and Proposition 2.1 we obtain that there is an algorithm $\widetilde{\Gamma} D(n)$ which performs n^3 independent calls to $\Gamma D(L^{=\ell}, L^{\neq \ell})$ such that for each call a new pair of lists $(L_i^{=\ell}, L_i^{\neq \ell})$ is used, where all lists are as defined in (3.6).

Before we proceed let us discuss a fundamental property of $\Gamma D(L^{=\ell}, L^{\neq \ell})$. Recall that the sampler accesses the two lists blockwise, i.e., it reads a complete block from a list, before it decides with rule $(R3)$ to possibly change its source of random values (recall that a block is a sequence of values ending at a two). Furthermore, note that due to Lemma 3.5, *iii)*, $\Gamma D(L^{=\ell}, L^{\neq \ell})$ needs exactly $n-1$ blocks to generate a dissection D with n vertices. Hence there is a value t such that $\Gamma D(L^{=\ell}, L^{\neq \ell})$ used the first t blocks from $L^{=\ell}$, and the remaining $n-1-t$ blocks from $L^{\neq \ell}$. Denote by t_m the number of blocks of size m among the t blocks. The fundamental property of $\Gamma D(L^{=\ell}, L^{\neq \ell})$ that we will exploit is the following.

$$(3.8) \qquad \begin{array}{c} t_m \text{ equals the number of vertices different from the end vertex} \\ \text{with postdegree } \ell \text{ and predegree } m \text{ in } D. \end{array}$$

To prove that (3.8) holds we shall apply Lemma 3.6. Recall that $\Gamma D(L^{=\ell}, L^{\neq \ell})$ imitates the execution of $\Gamma D(\rho_D)$, except for the fact that it reads the needed random variables from the two lists $L^{=\ell}$ and $L^{\neq \ell}$. In particular, the sequence $(\widehat{D}_i)_{1 \leq i \leq n}$ of partial dissections generated by $\Gamma D(L^{=\ell}, L^{\neq \ell})$ is the *same* as the sequence that is generated by $\Gamma D(\rho_D)$, if we assume that the output of both samplers is the same graph. Now, by applying Lemma 3.6 we observe the following. Consider any vertex u_j of D, where the notation is such that u_1 is the root vertex, and u_j is the neighbor of u_{j-1} in clockwise direction around the border of D. If the postdegree of u_j in D is ℓ, then it was due to *ii)* also ℓ in $\widehat{D}_j, \ldots, \widehat{D}_{|V(D)|} = D$. That is, after the $(j-1)$th block was read by the sampler, the postdegree of u_j did

not change any more. Moreover, due to *i)* we have that j is the only index x such that u_j is the first element of e_x. Hence, when the sampler reads the j-th block from one of the lists, we have that $e_j = \{u_j, v'_j\}$, for some other vertex v'_j in $\widehat{D}_j$, and according to the previous discussion, the postdegree of u_j will not change any more during the generation of D. By putting it all together we see that the number of blocks of size m read from $L^{=\ell}$ equals due to rule $(R3)$ the number of vertices in D with postdegree ℓ and predegree m.

We now define three events $\mathfrak{A}$, $\mathfrak{B}$, and $\mathfrak{C}$, which will imply (3.5).

($\mathfrak{A}$) $\widetilde{\Gamma}D(n) \neq \bot$.

($\mathfrak{B}$) $\mathsf{postd}\big(\ell; \widetilde{\Gamma}D(n)\big) \in (1 \pm \varepsilon/3)p_\ell n$, where $p_\ell = c_2(\rho_\mathcal{D})(1 - c_2(\rho_\mathcal{D}))^{\ell-1}$.

($\mathfrak{C}$) Define $N_{i,0} := 0$, and for $j \geq 1$ let $N_{i,j}$ be the position of the j-th occurrence of the value "2" in the list $L_i^{=\ell}$, and let $\Delta_{i,j} := N_{i,j} - N_{i,j-1}$. Furthermore, let $X_{i,j}$ be the indicator variable for the event "$\Delta_{i,j} = k - \ell$". Then for all $1 \leq i \leq n^3$ and all $(1 - \varepsilon/3)p_\ell n \leq n' \leq n$

$$X_i^{(n')} := \sum_{j=1}^{n'} X_{i,j} \in (1 \pm \varepsilon/3)p_{k-\ell}n'.$$

Now assume that the events $\mathfrak{B}$ and $\mathfrak{C}$ occur simultaneously. Clearly, $\mathfrak{B}$ implies $\mathfrak{A}$ and we know that thus there exists an i_0 such that $\Gamma D\big(L_{i_0}^{=\ell}, L_{i_0}^{\neq\ell}\big)$ outputs a dissection with precisely n vertices, which has $\mathsf{postd}\big(\ell; \Gamma D\big(L_{i_0}^{=\ell}, L_{i_0}^{\neq\ell}\big)\big) \in (1 \pm \varepsilon/3)p_\ell n$ vertices with postdegree ℓ. Then, due to property (3.8), the number of vertices different from the end vertex in the output of $\Gamma D\big(L_{i_0}^{=\ell}, L_{i_0}^{\neq\ell}\big)$ with predegree $k - \ell$ and postdegree ℓ equals the number of blocks of size $k - \ell$ among the first $\mathsf{postd}\big(\ell; \Gamma D\big(L_{i_0}^{=\ell}, L_{i_0}^{\neq\ell}\big)\big)$ blocks of $L_{i_0}^{=\ell}$. By applying event $\mathfrak{C}$ with $i = i_0$ and $n' = \mathsf{postd}\big(\ell; \Gamma D\big(L_{i_0}^{=\ell}, L_{i_0}^{\neq\ell}\big)\big) \in (1 \pm \varepsilon/3)p_\ell n$ we obtain that this number is in the interval

$$\left(1 \pm \frac{\varepsilon}{3}\right) \cdot p_{k-\ell} \cdot n' \subseteq \left(1 \pm \frac{\varepsilon}{3}\right)^2 p_{k-\ell} \cdot p_\ell \cdot n \stackrel{(p_\ell p_{k-\ell} = \frac{d_k}{k-1})}{\subseteq} \left(1 \pm \frac{7\varepsilon}{9}\right)\frac{d_k}{k-1}n.$$

From this we deduce that the total number of vertices with predegree $k - \ell$ and postdegree ℓ in the output of $\widetilde{\Gamma}D(n)$ is contained in

$$\left[\left(1 - \frac{7\varepsilon}{9}\right)\frac{d_k}{k-1}n, \left(1 + \frac{7\varepsilon}{9}\right)\frac{d_k}{k-1}n + 1\right] \subseteq (1 \pm \varepsilon)\frac{d_k}{k-1}n.$$

From Proposition 2.1 we know that $\mathbb{P}[\mathfrak{A}] \geq 1 - e^{-n}$. In the following we will show that $\mathbb{P}[\mathfrak{B} \,|\, \mathfrak{A}] \geq 1 - e^{-C'\varepsilon^2 p_\ell n}$ and $\mathbb{P}[\mathfrak{C}] \geq 1 - e^{-C''\varepsilon^2 d_k n}$, for suitably

chosen constants $C', C'' > 0$. Then the proof of the first part of the theorem follows immediately:

$$\mathbb{P}\left[(3.5)\right] = \mathbb{P}\left[\left|\mathsf{PRED}\left(k-\ell; \widetilde{\Gamma}D(n)\right) \cap \mathsf{POSTD}\left(\ell; \widetilde{\Gamma}D(n)\right)\right| \in (1\pm\varepsilon)\frac{d_k}{k-1}n \;\middle|\; \mathfrak{A}\right]$$

$$\geq \mathbb{P}\left[\mathfrak{B} \wedge \mathfrak{C} \mid \mathfrak{A}\right] \geq \mathbb{P}\left[\mathfrak{B} \mid \mathfrak{A}\right] - \mathbb{P}\left[\overline{\mathfrak{C}} \mid \mathfrak{A}\right] \geq \mathbb{P}\left[\mathfrak{B} \mid \mathfrak{A}\right] - \frac{\mathbb{P}\left[\overline{\mathfrak{C}}\right]}{\mathbb{P}\left[\mathfrak{A}\right]}$$

$$\geq 1 - e^{-C'\varepsilon^2 p_\ell n} - \frac{e^{-C''\varepsilon^2 \frac{d_k}{k}n}}{1 - e^{-n}} \geq 1 - e^{-C\varepsilon^2 \frac{d_k}{k}n},$$

for an appropriately chosen constant $C > 0$. To show the existence of a constant $C' > 0$ such that $\mathbb{P}\left[\mathfrak{B} \mid \mathfrak{A}\right] \geq 1 - e^{-C'\varepsilon^2 p_\ell n}$ we recall Lemma 3.9 and Lemma 3.8 and obtain

$$\begin{aligned} \mathbb{P}\left[\mathfrak{B} \mid \mathfrak{A}\right] &= \mathbb{P}\left[\mathsf{postd}\,(\ell; \mathsf{D_n}) \in \left(1 \pm \frac{\varepsilon}{3}\right) p_\ell n\right] \\ (3.9) \qquad & \overset{\text{(Lemma 3.9)}}{=} \mathbb{P}\left[\mathsf{pred}\,(\ell; \mathsf{D_n}) \in \left(1 \pm \frac{\varepsilon}{3}\right) p_\ell n\right] \overset{\text{(Lemma 3.8)}}{\geq} 1 - e^{-C'\varepsilon^2 p_\ell n}. \end{aligned}$$

To bound the probability of $\overline{\mathfrak{C}}$ we proceed precisely as in the proof of event $\mathfrak{B}$ in Lemma 3.8 – we observe that the values $\Delta_{i,1}, \ldots, \Delta_{i,n'}$ are independent, and that $\mathbb{P}[\Delta_{i,j} = k - \ell] = p_{k-\ell}$. Therefore $X_i^{(n')}$ is $\mathrm{Bin}(n', p_{k-\ell})$-distributed, and Lemma 3.7, the union bound and the lower bound for ε yield the desired upper bound for $\mathbb{P}[\overline{\mathfrak{C}}]$.

Now let us turn our attention to the case $k \in [k_0 + 1, 10 \log n]$. Here we cannot work directly with (3.5), as due to Lemma 3.8 it depends on the choice of ℓ whether the number of vertices with predegree ℓ is around $p_\ell n$ or is at most $(\log n)^2$. In the first case we can exploit the event $\mathfrak{B}$ from above, while in the latter we use the following modification of it.

$(\mathfrak{B}'_\ell)$ $\mathsf{postd}\left(\ell; \widetilde{\Gamma}D(n)\right) < (\log n)^2.$

Furthermore, we shall modify $\mathfrak{C}$ for every admissible ℓ as follows.

$(\mathfrak{C}'_\ell)$ $X_i^{(n')} < 20(\log n)^2$ for all $n' < (1+\varepsilon)p_\ell n + (\log n)^2 =: n'_0.$

Now, if ℓ is such that $p_\ell n < \frac{1}{9}(\log n)^2$, then with the same argument as in (3.9) we can show that $\mathbb{P}[\mathfrak{B}_\ell \mid \mathfrak{A}] \geq 1 - n^{-\log n}$. But then $\mathfrak{B}'_\ell$ implies with Property (3.8) that with probability at most $n^{-\log n}$, the number of vertices with postdegree ℓ and predegree $k-\ell$ is larger than $(\log n)^2$. Furthermore, if ℓ is such that $p_\ell n \geq \frac{1}{9}(\log n)^2$, then $\mathfrak{C}'_\ell$ implies that there are at most $20(\log n)^2$ vertices different from the end vertex with postdegree ℓ and predegree $k-\ell$. Summing over all admissible ℓ yields then that the total number of vertices with degree k is at most $k \cdot (20(\log n)^2 + 1) \leq (\log n)^4.$

It remains to show that $\mathbb{P}[\mathfrak{C}'_\ell] \geq 1 - n^{-\log n}$. To see this, observe that $X_i^{(n')}$ is distributed like $\mathrm{Bin}(n', p_{k-\ell})$. Observe that from $d_k n < (\log n)^3$ we can deduce that $k \geq \frac{1}{2} \log n$ and therefore $\frac{d_k n}{k} < 2(\log n)^2$. Hence

$$\mathbb{P}[\overline{\mathfrak{C}'_\ell}] = \mathbb{P}[\exists i, n' : X_i^{(n')} > 20(\log n)^2] \leq n^3 \sum_{n' \leq n'_0} \binom{n'}{20(\log n)^2} p_{k-\ell}^{20(\log n)^2}$$

$$\leq n^4 \binom{n'_0}{20(\log n)^2} p_{k-\ell}^{20(\log n)^2} \leq n^4 \binom{2p_\ell n + (\log n)^2}{20(\log n)^2} p_{k-\ell}^{20(\log n)^2}$$

$$\leq n^4 \left(\frac{e(2p_\ell n + (\log n)^2)p_{k-\ell}}{20(\log n)^2} \right)^{20(\log n)^2}$$

$$= n^4 \left(\frac{e}{20} \left(\frac{2p_\ell p_{k-\ell} n}{(\log n)^2} + p_{k-\ell} \right) \right)^{20(\log n)^2}$$

$$\overset{(p_\ell p_{k-\ell} = \frac{d_k}{k-1})}{\leq} n^4 \left(\frac{e}{20} \left(\frac{2 d_k n}{k(\log n)^2} + 1 \right) \right)^{20(\log n)^2} \leq \left(\frac{e}{4} \right)^{20(\log n)^2} \leq n^{-\log n}.$$

Finally, if $k > 10 \log n$, then observe that there is a vertex in the output of $\widetilde{\varGamma D}(n)$ such that either its predegree or its postdegree is larger than $5 \log n$. But this event does not occur with high probability, due to Lemmas 3.8 and 3.2. ∎

3.3. Number of Edges and of Small Subgraphs

In this section we shall prove Theorem 1.4. As previously, we will restrict our considerations to edge-rooted unlabeled dissections, as every edge-rooted dissection gives rise to precisely $\frac{(n-1)!}{2}$ labeled dissections.

We start by stating the following lemma that is a special case of Theorem 1.4. In fact we count the number of edges in a random dissection, and this will actually correspond to the induction basis in the proof of Theorem 1.4.

Lemma 3.10. *Let $e(G)$ denote the number of edges of a graph G and define $c_e := \frac{1}{2 - \sqrt{2}}$. Then for every $\frac{\log n}{\sqrt{n}} < \varepsilon = \varepsilon(n) < 1$ and sufficiently large n*

$$\mathbb{P}[e(\mathsf{D}_n) \in (1 \pm \varepsilon)c_e n] \geq 1 - e^{-\frac{\varepsilon^2}{2}n}.$$

Proof. We will prove that the number of edges of a dissection built by the sampler $\widetilde{\varGamma D}(n)$ belongs to $(1 \pm \varepsilon)c_e n$ with at least the claimed probability (recall that $\widetilde{\varGamma D}(n)$ is the sampler that calls $\varGamma D(\rho_D)$ n^3 times, and returns

a dissection on n vertices with probability at least $1-e^{-n}$). Then the lemma follows by applying Proposition 2.2.

By applying Lemma 3.5, *ii)* and *iii)*, we see that the number of edges in the output of $\widetilde{\Gamma}D(n)$ is equal to the number of random values used by the sampler during the run that returned this object, and this is precisely the number of values until the $(n-1)$st occurrence of the value "2" in this list. The position $P\left(i,\frac{1}{c_e}\right)$ of the ith "2" is distributed as the sum of i independent geometric variables with parameter $\frac{\rho_D^2}{D(\rho_D)}=\frac{1}{c_e}$ (the latter equality follows from (3.1) and the fact that $\rho_D=3-2\sqrt{2}$), thus we expect that the $(n-1)$st occurrence of a "2" will be at position $c_e(n-1)$. To bound the probability that the sum of independent geometric variables assumes a value far from its expected value, we make use of the following simple observation.

$$\mathbb{P}\left[P\left(i,\frac{1}{c_e}\right)>m\right]=\mathbb{P}\left[\mathrm{Bin}\left(m,\frac{1}{c_e}\right)<i\right]\quad\text{and}$$

$$\mathbb{P}\left[P\left(i,\frac{1}{c_e}\right)\leq m\right]=\mathbb{P}\left[\mathrm{Bin}\left(m,\frac{1}{c_e}\right)\geq i\right].$$

If we consider n^3 (infinite) lists of random values chosen according to (3.2), then by applying Lemma 3.7 we obtain that with probability at least $1-2n^3e^{-c_e\frac{\varepsilon^2}{3}n}$, in *every* list the $(n-1)$st occurrence of a "2" will be at a position in the interval $(1\pm\varepsilon)c_en$. Since the probability that $\widetilde{\Gamma}D(n)$ will output a dissection on n vertices is at least $1-e^{-n}$, we obtain with room to spare that for large n that with probability at least $1-e^{-\varepsilon^2n/2}$ the number of edges in this dissection will be in the interval $(1\pm\varepsilon)c_en$. ∎

To prove Theorem 1.4 we introduce some new notation. First, let S be an induced biconnected subgraph of a dissection D. We now describe how we uniquely *root* this subgraph. In the internal dual tree of D, the vertices corresponding to the faces of S form a subtree. We root S at the edge of D crossed by the last edge of the unique path connecting the root vertex of T_D to this subtree. Let us call this edge e_S, and orient it in such a way that the outer face of S is on the right side of e_S.

Now let H be a fixed unlabeled biconnected outerplanar graph. Let $\mathcal{H}$ be the set of edge-rooted dissections that can be obtained from H by rooting an edge that is contained in the external face – see Figure 5 for an illustration – and define r_H of Theorem 1.4 to be $|\mathcal{H}|$. Now let $\widehat{H}\in\mathcal{H}$. We say that an induced copy of H in D is a *rooted copy* of $\widehat{H}$ in D, if the dissection obtained by rooting that copy in the way we defined above is isomorphic to $\widehat{H}$. Note that every copy of H in D corresponds in this way to a rooted copy of exactly

Figure 4. Example of rooting induced copies of two different graphs S and S' in a dissection.


Figure 5. A subgraph H and its corresponding set $\mathcal{H}$.

one graph in $\mathcal{H}$ in D. An illustration of this is given in Figure 6, where the copies of the graph in Figure 5 are highlighted.


Figure 6. Different rooted copies of the subgraph H of Figure 5 in a dissection D.

For a dissection D, let $\mathsf{rcopy}(\widehat{H}; D)$ be the number of rooted copies of $\widehat{H}$ in D. Then the above discussion leads immediately to the following statement.

Lemma 3.11. Let $\mathcal{H} = \{\widehat{H}_1, \ldots, \widehat{H}_{r_H}\}$. Then

$$\mathsf{copy}\,(H; D) = \sum_{i=1}^{r_H} \mathsf{rcopy}(\widehat{H}_i; D).$$

Proof. Every induced copy of H in D can be uniquely rooted in the way we defined above. Therefore the set of induced copies of H in D is partitioned

into sets of rooted copies isomorphic to the elements of $\mathcal{H}$, and the statement follows. ∎

The next lemma is the crucial point of the proof of Theorem 1.4. It states that in order to count the number of induced copies of a given H in D_n it is enough to fix one element of $\mathcal{H}$ and to count the number of rooted copies of this element.

Lemma 3.12. *Let* $\widehat{H} \in \mathcal{H}$, $\widehat{n} := |V(\widehat{H})|$ *and* $q := \frac{2-\sqrt{2}}{2}$. *Then for every* $\frac{\log n}{\sqrt{n8^{-n_H}}} < \varepsilon = \varepsilon(n) < 1$

$$(3.10) \qquad \mathbb{P}\left[\operatorname{rcopy}(\widehat{H}; \mathsf{D}_n) \in (1 \pm \varepsilon)\frac{1}{2}q^{\widehat{n}-3}n \right] \geq 1 - e^{-2\left(\frac{1}{8}\right)^{\widehat{n}}\varepsilon^2 n}.$$

Proof. We prove that the statement is true for the object returned by $\widetilde{\Gamma}D(n)$ (recall that this is the sampler that calls $\Gamma D(\rho_D)$ n^3 times, and returns a dissection on n vertices with probability at least $1 - e^{-n}$). Then the lemma follows by applying Proposition 2.2. The proof has many similarities with the proof for the degree sequence (Section 3), and therefore we will only explain the main differences. We proceed by induction on the number of faces $f_{\widehat{H}}$ in $\widehat{H}$.

If $f_{\widehat{H}} = 0$, then $\widehat{H}$ is just a single edge, i.e., we want to count how many edges there are in the output of $\widetilde{\Gamma}D(n)$. (3.10) is in this case an immediate consequence of Lemma 3.10.

Let us now discuss the induction step. We assume that the statement is true for all subgraphs of $\widehat{H}$ that are dissections obtained by removing exactly one non-root face from $\widehat{H}$. Let $T_{\widehat{H}}$ be the internal dual tree corresponding to $\widehat{H}$ as defined in Section 3, and observe that $T_{\widehat{H}}$ is an *embedded* tree, i.e., if we change the order of the children of one vertex, we will obtain a different tree. Suppose that we always draw $T_{\widehat{H}}$ with its root at the bottom, and letting it grow upwards. We define a dissection $\widehat{H}'$ as follows. If $\widehat{H}$ is a single bounded face, then $\widehat{H}'$ is the dissection consisting of a single rooted edge. Else, let $T_{\widehat{H}'}$ be the tree obtained from $T_{\widehat{H}}$ by removing the rightmost leaf (which corresponds to deleting the "rightmost" face of $\widehat{H}$). $\widehat{H}'$ is then the dissection corresponding to $T_{\widehat{H}'}$. Notice that $\widehat{n}' := |V(\widehat{H}')| = \widehat{n} - (s-2)$, where s denotes the number of vertices in the face which was removed from $\widehat{H}$ to obtain $\widehat{H}'$. By applying the induction hypothesis on $\widehat{H}'$ we obtain that

$$(3.11) \quad \mathbb{P}\left[\operatorname{rcopy}(\widehat{H}'; \widetilde{\Gamma}D(n)) \in \left(1 \pm \frac{\varepsilon}{2}\right)\frac{1}{2}q^{\widehat{n}'-3}n \right] \geq 1 - \exp\left\{-\left(\frac{1}{8}\right)^{\widehat{n}'}\frac{\varepsilon^2}{4}n\right\}.$$

Now let us consider the execution of a slightly modified version of $\Gamma D(\rho_D)$. Let $\Gamma D(L^{=\widehat{H}}, L^{\neq\widehat{H}})$ be the sampler that receives as input two lists of the form

$$(3.12) \qquad L^{=\widehat{H}} = (\alpha_1, \alpha_2, \dots) \quad \text{and} \quad L^{\neq\widehat{H}} = (\beta_1, \beta_2, \dots),$$

namely two (infinite) sequences of variables drawn independently according to the distribution (3.2), where $x = \rho_D$. Then $\Gamma D(L^{=\widehat{H}}, L^{\neq\widehat{H}})$ proceeds as $\Gamma D(\rho_D)$, with the difference that every time that it has to make a random choice, it picks a variable from one of the two lists according to the following rules.

(R1) In the beginning the sampler draws from $L^{\neq\widehat{H}}$ if and only if $f_{\widehat{H}} \geq 2$.

(R2) Let D_i be the partial dissection generated by a call to $\Gamma D(L^{=\widehat{H}}, L^{\neq\widehat{H}})$ after having read in total exactly i values from $L^{=\widehat{H}}$ and $L^{\neq\widehat{H}}$. Moreover, suppose that the execution of the sampler is not finished, i.e., D_i is not the output of $\Gamma D(L^{=\widehat{H}}, L^{\neq\widehat{H}})$. Denote by $e_i = \{v_i, v_i'\}$ the edge on the border of D_i which is going to be identified with the root edge of the dissection that will be generated in the next recursive call to the sampler, such that v_i is identified with the root vertex of that dissection. Then $\Gamma D(L^{=\widehat{H}}, L^{\neq\widehat{H}})$ will draw the next value from $L^{=\widehat{H}}$ if and only if the edge e_i is such that attaching to it a cycle of size s would create a rooted copy of $\widehat{H}$ (in D_{i+1}).

As $\Gamma D(L^{=\widehat{H}}, L^{\neq\widehat{H}})$ just imitates the execution of $\Gamma D(\rho_D)$, due to Lemma 3.3 this is a Boltzmann sampler for $\mathcal{D}$.

The above rules and Lemma 3.5, i), imply that for every dissection D there is a quantity $t = t(D)$ such that $\Gamma D(L^{=\widehat{H}}, L^{\neq\widehat{H}})$ reads t values from list $L^{=\widehat{H}}$, and the remaining values from $L^{\neq\widehat{H}}$ such that it outputs D. Observe that t is exactly the number of rooted copies of $\widehat{H}'$ present in D, since there is exactly one edge in every rooted copy of $\widehat{H}'$ that could be replaced by a cycle so as to create a rooted copy of $\widehat{H}$. But then, due to (3.11) we have that $t \in (1 \pm \varepsilon/2) \frac{1}{2} q^{\widehat{n}'-3} n$ with probability at least $1 - \exp\{-\left(\frac{1}{8}\right)^{\widehat{n}'} \frac{\varepsilon^2}{4} n\}$. Moreover, notice that $\Gamma D(L^{=\widehat{H}}, L^{\neq\widehat{H}})$ creates one rooted copy of $\widehat{H}$ every time it reads the value s from list $L^{=\widehat{H}}$, which it accesses precisely t times.

To complete the proof, let $X_{s,n'}^{(i)}$ be the random variable counting the number of values among the first n' values in $L_i^{=\widehat{H}}$ that are equal to s. A straightforward application of Lemma 3.7 shows that for all $1 \leq i \leq n^3$ and $n' \in (1 \pm \varepsilon/2) \frac{1}{2} q^{\widehat{n}'-3} n$ we have $X_{s,n'}^{(i)} \in (1 \pm \varepsilon/3) \frac{1}{2} q^{\widehat{n}-3} n$, with probability larger than

$1-2n^4\exp\{-\frac{1}{3}\cdot\frac{\varepsilon^2}{9}\cdot\frac{1}{2}q^{\widehat{n}-3}n\}\geq 1-\exp\{-\frac{\varepsilon^2}{2}q^{\widehat{n}}n\}$. By putting all facts together we see that with probability at least $1-\exp\{-\left(\frac{1}{8}\right)^{\widehat{n}'}\frac{\varepsilon^2}{4}n\}-\exp\{-\frac{\varepsilon^2}{2}q^{\widehat{n}}n\}$ the sampler $\widetilde{\Gamma}D(n)$ will output an object that has $(1\pm\varepsilon)\frac{1}{2}q^{\widehat{n}-3}n$ rooted copies of $\widehat{H}$. Since $\widehat{n}'\leq\widehat{n}-1$ we obtain for large n

$$\exp\left\{-\left(\frac{1}{8}\right)^{\widehat{n}'}\frac{\varepsilon^2}{4}n\right\}+\exp\left\{-\frac{\varepsilon^2}{2}q^{\widehat{n}}n\right\}\leq 2\exp\left\{-\min\left\{2\left(\frac{1}{8}\right)^{\widehat{n}},\frac{1}{2}q^{\widehat{n}}\right\}\varepsilon^2 n\right\},$$

which is due to $q=\frac{2-\sqrt{2}}{2}$ and $\widehat{n}\geq 3$ smaller than $\exp\{-2\left(\frac{1}{8}\right)^{\widehat{n}}\varepsilon^2 n\}$. This proves (3.10). ∎

Proof of Theorem 1.4. Since the size of the interval $(1\pm\varepsilon)\frac{1}{2}q^{\widehat{n}-3}n$ depends only on the number of vertices of $\widehat{H}$ (for fixed n and ε), which is the same for all elements of $\mathcal{H}$, with Lemma 3.11 we obtain that the total number of induced copies of H in D_n is in $(1\pm\varepsilon)r_H\frac{1}{2}q^{\widehat{n}-3}n$ with probability at least $1-r_H e^{-2\left(\frac{1}{8}\right)^{n_H}\varepsilon^2 n}\geq 1-e^{-\left(\frac{1}{8}\right)^{n_H}\varepsilon^2 n}$. This concludes the proof of the theorem. ∎

4. Triangulations

A simplification of dissections leads us to *triangulations* of convex polygons, where each face (except for the outer face) is bounded by a triangle. The decomposition and the sampler for triangulations are therefore very similar to their corresponding counterparts for general dissections, and hence we are going to sketch very briefly the straightforward modifications in our proofs. The generating function for the class $\mathcal{T}$ of edge-rooted, unlabeled triangulations can be directly derived and is

$$T(z)=z^2+\frac{T(z)^2}{z}=\frac{z}{2}(1-\sqrt{1-4z}),$$

which becomes singular at $\rho_{\mathcal{T}}=1/4$.

To construct a random triangulation, one proceeds in a similar fashion as in the case of dissections: we begin with the root edge, and then "expand" it into a triangle with probability $\frac{T(x)}{x}$, and do the same for every newly created edge. Thus the sampler $\Gamma T(x)$ is essentially the same as $\Gamma D(x)$, with the only difference that the distribution in (3.2) has to be changed in the following way, so that it only generates values of size 2 or 3.

$$\mathbb{P}\left[\mathrm{Cyc}'(x)=\ell\right]=\begin{cases}\frac{x^2}{T(x)}, & \text{if } \ell=2\\\frac{T(x)}{x}, & \text{if } \ell=3\\0, & \text{otherwise.}\end{cases}$$

With this in mind, everything else in the proof for the degree sequence of dissections can be transferred to the simpler case of triangulations, as we have to deal only with triangles instead of cycles of arbitrary length. By observing that the expression $\frac{T(x)}{x}$ equals $\frac{1}{2}$ if we set $x = \rho_{\mathcal{T}} = \frac{1}{4}$, we obtain Theorem 1.5.

References

[1] L. COMTET: *Advanced combinatorics: The Art of Finite and Infinite Expansions*; D. Reidel Publishing Co., Dordrecht, enlarged edition, 1974.

[2] P. DUCHON, P. FLAJOLET, G. LOUCHARD and G. SCHAEFFER: Boltzmann samplers for the random generation of combinatorial structures, *Combin. Probab. Comput.* **13(4–5)** (2004), 577–625.

[3] P. FLAJOLET and M. NOY: Analytic combinatorics of non-crossing configurations, *Discrete Math.* **204(1–3)** (1999), 203–229.

[4] P. FLAJOLET and R. SEDGEWICK: *Analytic Combinatorics*, Cambridge University Press, 2009.

[5] É. FUSY: Quadratic exact size and linear approximate size random generation of planar graphs, in: *2005 International Conference on Analysis of Algorithms*, volume **AD** of *Discrete Mathematics and Theoretical Computer Science Proceedings*, pages 125–138. 2005.

[6] Z. GAO and N. C. WORMALD: The distribution of the maximum vertex degree in random planar maps, *J. Combin. Theory Ser. A* **89(2)** (2000), 201–230.

[7] Z. GAO and N. C. WORMALD: Sharp concentration of the number of submaps in random planar triangulations, *Combinatorica* **23(3)** (2003), 467–486.

[8] Z. GAO and N. C. WORMALD: Asymptotic normality determined by high moments, and submap counts of random maps; *Probab. Theory Related Fields* **130(3)** (2004), 368–376.

[9] M. HABIB, C. MCDIARMID, J. RAMIREZ-ALFONSIN and B. REED: *Probabilistic methods for algorithmic discrete mathematics*, volume **16** of *Algorithms and Combinatorics*, Springer-Verlag, Berlin, 1998.

Nicla Bernasconi, Konstantinos Panagiotou, Angelika Steger

Institute for Theoretical Computer Science
ETH Zürich
Universitätsstrasse 6
CH-8092 Zürich
Switzerland
{nicla,panagiok,steger}@inf.ethz.ch