

Asymptotic enumeration of non-crossing partitions on surfaces[☆]

Juanjo Rué^{a,*}, Ignasi Sau^b, Dimitrios M. Thilikos^c

^a CNRS, Laboratoire d'Informatique, École Polytechnique, 91128 Palaiseau Cedex, France

^b CNRS, LIRMM, Montpellier, France

^c Department of Mathematics, National and Kapodistrian University of Athens, Greece

ARTICLE INFO

Article history:

Received 14 May 2012

Received in revised form 4 October 2012

Accepted 11 December 2012

Available online 3 January 2013

Keywords:

Analytic combinatorics

Symbolic method

Map enumeration

Bijjective techniques

ABSTRACT

We generalize the notion of non-crossing partition on a disk to general surfaces with boundary. For this, we consider a surface Σ and introduce the number $C_{\Sigma}(n)$ of non-crossing partitions of a set of n points lying on the boundary of Σ . Our main result is an asymptotic estimate for $C_{\Sigma}(n)$. The proofs use bijective techniques arising from map enumeration, joint with the symbolic method and singularity analysis on generating functions. An outcome of our results is that the exponential growth of $C_{\Sigma}(n)$ is the same as the one of the n -th Catalan number, i.e., does not change when we move from the case where Σ is a disk to general surfaces with boundary.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

In combinatorics, a *non-crossing partition* of size n is a partition of the set $\{1, 2, \dots, n\}$ with the following property: if $1 \leq a < b < c < d \leq n$ and a subset of the non-crossing partition contains a and c , then no other subset contains both b and d . One can represent such a partition on a disk by placing n points on the boundary of the disk, labeled in cyclic order, and drawing each subset as a convex polygon (also called *block*) on the points belonging to the subset. Then, the *non-crossing* condition is equivalent to the fact that the drawing is plane and the blocks are pairwise disjoint. These combinatorial objects are important in many aspects; see for instance [8, 18]. The enumeration of non-crossing partitions of size n is one of the first nontrivial problems in enumerative combinatorics; it is well-known that the number of these structures (either by using direct root decompositions [9] or bijective arguments [19]) corresponds to Catalan numbers. More concretely, the number of non-crossing partitions of $\{1, 2, \dots, n\}$ on a disk is equal to the Catalan number $C(n) = \frac{1}{n+1} \binom{2n}{n}$. This paper deals with a generalization of the notion of non-crossing partition on surfaces of higher genus with boundary, either orientable or not.

In the elementary case where Σ is a disk, the enumeration of non-crossing partitions can be directly reduced by bijective arguments to the map enumeration framework (all partitions can be realized geometrically in such a way that regions are contractible). Therefore, in this case $C_{\Sigma}(n)$ is the n -th Catalan number. However, to generalize the notion of non-crossing partition to surfaces of higher genus is not straightforward and needs to be defined properly (see Section 3). In particular we are interested in non-crossing partitions of sets of size n , namely the total number of vertices over all the boundary

[☆] Most of the results of this paper were announced in the extended abstract “Dynamic programming for graphs on surfaces. Proc. of ICALP’2010, volume 6198 of LNCS, pages 372–383”, which is a combination of an algorithmic framework (whose full version can be found in [16]) and the enumerative results presented in this paper.

* Corresponding author.

E-mail addresses: juanjo.rue@gmail.com, rue1982@lix.polytechnique.fr (J. Rué), ignasi.sau@lirmm.fr (I. Sau), sedthilk@math.uoa.gr (D.M. Thilikos).

components is equal to n . Additionally, these vertices are marked in the way that is described in Section 3. The main difficulty is that there is no bijection between non-crossing partitions of a set of size n on a surface Σ and its geometric representation, which means that the same non-crossing partition can arise from different geometric representations (see Fig. 10 for an example).

In this paper, we study enumerative properties of these geometric representations. From this study we deduce asymptotic estimates for the subadjacent non-crossing partitions for every surface Σ . The main result of this paper is the following: let Σ be a surface with Euler characteristic $\chi(\Sigma)$ and whose boundary has $\beta(\Sigma)$ connected components. Then the number of non-crossing partitions on Σ , $C_\Sigma(n) = |\Pi_\Sigma(n)|$, for n large enough, verifies the asymptotic upper bound

$$|\Pi_\Sigma(n)| \leq \frac{c(\Sigma)}{\Gamma(-3/2\chi(\Sigma) + \beta(\Sigma))} n^{-3/2\chi(\Sigma) + \beta(\Sigma) - 1} 4^n, \quad (1)$$

where $c(\Sigma)$ is a function depending only on Σ (for a bound on $c(\Sigma)$, see Section 5), and Γ is the Gamma function: $\Gamma(u) = \int_0^\infty t^{u-1} e^{-t} dt$. This upper bound, together with the fact that every non-crossing partition on a disk admits a realization on Σ (and consequently $C(n) \leq C_\Sigma(n)$), gives the result

$$\lim_{n \rightarrow \infty} C_\Sigma(n)^{1/n} = \lim_{n \rightarrow \infty} C(n)^{1/n} = 4. \quad (2)$$

In other words, $C_\Sigma(n)$ has the same exponential growth as the Catalan numbers, no matter the surface Σ .

In order to get the upper bound (1), we argue in three levels: we start from a topological level, stating the precise definitions of the objects we want to study, and showing that we can restrict ourselves to the study of hypermaps and bipartite maps [5]. Once we restrict ourselves to the map enumeration framework, we use the main ideas of [2] in order to obtain combinatorial decompositions of the dual maps of the objects under study (see also [3,4]). Finally, once we have explicit expressions for the generating functions of these combinatorial families, we study generating functions (formal power series) as analytic objects. In the analytic step, we extract singular expansions of the counting series from the resulting generating functions. We derive asymptotic formulas from these singular expansions by extracting coefficients, using the Transfer Theorems of singularity analysis [11,10].

The asymptotic analysis carried out in this paper has important consequences in the design of algorithms for graphs on surfaces: the enumeration of non-crossing partitions has been used in [16] to build a framework for the design of $2^{O(k)} \cdot n^{O(1)}$ step dynamic programming algorithms to solve a broad class of NP-hard optimization problems for surface-embedded graphs on n vertices of branch width at most k . The approach is based on a new type of branch decomposition called *surface cut decomposition*, which generalizes sphere cut decompositions for planar graphs introduced by Seymour and Thomas [17] (see also [7]) and where dynamic programming should be applied for each particular problem. More precisely, the use of surface cut decompositions yields algorithms with running times with a *single-exponential dependence* on branch width, and allows to unify and improve all previous results in this active field of parameterized complexity [6,7,16]. The key idea is that the size of the tables of a dynamic programming algorithm over a surface cut decomposition can be upper-bounded in terms of the non-crossing partitions on surfaces with boundary. See [16] for more details and references.

Outline of the paper. In Section 2 we include all the definitions and the required background concerning topological surfaces, maps on surfaces, the symbolic method in combinatorics, and the singularity analysis on generating functions. In Section 3 we state the precise definition of non-crossing partition on a general surface, as well as the connection with the map enumeration framework. Upper bounds for the number of non-crossing partitions on a surface Σ with boundary are obtained in Section 4, and the main result is proved. A more detailed study of the constant $c(\Sigma)$ of Eq. (1) is done in Section 5.

2. Background and definitions

In this section we state all the necessary definitions and results needed in the sequel. In Section 2.1 we state the main results concerning topological surfaces, and in Section 2.2 we recall the basic definitions about maps on surfaces. Finally, in Section 2.3 we make a brief summary of the symbolic method in combinatorics, as well as the basic techniques in singularity analysis on generating functions.

2.1. Topological surfaces

In this work, surfaces are compact (hence closed and bounded) and their boundary is homeomorphic to a finite set (possibly empty) of disjoint simple circles. We denote by $\beta(\Sigma)$ the number of connected components of the boundary of a surface Σ . The Surface Classification Theorem [15] asserts that a compact and connected surface without boundary is determined, up to homeomorphism, by its Euler characteristic $\chi(\Sigma)$ and by its orientability. More precisely, orientable surfaces are obtained by adding $g \geq 0$ handles to the sphere \mathbb{S}^2 , obtaining a surface with Euler characteristic $2 - 2g$. Non-orientable surfaces are obtained by adding $h > 0$ cross-caps to the sphere, getting a non-orientable surface with Euler characteristic $2 - h$. Given a surface with boundary Σ , we denote by $\overline{\Sigma}$ the surface (without boundary) obtained from Σ by gluing a disk on each of the $\beta(\Sigma)$ components of the boundary of Σ . It is then easy to show that $\chi(\overline{\Sigma}) = \beta(\Sigma) + \chi(\Sigma)$. In other words, surfaces under study are determined, up to homeomorphism, by their orientability, their Euler characteristic, and the number of connected components of their boundary.




Fig. 1. A map with boundary and its dual.

A cycle on Σ is a topological subspace of Σ which is homeomorphic to a circle. We say that a cycle \mathbb{S}^1 separates Σ if $\Sigma \setminus \mathbb{S}^1$ has two connected components. The following result concerning a separating cycle is an immediate consequence of Proposition 4.2.1 in [15].

Lemma 2.1.1. *Let Σ be a surface with boundary and let \mathbb{S}^1 be a separating cycle on Σ . Let V_1 and V_2 be connected surfaces obtained by cutting Σ along \mathbb{S}^1 and gluing a disk on the newly created boundaries. Then $\chi(\Sigma) = \chi(V_1) + \chi(V_2) - 2$.*

2.2. Maps on surfaces and duality

Our main reference for maps is the monograph of Lando and Zvonkin [14]. A map on Σ is a partition of Σ into zero, one, and two dimensional sets homeomorphic to zero, one and two dimensional open disks, respectively (in this order, vertices, edges, and faces). The set of vertices, edges, and faces of a map M is denoted by $V(M)$, $E(M)$, and $F(M)$, respectively. We use $v(M)$, $e(M)$, and $f(M)$ to denote $|V(M)|$, $|E(M)|$, and $|F(M)|$, respectively. The degree $d(v)$ of a vertex v is the number of edges incident with v , counted with multiplicity (loops are counted twice). An edge of a map has two ends (also called half-edges), and either one or two sides, depending on the number of faces which is incident with.

A map is rooted if an edge and one of its half-edges and sides are distinguished as the root-edge, root-end, and root-side, respectively. This definition is equivalent to marking a corner of the skeleton of the object. Observe that rooting on orientable surfaces usually omits the choice of a root-side because the subjacent surface carries a global orientation, and maps are considered up to orientation-preserving homeomorphism. Our choice of a root-side is equivalent in the orientable case to the choice of an orientation of the surface. The root-end and -sides define the root-vertex and -face, respectively. Rooted maps are considered up to cell-preserving homeomorphisms preserving the root-edge, -end, and -side. In figures, the root-edge is indicated as an oriented edge pointing away from the root-end and crossed by an arrow pointing towards the root-side (this last, provides the orientation of the surface). For a map M , the Euler characteristic of M , which is denoted by $\chi(M)$, is the Euler characteristic of the underlying surface.

Duality. Given a map M on a surface Σ without boundary, the dual map of M , which we denote by M^* , is a map on Σ obtained by drawing a vertex of M in each face of M and an edge of M across each edge of M . If the map M is rooted, the root-edge e of M is defined in the natural way: the root-end and root-side of M correspond to the side and end of e which are not the root-side and root-end of M , respectively. This construction can be generalized to surfaces with boundary in the following way: for a map M on a surface Σ with boundary, notice that the (rooted) map M defines a (rooted) map \bar{M} on $\bar{\Sigma}$ by gluing a disk (which becomes a face of \bar{M}) along each boundary component of Σ . We call these faces of \bar{M} external. Then the usual construction for the dual map \bar{M}^* applies using the external faces. The dual of a map M on a surface Σ with boundary is the map on $\bar{\Sigma}$, denoted by M^* , constructed from \bar{M}^* by splitting each external vertex of \bar{M}^* (that is, vertices associated to external faces).

The new vertices that are obtained are called dangling leaves, which have degree one. Observe that we can reconstruct the map M from M^* , by pasting the dangling leaves incident with the same face, and applying duality. An example of this construction is shown in Fig. 1.

2.3. The symbolic method and analytic combinatorics

Our main reference in enumerative combinatorics is the book of Flajolet and Sedgewick [11]. The framework introduced in this book gives a language to translate combinatorial conditions between combinatorial classes into equations relating the associated generating functions. This is what is called the symbolic method in combinatorics. Later, we can treat these equations as relations between analytic functions. This point of view gives the possibility to use complex analysis techniques to obtain information about the combinatorial classes. This is the origin of the term analytic combinatorics.

Table 1
Constructions and translations into GFs.

Construction	GF
Union	$\mathcal{A} \cup \mathcal{B}$ $\mathbf{A}(x) + \mathbf{B}(x)$
Product	$\mathcal{A} \times \mathcal{B}$ $\mathbf{A}(x)\mathbf{B}(x)$
Sequence	$\text{Seq}(\mathcal{A})$ $\frac{1}{1-\mathbf{A}(x)}$
Pointing	\mathcal{A}^\bullet $x \frac{\partial}{\partial x} \mathbf{A}(x)$

The symbolic method. For a set \mathcal{A} of objects, let $|\cdot|$ be an application (called *size*) from \mathcal{A} to \mathbb{N} . We assume that the number of elements in \mathcal{A} with a fixed size is always finite. A pair $(\mathcal{A}, |\cdot|)$ is called a *combinatorial class*. Under these assumptions, we define the formal power series (called the *generating function* or *GF* associated with the class) $\mathbf{A}(x) = \sum_{a \in \mathcal{A}} x^{|a|} = \sum_{n=0}^{\infty} a_n x^n$. Conversely, we write $a_n = [x^n]\mathbf{A}(x)$. The *symbolic method* provides a direct way to translate combinatorial constructions between combinatorial classes into equations between GFs. The constructions we use in this work and their translation into the language of GFs are shown in Table 1.

The union $\mathcal{A} \cup \mathcal{B}$ of \mathcal{A} and \mathcal{B} refers to the disjoint union of the classes. The Cartesian product $\mathcal{A} \times \mathcal{B}$ of \mathcal{A} and \mathcal{B} is the set $\{(a, b) : a \in \mathcal{A}, b \in \mathcal{B}\}$. The sequence $\text{Seq}(\mathcal{A})$ of a set \mathcal{A} corresponds to the set $\mathcal{E} \cup \mathcal{A} \cup (\mathcal{A} \times \mathcal{A}) \cup (\mathcal{A} \times \mathcal{A} \times \mathcal{A}) \cup \dots$, where \mathcal{E} denotes the empty set. At last, the pointing operator \mathcal{A}^\bullet of a set \mathcal{A} consists in pointing one of the atoms of each element $a \in \mathcal{A}$. Notice that in the sequence construction, the expression $\mathcal{E} \cup \mathcal{A} \cup (\mathcal{A} \times \mathcal{A}) \cup (\mathcal{A} \times \mathcal{A} \times \mathcal{A}) \cup \dots$ translates into $\sum_{k=0}^{\infty} \mathbf{A}(x)^k$, which is a sum of a geometric series. In the case of pointing, note also that $x \frac{\partial}{\partial x} \mathbf{A}(x) = \sum_{n>0} n a_n x^n$.

Singularity analysis. The study of the asymptotic growth of the coefficients of GFs can be obtained by considering GFs as complex functions analytic around $z = 0$. This is the main idea of analytic combinatorics. The growth behavior of the coefficients depends only on the smallest positive singularity of the GF. Its *location* provides the *exponential growth* of the coefficients, and its *behavior* gives the *subexponential growth* of the coefficients.

More concretely, for real numbers $R > \rho > 0$ and $0 < \phi < \pi/2$, let $\Delta_\rho(\phi, R)$ be the set $\{z \in \mathbb{C} : |z| < R, z \neq \rho, |\text{Arg}(z - \rho)| > \phi\}$. We call a set of this type a *dented domain* or a *domain dented at ρ* . Let $\mathbf{A}(z)$ and $\mathbf{B}(z)$ be GFs whose smallest singularity is the real number ρ . We write $\mathbf{A}(z) \sim_{z \rightarrow \rho} \mathbf{B}(z)$ if $\lim_{z \rightarrow \rho} \mathbf{A}(z)/\mathbf{B}(z) = 1$. We obtain the asymptotic expansion of $[z^n]\mathbf{A}(z)$ by *transferring* the behavior of $\mathbf{A}(z)$ around its singularity from a simpler function $\mathbf{B}(z)$, from which we know the asymptotic behavior of their coefficients. This is the main idea of the so-called *Transfer Theorems* developed by Flajolet and Odlyzko [10]. These results allow us to deduce asymptotic estimates of an analytic function using its asymptotic expansion near its dominant singularity. In our work we use a mixture of Theorems VI.1 and VI.3 from [11]:

Proposition 2.3.1 (*Transfer Theorem*). *If $\mathbf{A}(z)$ is analytic in a dented domain $\Delta = \Delta_\rho(\phi, R)$, where ρ is unique singularity with smallest modulo of $\mathbf{A}(z)$, and*

$$\mathbf{A}(z) \underset{z \in \Delta, z \rightarrow \rho}{\sim} c \cdot \left(1 - \frac{z}{\rho}\right)^{-\alpha} + O\left(\left(1 - \frac{z}{\rho}\right)^{-\alpha+\gamma}\right),$$

for $\alpha \notin \{0, -1, -2, \dots\}$, and $\gamma > 0$ then

$$a_n = c \cdot \frac{n^{\alpha-1}}{\Gamma(\alpha)} \cdot \rho^{-n} (1 + O(n^{-\gamma})), \tag{3}$$

where Γ is the Gamma function: $\Gamma(u) = \int_0^\infty t^{u-1} e^{-t} dt$.

3. Non-crossing partitions on surfaces with boundary

In this section we introduce the precise definition of a non-crossing partition on a surface with boundary. The notion of a non-crossing partition on a general surface is not as simple as in the case of a disk, and must be stated in terms of objects more general than maps. Our strategy to obtain asymptotic estimates for the number of non-crossing partitions on surfaces consists of showing that we can restrict ourselves to the study of certain families of maps. More concretely, we show that the study of non-crossing partitions is a particular case of the study of hypermaps [5], which can be interpreted as bipartite maps.

The plan for this section is the following: in Section 3.1 we set up our notation and we define a non-crossing partition on a general surface. In Section 3.2 we show that we can restrict ourselves to the study of bipartite maps in which vertices belong to the boundary of the surface.

3.1. Bipartite subdivisions and non-crossing partitions

Let Σ be a connected surface with boundary, and let $\mathbb{S}_1^1, \mathbb{S}_2^1, \dots, \mathbb{S}_{\beta(\Sigma)}^1$ be the connected components of its boundary. For $1 \leq r \leq \beta(\Sigma)$, we denote by $V_r = \{1_r, 2_r, \dots, m_r\}$ a set of vertices over \mathbb{S}_r^1 , and $|V_r| = n_r$. We assume that there exists a




Fig. 2. Geometric representation of non-crossing partitions on a disk, on a cylinder, and on a Möbius band. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 3. Three bipartite subdivisions S_1 , S_2 , and S_3 . S_2 is regular but not irreducible, while S_3 is irreducible.

total number of n vertices on the boundary of Σ , hence $n_1 + \dots + n_{\beta(\Sigma)} = n$. Observe that vertices are distributed among all the components of the boundary of Σ . In particular, it is possible that a boundary component is not incident with any vertex. Vertices on each boundary component are labeled in counterclockwise order. In particular, boundary components are distinguishable. Observe that an equivalent way to label these vertices is distinguishing on each boundary component an edge-root, whose ends are vertices 1_r and 2_r . Hence, every connected component of the boundary of Σ is edge-rooted in counterclockwise order.

A bipartite subdivision with n vertices S of Σ is a decomposition of Σ (up to homeomorphisms of the surface) into zero-, one-, and two-dimensional open and connected subsets, where the set of vertices is equal to $V_1 \cup \dots \cup V_{\beta(\Sigma)}$, and there is a proper two-coloring (namely, using black and white colors) of the two-dimensional regions, in such a way that each vertex appears (possibly more than once) in the boundary of a unique black region. We also demand that the intersection of the vertex set and the boundary of each black region contains at least one vertex. See Fig. 2 for examples of geometric representations of bipartite subdivisions on different surfaces with boundary.

In general, a bipartite subdivision is not a map, as two-dimensional regions might not be contractible. Given a bipartite subdivision, we define its blocks as the closures of its black faces. The size of a block is the number of vertices appearing on its boundary, counting multiple appearances only once. A block of size k is regular if it is incident with exactly k vertices (namely, the contour walk along the border of the block meets each vertex exactly once) and it is contractible (i.e., homeomorphically equivalent to a disk). A bipartite subdivision is regular if each block is regular. A bipartite subdivision is irreducible if it is regular and all its white faces are contractible. We denote by $\mathcal{S}_\Sigma(n)$, $\mathcal{R}_\Sigma(n)$, and $\mathcal{P}_\Sigma(n)$ the set of general, regular, and irreducible bipartite subdivisions with n vertices of Σ , respectively. See Fig. 3 for examples of bipartite subdivisions. In particular, the darker blocks in the first bipartite subdivision are not regular.

Let S be a bipartite subdivision of Σ with n vertices and let X_1, \dots, X_s be the set of its blocks. Clearly, these blocks define the partition $\pi_\Sigma(S)$ of the vertex set $V_1 \cup \dots \cup V_{\beta(\Sigma)}$. We say that a partition of the vertex set is non-crossing on a surface Σ if it is equal to $\pi_\Sigma(S)$ for a certain bipartite subdivision S of Σ . A non-crossing partition is said to be regular (or irreducible) if it arises from a regular (or irreducible) bipartite subdivision. Observe that this definition generalizes the notion of a non-crossing partition on a disk. We define $\Pi_\Sigma(n)$ as the set of non-crossing partitions of Σ with n vertices and we set $C_\Sigma(n) = |\Pi_\Sigma(n)|$.

3.2. Reduction to the map framework

In this subsection we show that we can restrict ourselves to the study of bipartite maps in which vertices belong to the boundary of the surface. Later, this reduction allow us to study non-crossing partitions in the context of map enumeration.

Let Σ_1 and Σ_2 be surfaces with boundary. We write $\Sigma_2 \subset \Sigma_1$ if there exists a continuous injection $i : \Sigma_2 \hookrightarrow \Sigma_1$ such that $i(\Sigma_2)$ is homeomorphic to Σ_2 (in particular, the image by i of the boundary of Σ_2 is contained in the boundary of Σ_1).




Fig. 4. The operation of cutting the vertex v .




Fig. 5. The operation of joining boundaries.

If S is a bipartite subdivision of Σ_2 and $\Sigma_2 \subset \Sigma_1$, then the injection i induces a bipartite subdivision $i(S)$ on Σ_1 such that $\pi_{\Sigma_2}(S) = \pi_{\Sigma_1}(i(S))$. Roughly speaking, all bipartite subdivisions on Σ_2 can be realized on a surface Σ_1 which contains Σ_2 . One can write then that $\Pi_{\Sigma_2}(n) \subseteq \Pi_{\Sigma_1}(n)$ if $\Sigma_2 \subset \Sigma_1$, and then it holds that $|\Pi_{\Sigma_2}(n)| \leq |\Pi_{\Sigma_1}(n)|$. This proves the trivial bound $C(n) \leq |\Pi_{\Sigma}(n)|$ for all choices of Σ .

As the following lemma shows, regularity is preserved by injections of surfaces.

Lemma 3.2.1. *Let M_1 be a regular bipartite subdivision of Σ_1 , and let $\Sigma_1 \subset \Sigma$. Then M_1 defines a regular bipartite subdivision M over Σ such that $\pi_{\Sigma_1}(M_1) = \pi_{\Sigma}(M)$.*

Proof. Let $i : \Sigma_1 \hookrightarrow \Sigma$ be the corresponding injective application, and consider $M = i(M_1)$. In particular, a block X of M_1 is topologically equivalent to the block $i(X)$: i is a homeomorphism between Σ and $i(\Sigma)$. Hence $i(X)$ is regular (this is easily verified: if the size of X is equal to k , then $i(X)$ is contractible and its closure intersects the boundary exactly in k points, hence it is regular) and M is regular. \square

The following proposition allows us to reduce the problem to the study of regular bipartite subdivisions.

Proposition 3.2.2. *Let $S \in \mathcal{S}_{\Sigma}(n)$ be a bipartite subdivision of Σ and let $\pi_{\Sigma}(S)$ be the associated non-crossing partition on Σ . Then, there exists a regular bipartite subdivision $R \in \mathcal{R}_{\Sigma}(n)$ such that $\pi_{\Sigma}(R) = \pi_{\Sigma}(S)$.*

Proof. Let X be a block of S of size k . If X is not regular, then either its boundary meets at least one vertex several times or X is not contractible. We show that, in each of the previous cases, we can apply local operations – that may simplify the surface – and transform X into a new block with the same vertices. Hence the associated non-crossing partition remains the same.

Assume that the first case happens. Let v be a vertex incident several times with X . In this case we define the operation of *cutting a vertex* as follows: consider the intersection of a small ball of radius $\varepsilon > 0$ centered at v with the block X , namely $B_{\varepsilon}(v) \cap X$. Observe that $(B_{\varepsilon}(v) \cap X) \setminus \{v\}$ has several connected components (the same as the number of times the closure of X intersects v). We define the new block by deforming all except one of these components in such a way that they do not intersect the boundary of Σ . Next, we paste the vertex v to the unique component which has not been deformed (see Fig. 4). Then the resulting bipartite subdivision has the same associated non-crossing partition, and v is incident with the corresponding block exactly once. Applying this argument for each vertex of X we get a block which is incident once (as we move along its border) with every one of the vertices which define it.

Applying this operation at most n times the resulting block intersects each vertex either one or zero times. We can apply then this operation to all the blocks. Without loss of generality, using this operation the number of times needed, we can assume that our initial bipartite subdivision satisfies that each vertex appears at most once on the boundary of a block.

Assume now that X is a non-contractible block of S . Let S_X^1 be a non-contractible cycle contained in X . Two situations may happen: either $X \setminus S_X^1$ is connected or disconnected. We analyze both cases.

Assume first that S_X^1 disconnects X . We may assume that each component contains, at least, one boundary component (if not, the corresponding connected component of $X \setminus S_X^1$ could be substituted by a disk). In particular, invoking the Jordan’s curve theorem, there exists a pair of vertices in each connected component (and also the corresponding boundary components where these vertices belong). We define the operation of *joining boundaries* in the following way (see Fig. 5 in order to clarify this construction): we consider a path between these two vertices (in Fig. 5 it is the line joining the two vertices). This path exists as X is a connected open subset of Σ . Consider also two new paths inside X that join these two vertices around the initial path. We then define a new block X' by deleting from X the open region defined by these two paths (the region which contains the initial path between the two vertices).




Fig. 6. A non-crossing partition on a disk and the associated tree.

Applying the cut vertex operation over these two vertices we get a new block X'' , and the non-crossing partition associated with this new bipartite subdivision is the same as the initial one. Observe that there does not exist a cycle which disconnects the block in such a way that the considered boundary components lay, each of them, in a different connected component. As the number of boundary components is finite, we can apply these operations a finite number of times in order to get a new block without separating cycles.

Suppose now that $X \setminus S_X^1$ is connected. We cut the surface along S_X^1 and we paste either a disk or a pair of disks along its border, depending on whether S_X^1 is one- or two-sided on Σ . This operation decreases the genus of the surface and does not alter the vertices incident with the block (as the operation does not change the set of vertices met by each block). As the genus of the surface is bounded, we can apply this operation a finite number of times until the resulting block becomes contractible.

To conclude, after converting each block to a regular one, the resulting surface is $\Sigma_1 \subset \Sigma$. The resulting bipartite subdivision S' on Σ_1 is regular (since all the blocks are regular), and then by Lemma 3.2.1 there exists a regular bipartite subdivision R over Σ such that $\pi_\Sigma(R) = \pi_{\Sigma_1}(M')$, as claimed. \square

Notice that from Proposition 3.2.2 we also see that

$$|\Pi_\Sigma(n)| \leq |\mathcal{R}_\Sigma(n)|. \tag{4}$$

In the next section we reduce our study to the family of irreducible bipartite subdivisions. This permits us to upper-bound $|\mathcal{P}_\Sigma(n)|$ instead of dealing with the more complicated task of upper-bounding $|\mathcal{R}_\Sigma(n)|$. The reason why this also gives an asymptotic bound for $|\Pi_\Sigma(n)|$ is that the subfamily $\mathcal{P}_\Sigma(n)$ provides the main contribution to the asymptotic estimates for $\mathcal{R}_\Sigma(n)$.

4. Upper bounds for non-crossing partitions on surfaces

The plan for this section is the following: in Section 4.1 we introduce families of plane trees that arise by duality on non-crossing partitions on a disk. These combinatorial structures are used in Section 4.2 to obtain a tree-like decomposition which provides a way to obtain asymptotic estimates for the number of irreducible bipartite subdivisions of Σ with n vertices, namely $|\mathcal{P}_\Sigma(n)|$. These asymptotic estimates are found in Section 4.3 for irreducible bipartite subdivisions. Finally, we prove in Section 4.4 that the number of irreducible bipartite subdivisions is asymptotically equal to the number of regular bipartite subdivisions, hence the estimate obtained in Section 4.2 is an upper bound for the number of non-crossing partitions on surfaces. All previous steps are summarized in Section 4.5.

4.1. Planar constructions

The dual map of a non-crossing partition on a disk is a tree, which is called the *non-crossing partition tree associated with the non-crossing partition*. To abbreviate this notation, we simply say tree associated to the corresponding non-crossing partition. This tree corresponds to the notion of dual map for surfaces with boundary introduced in Section 2.2. Recall that vertices of degree one are called the *dangling leaves* of the tree. In trees associated to non-crossing partitions we have three types of vertices: vertices of the tree are called *block vertices* if they are associated with a block of the non-crossing partition. The remaining vertices are either *non-block vertices* or *danglings*. By construction, all vertices adjacent to a block vertex are non-block vertices. Conversely, each vertex adjacent to a non-block vertex is either a block vertex or a dangling. Graphically, we use the symbols \blacksquare for block vertices, \square for non-block vertices, and \circ for danglings. Non-crossing partitions trees are rooted: the root of a non-crossing partition tree is defined by the root of the initial non-crossing partition on a disk (i.e., the root from vertex 1 to vertex 2). The block vertex which carries the role of the root vertex of the tree is the one associated with the block containing vertex with label 2 (or equivalently, the end-vertex of the root). See Fig. 6 for an example of this construction.




Fig. 7. A double tree and its decomposition.

Let \mathcal{T} be the set of non-crossing partitions trees, and let $\mathbf{T} = \mathbf{T}(z, u) = \sum_{n,m \geq 0} t_{n,m} z^n u^m$ be the corresponding generating function, where the variable z marks danglings and u marks block vertices. We use an auxiliary family \mathcal{B} , defined as the set of trees which are rooted at a non-block vertex. Let $\mathbf{B} = \mathbf{B}(z, u) = \sum_{n,m \geq 0} b_{n,m} z^n u^m$ be the associated generating function. The next lemma gives the exact enumeration of \mathcal{T} and \mathcal{B} . In particular, this lemma implies the well-known Catalan numbers for non-crossing partitions on a disk.

Lemma 4.1.1. *The number of non-crossing trees counted by the number of danglings and block vertices is enumerated by the generating function*

$$\mathbf{T}(z, u) = \frac{1 - z(1 - u) - \sqrt{(z(1 - u) - 1)^2 - 4zu}}{2zu}. \tag{5}$$

Furthermore, $\mathbf{B}(z, u) = z\mathbf{T}(z, u)$.

Proof. We establish combinatorial relations between \mathcal{B} and \mathcal{T} from which we deduce the result. Observe that there is no restriction on the number of vertices incident with a given block. Hence the degree of every block vertex is arbitrary. This condition is translated symbolically via the relation

$$\mathcal{T} = \{\blacksquare\} \times \text{Seq}(\mathcal{B}).$$

Similarly, \mathcal{B} can be written in the form

$$\mathcal{B} = \{\circ\} \times \text{Seq}(\mathcal{T} \times \{\circ\}).$$

These combinatorial conditions translate using Table 1 into the system of equations

$$\mathbf{T} = \frac{u}{1 - \mathbf{B}}, \quad \mathbf{B} = \frac{z}{1 - z\mathbf{T}}.$$

Substituting the expression of \mathbf{B} in the first equation, one obtains that \mathbf{T} satisfies the relation $z\mathbf{T}^2 + (z(1 - u) - 1)\mathbf{T} + u = 0$. The solution to this equation with positive coefficients is (5). Solving the previous system of equations in terms of \mathbf{B} brings $\mathbf{B} = z\mathbf{T}$, as claimed. \square

Observe that writing $u = 1$ in \mathbf{T} and \mathbf{B} we obtain that $\mathbf{T}(z) = \mathbf{T}(z, 1) = \frac{1 - \sqrt{1 - 4z}}{2z}$, and $\mathbf{B}(z) = \mathbf{B}(z, 1) = z\mathbf{T}(z)$, deducing the well-known generating function for Catalan numbers.

We introduce another family of trees related to non-crossing partitions, which we call *double trees*. A double tree is defined in the following way: consider a path where we concatenate block vertices and non-block vertices. We consider the internal vertices of the path. A double tree is obtained by pasting on every block vertex of the path a pair of elements of \mathcal{T} (one at each side of the path), and a pair of elements of \mathcal{B} for non-block vertices. We say that a double tree is of type $\blacksquare-\blacksquare$, $\blacksquare-\square$, or $\square-\square$ depending on the ends of the path. An example of a double tree of type $\blacksquare-\blacksquare$ is shown in Fig. 7.

We denote these families by $\mathcal{T}_{\blacksquare-\blacksquare}$, $\mathcal{T}_{\blacksquare-\square}$, and $\mathcal{T}_{\square-\square}$, and the corresponding generating function by $\mathbf{T}_1(z, u) = \mathbf{T}_1$, $\mathbf{T}_2(z, u) = \mathbf{T}_2$, and $\mathbf{T}_3(z, u) = \mathbf{T}_3$, respectively. Recall that in all cases z marks danglings and u marks block vertices. A direct application of the symbolic method provides a way to obtain explicit expressions for the previously defined generating functions. The decomposition and the GFs of the three families are summarized in Table 2.

Table 2
GFs for double trees.

Family	Specification	Development	Compact expression
$\mathcal{T}_{\square-\blacksquare}$	$\text{Seq}(\mathcal{B}^2 \times \mathcal{T}^2)$	$1 + \frac{1}{u}\mathbf{B}^2\mathbf{T}^2 + \frac{1}{u^2}\mathbf{B}^4\mathbf{T}^4 + \dots$	$1/(1 - \mathbf{T}^2\mathbf{B}^2/u)$
$\mathcal{T}_{\blacksquare-\blacksquare}$	$\mathcal{B}^2 \times \text{Seq}(\mathcal{B}^2 \times \mathcal{T}^2)$	$\mathbf{B}^2 + \frac{1}{u}\mathbf{B}^4\mathbf{T}^2 + \frac{1}{u^2}\mathbf{B}^6\mathbf{T}^4 + \dots$	$\mathbf{B}^2/(1 - \mathbf{T}^2\mathbf{B}^2/u)$
$\mathcal{T}_{\square-\square}$	$\mathcal{T}^2 \times \text{Seq}(\mathcal{B}^2 \times \mathcal{T}^2)$	$\frac{1}{u}\mathbf{T}^2 + \frac{1}{u^2}\mathbf{B}^2\mathbf{T}^4 + \frac{1}{u^3}\mathbf{B}^4\mathbf{T}^6 + \dots$	$\frac{1}{u}\mathbf{T}^2/(1 - \mathbf{T}^2\mathbf{B}^2/u)$




Fig. 8. The construction of the scheme of an element in \mathcal{P}_Σ . We consider the dual of an irreducible bipartite subdivision (leftmost figure). After deleting vertices of degree one recursively and dissolving vertices of degree two, we obtain the associated scheme (rightmost figure). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

To conclude, the family of *pointed* non-crossing trees \mathcal{T}^\bullet is built by pointing a dangling on each non-crossing partition tree. In this case, the associated GF is $\mathbf{T}^\bullet = z \frac{\partial}{\partial z} \mathbf{T}$. Similar definitions can be done for the family \mathcal{B} . Pointing a dangling defines a unique path between this distinguished dangling and the root of the tree.

4.2. The scheme of an irreducible bipartite subdivision

In this subsection we generalize the construction of non-crossing partition trees introduced in Section 4.1. In order to characterize it, we exploit the dual construction for maps on surfaces (see Section 2.2). More concretely, for an element $M \in \mathcal{P}_\Sigma(n)$, let M^* be the dual map of M on $\overline{\Sigma}$. By construction, there is no incidence in M^* between either pairs of block vertices or pairs of non-block vertices.

From M^* we define a new rooted map (a root for each boundary component of Σ) on $\overline{\Sigma}$ in the following way: we start by deleting recursively vertices of degree one which are not roots. Then we continue *dissolving* vertices of degree two, that is, replacing the two edges incident to a vertex of degree two with a single edge. The resulting map has $\beta(\Sigma)$ faces and all vertices have degree at least three (apart from root vertices, which have degree one), and vertices of two colors (vertices of different colors could be end-vertices of the same edge). The resulting map is called the *scheme associated with M*; we denote it by s_M . See Fig. 8 for an example of this construction.

The previous decomposition can be constructed in the reverse way: duals of irreducible bipartite subdivision are constructed from a generic scheme s in the following way.

- (1) For an edge of s with both end-vertices of type \blacksquare , we paste a double tree of type $\blacksquare-\blacksquare$ along it. Similar operations are done for edges with end-vertices $\{\square, \blacksquare\}$ and $\{\square, \square\}$.
- (2) For vertex v of s of type \blacksquare we paste $d(v)$ elements of \mathcal{T} (identifying the roots of the trees with v), one on each corner of v . The same operation is done for vertices of s of type \square .
- (3) We paste an element of \mathcal{T}^\bullet along each one of the roots of s (the marked leaf determines the dangling root).

To conclude, this construction provides a way to characterize the set of schemes. Indeed, if we denote by \mathfrak{S}_Σ the set of maps on $\overline{\Sigma}$ with $\beta(\Sigma)$ faces with a root on each face and with vertices of two different colors (namely, vertices of type \blacksquare and \square), then $|\mathfrak{S}_\Sigma|$ is finite, since the surface is fixed, the number of faces of each element in \mathfrak{S}_Σ is equal to $\beta(\Sigma)$ and we are not dealing with vertices of degree 1 and 2. In fact, \mathfrak{S}_Σ is the set of all possible schemes: from an arbitrary element $s \in \mathfrak{S}_\Sigma$ we can construct a map on $\overline{\Sigma}$ with $\beta(\Sigma)$ faces by pasting double trees along each edge of s (according to the end-vertices of each edge). In other words, given $M \in \mathcal{P}_\Sigma(n)$ and s_M , M^* can be reconstructed by pasting on every edge of s_M a double tree, depending on the nature of the end-vertices of each edge of s_M . See Fig. 9 for an example.

4.3. Asymptotic enumeration

The decomposition introduced in Section 4.2 can be exploited in order to get asymptotic estimates for $|\mathcal{P}_\Sigma(n)|$, and consequently upper bounds for $|\mathcal{N}_\Sigma(n)|$. In this subsection we provide estimates for the number of irreducible bipartite




Fig. 9. The decomposition into bicolored trees and the associated scheme. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

subdivisions. We obtain these estimates directly for the surface Σ , while the usual technique consists in reducing the enumeration to surfaces of smaller genus, and returning back to the initial one by topological “pasting” arguments. The main point consists in exploiting tree structures of the dual graph associated with an irreducible bipartite subdivision. The main ideas are inspired by [2], where the authors find the asymptotic enumeration of simplicial decompositions of surfaces with boundaries without interior points.

We use the notation and definitions introduced in Section 4.1 (i.e., families of trees, double trees and pointed trees, and the corresponding GFs), joint with the decomposition introduced in Section 4.2. Let us now introduce some extra notation.

We denote by $\mathcal{P}_\Sigma(n, m)$ the set of irreducible bipartite subdivisions of Σ with n vertices and m blocks. We write $p_{n,m}^\Sigma$ for the cardinality of this set and $\mathbf{P}_\Sigma(z, u) = \sum_{n,m \geq 0} p_{n,m}^\Sigma z^n u^m$. Let $p_n^\Sigma = \sum_{m \geq 0} p_{n,m}^\Sigma = [z^n] \mathbf{P}_\Sigma(z, 1)$. Let $\mathfrak{s} \in \mathfrak{S}_\Sigma$. Denote by $v_1(\mathfrak{s})$ and $v_2(\mathfrak{s})$ the set of vertices of type \blacksquare and \square of \mathfrak{s} , respectively. Write $b(\mathfrak{s}), w(\mathfrak{s})$ for the number of roots which are incident with a vertex of type \blacksquare and \square , respectively. In particular, $b(\mathfrak{s}) + w(\mathfrak{s}) = \beta(\Sigma)$. Denote by $e_1(\mathfrak{s})$ the number of edges in \mathfrak{s} of type $\blacksquare\text{--}\blacksquare$. We similarly define $e_2(\mathfrak{s})$ and $e_3(\mathfrak{s})$ for edges of type $\square\text{--}\blacksquare$ and $\square\text{--}\square$, respectively. Observe that $e_1(\mathfrak{s}) + e_2(\mathfrak{s}) + e_3(\mathfrak{s}) + b(\mathfrak{s}) + w(\mathfrak{s})$ is equal to the number of edges of \mathfrak{s} , namely $e(\mathfrak{s})$. For a vertex v of \mathfrak{s} , denote by $r(v)$ the number of roots which are incident with it. Finally, denote by $\mathfrak{C}_\Sigma \subset \mathfrak{S}_\Sigma$ the set of maps on \mathfrak{S}_Σ whose vertex degree is equal to three (namely, cubic maps on Σ with $\beta(\Sigma)$ faces).

The decomposition discussed in Section 4.2 together with Proposition 2.3.1 give the following:

Lemma 4.3.1. *Let Σ be a surface with boundary. Then*

$$[z^n] \mathbf{P}_\Sigma(z, 1) = p_n^\Sigma = \frac{c(\Sigma)}{\Gamma(-3\chi(\Sigma)/2 + \beta(\Sigma))} n^{-3\chi(\Sigma)/2 + \beta(\Sigma) - 1} 4^n (1 + O(n^{-1/2})), \tag{6}$$

where $c(\Sigma)$ is a nonzero value depending only on Σ .

Proof. If Σ is the disk we obtain the well known asymptotic estimate for Catalan numbers, which fits into the general expression in (6). Let us assume that Σ it is a surface with boundary different from the disk.

According to the decomposition introduced in Section 4.2, $\mathbf{P}_\Sigma(z, u)$ can be written in the following form: for each $\mathfrak{s} \in \mathfrak{S}_\Sigma$, we replace edges (not roots) with double trees, roots with pointed trees, and vertices with sets of trees. More concretely,

$$\mathbf{P}_\Sigma(z, u) = \sum_{\mathfrak{s} \in \mathfrak{S}_\Sigma} u^{|v_1(\mathfrak{s})|} \mathbf{T}_1^{e_1(\mathfrak{s})} \mathbf{T}_2^{e_2(\mathfrak{s})} \mathbf{T}_3^{e_3(\mathfrak{s})} \left(\frac{\mathbf{T}}{u}\right)^{\sum_{x \in v_1(\mathfrak{s})} (d(x) - 2r(x))} \times \mathbf{B}^{\sum_{y \in v_2(\mathfrak{s})} (d(y) - 2r(y))} \left(\frac{\mathbf{T}^\bullet}{u}\right)^{B(\mathfrak{s})} (\mathbf{B}^\bullet)^{W(\mathfrak{s})}. \tag{7}$$

Observe in the previous expression that terms \mathbf{T} and \mathbf{T}^\bullet appear divided by u : blocks on the dual map are considered in the term $u^{|v_1(\mathfrak{s})|}$, so we do not consider the root of the different non-crossing trees.

To obtain the asymptotic behavior in terms of the number of danglings, we write $u = 1$ in Eq. (7). To study the resulting GF, we need the expression of each factor of Eq. (7) when we write $u = 1$; all these expressions are shown in Table 3. This table is built from the expressions for \mathbf{T} and \mathbf{B} deduced in Lemma 4.1.1 and the expressions for double trees in Table 2.

Table 3
Univariate GF for all families of trees.

GF	Expression
$\mathbf{T}_1(z, 1)$	$1/16(1 - 4z)^{-1/2} - 1/8(1 - 4z)^{1/2} + 1/16(1 - 4z)^{3/2}$
$\mathbf{T}_2(z, 1)$	$1/4(1 - 4z)^{-1/2} + 1/2 + (1 - 4z)^{1/2}$
$\mathbf{T}_3(z, 1)$	$z^2 (1/16(1 - 4z)^{-1/2} - 1/8(1 - 4z)^{1/2} + 1/16(1 - 4z)^{3/2})$
$\mathbf{T}(z, 1)$	$(1 - (1 - 4z)^{1/2})/(2z)$
$\mathbf{B}(z, 1)$	$(1 - (1 - 4z)^{1/2})/2$
$\mathbf{T}^*(z, 1)$	$(1 - 4z)^{-1/2}/z - (1 - (1 - 4z)^{-1/2})/(2z^2)$
$\mathbf{B}^*(z, 1)$	$(1 - 4z)^{-1/2}$

The GF in Eq. (7) is a finite sum (a total of $|\mathfrak{S}_\Sigma|$ terms), so its singularity is located at $z = 1/4$ (since each addend has a singularity at this point). For each choice of \mathfrak{s} ,

$$\mathbf{T}(z, 1)^{\sum_{x \in v_1(\mathfrak{s})} (d(x) - 2r(x))} \mathbf{B}(z, 1)^{\sum_{y \in v_2(\mathfrak{s})} (d(y) - 2r(y))} = \sum_{n=0}^{f(\mathfrak{s})} f_n(z) (1 - 4z)^{n/2}, \tag{8}$$

where the positive integer $f(\mathfrak{s})$ depends only on \mathfrak{s} , $f_n(z)$ are functions analytic at $z = 1/4$, and $f_0(z) \neq 0$ at $z = 1/4$. For the other multiplicative terms, we obtain

$$\mathbf{T}_1(z, 1)^{e_1(\mathfrak{s})} \mathbf{T}_2(z, 1)^{e_2(\mathfrak{s})} \mathbf{T}_3(z, 1)^{e_3(\mathfrak{s})} \mathbf{T}^*(z, 1)^{B(\mathfrak{s})} \mathbf{B}^*(z, 1)^{W(\mathfrak{s})} = G_{\mathfrak{s}}(z) (1 - 4z)^{-\frac{e(\mathfrak{s})}{2}} + \dots, \tag{9}$$

where $G_{\mathfrak{s}}(z)$ is a function analytic at $z = 1/4$. The reason for this fact is that each factor in Eq. (9) can be written in the form $p(z)(1 - 4z)^{-1/2} + \dots$, where $p(z)$ is a function analytic at $z = 1/4$, and $e_1(\mathfrak{s}) + e_2(\mathfrak{s}) + e_3(\mathfrak{s}) + b(\mathfrak{s}) + w(\mathfrak{s})$ is the total number of edges. Multiplying Expressions (8) and (9) we obtain the contribution of a map \mathfrak{s} in $\mathbf{P}_\Sigma(z, 1)$. More concretely, the contribution of a single map \mathfrak{s} to Eq. (7) can be written in the form

$$g_{\mathfrak{s}}(z) (1 - 4z)^{-e(\mathfrak{s})/2} + \dots,$$

where $g_{\mathfrak{s}}(z)$ is a function analytic at $z = 1/4$. Looking at (3) from Proposition 2.3.1, we deduce that the maps giving the main contribution to the asymptotic estimate of p_k^Σ are the ones maximizing the value $e(\mathfrak{s})$. Applying Euler’s formula (recall that all maps in \mathfrak{S}_Σ have $\beta(\Sigma)$ faces) on $\overline{\Sigma}$ gives that these maps are precisely the maps in \mathfrak{C}_Σ . In particular, maps in \mathfrak{C}_Σ have $2\beta(\Sigma) - 3\chi(\Sigma)$ edges. Hence, the singular expansion of $\mathbf{P}_\Sigma(z, 1)$ at $z = 1/4$ is

$$\mathbf{P}_\Sigma(z, 1) \sim_{z \rightarrow 1/4} c(\Sigma) (1 - 4z)^{3\chi(\Sigma)/2 - \beta(\Sigma)} (1 + O((1 - 4z)^{1/2})), \tag{10}$$

where $c(\Sigma) = \sum_{\mathfrak{s} \in \mathfrak{C}_\Sigma} g_{\mathfrak{s}}(1/4)$ counts the number of such maps (so, in particular, this value is not zero). Applying Proposition 2.3.1 on this expression yields the result as claimed. \square

4.4. Irreducibility vs reducibility

For conciseness, in this subsection we write

$$a(\Sigma) = \frac{c(\Sigma)}{\Gamma(-3\chi(\Sigma)/2 + \beta(\Sigma))}$$

to denote the constant term which appears in Eq. (6) from Lemma 4.3.1. Observe that for a non-irreducible bipartite subdivision M of \mathcal{R}_Σ , there is a non-contractible cycle \mathbb{S}^1 contained in a white two-dimensional region of M . Additionally, M induces a regular bipartite subdivision on the surface $\Sigma \setminus \mathbb{S}^1 = \Sigma'$, which can be irreducible or not. By Lemma 3.2.1, each element of $\mathcal{R}_{\Sigma'}$ defines an element of \mathcal{R}_Σ . To prove that irreducible bipartite subdivisions over Σ give the maximal contribution to the asymptotic, we apply a double induction argument on the pair $(\chi(\Sigma), \beta(\Sigma))$. The critical point is the initial step, namely genus equal to 0 and several boundaries, which corresponds to the case where $\overline{\Sigma}$ is the sphere. The details are shown in the following lemma.

Lemma 4.4.1. *Let Σ be a surface obtained from the sphere by deleting β disjoint disks. Then*

$$|\mathcal{R}_\Sigma(n) \setminus \mathcal{P}_\Sigma(n)| = o(|\mathcal{P}_\Sigma(n)|).$$

Proof. We apply induction on β (in this proof the genus of $\overline{\Sigma}$ is always equal to 0). The case $\beta = 1$ corresponds to a disk, where $\mathbf{T}(z, u)$ in Eq. (5) is the same as $\mathbf{P}_\Sigma(z, u)$. In this case, the equality $\mathcal{R}_\Sigma(n) = \mathcal{P}_\Sigma(n)$ holds for every value of n .

Let us proceed with the inductive step. Let $\beta > 1$ be the number of boundary connected components of Σ . Applying Jordan’s Theorem in the sphere, a non-contractible cycle always separates Σ into two connected components, namely Σ_1 and Σ_2 . By induction hypothesis,

$$|\mathcal{R}_{\Sigma_j}(n) \setminus \mathcal{P}_{\Sigma_j}(n)| = o(|\mathcal{P}_{\Sigma_j}(n)|),$$

for $j = 1, 2$. Hence, we only need to deal with irreducible decompositions of Σ_1 and Σ_2 (as the number of boundary components of Σ_1 and Σ_2 is smaller than β and its genus is 0). The GF of the regular bipartite subdivisions that reduce to decompositions over Σ_1 and Σ_2 has the same asymptotic as $\mathbf{P}_{\Sigma_1}(z, 1) \cdot \mathbf{P}_{\Sigma_2}(z, 1)$. The estimate of its coefficients is

$$[z^n]\mathbf{P}_{\Sigma_1}(z, 1)\mathbf{P}_{\Sigma_2}(z, 1) = a(\Sigma_1)a(\Sigma_2)[z^n](1 - 4z)^{-5/2\beta(\Sigma_1)+3}(1 - 4z)^{-5/2\beta(\Sigma_2)+3}.$$

Applying Proposition 2.3.1 gives the estimate $[z^n]\mathbf{P}_{\Sigma_1}(z, 1)\mathbf{P}_{\Sigma_2}(z, 1) = O(n^{5/2\beta-7}4^n)$. Consequently, when n is large enough the above term is smaller than $p_n^\Sigma = O(n^{5/2\beta-4}4^n)$, and the result follows. \square

The next step consists in adapting the previous argument to surfaces of arbitrary genus but with just one boundary component. This is proved in the next lemma:

Lemma 4.4.2. *Let Σ be a surface such that the number of connected components of its boundary is equal to 1. Then,*

$$|\mathcal{R}_\Sigma(n) \setminus \mathcal{P}_\Sigma(n)| = o(|\mathcal{P}_\Sigma(n)|).$$

Proof. We apply induction on the Euler genus of the surface. For Euler genus 0 the result is the one in the disk, and the claim is obvious. Let us assume the induction hypothesis and suppose that the Euler genus of the surface is greater than or equal to 1. For a regular (but not irreducible) bipartite subdivision, consider a non-contractible cycle C on one of its white faces. Observe that C can be either one- or two-sided. Let \mathcal{Y} be the surface obtained from $\Sigma \setminus C$ by pasting a disk (or two disks) along the cut (depending on whether C is one- or two-sided). As the boundary has exactly one connected component, either the resulting surface \mathcal{Y} is connected and we set $\Sigma' = \mathcal{Y}$ or \mathcal{Y} has two connected components and we define Σ' as the one that contains the boundary. Hence, we have the relation

$$\mathcal{R}_\Sigma(n) \setminus \mathcal{P}_\Sigma(n) \subseteq \bigcup_{\Sigma'} \mathcal{R}_{\Sigma'}(n),$$

where the union runs over all connected surfaces Σ' with one boundary component with genus strictly smaller than the one of Σ . By induction hypothesis, $|\mathcal{R}_{\Sigma'}(n) \setminus \mathcal{P}_{\Sigma'}(n)| = o(|\mathcal{P}_{\Sigma'}(n)|)$, and $|\mathcal{R}_{\Sigma'}(n)| \sim |\mathcal{P}_{\Sigma'}(n)|$.

The main contribution to the asymptotic of $|\bigcup_{\Sigma'} \mathcal{P}_{\Sigma'}(n)|$ arises from the surface whose genus is a unit less than the one of Σ , which corresponds to

$$|\mathcal{P}_{\Sigma'}(n)| = a(\mathcal{Y})n^{-3/2\chi(\mathcal{Y})-3/2}4^n(1 + O(n^{-1/2})) = o(n^{-3/2\chi(\Sigma)}4^n),$$

and the result holds. \square

This last step is done in the following proposition, in which we deal with the general situation:

Proposition 4.4.3. *Let Σ be a surface with boundary. Then*

$$|\mathcal{R}_\Sigma(n) \setminus \mathcal{P}_\Sigma(n)| = o(|\mathcal{P}_\Sigma(n)|).$$

Proof. Let Σ be a surface with boundary and Euler characteristic $\chi(\Sigma)$. Consider a non-contractible cycle C contained in a two-dimensional white region. Observe that C can be either one- or two-sided. Let \mathcal{Y} be the surface obtained from $\Sigma \setminus C$ by pasting a disk (or two disks) along the cut (depending on whether C is one- or two-sided). The two situations that could occur are the following:

- (1) \mathcal{Y} is connected and $\beta(\mathcal{Y}) = \beta(\Sigma)$. In this case, the Euler characteristic has been increased by either one if the cycle is one-sided or by two if the cycle is two-sided (see Lemma 4.2.4 in [15]).
- (2) The resulting surface has two connected components, namely $\mathcal{Y} = \mathcal{Y}_1 \sqcup \mathcal{Y}_2$. In this case, the total number of boundaries is $\beta(\mathcal{Y}) = \beta(\mathcal{Y}_1) + \beta(\mathcal{Y}_2)$. By Lemma 2.1.1, $\chi(\Sigma) = \chi(\mathcal{Y}_1) + \chi(\mathcal{Y}_2) - 2$.

Clearly, the base of the induction is given by Lemmas 4.4.1 and 4.4.2. The induction argument distinguishes between the following two cases:

Case 1. \mathcal{Y} is connected, by induction on the genus, $|\mathcal{R}_\mathcal{Y}(n) \setminus \mathcal{P}_\mathcal{Y}(n)| < |\mathcal{P}_\mathcal{Y}(n)|$ when $n \rightarrow \infty$. Additionally, by Relation (6), an upper bound for $|\mathcal{P}_\mathcal{Y}(n)|$ is

$$[z^n]\mathbf{P}_\mathcal{Y}(z, 1) = a(\mathcal{Y})n^{-3/2\chi(\mathcal{Y})+\beta(\mathcal{Y})-1-3/2}4^n(1 + O(n^{-1/2})) = o(n^{-3/2\chi(\Sigma)+\beta(\Sigma)-1}4^n).$$

Case 2. \mathcal{Y} is not connected. Then $\mathcal{Y} = \mathcal{Y}_1 \sqcup \mathcal{Y}_2$, $\beta(\Sigma) = \beta(\mathcal{Y}) = \beta(\mathcal{Y}_1) + \beta(\mathcal{Y}_2)$, and $\chi(\Sigma) = \chi(\mathcal{Y}_1) + \chi(\mathcal{Y}_2) - 2$. Again, by induction hypothesis we only need to consider, in both surfaces, irreducible bipartite subdivisions (which are the ones with the main contribution in the asymptotic). Consequently,

$$[z^n]\mathbf{P}_{\mathcal{Y}_1}(z, 1)\mathbf{P}_{\mathcal{Y}_2}(z, 1) = a(\mathcal{Y}_1)a(\mathcal{Y}_2)[z^n](1 - 4z)^{3/2(\chi(\mathcal{Y}_1)+\chi(\mathcal{Y}_2))-(\beta(\mathcal{Y}_1)+\beta(\mathcal{Y}_2))+3}.$$

The exponent of $(1 - 4z)$ in the last equation can be written as $3/2\chi(\Sigma) - \beta(\Sigma) + 3$. Consequently, the value $[z^n]\mathbf{P}_{\mathcal{Y}_1}(z, 1)\mathbf{P}_{\mathcal{Y}_2}(z, 1)$ is bounded, for n large enough, by

$$n^{-3/2\chi(\Sigma)+\beta(\Sigma)-3-1}4^n = n^{-3/2\chi(\Sigma)+\beta(\Sigma)-4}4^n = o(n^{-3/2\chi(\Sigma)+\beta(\Sigma)-1}4^n).$$

Hence the contribution is smaller than the one given by $|\mathcal{P}_\Sigma(n)|$, as claimed. \square

Table 4
A simplification of Table 3 used in Lemma 4.3.1.

GF	Expression	Development at $z = 1/4$
$T_1(z)$	$(1 - 4z)^{-1/2}/16 + \dots$	$1/16(1 - 4z)^{-1/2} + \dots$
$T_2(z)$	$(1 - 4z)^{-1/2}/4 + \dots$	$1/4(1 - 4z)^{-1/2} + \dots$
$T_3(z)$	$z^2/16(1 - 4z)^{-1/2} + \dots$	$1/256(1 - 4z)^{-1/2} + \dots$
$T(z)$	$1/(2z) + \dots$	$2 + \dots$
$B(z)$	$1/2 + \dots$	$1/2 + \dots$
$T^\bullet(z)$	$(1 - 4z)^{-1/2}/z + \dots$	$4(1 - 4z)^{-1/2} + \dots$
$B^\bullet(z)$	$(1 - 4z)^{-1/2}$	$(1 - 4z)^{-1/2}$

4.5. Upper bounds for non-crossing partitions

In this subsection we summarize all the steps in the previous subsections of this section. Our main result is the following:

Theorem 4.5.1. *Let Σ be a surface with boundary. Then the number $|\Pi_\Sigma(n)|$ verifies, for $n \rightarrow \infty$*

$$|\Pi_\Sigma(n)| \leq \frac{c(\Sigma)}{\Gamma(-3/2\chi(\Sigma) + \beta(\Sigma))} n^{-3/2\chi(\Sigma) + \beta(\Sigma) - 1} 4^n (1 + O(n^{-1/2})), \tag{11}$$

where $c(\Sigma)$ is a function depending only on Σ .

Proof. We have seen in Proposition 3.2.2 that in fact $|\Pi_\Sigma(n)| \leq |\mathcal{R}_\Sigma(n)|$, as each bipartite subdivision can be reduced to a regular bipartite subdivision. We partition the set $\mathcal{R}_\Sigma(n)$ using the notion of irreducibility in the form

$$\mathcal{R}_\Sigma(n) = \mathcal{P}_\Sigma(n) \cup (\mathcal{R}_\Sigma(n) \setminus \mathcal{P}_\Sigma(n)).$$

Estimates for $|\mathcal{P}_\Sigma(n)|$ are obtained in Lemma 4.3.1, getting the bound stated in Eq. (6). In Proposition 4.4.3 we prove that $|\mathcal{R}_\Sigma(n) \setminus \mathcal{P}_\Sigma(n)| = o(|\mathcal{P}_\Sigma(n)|)$, hence the estimate in Eq. (11) holds. \square

5. Bounding $c(\Sigma)$ in terms of cubic maps

In this section we obtain upper bounds for $c(\Sigma)$ by doing a more refined analysis over functions $g_\varepsilon(z)$ (recall the notation used in Section 4.3). This is done in the following proposition.

Lemma 5.0.2. *The function $c(\Sigma)$ defined in Lemma 4.3.1 satisfies*

$$c(\Sigma) \leq 2^{\beta(\Sigma)} |\mathcal{C}_\Sigma|. \tag{12}$$

Proof. For each $\varepsilon \in \mathcal{C}_\Sigma$, we obtain bounds for $g_\varepsilon(1/4)$. We use Table 4, which is a simplification of Table 3. Now we are only concerned about the constant term on each GF. Table 4 brings the following information: the main contribution from double trees, trees, and families of pointed trees comes from $\mathcal{T}_{\square-\blacksquare}$, \mathcal{T} , and \mathcal{T}^\bullet , respectively. The constants are $1/4$, 2 , and 4 , respectively. Each cubic map has $-3\chi(\Sigma) + 2\beta(\Sigma)$ edges ($\beta(\Sigma)$ of them being roots) and $-2\chi(\Sigma) + \beta(\Sigma)$ vertices ($\beta(\Sigma)$ of them being incident with roots). This characterization provides the following upper bound for $g_\varepsilon(1/4)$:

$$g_\varepsilon(1/4) \leq \left(\frac{1}{4}\right)^{2\beta(\Sigma) - 3\chi(\Sigma) - \beta(\Sigma)} 2^{-3 \cdot 2\chi(\Sigma) + \beta(\Sigma)} 4^{\beta(\Sigma)} = 2^{\beta(\Sigma)}. \quad \square \tag{13}$$

The value of \mathcal{C}_Σ can be bounded using the results in [1,12]. Indeed, Gao shows in [12] that the number of rooted cubic maps with n vertices in an orientable surface of genus¹ g is asymptotically equal to

$$t_g \cdot n^{5(g-1)/2} \cdot (12\sqrt{3})^n,$$

where the constant t_g tends to zero as g tends to infinity [1]. A similar result is also stated in [12] for non-orientable surfaces. By duality, the number of rooted cubic maps on a surface $\bar{\Sigma}$ of genus $g(\Sigma)$ with $\beta(\Sigma)$ faces is asymptotically equal to $t_{g(\Sigma)} \cdot \beta(\Sigma)^{5(g(\Sigma)-1)/2} \cdot (12\sqrt{3})^{\beta(\Sigma)}$.

To conclude, we observe that the elements of \mathcal{C}_Σ are obtained from rooted cubic maps with $\beta(\Sigma)$ faces by adding a root on each face different from the root face. Observe that each edge is incident with at most two faces, and that the total number of edges is $-3\chi(\Sigma)$. Consequently, the number of ways of rooting a cubic map with $\beta(\Sigma) - 1$ unrooted faces is bounded by $\binom{-6\chi(\Sigma)}{\beta(\Sigma)-1}$.

Lemma 5.0.2, together with the discussion above, yields the following bound for $c(\Sigma)$.

¹ The genus $g(\Sigma)$ of an orientable surface Σ is defined as $g(\Sigma) = 1 - \chi(\Sigma)/2$ (see [15]).




Fig. 10. Two different representations of the same partition.

Proposition 5.0.3. *The constant $c(\Sigma)$ verifies, for $-\chi(\Sigma) \rightarrow \infty$*

$$c(\Sigma) < t_{1-\chi(\Sigma)/2} \cdot \beta(\Sigma)^{-5\chi(\Sigma)/2} \cdot (12\sqrt{3})^{\beta(\Sigma)} \cdot \binom{-6\chi(\Sigma)}{\beta(\Sigma) - 1} \cdot 2^{\beta(\Sigma)}.$$

Further research. In this article, we provided upper bounds for $|\Pi_{\Sigma}(n)|$. This upper bound is exact for the exponential growth (recall Section 1). However, we cannot assure exactness for the subexponential growth. The main problem in order to state asymptotic equalities is that $|\Pi_{\Sigma}(n)| \neq |\mathcal{P}_{\Sigma}(n)|$: there exist different irreducible bipartite subdivisions with n vertices which define the same non-crossing partition (see Fig. 10 for an example). However, we conjecture that the upper bound we have obtained is tight.

A natural strategy to find lower bounds with the same subexponential growth could be to consider only a subset of irreducible subdivisions giving rise to different non-crossing partitions. On surfaces obtained from the sphere this subset might be defined as the irreducible subdivisions such that there are at least two black vertices on each edge of the corresponding scheme (intuitively, such irreducible subdivisions do not allow the rotational symmetry around a boundary that occurs in the example of Fig. 10). However, this condition seems not to be sufficient on surfaces of higher genus. Hence, an open problem in this context is finding precise definitions of subfamilies of irreducible subdivision whose enumeration matches our asymptotic upper bound.

Another interesting problem is based on generalizing the notion of k -triangulation to the partition framework and getting the asymptotic enumeration: the enumeration of k -triangulations on a disk was found using algebraic methods in [13]. This notion can be easily translated to the non-crossing partition framework on a disk, and the *exact* enumeration in this case seems to be more involved. In the same way as non-crossing partitions on surfaces play a crucial role for designing algorithms for graphs on surfaces (see [16]), it turns out that the enumeration mentioned above is of capital importance in order to design algorithm for families of graphs defined by excluding minors.

Acknowledgments

The authors would like to thank Marc Noy for pointing them to Refs. [1,12], and the anonymous referees for suggesting interesting improvements in the presentation of the results.

The first author is supported by the European Research Council under the European Community's 7th Framework Programme, ERC grant agreement no. 208471—ExploreMaps project. The second author is supported by French projects AGAPE (ANR-09-BLAN-0159) and GRATOS (ANR-09-JCJC-0041-01). The third author was co-financed by the European Union (European Social Fund—ESF) and Greek national funds through the Operational Program “Education and Lifelong Learning” of the National Strategic Reference Framework (NSRF)—Research Funding Program: “Thales. Investing in knowledge society through the European Social Fund”.

References

- [1] E.A. Bender, Z. Gao, L.B. Richmond, The map asymptotics constant t_g , *The Electronic Journal of Combinatorics* 15 (1) (2008) R51. p. 8.
- [2] O. Bernardi, J. Rué, Enumerating simplicial decompositions of surfaces with boundaries, *European Journal of Combinatorics* 33 (3) (2012) 302–325.
- [3] G. Chapuy, Asymptotic enumeration of constellations and related families of maps on orientable surfaces, *Combinatorics, Probability and Computing* 18 (4) (2009) 477–516.
- [4] G. Chapuy, M. Marcus, G. Schaeffer, A bijection for rooted maps on orientable surfaces, *SIAM Journal on Discrete Mathematics* 23 (3) (2009) 1587–1611.
- [5] R. Cori, Indecomposable permutations, hypermaps and labeled dyck paths, *Journal of Combinatorial Theory, Series A* 116 (8) (2009) 1326–1343.
- [6] F. Dorn, F.V. Fomin, D.M. Thilikos, Fast subexponential algorithm for non-local problems on graphs of bounded genus, in: *Proc. of the 10th Scandinavian Workshop on Algorithm Theory, SWAT*, in: LNCS, vol. 4059, 2006, pp. 172–183.
- [7] F. Dorn, E. Penninx, H.L. Bodlaender, F.V. Fomin, Efficient exact algorithms on planar graphs: exploiting sphere cut branch decompositions, in: *Proc. of the 13th Annual European Symposium on Algorithms, ESA*, in: LNCS, vol. 3669, 2005, pp. 95–106.
- [8] P.H. Edelman, R. Simion, Chains in the lattice of noncrossing partitions, *Discrete Mathematics* 126 (1–3) (1994) 107–119.

- [9] P. Flajolet, M. Noy, Analytic combinatorics of non-crossing configurations, *Discrete Mathematics* 204 (1) (1999) 203–229.
- [10] P. Flajolet, A. Odlyzko, Singularity analysis of generating functions, *SIAM Journal on Discrete Mathematics* 3 (2) (1990) 216–240.
- [11] P. Flajolet, R. Sedgewick, *Analytic Combinatorics*, Cambridge University Press, 2008.
- [12] Z. Gao, The number of rooted triangular maps on a surface, *Journal of Combinatorial Theory, Series B* 52 (2) (1991) 236–249.
- [13] J. Jonsson, Generalized triangulations and diagonal-free subsets of stack polyominoes, *Journal of Combinatorial Theory, Series A* 112 (1) (2005) 117–142.
- [14] S.K. Lando, A.K. Zvonkin, *Graphs on Surfaces and their Applications*, in: *Encyclopaedia of Mathematical Sciences: Lower-Dimensional Topology II*, vol. 141, Springer-Verlag, 2004.
- [15] B. Mohar, C. Thomassen, *Graphs on Surfaces*, John Hopkins University Press, 2001.
- [16] J. Rué, I. Sau, D.M. Thilikos, Dynamic programming for graphs on surfaces, *ACM Transactions on Algorithms* (in press). Available online at: <http://arxiv.org/abs/1104.2486>.
- [17] P. Seymour, R. Thomas, Call routing and the ratcatcher, *Combinatorica* 14 (2) (1994) 217–241.
- [18] R. Simion, Combinatorial statistics on noncrossing partitions, *Journal of Combinatorial Theory, Series A* 66 (2) (1994) 270–301.
- [19] R.P. Stanley, *Enumerative Combinatorics, Vol. 2*, in: *Cambridge Studies in Advanced Mathematics*, vol. 62, Cambridge University Press, Cambridge, 1999.