
Volume 78 Issue 5 September 2012 ISSN 0022-0000
ELSEVIER
JOURNAL OF COMPUTER AND SYSTEM SCIENCES
JCSS Special Issue: Cloud Computing 2011
Guest Editors: Jinjun Chen and Lizhe Wang
JCSS Special Issue: Learning Theory 2009
Guest Editor: Alexander Rakhlin
Managing Editor
E. K. Blum
Editors
Alfred V. Aho
Jin-Yi Cai
Ronald Fagin
J. Hartmanis
John Hopcroft
R. M. Karp
D. E. Knuth
Raymond E. Miller
Christos Papadimitriou
Michael Rabin
Hartley Rogers, Jr.
Dana Scott
J. D. Ullman
L. A. Zadeh
Available online at www.sciencedirect.com
SciVerse ScienceDirect

Journal of Computer and System Sciences 78 (2012) 1606–1622

Contents lists available at SciVerse ScienceDirect

Journal of Computer and System Sciences

www.elsevier.com/locate/jcss

Catalan structures and dynamic programming in H-minor-free graphs ☆

Frederic Dorn [b,*], Fedor V. Fomin [a], Dimitrios M. Thilikos [c]

[a] *Department of Informatics, University of Bergen, Bergen, Norway*
[b] *SINTEF Energy Research, Trondheim, Norway*
[c] *Department of Mathematics, National & Kapodistrian University of Athens, Greece*

ARTICLE INFO

Article history:
Received 5 November 2009
Received in revised form 19 January 2012
Accepted 24 February 2012
Available online 1 March 2012

Keywords:
Parameterized complexity
Longest path
Minor-free graphs
Catalan structure

ABSTRACT

We give an algorithm that, for a fixed graph H and integer k, decides whether an n-vertex H-minor-free graph G contains a path of length k in $2^{O(\sqrt{k})} \cdot n^{O(1)}$ steps. Our approach builds on a combination of Demaine–Hajiaghayi's bounds on the size of an excluded grid in such graphs with a novel combinatorial result on certain branch decompositions of H-minor-free graphs. This result is used to bound the number of ways vertex disjoint paths can be routed through the separators of such decompositions. The proof is based on several structural theorems from the Graph Minors series of Robertson and Seymour. With a slight modification, similar combinatorial and algorithmic results can be derived for many other problems. Our approach can be viewed as a general framework for obtaining time $2^{O(\sqrt{k})} \cdot n^{O(1)}$ algorithms on H-minor-free graph classes.

1. Introduction

One of the motivations of this paper was the seminal result of Alon, Yuster, and Zwick in [3] that proved that a path of length $\log n$ can be found in polynomial time, answering to a question by Papadimitriou and Yannakakis in [32]. One of the open questions left in [3] was: "*Is there a polynomial time (deterministic or randomized) algorithm for deciding if a given graph G contains a path of length, say,* $\log^2 n$?". Of course, a $2^{O(\sqrt{k})} \cdot n^{O(1)}$ step algorithm for checking if a graph contains a path of length k would resolve this question. However, an algorithm of running time $2^{o(k)} \cdot n^{O(1)}$ for this problem, even for sparse graphs, would contradict the widely believed exponential time hypothesis, i.e. would imply that 3-SAT can be solved in subexponential time [25]. In this paper, we devise a $2^{O(\sqrt{k})} \cdot n^{O(1)}$ step algorithm for this problem on H-minor-free graphs, implying a polynomial-time algorithm for a $\log^2 n$-length path. This result is tight, because, according to Deĭneko, Klinz, and Woeginger [8], the existence of a $2^{o(\sqrt{k})} \cdot n^{O(1)}$ step algorithm, even for planar graphs, would again violate the exponential time hypothesis.

Our work is also motivated by the paradigm of parameterized algorithms [21,22,31]. A common technique in parameterized algorithms for problems asking for the existence of vertex/edge subsets of size k with certain properties is based on branchwidth (treewidth) and involves the following two ingredients: The first is a combinatorial proof that if the branchwidth of the input graph is at least $f(k)$ (where f is some function of k) then the answer to the problem is directly implied. The second is a $g(\mathbf{bw}(G)) \cdot n^{O(1)}$ step dynamic programming algorithm for the problem (here $\mathbf{bw}(G)$ is the branchwidth of the input graph G). For obtaining a $2^{O(\sqrt{k})} \cdot n^{O(1)}$ step algorithm out of this, we further require that

☆ An extended abstract of this paper appeared in SODA'08 (Dorn et al., 2008 [18]). The research of the last author has been supported by the project "Kapodistrias" (ΑΠ 736/24.3.2006 and 02839/28.07.2008) of the National and Kapodistrian University of Athens (project code: 70/4/8757).

* Corresponding author.
E-mail addresses: Frederic.Dorn@sintef.no (F. Dorn), fomin@ii.uib.no (F.V. Fomin), sedthilk@math.uoa.gr (D.M. Thilikos).

0022-0000/$ – see front matter © 2012 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcss.2012.02.004

(a) $f(k) = O(\sqrt{k})$, and
(b) $g(k) = 2^{O(\mathbf{bw}(G))}$.

For planar graphs (and also for H-minor-free graphs or apex-minor-free graphs – see [13] and [9]) **(a)** can be proved systematically using the idea of Bidimensionality [12,19]. However, not an equally general theory exists for **(b)**. On the positive side, **(b)** holds for several combinatorial problems. Typical problems in NP that fall in this category are Vertex Cover, Dominating Set or Edge Dominating Set, where no global conditions are imposed on their certificates [1,2,10,23]. This implies that the existence of such a set of size $\log^2 n$ can be decided in polynomial time and this answers positively the analogue of the question in [3] for these problems on H-minor-free graphs. The bad news is that, for many combinatorial problems, a general algorithm for proving **(b)** is missing. Longest path is a typical example of such a problem. Here the certificate of a solution should satisfy a global connectivity requirement. For this reason, the dynamic programming algorithm must keep track of all the ways the required path may traverse the corresponding separator of the decomposition, that is $\Omega(\ell^{\ell})$ on the size ℓ of the separator and therefore of treewidth/branchwidth. The same problem in designing dynamic programming algorithms appears for many other combinatorial problems in NP whose solution certificates are restricted by global properties such as connectivity. Other examples of such problems are Longest Cycle, Connected Dominating Set, Feedback Vertex Set, Hamiltonian Cycle and Graph Metric Traveling Salesman Problem (TSP).

Recently, [20] overcame the above deadlock for the class of planar graphs. Later, a similar result was given in [17] for graphs of bounded genus. The proofs in [20,17] are heavily based on arguments about non-crossing paths in graphs embedded in topological surfaces. This makes it possible to construct special types of graph decompositions of the input graph where the number of ways a path (or a cycle) traverses a separator of the decomposition is linearly bounded by the Catalan number of the separator size (which yields the desired single exponential dependance on treewidth or branchwidth). We refer to Stanley's book [40] for more information and different applications of Catalan numbers. It is not clear, a priori, whether this type of arguments can be extended to graphs excluding a minor. Another example of a technique, where the extension from planar and bounded genus to H-minor-free graphs is not clear, is given in [27,14].

In this paper, we provide a general framework for the design of dynamic programming algorithms on H-minor-free graphs. For this, it is necessary to go through the entire characterization of H-minor-free graphs given by Robertson and Seymour in their Graph Minors project (in particular, in [37]) to prove counting lemmata that can suitably bound the amount of information required in each step of the dynamic programming algorithm.

The main combinatorial result of this paper is Theorem 2, concerning the existence of suitably structured branch decompositions of H-minor-free graphs. While the grid excluding part follows directly from [13], the construction of the branch decomposition of Theorem 2 is quite involved. Indeed it uses the fact, proven by Robertson and Seymour in [37], that any H-minor-free graph can roughly be obtained by identifying in a tree-like way small cliques of a collection of components that are almost embeddable on bounded genus surfaces. The main proof idea is based on a procedure of "almost"-planarizing the components of this collection. However, we require a planarizing set with certain topological properties, able to reduce the high genus embeddings to planar ones where the planarizing vertices are cyclically arranged in the plane. This makes it possible to use a special type of planar branch decomposition, invented in [38], that permits to view collections of paths that may pass through a separator as non-crossing pairings of the vertices of a cycle. This provides the so-called *Catalan structure* of the decomposition and permits us to suitably bound the ways a path may cross its separators. Let us remark that similar ideas were also used in parameterized and approximation algorithms on planar graphs [8,11,27]. This decomposition is used to build a decomposition on the initial almost embeddable graph. Then using the tree-like way these components are linked together, we build a branch decomposition of the entire graph. The most technical part of the proof is to show that each step of this construction, from the almost planar case to the entire graph, maintains the Catalan structure, yielding the claimed upper bound.

Almost immediately, Theorem 2 implies the main algorithmic result of this paper. If a graph G on n vertices contains a $(\sqrt{k} \times \sqrt{k})$-grid, then G has a path of length k. Otherwise, by Theorem 2, it has a branch decomposition of width $O(\sqrt{k})$ with the Catalan structure. By standard dynamic programming on this branch decomposition (e.g., see [6]) we find the longest path in G. We stress that the dynamic programming algorithm is not different from the standard one. It is the special branch decomposition of Theorem 2 that accelerates its running time because the number of states at each step of the dynamic programming is bounded by $2^{O(\sqrt{k})}$. Thus the running time of the algorithm is $2^{O(\sqrt{k})} \cdot n^{O(1)}$.

Organization of the paper. First, we give some preliminaries in Section 2, where we restate the structural theorem on H-minor-free graphs for our purposes. In Section 3, we state our main theorem and give the algorithm for computing branch decompositions with Catalan structure, which we employ in Section 5 for solving the problem of finding a path of length k in H-minor-free graphs. In Section 4 we provide technical details of the correctness proof of the algorithm. We end (Section 6) with concluding remarks and open problems.

2. Preliminaries

Surface embeddable graphs. We refer to the book of Diestel [16] for the basic Graph Theory terminology and to the book of Mohar and Thomassen [29] on basics of Topological Graph Theory. We use the notation $V(G)$ and $E(G)$, for the set of the vertices and edges of G. A *surface* Σ is a compact 2-manifold without boundary (we always consider connected surfaces).

A *line* in Σ is a subset homeomorphic to $[0,1]$. An O-arc is a subset of Σ homeomorphic to a circle. Whenever we refer to a *Σ-embedded graph* G we consider a 2-cell embedding of G in Σ. To simplify notations we do not distinguish between a vertex of G and the point of Σ used in the drawing to represent the vertex or between an edge and the line representing it. We also consider G as the union of the points corresponding to its vertices and edges. That way, a subgraph H of G can be seen as a graph H where $H \subseteq G$. We call a *region* of G any connected component of $(\Sigma \setminus E(G)) \setminus V(G)$. (Every region is an open disk.) A subset of Σ meeting the drawing only in vertices of G is called *G-normal*. If an O-arc is G-normal, then we call it *noose*. The length of a noose N is the number of its vertices and we denote it by $|N|$. If the intersection of a noose with any region results into a connected subset of the surface, then we call such a noose *tight*. In that case the subset is a line in Σ. Let Δ be a closed disk and the open disk $\mathbf{int}(\Delta)$ its interior and $\mathbf{bor}(\Delta)$ its boundary. Then $\Delta = \mathbf{int}(\Delta) \cup \mathbf{bor}(\Delta)$. We say two closed disks Δ_1, Δ_2 *touch* if $\Delta_1 \cap \Delta_2 \neq \emptyset$ but $\mathbf{int}(\Delta_1) \cap \mathbf{int}(\Delta_2) = \emptyset$.

Surface cutting. We need to define the graph obtained by *cutting along* a non-contractible tight noose N. We suppose that for any $v \in N \cap V(G)$, there exists an open disk Δ containing v and such that for every edge e adjacent to v, $e \cap \Delta$ is connected. We also assume that $\Delta \setminus N$ has two connected components Δ_1 and Δ_2. Thus we can define a partition of $N(v) = N_1(v) \cup N_2(v)$, where $N_1(v) = \{u \in N(v)\colon \{u,v\} \cap \Delta_1 \neq \emptyset\}$ and $N_2(v) = \{u \in N(v)\colon \{u,v\} \cap \Delta_2 \neq \emptyset\}$. For each $v \in N \cap V(G)$ we *duplicate* v: (a) remove v and its incident edges, (b) introduce two new vertices v^1, v^2, and (c) connect v^i with the vertices in N_i, $i = 1, 2$. v^1 and v^2 are vertices of the new G-normal O-arcs N_1 and N_2 that meet the border Δ_1 and Δ_2, respectively. We say that the vertices v, v^1, v^2 are *relatives* while, after any further cutting, the relation of being "relative" is inherited to new vertices that may occur by splitting v^1 or v^2. We call N_1 and N_2 *cut-nooses*. We can see the operation of "cutting G along a non-contractible noose N" as "sawing" the surface where G is embedded. This helps us to embed the resulting graph to the surface(s) that result after adding to the sawed surface two disks, one for each side of the splitting. We call these disks *holes* and we will treat them as closed disks. Clearly, in the new embedding(s) the duplicated vertices will all lay on the borders of these holes.

Branch and linear decompositions. Let G be a graph and let $E \subseteq E(G)$. We define ∂E as the set of vertices in G that are incident to edges in both E and $E(G) - E$. We call the pair (T, τ) *branch decomposition* of G if T is a ternary tree and τ is a bijection mapping the edges of G to the leaves of T. For each edge of T we define $\omega(e)$ as the vertex set ∂F_e where F_e are all the preimages of the leaves of one of the connected components of $T - e$. The *width* of a branch decomposition is the maximum $|\omega(e)|$ over all edges of T. The *branchwidth* of a graph is the minimum width over all branch decompositions of G.

By modifying the definition of a branch decomposition such that T is a caterpillar, which is a path with attached pendant edges, we obtain a linear variant of branchwidth. The occurring width parameter is essentially the same (and actually the same for values bigger than one) as the notion of a *linear-width* of a graph and τ maps the edges of G to the edges of T. Let $e_1, \ldots, e_q$ be an ordering of edges of graph G, the sets $X_i = \partial(\{e_1, \ldots, e_i\})$ form a linear ordering $\mathcal{X} = (X_1, \ldots, X_q)$, called *linear decomposition*. For convenience, we will use such ordered sets to denote linear decompositions and, in order to include all vertices of G in the sets of $\mathcal{X}$, we will often consider linear decompositions of $\hat{G}$ that is G with loops added to all its vertices (this operation does not increase the linear-width by more than one). Sets X_i correspond to sets $\omega(e)$, $e \in E(T)$.

Sphere cut decompositions. For a graph G embedded in the sphere, we define a *sphere cut decomposition* or *sc-decomposition* (T, τ, π) as a branch decomposition such that for every edge e of T and E_e^1 and E_e^2, the two sets of preimages of the two components of $T - e$, there exists a tight noose N bounding two open disks Δ_1 and Δ_2 such that $E_e^i \subseteq \Delta_i \cup N$, $1 \leqslant i \leqslant 2$. Thus N meets G only in ∂E_e and its length is $|\partial E_e|$. Clockwise traversal of N in the drawing of G defines the cyclic order π of ∂E_e. We always assume that in an sc-decomposition the vertices of $\partial E_e = E_e^1 \cap E_e^2$ are enumerated according to π. According to the celebrated ratcatcher algorithm, due to Seymour and Thomas [38] (improved by [24]), there is a time $O(n^3)$ algorithm finding an optimal branch decomposition of a planar graph. At the same time, shown by [20], this is also an optimal sc-decomposition.

A cornerstone theorem of Graph Minors. We say that H is a *minor* of G if H is obtained from a subgraph of G by contracting edges. Let be given two graphs G_1 and G_2 and two h-cliques $S_i \subseteq V(G_i)$ $(i = 1, 2)$. We obtain graph G by identifying S_1 and S_2 and deleting none, some or all clique-edges. Then, G is called the *h-clique-sum* of the *clique-sum components* G_1 and G_2. Note that the clique-sum gives many graphs as output depending on the edges of the clique that are deleted. According to Lemma 21 (proved in Section 4), given a graph G with branch-decomposition (T, τ), for any clique with vertex set S there exists a node $t \in T$ such that $S \subseteq \omega(\{t,a\}) \cup \omega(\{t,b\}) \cup \omega(\{t,c\})$ where a, b, c are the neighbors of t in T. We call such a vertex of T a *clique node* of S.

Let Σ be a surface. We denote as Σ^{-r} the subspace of Σ obtained if we remove from Σ the interiors of r disjoint closed disks (we will call them *vortex disks*). Clearly, the boundary $\mathbf{bor}(\Sigma^{-r})$ of Σ^{-r} is the union of r disjoint cycles. We say that G is *h-almost embeddable* in Σ if there exists a set $A \subseteq V(G)$ of vertices, called *apices of* G, where $|A| \leqslant h$ and such that $G - A$ is isomorphic to $G_u \cup G_1 \cup \cdots \cup G_r$, $r \leqslant h$, in a way that the following conditions are satisfied (the definition below is not the original one from [37] but equivalent, slightly adapted for the purposes of our paper):

- There exists an embedding $\sigma : G_u \to \Sigma^{-r}$, $r \leqslant h$, such that only vertices of G_u are mapped to points of the boundary of Σ^{-r}, i.e. $\sigma(G_u) \cap \mathbf{bor}(\Sigma^{-r}) \subseteq V(G)$ (we call G_u the *underlying graph* of G).
- The graphs $G_1, \ldots, G_r$ are pairwise disjoint (called *vortices* of G). Moreover, for $i = 1, \ldots, r$, if $E_i = E(G_i) \cap E(G_u)$ and $B_i = V(G_i) \cap V(G_u)$ (we call B_i *base set* of the vortex G_i and its vertices are the *base vertices* of G_i), then $E_i = \emptyset$ and $\sigma(B_i) \subseteq C_i$ where $C_1, \ldots, C_r$ are the cycles of $\mathbf{bor}(\Sigma^{-r})$.
- For every $i = 1, \ldots, r$, there is a linear decomposition $\mathcal{X}_i = (X^i_1, \ldots, X^i_{q_i})$ of the vortex G_i with width at most h and a subset $J_i = \{j^i_1, \ldots, j^i_{|B_i|}\} \subseteq \{1, \ldots, q_i\}$ such that $\forall_{k=1,\ldots,|B_i|}\ u^i_k \in X^i_{j^i_k}$ for some respectful ordering $(u^i_1, \ldots, u^i_{|B_i|})$ of B_i. (An ordering $(u^i_1, \ldots, u^i_{|B_i|})$ is called *respectful* if the ordering $(\sigma(u^i_1), \ldots, \sigma(u^i_{|B_i|}))$ follows the cyclic ordering of the corresponding cycle of $\mathbf{bor}(\Sigma^{-r})$.) For every vertex $u^i_k \in B_i$, we call $X^i_{j^i_k}$ the *overlying set* of u^i_j and we denote it by $\mathbf{X}(u^i_j)$.

If in the above definition $A = \emptyset$, then we say that G is *smoothly h-almost embeddable* in Σ. Moreover, if $r = 0$, then we just say that G is *embeddable* in Σ. Note that the underlying graph G_u of a smoothly h-almost embeddable graph G in Σ is itself embeddable in Σ.

For reasons of uniformity, we will extend the notion of the overlying set of a vertex in B_i to any other vertex v of the underlying graph G_u by defining its *overlying set* as the set consisting only of v. For any $U \subseteq V(G_u)$, the *overlying set* of U is defined by the union of the overlying sets of all vertices in U and it is denoted as $\mathbf{X}(U)$.

We will strongly use the following structural theorem of Robertson and Seymour (see [37],) characterizing H-minor-free graphs.

Theorem 1. *(See [37].) Let $\mathcal{G}$ be the graph class not containing a graph H as a minor. Then there exists a constant h, depending only on H, such that any graph $G \in \mathcal{G}$ is the* (*repeated*) *h-clique-sum of h-almost embeddable graphs* (*we call them* clique-sum components) *in a surface Σ of genus at most h.*

That is, beginning with an h-almost embeddable graph G, we repeatedly construct the h-clique-sum of G with another h-almost embeddable graph. As a result we obtain a tree-like structure of h-almost embeddable graphs connected by h-clique-sums. It was proved in [26] that such a decomposition can be computed in $O_H(n^{O(1)})$ steps.

Path collections. Let G be a graph and let $E \subseteq E(G)$ and $S \subseteq V(G)$. We will consider collections of internally vertex disjoint paths using edges from E and having their endpoints in S. We use the notation $\mathbf{P}$ to denote such a path collection and we define $\mathbf{paths}_G(E, S)$ as the collection of all such path collections. Define the equivalence relation $\sim$ on $\mathbf{paths}_G(E, S)$: for $\mathbf{P}_1, \mathbf{P}_2 \in \mathbf{paths}_G(E, S)$, $\mathbf{P}_1 \sim \mathbf{P}_2$ if there is a bijection between $\mathbf{P}_1$ and $\mathbf{P}_2$ such that corresponding paths have the same endpoints. We denote by $\texttt{q-paths}_G(E, S) = |\mathbf{paths}_G(E, S)/\sim|$, i.e. the cardinality of the quotient set of $\mathbf{paths}_G(E, S)$ by $\sim$.

Let G be a planar graph, S be a cyclic separator, i.e. the set of vertices lying on a noose N meeting G only in S, and E be the set of edges inside one of the disks bounded by N. Then every equivalence class $\mathbf{paths}_G(E, S)/\sim$ corresponds to a set of non-crossing matchings, which is a set of internally disjoint arcs drawn inside the disc and connecting vertices of S lying on the border of the disk. In this case the number $\texttt{q-paths}_G(E, S)$ is at most

$$\mathrm{CN}\left(\frac{n}{2}\right) \sim \frac{2^n}{\sqrt{\pi}(\frac{n}{2})^{\frac{3}{2}}} \approx 2^n,$$

where $n = |S|$ and $\mathrm{CN}(n)$ is the n-th Catalan number: $\mathrm{CN}(n) = \frac{1}{n+1}\binom{2n}{n} \sim \frac{4^n}{\sqrt{\pi}n^{\frac{3}{2}}} \approx 4^n$.

3. Main result and the algorithm

Before we state our main result, we need some notation especially for the context of our algorithm. Given a graph H and a function f we use the notation $O_H(f)$ to denote $O(f)$ while emphasizing that the hidden constants in the big-O notation depend exclusively on the size of H. We also define analogously the notation $\Omega_H(f)$.

Given a graph G and a branch decomposition (T, τ) of G we say that (T, τ) has the *Catalan structure* if

$$\text{for any edge } e \in E(T), \quad \texttt{q-paths}(E_e, \partial E_e) = 2^{O_H(|\partial E_e|)}.$$

Our main result is the following.

Theorem 2. *For any H-minor-free graph class $\mathcal{G}$, the following holds*:

For every graph $G \in \mathcal{G}$ and any positive integer w, there is an algorithm that in time $O_H(n^{O(1)})$ outputs one of the following:

1. *A correct report that G contains a $(w \times w)$-grid as a minor.*
2. *A branch decomposition (T, τ) with the Catalan structure of width $O_H(w)$.*

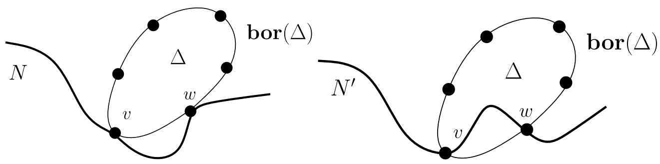

Fig. 1. Re-routing a noose.

The proof of Theorem 2 is algorithmic. In the remaining part of this section we provide the description of the algorithm used to prove Theorem 2. The correctness of the algorithm is given in Section 4. While the first statement of the theorem follows almost directly from [15], our main contribution is the proof of statement 2 and, more concretely, the existence of a branch decomposition of this particular type.

The algorithm starts with decomposing the graph into a collection of almost embeddable graphs whose clique-sum decomposition forms the input graph G. For every such almost embeddable graph we show how to "planarize" it by removing some vertices and cutting along some curves in the surface such that the planar branch decomposition of the resulting graph can be transformed into a branch decomposition with Catalan structure of the almost embeddable graph. This part (Step 2) is the most technical part of our paper. Finally (Step 3), we glue together the computed branch decompositions of the almost embeddable components to form a branch decomposition with Catalan structure of G. The detailed description of the algorithm follows.

Step 1. Use the time $O_H(n^{O(1)})$ algorithm of [26], see also [7,15], to decompose the input graph into a collection $\mathcal{C}$ of clique-sum components as in Theorem 1. Every graph in $\mathcal{C}$ is a γ_H-almost embeddable graph to some surface of genus $\leqslant \gamma_H$ where $\gamma_H = O_H(1)$.

Step 2. For every $G^a \in \mathcal{C}$, do

Step 2.a. Let G^s be the graph G^a without the apex vertices A (i.e. G^s is smoothly γ_H-almost embeddable in a surface of genus γ_H). Denote by G^s_u the underlying graph of G^s.

Step 2.b. Set $G^{(1)}_u \leftarrow G^s_u$, $G^{(1)} \leftarrow G^s$, and $i \leftarrow 1$.

Step 2.c. Apply the following steps as long as $G^{(i)}_u$ is non-planar.

Step 2.c.i. Find a non-contractible noose N in $G^{(i)}_u$ of minimum length, using the polynomial time algorithm in [41].

Step 2.c.ii. If $|N| \geqslant 2^{i-1} f(H) w$, then output "*The input graph contains a* $(w \times w)$*-grid as a minor*" and stop. The estimation of the value $f(H)$ comes from the results in [13]. The correctness of this step is given in Lemmata 4 and 5.

We will see in Lemma 6 that N is tight, which means that N cannot intersect the interior of a hole or a vortex disc Δ more than once. In Lemma 7 we show that N intersects $\mathbf{bor}(\Delta)$ in at most two vertices. If $\mathbf{int}(\Delta) \cap N = \emptyset$ and $\mathbf{bor}(\Delta) \cap N = \{v, w\}$, then v and w are both incident to one more region different than Δ which is intersected by N, as we will show in Lemma 8. In this case, we re-route this portion of N so that it crosses the interior of Δ (see Fig. 1).

Step 2.c.iii. As long as N intersects some hole (initially the graph $G^{(1)}$ does not contain holes but they will appear later in $G^{(i)}$'s for $i \geqslant 2$) or some vortex disc of $G^{(i)}_u$ in *only* one vertex v, update $G^{(i)}$ by removing v and the overlying set of all its relatives (including $\mathbf{X}(v)$) from $G^{(1)}, \ldots, G^{(i)}$. To maintain the $O_H(1)$-almost embeddibility of G^a, compensate this loss of vertices in the initial graph $G^s = G^{(1)}$, by moving in A the overlying set of the relatives of v in $G^{(1)}$ (as the number of vortex disks and holes depends only on H, the updated apex set has again size depending on H). Notice that after this update, all cut-nooses found so far, either remain intact or they become smaller. The disks and vortices in $G^{(1)}, \ldots, G^{(i)}$ may also be updated as before and can only become smaller. We observe that after this step, if a noose N intersects a hole or vortex disc Δ it also intersects its interior and therefore it will split Δ into two parts Δ_1 and Δ_2.

Step 2.c.iv. We cut $G^{(i)}_u$ along N and call the two disks created by the corresponding cut of the surface by *holes* of the new embedding. We go through the same cut in order to "saw" $G^{(i)}$ along N as follows: If the base set of a vortex is crossed by N then we also split the vortex according to the two sides of the noose; this creates two vortices in $G^{(i+1)}$. For this, consider a vortex G^v and a linear decomposition $\mathcal{X} = (X_1, \ldots, X_q)$ of G^v. Let also a, b be the vertices of the base set B of G^v that are intersected by N and let $a \in X_{j_a}$, $b \in X_{j_b}$, where w.l.o.g. we assume that $a < b$. When we split, the one vortex is the subgraph of G^v induced by $X_1 \cup \cdots \cup X_{j_a} \cup X_{j_b} \cup \cdots \cup X_q$ the other is the subgraph of G^v induced by $X_{j_a} \cup \cdots \cup X_{j_b}$ (notice that the vertices that are duplicated are those in X_{j_a} and X_{j_b}). Let $G^{(i+1)}$ be the graph embedding that is created that way and let $G^{(i+1)}_u$ be its underlying graph. Recall that, from the previous steps, a vortex disc or a hole Δ (if divided)

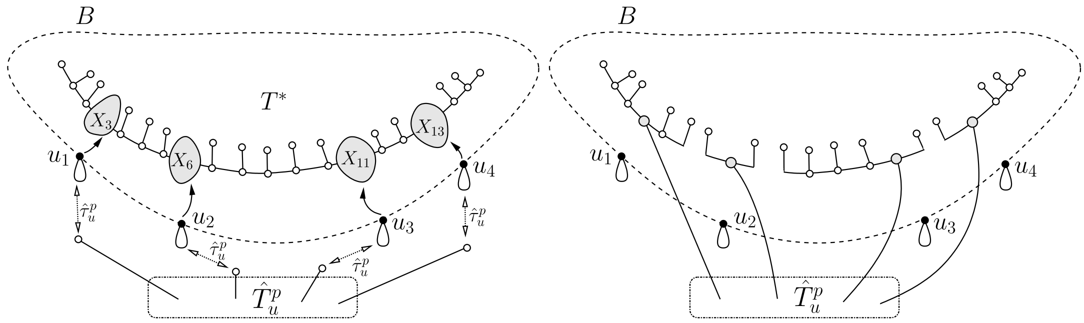

Fig. 2. The procedure of enhancing the branch decomposition (T_u^p, τ_u^p) of G_u^p to a branch decomposition of G^p.

is divided into two parts Δ_1 and Δ_2 by N. That way, the splitting of a vortex in $G^{(i)}$ creates two vortices in $G^{(i+1)}$. As the number of vortices in $G^{(i)}$ is $O_H(1)$, the same holds also for the number of vortices in $G^{(i+1)}$. If N splits a hole of $G^{(i+1)}$, then the two new holes Δ'_1, Δ'_2, that the splitting creates in $G^{(i+1)}$, are augmented by the two parts Δ_1 and Δ_2 of the old hole Δ (i.e. $\Delta'_j \leftarrow \Delta_j \cup \Delta'_j$, $j = 1, 2$).

Step 2.c.v. $i \leftarrow i + 1$. The loop of Step **2.c** ends up with a planar graph $G_u^{(i)}$ after $O_H(1)$ splittings because the genus of $G_u^{(1)}$ is $O_H(1)$ (each step creates a graph of smaller Euler genus – see [29, Propositions 4.2.1 and 4.2.4]). This implies that the number of holes or vortex disks in each $G_u^{(i)}$ remains $O_H(1)$. Therefore, $G^{(i)}$ is a smoothly $O_H(1)$-almost embeddable graph in the sphere. The total length of the holes of $G_u^{(i)}$ is at most the sum of the lengths of the nooses we cut along, which is $O_H(w)$.

Step 2.d. Set $G^p \leftarrow G^{(i)}$ and $G_u^p \leftarrow G_u^{(i)}$ and compute an optimal sphere-cut branch decomposition (T_u^p, τ_u^p) of G_u^p, using the polynomial algorithm from [38].

Step 2.e. If $\mathbf{bw}(G_u^p) \geqslant 2^{\gamma_H - 1} \cdot f(H) \cdot w = \Omega_H(w)$, then output "*The input graph contains a* $(w \times w)$*-grid as a minor*" and stop. The correctness of this step is given in Lemma 9.

Step 2.f. Enhance (T_u^p, τ_u^p), so that the edges of the vortices of G^p are included in it, as follows: Let G^v be a vortex of G^p with base set $B = \{u_1, \ldots, u_m\}$ ordered in a respectful way such that $\forall_{k=1,\ldots,m}\ u_k \in \omega(f_{j_k})$ where the ordering $f_{j_1}, \ldots, f_{j_m}$ contains the edges of a longest path of the tree T^* of some linear decomposition (T^*, τ^*) of G^v. Update (T_u^p, τ_u^p) to a branch decomposition $(\hat{T}_u^p, \hat{\tau}_u^p)$ of $\hat{G}^p$ (if the branchwidth of a graph is more than 1, it does not change when we add loops). Let $l_1, \ldots, l_m$ be the leaves of $(\hat{T}_u^p, \hat{\tau}_u^p)$ corresponding to the loops of the base vertices of G^v. We subdivide each f_{j_k} in T^* and we identify the subdivision vertex with l_k for any $k = 1, \ldots, m$. We transform the resulting graph into a ternary tree, by removing a minimum number of edges in T^* and desolving their endpoints in the resulting forest. That way, we construct a branch decomposition of $\hat{G}_u^p \cup G^v$ which, after discarding the leaves mapped to loops, gives a branch decomposition of $G_u^p \cup G^v$ (see Fig. 2).

Applying this transformation for each vortex G^v of G^p, we construct a branch decomposition of G^p. Moreover, the width of this decomposition is $O_H(\mathbf{bw}(G_u^p))$. The correctness of this construction and the proof that $\mathbf{bw}(G^p) = O_H(\mathbf{bw}(G_u^p))$ is given in Lemma 10 in Section 4.2.

Step 2.g. Notice that, while successively splitting G^s during the loop of Step **2.c**, all edges remain topologically intact and only vertices may be duplicated. This establishes a bijection between $E(G^s)$ and $E(G^p)$, which allows us to transform (T^p, τ^p) to a branch decomposition (T^s, τ^s) where $T^s = T^p$. The branch decomposition (T^s, τ^s) of $G^s = G^{(1)}$ which is a smoothly $O_H(1)$-almost embeddable in a surface of higher genus, is a decomposition with the Catalan structure of width $O_H(w)$. The proof of this claim is non-trivial, and we provide it in Section 4.3.

Step 2.h. Construct a branch decomposition (T^a, τ^a) of G^a by adding in (T^s, τ^s) the edges incident to the apices of G^a. To do this, for every apex vertex a and for every neighbor v, choose an arbitrary edge e of T^s, such that $v \in \partial E_e$. Subdivide e and add a new edge to the new node and set $\tau(\{a, v\})$ to be the new leaf. The proof that $\mathbf{bw}(G^a) = O_H(\mathbf{bw}(G^s))$ is easy, as G^s contains only $O_H(1)$ more vertices (Lemma 12 in Section 4.3). With more effort (and for the same reason) we prove that the Catalan structure for G^s implies the Catalan structure for G^a (Lemma 20 in Section 4.5).

Step 3. For any $G^a \in \mathcal{C}$, merge the branch decompositions constructed above according to the way they are joined by clique sums and output the resulting branch decomposition of the input graph G. In particular, if (T_1^a, τ_1^a) and (T_2^a, τ_2^a) are two branch decompositions of two graphs G_1^a and G_2^a with cliques S_1 and S_2 respectively and $|S_1| = |S_2|$, we construct a branch-decomposition (T', τ') of the graph G', taken after a clique sum of G_1^a and G_2^a, as follows: Let t^i be a clique-node of S_i in (T_i, τ_i), $i = 1, 2$. Then, the branch decomposition (T', τ') of G' is obtained by first subdividing an incident edge e_{t^i}, $i = 1, 2$, and then connecting the new nodes together. Secondly, remove each leaf l of T' that corresponds to an edge that has a parallel edge or is deleted in the clique-sum operation and finally contract an incident edge in T' of each degree-two node. We prove (Lemma 24 in Section 4.6) that this merging does not harm neither the bounds for branchwidth nor the Catalan structure of the obtained branch decomposition and this finally holds for the input graph G, justifying Theorem 2.

4. Correctness of the algorithm

4.1. Correctness of Steps **2.c.ii** and **2.e**: exit conditions for the algorithm

In this subsection, we give a lower bound on the branchwidth of the input graph G, that is, we give two exit conditions on which the algorithm terminates and fulfills the first part of Theorem 2, namely to give a certificate that G has large branchwidth. Representativity [35] is a measure how densely a non-planar graph is embedded on a surface. The *representativity* (or *face-width*) $\mathbf{rep}(G)$ of a graph G embedded in surface $\Sigma \neq \mathbb{S}_0$ is the smallest length of a non-contractible noose in Σ. The following lemma follows from Theorem 4.1 of [36].

Lemma 3. *Let G be a graph embedded on a surface which is not a sphere. Then* $\mathbf{rep}(G) \leqslant \mathbf{bw}(G)$.

The function $f(H)$ in Step **2.c.ii** of the algorithm is defined by the following lemma that follows from [13, Lemmata 5, 6, and 7].

Lemma 4. *Let G be an H-minor-free graph and let G_u^s be as in Step* **2.a***. Then, there exists a function $f(H)$ such that if* $\mathbf{bw}(G_u^s) \geqslant f(H) \cdot w$*, then G contains a $(w \times w)$-grid as a minor.*

The following lemma justifies the first terminating condition for the algorithm, depending on the value of $f(H)$ estimated in Lemma 4.

Lemma 5. *If in the x-th application of Step* **2.c.ii** $|N| \geqslant 2^{x-1} f(H) \cdot w$*, then $G_u^s = G_u^{(1)}$ has branchwidth at least $f(H) \cdot w$.*

Proof. Let $N_1, \ldots, N_{x-1}$ be the nooses along which we cut the graphs $G_u^{(1)}, \ldots, G_u^{(x-1)}$ in Step **2.c.iv** towards creating $G_u^{(x)}$. We have that

$$\sum_{j=1,\ldots,x-1} |N_j| \leqslant \sum_{j=1,\ldots,x-1} 2^{j-1} f(H) \cdot w = \left(2^{x-1} - 1\right) f(H) \cdot w.$$

We also observe that G_u^{j-1} contains as a subgraph the graph taken from G_u^j if we remove one copy of each of its $|N_{j-1}|$ duplicated vertices. This implies that

$$\mathbf{bw}\left(G_u^{j-1}\right) \geqslant \mathbf{bw}\left(G_u^j\right) - |N_{j-1}|, \quad j = 2, \ldots, x.$$

Inductively, we have

$$\mathbf{bw}\left(G_u^1\right) \geqslant \mathbf{bw}\left(G_u^x\right) - \sum_{j=1,\ldots,x-1} |N_j| \geqslant \mathbf{bw}\left(G_u^x\right) - \left(2^{x-1} - 1\right) f(H) \cdot w.$$

We set $N = N_x$. By Lemma 3, $\mathbf{bw}(G_u^x) \geqslant \mathbf{rep}(G_u^x)$ and $\mathbf{rep}(G_u^x) \geqslant |N_x| \geqslant 2^{x-1} f(H) \cdot w$. Thus, we conclude that $\mathbf{bw}(G_u^1) \geqslant f(H) \cdot w$. $\square$

The following lemmata help in finding non-contractible nooses eligible for planarizing the embeddable graph.

Lemma 6. *Let N be a non-contractible noose of mimimum length $(\geqslant 3)$ in an embeddable graph. Then N is tight.*

Proof. This is a proof by contradiction on the choice of minimality of the non-contractible noose. Let us assume N is not tight. Then there exists a region Δ where we have that $\Delta \cap N$ is a union of more than one connected subset of the surface. Let us pick two such subsets S_1 and S_2 such that $\Delta \setminus (S_1 \cup S_2)$ contains at least one component Δ_S with $\mathbf{int}(\Delta_S) \cap N = \emptyset$. Next we add a line L to $\mathbf{int}(\Delta_S)$ connecting one point of $\mathbf{bor}(\Delta) \cap S_1$ with one point of $\mathbf{bor}(\Delta) \cap S_2$. L always separates N into two lines L_1 and L_2 and thus forms two news nooses $N_1 := L \cup L_1$ and $N_2 := L \cup L_2$, respectively. We will now argue using the 3-path-condition (see for example [29]) that one of the nooses N_1 and N_2 is a non-contractible noose which is shorter than N. The 3-path-condition applies to cycles in embeddable graphs. The property of interest for us is that: let three internally vertex-disjoint paths P_1, P_2, P_3 have common two endvertices. If any two cycles formed by $P_i \cup P_j$ $(1 \leqslant i < j \leqslant 3)$ are contractible so is the third cycle. The 3-path-condition also applies to non-contractible nooses.[1] Since N is non-contractible, we have according to the 3-path-condition that at least one of N_1 and N_2 must be non-contractible, too, and we get a contradiction on the minimality assumption. □

Lemma 7. *Let N be a non-contractible noose of mimimal length in an embeddable graph G and Δ a region of G. Then N intersects $\mathbf{bor}(\Delta)$ in at most two vertices.*

Proof. The proof idea is the same as in the proof of Lemma 6. Let us assume a region Δ with $\mathbf{bor}(\Delta) \cap N$ containing more than two points. Since N is tight, at least one of those points, say x, is not an endpoint of line $L := \mathbf{int}(\Delta) \cap N$ (in the case $L \neq \emptyset$). We can then again achieve a shorter non-contractible noose when connecting x to another point y of $\mathbf{bor}(\Delta) \cap N$ by a line through the connected surface component of $\mathbf{int}(\Delta) \setminus N$ to which both x and y are incident. □

Lemma 8. *Let N be a non-contractible noose of mimimal length in an embeddable graph G and Δ a region of G. If $\mathbf{int}(\Delta) \cap N = \emptyset$ and N intersects $\mathbf{bor}(\Delta)$ in two vertices u, v, then u, v are both incident to one more region different than Δ which is intersected by N.*

Proof. As before we will use the 3-path-condition to prove this statement. We simply draw a line L through $\mathbf{int}(\Delta)$ connecting the points u and v. Then if both u and v are incident only to Δ, we argue that L forms a shorter non-contractible noose and get a contradiction. □

The following lemma justifies the first terminating condition for the algorithm, depending on the value of $f(H)$ and the genus γ_H.

Lemma 9. *If in Step* **2.e** $\mathbf{bw}(G_u^p) \geqslant 2^{\gamma_H - 1} f(H) \cdot w$, *then* $G_u^s = G_u^{(1)}$ *has branchwidth at least* $f(H) \cdot w$.

Proof. The proof is the same as the proof of Lemma 5 if we set $x = i$ and with deference that, in the end, we directly have that $G_u^{(i)} = G_u^p$ has branchwidth at least $2^{i-1} f(H) \cdot w$. The result follows as the genus of G_u^s is bounded by γ_H and therefore $i \leqslant \gamma_H$. □

4.2. Correctness of Step **2.f**: enhancing the branch decomposition of G_u^p

We want to prove that the enhancement of the branch decomposition of G_u^p by the linear decompositions of the vortices performed in Step **2.f** can add $O_H(1)$ vertices for each vertex in $\omega(e)$, $e \in E(T_u^p)$, and therefore, $\mathbf{bw}(G^p) = O_H(\mathbf{bw}(G_u^p))$.

Lemma 10. *Let (T_u^p, τ_u^p) be a branch decomposition of G_u^p and let (T^p, τ^p) be the branch decomposition of G^p constructed in Step* **2.f**. *Then the width of (T^p, τ^p) is bounded by the width w of (T_u^p, τ_u^p) plus some constant that depends only on H.*

Proof. By the construction of (T^p, τ^p), for any $e \in E(T^p)$, we have that $\partial E_e(G^p) \subseteq \mathbf{X}(\partial E_e(G_u^p))$. As the vertices of $\partial E_e(G_u^p)$ are the vertices of some tight noose N_e of $\mathbb{S}_0$, and this noose meets at most $r \leqslant h$ vortex disks we have that there are at most $2r \leqslant 2h$ vertices of $\partial E_e(G_u^p)$ that are members of some base sets B. Therefore, for any $e \in E(T^p)$, $|\partial E_e(G^p)| \leqslant w + 2h^2$. We conclude that the width of (T^p, τ^p) is at most $w + 2h^2$. □

4.3. Correctness of Step **2.g**: towards Catalan structure

The next two subsections are devoted to the proof that the branch decomposition constructed for a smoothly almost-embeddable graph has the Catalan structure. We start with an overview of the proof and then proceed with a sequence of lemmata.

We want to prove that if the bounds of Theorem 2 hold for the graph $G^p = G^{(i)}$ (a graph that is smoothly $O_H(1)$-almost embeddable in the sphere), then they also hold for the graph $G^s = G^{(1)}$ that is smoothly $O_H(1)$-almost embeddable in a surface of higher genus. We prove that $\mathbf{bw}(G^s) = O_H(\mathbf{bw}(G^p))$ with the help of Lemma 11.a. However, what is far more complicated is to prove that (T^s, τ^s) has the Catalan structure. For this, we first prove (Lemma 11.b) that for any edge e of $T^s = T^p$, it holds that

$$\texttt{q-paths}_{G^s}(E_e, \partial E_e) \leqslant \texttt{q-paths}_{G^p}(E_e, \partial E_e \cup D_e), \tag{1}$$

[1] In our graph, by definition, we add an edge for every line segment of N intersecting a region of the graph. We also add an edge for L. Then we have that N is a cycle as well as N_1 and N_2.

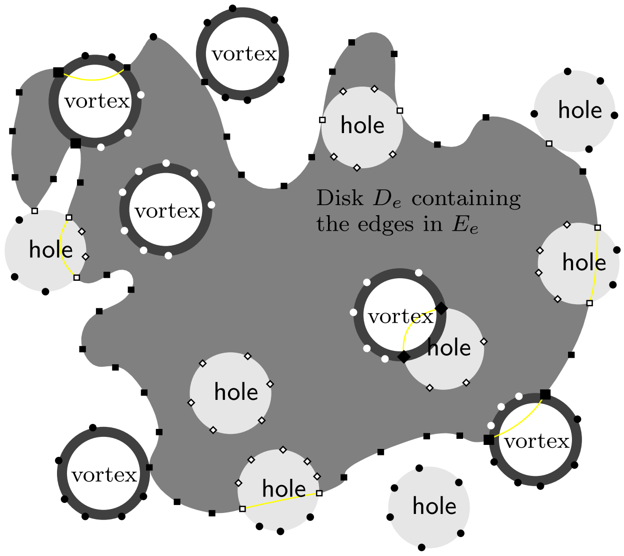

Fig. 3. Vortices and holes around the disc Δ_e.

where D_e is the set of all vertices of the holes of G^p that are endpoints of edges in E_e. Intuitively, while splitting the graph G^s along non-contractible nooses, the split vertices in the nooses (i.e., the vertices in D_e) may separate paths counted in the left side of Eq. (1). Therefore, in order to count them, we have to count equivalence classes of collections of internally vertex disjoint paths in the planar case allowing their endpoints to be not only in ∂E_e but also in D_e. That way, we reduce the problem of proving that (T^s, τ^s) has the Catalan structure, to the following problem: to bound the number of equivalence classes of collections of vertex disjoint paths whose endpoints may be:

a) Vertices of the disc Δ_e bounding the edges E_e in the sphere-cut decomposition (T^p, τ^p) of G_u^p (see Lemma 17 in Section 4.4), along with their overlying sets. By Lemma 10, the number of such vertices is $\mathbf{bw}(G^p) = O_H(w)$.
b) Vertices (and their overlying sets) of $2 \cdot \gamma_H = O_H(1)$ disjoint holes created by cutting along nooses. We prove in Lemma 18 in Section 4.4 (and by making use of Lemma 5) that the total amount of such vertices is at most $\gamma_H \cdot (2^{\gamma_H} - 1) \cdot f(H) \cdot w = O_H(w)$.

Notice that these paths can be routed also via at most $\gamma_H \cdot 2^i \leqslant \gamma_H \cdot 2^{\gamma_H} = O_H(1)$ vortices (because, initially, G_u^s had γ_H vortices and, in the worst case, each noose can split every vortex into two parts), each of unbounded size (see Lemma 19 in Section 4.4). Recall that the holes and the vortex disks of G_u^p do not touch (recall the notion *touch*) because of the simplification in Step **2.c.iii**, however, they may have common interiors. Finally, the boundary of Δ_e can touch any number of times a vortex disc or a hole but can traverse it only once (recall that by the definition of sc-decompositions $\mathbf{bor}(\Delta_e)$ should be a tight noose). (For an example of the situation of the holes and vortices around the disc bounding the edges E_e, see Fig. 3.) Now our target is to relate $\texttt{q-paths}_{G^p}(E_e, \partial E_e \cup D_e)$ to the classical Catalan structure of non-crossing partitions on a cycle. This is done in Lemma 12. The proof is technical and is performed in several steps in Section 4.4. The first two steps are to "force" holes and vortex disks not to touch the boundary of Δ_e and to "force" vortex disks not to intersect with holes or with $\mathbf{bor}(\Delta_e)$. For each of these two steps, we bound $\texttt{q-paths}_{G^p}(E_e, \partial E_e \cup D_e)$ by its counterpart in a "normalized" instance of the same counting problem (related to the original one by a "rooted minor" relation). That way, the problem is reduced to counting equivalence classes of collections of vertex disjoint paths with endpoints (recall that there are $O_H(w)$ such endpoints) on the boundary of $O_H(1)$ disjoint holes (the disc taken if we remove from Δ_e all holes that intersect it, is also considered as one of these holes). However, we still do not have to count equivalence classes of non-crossing collections of paths because of the presence of the vortices that may permit crossing paths. At that point, we prove that no more than $O_H(\beta)$ paths can mutually cross, where $\beta = O_H(1)$ is the number of vortices. Using this observation, we prove that each equivalence class is the superposition of $O_H(1)$ equivalence classes of non-crossing collections of disjoint paths. Because of this, the number of equivalence classes of collections of disjoint paths are in total $2^{O_H(w)}$ and, that way, we bound $\texttt{q-paths}_{G^p}(E_e, \partial E_e \cup D_e)$ as required.

With the following lemma, we can inductively show how to get a branch decomposition with the Catalan structure of a smoothly almost-embeddable graph G_u^s embedded on a higher surface when starting with a branch decomposition of its planarized version G_u^p.

Lemma 11. *Let G, G' be two almost embeddable graphs created successively during Step* **2.c** ($i \geqslant 2$)*, let N be the noose along which G_u was cut towards constructing G_u' and G', and let N_1 and N_2 be the boundaries of the two holes of G_u' created after this splitting during Step* **2.c.iv**. *Let also (T', τ') be a branch decomposition of G' and let (T, τ) be the branch decomposition of G defined as*

$T = T'$ and $\tau = \tau' \circ \sigma$, where $\sigma : E(T) \to E(T')$ is the bijection between topologically equivalent edges in G and G'. Then for every $e \in E(T) = E(T')$, the following hold:

a. $|\omega_G(e)| \leqslant |\omega_{G'}(\sigma(e))| + |\mathbf{X}(N)|$.
b. *For $S \subseteq V(E_e)$, the number $|\mathbf{paths}_G(E_e, \partial E_e \cup S)|$ is at most*

$$\left|\mathbf{paths}_{G'}\left(\sigma(E_e), \partial\sigma(E_e) \cup S \cup \mathbf{X}\left(N_1 \cap V\left(\sigma(E_e)\right)\right) \cup \mathbf{X}\left(N_2 \cap V\left(\sigma(E_e)\right)\right)\right)\right|.$$

Proof. To see $|\omega_G(e)| \leqslant |\omega_{G'}(e)| + |\mathbf{X}(N)|$, it is enough to observe that the identification of vertices in a graph may only add identified vertices in the border of an edge set and that any overlying set is a separator of vortices.

The second relation follows from the fact that any path in $G[E_e]$ connecting endpoints in $\partial E_e \cup S$ is the concatenation of a set of paths in $G'[E_e]$ connecting endpoints in $\partial E_e \cup S$ but also the vertices of cut-noose N_1, N_2 that are endpoints of E_e (along with their overlying sets). $\square$

Section 4.4 is devoted to the proof of the following lemma:

Lemma 12. *Let G^p be a smoothly $O_H(1)$-almost embeddable graph in the sphere and let $\Delta_1, \ldots, \Delta_r$ ($r = O_H(1)$) be disjoint closed disks (holes) of the sphere whose interiors do not intersect the underlying graph G^p_u. Assume also that, if a vortex disc and a hole intersect, then they have common interior points. Let Δ_e be a closed disc of the sphere whose boundary is a tight noose touching G^p_u in vertex set ∂E_e. We denote as D_e the set containing all points on the boundary of the disks $\Delta_1, \ldots, \Delta_r$ that are endpoints of edges in $G^p_u \cap \Delta_e$ and as E_e the set of edges in $G^p \cap \Delta_e$. In G^p, let $\mathbf{X}(\partial E_e \cup D_e)$ be the overlying set of vertex set $\partial E_e \cup D_e$. If $|\mathbf{X}(\partial E_e \cup D_e)| = O_H(w)$, then*

$$\texttt{q-paths}_{G^p}\left(E_e, \mathbf{X}(\partial E_e \cup D_e)\right) = 2^{O_H(w)}.$$

4.4. Proof of Lemma 12: fixing paths in smoothly h-almost embeddable graphs on the sphere

In the following and for an easier estimation on $\texttt{q-paths}_{G^p}(E_e, \mathbf{X}(\partial E_e \cup D_e))$, we stepwise transform the graph in a way such that neither of the holes, the vortices and ∂E_e mutually intersect, while simultaneously nondecreasing the number of sets of paths.

Inverse edge contractions. From now on we will use the notation $\mathbf{V_e}$ for the vertex set $\mathbf{X}(\partial E_e \cup D_e)$.

The operation of *inverse edge contraction* is defined by duplicating a vertex v and connecting it to its duplicate v' by a new edge. However, we have that v maintains all its incident edges.

If two closed disks Δ_1 and Δ_2 touch in vertices, we call these vertices *touching vertices* of the closed disks Δ_1 and Δ_2.

In order to simplify the structure of the planar embedding of G^p we will apply a series of inverse edge contractions to the touching vertices between the boundary of Δ_e and the vortex disks and holes.

Also, we assume that if we apply inverse edge contraction on a base vertex v of a vortex, v keeps all its incident edges and the duplicate of the respective boundary of a hole and of Δ_e has degree one. This creates a new graph G^{p*} that contains G^p as a minor and thus, each set of paths in G^p corresponds to a set of paths in G^{p*}.

We obtain the following:

Lemma 13. *Let $E_e, \mathbf{V_e}$ be as above. Then,*

$$\texttt{q-paths}_{G^p}(E_e, \mathbf{V_e}) \leqslant \texttt{q-paths}_{G^{p*}}\left(E^*_e, \mathbf{V}^*_\mathbf{e}\right),$$

where E^*_e and $\mathbf{V_e}^*$ are the enhanced sets in G^{p*}.

The red (in the web version) lines in the diagram in Fig. 4 emphasize inverse edge contractions.

Notice that each splitting creates duplicates of some vertex of $\mathbf{V_e}$. Therefore,

Lemma 14. $|\mathbf{V_e}^*| \leqslant 2|\mathbf{V_e}|$.

On the left of Fig. 5, we have the now resulting graph of Fig. 4 where the gray part is $G^{p*}[E_e]$. On the right, we only emphasize $G^{p*}[E_e]$ as the part where the sets of $\mathbf{paths}_{G^{p*}}(E^*_e, \mathbf{V_e}^*)/\sim$ should be drawn.

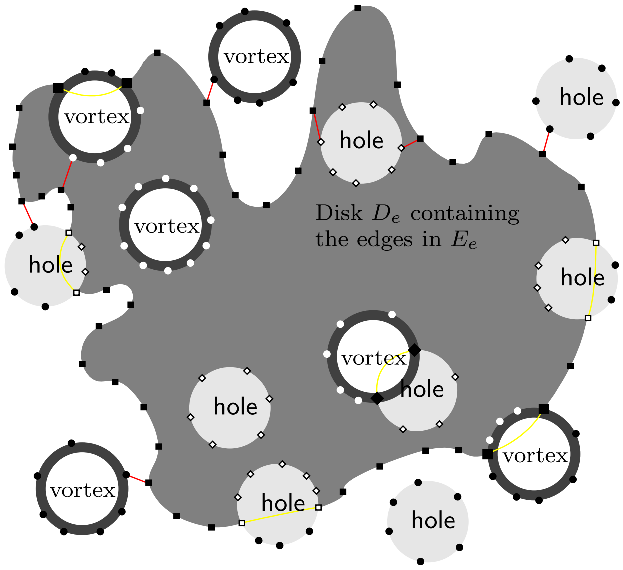

Fig. 4. Vortices and holes not intersecting the disc Δ_e (first normalization).

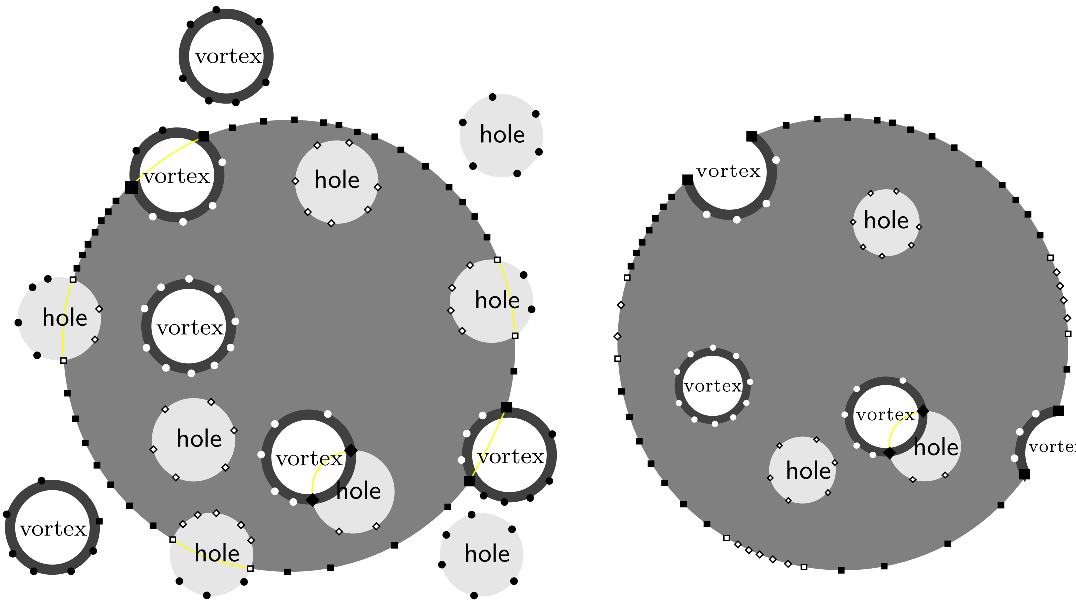

Fig. 5. Another way to see Fig. 4.

Vortex pattern. In order to have a more uniform image on how paths cross a vortex, we define the graph $R_{h,s}$ so that

$$V(R_{h,s}) = V_1 \cup \cdots \cup V_s \quad \text{for mutually disjoint vertex sets } V_i \text{ with } |V_i| = h \quad \text{and}$$

$$E(R_{h,s}) = \bigl\{\{x_j, x_k\} \bigm| x_j, x_k \in V_i,\ 1 \leqslant j \neq k \leqslant h,\ 1 \leqslant i \leqslant s\bigr\} \cup \bigl\{\{x_j, y_j\} \bigm| x_j \in V_{i-1} \wedge y_j \in V_i,\ 1 \leqslant j \leqslant h,\ 2 \leqslant i \leqslant s\bigr\}.$$

In $R_{h,s}$ we also distinguish a subset $S \subseteq V(R_{h,s})$ containing exactly one vertex from any V_i. We call the pair $(R_{h,s}, S)$ an (h,s)-vortex pattern. See Fig. 6 for an example of a normalized vortex.

We now prove the following:

Lemma 15. *Any vortex of an h-almost embeddable graph with base set B is a minor of an (h,s)-vortex pattern $(R_{h,s}, S)$ where the minor operations map bijectively the vertices of S to the vertices in B in a way that the order of the vortex and the cyclic ordering of S induced by the indices of its elements is respected. Moreover, $R_{h,s}$ has linear-width at most h.*

Proof. We show how any vortex with base set B and linear decomposition $\mathcal{X} = (X_1, \ldots, X_{|B|})$ of width $< h$ is a minor of some (h,s)-vortex pattern $(R_{h,s}, S)$ with $V(R_{h,s}) = V_1 \cup \cdots \cup V_s$ and $|V_i| = h$ $(1 \leqslant i \leqslant s)$. Choose $s = |B|$ and set $S = B$. Start with the vertices in X_1: set $V_1 = X_1$ plus some additional vertices to make $|V_1| = h$ and make $G[V_1]$ complete. Iteratively,

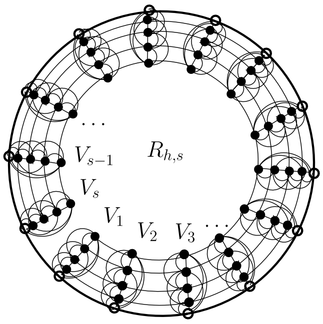

Fig. 6. Normalizing vortices. An example of a $(5, 14)$-vortex pattern.

set $V_i = X_i \setminus X_{i-1}$. Apply inverse edge contraction for all vertices in $X_i \cap X_{i-1}$ and add the new vertices to V_i. Again, add additional vertices to make $|V_i| = h$. Make $G[V_i]$ complete and add all missing edges between V_i and V_{i-1} in order to obtain a matching. $\square$

Using Lemma 15, we can replace G^p by a new graph $G^{p\prime}$ where any vortex of G^p is replaced by a suitably chosen $R_{h,s}$. The bijection of the lemma indicates where to stick the replacements to the underlying graph G^p_u. We adopt the same notions as for vortices. I.e, we denote S as base set consisting of base vertices, etc. In the remainder of the paper we will refer to h-almost embeddable graphs as graphs with (h, s)-vortex pattern instead of vortices, unless we clearly state differently and we will use the term vortex for (h, s)-vortex pattern eventually.

Normalizing vortices. We can now apply one more transformation in order to have both all vortex disks entirely inside $\mathbf{int}(\Delta_e)$ and no vortices intersecting holes. In fact, we can handle both cases in exactly the same way, because $\Sigma \setminus D_e$ can be treated as any hole Δ where ∂E_e corresponds to $\mathbf{bor}(\Delta)$. Thus, assume that for any (h, s)-vortex pattern we have that V_i and V_j are the two vertex sets of $R_{h,s} \cap \mathbf{bor}(\Delta)$, where Δ is one of the holes $\Delta_1, \dots, \Delta_r$. For every vertex in V_i and V_j, we create new vertex sets which we connect in the same way as in the construction of vortex patterns. Thus we create $2h - 2$ new vertex sets $V_i^1, \dots, V_i^{h-1}$ and $V_j^1, \dots, V_j^{h-1}$ to obtain a new $(h, s + 2h - 2)$-vortex pattern. We then apply inverse edge contraction on the base vertices of V_i and V_j and the new sets. This transformation is shown in the next figure, where contracting the red (in the web version) edges illustrated to the right leads to the minor illustrated to the left.

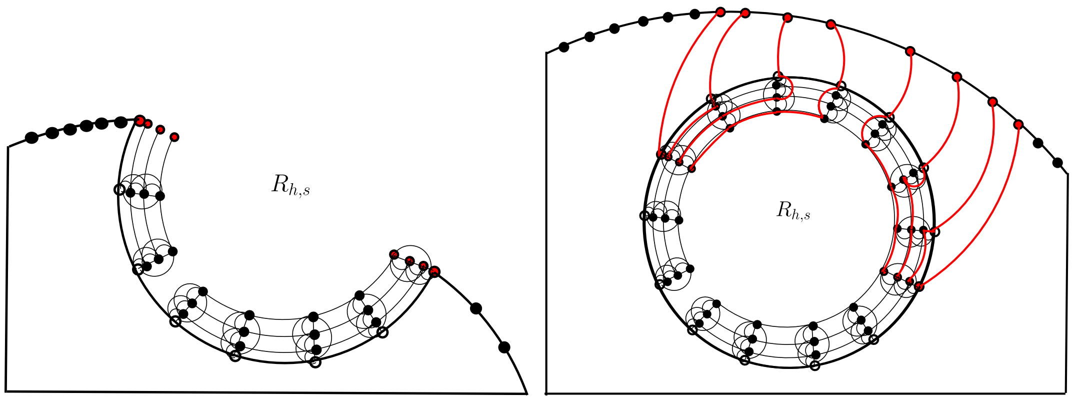

Again we can rename the graph before this new transformation G^p and the graph produced $G^{p\prime}$ and prove the following using the fact the later contains the former as a minor.

Lemma 16. *Let G^p and $G^{p\prime}$ be as above. Then*

$$\texttt{q-paths}_{G^p}(E_e, \mathbf{V_e}) \leqslant \texttt{q-paths}_{G^{p\prime}}\left(E'_e, \mathbf{V'_e}\right).$$

Final setting. To have an idea of how Δ_e looks like after the previous transformation, see the left part of the next figure. The outer face can now be seen as a hole and we redraw the whole embedding in a sphere as indicated by the second part of the same figure. We will denote by $\Delta_{\partial E_e}$ the new disk in the graph embedding that is bounded by the union ∂E_e and some intersecting holes of $\Delta_1, \ldots, \Delta_r$.

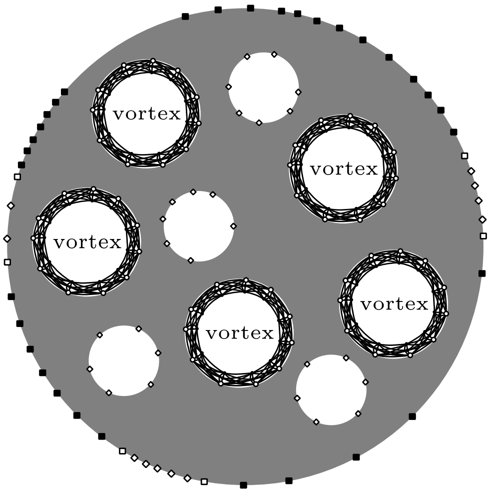

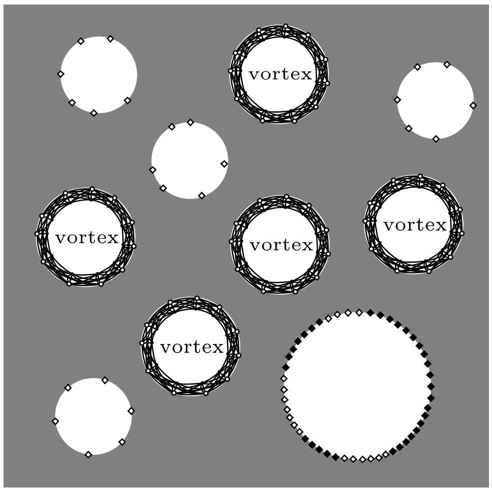

In the current setting, we have a collection of holes in the sphere with ℓ vertices on their borders. We will now count $\texttt{q-paths}_{G^{p'}}(E'_e, \mathbf{V'_e})$ for the sets of vertex disjoint paths $\mathbf{paths}_{G^{p'}}(E'_e, \mathbf{V_e}')$ between these ℓ vertices.

Tree structure for fixing paths. Before we are ready to prove Lemma 12, namely that

$$\texttt{q-paths}_{G^p}(E_e, \mathbf{V_e}) = 2^{O_H(w)},$$

we need some auxiliary lemmas:

Non-crossing matchings

Lemma 17. *(See [28].) Let $P(n) = \{P_1, \ldots, P_{\frac{n}{2}}\}$ be a partition of an ordered set $S = \{x_1, \ldots, x_n\}$ into tuples, such that there are no elements $x_i < x_j < x_k < x_\ell$ with $\{x_i, x_k\}$ and $\{x_j, x_\ell\}$ in $P(n)$. Let $\mathcal{P}_S$ be the collection of all such partitions of S. Then*

$$|\mathcal{P}_S| = O(2^n).$$

The partitions of Lemma 17 are called *non-crossing matchings*. A non-crossing matching can be visualized by placing n vertices on a cycle, and connecting matching vertices by non-crossing arcs at one side of the cycle. In a graph G, each element P of $\mathbf{paths}_G(E, S)/\sim$ can be seen as a set of arcs with endpoints in S. If every P is a non-crossing matching, we say that the paths in $P_i \in \mathbf{paths}_G(E, S)$ with $P \sim P_i$ are *non-crossing* and S has a Catalan structure.

(n, r)-non-crossing matchings. The next lemma gives an estimation on the non-crossing matchings arising from the case that one has several cycles in the plane that are connected by non-crossing arcs. After planarizing the graph of bounded genus by cutting along non-contractible nooses, the situation arises, that components and paths, respectively, have their endpoints in several such "cut-nooses".

Lemma 18. *Let r disjoint empty disks $\Delta_1, \ldots, \Delta_r$ be embedded on the sphere $\mathbb{S}_0$ where each disc is bounded by a cycle of at most n vertices. Let P be a set of arcs connecting the vertices, such that P can be embedded onto $\mathbb{S}_0 - \{\Delta_1, \ldots, \Delta_r\}$ without arcs crossing. Let $\mathcal{P}_{n,r}$ be the collection of all such P. Then,*

$$|\mathcal{P}_{n,r}| \leqslant r^{r-2} \cdot n^{2r} \cdot 2^{rn}.$$

Proof. We show how to reduce the counting of $|\mathcal{P}_{n,r}|$ to non-crossing matchings. Here we deal with several open disks and our intention is to transform them into one single disc in order to apply Lemma 17.

First let us assume $r = 2$. Choose a set $P \in \mathcal{P}_{n,2}$. Assume two vertices x and y on the boundary of two different disks Δ_1 and Δ_2 in being two endpoints of an arc $\{x, y\}$ in P. We observe that no other arc in P crosses $\{x, y\}$ in the $\mathbb{S}_0$-embedding of P. So we are able to 'cut' the sphere $\mathbb{S}_0$ along $\{x, y\}$ and, that way, create a "tunnel" between Δ_1 and Δ_2 unifying them to a single disc and thus reduced the problem to counting non-crossing matchings. That is, for obtaining a rough upper bound on $|\mathcal{P}_{n,2}|$, one fixes every pair of vertices $x \in \Delta_1$ and $y \in \Delta_2$ and we obtain

$$|\mathcal{P}_{n,2}| = O(n^2 \cdot 2^{2n}).$$

The next difficulty is that all disks in $\Delta_1, \ldots, \Delta_r$ are connected by arcs of $P \in \mathcal{P}_{n,r}$ in an arbitrary way. We use a tree structure in order to cut the sphere along that structure. Given such a tree structure, we create tunnels in order to connect the open disks and to merge them to one disk.

Consider all $\leqslant n^{n-2}$ possible spanning trees on n vertices [4]. Here, we have a spanning tree over r vertices, representing the r disks in $\Delta_1, \ldots, \Delta_r$. Then the boundary of each disc has length $\leqslant n$. Hence, there are $O(n^2)$ possible fixed arcs between the boundaries of each two disks. Then we obtain a rough upper bound of n^{2r} on the number of possible fixed arcs between the disks in a given tree-structure. We obtain $r^{r-2} \cdot n^{2r}$ possibilities for the above concatenation and tunneling of $\Delta_1, \ldots, \Delta_r$. We argue that P has a Catalan structure when tunneling the disks in this way. Thus,

$$|\mathcal{P}_{n,r}| \leqslant r^{r-2} \cdot n^{2r} \cdot 2^{rn}. \qquad \square$$

We call an element of $\mathcal{P}_{n,r}$ an (n,r)*-non-crossing matching*. If for a graph G, each element of $\mathbf{paths}_G(E,S)/\sim$ is an (n,r)-non-crossing matching then S has Catalan structure.

h-almost-(n,r)-non-crossing matchings. For obtaining a Catalan structure on H-minor-free graphs, we extend the case of bounded genus graphs with additional h "areas of non-planarity", the vortices.

Lemma 19. *Let r disjoint empty disks $\Delta_1, \ldots, \Delta_r$ be embedded on the sphere $\mathbb{S}_0$ where each disc is bounded by a cycle of at most n vertices. Let $\mathbb{S}_0 - \{\overline{\Delta}_1, \ldots, \overline{\Delta}_r\}$ contain $\leqslant h$ disjoint disks $R_1, \ldots, R_h$.*

Let P be a set of arcs connecting the vertices of $\mathbf{bor}(\Delta_1), \ldots, \mathbf{bor}(\Delta_r)$, such that

- *P can be embedded onto $\mathbb{S}_0 - \{\Delta_1, \ldots, \Delta_r\}$ with arcs crossing only inside $R_1, \ldots, R_h$,*
- *at most h arcs of $P \cap R_j$ cross mutually for every j $(1 \leqslant j \leqslant h)$.*

Then P is a superposition of $O((h+r)^h)$ many (n,r)-non-crossing matchings. Let $\mathcal{P}^h_{n,r}$ be the collection of all such P. Thus,

$$\left|\mathcal{P}^h_{n,r}\right| \leqslant \left(r^{r-2} \cdot n^{2r} \cdot 2^{rn}\right)^{[(h+r)^h]}.$$

Proof. We interpret $P \cap R_j$ as a circular arc graph C_j with maximum clique size h. That is, the arcs in $P \cap R_j$ form the vertex set of C_j and we connect two vertices by an edge if the corresponding arcs intersect. As there are only up to h arcs mutually intersecting, C_j may only contain cliques of size at most h. By [42] the chromatic number of a circular arc graph is bounded by twice the clique number, here $2h$.

But looking at an entire arc α of P connecting two disjoints disks Δ_i and Δ_j $(1 \leqslant i < j \leqslant r)$, α may intersect R_j arbitrarily often. Thus we may have more than h mutually crossings of arcs in P when intersecting R_j.

Together with the observation that α cuts $\mathbb{S}_0 - \{\Delta_1, \ldots, \Delta_r\}$ into several disks, we conclude that there are less than $2h + r - 1$ arcs that mutually cross in R_j. We color the arcs of P such that no two arcs of the same color class cross. For arcs crossing in one R_j we thus need up to $2h + r - 1$ colors.

Furthermore, we observe that two arcs of P may be assigned the same color in R_j but cross in another R_i, etc. Hence, we have a rough upper bound of $(h+r)^h$ colors, that is every arc can be assigned by $2h + r - 1$ colors per R_j and is thus assigned by an h-vector of colors for all R_j $(1 \leqslant j \leqslant h)$. With Lemma 18, we count for every color class the number of (n,r)-non-crossing matchings and we get that the overall size of $\mathcal{P}^h_{n,r}$ is bounded by $(r^{r-2} \cdot n^{2r} \cdot 2^{rn})^{[(h+r)^h]}$. $\square$

We can apply Lemma 19 to our terminology:

We say two paths $P_1, P_2 \in \mathbf{paths}_{G^p}(E_e, \mathbf{V_e})$ *cross inside an (h,s)-vortex pattern* $(R_{h,s}, S)$ if there is a vertex set $V_i \in V(R_{h,s,})$ that is used by P_1 and P_2.

Each element of $\mathcal{P}^{O_H(1)}_{O_H(w), O_H(1)}$ is an equivalence class of the paths in the set $\mathbf{paths}_{G^p}(E_e, \mathbf{V_e})/\sim$ with $O_H(w)$ endpoints in $\mathbf{V_e}$ crossing inside $O_H(1)$ vortex patterns. Thus, we have proven Lemma 12, namely that

$$\texttt{q-paths}_{G^p}(E_e, \mathbf{V_e}) = \left|\mathcal{P}^{O_H(1)}_{O_H(w), O_H(1)}\right| = 2^{O_H(w)}. \qquad \square$$

4.5. Proof of Step **2.h***: dealing with the apices*

So far, we considered smoothly h-almost Σ-embeddable graphs G^s without apices. To include the apices, we enhance the branch decomposition (T^s, τ^s) of G^s of width $O_H(w)$ so that each middle set contains at most all $O_H(1)$ apices. We construct an enhanced branch decomposition (T, τ) of an h-almost Σ-embeddable graph G^a as follows: For every apex vertex α and for every neighbor v, choose an arbitrary edge e of T^s, such that $v \in \partial E_e$. Subdivide e and add a new edge to the new node and set $\tau(\{\alpha, v\})$ to be the new leaf. In this way, the enhanced branch decomposition (T, τ) of G^a has width $O_H(w) + O_H(1)$. We obtain the following:

Lemma 20. *Let G be an h-almost embeddable graph and let G^a be the smoothly h-almost embeddable graph G^s obtained after removing the apices. Suppose also that G^s has a branch decomposition (T^s, τ^s) of width $O_H(w)$. Then the enhanced branch decomposition (T, τ) of G^a has width $O_H(w)$ and*

$$\texttt{q-paths}_{G^a}(E_e, \partial E_e) \leqslant w^{O_H(1)} \cdot \texttt{q-paths}_{G^s}\big(E^s_e, \partial E^s_e\big).$$

Proof. For an edge $e \in T$, we observe: Any path in $G^a[E_e]$ going through an apex vertex α one or two neighbors of α (depending whether a is an internal vertex of the path or not). If $G^a[E_e]$ does not contain any neighbor of any $\alpha \in G^a[E_e]$ then $\texttt{q-paths}_{G^a}(E_e, \partial E_e) \leqslant w^{O_H(1)} \cdot \texttt{q-paths}_{G^s}(E^s_e, \partial E^s_e)$. If $G^a[E_e]$ contains some neighbor, then α may be connected by a path to one or two vertices in ∂E^s_e. I.e., any apex vertex can only contribute to one path P in $G^a[E_e]$. Thus, for one α we count $O_H(w^2)$ different possible endpoint of P in ∂E^s_e, and for all $O_H(1)$ apices $w^{O_H(1)}$. $\square$

*4.6. Proof of Step **3**: dealing with the clique-sums*

Let be given a graph an H-minor-free graph G. By Theorem 1, G can be decomposed in a tree-like way into several h-almost embeddable graphs by reversing the clique-sum operation. That is, we obtain a collection $\mathcal{C} = \{G^a_1, \ldots, G^a_n\}$ with each G^a_i being h-almost embeddable graphs with up to n (possibly intersecting) h-cliques that contributed to the clique-sum operation.

Lemma 21. *Let G be a graph with a branch decomposition (T, τ). For any k-clique S in G, there are three adjacent edges e, f, g in T such that $S \subseteq \partial E_e \cup \partial E_f \cup \partial E_g$.*

Proof. Say, for node t incident to above e, f, g, $\partial E_t = \partial E_e \cup \partial E_f \cup \partial E_g$. We will prove the lemma inductively. Let a 3-clique consist of the vertices u, v, w. In (T, τ), we consider path $P_u \in T$ to be the path between the leaves $\tau(\{u, v\})$ and $\tau(\{u, w\})$. By definition, $u \in \partial E_e$ for all $e \in P_u$. Let node $t \in P_u$ be an endpoint of the path in $T \setminus P_u$ with other endpoint $\tau(\{v, w\})$. Then, $\{u, v, w\} \subseteq \partial E_t$. For an $i \leqslant k$, let $S_i \subset S$ be an i-clique for which there is a $t \in T$ with $S_i \subseteq \partial E_t$. Let $T_i \subseteq T$ be the tree induced by the paths between the leaves corresponding to the edges of S_i. Let $z \in S \setminus S_i$ and $T_z \subseteq T$ be the subtree induced by the paths connecting the leaves corresponding to edges between z and S_i. Then, we differ two cases: either $t \in T_i \cap T_z$ or there is a path in T_i connecting t and the closest node t_z in T_z. In the first case, under the assumption that $S_i \subseteq \partial E_t$ we obtain that $S_i \cup \{z\} \subseteq \partial E_t$. In the second case, since $S_i \subseteq \partial E_t$ and each vertex of S_i is an endpoint of some edge in a leaf of T_z, we get that $S_i \subseteq \partial E_{t_z}$. By definition $z \in \partial E_{t_z}$ and we are done. $\square$

We define the node t incident to above edges $e, f, g \in T$ as a *k-clique-node*.

Since for any edge $e \in T$ for a branch decomposition (T, τ), the vertex set ∂E_e separates the graph into two parts, we obtain the following lemma:

Lemma 22. *Let be given a graph G and a branch decomposition (T, τ). For any edge $e \in T$ if $\texttt{q-paths}_G(E_e, \partial E_e) \leqslant q$ and $\texttt{q-paths}_G(\overline{E_e}, \partial \overline{E_e}) \leqslant q$ then $\texttt{q-paths}_G(E(G), \partial E_e) \leqslant q^2$. For any three adjacent edges $e_1, e_2, e_3 \in T$, if $\texttt{q-paths}_G(E_{e_i}, \partial E_{e_i}) \leqslant q$ and $\texttt{q-paths}_G(\overline{E_{e_i}}, \partial \overline{E_{e_i}}) \leqslant q$ for $i = 1, 2, 3$ then $\texttt{q-paths}_G(E(G), \bigcup_{i=1,2,3} \partial E_{e_i}) \leqslant q^3$.*

We now show how to construct the branch decomposition of an h-clique-sum by connecting the branch decompositions of the two clique-sum components at some h-clique-nodes that correspond to the involved h-clique: Let G^a_1 and G^a_2 be the two clique-sum components with the cliques $S_i \subseteq V(G^a_i)$ $(i = 1, 2)$ together with the branch decompositions (T^a_i, τ^a_i) and an h-clique-node t^i. Then, the branch decomposition (T', τ') of the clique-sum G' is obtained by first subdividing an incident edge e_{t^i} and connecting the new nodes together. Secondly, remove each leaf l of T' that corresponds to an edge that has a parallel edge or is deleted in the clique-sum operation, and finally contract an incident edge in T' of each degree-two node.

Lemma 23. *Let G^a_1 and G^a_2 have branch decompositions (T^a_1, τ^a_1), (T^a_2, τ^a_2) with maximum width w and for all edges $e \in T'_1 \cup T'_2$ let $\texttt{q-paths}_{G'}(E_e, \partial E_e) \leqslant q$. The above construction of the branch decomposition (T', τ') of the h-clique-sum G' has width $\leqslant w + h$ and for all edges $e \in T'$ $\texttt{q-paths}_{G'}(E_e, \partial E_e) \leqslant q^2$.*

Proof. For all $e \in T'$, ∂E_e has the same cardinality as in $T^a_1 \cup T^a_2$. Only for the edges e_{t^i}, we have that $\partial E_{e_{t_i}} \subseteq \partial E_{t^i}$. Hence, the width increases by at most h.

For any tree edge $e \in T'$, let L be a set of leaves in the subtree inducing E_e corresponding to the edges of the cliques S_i in E_e. Then, for all $\tau'(\{u, v\}) \in L$, both endpoints u, v are vertices in ∂E_e. Let t^i be in the subtree inducing $\overline{E_e}$. Since in T^a_i, $\texttt{q-paths}_{G^a_i}(E_e, \partial E_e) \leqslant q$ and $\texttt{q-paths}_{G^a_i}(\overline{E_e}, \partial \overline{E_e}) \leqslant q$, and also $\texttt{q-paths}_{G^a_i}(E_e \cap E(S_i), \partial E_e \cap S_i) \leqslant |L|^{|L|}$, we have that in T $\texttt{q-paths}_{G'}(\overline{E_e}, \partial \overline{E_e}) \leqslant q \cdot |L|^{|L|} \leqslant q^2$ for $e \neq e_{t^i}$. From Lemma 22, and since $\partial E_{e_{t_i}} \subseteq \partial E_{t^i}$, we get $\texttt{q-paths}_{G'}(E_{e_{t_i}}, \partial E_{e_{t_i}}) \leqslant q^2$ and $\texttt{q-paths}_{G'}(\overline{E_{e_{t_i}}}, \partial \overline{E_{e_{t_i}}}) \leqslant q^2$. Deleting leaves from (T', τ') does neither increase the width nor increase the number of path collections. $\square$

In this way, we construct the branch decomposition (T, τ) of an H-minor-free graph G out of the branch decompositions $(T_1^a, \tau_1^a), \ldots, (T_n^a, \tau_n^a)$ of $\leqslant 1.5$ times the maximum width $O_H(w)$ of the h-almost embeddable graphs $G_1^a, \ldots, G_n^a$. Extending the arguments of the previous proof, and taking account that the additional value of h that is added in Lemma 23 for each $1 \leqslant i < j \leqslant n$, comes from the "gluing" clique-sum vertices between G_i^a and G_j^a that can be added in the resulting branch decomposition without contributing more than h vertices in its overall width, we obtain the following lemma.

Lemma 24. *Let be given above h-almost embeddable graphs $G_1^a, \ldots, G_n^a$ and branch decompositions $(T_1^a, \tau_1^a), \ldots, (T_n^a, \tau_n^a)$ of maximum width $O_H(w)$ and a collection $\mathcal{S}$ of h-cliques, each in one of $G_1, \ldots, G_n$. Then, the new branch decomposition (T, τ) of G has width $O_H(w)$ and for all edges $e \in T$, we have $\texttt{q-paths}_G(E_e, \partial E_e) \leqslant 2^{O_H(w)}$.*

Proof. Let L be the set of leaves for one branch decompositions (T_j^a, τ_j^a) defined as above for all h-cliques of the h-clique-sum operation for G_j^a. Since the edges corresponding to L contribute already to the sets of $\texttt{q-paths}_G(E_e, \partial E_e)$ for all $e \in T$ with $e \neq e_{tj+1}$, we get that $\texttt{q-paths}_G(\overline{E_e}, \partial \overline{E_e}) \leqslant q^2$. By Lemma 22 combined with $\partial E_{e_{tj+1}} \subseteq \partial E_{tj+1}$, we have that $\texttt{q-paths}_G(E_{e_{tj+1}}, \partial E_{e_{tj+1}}) \leqslant q^3$. $\square$

5. Algorithmic consequences

The first application of Theorem 2 is the following.

Corollary 25. *The problem of checking whether an n-vertex H-minor-free graph G has a path of length k can be solved in time $2^{O_H(\sqrt{k})} \cdot n^{O(1)}$.*

Proof. We apply the algorithm of Theorem 2 for $w = \sqrt{k}$. If it reports that G contains a $(\sqrt{k} \times \sqrt{k})$-grid as a minor, then G also contains a path of length k. If not, then the algorithm outputs in time $O_H(n^{O(1)})$ a branch decomposition (T, τ) of width $O_H(\sqrt{k})$ with the Catalan structure. In this case we use standard dynamic programming on (T, τ) (see e.g. [6]). At every step of the dynamic programming for each $e \in E(T)$, we keep track of all the ways the required path (or cycle) can cross $\omega(e) = \partial E_e$. In other words, we count all ways these paths can be rooted through ∂E_e. But this is proportional to $\texttt{q-paths}_G(E_e, \partial E_e)$. As $\texttt{q-paths}_G(E_e, \partial E_e) = 2^{O_H(\sqrt{k})}$, dynamic programming can be performed in time $O(2^{O_H(\sqrt{k})} n)$. Thus the running time of the algorithm is $O_H(n^{O(1)}) + O(2^{O_H(\sqrt{k})} n)$. $\square$

Note, that for $k = \log^2 n$, Corollary 25 gives a polynomial time algorithm for checking if an n-vertex graph has a path of length $\log^2 n$.

6. Conclusion

In this paper we have shown that every H-minor-free graph G of branch-width w has a branch decomposition of width $O(w)$ with the Catalan structure. Based on this combinatorial result we obtain a time $2^{O(\sqrt{k})} n^{O(1)}$ algorithm deciding if an H-minor-free n-vertex graph contains a k-vertex path.

Other problems that can be solved in $2^{O_H(\sqrt{k})} \cdot n^{O(1)}$ steps in H-minor-free graph classes, applying simple modifications to our technique, are the standard parameterizations of Longest Cycle, Cycle Packing, and Cycle/Path Cover (parameterized either by the total length of the cycles/paths or the number of the cycles/paths).

For a wider family of problems, where the tables of the dynamic programming are encoding packings instead of pairings, Rué et al. [34] obtained subexponential algorithms for graphs of bounded genus (see also [39]). Recently, in a breakthrough paper, Cygan et al. [5] gave a randomized algorithm solving many "connectivity" problems, including Longest Path, in time $2^{O(w)} n^{O(1)}$, where w is the branch-width of a graph. See also the work of Pilipczuk [33] on a logical characterization of such problems.

We conclude with a number of open questions.

- Our algorithmic framework is heavily based on the Robertson and Seymour's structural theorem and the suitable (topological) extension of the concept of sphere-cut decomposition. It is a challenge to prove the existence of $2^{o(k)} n^{O(1)}$ step algorithms for H-minor free graphs without using this machinery. A *combinatorial bound* to the size of the tables of the dynamic programming would be a significant step in this direction.
- A natural question is if our combinatorial result for H-minor-free graphs can be extended to larger graph classes, such as graphs locally excluding some graph as a minor [7] or graphs of bounded expansion [30]. Our result is based on bidimensionality and Robertson and Seymour's structural theorem. It is not clear if one can extend these concepts to the aforementioned larger graph classes. Even the existence of a subexponential FPT algorithm, i.e. of running time $2^{o(k)} n^{O(1)}$, solving the k-Path on these classes of graphs is unknown.

- The longest path in a (general) n-vertex graph of branch-width w can be found in time $2^{O(w \log w)} n^{O(1)}$ by making use of the standard dynamic programming on graphs of bounded branch-width. On the other hand, by a recent result of Cygan et al. [5], the problem is solvable in $2^{O(w)} n^{O(1)}$ randomized time. The natural question here is to find a single exponential deterministic algorithm for this problem.
- It is a remaining challenge to extend the framework of [34] from bounded genus graphs to H-minor free graphs.

References

[1] J. Alber, H.L. Bodlaender, H. Fernau, T. Kloks, R. Niedermeier, Fixed parameter algorithms for dominating set and related problems on planar graphs, Algorithmica 33 (2002) 461–493.
[2] J. Alber, H. Fernau, R. Niedermeier, Parameterized complexity: exponential speed-up for planar graph problems, J. Algorithms 52 (2004) 26–56.
[3] N. Alon, R. Yuster, U. Zwick, Color-coding, J. ACM 42 (1995) 844–856.
[4] A. Cayley, A theorem on trees, Quart. J. Pure Appl. Math. 23 (1889) 26–28.
[5] M. Cygan, J. Nederlof, M. Pilipczuk, M. Pilipczuk, J.M.M.v. Rooij, J.O. Wojtaszczyk, Solving connectivity problems parameterized by treewidth in single exponential time, in: Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS 2011), IEEE Computer Society, 2005, pp. 150–159.
[6] W. Cook, P. Seymour, Tour merging via branch-decomposition, INFORMS J. Comput. 15 (2003) 233–248.
[7] A. Dawar, M. Grohe, S. Kreutzer, Locally excluding a minor, in: Proceedings of the 22nd IEEE Symposium on Logic in Computer Science (LICS 2007), IEEE Computer Society, 2007, pp. 270–279.
[8] V.G. Deĭneko, B. Klinz, G.J. Woeginger, Exact algorithms for the Hamiltonian cycle problem in planar graphs, Oper. Res. Lett. 34 (2006) 269–274.
[9] E.D. Demaine, F.V. Fomin, M. Hajiaghayi, D.M. Thilikos, Bidimensional parameters and local treewidth, SIAM J. Discrete Math. 18 (2004/2005) 501–511.
[10] E.D. Demaine, F.V. Fomin, M. Hajiaghayi, D.M. Thilikos, Subexponential parameterized algorithms on graphs of bounded genus and H-minor-free graphs, J. ACM 52 (2005) 866–893.
[11] E.D. Demaine, M. Hajiaghayi, Bidimensionality: new connections between FPT algorithms and PTASs, in: Proceedings of the 16th Annual ACM–SIAM Symposium on Discrete Algorithms (SODA 2005), ACM–SIAM, New York, 2005, pp. 590–601.
[12] E.D. Demaine, M. Hajiaghayi, The bidimensionality theory and its algorithmic applications, Comput. J. 51 (2007) 292–302.
[13] E.D. Demaine, M. Hajiaghayi, Linearity of grid minors in treewidth with applications through bidimensionality, Combinatorica 28 (2008) 19–36.
[14] E.D. Demaine, M. Hajiaghayi, B. Mohar, Approximation algorithms via contraction decomposition, in: Proceedings of the 18th Annual ACM–SIAM Symposium on Discrete Algorithms (SODA 2007), 2007, pp. 278–287.
[15] E.D. Demaine, M.T. Hajiaghayi, K. Kawarabayashi, Algorithmic graph minor theory: Decomposition, approximation, and coloring, in: Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005), IEEE Computer Society, 2005, pp. 637–646.
[16] R. Diestel, Graph Theory, third ed., Grad. Texts in Math., vol. 173, Springer-Verlag, Berlin, 2005.
[17] F. Dorn, F.V. Fomin, D.M. Thilikos, Fast subexponential algorithm for non-local problems on graphs of bounded genus, in: Proceedings of the 10th Scandinavian Workshop on Algorithm Theory (SWAT 2006), in: Lecture Notes in Comput. Sci., vol. 4059, Springer, Berlin, 2006, pp. 172–183.
[18] F. Dorn, F.V. Fomin, D.M. Thilikos, Catalan structures and dynamic programming in H-minor-free graphs, in: Proceedings of the 19th ACM–SIAM Symposium on Discrete Algorithms (SODA 2008), ACM–SIAM, 2008, pp. 631–640.
[19] F. Dorn, F.V. Fomin, D.M. Thilikos, Subexponential parameterized algorithms, Comput. Sci. Rev. 2 (2008) 29–39.
[20] F. Dorn, E. Penninkx, H. Bodlaender, F.V. Fomin, Efficient exact algorithms on planar graphs: Exploiting sphere cut decompositions, Algorithmica 58 (3) (2010) 790–810.
[21] R.G. Downey, M.R. Fellows, Parameterized Complexity, Springer-Verlag, New York, 1999.
[22] J. Flum, M. Grohe, Parameterized Complexity Theory, Texts Theoret. Comput. Sci. EATCS Ser., Springer-Verlag, Berlin, 2006.
[23] F.V. Fomin, D.M. Thilikos, Dominating sets in planar graphs: Branch-width and exponential speed-up, SIAM J. Comput. 36 (2006) 281–309.
[24] Q.-P. Gu, H. Tamaki, Optimal branch-decomposition of planar graphs in $O(n^3)$ time, ACM Trans. Algorithms 4 (2008).
[25] R. Impagliazzo, R. Paturi, F. Zane, Which problems have strongly exponential complexity, J. Comput. System Sci. 63 (2001) 512–530.
[26] K. Kawarabayashi, P. Wollan, A simpler algorithm and shorter proof for the graph minor decomposition, in: Proceedings of the 43rd Annual ACM Symposium on Theory of Computing (STOC 2011), ACM, 2011, pp. 451–458.
[27] P.N. Klein, A linear-time approximation scheme for TSP in undirected planar graphs with edge-weights, SIAM J. Comput. 37 (2008) 1926–1952.
[28] G. Kreweras, Sur les partition non croisées d'un circle, Discrete Math. 1 (1972) 333–350.
[29] B. Mohar, C. Thomassen, Graphs on Surfaces, Johns Hopkins University Press, Baltimore, MD, 2001.
[30] J. Nešetřil, P. Ossona de Mendez, Linear time low tree-width partitions and algorithmic consequences, in: Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC 2006), ACM, 2006, pp. 391–400.
[31] R. Niedermeier, Invitation to Fixed-Parameter Algorithms, Oxford Lecture Ser. Math. Appl., vol. 31, Oxford University Press, Oxford, 2006.
[32] C.H. Papadimitriou, M. Yannakakis, On limited nondeterminism and the complexity of the V-C dimension, J. Comput. System Sci. 53 (1996) 161–170.
[33] M. Pilipczuk, Problems parameterized by treewidth tractable in single exponential time: a logical approach, in: Proceedings of the 36th International Symposium Mathematical Foundations of Computer Science (MFCS 2011), in: Lecture Notes in Comput. Sci., vol. 6907, Springer, 2011, pp. 520–531.
[34] J. Rué, I. Sau, D.M. Thilikos, Dynamic programming for graphs on surfaces, in: Proceedings of the 37th International Colloquium Automata, Languages and Programming (ICALP 2010), in: Lecture Notes in Comput. Sci., vol. 6198, Springer, 2010, pp. 372–383.
[35] N. Robertson, P.D. Seymour, Graph minors. VII. Disjoint paths on a surface, J. Combin. Theory Ser. B 45 (1988) 212–254.
[36] N. Robertson, P.D. Seymour, Graph minors. XI. Circuits on a surface, J. Combin. Theory Ser. B 60 (1994) 72–106.
[37] N. Robertson, P.D. Seymour, Graph minors. XVI. Excluding a non-planar graph, J. Combin. Theory Ser. B 89 (2003) 43–76.
[38] P.D. Seymour, R. Thomas, Call routing and the ratcatcher, Combinatorica 14 (1994) 217–241.
[39] I. Sau, D.M. Thilikos, Subexponential parameterized algorithms for degree-constrained subgraph problems on planar graphs, J. Discrete Algorithms 8 (3) (2010) 330–338.
[40] R.P. Stanley, Enumerative Combinatorics, vol. 1, Cambridge Stud. Adv. Math., vol. 49, Cambridge University Press, Cambridge, 1997.
[41] C. Thomassen, Embeddings of graphs with no short noncontractible cycles, J. Combin. Theory Ser. B 48 (1990) 155–177.
[42] A. Tucker, Coloring a family of circular arcs, SIAM J. Appl. Math. 29 (3) (1975) 493–502.