REPORTS IN INFORMATICS

ISSN 0333-3590

Catalan Structures and Dynamic Programming in H-minor-free graphs

Frederic Dorn   Fedor V. Fomin
Dimitrios M. Thilikos

REPORT NO 357   July 2007

Department of Informatics
UNIVERSITY OF BERGEN
Bergen, Norway

This report has URL http://www.ii.uib.no/publikasjoner/texrap/pdf/2007-357.pdf

Reports in Informatics from Department of Informatics, University of Bergen, Norway, is available at
http://www.ii.uib.no/publikasjoner/texrap/.

Requests for paper copies of this report can be sent to:

Department of Informatics, University of Bergen, Høyteknologisenteret,
P.O. Box 7800, N-5020 Bergen, Norway

Catalan Structures and Dynamic Programming in H-minor-free graphs

Frederic Dorn* Fedor V. Fomin* and Dimitrios M. Thilikos†

Department of Informatics
University of Bergen
PO Box 7800, 5020 Bergen, Norway
{dorn,fomin}@ii.uib.no

Department of Mathematics
National & Capodistrian University of Athens
Panepistimioupolis, GR-15784, Athens, Greece
sedthilk@math.uoa.gr

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.

*Additional support by the Research Council of Norway.

†Supported by the Project "Capodistrias" (AΠ 736/24.3.2006) of the National and Capodistrian University of Athens (project code: 70/4/8757).

1 Introduction

Our research has been motivated by the seminal result of Alon, Yuster, and Zwick in [4] that proved that a path of length $\log n$ can be found in polynomial time, answering to a question by Papadimitriou and Yannakakis in [30]. One of the open questions left in [4] 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 [9], 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 [20, 21, 29]. 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 **(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 [10]) **(a)** can be proved systematically using the idea of Bidimensionality [12]. 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 to the distribution of their certificates in the graph ([1, 2, 11, 22]). 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 [4] 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 TSP.

Recently, [19] overcommed the above deadlock for the class of planar graphs. Later, a similar result was given in [18] for graphs of bounded genus. The proofs in [19, 18] 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). It is not clear, a priory, if these type of arguments can be extended to graphs excluding a minor. We stress that the lack of such arguments was explicitly named by Grigni [23] as the main obstacle of obtaining truly polynomial time approximation scheme for TSP on H-minor-free graphs, while this was possible for planar [5] or bounded genus graphs [23]. (For another example of a technique where the extension from bounded genus to H-minor-free graphs is not clear, see [15].)

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 [33]) to prove counting lemmata that can suitably bound the amount of information required in each step of the dynamic programming algorithm. Among the problems that are amenable to our approach, we drive our presentation using the LONGEST PATH problem.

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 [33], 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 [34], 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 for parameterized and approximation algorithms in [9,14,26]. This decomposition is used to build a decomposition on the initial almost embeddible 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 [7]) we find the longest path in G. We stress that the dynamic programming algorithm is not different than 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 total running time of the algorithm is $2^{O(\sqrt{k})} \cdot n^{O(1)}$.

Finally, let us remark that analogues of Theorem 2 can be proved for other problems without any substantial change of our methodology. This implies the existence of $2^{O(\sqrt{k})} \cdot n^{O(1)}$-step algorithms for several parameterized problems on H-minor-free graphs. Such problems are the natural parameterizations for LONGEST CYCLE, FEEDBACK VERTEX SET and different parameterizations of PATH COVER and CYCLE COVER (for the concept of parameterization and other concepts of parameterized complexity, see [20]). Also, we can prove $2^{O(\sqrt{k})} \cdot n^{O(1)}$-step algorithms for apex-minor-free graphs for CONNECTED DOMINATING SET and MAX LEAF TREE. Finally, our results imply $2^{O(\sqrt{n})}$ step exact algorithms for all aforementioned problems as well as GRAPH METRIC TSP, STEINER TREE and MAXIMUM FULL DEGREE SPANNING TREE.

2 Preliminaries

Surface embeddable graphs 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 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, then we call such a noose *tight*. Let Δ be a closed disk and the open disk $\mathbf{int}(\Delta)$ its interior and $\mathbf{bor}(\Delta)$ its boundary. Then $\Delta = \mathbf{int}(\Delta) \cup \mathbf{bor}(\Delta)$. If $\mathbf{int}(\Delta)$ is subset of a region of G, then $\mathbf{bor}(\Delta)$ is a noose.

Surface cutting. We need to define the graph obtained by *cutting along* a noncontractible 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 trunk decompositions Let G be a graph and let $E \subseteq E(G)$. We define ∂E as the set of vertices in G that are endpoints of edges in E and of edges in $E(G) - E$. 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 ∂E_e where E_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. If in the definition of a branch decomposition we further demand the ternary tree T to be a caterpillar, then we define the notion of a *trunk decomposition* and the parameter of the *trunkwidth* of a graph. For a longest path with edges $e_1, \ldots, e_q$ of such a caterpillar, the sets $X_i = \omega(e_i)$ form a linear ordering $\mathcal{X} = (X_1, \ldots, X_q)$. For convenience, we will use ordered sets to denote a trunk decompositions and, in order to include all vertices of G in the sets of $\mathcal{X}$, we will often consider trunk decompositions of $\hat{G}$ that is G with loops added to all its vertices (this operation cannot increase the width by more than one).

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, there exists a tight noose N bounding two open disks Δ_1 and Δ_2 such that $E_e^i \subseteq \Delta_i \cup N$, $1 \leq i \leq 2$. Thus N meets G only in ∂E_e and its length is $|\partial E_e|$. Clockwise traversing of N in the drawing 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 [34] (improved by [24]), there is a $O(n^3)$ algorithm finding an optimal branch decomposition of a planar graph.

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. 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 19 (proved in the Appendix), 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 $S = \omega(\{t, a\}) \cup \omega(\{t, b\}) \cup \omega(\{t, c\})$ 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 $|A| \leq h$ and such that $G - A$ is isomorphic to $G_u \cup G_1 \cup \cdots \cup G_r, r \leq h$ in a way that the following conditions are satisfied (the definition below is not the original one from [33] but equivalent, slightly adapted for the purposes of our paper):

• There exists an embedding $\sigma : G_u \to \Sigma^{-r}, r \leq h$ 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 trunk decomposition $\mathcal{X}_i = (X_1^i, \ldots, X_{q_i}^i)$ of the vortex G_i with width at most h and a subset $J_i = \{j_1^i, \ldots, j_{|B_i|}^i\} \subseteq \{1, \ldots, q_i\}$ such that $\forall_{k=1,\ldots,|B_i|}\ u_k^i \in X_{j_k^i}^i$ for some respectful ordering $(u_1^i, \ldots, u_{|B_i|}^i)$ of B_i. (An ordering $(u_1^i, \ldots, u_{|B_i|}^i)$ is called *respectful* if the ordering $(\sigma(u_1^i), \ldots, \sigma(u_{|B_i|}^i))$ follows the cyclic ordering of the corresponding cycle of $\mathbf{bor}(\Sigma^{-r})$.) For every vertex $u_k^i \in B_i$, we call $X_{j_k^i}^i$ *overlying set* of u_j^i and we denote it by $\mathbf{X}(u_j^i)$.

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.

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 [33],) characterizing H-minor-free graphs.

Proposition 1 ([33]). *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.

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 (different) 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 $\sim$.

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

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

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, it is possible to construct a time $O_H(1) \cdot n^{O(1)}$ algorithm that 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 and of width $O_H(w)$.*

In what follows, we will give the description of the algorithm of Theorem 2 and we will sketch the main arguments supporting its correctness. While the first statement of the theorem follows almost directly from [16], our main contribution lies in the proof of statement 2. The full proof is lengthy and complicated and all lemmata supporting its correctness have moved to the Appendix.

1. Use the time $O_H(1) \cdot n^{O(1)}$ algorithm of [16] (see also [8]) to decompose the input graph into a collection $\mathcal{C}$ of clique-sum components as in Proposition 1. Every graph in $\mathcal{C}$ is a γ_H-almost embeddable graph to some surface of genus $\leq \gamma_H$ where $\gamma_H = O_H(1)$.

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

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.

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

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

i. Find a non-contractible noose N in $G^{(i)}_u$ of minimum length, using the polynomial time algorithm in [35].

ii. If $|N| \geq 2^{i-1} f(H) \cdot w$ then output "*The input graph contains a* $(w \times w)$*-grid as a minor*" and stop. The estimation of $f(H)$ comes from the results in [13], presented in Lemma 5. Along with Lemma 6 follows the correctness of this step.

Notice that, by minimality, N cannot intersect the interior of a hole or a vortex disk Δ more than once and, for the same reason, it can intersect $\mathbf{bor}(\Delta)$ in at most two vertices. If $\mathbf{int}(\Delta) \cap N = \emptyset$ and $\mathbf{bor}(\Delta) \cap N = \{v, w\}$, again from minimality, v and w should be successive in $\mathbf{bor}(\Delta)$. In this case, we re-route this portion of N so that it crosses the interior of Δ (see Figure 2 in the Appendix).

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 $i \geq 2$) or some vortex disk of $G_u^{(i)}$ 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 relative 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 disk Δ it also intersects its interior and therefore it will split Δ into two parts Δ_1 and Δ_2.

iv. Cut $G_u^{(i)}$ along N and call the two disks created by the corresponding cut of the surface *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 trunk 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 $G_u^{(i+1)}$ be its underlying graph. Recall that, from the previous steps, a vortex disk or a hole Δ (if divided) 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. $\Delta_j' \leftarrow \Delta_j \cup \Delta_j', j = 1, 2$).

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 lower Euler genus – see [28, 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 to the sphere. Also notice that the total length of the holes of $G_u^{(i)}$ is upper bounded by the sum of the lengths of the nooses we cut along, which is $O_H(w)$.

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 [34].

e. If $\mathbf{bw}(G_u^p) \geq 2^{\gamma_H} \cdot f(H) \cdot w = \Omega_H(w)$, then output "*The input graph contains a* $(w \times w)$*-grid as a minor*" and stop. This step is justified by Lemma 7.

f. Enhance (T_u^p, τ_u^p), so that the edges of the vortices of G^p are included to 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 trunk 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 make the resulting graph 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 Figure 1 in the Appendix). Applying this transformation for each vortex G^v of G^p, we construct a branch decomposition of G^p. In Lemma 8, we prove that this enhancement of the branch decomposition of G_u^p can add $O_H(1)$ vertices for each vertex in $\omega(e), e \in E(T_u^p)$, therefore, $\mathbf{bw}(G^p) = O_H(\mathbf{bw}(G_u^p))$.

g. Notice that, while successively splitting G^s during the loop of step **2.c.**, all edges remain topologically intact (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 $T^s = T^p$. At this point, we have to prove that if the bounds of Theorem 2 holds for the graph $G^p = G^{(i)}$ (a graph that is smoothly $O_H(1)$-almost embeddible in the sphere), then they also hold for the graph $G^s = G^{(1)}$ that is a smoothly $O_H(1)$-almost embeddible in a surface of higher genus. We prove that $\mathbf{bw}(G^s) = O_H(\mathbf{bw}(G^p))$ with the help of Lemma 9.a (using induction). However, what is far more complicated is to prove that (T^s, τ^s) has the Catalan structure. For this, we first prove (using inductively Lemma 9.b) that for any edge e of $T^s = T^p$, holds that

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

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-contracible nooses, the split vertices in the nooses (i.e., the vertices in D_e) may separate paths counted in the left side of Equation (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: find a bound for the number of equivalent classes of collections of vertex disjoint paths whose endpoints may be a) vertices of the disk Δ_e bounding the edges E_e in the sphere-cut decomposition (T^p, τ^p) of G^p_u, along with their overlying sets (all-together, these ends are at most $\mathbf{bw}(G^p) = O_H(w)$, because of Lemma 8) and b) vertices (and their overlying sets) of $2 \cdot \gamma_H = O_H(1)$ disjoint holes (created by cutting through nooses) of total size $\leq \gamma_H \cdot (2^{\gamma_H} - 1) \cdot f(H) \cdot w = O_H(w)$ (see the proof of Lemma 6). Notice that these paths can be routed also via $\leq \gamma_H \cdot 2^i \leq \gamma_H \cdot 2^{\gamma_H} = O_H(1)$ vortices (because, initially, G^s_u had γ_H vortices and, in the worst case, each noose can split every vortex into two parts), each of unbounded size. Recall that the holes and the vortex disks of G^p_u do not touch (i.e. they intersect but their interiors do not) 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 disk 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 disk bounding the edges E_e, see Figure 3.) 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. As this proof is quite technical, we moved it to the Appendix (Lemma 10) and, here, we will give just a sketch. Our 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 equivalent 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 disk 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 equivalent 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 equivalent class is the superposition of $O_H(1)$ equivalent classes of non-crossing collections of disjoint paths. Because of this, the number of equivalent 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.

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 $O_H(1)$ vertices more (Lemma 10). 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 18).

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^a_1, τ^a_1) and (T^a_2, τ^a_2) are two branch decompositions of two graphs G^a_1 and G^a_2 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^a_1 and G^a_2, as follows: Let t^i be a clique-node of S_i in $(T_i, \tau_i), i = 1, 2$. 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 22) 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 Algorithmic consequences

A first application of Theorem 2 is the following.

Corollary 3. *The problem of checking whether there is a path of length k on H-minor-free graphs can be solved in $2^{O_H(\sqrt{k})} \cdot n^{O(1)}$ steps.*

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, then G also contains a path of length k. If not, then the algorithm outputs a branch decomposition (T, τ) of width $O_H(\sqrt{k})$, as in Theorem 2. By applying dynamic programming on (T, τ) we have, for each $e \in E(T)$, to keep track of all the ways the required path (or cycle) can cross $\omega(e) = \partial E_e$. This is proportional to $\texttt{q-paths}_G(E_e, \partial E_e)$ (counting all ways these paths can be rooted through ∂E_e). As $\texttt{q-paths}_G(E_e, \partial E_e) = 2^{O_H(\sqrt{k})}$ we have the claimed bounds. □

Note, that for $k = \log^2 n$, Corollary 3 gives a polynomial time algorithm for checking if a n-vertex graph has a path of length $\log^2 n$.

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, FEEDBACK VERTEX SET, and CYCLE/PATH COVER (parameterized either by the total length of the cycles/paths or the number of the cycles/paths).

Moreover, combining Theorem 2, with the results in [17] we can derive time $2^{O_H(\sqrt{k})} \cdot n^{O(1)}$ algorithms for problems emerging from contraction closed parameters for apex-minor-free graph classes (a graph is an *apex* graph the removal of one of its vertices creates a planar graph). The most prominent examples of such problems are CONNECTED DOMINATING SET and MAX LEAF TREE. (The best previous algorithm for these problems for apex-minor-free graph classes was a $2^{O_H(\sqrt{k} \cdot \log k)} \cdot n^{O(1)}$ step algorithm given in [14].)

Our technique can also be used to design fast subexponential exact algorithms. Notice that the branchwidth of any H-minor-free graph is at most $O_H(\sqrt{n})$ [3]. The algorithm of Theorem 2 will output a branch decomposition of width $O_H(\sqrt{n})$ that, using an adequate definition of Catalan structure, can be used to derive $2^{O_H(\sqrt{n})}$ step algorithms for several problems. Consider for example WEIGHTED GRAPH METRIC TSP (TSP with the shortest path metric of G as distance metric). It is shown in [19] how to solve GRAPH METRIC TSP on planar graphs and in [18] on bounded genus graphs. The basic idea is that any solution to GRAPH METRIC TSP can be reduced to finding a minimum weight spanning Eulerian subgraph. In this case, instead of having collections of paths $\mathbf{paths}_G(E_e, \partial E_e)$ we deal with connected components, say $\mathbf{comp}_G(E_e, \partial E_e)$. Nevertheless, we can use the Catalan structures argument and extend our counting results about $\texttt{q-paths}_G(E_e, \partial E_e)$. Apart from the problems that we have already mentioned above, $2^{O_H(\sqrt{n})}$ step exact algorithms can be designed for STEINER TREE, MAXIMUM FULL DEGREE SPANNING TREE, and other types of spanning tree problems.

5 Conclusion

When applying our technique on different problems we define, for each one of them, an appropriate analogue of $\texttt{q-paths}_G(E_e, \partial E_e)$ and prove that it also satisfies the Catalan structure property (i.e. is bounded by $2^{O_H(|\partial E_e|)}$). It would be challenging to find a classification criterion (logical or combinatorial) for the problems that are amenable to this approach.

References

[1] J. ALBER, H. L. BODLAENDER, H. FERNAU, T. KLOKS, AND R. NIEDERMEIER, *Fixed parameter algorithms for dominating set and related problems on planar graphs*, Algorithmica, 33 (2002), pp. 461–493.

[2] J. ALBER, H. FERNAU, AND R. NIEDERMEIER, *Parameterized complexity: exponential speed-up for planar graph problems*, J. Algorithms, 52 (2004), pp. 26–56.

[3] N. ALON, P. SEYMOUR, AND R. THOMAS, *A separator theorem for nonplanar graphs*, J. Amer. Math. Soc., 3 (1990), pp. 801–808.

[4] N. ALON, R. YUSTER, AND U. ZWICK, *Color-coding*, J. Assoc. Comput. Mach., 42 (1995), pp. 844–856.

[5] S. ARORA, M. GRIGNI, D. KARGER, P. KLEIN, AND A. WOLOSZYN, *A polynomial-time approximation scheme for weighted planar graph TSP*, in Proceedings of the Ninth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 1998), New York, 1998, ACM, pp. 33–41.

[6] A. CAYLEY, *A theorem on trees*, Quart J. Pure Appl. Math., 23 (1889), pp. 26–28.

[7] W. COOK AND P. SEYMOUR, *Tour merging via branch-decomposition*, INFORMS Journal on Computing, 15 (2003), pp. 233–248.

[8] A. DAWAR, M. GROHE, AND S. KREUTZER, *Locally excluding a minor*, preprint NI07001-LAA, Isaac Newton Institute of Mathematical Sciences, (2007).

[9] V. G. DEǏNEKO, B. KLINZ, AND G. J. WOEGINGER, *Exact algorithms for the Hamiltonian cycle problem in planar graphs*, Oper. Res. Lett., 34 (2006), pp. 269–274.

[10] E. D. DEMAINE, F. V. FOMIN, M. HAJIAGHAYI, AND D. M. THILIKOS, *Bidimensional parameters and local treewidth*, SIAM J. Discrete Math., 18 (2004/05), pp. 501–511 (electronic).

[11] ———, *Subexponential parameterized algorithms on graphs of bounded genus and H-minor-free graphs*, Journal of the ACM, 52 (2005), pp. 866–893.

[12] E. D. DEMAINE AND M. HAJIAGHAYI, *The bidimensionality theory and its algorithmic applications*, Computer Journal, to appear.

[13] ———, *Linearity of grid minors in treewidth with applications through bidimensionality*, Combinatorica, to appear.

[14] ———, *Bidimensionality: new connections between FPT algorithms and PTASs*, in Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2005), New York, 2005, ACM-SIAM, pp. 590–601.

[15] E. D. DEMAINE, M. HAJIAGHAYI, AND 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.

[16] E. D. DEMAINE, M. T. HAJIAGHAYI, AND 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.

[17] E. D. DEMAINE, M. T. HAJIAGHAYI, AND D. M. THILIKOS, *Exponential speedup of fixed-parameter algorithms for classes of graphs excluding single-crossing graphs as minors.*, Algorithmica, 41 (2005), pp. 245–267.

[18] F. DORN, F. V. FOMIN, AND 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), vol. 4059 of LNCS, Springer, Berlin, 2006, pp. 172–183.

[19] F. DORN, E. PENNINKX, H. BODLAENDER, AND F. V. FOMIN, *Efficient exact algorithms on planar graphs: Exploiting sphere cut branch decompositions*, in Proceedings of the 13th Annual European Symposium on Algorithms (ESA 2005), vol. 3669 of LNCS, Springer, Berlin, 2005, pp. 95–106.

[20] R. G. DOWNEY AND M. R. FELLOWS, *Parameterized complexity*, Springer-Verlag, New York, 1999.

[21] J. FLUM AND M. GROHE, *Parameterized Complexity Theory*, Texts in Theoretical Computer Science. An EATCS Series, Springer-Verlag, Berlin, 2006.

[22] F. V. FOMIN AND D. M. THILIKOS, *Dominating sets in planar graphs: Branch-width and exponential speed-up*, SIAM J. Comput., 36 (2006), pp. 281–309.

[23] M. GRIGNI, *Approximate TSP in graphs with forbidden minors.*, in proceedings of the 27th International ColloquiumAutomata, Languages and Programming (ICALP 2000), vol. 1853 of Lecture Notes in Computer Science, Springer, 2000, pp. 869–877.

[24] Q.-P. GU AND H. TAMAKI, *Optimal branch-decomposition of planar graphs in O(n^3) time*, in Proceedings of the 32nd International Colloquium on Automata, Languages and Programming (ICALP 2005), vol. 3580 of LNCS, Springer, Berlin, 2005, pp. 373–384.

[25] R. IMPAGLIAZZO, R. PATURI, AND F. ZANE, *Which problems have strongly exponential complexity*, Journal of Computer and System Sciences, 63 (2001), pp. 512–530.

[26] P. N. KLEIN, *A linear-time approximation scheme for planar weighted TSP*, in Proceedings of 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005), IEEE Computer Society, 2005, pp. 647–657.

[27] G. KREWERAS, *Sur les partition non croisées d'un circle*, Discrete Mathematics, 1 (1972), pp. 333–350.

[28] B. MOHAR AND C. THOMASSEN, *Graphs on surfaces*, Johns Hopkins University Press, Baltimore, MD, 2001.

[29] R. NIEDERMEIER, *Invitation to fixed-parameter algorithms*, vol. 31 of Oxford Lecture Series in Mathematics and its Applications, Oxford University Press, Oxford, 2006.

[30] C. H. PAPADIMITRIOU AND M. YANNAKAKIS, *On limited nondeterminism and the complexity of the v-c dimension*, J. Comput. Syst. Sci., 53 (1996), pp. 161–170.

[31] N. ROBERTSON AND P. D. SEYMOUR, *Graph minors. VII. Disjoint paths on a surface*, J. Combin. Theory Ser. B, 45 (1988), pp. 212–254.

[32] ———, *Graph minors. XI. Circuits on a surface*, J. Combin. Theory Ser. B, 60 (1994), pp. 72–106.

[33] ———, *Graph minors. XVI. Excluding a non-planar graph*, J. Combin. Theory Ser. B, 89 (2003), pp. 43–76.

[34] P. D. SEYMOUR AND R. THOMAS, *Call routing and the ratcatcher*, Combinatorica, 14 (1994), pp. 217–241.

[35] C. THOMASSEN, *Embeddings of graphs with no short noncontractible cycles*, J. Combin. Theory Ser. B, 48 (1990), pp. 155–177.

Appendix

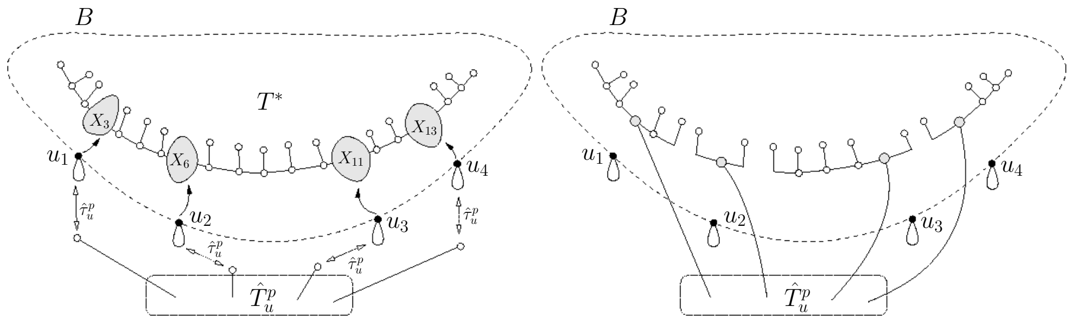

Figure 1: The procedure of enhancing the branch decomposition (T_u^p, τ_u^p) of G_u^p to a branch decomposition of G^p.

Proof of Step 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 [31] 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 noncontractible noose in Σ. The following lemma follows from Theorem 4.1 of [32].

Lemma 4. *For any graph embeddable in a non planar surface, it holds that* $\mathbf{rep}(G) \leq \mathbf{bw}(G)$.

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 5. *Let G be an H-minor-free graph and let G_u^s as in Step* **2.a**. *Then, there exists a function $f(H)$ such that if* $\mathbf{bw}(G_u^s) \geq f(H) \cdot w$*, 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 5.

Lemma 6. *If in the x-th application of Step* **2.c.ii**, $|N| \geq 2^{x-1} f(H) \cdot w$*, 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

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

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

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

Inductively, we have

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

We set $N = N_x$. By Lemma 4, $\mathbf{bw}(G_u^x) \geq \mathbf{rep}(G_u^x)$ and $\mathbf{rep}(G_u^x) \geq |N_x| \geq 2^{x-1} f(H) \cdot w$. Thus, we conclude that $\mathbf{bw}(G_u^1) \geq f(H) \cdot w$. □

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

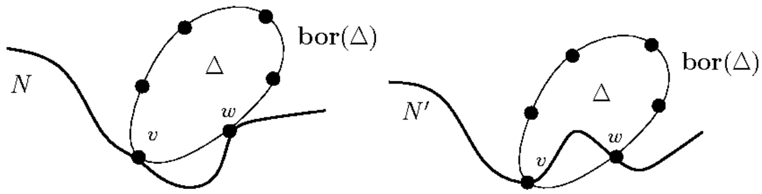

Figure 2: Re-routing a noose.

Lemma 7. *If in Step* **2e**, $\mathbf{bw}(G_u^p) \geq 2^{\gamma_H - 1} f(H) \cdot w$*, 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 6 if we set $x = i$ and with deference that, in the end, we directly we 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 $i \leq \gamma_H$. □

Proof of Step 2.f.: Enhancing the branch decomposition of G_u^p

In the following lemma, we state how to enhance the branch decomposition of G_u^p by the trunk decompositions of the vortices in order to obtain a branch decomposition of G^p.

Lemma 8. *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 $e \in E(T^p)$, $\partial E_e(G^p) \subseteq \mathbf{X}(\partial E_e(G_u^p))$. 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 $r \leq h$ vortex disks we have that there are at most $2r \leq 2h$ vertices of $\partial E_e(G_u^p)$ that are members of some base sets B. Therefore, for any $e \in E(T^p)$, $|\partial E_e(G^p)| \leq w + 2h^2$. We conclude that the width of (T^p, τ^p) is at most $w + 2h^2$. □

Proof of Step 2.g.: Towards Catalan structure

The whole subsection is devoted to proving why the branch decomposition we constructed for a smoothly almost-embeddable graph has the Catalan structure. With the following lemma, we can inductively show how we obtain a branch decomposition with the Catalan structure for a smoothly almost-embeddable graph to a higher surfaces from the branch decomposition of its planarized version.

Lemma 9. *Let* G, G' *be two almost embeddedible graphs created successively during Step* **2.c** $(i \geq 2)$*, 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 wholes 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 if* $T = T'$ *and* $\tau = \tau' \circ \sigma$ *where* $\sigma : E(T) \to E(T')$ *is the bijection pairing topologically equivalent edges in* G *and* G'. *For any* $e \in E(T) = E(T')$*, the following hold:*

a. $|\omega_G(e)| \leq |\omega_{G'}(\sigma(e))| + |\mathbf{X}(N)|$.

b. if $S \subseteq V(E_e)$*, the number* $|\mathbf{paths}_G(E_e, \partial E_e \cup S)|$ *is bounded by*

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

Proof. To see $|\omega_G(e)| \leq |\omega_{G'}(e)| + |\mathbf{X}(N)|$, 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.

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). □

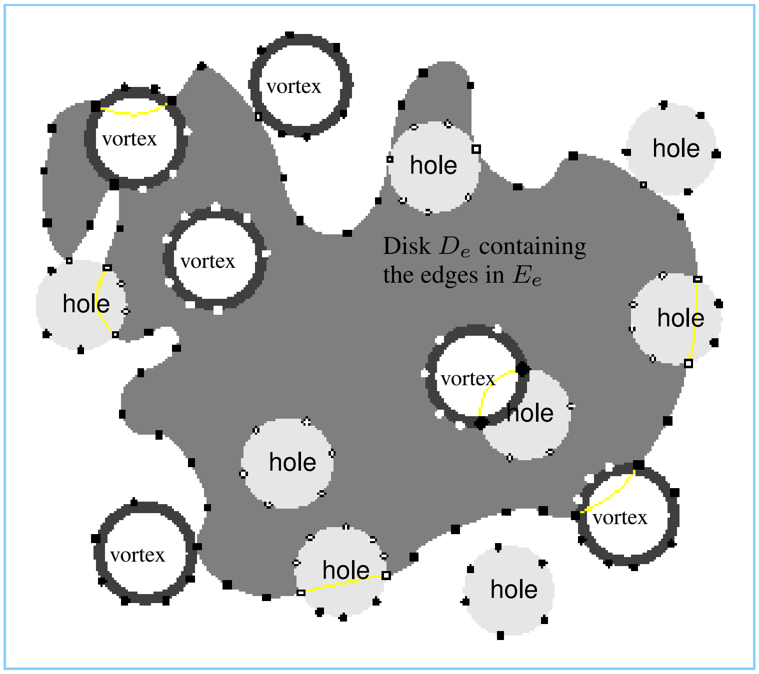

Figure 3: Vortices and holes around the disk Δ_e.

Fixing paths in smoothly h-almost embeddible graphs on the sphere. The remainder of this subsection is devoted to proving the following lemma:

Lemma 10. *Let G^p be a smoothly $O_H(1)$-almost embeddible graph in the sphere and let $\Delta_1, \ldots, \Delta_r$ ($r = O_H(1)$) be disjoint closed disks (holes) of the sphere whose interior does not intersect the underlying graph G^p_u. Assume also that, if a vortex disk and a hole intersect, then they have common interior points. Let Δ_e be a closed disk 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*

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

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, by 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.

We say that two closed disks Δ_1 and Δ_2 *touch* if their interiors are disjoint and they have common points that are vertices. These vertices are called *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 11. *Let $E_e, \mathbf{V_e}$ be as above. Then,*

$$\texttt{q-paths}_{G^p}(E_e, \mathbf{V_e}) \leq \texttt{q-paths}_{G^{p*}}(E_e^*, \mathbf{V_e^*}).$$

Where E_e^* and ${\mathbf{V_e}}^*$ are the enhanced sets in G^{p*}.

The red lines in the diagram in Figure 4 emphasize inverse edge contractions.

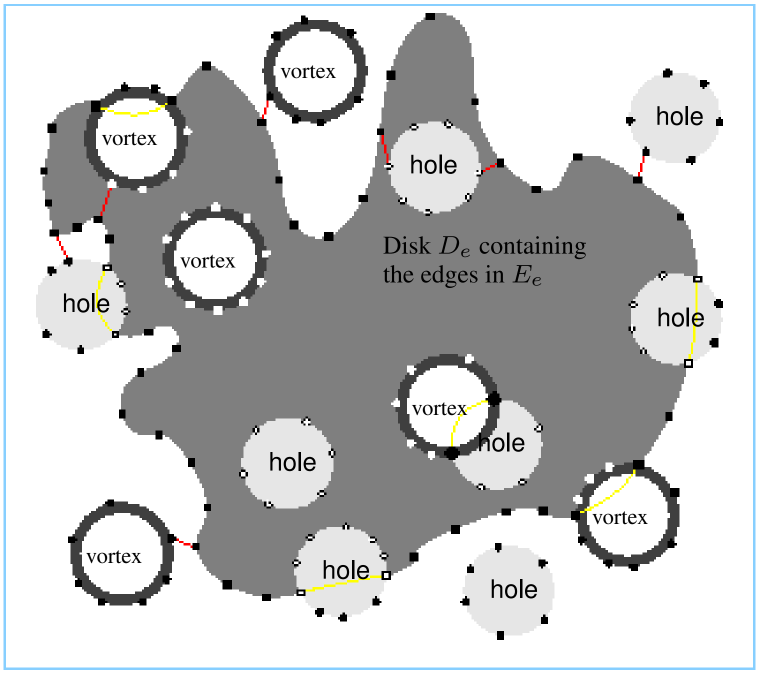

Figure 4: Vortices and holes not intersecting the disk Δ_e (first normalization).

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

Lemma 12. $|\mathbf{V_e}^*| \leq 2|\mathbf{V_e}|$.

On the left of Figure 5, we have the now resulting graph of Figure 4 where the grey 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.

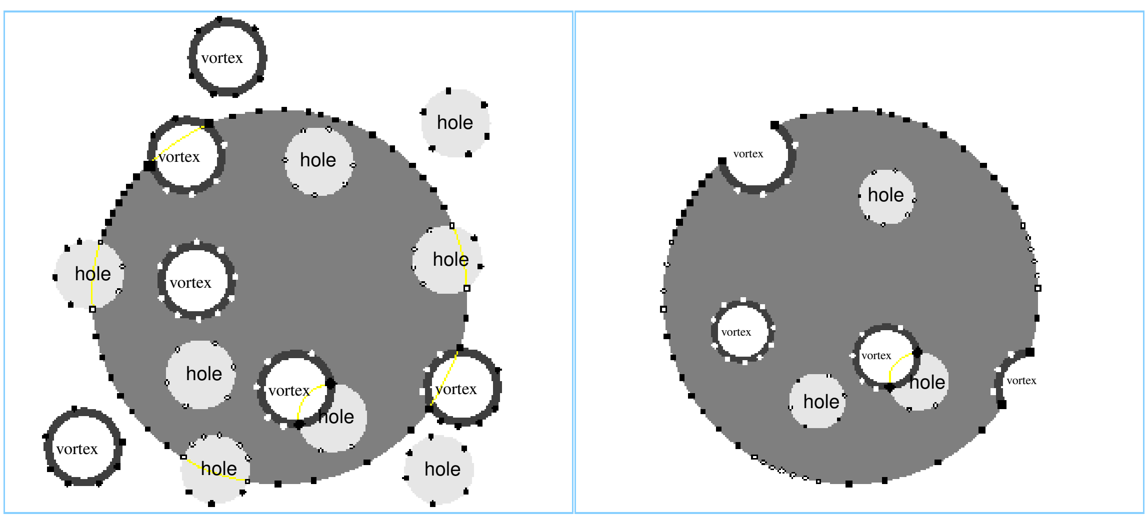

Figure 5: Another way to see Figure 4.

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

$$V(R_{h,s}) = V_1 \cup \cdots \cup V_s \text{ with } |V_i| = h \text{ and}$$

$$E(R_{h,s}) = \{\{x_j, x_k\} \mid x_j, x_k \in V_i, 1 \leq j \neq k \leq h, 1 \leq i \leq s\} \cup$$

$$\{\{x_j, y_j\} \mid x_j \in V_{i-1} \wedge y_j \in V_i, 1 \leq j \leq h, 1 \leq i \leq s\}.$$

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 $(R_{h,s}, S)$ a (h, s)-vortex pattern. See Figure 6 for an example of a normalized vortex.

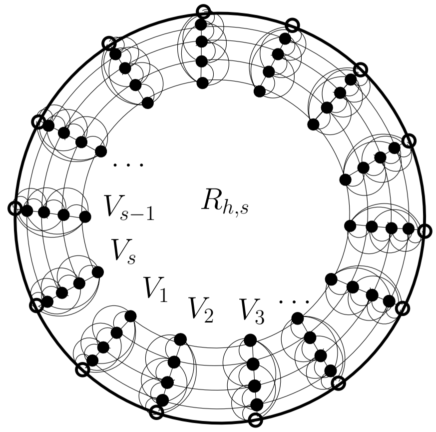

Figure 6: **Normalizing vortices.** An example of a $(6, 14)$-vortex pattern.

We now prove the following:

Lemma 13. *Any vortex of a h-almost embeddable graph with base set B is a minor of a (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. $R_{h,s}$ has trunkwidth h.*

Proof. We show how any vortex with base set B and trunk decomposition $\mathcal{X} = (X_1, \dots, X_{|B|})$ of width $< h$ is a minor of some (h, s)-vortex pattern $(R_{h,s}, S)$ with $V(R_{h,s}) = V_1 \cup \dots \cup V_s$ and $|V_i| = h (1 \leq i \leq s)$. 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, 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.

□

Using Lemma 13, 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 all vortex disks inside Δ_e and no vortices intersecting holes. Assume that for any (h, s)-vortex pattern we have that V_i and V_j are the two vertex sets of $R_{h,s} \cap \partial E_e$ (and $R_{h,s} \cap \mathbf{bor}(\Delta)$, Δ one of the holes $\Delta_1, \dots, \Delta_r$, respectively.) 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 show in the next figure.

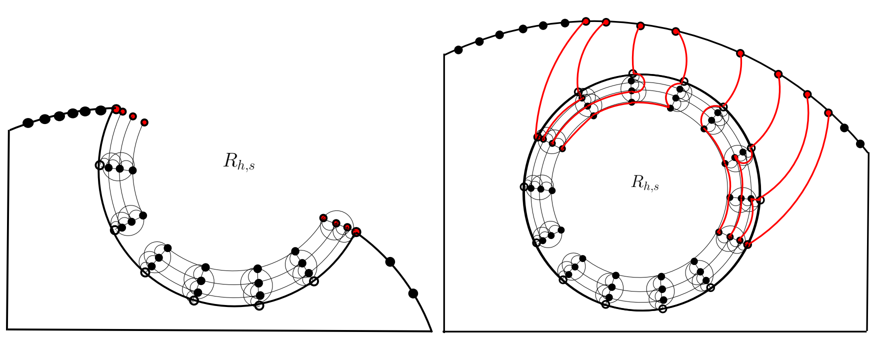

Again we can rename the graph before this new transformation G^p and the graph produced $G^{p\prime}$ and prove the following using the fact the later contains the former as a minor.

Lemma 14. *Let G^p and $G^{p\prime}$ be as above. Then*

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

Final setting. To have an idea of how Δ_e looks like after the previous tranformation, see the left part of the next figure. Clearly, 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 disc in the graph embedding that is bounded by the union ∂E_e and some intersecting holes of $\Delta_1, \ldots, \Delta_r$.

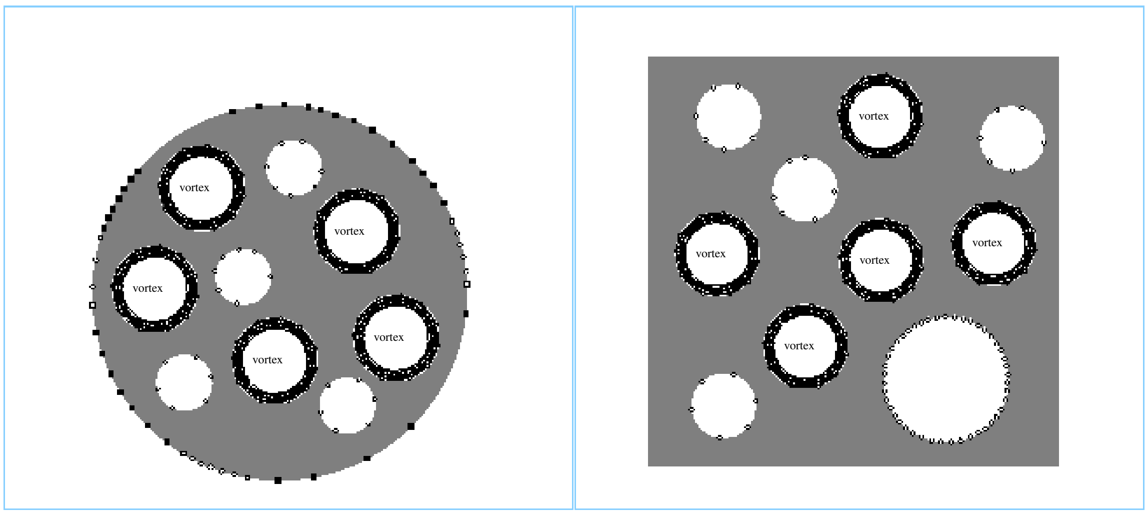

In the current setting, we have a collection of holes in the sphere with ℓ vertices on their borders. We will now count $\texttt{q-paths}_{G^{p\prime}}(E'_e, \mathbf{V'_e})$ for the sets of vertex disjoint paths $\mathbf{paths}_{G^{p\prime}}(E'_e, \mathbf{V_e}')$ between these ℓ vertices.

Tree structure for fixing paths Before we are ready to prove Lemma 10, namely that

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

we need some auxiliary lemmas:

Lemma 15. *[27] 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 15 are called *non-crossing matchings*. A non-crossing matching can be visualized by placing n vertices on a cycle , and connecting matching vertices by 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 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.

Lemma 16. *Let r disjoint empty discs $\Delta_1, \ldots, \Delta_r$ 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,*

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

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 disk in order to apply Lemma 15.

First lets 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 disk 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 $\leq n^{n-2}$ possible spanning trees on n vertices [6]. Here, we have a spanning tree over r vertices, representing the r disks in $\Delta_1, \ldots, \Delta_r$. Then the boundary of each disk has length $\leq n$. 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 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,

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

□

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

Lemma 17. *Let r disjoint empty discs $\Delta_1, \ldots, \Delta_r$ 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 $\leq h$ disjoint discs $R_1, \ldots, R_h$.*

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 $R_1, \ldots, R_h$ such that $P \cap R_j$ $(1 \leq j \leq h)$ is a superposition of h non-crossing matchings. 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,

$$|\mathcal{P}^h_{n,r}| \leq (r^{r-2} \cdot n^{2r} \cdot 2^{rn})^{(h+r)^h}.$$

Proof. We observe that only h arcs of $P \cap R_j$ may cross mutually inside one of R_j $(1 \leq j \leq h)$, but since the entire arc α of P may enter and leave R_j arbitrarily often, we may have more than h mutually crossings of P in R_j. However, we observe that α then cuts the $\mathbb{S}_0 - \{\Delta_1, \ldots, \Delta_r\}$ into several discs. It follows, together with the Helly property of circular arcs, that there are roughly $\frac{3}{2}h + r - 1$ 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 $\frac{3}{2}h + r - 1$ colors.

Furthermore, we observe that two arcs of P may be assigned with 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 $\frac{3}{2}h + r - 1$ colors per R_j and is thus assigned by a h-vector of colors for all R_j $(1 \leq j \leq h)$. With Lemma 16, 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 $(r^{r-2} \cdot n^{2r} \cdot 2^{rn})^{(h+r)^h}$.

□

We can apply Lemma 17 to our terminology:

We say two paths $P_1, P_2 \in \mathbf{paths}_{G^p}(E_e, \mathbf{V_e})$ *cross inside a (h,s)-vertex 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 paths $\mathbf{paths}_{G^p}(E_e, \mathbf{V_e})/\sim$ with $O_H(w)$ endpoints in $\mathbf{V_e}$ crossing inside $O_H(1)$ vertex patterns. Thus, we have proven Lemma 10, namely that

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

Proof of Step 2.h: Taming 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 a 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 18. *Given a h-almost embeddable graph G^a and its smoothly h-almost embeddable graph G^s after removing the apices with 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*

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

Proof. For an edge $e \in T$, we observe: Any path in $G^a[E_e]$ through an apex vertex α passes exactly two neighbors of α. If $G^a[E_e]$ does not contain any neighbor of any $\alpha \in G^a[E_e]$ then $\texttt{q-paths}_{G^a}(E_e, \partial E_e) \leq w^{O_H(1)} \cdot \texttt{q-paths}_{G^s}(E^s_e, \partial E^s_e)$. 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 $w^{O_H(1)}$.

□

Proof of Step 3: Taming the clique-sums

Given a graph an H-minor-free graph G. By Proposition 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 19. *Given a graph G with 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 $i \leq k$, 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 vertices of S_i is 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.

□

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

We now show how to construct the branch decomposition of a 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 $S_i \subseteq$

$V(G_i^a), (i = 1, 2)$ together with the branch decompositions (T_i^a, τ_i^a) and a 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 21. *Let G_1^a and G_2^a have branch decompositions $(T_1^a, \tau_1^a), (T_2^a, \tau_2^a)$ with the maximum width w and for all edges $e \in T_1' \cup T_2'$ let $\texttt{q-paths}_{G'}(E_e, \partial E_e) \leq q$. The previous construction of the branch decomposition (T', τ') of the h-clique-sum G' has width $\leq w + h$ and for all edges $e \in T'$ $\texttt{q-paths}_{G'}(E_e, \partial E_e) \leq q^2$.*

Proof. For all $e \in T'$, ∂E_e has the same cardinality as in $T_1^a \cup T_2^a$. 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_i^a, $\texttt{q-paths}_{G_i^a}(E_e, \partial E_e) \leq q$ and $\texttt{q-paths}_{G_i^a}(\overline{E_e}, \partial\overline{E_e}) \leq q$, and also $\texttt{q-paths}_{G_i^a}(E_e \cap E(S_i), \partial E_e \cap S_i) \leq |L|^{|L|}$, we have that in T $\texttt{q-paths}_{G'}(\overline{E_e}, \partial\overline{E_e}) \leq q \cdot |L|^{|L|} \leq q^2$ for $e \neq e_t$. With Lemma 20, and since $\partial E_{e_{t^i}} \subseteq \partial E_{t^i}$, we get $\texttt{q-paths}_{G'}(E_{e_{t^i}}, \partial E_{e_{t^i}}) \leq q^2$ and $\texttt{q-paths}_{G'}(\overline{E_{e_{t^i}}}, \partial\overline{E_{e_{t^i}}}) \leq q^2$.Deleting leaves from (T', τ') does neither increase the width nor increase the number of path collections. □

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 ≤ 1.5 times the maximum width $O_H(w)$ of the h-almost embeddable graphs $G_1^a, \ldots, G_n^a$. Extending the same arguments of the previous proof, we get the following lemma.

Lemma 22. *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 $e \in T$, $\texttt{q-paths}_G(E_e, \partial E_e) \leq 2^{O_H(w)}$.*

Proof. (Sketch) 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 $e \neq e_{t^{j+1}}$, we get that $\texttt{q-paths}_G(\overline{E_e}, \partial\overline{E_e}) \leq q^2$. With Lemma 20, and with $\partial E_{e_{t^{j+1}}} \subseteq \partial E_{t^{j+1}}$, $\texttt{q-paths}_G(E_{e_{t^{j+1}}}, \partial E_{e_{t^{j+1}}}) \leq q^3$.

□