



HAL
open science

Killing a vortex

Dimitrios M. Thilikos, Sebastian Wiederrecht

► **To cite this version:**

Dimitrios M. Thilikos, Sebastian Wiederrecht. Killing a vortex. FOCS 2022 - 63rd IEEE Annual Symposium on Foundations of Computer Science, Oct 2022, Denver, United States. pp.1069-1080, 10.1109/FOCS54457.2022.00104 . hal-04043027

HAL Id: hal-04043027

<https://hal.science/hal-04043027v1>

Submitted on 17 Oct 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Killing a Vortex¹

Dimitrios M. Thilikos^{2,3}

Sebastian Wiederrecht^{4,5,6}

Abstract

The Graph Minors Structure Theorem of Robertson and Seymour asserts that, for every graph H , every H -minor-free graph can be obtained by clique-sums of “almost embeddable” graphs. Here a graph is “almost embeddable” if it can be obtained from a graph of bounded Euler-genus by pasting graphs of bounded pathwidth in an “orderly fashion” into a bounded number of faces, called the *vortices*, and then adding a bounded number of additional vertices, called *apices*, with arbitrary neighborhoods. Our main result is a full classification of all graphs H for which the use of vortices in the theorem above can be avoided. To this end we identify a (parametric) graph \mathcal{S}_t and prove that all \mathcal{S}_t -minor-free graphs can be obtained by clique-sums of graphs embeddable in a surface of bounded Euler-genus after deleting a bounded number of vertices. We show that this result is tight in the sense that the appearance of vortices cannot be avoided for H -minor-free graphs, whenever H is not a minor of \mathcal{S}_t for some $t \in \mathbb{N}$.

Using our new structure theorem, we design an algorithm that, given an \mathcal{S}_t -minor-free graph G , computes the generating function of all perfect matchings of G in polynomial time. Our results, combined with known complexity results, imply a complete characterization of minor-closed graph classes where the number of perfect matchings is polynomially computable: They are exactly those graph classes that do not contain every \mathcal{S}_t as a minor. This provides a *sharp* complexity dichotomy for the problem of counting perfect matchings in minor-closed classes.

Keywords: Perfect Matchings, Permanent, Pfaffian Orientations, Graph Minors, Counting Algorithms, Graph Parameters.

¹An earlier version of this paper has appeared at FOCS 2022 [53].

²LIRMM, Université de Montpellier, CNRS, Montpellier, France. Email: sedthilk@thilikos.info.

³Supported by the ANR projects DEMOGRAPH (ANR-16-CE40-0028), ESIGMA (ANR-17-CE23-0010), and the French-German Collaboration ANR/DFG Project UTMA (ANR-20-CE92-0027).

⁴LIRMM, Université de Montpellier, CNRS, Montpellier, France and Discrete Mathematics Group, Institute for Basic Science, Daejeon, South Korea. Email: sebastian.wiederrecht@gmail.com

⁵Supported by the ANR project ESIGMA (ANR-17-CE23-0010).

⁶Supported by the Institute for Basic Science (IBS-R029-C1).

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 3 |
| 1.1 | Some history | 3 |
| 1.2 | Counting matchings in minor-closed graph classes | 4 |
| 1.3 | Our main result | 5 |
| 1.4 | The vga-lattice | 6 |
| 1.5 | Our approach | 9 |
| 2 | Definitions and preliminary results | 11 |
| 2.1 | The Graph Minors Structure Theorem | 11 |
| 2.2 | A refined version of Proposition 20 | 19 |
| 3 | Combinatorial results | 22 |
| 3.1 | Advancing through a nest | 22 |
| 3.2 | The death of a vortex | 25 |
| 3.3 | Excluding a shallow vortex minor | 37 |
| 3.4 | Proof of the combinatorial lower bound | 40 |
| 4 | Algorithmic consequences | 43 |
| 4.1 | Dynamic programming for the generating function | 43 |
| 4.2 | The complexity lower bound | 60 |
| 5 | Conclusion | 61 |

1 Introduction

We consider the problem $\#\text{PERFECT MATCHING}$, asking for the number of perfect matchings, on minor-closed graph classes. The first polynomial algorithm for this problem was given for the class of planar graphs by Kasteleyn in 1961 [26], motivated by the dimer problem in Theoretical Physics [25–27, 51] (see also the results by Temperley and Fisher [51]). For this algorithm, Kasteleyn introduced the celebrated FKT-method and the concept of Pfaffian orientations.

Using these concepts as a departure point, the tractability horizon has been extended to several minor-closed graph classes, further than planar graphs, and it was an open problem whether this horizon contained all minor-closed graph classes. A negative answer to this question was given by Curticapean and Xia in [11] who proved that the classic result of Valiant on the $\#\text{P}$ -completeness of $\#\text{PERFECT MATCHING}$ holds even when restricted to K_8 -minor free graphs.

In this paper, we completely resolve the complexity of $\#\text{PERFECT MATCHING}$ on minor-closed graph classes by providing a sharp characterization of the classes for which the problem is tractable.

1.1 Some history

Given a $n \times n$ matrix $A = (a_{i,j})$ the *permanent* and the *determinant* of A are defined as

$$\text{perm}(A) = \sum_{\sigma \in S_n} \prod_{i=1}^n a_{i,\sigma(i)} \quad \text{and} \quad \det(A) = \sum_{\sigma \in S_n} \text{sgn}(\sigma) \cdot \prod_{i=1}^n a_{i,\sigma(i)}$$

respectively, where S_n is the set of all possible permutations of the set $[n] = \{1, \dots, n\}$ and $\text{sgn}(\sigma)$ is the sign of the permutation $\sigma \in S_n$. The permanent is closely related to the $\#\text{PERFECT MATCHING}$ problem as the number of perfect matchings of a bipartite graph B is equal to $\text{perm}(A'_B)$ where A'_B is the biadjacency matrix¹ of B .

In 1913 György Pólya [37] asked when it is possible to change the signs of the entries of a binary $n \times n$ matrix $A = (a_{i,j})$ and obtain a new matrix A' where $\text{perm}(A) = \det(A')$. Notice that, in such cases, the computation of the permanent is reduced to the computation of the determinant that, in turn, can be computed in polynomial time (see [1] for an exposition on the permanent versus determinant problem).

Kasteleyn in 1961, in an attempt to solve the dimer problem (originated from Statistical Physics [26, 27, 51]) introduced the concept of a Pfaffian graph: a matchable² graph G is *Pfaffian* if it admits an orientation \vec{G} such that every conformal² cycle C of G has an odd number of directed edges in agreement to the orientation \vec{G} , when traversed clockwise. Such an orientation of G is called *Pfaffian*.

A Pfaffian orientation of a graph G implies a scheme to change the signs of the adjacency matrix $A(G)$ such that the determinant of the resulting matrix still yields the number of perfect matchings of G . As a special case of Kasteleyn's approach, Pólya's question can be answered affirmatively for a matrix A if and only if the bipartite graph B that has A as its biadjacency matrix is Pfaffian

¹The *biadjacency matrix* of a bipartite graph $B = (X \dot{\cup} Y, E)$ where $X = \{x_1, \dots, x_n\}$, $Y = \{y_1, \dots, y_n\}$, and $E \subseteq X \times Y$, is the binary $n \times n$ matrix $A = (a_{i,j})$ where $a_{i,j} = 1$ if and only if $\{x_i, y_j\} \in E(B)$.

²A graph is *matchable* if it contains at least one perfect matching. A cycle C of a graph G is *conformal* if it is even and $G - C$ is matchable.

(see [33]). In particular, the requested change of signs follows directly from a Pfaffian orientation \vec{B} of B . Moreover, Kasteleyn proved that planar (matchable) graphs are Pfaffian and gave a polynomial algorithm for computing a Pfaffian orientation of a planar graph. This algorithm is widely known as the FKT-method (making also reference to the authors of [51]) and implies that $\#\text{PERFECT MATCHING}$ is polynomially solvable in planar graphs.

In 1972, Little treated Pólya’s problem by giving a *complete* combinatorial characterization of Pfaffian bipartite graphs. Later, McCuaig, Robertson, Seymour, and Thomas gave a structural characterization of Little’s condition that permitted the design of a polynomial algorithm checking whether a bipartite graph is Pfaffian [34, 38]. This immediately implied a polynomial algorithm for $\#\text{PERFECT MATCHING}$ on Pfaffian bipartite graphs. On the negative side (motivated by the permanent vs determinant problem) Valiant introduced the counting complexity class $\#P$ and proved that $\#\text{PERFECT MATCHING}$ is $\#P$ -complete [57].

See Figure 1 for a timeline of the known results on the $\#\text{PERFECT MATCHING}$ problem.

1.2 Counting matchings in minor-closed graph classes

In what concerns general (i.e. non-bipartite) graphs, Kasteleyn claimed in [26, 27] that his algorithm for $\#\text{PERFECT MATCHING}$ on planar graphs can be extended to graphs of bounded Euler-genus. This was proved by Galluccio and Loebel [17] for orientable surfaces and by Tesler [52] for non-orientable surfaces. The later results were based on a reduction of the problem to the computation of 2^g orientations³, where g is the Euler-genus of the input graph. A unified algorithm for counting the number of perfect matchings on graphs of bounded Euler-genus without the use of orientations was given by Curticapean and Xia in [9].

Notice that graphs of bounded Euler-genus are minor-closed. The emerging question is whether $\#\text{PERFECT MATCHING}$ is polynomially solvable for all minor-closed graph classes. For this, given a finite set of graphs \mathcal{F} , we introduce the notation $\#\text{PERFECT MATCHING}(\mathcal{F})$ for the $\#\text{PERFECT MATCHING}$ problem restricted to graphs excluding all graphs in \mathcal{F} as minors. We stress that every minor-closed graph class can be characterized by the minor-exclusion of some *finite* \mathcal{F} , because of Robertson and Seymour’s theorem [46]. Observe also that $\#\text{PERFECT MATCHING}$ is the same as $\#\text{PERFECT MATCHING}(\emptyset)$, while $\#\text{PERFECT MATCHING}$ on planar graphs is $\#\text{PERFECT MATCHING}(\{K_5, K_{3,3}\})$. Using this notation, advances on the $\#\text{PERFECT MATCHING}$ problem can be described as follows: Valiant proved in [59] that $\#\text{PERFECT MATCHING}(\{K_{3,3}\})$ is polynomially solvable, Straub, Thierauf, and Wagner proved in [50] that $\#\text{PERFECT MATCHING}(\{K_5\})$ is polynomially solvable and later Curticapean in [7] and Eppstein and Vazirani in [16] proved that $\#\text{PERFECT MATCHING}(\mathcal{F})$ is polynomially solvable for every \mathcal{F} containing a minor of a graph that admits a drawing in the plane with at most one crossing.

Recently, Curticapean and Xia strengthened Valiant’s original hardness result by proving that $\#\text{PERFECT MATCHING}(\{K_8\})$ is $\#P$ -complete [10]. This means that the tractability horizon of $\#\text{PERFECT MATCHING}$ does not include all minor-closed graph classes and it lies somewhere

³The original result suggested by Kasteleyn used 4^g many orientations. This was since Kasteleyn worked with orientable genus, hence the number h of handles, while the result we present here refers to the Euler-genus which is $2h + c$, where c is the number of crosscaps involved.

The results in blue concern (weighted) bipartite graphs, corresponding to general matrices. All other of the results concern algorithmic (green) and hardness (red) results on general (weighted) graphs and their adjacency matrices.

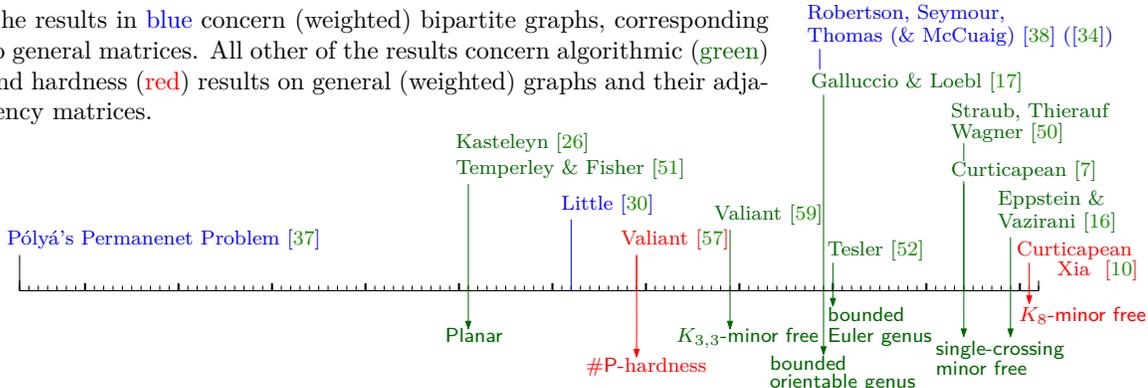


Figure 1: Timeline of the results on the complexity of the $\#$ PERFECT MATCHING problem.

“above” single-crossing minor free graphs and “below” K_8 -minor free graphs. In this paper we completely determine this tractability horizon.

1.3 Our main result

Let G be a graph and let $\mathbf{w} : E(G) \rightarrow \mathbb{Z}$ be a function assigning weights to the edges of G such that $\max\{|\mathbf{w}(e)| \mid e \in E(G)\}$ is bounded by some polynomial function of $|G|$ (that is the number of vertices of G). We refer to such a pair (G, \mathbf{w}) as an *edge-weighted graph* (or simply *weighted graph*). We use $\mathcal{M}(G)$ for the set of all perfect matchings of G . We define the *generating function of perfect matchings* of the weighted graph (G, \mathbf{w}) as

$$\text{PerfMatch}(G, \mathbf{w}) := \sum_{M \in \mathcal{M}(G)} \prod_{e \in M} x^{\mathbf{w}(e)}.$$

Notice that if $\mathbf{1}$ is the weighing function assigning unit weights on the edges of G , then $\text{perm}(A_G) = \text{PerfMatch}(G, \mathbf{1}) = c \cdot x^{|G|/2}$ where $c = |\mathcal{M}(G)|$ is the number of perfect matchings of G . Therefore, any algorithm computing PerfMatch in polynomial time can also serve as a polynomial algorithm for the $\#$ PERFECT MATCHING problem. Moreover, the computation of $\text{PerfMatch}(G, \mathbf{w})$ also permits to solve the EXACT PERFECT MATCHING problem: given an edge-weighted graph and some non-negative integer k , decide whether there is a perfect matching of total weight exactly k . EXACT PERFECT MATCHING was defined by Papadimitriou and Yannakakis in [35], has been extensively studied [21, 60], with applications on DNA sequencing [3] and storage load balancing in blockchain networks [31].

Definition 1 (Shallow vortex grids). The *shallow vortex grid* of order k is the graph \mathcal{S}_k obtained from the product of a cycle $C = (u_1, u_2, \dots, u_{4k})$ on $4k$ vertices with a path $P = v_1, \dots, v_k$ on k vertices (that we call $(k, 4k)$ -cylindrical grid) by adding the edges $\{(u_{4(i-1)+1}, v_1), (u_{4(i-1)+3}, v_1)\}$ and $\{(u_{4(i-1)+2}, v_1), (u_{4(i-1)+4}, v_1)\}$ for every $i \in \{1, \dots, k\}$.

In general, we see the *shallow vortex grid* as a *parametric graph*, that is the sequence $\mathcal{S} = \langle \mathcal{S}_k \rangle_{k \in \mathbb{N}}$. To get a fairly good idea of \mathcal{S}_k , see Figure 2 for a drawing of \mathcal{S}_6 . We also define \mathcal{S} as the graph class consisting of all minors of shallow vortex grids, i.e., all minors of the graphs in \mathcal{S} . (For more on parametric graphs, see [36].)

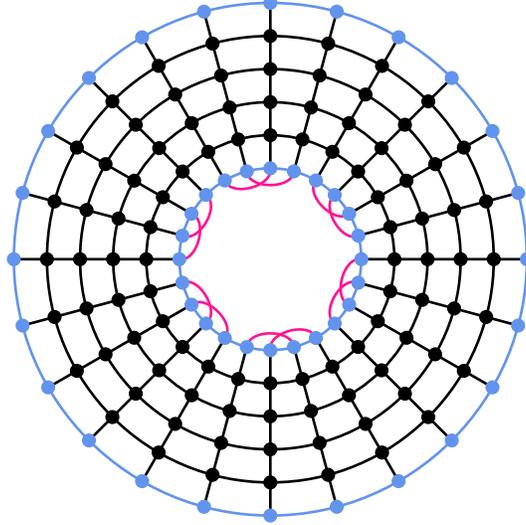


Figure 2: The shallow vortex grid H_6 . The additional six pairs of crossed edges are depicted in red. The two blue cycles are the two extremal cycles of the $(6, 24)$ -cylindrical grid.

Our main result is a complexity dichotomy for $\#\text{PERFECT MATCHING}(\mathcal{F})$, based on the class \mathcal{S} .

Theorem 2. *Let \mathcal{F} be some finite set of graphs. Then $\#\text{PERFECT MATCHING}(\mathcal{F})$ can be solved in polynomial time if $\mathcal{F} \cap \mathcal{S} \neq \emptyset$; otherwise it is $\#\text{P}$ -complete.*

As an example of an application of [Theorem 2](#), K_7 is a minor of \mathcal{S}_{18} (see [Subsection 3.4](#)), therefore $K_7 \in \mathcal{S}$. This implies that $\#\text{PERFECT MATCHING}(\{K_r\})$ is polynomially solvable for $r \leq 7$ (which already answers the open question in [\[10\]](#)) and, as proved in [\[10\]](#), is $\#\text{P}$ -complete for $r \geq 8$. For the general minor-exclusion of a finite set \mathcal{F} , containing graphs of size at most h , one needs to check whether some graph in \mathcal{F} is one of the graphs in \mathcal{S} with at most h vertices.

In order to give some intuition why the minors of graphs as the one in [Figure 2](#) provide the correct dichotomy criterion, we need a brief introduction to the Graph Minors Structure Theorem, proven by Robertson and Seymour in [\[45\]](#) (in this paper we use the notation used in the simpler proof of this theorem that was proposed recently by Kawarabayashi, Thomas, and Wollan in [\[23\]](#) – see also [\[28\]](#)).

1.4 The vga-lattice

A *graph parameter* is a function mapping graphs to non-negative integers. Let \mathfrak{p} and \mathfrak{p}' be two graph parameters. We write $\mathfrak{p} \preceq \mathfrak{p}'$ if there is a function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that, for every graph G , it holds that $\mathfrak{p}(G) \leq f(\mathfrak{p}'(G))$. We also say that \mathfrak{p} and \mathfrak{p}' are *asymptotically equivalent* if $\mathfrak{p} \preceq \mathfrak{p}'$ and $\mathfrak{p}' \preceq \mathfrak{p}$.

Definition 3 (Treewidth). Let G be a graph. A *tree decomposition* of G is a tuple (T, β) where T is a tree and $\beta : V(T) \rightarrow 2^{V(G)}$ is a function, called the *bags* of (T, β) , such that

- i) $\bigcup_{t \in V(T)} \beta(t) = V(G)$,
- ii) for every $e \in E(G)$ there exists $t \in V(T)$ with $e \subseteq \beta(t)$, and
- iii) for every $v \in V(G)$ the set $\{t \in V(T) \mid v \in \beta(t)\}$ induces a subtree of T .

The *width* of a tree decomposition is defined as $\max_{t \in V(T)} |\beta(t)| - 1$ and the *treewidth* of G , denoted by $\text{tw}(G)$, is the minimum width over all tree decompositions for G . The *adhesion* of (T, β) is $\max_{d, t \in E(T)} |\beta(d) \cap \beta(t)|$. For a vertex $t \in V(T)$, the *torso of G at t* is the graph G_t obtained from $G[\beta(t)]$ by turning the sets $\beta(t) \cap \beta(d)$ into cliques, for all neighbors $d \in V(T)$ of t .

Definition 4. Given a graph G , we define $\mathfrak{p}_{\text{vga}}(G)$ as the minimum k such that G has a tree decomposition $\mathcal{D} = (T, \beta)$ where, for every torso G_t of \mathcal{D} , if $|G_t| > k$, then the following holds: There is a set $A \subseteq V(G_t)$, called *apex set*, a surface Σ , and a Σ -decomposition Δ of $G_t - A$ such that:

- (v) Δ has at most $k_{\mathbf{v}}$ vortices, each of depth at most $k_{\mathbf{v}}$,
- (g) Σ has Euler-genus at most $k_{\mathbf{g}}$, and
- (a) $|A| \leq k_{\mathbf{a}}$,

where $\max\{k_{\mathbf{v}}, k_{\mathbf{g}}, k_{\mathbf{a}}\} \leq k$. We postpone the formal definition of a Σ -decomposition of a graph as well the definition of a vortex and its depth to [Subsection 2.1](#) (see [Definition 13](#)). Intuitively, G_t has a Σ -decomposition with k vortices if $G_t = G^{(0)} \cup G^{(1)} \cup \dots \cup G^{(k)}$ where $G^{(0)}$ is a graph embedded in Σ and each vortex $G^{(i)}$ is a graph of bounded pathwidth “attached around” some of the vertices of some face of the embedding of $G^{(0)}$.

[Definition 4](#) has several variants $\mathfrak{p}_{\mathbf{w}}$. Where \mathbf{w} is string of length three obtained from vga by replacing some of “v”, “g”, or “a” with “-”. Given such a string \mathbf{w} the parameter $\mathfrak{p}_{\mathbf{w}}$ is defined via altering [Definition 4](#) by replacing for each $x \in \{\mathbf{v}, \mathbf{g}, \mathbf{a}\}$, where x has been replaced with “-” in \mathbf{w} , the number k_x from the corresponding condition in [Definition 4](#) with 0.

In addition, we derive one more parameter from [Definition 4](#). That is the parameter $\mathfrak{p}_{\text{false}}$ which we obtain by replacing whatever follows the statement “the following holds:” by some false statement. We wish to stress here that $\mathfrak{p}_{\text{---}}$ and $\mathfrak{p}_{\text{false}}$ are different parameters. To see this observe that $\mathfrak{p}_{\text{---}}$ is zero on all planar graphs while $\mathfrak{p}_{\text{false}}(G) = \text{tw}(G) + 1$ for all graph G . This generates nine variants of parameters whose relation with respect to \preceq is depicted in [Figure 3](#). We refer to this lattice of parameters as the *vga-lattice*. We will later fix our attention to $\mathfrak{p}_{\text{-ga}}$, where vortices disappear and therefore $G_t - A$ is just required to be embeddable to a surface of Euler-genus at most k .

On a lower level, the eight parameters above $\mathfrak{p}_{\text{false}}$ can be understood as the different variants for the structure of the “area” which is controlled by a large order *tangle* in the graph. To clarify, tangles are dual objects to *branchwidth*, a parameter which is asymptotically equivalent to treewidth which was introduced in Graph Minors X by Robertson and Seymour [41] to identify substructures in the graph which obstruct small treewidth. The Grid Theorem asserts that every tangle of large order “controls” a large grid minor and the subsequent structure theorems, in essence, describe how the rest of the graph attaches to such a grid minor. This form of attachment happens in two ways; via the highly representative infrastructure of a surface, or through small order separations which

are represented by the decompositions of vortices, the apex set, and the ≤ 3 -clique sums generated by applications of the *Two Paths Theorem*⁴. So, in a sense, the *vga*-lattice classifies graphs with respect to the structure of their high order tangles. This is another way how the special role of p_{false} becomes apparent as this measures the situation where G does not have any high-order tangles. For more information on the lower levels of the *vga*-lattice we refer the reader to [54, 55].

In Figure 3 we depict, in relation to the *vga*-lattice, two more parameters that, when bounded, allow for $\#\text{PERFECT MATCHING}$ to be solved in polynomial time. The first is *apex* where $\text{apex}(G)$ is the minimum number of vertices whose removal from G yields a planar graph and the second is *genus* where $\text{genus}(G)$ is the minimum Euler-genus of a surface where G can be embedded.

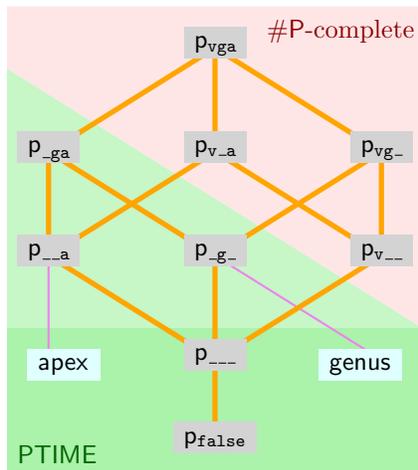


Figure 3: The *vga*-lattice of parameters and the position of the parameters *apex* and *genus* in it. If p and p' are parameters in the above diagram then $p \preceq p'$ if and only if there is a path between p and p' that is “above” p' . The two green/pink-colored areas indicate the complexity of $\#\text{PERFECT MATCHING}$ when restricted to graphs where each of the depicted parameters is bounded. The lower dark green area indicates the current state of the art on polynomial algorithms for $\#\text{PERFECT MATCHING}$.

The *Hadwiger number* of a graph G is the maximum size of a clique minor of G , denoted by $\text{hw}(G)$. The Graph Minors Structure Theorem [23, 45] (in short GMST) can be stated as follows.

Proposition 5. *The graph parameters hw and p_{vga} are asymptotically equivalent.*

(For an alternative (non-parametric) statement of GMST, see Proposition 21.) Notice that all parameters of the *vga*-lattice can be seen as generalizations of treewidth, starting from the lowest level parameter p_{false} , that is asymptotically equivalent to treewidth, to the highest level parameter p_{vga} , that is asymptotically equivalent to the Hadwiger number.

In [43], Robertson and Seymour proved that $p_{\text{---}}(G)$ is asymptotically equivalent to the maximum size of an internally 4-connected⁵ single-crossing⁶ minor of G . Using this, the results in [16]

⁴We discuss these notions in greater detail in Subsection 2.1

⁵A graph is *internally 4-connected* if it is 3-connected and for every separation (A, B) either $|A \setminus B| \leq 1$ or $|B \setminus A| \leq 1$.

⁶A *single-crossing graph* is one that can be drawn in the plane with only one crossing.

and in [7] imply an algorithm that, given a weighted graph (G, \mathbf{w}) , outputs $\text{GenPM}(G, \mathbf{w})$ in time⁷ $\mathcal{O}_k(|G|^{\mathcal{O}(1)})$, where $k = \mathbf{p}_{\text{---}}(G)$. This positions the results of [7, 16] to the second lower level of the \mathbf{vga} -lattice (just above $\mathbf{p}_{\text{false}}$). Apart from this result, the results in [17, 18, 52] imply that $\text{PerfMatch}(G, \mathbf{w})$ can be computed in time $\mathcal{O}_k(|G|^{\mathcal{O}(1)})$ when $k = \text{genus}(G)$. Also, it is easy to see that $\text{PerfMatch}(G, \mathbf{w})$ can be computed in time $|G|^{\mathcal{O}(k)}$ when $k = \text{apex}(G)$ [8, 9]. These three results are the best, so far, algorithmic results about when $\#\text{PERFECT MATCHING}$ can be solved in polynomial time (corresponding to the dark green rectangle of the diagram in Figure 3).

1.5 Our approach

We define a new parameter \mathbf{p} , based on shallow vortex grids \mathcal{S}_k , as follows

$$\mathbf{p}(G) = \max\{k \mid \mathcal{S}_k \text{ is a minor of } G\}.$$

Our main combinatorial result is a vortex-free refinement of the GMST (Proposition 5). In particular \mathbf{p} can be seen as the vortex-free analogue of the Hadwiger number.

Theorem 6. *The graph parameters \mathbf{p} and $\mathbf{p}_{\text{-ga}}$ are asymptotically equivalent.*

Theorem 6 is a min-max theorem indicating that shallow vortex grids can be seen as “universal obstructions” for the parameter $\mathbf{p}_{\text{-ga}}$. Our proof implies that there is a function $f : \mathbb{N} \rightarrow \mathbb{N}$ and an algorithm that, given a graph G and an integer k , either finds the shallow vortex grid \mathcal{S}_k as a minor of G or outputs a tree decomposition of G certifying that $\mathbf{p}_{\text{-ga}}(G) \leq f(k)$. Moreover, this algorithm runs in time $\mathcal{O}_k(|G|^3)$ (see Corollary 45) and, moreover, f is a function that is exponential to some polynomial of k .

The proof of Theorem 6 is the main technical part of this paper. The fact that $\mathbf{p}_{\text{-ga}} \preceq \mathbf{p}$ follows from the following result.

Theorem 7. *There exist functions $\alpha, \gamma : \mathbb{N} \rightarrow \mathbb{N}$ such that every graph G excluding some graph $H \in \mathcal{S}$ as a minor has a tree decomposition (T, β) of adhesion at most $4\alpha(|V(H)|)$ such that for every $t \in V(T)$, if G_t is the torso of G at t , then there exists a set $A \subseteq V(G_t)$ with $|A| \leq \alpha(|V(H)|)$ such that $G_t - A$ has Euler-genus at most $\gamma(|V(H)|)$. Moreover, $\gamma(x) = \text{poly}(x)$, i.e., γ is a polynomial function, while $\alpha(x) = 2^{\text{poly}(x)}$.*

The proof of Theorem 7 is presented in the first three subsections of Section 3. The proof of Theorem 6 is completed in Subsection 3.4, where we show that $\mathbf{p} \preceq \mathbf{p}_{\text{-ga}}$ (Lemma 50).

We prove Theorem 7 by first using a local version of Proposition 5 and then applying the following key observation. Indeed, one can get a stronger version of Proposition 5 which equips each vortex with a sequence of concentric cycles in the obtained Σ -decomposition. These cycles provide quite a lot of infrastructure and are commonly referred to as a *nest*. Now consider the following sequence of observations.

1. If there are not many pairwise disjoint paths from the outer part of the nest to the boundary of the vortex itself, there exists a small set of vertices which may be added to the apex set in order to fully “remove” the vortex from the almost embedding by “pushing” the corresponding subgraph deeper into the tree-decomposition.

⁷We use the notation $h(k, n) = \mathcal{O}_k(g(n))$ to denote that $h(k, n) = \mathcal{O}(f(k) \cdot g(n))$, for some function f .

2. Hence, we may assume that there exist many pairwise disjoint paths which meet all cycles in the nest and end on the boundary of the vortex. If now we cannot find many pairwise disjoint pairs of crossings within the vortex such that these crossings are arranged sequentially on the vortex-boundary we may use the fact that the vortex has a path-decomposition of bounded adhesion to delete a small number of vertices and remove all crossings on the boundary of the vortex which lie inside the vortex. By doing so, however, the parts of the vortex that remain fully connected to the nest do not have any crossings and thus may be fully incorporated into the embedded part of the Σ -decomposition. This procedure is another way of essentially removing the vortex.
3. On the other hand, if we find many pairwise disjoint paths meeting all cycles in the nest and ending on the vortex boundary such that there are many pairwise disjoint pairs of crossings on the endpoints of these paths which are arranged sequentially along the vortex boundary we have, essentially, found a minor model of a shallow vortex grid as depicted in [Figure 2](#).

By applying these three observations to each of the vortices of the Σ -decomposition, we either find a way to enrich the apex set and thereby get rid of all vortices completely, or one of the vortices witnesses the third case above which provides us with the minor we excluded in the assumption.

Based on [Theorem 7](#), we next prove our main algorithmic result.

Theorem 8. *There are an algorithm and a function $f : \mathbb{N} \rightarrow \mathbb{N}$ that, given a weighted graph (G, \mathbf{w}) , where the maximum absolute value of \mathbf{w} is bounded by some polynomial in $|G|$, outputs $\text{PerfMatch}(G, \mathbf{w})$ in time $\mathcal{O}(|G|^{f(k)})$, where $k = \mathfrak{p}_{\text{-ga}}(G)$. Our algorithm assumes that arithmetic operations are performed in constant time.*

The algorithm of [Theorem 8](#) is presented in [Subsection 4.1](#). It performs dynamic programming on the tree decomposition provided by [Theorem 7](#) and combines all the algebraic tools that have been employed so far around Pfaffian orientations for the $\#\text{PERFECT MATCHING}$ problem [[17, 18, 26, 27, 51, 52](#)].

Notice that $\mathcal{F} \cap \mathcal{S} \neq \emptyset$ if and only if there is a constant $c_{\mathcal{F}}$ such that for every \mathcal{F} -minor free graph G it holds that $\mathfrak{p}(G) \leq c_{\mathcal{F}}$. This fact, combined with [Theorem 6](#) and [Theorem 8](#) implies the positive part of [Theorem 2](#).

On the negative side, in [Subsection 4.2](#), we prove the following using as a departing point the complexity lower bound in [[11](#)].

Theorem 9. *For every graph class \mathcal{G} where $\mathcal{S} \subseteq \mathcal{G}$, $\#\text{PERFECT MATCHING}$ is $\#\text{P}$ -complete when its inputs are restricted to the graphs in \mathcal{G} .*

Notice that $\mathcal{F} \cap \mathcal{S} = \emptyset$ if and only if all graphs in \mathcal{S} are \mathcal{F} -minor free. This, combined with [Theorem 9](#), yields the negative part of [Theorem 2](#).

One may ask whether the algorithm of [Theorem 8](#) can be improved to a *fixed parameter* one, that is an algorithm running in time $\mathcal{O}_k(|G|^{\mathcal{O}(1)})$ where $k = \mathfrak{p}_{\text{-ga}}(G)$ (we already mentioned that this is the case when $k = \mathfrak{p}_{\text{---}}(G)$ [[7, 16](#)] or when $k = \text{genus}(G)$ [[17, 18, 52](#)]). Unfortunately, this is not something to expect even for $k = \mathfrak{p}_{\text{--a}}(G)$ (which is lower than $\mathfrak{p}_{\text{-ga}}$ in the vga -lattice). Indeed, it was proved in [[9](#)] that $\#\text{PERFECT MATCHING}$ is $\#\text{W}[1]$ -hard when parameterized by $\text{apex}(G)$. As $\mathfrak{p}_{\text{-ga}} \preceq \mathfrak{p}_{\text{--a}} \preceq \text{apex}$, this hardness result holds also when the parameter is $k = \mathfrak{p}_{\text{--a}}(G)$ or, even

more, $k = \mathfrak{p}\text{-ga}(G)$. This indicates that, under standard computational complexity assumptions, the algorithm of [Theorem 8](#) is asymptotically optimal from the parameterized complexity point of view. We stress that the dependency of the degree of the polynomial on the excluded minor, while it cannot be removed, could however possibly be improved.

2 Definitions and preliminary results

We denote by \mathbb{Z} the set of integers and by \mathbb{R} the set of reals. Given two integers $a, b \in \mathbb{Z}$ we denote the set $\{z \in \mathbb{Z} \mid a \leq z \leq b\}$ by $[a, b]$. In case $a > b$ the set $[a, b]$ is empty. For an integer $p \geq 1$, we set $[p] = [p]$ and $\mathbb{N}_{\geq p} = \mathbb{N} \setminus [0, p - 1]$. Whenever we need a closed interval over the reals we use $[x, y]_{\mathbb{R}}$ to avoid ambiguity. Please note that this only happens on rare occasions.

All graphs considered in this paper are undirected, finite, and without loops or multiple edges. We use standard graph-theoretic notation and we refer the reader to [\[12\]](#) for any undefined terminology.

2.1 The Graph Minors Structure Theorem

We rely heavily on the more refined versions of the GMST from [\[23\]](#) and [\[13\]](#) instead of the original result proven by Robertson and Seymour in the Graph Minors series. As the involved definitions and concepts are highly technical we dedicate this section almost entirely just to their statements and some explanations. In the proof of our main theorem we actually need a stronger version of GMST which has, implicitly, already been proven in [\[23\]](#) and which resembles the main theorem of [\[13\]](#). However, for the purpose of our proofs, in particular regarding algorithmic applications, it is more convenient to derive a synthesis of the results in [\[13\]](#) and those in [\[23\]](#). The resulting statement makes the structure of the vortices within the original structure theorem more accessible and this might be of use to applications other than our own as well.

Our proof strategy is to “exclude” a shallow vortex grid as a minor. This means that whenever we find a clique minor on the same number of vertices as our shallow vortex grid we have found a contradiction. Hence, we may immediately apply the GMST to obtain a, so to speak, preliminary structure which we then can refine. The following definitions come from the framework of [\[23, 28\]](#). Their main purpose is to provide a formal environment in which we can apply our refinement strategy. This is necessary specifically because we will need to work both inside the vortices and outside of them, that is, in the parts of the graph that are properly embedded in some surface.

The Two Paths Theorem. We begin by introducing the *Two Paths Theorem*. This theorem can be seen as the central bridging element that combines a notion of embeddability with a way disjoint paths can be routed within the graph.

Let G be a graph and let $s_1, s_2, t_1, t_2 \in V(G)$. The TWO DISJOINT PATHS PROBLEM (TDPP) with *terminals* s_1, s_2, t_1, t_2 is the question for the existence of two paths P_1 and P_2 such that for both $i \in [2]$ P_i joins s_i and t_i and P_1 and P_2 are vertex disjoint. The characterization for the yes-instances of the TDPP known as the *Two Paths Theorem* plays an integral role in structural graph theory. The statement of the Two Paths Theorem we present here makes use of the concept

of the so called “societies” which play a focal role in [23] and are used extensively in our proofs as well.

There exists a number of different forms for stating the Two Paths Theorem. The version we present here is the one used by Kawarabayashi, Thomas, and Wollan in their proof of the GMST in [23] and it is based on the concept of societies. Roughly speaking, a society is a graph G together with a vertex set $X \subseteq (G)$ which is equipped with a cyclic ordering Ω . This cyclic ordering is meant to indicate that we are interested in some particular topological properties of the set X with regard to its connectivity within G . Through the lens of the Two Paths Theorem, a society is meant to encode a specific subgraph which we would like to embed in a (closed) disk Δ while drawing exactly the vertices of X onto the boundary of Δ . If this is possible, then X forms a noose in the resulting embedding of the entire graph, otherwise the society represents an obstruction to embedability which, if such a thing occurs to often, would allow us to build a large clique minor.

Definition 10 (Society). Let Ω be a cyclic ordering of the elements of some set which we denote by $V(\Omega)$. A *society* is a pair (G, Ω) , where G is a graph and Ω is a cyclic ordering with $V(\Omega) \subseteq V(G)$. A *cross* in a society (G, Ω) is a pair (P_1, P_2) of disjoint paths⁸ in G such that P_i has endpoints $s_i, t_i \in V(\Omega)$ and is otherwise disjoint from $V(\Omega)$, and the vertices s_1, s_2, t_1, t_2 occur in Ω in the order listed.

Hence, (G, s_1, s_2, t_1, t_2) is a yes-instance of TDPP if and only if the society (G, Ω) , where $V(\Omega) = \{s_1, s_2, t_1, t_2\}$ and the vertices occur in Ω in the order listed, has a cross. See Figure 4 for an illustration of a society with a cross.

We will make use of this concept later when we start refining the vortices of our almost embeddings as follows: The entire vortex will be encoded as a society (G, Ω) equipped with some additional infrastructure. We will move along Ω trying to isolate pairwise disjoint crosses over “segments” of consecutive vertices of Ω . In case we can find many such segments, we can use these crosses together with the afore mentioned infrastructure to construct a shallow vortex grid. Otherwise, only few such segments can be found which will allow us to remove all crosses on (G, Ω) by deleting a small vertex set. The Two Paths Theorem will then allow us to embed the remaining parts of (G, Ω) in a way that is compatible with the already embedded part of our graph.

To fully present the Two Paths Theorem, we need to introduce some topological concepts as well.

By a *surface* we mean a compact 2-dimensional manifold with or without boundary. By the classification theorem of surfaces [14], every surface is homeomorphic to the sphere with h handles and c cross-caps added, and the interior of d disjoint closed disks $\Delta_1, \dots, \Delta_d$ removed, in which case the *Euler-genus* of the surface is defined to be $2h + c$. We call the union of the boundaries of the disks Δ_i the *boundary* of the surface and each such boundary is a *boundary component* of the surface. For more information on surfaces and the Classification Theorem see for example [32].

Definition 11. (Drawing on a surface) A *drawing* (with crossings) *in a surface* Σ is a triple $\gamma = (U, V, E)$ such that

- V and E are finite,
- $V \subseteq U \subseteq \Sigma$,

⁸When we say two paths are *disjoint* we mean that their vertex sets are disjoint.

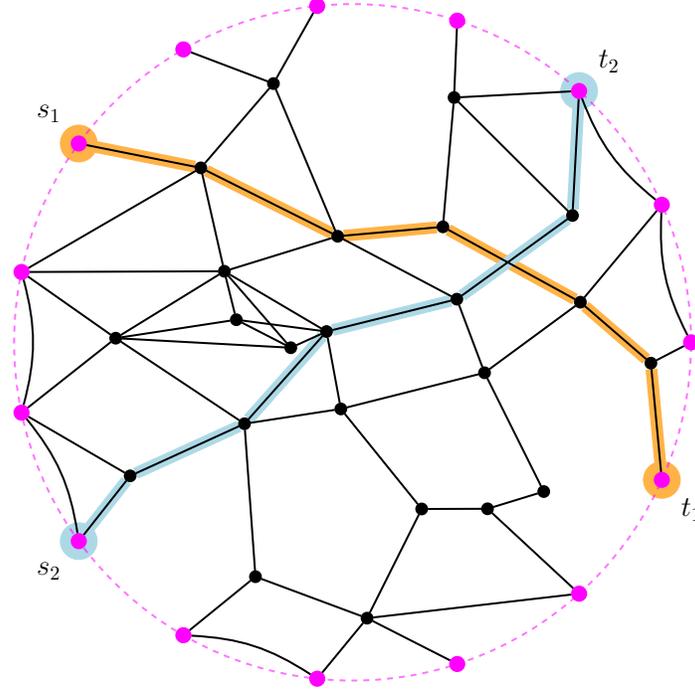


Figure 4: A society (G, Ω) with a cross. The magenta vertices are the vertices of $V(\Omega)$.

- $V \cup \bigcup_{e \in E} e = U$ and $V \cap (\bigcup_{e \in E} e) = \emptyset$,
- for every $e \in E$, either $e = h([0, 1]_{\mathbb{R}}) \setminus \{h(0), h(1)\}$, where $h: [0, 1]_{\mathbb{R}} \rightarrow U$ is a homeomorphism onto its image with $h(0), h(1) \in V$, or $e = h(\mathbb{S}^2 - (1, 0))$, where $h: \mathbb{S}^2 \rightarrow U$ is a homeomorphism onto its image with $h((0, 1)) \in V$, and
- if $e, e' \in E$ are distinct, then $|e \cap e'|$ is finite.

We call the set V , sometime referred to by $V(\gamma)$, the *vertices of γ* and the set E , referred to by $E(\gamma)$, the *edges of γ* . If G is graph and $\gamma = (U, V, E)$ is a drawing with crossings in a surface Σ such that V and E naturally correspond to $V(G)$ and $E(G)$ respectively, we say that γ is a *drawing of G in Σ (possibly with crossings)*. In the case where no two edges in $E(\gamma)$ have a common point, we say that γ is a *drawing of G in Σ without crossings*. In this last case, the connected components of $\Sigma \setminus U$, are the *faces* of γ .

Next we need a more intricate formalization of what an “almost embedding” of a graph is supposed to be. The main issue is that large portions of the graph that we are working with, might be hidden behind 3-clique sums produced by repeated applications of the Two Paths Theorem. To fully encode the almost embedding of a part of the graph G , while keeping the entirety of G accessible, Kawarabayashi, Thomas, and Wollan developed Σ -decompositions in [23, 28].

Definition 12 (Σ -Decomposition). Let Σ be a surface. A Σ -decomposition of a graph G is a pair $\Delta = (\gamma, \mathcal{D})$, where γ is a drawing of G in Σ with crossings, and \mathcal{D} is a collection of closed disks, each a subset of Σ such that

- i) the disks in \mathcal{D} have pairwise disjoint interiors,
- ii) the boundary of each disk in \mathcal{D} intersects γ in vertices only,
- iii) if $\Delta_1, \Delta_2 \in \mathcal{D}$ are distinct, then $\Delta_1 \cap \Delta_2 \subseteq V(\gamma)$, and
- iv) every edge of γ belongs to the interior of one of the disks in \mathcal{D} .

Let N be the set of all vertices of γ that do not belong to the interior of any of the disks in \mathcal{D} . We refer to the elements of N as the *nodes* of Δ . If $\Delta \in \mathcal{D}$, then we refer to the set $\Delta - N$ as a *cell* of Δ . We denote the set of nodes of Δ by $N(\Delta)$ and the set of cells by $C(\Delta)$. For a cell $c \in C(\Delta)$ the set of nodes that belong to the closure of c is denoted by \tilde{c} . For a cell $c \in C(\Delta)$ we define $\sigma_\Delta(c)$, or $\sigma(c)$ if Δ is clear from the context, to be the subgraph of G consisting of all vertices and edges drawn in the closure of c .

We define $\pi_\Delta: N(\Delta) \rightarrow V(G)$ to be the mapping that assigns to every node in $N(\Delta)$ the corresponding vertex of G .

For illustrations of Σ -decompositions consider [Figure 5](#) and [Figure 7](#). In both figures the vertices in magenta and black are the nodes while the grey vertices are drawn in the interiors of cells which means they either sit behind h -clique sums for $h \leq 3$ or belong to a vortex. The cells are depicted as blue shapes and every edge of G , including those between nodes, is drawn within a cell.

Isomorphisms between two Σ -decompositions are defined in the natural way. That is given Σ -decompositions $\Delta = (\gamma, \mathcal{D})$ and $\Delta' = (\gamma', \mathcal{D}')$, the drawing $\gamma = (U, V, E)$ is mapped to a drawing $\gamma' = (U', V', E')$ where the elements of V are in bijection with the elements of V' , similarly for E and E' such that the map between the corresponding graphs is a graph isomorphism, and the elements of \mathcal{D} are mapped to the disks in \mathcal{D}' while agreeing with the map between γ and γ' .

Notice that, in the definition of a Σ -decomposition Δ , the cells c of Δ with $\tilde{c} \neq \emptyset$ correspond to the hyperedges of a hypergraph with vertex set $N(\Delta)$ where \tilde{c} is the set of vertices incident with c . Moreover, this hypergraph can be embedded in Σ such that hyperedges only intersect in common vertices and all vertices are drawn on the boundaries of their hyperedges. As examples of this consider [Figure 5](#) and [Figure 7](#). In both figures, the light blue areas mark the cells which become the hyperedges of some hypergraph. Please note that, in order to not overload the figures, the graphs in the interiors of the cells are chosen to be somewhat minimal with the property that, together with the remaining planar part of the graph, they form an obstruction to planarity. In general no such restriction exists, the graphs within cells can also be planar or arbitrarily complex.

Definition 13 (Vortex). Let G be a graph, Σ be a surface and $\Delta = (\gamma, \mathcal{D})$ be a Σ -decomposition of G . A cell $c \in C(\Delta)$ is called a *vortex* if $|\tilde{c}| \geq 4$. Moreover, we call Δ *vortex-free* if no cell in $C(\Delta)$ is a vortex.

See [Figure 7](#) for an illustration of a part of some Σ -decomposition which includes a vortex.

The reason why the threshold for the boundary size of a vortex is four lies hidden in the *Two Paths Theorem* which we present below ([Proposition 15](#)). Whenever the society defined by the boundary of a cell has at most three vertices, it is impossible to have a cross. This means, in particular, that, if a Σ -decomposition Δ of a graph G has no vortex, one could forget about the interiors of the cells of Δ , for each cell of Δ transform the vertices drawn on its boundary into a clique, and thereby obtain a graph on the vertex set $\pi(N(\Delta))$ which is drawn in Σ without crossings. This is, roughly, the intuition how Σ -decompositions encode a torso.

Definition 14 (Rendition). Let (G, Ω) be a society, and let Σ be a surface with one boundary component B . A *rendition* in Σ of G is a Σ -decomposition ρ of G such that the image under π_ρ of $N(\rho) \cap B$ is $V(\Omega)$ and Ω is one of the two cyclic orderings of $V(\Omega)$ defined by the way the points of $\pi_\delta(V(\Omega))$ are arranged in the boundary B .

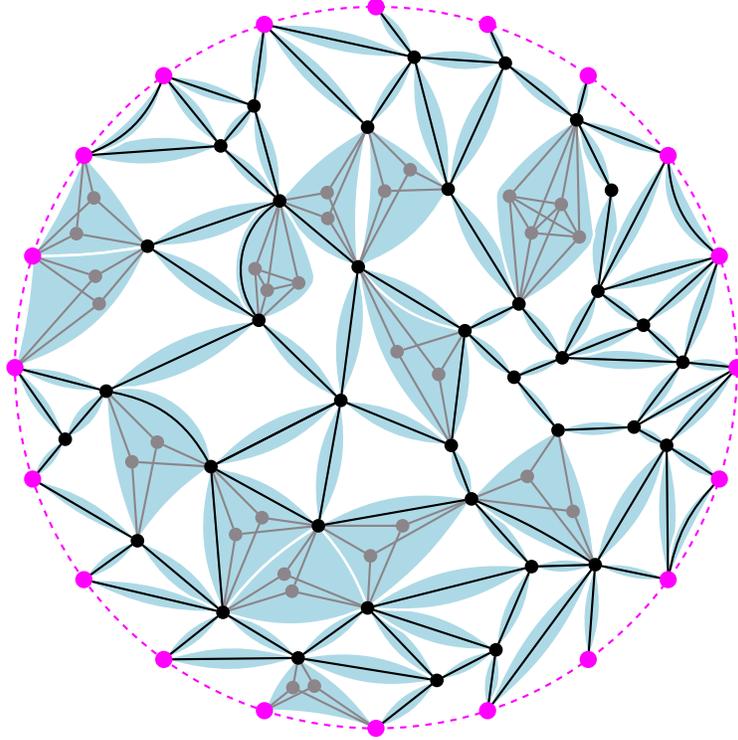


Figure 5: A vortex-free rendition of a society (G, Ω) in the disk.

These technical definitions allow us to state the Two Paths Theorem in the general context of the Graph Minors Structure Theorem as follows. For an illustration of a vortex-free rendition and the absence of a cross on a society (G, Ω) , even when G is non-planar, see [Figure 5](#).

Proposition 15 (Two Paths Theorem, [\[24, 40, 47, 49, 56\]](#)). *A society (G, Ω) has no cross if and only if it has a vortex-free rendition in a disk.*

The Flat Wall Theorem. This section is dedicated to a weaker version of the structure theorem for K_t -minor-free graphs known as the *Flat Wall Theorem* [\[28, 44\]](#). While we do not need the Flat Wall Theorem directly in this paper, it is useful to keep in mind as it is the primary source for the infrastructure we need to construct our shallow vortex grid after processing the vortices of a Σ -decomposition. Indeed, it acts as the base of the construction of a rendition with a bounded number of bounded depth vortices for any K_t -minor-free graph of large treewidth and, as such, it is needed for the statements of the slightly altered versions of lemmas and theorems that we extract from [\[23\]](#).

A *separation* in a graph G is a pair (A, B) of vertex sets such that $A \cup B = V(G)$ and there is no edge in G with one endpoint in $A \setminus B$ and the other in $B \setminus A$. The *order* of (A, B) is $|A \cap B|$.

An $(n \times m)$ -grid is the graph $G_{n,m}$ which is the product of a path $P = u_1, u_2, \dots, u_n$ with n vertices and a path $Q = v_1, v_2, \dots, v_m$ with m vertices. We call the copies of Q in $G_{n,m}$ the *rows* and the copies of P the *columns*. If L is a row of the form $(u_i, v_1), (u_i, v_2), \dots, (u_i, v_m)$ we call it the i th row and for $j \in [m]$ we say that (u_i, v_j) is the j th vertex of the i th row while the edge $\{(u_i, v_j)(u_i, v_{j+1})\}$ is the j th edge of the i th row. For columns we define analogous terminology. An *elementary k -wall*, $k \geq 3$, is obtained from a $(k \times 2k)$ -grid by deleting every odd edge in every odd column and every even edge in every even column. An elementary k -wall W has a unique face whose boundary contains more than six vertices. The *perimeter* of an elementary k -wall is defined to be the subgraph of W induced by all vertices that lie on the unique face with more than six vertices. A k -wall W' is a subdivision of an elementary k -wall W . In other words, W' is obtained by the k -wall W after subdividing each edge of W an arbitrary (possibly zero) number of times. The *perimeter* of W' , denoted by $\text{Per}(W')$, is the subgraph of W' induced by the vertices of the perimeter of W together with the subdivision vertices of the edges of the perimeter of W .

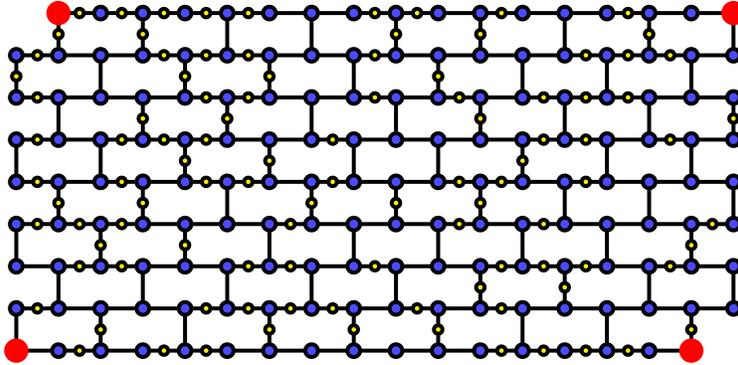


Figure 6: A 9-wall. The subdivision vertices are in yellow and the four corners are in red.

Let G be a graph and $W \subseteq G$ be a wall (here \subseteq denotes the subgraph relation). The *compass* of W in G , denoted by $\text{Compass}_G(W)$, is the subgraph of G induced by the vertices of $\text{Per}(W)$ together with the vertices of the unique component of $G - \text{Per}(W)$ that contains $W - \text{Per}(W)$. The *corners* of an elementary r -wall W are first and the last vertices of the first and the last row respectively. Notice that if W is an r -wall its corners are not uniquely determined. However, we assume to always be given some choice of corners which we identify with the names from above (see for instance the red vertices in Figure 6). The *corner society* of an r -wall W is $(\text{Compass}_G(W), \Omega_W)$ where Ω_W is a cyclic ordering of the corners that agrees with a cyclic ordering of $\text{Per}(W)$.

Definition 16 (Flat Wall). Let $r \geq 2$ be an integer. Let G be a graph and W be an r -wall in G . We say that W is *flat* if $(\text{Compass}_G(W), \Omega_W)$ has a vortex-free rendition in the disk.

We wish to stress that this definition is not exactly the definition of flatness used in [23] and [28]. However, if W is a flat wall in a graph G , then $W - \text{Per}(W)$ can be seen to satisfy the stronger requirements for the flatness in [23, 28] and any wall which is flat in this stronger sense also must be flat in our sense.

To state the Flat Wall Theorem in the terminology of [23, 28] we need to define what it means for a minor to be attached to the infrastructure provided by a given wall.

Let G and H be graphs. If H is a minor of G , then it is possible to find connected subgraphs X_v of G for each $v \in V(H)$ such that $X_u \cap X_v = \emptyset$ if $u \neq v$, and if $uv \in E(H)$, there exists an edge e in G with one end in X_u and the other in X_v . We say that $\{X_v \mid v \in V(H)\}$ for a *minor model* (or simply a *model*) of H in G .

Let W be a wall in G . We say that W *grasps an H minor* if there exists a model $\{X_v \mid v \in V(H)\}$ of H in G together with indices i_v, j_v for each $v \in V(H)$ such that X_v meet the intersection of the i_v th row and j_v th column of W .

Proposition 17 (Flat Wall Theorem [28, 44]). *Let $r, t \geq 1$ be integers, $R := 49152t^{24}(40t^2 + r)$, let G be a graph and let W be an R -wall in G . Then either G has a model of a K_t -minor grasped by W , or there exist a set $A \subseteq V(G)$ of size at most $12288t^{24}$ and an r -subwall $W' \subseteq W - A$ which is flat in $G - A$.*

Paths, Transactions, and Societies of Bounded Depth. In order to be able to apply our refinement strategy to a given vortex together with its infrastructure provided by the refined version of the GMST in [23], we require a particular situation for that infrastructure itself. To achieve this situation we will first apply a preprocessing step which allows us to “push” the infrastructure as close to the actual vortex as possible. For this, we need some additional terminology to handle large linkages between segments of a society.

If P is a path and x and y are vertices on P , we denote by xPy the subpath of P with endpoints x and y . Moreover, if s and t are the endpoints of P and we have fixed an order of the vertices of P , say s is the first and t the last vertex, then xP denotes the path xPt and Px denotes the path sPx . Let P be a path from s to t and Q be a path from q to p . If x is a vertex in $V(P) \cap V(Q)$, then PxQ is the path obtained from the union of Px and xQ . Let $X, Y \subseteq V(G)$. An X -*path* is a path of length at least one with both endpoints in X and internally disjoint from X . In a society (G, Ω) , we write Ω -*path* as a shorthand for a $V(\Omega)$ -*path*. A path is an X - Y -*path* if it has one endpoint in X and the other in Y . Whenever we consider X - Y -paths we implicitly assume them to be ordered starting in X and ending in Y , except if stated otherwise.

Let (G, Ω) be a society. A *segment* of Ω is a set $S \subseteq V(\Omega)$ such that there do not exist $s_1, s_2 \in S$ and $t_1, t_2 \in V(\Omega) \setminus S$ such that s_1, t_1, s_2, t_2 occur in Ω in the order listed, i.e., vertices of $V(\Omega)$ that appear consecutively in Ω . A vertex $s \in S$ is an *endpoint* of the segment S if there is a vertex $t \in V(\Omega) \setminus S$ which immediately precedes or immediately succeeds s in the order Ω . For vertices $s, t \in V(\Omega)$ we denote by $s\Omega t$ the uniquely determined segment with first vertex s and last vertex t . In case t immediately precedes s , we define $s\Omega t$ to be the *trivial segment* $V(\Omega)$.

Let G be a graph. A *linkage* in G is a set of pairwise vertex disjoint paths. In slight abuse of notation, if \mathcal{L} is a linkage, we use $V(\mathcal{L})$ and $E(\mathcal{L})$ to denote $\bigcup_{L \in \mathcal{L}} V(L)$ and $\bigcup_{L \in \mathcal{L}} E(L)$ respectively. Given two sets A and B we say that a linkage \mathcal{L} is a A - B -*linkage* if every path in \mathcal{L} has one endpoint in A and one endpoint in B .

Let (G, Ω) be a society. A *transaction* in (G, Ω) is a linkage \mathcal{L} of Ω -paths in G such that there exist disjoint segments A, B of Ω where the paths in \mathcal{L} are A - B -paths. We define the *depth* of (G, Ω) as the maximum cardinality of a transaction in (G, Ω) .

Let \mathcal{T} be a transaction in a society (G, Ω) . We say that \mathcal{T} is *planar* if no two members of \mathcal{T} form a cross in (G, Ω) . An element $P \in \mathcal{T}$ is *peripheral* if there exists a segment X of Ω containing

both endpoints of P and no endpoint of another path in \mathcal{T} . A transaction is *crooked* if it has no peripheral element.

Finally we will need the following proposition.

Proposition 18 ([23]). *Let (G, Ω) be a society and $p \geq 1$, $q \geq 2$ positive integers. Let \mathcal{P} be a transaction in (G, Ω) of order $p + q - 2$. Then there exists $\mathcal{P}' \subseteq \mathcal{P}$ such that \mathcal{P}' is either a planar transaction of order p or a crooked transaction of order q .*

The Graph Minors Structure Theorem. Next we present two different statements, both fit to capture the global structure of H -minor-free graphs. The first one focusses on the structure relative to a wall and thus can be seen as a local extension of the Flat Wall Theorem, hence we call this one the *Local Structure Theorem*. The second one is the Graph Minors Structure Theorem that completely describes the structure of H -minor-free graphs in terms of graphs of bounded Euler-genus with a bounded number of bounded depth vortices, clique sums and apex vertices. The corresponding graph parameter is asymptotically equivalent to $\mathfrak{p}_{\text{vga}}$ and represents the global maximum of the vga-lattice (see Figure 3).

Definition 19 (Vortex Societies and Breadth and Depth of a Σ -Decomposition). Let Σ be a surface and G be a graph. Let $\Delta = (\gamma, \mathcal{D})$ be a Σ -decomposition of G . Every vortex c defines a society $(\sigma(c), \Omega)$, called the *vortex society* of c , by saying that Ω consists of the vertices $\pi_{\Delta}(n)$ for $n \in \tilde{c}$ in the order given by γ . (There are two possible choices of Ω , namely Ω and its reversal. Either choice gives a valid vortex society.). The *breadth* of Δ is the number of cells $c \in C(\Delta)$ which are a vortex and the *depth* of Δ is the maximum depth of the vortex societies $(\sigma(c), \Omega)$ over all vortex cells $c \in C(\Delta)$.

Next we need to combine our definition of flat walls with the idea of Σ -decompositions. This is a necessary step so to be able to relate flat renditions of walls to the drawings provided by Σ -decompositions and thus to impose additional structure onto these drawings.

Let G be a graph and W be a wall in G . We say that W is *flat in a Σ -decomposition Δ of G* if there exists a closed disk $\Delta \subseteq \Sigma$ such that

- the boundary of Δ does not intersect any cell of Δ ,
- $\pi(N(\delta)) \cap \text{Boundary}(\Delta) \subseteq V(\text{Per}(W))$
- for each degree-three vertex v of W such that v is not mapped to a member of $N(\Delta)$ by π , let $c_v \in C(\Delta)$ be the cell with $v \in V(\sigma(c_v))$. Then, for all pairs of such distinct degree-three vertices u, v of W the cells c_u and c_v are disjoint and $c_v \subseteq \Delta$,
- no cell $c \in C(\Delta)$ with $c \subseteq \Delta$ is a vortex, and
- $W - V(\text{Per}(W))$ is a subgraph of $\bigcup \{\sigma(c) \mid c \subseteq \Delta\}$.

Let G be a graph, let $r \geq 1$ be an integer, let W' be an $r + 2$ -wall in G , and let $W \subseteq W'$ be the unique r -subwall of W' disjoint from its perimeter. If (A, B) is a separation of G of order at most $r - 1$, then exactly one of the sets $A \setminus B$ and $B \setminus A$ includes the vertex set of a column and a row of W . If it is the set A , we say that A is the *W -majority side of the separation (A, B)* ; otherwise,

we say that B is the W -majority side. For those readers familiar with the concept of tangles [41] (see Definition 42), the orientation of all separations of order at most $r - 1$ induced by the majority side of a wall is exactly the tangle induced by the wall.

To apply the intuition of this “orientation” of separations to Σ -decompositions recall that the boundary of every non-vortex cell contains at most three vertices. Moreover, these boundaries define separations between the part of G which is drawn in the interior of a single cell and the part of G which is drawn outside. To link a fixed wall W in a graph G with a Σ -decomposition we make use of this observation as follows. Let Σ be a surface and $\Delta = (\gamma, \mathcal{D})$ be a Σ -decomposition of G . We say that Δ is W -central if there is no cell $c \in C(\Delta)$ such that $V(\sigma(c))$ includes the W -majority side of a separation of G of order at most $r - 1$. Similarly, let $Z \subseteq V(G)$, $|Z| \leq r - 1$, let Σ' be a surface and Δ' be a Σ' -decomposition of $G - Z$. Then Δ' is a W -central decomposition of $G - Z$ if for all separations (A, B) of order at most $r - |Z| - 1$ such that $B \cup Z$ is the majority side of the separation $(A \cup Z, Z \cup B)$ of G , there is no cell $c \in C(\Delta')$ such that $V(\sigma_{\Delta'}(c))$ contains B .

Proposition 20 (Local Structure Theorem [23, 45]). *Let $r, p \geq 0$ be integers, and let $R := 49152p^{24}r + p^{10^7 p^{26}}$. Let G be a graph and let W be an R -wall in G . Then either G has a model of a K_p -minor grasped by W , or there exists a set $A \subseteq V(G)$ of size at most $p^{10^7 p^{26}}$, a surface Σ of Euler-genus at most $p(p + 1)$, a W -central Σ -decomposition Δ of $G - A$ of depth at most $p^{10^7 p^{26}}$ and breadth at most $2p^2$, and an r -subwall $W' \subseteq W - A$ which is flat in Δ .*

From the local structure theorem, a global version can be derived. The way this is usually done is following a balanced separator argument. That is, one fixes a small set of vertices X and either finds a small balanced separator for it, which allows to continue the construction in each of the resulting components, or no such separator exists. In the second case, X witnesses large treewidth which allows one to deduce the existence of a big wall and thus, Proposition 20 can be applied. See the proof of Theorem 44 for a variant of this proof in full detail.

An α -near embedding of a graph G in a surface Σ of depth k and breadth t is a pair (Δ, A) such that $A \subseteq V(G)$, $|A| \leq \alpha$, and Δ is a Σ -decomposition of $G - A$ of depth⁹ k and breadth t such that for every $c \in C(\Delta)$ which is not a vortex, $V(\sigma(c))$ induces a clique in G and $V(\sigma(c)) \subseteq \pi(N(\Delta))$.

Proposition 21 (Global Structure Theorem [23, 45]). *There exists a constant c that satisfies the following. Let $p \geq 1$ be a positive integer and let G be a graph which does not contain K_p as a minor. Let $\alpha := p^{18 \cdot 10^7 p^{26} + c}$. Then G has a tree decomposition (T, β) of adhesion at most 4α such that for all $t \in V(T)$, if G' is the torso of G at t then G' has an α -near embedding of breadth at most $2p^2$ and depth at most α in a surface of Euler-genus at most $p(p + 1)$.*

2.2 A refined version of Proposition 20

The next step is to derive slightly refined versions of Proposition 20. Towards this goal we first introduce additional definitions from [23] and then describe in some detail how the refined versions follow from the proofs in [23]. A similar variant as the one we state below has been proven by Diestel, Kawarabayashi, Müller, and Wollan in [13]. While their theorem exposes the vortices and

⁹Please note that in [23] there is a difference between the “depth” of a Σ -decomposition and the “width” of an α -near embedding. This difference arises from the fact that after resolving the clique sums, each vortex can be decomposed into a path decomposition of bounded width. This width however is related to the depth of the vortex by a small constant factor.

the corresponding infrastructure in a similar way, for our purposes [Proposition 24](#) below is more convenient to work with. One reason for this is that the new statement is better suited to be incorporated into the balanced separator argument we touched upon above.

Societies and Nests. The proof of [Proposition 20](#) in [\[23\]](#) is based on the systematic study of societies. We start by introducing further definitions.

Definition 22 (Cylindrical Rendition). Let (G, Ω) be a society, $\rho = (\gamma, \mathcal{D})$ be a rendition of (G, Ω) in a disk, and let $c_0 \in C(\rho)$ be such that no cell in $C(\rho) \setminus \{c_0\}$ is a vortex. In those circumstances we say that the triple $(\gamma, \mathcal{D}, c_0)$ is a *cylindrical rendition* of (G, Ω) around c_0 .

Let $\rho = (\gamma, \mathcal{D})$ be a rendition of a society (G, Ω) in a surface Σ . For every cell $c \in C(\rho)$ with $|\tilde{c}| = 2$, we select one of the components of $\text{Boundary}(c) - \tilde{c}$. This selection will be called a *tie-breaker* in ρ , and we will assume that every rendition comes equipped with a tie-breaker. Let Q be either a cycle or a path in G that uses no edge of $\sigma(c)$ for every vortex $c \in C(\rho)$. We say that Q is *grounded* in ρ if either Q is a non-zero length path with both endpoints in $\pi_\rho(N(\rho))$, or Q is a cycle and it uses edges of $\sigma(c_1)$ and $\sigma(c_2)$ for two distinct cells $c_1, c_2 \in C(\rho)$. If Q is grounded we define the *trace* of Q as follows. Let P_1, \dots, P_k be distinct maximal subpaths of Q such that P_i is a subgraph of $\sigma(c)$ for some cell c . Fix an index i . The maximality of P_i implies that its endpoints are $\pi(n_1)$ and $\pi(n_2)$ for distinct nodes $n_1, n_2 \in N(\rho)$. If $|\tilde{c}| = 2$, define L_i to be the component of $\text{Boundary}(c) - \{n_1, n_2\}$ selected by the tie-breaker, and if $|\tilde{c}| = 3$, define L_i to be the component of $\text{Boundary}(c) - \{n_1, n_2\}$ that is disjoint from \tilde{c} . Finally, we define L'_i by pushing L_i slightly so that it is disjoint from all cells in $C(\rho)$. We define such a curve L'_i for all i , maintaining that the curves intersect only at a common endpoint. The *trace* of Q is defined to be $\bigcup_{i \in [k]} L'_i$. So the trace of a cycle is the homeomorphic image of the unit circle, and the trace of a path is an arc in Δ with both endpoints in $N(\rho)$.

See [Figure 7](#) for an illustration of some part of a Σ -decomposition. In this figure we have marked three cycles together with their traces which are closed curves, depicted as dashed red lines. The same figure also illustrates the idea of the tie-breaker as one can see that the traces of the three cycles may alter between “inside” and “outside” the actual drawing of their respective cycle.

Definition 23 (Nest). Let $\rho = (\gamma, \mathcal{D})$ be a rendition of a society (G, Ω) in a surface Σ and let $\Delta \subseteq \Sigma$ be an arcwise connected set¹⁰. A *nest in ρ around Δ of order s* is a sequence $\mathcal{C} = (C_1, C_2, \dots, C_s)$ of disjoint cycles in G such that each of them is grounded in ρ , and the trace of C_i bounds a closed disk Δ_i in such a way that $\Delta \subseteq \Delta_1 \subsetneq \Delta_2 \subsetneq \dots \subsetneq \Delta_s \subseteq \Sigma$.

For an simple example of a cylindrical rendition one may readily consider as the shallow vortex grid \mathcal{S}_6 depicted in [Figure 2](#). The “outer face” of the $(6, 24)$ -cylindrical grid of \mathcal{S}_6 is defining a cyclic ordering Ω of the vertices in its boundary. The closure of the face where the crossings are drawn, is the (unique) vortex cell c_0 of the cylindrical rendition ρ , where other cells correspond to the edges of the $(6, 24)$ -cylindrical grid around c_0 . The disjoint cycles of this cylindrical grid are the cycles of a nest in ρ around c_0 of order six. The depth of the society defined by the cell c_0 is two. Notice that this is essentially the minimum possible depth of a vortex as the society defined by a vortex

¹⁰That is, for every pair of points x and y in Δ there exists a curve ζ which lies completely in Δ and which contains x and y .

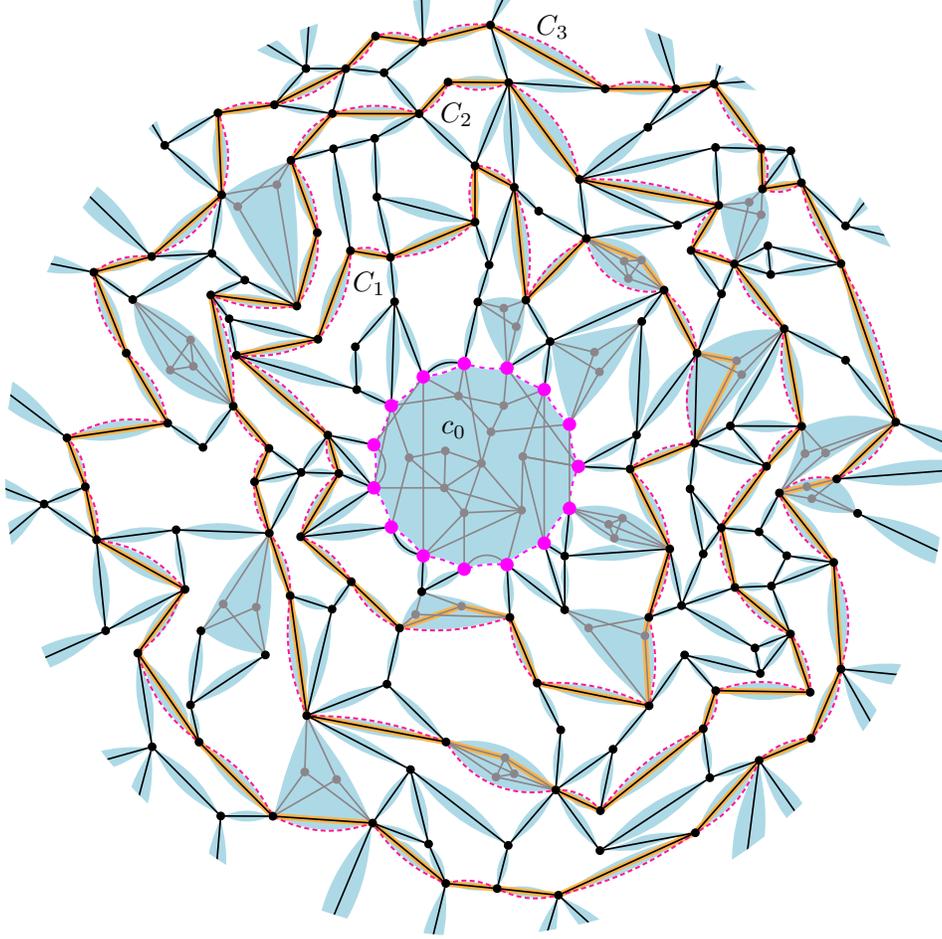


Figure 7: A partial depiction of a Σ -decomposition focussing on a disk around a single vortex cell c_0 together with a nest of order three. Each of the three cycles C_1 , C_2 , and C_3 of the nest comes with its trace (the dashed red lines “mirroring” the behavior of the cycles on the boundary of the cells).

of depth less than two has already a vortex-free rendition in a disk by the Two Paths Theorem (Proposition 15). This justifies the term “shallow” in the term shallow vortex grid that we use.

A more involved example of a nest can be found in Figure 7 where we depict a nest of order three around the vortex cell c_0 .

We are finally ready to state the refined version of Proposition 20. The purpose of this refinement is to establish the nest structure around the vortices in Proposition 20 in such a way that we are able to guarantee large nests for all vortices while also keeping these nests disjoint. A similar statement can be found in [13] presented there in different terminology without explicit bounds, and with a slightly different definition for the depth of a vortex.

Proposition 24 (Local Structure Theorem (with vortex infrastructure) [23, 45]). *Let $r, p \geq 0$ be integers, and let $R := 49152p^{24}r + p^{107p^{26}}$. Let G be a graph and let W be an R -wall in G . Then either G has a model of a K_p -minor grasped by W , or there exists*

- i) a set $A \subseteq V(G)$ of size at most $p^{10^7 p^{26}}$,
- ii) a surface Σ of Euler-genus at most $p(p+1)$,
- iii) a W -central Σ -decomposition $\Delta = (\gamma, \mathcal{D})$ of $G - A$ of depth at most $p^{10^7 p^{26}}$ and breadth at most $2p^2$, and
- iv) an r -subwall $W' \subseteq W - A$ which is flat in Δ .

Moreover, for each vortex cell $c \in C(\Delta)$,

- v) there exists a nest \mathcal{C}_c of order $10^{21} p^{100}$ in Δ around the unique disk $\Delta \in \mathcal{D}$ corresponding to c , and
- vi) let $\Delta_c \subseteq \Sigma$ be the disk bounded by the trace of $C_{10^{21} p^{100}} \in \mathcal{C}_c$, then for each pair of distinct vortex cells $c, c' \in C(\Delta)$ we have that $\Delta_c \cap \Delta_{c'} = \emptyset$.

The reason why [Proposition 24](#) is more convenient to work with than the main theorem of [\[13\]](#) lies in the fact that the latter was formulated for graphs of high treewidth only, while the above theorem only requires the existence of a large wall. While high treewidth and the existence of large walls are famously asymptotically equivalent, it is much easier to derive a global version for all graphs using [Proposition 24](#). Thus, if one is after a global description of graphs excluding a certain minor while maintaining some kind of infrastructure for each vortex, the above theorem is a bit more convenient.

3 Combinatorial results

In [Section 2](#) we presented all concepts and theorems that are necessary for the proof of our main combinatorial result, that is [Theorem 6](#). In this section we prove this result. The direction $\mathfrak{p}\text{-ga} \preceq \mathfrak{p}$ follows from [Theorem 7](#), whose proof spans the first three subsections and the proof of $\mathfrak{p} \preceq \mathfrak{p}\text{-ga}$ is given in [Subsection 3.4](#).

3.1 Advancing through a nest

In this subsection we are going to further refine [Proposition 24](#). Our goal is to describe the structure surrounding a given vortex.

Definition 25 (Rendition of a Nest). Let (G, Ω) be a society with a cylindrical rendition $\rho = (\gamma, \mathcal{D}, c_0)$ and let $\Delta \in \mathcal{D}$ be the disk corresponding to c_0 . Let $\mathcal{C} = \{C_1, \dots, C_s\}$ be a nest in ρ around Δ of order $s \geq 1$. We say that (ρ, G, Ω) is the *rendition of \mathcal{C} around c_0 in ρ* if $V(C_s) = V(\Omega)$.

Suppose $\Delta = (\gamma', \mathcal{D}')$ is a Σ' -decomposition of some graph G' and $c \in C(\Delta)$ is a vortex cell with nest $\mathcal{C}' = \{C'_1, \dots, C'_h\}$ such that the disk Δ' containing c defined by the trace of C'_h does not contain a vortex cell besides c . We denote by $(\Delta_{\mathcal{C}'}, H_{\mathcal{C}'}, \Omega_{\mathcal{C}'})$ the rendition of \mathcal{C} around c induced by Δ' in Δ .

Definition 26 (Tight Nests). Let θ be a positive integer. Let (ρ, G, Ω) be a rendition of a nest \mathcal{C} around a vortex c of depth at most θ in a disk Δ . We say that (ρ, G, Ω) is θ -tight if one of the following is true

- i) there exists a set $Z \subseteq V(G)$ such that $|Z| \leq \theta$ and every $V(\Omega)$ - \tilde{c} -path in G intersects Z , or
- ii) there exists a disk $\Delta' \subseteq \Delta$ whose boundary intersects γ only in nodes with $X = \text{Boundary}(\Delta') \cap N(\rho)$ such that
- Δ' contains c ,
 - there exists a family \mathcal{P} of pairwise disjoint $V(\Omega)$ - \tilde{c} -paths with $|\mathcal{P}| = |\pi(X)| \geq \theta$, and
 - if Ω' is an ordering of $\pi(X)$ induced¹¹ by Δ' , then the society $(G[V(G) \cap \sigma_\rho(\Delta')], \Omega')$ has depth at most 3θ , where $\sigma_\rho(\Delta')$ denotes $\bigcup_{c \in C(\rho), c' \subseteq \Delta'} V(\sigma_\rho(c'))$.

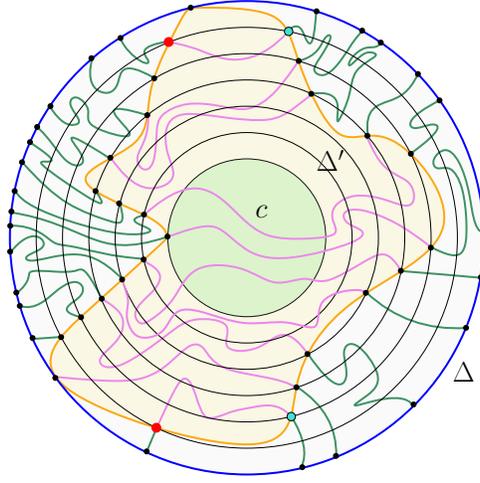


Figure 8: A visualization of the proof of [Theorem 27](#) which finds a tight nest. The **green** paths are the portions of the paths in \mathcal{P} , each joining some vertex of Ω with some vertex of Ω' . The disk Δ' is **orange** and, in the first case, the set Z can be seen as the vertices in its boundary. In the second case, the graph drawn inside the closed disk Δ' , along with the ordering of Ω' , should form a society of depth at most 3θ . Towards a contradiction we assume the existence of $3\theta + 1$ violet paths forming a linkage between two disjoint segments of Ω' (the one segment is between the two **red** vertices and the other between the two **blue** vertices).

The remainder of this section is dedicated to the proof of the following theorem.

Theorem 27 (Local Structure Theorem (with tight vortex infrastructure)). *Let $r, p \geq 0$ be integers, and let $R := 49152p^{24}r + p^{107}p^{26}$. Let G be a graph and let W be an R -wall in G . Then either G has a model of a K_p -minor grasped by W , or the second outcome of [Proposition 24](#) applies along with the following additional condition:*

- vii) for each vortex cell $c \in C(\Delta)$ the rendition $(\Delta_{C_c}, H_{C_c}, \Omega_{C_c})$ of the nest C_c is $p^{107}p^{26}$ -tight.

Towards a proof of [Theorem 27](#), we first show a slightly more general lemma that allows us to choose a new nest “closer” to the vortex whenever none of the two cases of the definition

¹¹This means that Ω' is either the ordering obtained by following along the boundary of Δ' in clockwise order, or its reversal.

of a tight nest holds. [Theorem 27](#) will then follow from iterative applications of this lemma. Moreover, the lemma is constructive in the sense that it only uses some re-routing arguments and Menger's Theorem, hence it can be applied in polynomial time. Therefore, from the polynomial time algorithm in [\[23\]](#) it follows that the Σ -decomposition of [Theorem 27](#) can be found in polynomial time.

Let $\rho = (\gamma, \mathcal{D}, c_0)$ be a cylindrical rendition of a society (G, Ω) and let $\mathcal{C} = \{C_1, \dots, C_s\}$ be a nest around c_0 in ρ . We associate a vector $\mathbf{v}_{\mathcal{C}} \in \mathbb{N}^s$ with \mathcal{C} as follows. For each $i \in [0, s-1]$ let \mathbf{v}_i be the number of nodes of ρ which are contained in the disk that contains c_0 and is bounded by the trace of C_{i+1} . The vector $\mathbf{v}_{\mathcal{C}'}$ is defined analogously.

Let $\mathcal{C}' = \{C'_1, \dots, C'_s\}$ be a nest around c_0 in ρ . We now write $\mathcal{C} < \mathcal{C}'$ if $\mathbf{v}_{\mathcal{C}} <_{\text{lex}} \mathbf{v}_{\mathcal{C}'}$, that is if $\mathbf{v}_{\mathcal{C}}$ is lexicographically smaller than $\mathbf{v}_{\mathcal{C}'}$.

Lemma 28. *Let $\theta \geq 2$ and s be positive integers with $s \leq \frac{\theta}{2}$. Let $(\rho = (\gamma, \mathcal{D}, c_0), G, \Omega)$ be a rendition of the nest \mathcal{C} of order s around the vortex c_0 of depth at most θ in ρ . Then either \mathcal{C} is θ -tight, or there exists a nest $\mathcal{C}' < \mathcal{C}$ of order s around c_0 within G .*

Proof. Let \mathcal{P} be a maximum family of pairwise disjoint $V(\Omega)$ - \tilde{c}_0 -paths in G . If $|\mathcal{P}| \leq \theta$ then, by Menger's Theorem, there exists a set $Z \subseteq V(G)$ of size at most θ which meets all $V(\Omega)$ - \tilde{c}_0 -paths in G and thus (ρ, G, Ω) is θ -tight.

Hence, we may assume $|\mathcal{P}| \geq \theta + 1$. Again by Menger's Theorem we can find a set $X \subseteq V(G)$ with $|X| = |\mathcal{P}|$ since \mathcal{P} such that X contains a vertex of every path in \mathcal{P} since \mathcal{P} is maximum. In particular, we may assume $X \subseteq \pi(N(\rho))$, that is, no vertex of X is drawn in the interior of a cell. In the following, in a slight abuse of notation, we will identify the sets X and $\pi^{-1}(X)$. Since $G - c_0$ has a vortex-free rendition in the disk and there is no $V(\Omega)$ - \tilde{c}_0 -path in $G - X - c_0$ we may find a disk Δ whose boundary intersects $N(\rho)$ exactly in X , is otherwise disjoint from γ , and $c_0 \subseteq \Delta$. Let Ω' be an ordering of the vertices in X obtained by traversing the boundary of Δ in an arbitrarily chosen direction. If the society $(G[V(G) \cap \sigma_\rho(\Delta)], \Omega')$ has depth at most 3θ we are done since this would mean that (ρ, G, Ω) is θ -tight.

Hence, we may assume that there exist disjoint segments I'_1 and I'_2 of Ω' and a family \mathcal{L} of at least $3\theta + 1$ pairwise disjoint I'_1 - I'_2 -paths in $G' := G[V(G) \cap \sigma_\rho(\Delta)]$ (see [Figure 8](#)). Suppose \mathcal{L} contains a crooked transaction of order $\theta + 1$. This means that we can select subpaths of some paths in \mathcal{L} to obtain a crooked transaction on the vortex society $(\sigma(c_0), \Omega_0)$ and therefore contradicts the assumption that c_0 is of depth at most θ . Hence, by [Proposition 18](#), \mathcal{L} contains a planar transaction \mathcal{L}_1 of order at least $2\theta + 1$. Moreover, if $\theta + 1$ paths of \mathcal{L}_1 contain an edge of $\sigma(c_0)$, then we can find a transaction of order $\theta + 1$ in $(\sigma(c_0), \Omega_0)$. Hence, there exists a family $\mathcal{L}_2 \subseteq \mathcal{L}_1$ of size at least $\theta + 1$ which does not contain an edge of $\sigma(c_0)$ nor a vertex of $\sigma(c_0 - \tilde{c}_0)$. Note that \mathcal{L}_2 must be a planar transaction on $(G[V(G) \cap V(\sigma(\Delta))], \Omega')$ since $G - c_0$ has a vortex-free rendition in the disk. Moreover, using the paths from \mathcal{P} and the fact that each vertex of X is an endpoint of some $V(\Omega)$ - X -path that is a subpath of some path in \mathcal{P} we can extend \mathcal{L}_2 to be a planar transaction \mathcal{L}_3 of size $h \geq \theta + 1$ of (G, Ω) .

We now describe how to use \mathcal{L}_3 to obtain the nest \mathcal{C}' . Let I_1 and I_2 be the two segments of Ω projected from I'_1 and I'_2 as follows. For each vertex of Ω' there is a unique vertex of Ω such that these two are linked via a subpath of some path in \mathcal{P} . Given some segment I' of Ω' we say its *projection to Ω* is the smallest segment I of Ω containing all vertices of Ω which are linked to I' by subpaths of paths in \mathcal{P} . Let us number the paths in $\mathcal{L}_3 = \{L_1, L_2, \dots, L_h\}$ according to the

appearance of their starting point on I_1 . From here on we declare the *inside* of a cycle $C \in \mathcal{C}$ to be everything drawn by γ onto the closed disk bounded by the trace of C that contains c_0 , but none of the vertices or edges of C itself. Now let $i \in [s]$ be the smallest number such that the inside of C_i contains a subpath P of some path of \mathcal{L}_3 such that all of P except its endpoints is drawn on the inside of C_i , both endpoints of P lie on C_i , and P is disjoint from all cycles in $\mathcal{C} \setminus \{C_i\}$.

Suppose the number i exists. Then the graph $C_i \cup P$ contains a unique cycle that uses edges of P and whose trace bounds a disk that contains c_0 . Let C'_i be this cycle. Note that, since G is a simple graph, there must exist a vertex $v \in V(C_i)$ which is not contained in C'_i and also not on the inside of C'_i . Hence, $\mathcal{C}' = \{C_1, \dots, C_{i-1}, C'_i, C_{i+1}, \dots, C_s\} < \mathcal{C}$ and we are done.

So we may assume that i does not exist. Let $k \in [h]$ be any number with $L_k \in \mathcal{L}_3$ and let $j_k \in [s]$ be the smallest integer such that L_k contains vertices of C_{j_k} . By assumption L_k cannot contain any edges or vertices drawn on the inside of C_{j_k} since otherwise the number i would exist.

We claim that either the paths in $\{L_z \mid z \in [k-1]\}$ or the paths in $\{L_z \mid z \in [k+1, h]\}$ are disjoint from C_{j_k} . Suppose the claim is false and there exist $a \in [k-1]$ and $b \in [k+1, h]$ such that L_a and L_b both meet C_{j_k} . Let x_1 be the first vertex of L_k on C_{j_k} and let x_2 be the last vertex of L_k on C_{j_k} . Then C_{j_k} is divided into two paths, say Q_1 and Q_2 , both with endpoints x_1 and x_2 . Since \mathcal{L}_3 is a planar transaction in a vortex-free cylindrical rendition of $(G - c_0, \Omega)$ exactly one of the paths, say L_a , can intersect Q_1 and thus L_b must intersect Q_2 . Let Q'_1 be a maximal subpath of $L_a \cap Q_1$ and let Q'_2 be a maximal subpath of $L_b \cap Q_2$. Observe that L_k must avoid both Q'_1 and Q'_2 . Moreover, L_k always stays “in between” L_a and L_b , meaning that the graph $C_h \cup L_a \cup L_b$ contains a unique cycle containing both paths L_a and L_b whose trace bounds a disk that fully contains L_k . It follows that L_k must contain a non-trivial subpath that lies on the inside of C_{j_k} contradicting our observation above. So our claim follows.

For each $j \in [s]$ let $\mathcal{R}_j \subseteq \mathcal{L}_3$ be the collection of all paths $L_k \in \mathcal{L}_3$ such that $j = j_k$. From the previous discussion it follows that $|\mathcal{R}_j| \leq 2$ for all $j \in [s]$. However, since $2s \leq \theta$ and $|\mathcal{L}_3| \geq \theta + 1$ there must exist some $\ell \in [s]$ such that $|\mathcal{R}_\ell| \geq 3$. As this is a contradiction, the case where the number i does not exist cannot occur and our claim follows. \square

Proof of Theorem 27. Let G be a graph and W be an R -wall in G . By Proposition 24 we know that either G contains a K_p -minor model grasped by W , or there exists a set $A \subseteq V(G)$, a surface Σ and a W -central Σ -decomposition Δ satisfying items *i.-vi.* of our theorem. We may assume the nests \mathcal{C}_c of all vortex cells $c \in C(\Delta)$ to be chosen such that $\mathcal{C}_c \leq \mathcal{C}$ for all other nests of order $10^{21}p^{100}$ around c that satisfy point *vi*) of the theorem.

Suppose there exists $c \in C(\Delta)$ such that $\mathcal{C}_c = \{C_1, \dots, C_{p^{107}p^{26}}\}$ is not $p^{107}p^{26}$ -tight. Then Lemma 28 yields the existence of a nest $\mathcal{C} < \mathcal{C}_c$ around c which is contained in the disk bounded by the trace of $C_{p^{107}p^{26}}$ that contains c . Since $\{\mathcal{C}_x \mid x \in C(\Delta) \text{ and } x \text{ is a vortex}\}$ satisfies *vi*) this means that $\{\mathcal{C}_x \mid x \in C(\Delta), x \neq c, \text{ and } x \text{ is a vortex}\} \cup \{\mathcal{C}\}$ must also satisfy *vi*). Hence, we have reached a contradiction to our previous assumption. \square

3.2 The death of a vortex

To create a minor model of a shallow vortex grid we need to be able to find a large number of crosses on the society of a nest such that these crosses are pairwise disjoint and occur in a sequential fashion on the society. In this section we discuss how to either separate the non-planar part of a vortex

completely from its nest by a bounded size set of vertices, and therefore “kill” the vortex, or to find these crosses to construct our minor.

Definition 29 (Cross Over a Segment). Let (G, Ω) be a society and S be a segment of Ω . We say that a pair of $V(G)$ -paths (P_1, P_2) form a *cross over* S if they form a cross over (G, Ω) and all of their four endpoints lie in S .

Definition 30 (Consecutive Crosses). Let (ρ, G, Ω) be a rendition of a nest \mathcal{C} around a vortex c in a disk Δ in ρ . A family $\mathcal{C} = \{(L_1, R_1), \dots, (L_h, R_h)\}$ of crosses over (G, Ω) is said to be *consecutive* if there exist segments I_1, \dots, I_h of Ω such that

- i) for each $i \in [h - 1]$ the last vertex of I_i comes before the first vertex of I_{i+1} and $I_1 \cap I_h = \emptyset$,
- ii) $\bigcup_{i \in [h]} \{L_i, R_i\}$ is a family of pairwise disjoint paths, and
- iii) for each $i \in [h]$, the pair (L_i, R_i) is a cross over I_i .

The main result of this section provides a duality that certifies for every vortex with a large enough nest around it that either this vortex can be completely separated from the properly embedded part of the graph with a small set of vertices or we find the desired minor.

Lemma 31. *Let $t \leq \theta$ be positive integers. There exists a positive universal constant c such that, if $(\rho = (\gamma, \mathcal{D}, c), G, \Omega)$ is a θ -tight rendition of a nest \mathcal{C} , with $|\mathcal{C}| \geq 12t^2 + c$, around a vortex c of depth at most θ in a disk Δ , then one of the following holds.*

- i) *There exists a separation (A, B) of order at most $12\theta(t - 1)$ with $V(\Omega) \cap B \subseteq A \cap B$, such that, if Ω' is the restriction of Ω to $A \setminus B$ then $(G[A \setminus B], \Omega')$ has a vortex-free rendition in the disk, or*
- ii) *G contains the shallow vortex grid of order t as a minor.*

Definition 32 (Patches). Let θ be a positive integer. Let $(\rho = (\gamma, \mathcal{D}, c), G, \Omega)$ be a θ -tight rendition of a nest \mathcal{C} around a vortex c of depth at most θ in a disk Δ in ρ . Moreover, assume that there exists a disk $\Delta' \subseteq \Delta$ with $c \subseteq \Delta'$, whose boundary intersects γ in a set X of nodes only such that

- there exists a family \mathcal{P} of pairwise disjoint $V(\Omega)$ - \tilde{c} -paths with $|\mathcal{P}| = |\pi(X)| \geq \theta$, and
- the society $(G' := G[V(G) \cap V(\sigma(\Delta'))], \Omega')$ has depth at most 3θ where Ω' is an ordering of $\pi(X)$ induced by Δ' .

Finally, let S be a segment of Ω' and $Z \subseteq V(G')$. Let Y be the collection of all vertices in $V(G')$ which are contained in a connected component of $G' - Z$ with a vertex of S , we call Y the *patch of (G', Ω') cut at S by Z* . Let $\hat{G} := G[V(G) \setminus V(G' - Y)]$ and let $\hat{\Omega}$ be the cyclical ordering of $V(\Omega) \cap V(\hat{G})$ induced by Ω . We call $(\hat{G}, \hat{\Omega})$ the *society from (G, Ω) cut by Z and patched at S (with the patch Y)*.

We say that a segment S of Ω' has a cross at a set $Z \subseteq V(G')$ if the society from (G, Ω) cut by Z and patched at S has a cross. (See [Figure 9](#) for a visualization of the above defined concepts.)

Let θ be a positive integer. Let $(\rho = (\gamma, \mathcal{D}, c), G, \Omega)$ be a θ -tight rendition of a nest \mathcal{C} around a vortex c of depth at most θ in a disk Δ in ρ . Moreover, assume that there exists a disk $\Delta' \subseteq \Delta$ whose boundary intersects γ only in vertices with $X = \text{Boundary}(\Delta') \cap N(\rho)$ such that

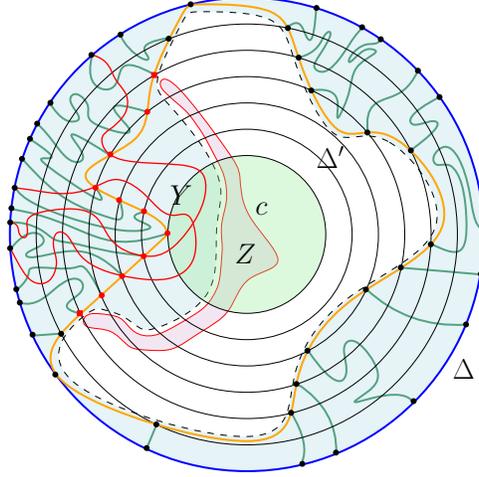


Figure 9: A visualization of the concepts of Definition 32. The green paths are the portions of the paths in \mathcal{P} , each joining some vertex of Ω with some vertex of Ω' . The graph G' is what is cropped by the disk Δ' whose boundary is the orange cycle. The segment S contains the 10 red vertices. The graph \hat{G} contains everything that is not inside the dashed cycle (that is, the graph inside the light blue area). The two red paths form a cross of the segment S at Z .

- there exists a family \mathcal{P} of pairwise $V(\Omega)$ - \tilde{c} -paths with $|\mathcal{P}| = |\pi(X)| \geq \theta$, and
- the society $(G' := G[V(G) \cap V(\sigma(\Delta'))], \Omega')$ has depth at most 3θ where Ω' is an ordering of $\pi(X)$ induced by Δ' .

We say that the tuple $(\theta, \rho, G, \Omega, \mathcal{C}, \Delta, \Delta', \mathcal{P})$ is a θ -suspension of c in (ρ, G, Ω)

Lemma 33. *Let θ be a positive integer. Let $(\rho = (\gamma, \mathcal{D}, c), G, \Omega)$ be a θ -tight rendition of a nest \mathcal{C} around a vortex c of depth at most θ in a disk Δ in ρ and let $(\theta, \rho, G, \Omega, \mathcal{C}, \Delta, \Delta', \mathcal{P})$ be a θ -suspension of c in (ρ, G, Ω) .*

Now let S_1, \dots, S_ℓ be pairwise disjoint segments of Ω' such that for each $i \in [\ell]$ there exists a set $Z_i \subseteq V(G')$ separating S_i from $V(\Omega') \setminus S_i$. For each $i \in [\ell]$ let (G'_i, Ω'_i) be the society from (G', Ω') cut by $Z := \bigcup_{i \in [\ell]} Z_i$ and patched at S_i with the patch Y_i .

If for each $i' \in [\ell]$ the society $(G'_{i'}, \Omega'_{i'})$ has a cross, then for every $i \in [\ell]$ there exists a segment I_i of Ω and a pair of $V(\Omega)$ -paths (L_i, R_i) such that

- i) if $i \neq j \in [\ell]$ then $I_i \cap I_j = \emptyset$,*
- ii) if $i \neq j \in [\ell]$ then $(V(L_i) \cup V(R_i)) \cap (V(L_j) \cup V(R_j)) = \emptyset$,*
- iii) (L_i, R_i) forms a cross over I_i , and*
- iv) L_i and R_i intersect every cycle in \mathcal{C} .*

See Figure 10 for an illustration of the statement of the above lemma.

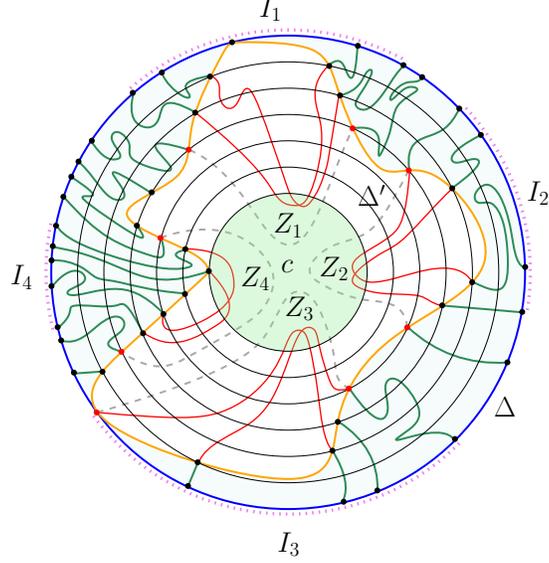


Figure 10: A visualization of Lemma 33 for $\ell = 4$. The green paths are the portions of the paths in \mathcal{P} , each joining some vertex of Ω with some vertex of Ω' . The segments of Ω are delimited by the red vertices, the dashed grey lines represent the separating sets Z_1, \dots, Z_4 , the crossing paths in $(G'_i, \Omega'_i), i \in [4]$ are red and the magenta dotted lines indicate the segments I_1, \dots, I_4 .

Proof. Since (G'_i, Ω'_i) has a rendition in the disk with a single vortex, for every $i \in [\ell]$ we find paths Q_1^i and Q_2^i for both $i \in [2]$ such that (Q_1^i, Q_2^i) forms a cross over (G'_i, Ω'_i) . Note that this means that Q_1^i and Q_2^i both contain an edge that is drawn in the interior of c . Let J'_i be the smallest segment of Ω'_i such that (Q_1^i, Q_2^i) is a cross over J'_i and let J_i be the smallest segment of Ω' containing J'_i which is disjoint from all $J'_j, i \neq j \in [\ell]$. Finally, for each $i \in [\ell]$ and each $j \in [2]$ we denote by $x_{i,j}$ and $y_{i,j}$ the two endpoints of Q_j^i such that $x_{i,j}$ comes before $y_{i,j}$ in the order Ω' . Moreover, we may assume Q_1^i and Q_2^i to be numbered such that $x_{i,1}$ comes before $x_{i,2}$ on Ω' .

Notice that $|J_i| \geq 4$ for all $i \in [\ell]$. Moreover, since (G, Ω) is θ -tight, from every vertex $v \in V(\Omega') = X$ there is a path $P_v \in \mathcal{P}$ joining v to a vertex of $V(\Omega) = V(C_{|C|})$ and all of these paths are internally disjoint from G' .

For every $i \in [\ell]$ let us choose the paths $\mathcal{P}_i := \{P_{x_{i,1}}, P_{y_{i,1}}, P_{x_{i,2}}, P_{y_{i,2}}\} \subseteq \mathcal{P}$. Notice that $\bigcup_{i \in [\ell]} \mathcal{P}_i$ is still a family of pairwise disjoint paths. Moreover, $Y_i \cap Y_j = \emptyset$ for all $i \neq j \in [\ell]$ by the definition of the Y_i and thus $\bigcup_{i \in [\ell]} \{Q_1^i, Q_2^i\}$ is also a family of pairwise disjoint paths. At last, notice that the endpoints of $P_{x_{i,1}}$ and $P_{y_{i,1}}$ on $V(\Omega)$ define two segments of Ω , each of which containing an endpoint of one of the two paths $P_{x_{i,2}}$ and $P_{y_{i,2}}$.

Now let I_i be the segment of Ω with endpoints $x_{i,1}$ and $y_{i,2}$ that contains $y_{i,1}$ and $x_{i,2}$. Then let $L_i := P_{x_{i,1}}x_{i,1}Q_1^iy_{i,1}P_{y_{i,1}}$ and $R_i := P_{x_{i,2}}x_{i,2}Q_2^iy_{i,2}P_{y_{i,2}}$. From the observations above it is now clear that each (L_i, R_i) forms a cross over I_i and that $\bigcup_{i \in [\ell]} \{L_i, R_i\}$ is a family of pairwise disjoint paths. Moreover, each L_i and each R_i must contain an edge that is drawn in the interior of c . Since ρ is a rendition in the disk and both L_i and R_i have their endpoints on Ω this implies that each such path meets all cycles of \mathcal{C} . At last, observe that the segments J_i and J_j on Ω' are disjoint if $i \neq j$. Suppose there are $i \neq j \in [\ell]$ such that $I_i \cap I_j \neq \emptyset$. Without loss of generality let us assume that

$x_{i,1}$ comes before $x_{j,1}$ on Ω . Then we have $y_{i,2} \in I_j$ and $x_{j,1} \in I_i$. Let F_1 be the path obtained by starting in $y_{i,2}$, following along $P_{y_{i,2}}$ until we meet $C_1 \in \mathcal{C}$ in the vertex u_1 for the first time, then following along C_1 within $Y_i \cup Z$ until we meet a vertex v_1 of $P_{x_{i,1}}$ and from here following along $P_{x_{i,1}}$ until we meet $x_{i,1}$. By the discussion above F_1 must be completely disjoint from the paths L_j and R_j since we only used vertices in $Y_i \cup Z$ and from the paths L_i and L_j . Let us construct the path F_2 in a similar way by starting in $x_{j,1}$ and ending in $y_{j,2}$ and let u_2 and v_2 be defined analogously to u_1 and v_1 .

It follows that any vertex of $V(F_1) \cap V(F_2)$ must belong to $Z \cap V(C_1)$. To see that this is impossible let us consider two subpaths of C_1 . Let B_1 be the u_1 - v_1 subpath of C_1 that does not contain a vertex of $\bigcup_{h \in [\ell] \setminus \{i\}} Y_h$. Similarly let B_2 be the u_2 - v_2 subpath of C_1 that does not contain a vertex of $\bigcup_{h \in [\ell] \setminus \{j\}} Y_h$. These paths must exist since there always exists an u_h - v_h -path on C_1 that lies in $Y_h \cup Z$ for both $h \in [2]$. The only way where these paths can intersect is, if their endpoints appear on C_1 in the order u_1, u_2, v_1, v_2 or u_1, v_2, v_1, u_2 . In both cases, however, it would follow that $Y_i \cap Y_j \neq \emptyset$ which contradicts our assumptions. Hence, F_1 and F_2 are disjoint. This implies the existence of a cross over Ω which does not contain a single edge drawn in the interior of c , contradicting the fact that ρ is a rendition of (G, Ω) in the disk. Hence, $I_i \cap I_j = \emptyset$ follows. \square

The next step would be to show that, given some integer t , we can either find t crosses as in the lemma above or completely remove the vortex by deleting a number of vertices bound by a function of t and θ . That is, we want to show that in the situation where we cannot find t pairwise disjoint patches, each hosting a cross, we find another set of vertices of bounded size whose deletion allows us to “flatten” all patches. If this case does not occur and we find many patches, each hosting a cross, we would like to “project” these crosses onto the original society (G, Ω) while maintaining vertex-disjointness. The following lemma is the key tool towards the structural results of this paper.

Lemma 34. *Let $t \leq \theta$ be positive integers. Let $(\rho = (\gamma, \mathcal{D}, c), G, \Omega)$ be a θ -tight rendition of a nest \mathcal{C} around a vortex c of depth at most θ in a disk Δ in ρ . Then either*

1. *there exists a separation (A, B) of order at most $6\theta(t-1)$ such that $(G[A \setminus B], \Omega^*)$ has a vortex-free rendition in the disk, where Ω^* is the restriction of Ω to $A \setminus B$ and $V(\Omega) \cap B \subseteq A \cap B$, or*
2. *there exists a θ -suspension $(\theta, \rho, G, \Omega, \mathcal{C}, \Delta, \Delta', \mathcal{P})$ of c in (ρ, G, Ω) and a consecutive family $\{\mathcal{Q}_1, \dots, \mathcal{Q}_t\}$ of t crosses over (G, Ω) together with t pairwise disjoint segments S_1, \dots, S_t of Ω' and a set $Z' \subseteq V(G')$ such that for each $i \in [t]$ the cross \mathcal{Q}_i intersects G' exactly in a cross of Ω' at Z' .*

Proof. As (ρ, G, Ω) is θ -tight, one of two cases may hold. Let us first assume that case *i*) holds. Then there exists a set $Z \subseteq V(G)$ with $|Z| \leq \theta$ such that Z separates $V(\Omega)$ from \tilde{c} . Let A be the collection of all vertex sets of components of $G - Z$ that contain a vertex of $V(\Omega)$ and let $B := V(G) \setminus (A \cup Z)$. Moreover, let Ω^* be the restriction of Ω to A . Then the restriction of ρ to $G[A]$ is a vortex-free rendition of $(G[A], \Omega^*)$ in a disk and $(A \cup Z, Z \cup B)$ is a separation of order at most $\theta \leq 6\theta(t-1)$. Hence, in this case we are done.

So from now on we may assume that case *ii*) from the definition of θ -tightness holds. Hence, we may assume that there exists a disk $\Delta' \subseteq \Delta$ whose boundary intersects γ only in vertices with $X = \text{Boundary}(\Delta') \cap N(\rho)$ such that

- there exists a family \mathcal{P} of pairwise $V(\Omega)$ - \tilde{c} -paths with $|\mathcal{P}| = |\pi(X)| \geq \theta$, and
- the society $(G' := G[V(G) \cap V(\sigma(\Delta'))], \Omega')$ has depth at most 3θ where Ω' is an ordering of $\pi(X)$ induced by Δ'

as in the definition of θ -suspensions.

In what follows we will iteratively construct segments J_i of Ω' together with sets $Z_i \subseteq V(G')$ and $S_i \subseteq V(G')$ such that

- for all $j \in [i]$ we have $|Z_j|, |S_j| \leq 3\theta$,
- for each $j \in [i-1]$ the last vertex of J_j comes before the first vertex of J_{j+1} and $J_1 \cap J_i = \emptyset$,
- for all $j \in [i]$ the set $Z^j := \bigcup_{h \in [j]} Z_h$ separates J_j from $V(\Omega') \setminus J_j$ in G' ,
- for all $j \in [i]$ the society from (G', Ω') cut by Z^i and patched at J_j has a cross, and
- if J^j is the segment of Ω' whose first vertex is the first vertex of J_1 and whose last vertex is the last vertex of J_j , while $S^j := \bigcup_{h \in [j]} S_h$, then for all $j \in [i]$ the society from (G', Ω') cut by $Z^j \cup S^j$ and patched at J^j has a vortex-free rendition in the disk.

In case this iterative process stops before we complete the step $i = t$ we will find the required separation. Otherwise, an application of [Lemma 33](#) yields the required t consecutive crosses.

Let λ be a linear ordering of $V(\Omega')$ obtained from Ω' and let us denote by x_1 the smallest vertex with respect to λ . If λ' is any restriction of λ to some vertex set U we write $V(\lambda')$ for the set U . For any restriction λ' of λ and any property defined for a vertex we say that a vertex u is the *smallest* with this property in λ' if no vertex $v \in V(\lambda')$ with $\lambda'(v) < \lambda'(u)$ has the property, but u does. For any two vertices $u, v \in V(\lambda)$ we define $J_{u,v} := \{w \in V(\lambda) \mid \lambda(u) \leq \lambda(w) \leq \lambda(v)\}$. Note that $J_{u,v}$ is empty if $\lambda(v) < \lambda(u)$. Moreover, $J_{u,v}$ always defines a segment of Ω' .

Since (G', Ω') has depth at most 3θ , for any segment I of Ω' there exists a set Y of size at most 3θ separating I from $V(\Omega') \setminus I$ in G' . For any pair $u, v \in V(\Omega')$ and an already given set Z^h , where $h \in \mathbb{Z}$, we denote by $Y_{u,v}^h$ a set of order at most 3θ such that $Z^h \cup Y_{u,v}^h$ separates $J_{u,v}$ from $V(\Omega') \setminus J_{u,v}$ and $Y_{u,v}^h$ only contains vertices of $G' - Z^h$ that belong to a component of $G' - Z^h$ which contains a vertex of $J_{u,v}$.

Let us fix $Z^0 := \emptyset$. Now let z_1 be the smallest vertex of $V(\lambda)$ such that the society from (G', Ω') cut by Y_{x_1, z_1}^0 and patched at J_{x_1, z_1} has a cross and let y_1 be the immediate predecessor of z_1 with respect to λ . Note that in case z_1 does not exist (G', Ω') has a vortex-free rendition in the disk and thus so does (G, Ω) which would complete our proof with the separation $(V(G), \emptyset)$. Let us now set $J_1 := J_{x_1, z_1}$, $Z_1 := Y_{x_1, z_1}^0$, and $S_1 := Y_{x_1, y_1}^0$. It follows from the minimality of z_1 that the society from (G', Ω') cut by $Z_1 \cup S_1$ and patched at J_1 has a vortex-free rendition in the disk. Hence, the first iteration is complete.

Now assume for some $\ell \leq t-1$ the vertices x_i, y_i, z_i together with the segments J_{x_i, z_i} and the sets $Z_i = Y_{x_i, z_i}^{i-1}$ and $S_i = Y_{x_i, y_i}^{i-1}$ have already been constructed for all $i \in [\ell]$ meeting the requirements from above. Let \bar{J} be the segment of Ω' obtained by deleting the vertices of the segment J^ℓ .

Suppose the society from (G', Ω') cut by $Z^\ell \cup S^\ell$ and patched at \bar{J} has no cross. Let A be the collection of all vertices of G that are contained in a component of $G - (Z^\ell \cup S^\ell)$ that contains a

vertex of $V(\Omega)$, let Ω^* be the restriction of Ω to A , and let $B := V(G) \setminus (A \cup Z^\ell \cup S^\ell)$. As $\ell \leq t-1$ and $|Z_i|, |S_i| \leq 3\theta$ for all $i \in [\ell]$ we have $|Z^\ell \cup S^\ell| \leq 6\theta(t-1)$. Hence, $(A \cup Z^\ell \cup S^\ell, Z^\ell \cup S^\ell \cup B)$ is a separation of order at most $6\theta(t-1)$ and $(G[A], \Omega^*)$ has a vortex-free rendition in the disk. Thus, we are done with this case.

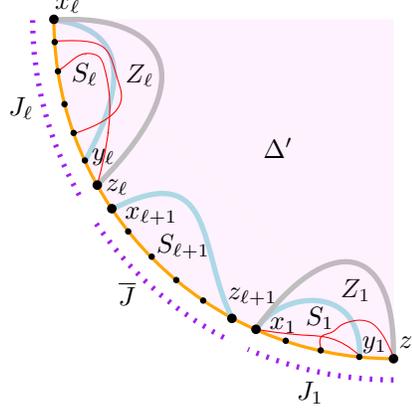


Figure 11: A visualization of **Rule 2** in the proof of [Lemma 34](#). The fat grey and light blue lines represent the separators Z_i and S_i , and the magenta dotted lines are the corresponding segments J_i .

From now on we may assume that the society from (G', Ω') cut by $Z^\ell \cup S^\ell$ and patched at \bar{J} has a cross.

Let $x_{\ell+1}$ be the immediate successor of z_ℓ with respect to λ and let $z_{\ell+1}$ be chosen according to one of the following rules:

Rule 1 Let $z_{\ell+1}$ be the smallest vertex of $V(\lambda)$ with $\lambda(x_{\ell+1}) \leq \lambda(z_{\ell+1})$ such that the society from (G', Ω') cut by $Y_{x_{\ell+1}, z_{\ell+1}}^\ell \cup Z^\ell$ and patched at $J_{x_{\ell+1}, z_{\ell+1}}$ has a cross.

Rule 2 If a choice according to **Rule 1** is not possible select $z_{\ell+1}$ such that $J_{x_{\ell+1}, z_{\ell+1}} = \bar{J}$ and the society from (G', Ω') cut by $Z^\ell \cup S^\ell$ and patched at \bar{J} has a cross.

By our assumption, we may consecutively apply **Rule 1** and, when this is not any more possible, either the procedure stops, which means that the society from (G', Ω') cut by $Z^\ell \cup S^\ell$ and patched at \bar{J} has a vortex free rendition in the disk, or it stops after applying **Rule 2** once. In the later case, the society from (G', Ω') cut by $(Z^\ell \cup S^\ell) \cup Y_{x_{\ell+1}, z_{\ell+1}}^\ell$ and patched at $V(\Omega')$ has a vortex free rendition in the disk (see [Figure 11](#)).

Either way let $y_{\ell+1}$ be the immediate predecessor of $z_{\ell+1}$ under λ .

In case z_ℓ was chosen by **Rule 2** we set $Z_{\ell+1} := \emptyset$ and $S_{\ell+1} := Y_{x_{\ell+1}, z_{\ell+1}}^\ell$. Note that in this case $J_{\ell+1} = V(\lambda)$. Hence, we satisfy all five rules of our iterative process simply by assumption and may continue with step $\ell+2$ or we are in the case $\ell+1 = t$ which will be treated later.

So we may further assume that $z_{\ell+1}$ was chosen according to **Rule 1**. Now we set $J_{\ell+1} := J_{x_{\ell+1}, z_{\ell+1}}$, $Z_{\ell+1} := Y_{x_{\ell+1}, z_{\ell+1}}^\ell$, and $S_{\ell+1} := Y_{x_{\ell+1}, y_{\ell+1}}^\ell$. We then have, by the discussion above, $|Z_{\ell+1}|, |S_{\ell+1}| \leq 3\theta$ and $J_{\ell+1}$ is disjoint from all other segments J_i , $i \in [\ell]$. Moreover, $Z_{\ell+1}$ separates $J_{\ell+1}$ from $V(\Omega') \setminus J_{\ell+1}$ within G' by choice of $Z_{\ell+1}$. By **Rule 1** the society from (G', Ω') cut by

$Z^{\ell+1}$ and patched at $J_{x_{\ell+1}, z_{\ell+1}}$ has a cross and finally, by the minimality of $z_{\ell+1}$ we have that the society from (G', Ω') cut by $Z^{\ell+1} \cup S^{\ell+1}$ and patched at $J^{\ell+1}$ has a vortex-free rendition in the disk. Thus, all five requirements for our iteration are met and we may continue or have entered the case where $\ell + 1 = t$.

To complete the proof let us discuss what happens in the case where we have successfully completed iteration round t without finding the separation. We need to show that this implies the existence of t consecutive crosses over (G, Ω) . Since we have successfully completed round t we have found segments J_1, \dots, J_t such that the first vertex of J_{i+1} comes before the last vertex of J_i for all $i \in [t-1]$ and $J_1 \cap J_t = \emptyset$. For each $i \in [t]$ let Y_i be the patch of (G', Ω') cut at J_i by Z^t and let (G_i, Ω_i) be the restriction of G and Ω to the graph $G - (V(G') \setminus Y_i)$. Since the society from (G', Ω') cut by Z^t at J_i has a cross, the society (G_i, Ω_i) cannot have a vortex-free rendition in the disk and thus, by [Proposition 15](#), it must have a cross. Therefore, we may apply [Lemma 33](#) to obtain our t consecutive crosses. \square

So far we have seen that, in the second case of θ -tightness, we can find disjoint crosses over our society that meet all cycles in the nest or we can completely eliminate the vortex. To construct a shallow vortex minor however, we need the paths of these crosses to be particularly well behaved with a large portion of the nest.

Our goal is to obtain a still large number of concentric cycles which behaves in an ‘‘orthogonal’’ fashion with respect to the t -consecutive crosses [Lemma 34](#) provides us with in case we cannot remove a vortex with a small set of vertices. To realise this plan we will take several steps. Let (ρ, G, Ω) be a rendition of a nest \mathcal{C} around a vortex c and let (P_1, P_2) be a cross over (G, Ω) . Clearly each of the two paths P_i needs to contain an edge that is drawn in the interior of c . However, we can not necessarily guarantee a bound on the number of components in $(P_1 \cup P_2) - \tilde{c}$. That is, the paths P_i may leave and re-enter the vortex c many times before eventually making progress towards $V(\Omega)$. This poses an issue as any \tilde{c} subpath of one of the P_i whose interior is disjoint from the interior of c cannot necessarily be used for a re-routing of the cycles in \mathcal{C} that eventually achieves ‘‘orthogonality’’. To cope with this issue we proceed to prove that there exists some constant c such that any cross (P_1, P_2) can be rerouted to a cross (P'_1, P'_2) with the same endpoints as (P_1, P_2) but every \tilde{c} -subpath of a path P'_i is disjoint from all cycles $C_j \in \mathcal{C}$ with $j > c$. For instance this property does not hold for the crossing paths in the left-side of [Figure 12](#) while it holds for the right-side one. To achieve this we need some further definitions.

In the following we will identify a linkage \mathcal{P} with the graph $\bigcup_{P \in \mathcal{P}} P$. Let G be a graph and let \mathcal{P} be a linkage in G . The set $\{\{s, t\} \mid \text{some path in } \mathcal{P} \text{ has endpoints } s \text{ and } t\}$ is called the *pattern* of \mathcal{P} . Two linkages $\mathcal{Q}_1, \mathcal{Q}_2$ are said to be *equivalent* if they have the same pattern. A linkage \mathcal{Q} is said to be *vital* if $V(\mathcal{Q}) = V(G)$ and there exists no other linkage in G equivalent to \mathcal{Q} . The famous *unique linkage theorem* of Robertson and Seymour [[29](#), [39](#)] states that, if \mathcal{P} is vital in a graph G , then the treewidth of G is bounded in some function of $|\mathcal{P}|$. Instead of the unique linkage theorem however, we make use of a slightly more convenient version from [[20](#)].

Definition 35 (LB-Pair). Given a graph G , a LB-pair of G is a pair (\mathcal{L}, B) where B is a subgraph of G with maximum degree two and \mathcal{L} is a linkage in G . We define $\text{dis}(\mathcal{L}, B) := |E(\mathcal{L}) \setminus E(B)|$ to be the number of edges from \mathcal{L} on which \mathcal{L} and B disagree.

Proposition 36 ([[20](#)]). *There exists a function $\text{link}: \mathbb{N} \rightarrow \mathbb{N}$ such that for all graphs G and all LB-pairs (\mathcal{L}, B) of G , if $\text{tw}(\mathcal{L} \cup B) > \text{link}(|\mathcal{L}|)$ then G contains a linkage \mathcal{R} such that*

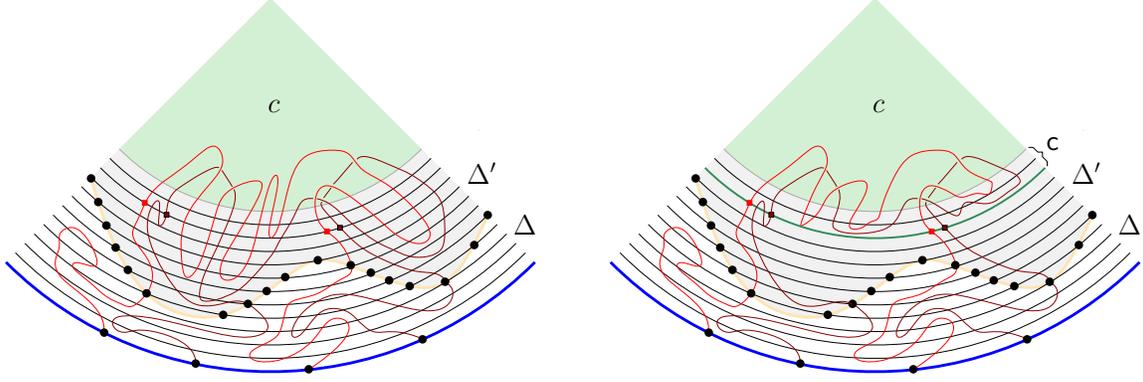


Figure 12: Two ways crossing paths P_1, P_2 (in red and dark red) may be rooted inside Δ' . Notice that in the picture on the right every \tilde{c} -subpath of a path P_i is disjoint from all cycles $C_j \in \mathcal{C}$ with $j > c$ while this is not the case in the picture on the left.

- i) $\text{dis}(\mathcal{R}, B) < \text{dis}(\mathcal{L}, B)$,
- ii) \mathcal{L} and \mathcal{R} are equivalent, and
- iii) $\mathcal{R} \subseteq \mathcal{L} \cup B$.

Given an LB-pair (\mathcal{L}, B) of a graph G we say that (\mathcal{L}, B) is of *maximum consensus* if for all families of pairwise disjoint paths \mathcal{R} which are equivalent to \mathcal{L} we have $\text{dis}(\mathcal{L}, B) \leq \text{dis}(\mathcal{R}, B)$. It follows from [Proposition 36](#) that $\text{tw}(\mathcal{L} \cup B) \leq \text{link}(|\mathcal{L}|)$ for all LB-pairs (\mathcal{L}, B) of maximum consensus.

Lemma 37. *Let θ be a positive integer. Let $(\rho = (\gamma, \mathcal{D}, c), G, \Omega)$ be a θ -tight rendition of a nest \mathcal{C} , with $|\mathcal{C}| \geq 2 \cdot \text{link}(2) + 4$, around a vortex c of depth at most θ in a disk Δ in ρ . Moreover, let $(\theta, \rho, G, \Omega, \mathcal{C}, \Delta, \Delta', \mathcal{P})$ be a θ -suspension of c in (ρ, G, Ω) . Let $Z \subseteq V(G')$ be a set of vertices and let $S \subseteq V(\Omega')$ be a segment of Ω' such that S has a cross (L, R) at Z . Then there exists a pair of paths (L', R') such that*

- i) (L', R') is a cross of S at the set Z ,
- ii) L' and L have the same endpoints,
- iii) R' and R have the same endpoints, and
- iv) any subpath of L' or R' with both endpoints in \tilde{c} intersects at most $c := 2\text{link}(2) + 2$ many cycles of \mathcal{C} .

Towards a proof for [Lemma 37](#) we need a way to certify large treewidth. To do this we make use of the well known concept of brambles.

Let G be a graph and H_1, H_2 be two connected subgraphs of G . We say that H_1 and H_2 *touch* if $V(H_1) \cap V(H_2) \neq \emptyset$, or there is an edge uv with $u \in V(H_1)$ and $v \in V(H_2)$. A set $S \subseteq V(G)$ is a *hitting set* or *cover* for a family \mathcal{S} of subgraphs of G , if $V(H) \cap S \neq \emptyset$ for all $H \in \mathcal{S}$.

Definition 38 (Bramble). Let G be a graph. A *bramble* $\mathcal{B} = \{B_1, B_2, \dots, B_\ell\}$ of G is a family of connected and pairwise touching subgraphs B_i of G . The *order* of \mathcal{B} is the size of a minimum hitting set for \mathcal{B} .

Proposition 39 ([48]). Let G be a graph, and $k \in \mathbb{N}$ a positive integer. Then G contains a bramble of order k if and only if $\text{tw}(G) \geq k - 1$.

Proof of Lemma 37. Let $(\hat{G}, \hat{\Omega})$ be the society from (G', Ω') cut by Z and patched at S . Let U be some subpath of $L \cup R$ with both endpoints in \tilde{c} that is internally disjoint from \tilde{c} and let U' be its trace. We call such a subpath of $L \cup R$ an *arc* of (L, R) . Note that U' together with the boundary of c divides Δ' into three different areas, the disk bounded by c and two disks which are subsets of $\Delta' - (c - \text{Boundary}(c))$, exactly one of these disks contains all vertices of L and R that lie on the boundary of Δ' . Let Δ_U be the remaining disk. Observe that every maximal subpath of $L \cup R$ which is drawn in Δ_U must have both of its endpoints in \tilde{c} . We call Δ_U a *mountain* if there does not exist another arc W of (L, R) such that $\Delta_U \subset \Delta_W$, moreover, if Δ_U is a mountain, we call U its *outline*.

Now every subpath of $L \cup R$ that intersects some cycle from \mathcal{C} and has both endpoints in \tilde{c} must be completely contained in \tilde{c} together with the union of all arcs of (L, R) . Hence, every arc of (L, R) must either be an outline of some mountain or drawn in the interior of some mountain. Moreover, if U and W are distinct outlines, then their corresponding mountains are disjoint with the possible exception for the endpoints of U and W . For every arc W of (L, R) let z_W be the maximum number of cycles from \mathcal{C} met by W . Let $n = |\mathcal{C}|$. Given a mountain Δ_U , we may associate a vector $\mathbf{v}(U) \in \mathbb{N}^n$ such that for all $i \in [n]$,

$$\mathbf{v}(U)_{n-i+1} = |\{W \mid W \text{ is an arc drawn in } \Delta_U \text{ with } z_W = i\}|.$$

We call $\mathbf{v}(U)$ the *characteristic vector* of the mountain Δ_U . Note that $\mathbf{v}(U)_i$ is the number of arcs drawn in Δ_U that meet the cycle C_{n-i+1} . We now proceed to show that we can find a cross (L', R') with properties *i.*, *ii.*, and *iii.*, such that no mountain of (L', R') contains a vertex of $c + 1$ cycles.

Let (L', R') be chosen to first minimize the number of mountains, second minimize the number of edges on which the arcs within a mountain disagree with the cycles from \mathcal{C} , and thirdly to lexicographically minimize the characteristic vectors of all mountains. Suppose there exists an outline U such that the corresponding mountain Δ_U meets at least $h \geq c + 1$ cycles. For each $i \in [h]$ let B_i be the subpath of C_i which is completely drawn into the disk Δ_U , and let B be the union of all of these subpaths. Let \mathcal{U} be the union of all arcs of (L, R) which are drawn in Δ_U . Note that the trace of any arc $W \in \mathcal{U}$, distinct from U , separates Δ_U into two disks, one of them containing all vertices of U , let us call this disk the *upper part of W* , the other disk is called the *lower part of W* . Moreover, every other arc is completely contained in one of these two disks. Let us assume $\mathcal{U} = \{U_1, \dots, U_\ell\}$ is numbered such that $U_1 = U$ and for all $i \in [2, \ell]$ the upper part of U_i contains the arc U_{i-1} . Please note that this ordering is not uniquely determined.

In the following we will define a bramble as follows. Let $K_1 := B_h \cup U_1$ and let x_1, y_1 be the two endpoints of U_1 . Now suppose U_2 does not meet B_{h-1} . In this case let P_1 be the shortest x_1 - B_{h-1} -subpath of U_1 and let P_2 be the shortest B_{h-1} - y_1 -subpath of U_1 . Finally, let P_3 be the shortest subpath of B_{h-1} that joins the two endpoints of P_1 and P_2 . We may now exchange the path U_1 in $L' \cup R'$ with the path $P_1 \cup P_3 \cup P_2$. This yields an immediate contradiction to the choice of (L', R') as this exchange yields a new pair of paths whose mountains are the same as the mountains of

(L', R') with the sole exception of U , here the new pair has a mountain whose characteristic vector is lexicographically smaller than the one of U . Moreover, we claim that the numbering of the U_i can be chosen such that for any $i \in [2, h]$ the path U_i meets B_{h-i+1} . To see this, suppose there is some i such that for all $j < i$ the choice was possible, but for i itself it is not. This means that for all $j \in [2, i-1]$ the arc U_j meets B_{h-i+2} and every arc in the lower part of U_j avoids B_{h-i+1} . Hence, we may use B_{h-i+1} to obtain a cross with the same number of mountains as (L', R') , but which contradicts the lexicographic minimality of the characteristic vectors as before. Hence, we may define $K_i := B_{h-i+1} \cup U_i$ for all $i \in h$.

Now observe that for any $j \in [\ell]$ and $i \in [h]$, if U_j meets B_i , then U_j intersects all $B_{i'}$ with $i' \leq i$. Hence, K_i and K_j intersect for all choices of $i, j \in [h]$. However, since the B_i are pairwise disjoint and the U_j are pairwise disjoint, no vertex can be contained in more than two of the K_i at once. Hence, any cover of $\mathcal{K} := \{K_i \mid i \in [h]\}$ must have size at least $\lceil \frac{h}{2} \rceil \geq \text{link}(2) + 2$, making it a bramble of order at least $\text{link}(2) + 1$. Hence, by [Proposition 39](#) the treewidth of $\{U_i \mid i \in [h]\} \cup B$ is at least $\text{link}(2) + 1$. By [Proposition 36](#) this is a contradiction to the second minimization parameter we used for the choice of (L', R') and thus our proof is complete. \square

Towards extracting the structure of a shallow vortex grid, we need that all crossing paths implied by [Lemma 34](#) traverse the cycles of the nest ‘‘orthogonally’’. This was formalized in [\[23\]](#) with the following definition.

Definition 40 (Orthogonal Linkage and Crosses). Let (ρ, G, Ω) be a rendition of a nest \mathcal{C} around a vortex c in a disk Δ in ρ and let \mathcal{P} be a $V(\Omega)$ - \tilde{c} linkage. We say that \mathcal{P} is *orthogonal* to \mathcal{C} if for every $C \in \mathcal{C}$ and every $P \in \mathcal{P}$ the graph $P \cap C$ is a path.

Let \mathcal{U} be a collection of crosses over (G, Ω) and let \mathcal{Q} be the collection of all $V(\Omega)$ - \tilde{c} -subpaths of the paths of the crosses in \mathcal{C} . We say that \mathcal{U} is *orthogonal* to \mathcal{C} if \mathcal{Q} is orthogonal to \mathcal{C} .

Lemma 41. *Let $t \leq \theta$ be positive integers. Let $(\rho = (\gamma, \mathcal{D}, c), G, \Omega)$ be a θ -tight rendition of a nest \mathcal{C} , with $|\mathcal{C}| \geq 8t^2 + 2\text{link}(2) + 4$, around a vortex c of depth at most θ in a disk Δ in ρ . Moreover, assume that $(\theta, \rho, G, \Omega, \mathcal{C}, \Delta, \Delta', \mathcal{P})$ is θ -suspension of c in (ρ, G, Ω) and that there is a consecutive family $\mathcal{Q} := \{\mathcal{Q}_1, \dots, \mathcal{Q}_t\}$ of t crosses over (G, Ω) together with t pairwise disjoint segments S_1, \dots, S_t of Ω'' and a set $Z'' \subseteq V(G'')$ such that for each $i \in [t]$ the cross \mathcal{Q}_i intersects G'' exactly in a cross of Ω'' at Z'' .*

Then there exists a family \mathcal{C}' of t pairwise vertex disjoint cycles whose traces all separate c from $V(\Omega)$ and such that \mathcal{Q} is orthogonal to \mathcal{C}' .

Proof. Let $\mathcal{C} = \{C_1, \dots, C_\ell\}$. By [Lemma 37](#) we may assume that no arc of \mathcal{Q}_i meets the cycle $C_{2\text{link}(2)+3}$ for any $i \in [t]$. For every $i \in [t]$ let $\mathcal{Q}_i = (L_i, R_i)$, and let $s_i^L, s_i^R, t_i^L, t_i^R$ be the four endpoints of the paths L_i and R_i respectively in the order they appear on Ω . Moreover, for every $i \in [t]$ and every $W \in \{L, R\}$ let x_i^W be the first vertex of \tilde{c} encountered when traversing W_i starting in s_i^W , and let y_i^W be the first vertex of \tilde{c} encountered when traversing W_i starting in t_i^W . Finally, let $\mathcal{P}_i := \{s_i^L L_i x_i^L, y_i^L L_i t_i^L, s_i^R R_i x_i^R, y_i^R R_i t_i^R\}$, $\mathcal{P} := \bigcup_{i=1}^t \mathcal{P}_i$, $T = \bigcup_{i \in [t]} \{x_i^L, y_i^L, x_i^R, y_i^R\}$, $\tilde{\Omega}$ be the cyclic ordering of the vertices in T obtained by traversing the boundary of c in clockwise direction, let $Q := \{C_{2\text{link}(2)+5}, \dots, C_{2\text{link}(2)+4+8t^2}\} \cup \mathcal{P}$, and $(Q, \tilde{\Omega})$ be the society defined by Q and $\tilde{\Omega}$. Note that it follows from our assumptions that $(Q, \tilde{\Omega})$ has a vortex-free rendition in the disk.

For better readability let us rename $O_i := C_{2\text{link}(2)+4+i}$ for every $i \in [4t^2]$. Note that for each $i \in [t]$ the vertices $x_i^L, x_i^R, y_i^L, y_i^R$ appear on $\tilde{\Omega}$ in the order listed. Moreover, the vertices of $V(\Omega)$ which are present in Q appear in the rendition of $(Q, \tilde{\Omega})$ in the same order as they appear on Ω . Now, for every $i \in [t-1]$ and $j \in [t]$

- let $B_{i,j,L,1}$ be a shortest $s_i^L L_i x_i^L - s_i^R R_i x_i^R$ subpath of the cycle $O_{t(j-1)+4(i-1)+1}$,
- let $B_{i,j,R,1}$ be a shortest $s_i^R R_i x_i^R - t_i^L L_i y_i^L$ subpath of the cycle $O_{t(j-1)+4(i-1)+2}$,
- let $B_{i,j,L,2}$ be a shortest $t_i^L L_i y_i^L - t_i^R R_i y_i^R$ subpath of the cycle $O_{t(j-1)+4(i-1)+3}$, and
- let $B_{i,j,R,2}$ be a shortest $t_i^R R_i y_i^R - s_{i+1}^L L_{i+1} y_{i+1}^L$ subpath of the cycle $O_{t(j-1)+4(i-1)+4}$.

Moreover,

- let $B_{t,j,L,1}$ be a shortest $s_t^L L_t x_t^L - s_t^R R_t x_t^R$ subpath of the cycle $O_{t(j-1)+4(t-1)+1}$,
- let $B_{t,j,R,1}$ be a shortest $s_t^R R_t x_t^R - t_t^L L_t y_t^L$ subpath of the cycle $O_{t(j-1)+4(t-1)+2}$,
- let $B_{t,j,L,2}$ be a shortest $t_t^L L_t y_t^L - t_t^R R_t y_t^R$ subpath of the cycle $O_{t(j-1)+4(t-1)+3}$, and
- let $B_{t,j,R,2}$ be a shortest $t_t^R R_t y_t^R - s_1^L L_1 y_1^L$ subpath of the cycle $O_{t(j-1)+4t}$.

Note that the paths of the form $B_{i,j,W,h}$ are pairwise internally disjoint and for each $j \in [t]$, every path of $P \in \mathcal{P}$ meets exactly two paths of the form $B_{i,j,W,h}$. Indeed, both of these paths are met exactly in their endpoints. These endpoints can be joined by a unique subpath of P each, resulting in total in $2t$ pairwise disjoint cycles D_1, \dots, D_t . Note that it immediately follows that the trace of each D_i separates c from $V(\Omega)$ in ρ .

By the choice of the O_i and our assumption from before it follows that each $B_{i,j,W,h}$ is internally disjoint from all paths from the crosses of \mathcal{Q} and thus \mathcal{Q} is orthogonal to $\{D_1, \dots, D_{2t}\}$. \square

We are finally ready to prove the main result of this section.

Proof of Lemma 31. Let $c := 2\text{link}(2) + 4$. We start by applying Lemma 34. This either yields the desired separation or we find a θ -suspension $(\theta, \rho, G, \Omega, \mathcal{C}, \Delta, \Delta', \mathcal{P})$ of c in (ρ, G, Ω) . Moreover, if $(G'' := G[V(G) \cap \Delta'], \Omega'')$ is the society defined by Δ' in the definition of θ -suspensions, then there exists a consecutive family $\mathcal{Q} = \{\mathcal{Q}_1, \dots, \mathcal{Q}_{2t}\}$ of $2t$ crosses over (G, Ω) together with $2t$ pairwise disjoint segments S_1, \dots, S_{2t} of Ω'' and a set $Z'' \subseteq V(G'')$ such that for each $i \in [2t]$ the cross \mathcal{Q}_i intersects G'' exactly in a cross of Ω'' at Z'' .

Now an application of Lemma 41 yields the existence of a family \mathcal{D} of $\frac{1}{8} \cdot 8 \cdot \sqrt{4t^2} = 2t$ pairwise disjoint cycles, all of which have traces that separate $V(\Omega)$ from c such that \mathcal{D} is orthogonal to \mathcal{Q} . Hence, the graph consisting of the cycles of \mathcal{D} together with the paths of the crosses from \mathcal{Q} is a minor model of a shallow vortex grid of order $2t$ as desired. \square

3.3 Excluding a shallow vortex minor

In the previous section we have seen that, under the absence of a shallow vortex grid, any vortex can be completely separated from the rest of a Σ -decomposition with a small set of vertices. Towards a complete structural description of graphs excluding a fixed shallow vortex minor we need to push this idea one step further. One way to do this would be to directly evoke the global structure theorem for H -minor-free graphs of Robertson and Seymour and afterwards remove all vortices while slightly increasing the apex set. However, this approach has the problem that the process of “killing” the vortices has the potential to change the decomposition within the collection of subtrees that attached to the linear decomposition of a vortex. To avoid the backtracking such an approach would require, it is more convenient to instead give a slightly altered version of the proof of the global structure theorem from [23] which builds the desired decomposition inductively.

What remains of this section is now dedicated to the proof of [Proposition 21](#). As described above, we heavily lean on the proof of the global structure theorem from [23] which itself is based on a similar proof in [13] and inspired by a preliminary version of the original theorem that can be found in [42]. The central tool for these proofs is an object dual to small treewidth called a *tangle*.

Definition 42 (Tangle). Let G be a graph and k be a positive integer. We denote by \mathcal{S}_k the collection of all tuples (A, B) where $A, B \subseteq V(G)$ and (A, B) is a separation of order $< k$ in G . An *orientation* of \mathcal{S}_k is a set \mathcal{O} such that for all $(A, B) \in \mathcal{S}_k$ exactly one of (A, B) and (B, A) belongs to \mathcal{O} .

A *tangle* of order k in G is an orientation \mathcal{T} of \mathcal{S}_k such that for all $(A_1, B_1), (A_2, B_2), (A_3, B_3) \in \mathcal{T}$ we have $A_1 \cup A_2 \cup A_3 \neq V(G)$.

Let G and H be graphs and let $p := |V(H)|$. Let \mathcal{T} be a tangle in G . A minor model of H in G is *controlled* by \mathcal{T} if for every separation $(A, B) \in \mathcal{T}$ with $|A \cap B| < p$ there does not exist a branch set of the model contained in $A \setminus B$. For a set $Z \subseteq V(G)$ we denote by $\mathcal{T} - Z$ the tangle in $G - Z$ formed by the separations (A, B) in $G - Z$ such that $(A \cup Z, B \cup Z) \in \mathcal{T}$. Given a surface Σ and a Σ -decomposition Δ of G , we say that the decomposition is *\mathcal{T} -central* if for every separation $(A, B) \in \mathcal{T}$, there does not exist a cell $c \in \mathcal{C}(\Delta)$ such that $B \subseteq V(\sigma_\Delta(c))$.

We will need the following theorem.

Proposition 43 ([23]). *There exists an absolute constant c which satisfies the following. Let $p \geq 1$ be a positive integer. Let \mathcal{T} be a tangle of order Θ in a graph G with*

$$\Theta \geq p^{18 \cdot 10^7 p^{26} + c}.$$

Then \mathcal{T} either controls a K_p -minor or there exists $A \subseteq V(G)$, $|A| \leq 5p^2 \cdot p^{10^7 p^{26}}$, a surface Σ of Euler-genus at most $p(p+1)$, and a Σ -decomposition Δ of $G - A$ of breadth at most $2p^2$ and depth at most $p^{10^7 p^{26}}$ which is $(\mathcal{T} - A)$ -central.

Instead of [Theorem 7](#), we prove a slightly stronger statement which will imply our main theorem.

[Theorem 44](#) is stronger than [Theorem 7](#) in two ways. First, it gives much more intricate information on the different parameters of the tree decomposition produced including the information that any adhesion set can avoid the apex set of any of its neighbouring bags in at most three vertices. Second, it essentially shows that it does not matter where we begin to construct the

decomposition. This is facilitated by the free choice of the set Z which acts as a way to “anchor” the decomposition. In each step of the inductive proof, we try to find a balanced separator for Z of small order. If we can find such a separator we may split the graph and apply induction on each of the pieces, otherwise Z does not have a small balanced separator. In the latter case we get a tangle of large order from Z which allows us to use [Proposition 43](#) to either find a large clique minor or a Σ -decomposition central to the tangle. If we find the clique we are done as it contains a shallow vortex grid, otherwise we apply [Lemma 31](#) to each of the vortices in order to obtain a vortex-free Σ -decomposition Δ . By definition of Σ -decompositions, after deleting the resulting apex set, any non-node vertex can be separated from the nodes of Δ by deleting at most three vertices. This allows us to continue the induction within each of the cells with non-empty interiors. Such an approach to constructing decompositions for global decompositions has already been employed for the proof of [Theorem 11.3](#) in [\[41\]](#) and was also used for the proof of the global structure theorem in [\[23\]](#).

Theorem 44. *Let $t \geq 1$ be a positive integer and let H be a minor of the shallow vortex grid of order t . Let G be a graph that does not contain H as a minor. Let $\alpha := t^{18 \cdot 10^8 t^{78} + c} + 96t^{10^8 t^{78} + 5} + 1$ and $\gamma = 4t^4 + 2t^2$. Let $Z \subseteq V(G)$ be such that $|Z| \leq 3\alpha$. Then G has a tree decomposition (T, β) with a root $r \in V(T)$ and adhesion at most α such that for every $d \in V(T)$, the torso G_d of G at d has a set $A_d \subseteq V(G_d)$ of size at most 4α for which the graph $G_d - A_d$ has Euler-genus at most γ . Moreover, we have $Z \subseteq A_r$, for every $(d_1, d_2) \in E(T)$ we have $|(\beta(d_1) \setminus A_{d_1}) \cap (\beta(d_2) \setminus A_{d_2})| \leq 3$, and if $|(\beta(d_1) \setminus A_{d_1}) \cap (\beta(d_2) \setminus A_{d_2})| = 3$ and $\beta(d_1)$ is larger than 4α , then $(\beta(d_1) \setminus A_{d_1}) \cap (\beta(d_2) \setminus A_{d_2})$ induces a triangle in $G_{d_1} - A_{d_1}$ which bounds a face.*

Proof. Let us assume the assertion is false. We fix G , t and Z to form a counter example minimizing $|G| + |G \setminus Z|$.

Observe that, in case $t = 1$ we have that H is a single crossing minor and thus the claim follows immediately from the main result of [\[43\]](#). Hence, we may assume $t \geq 2$. Furthermore, we may assume $|G| > 4\alpha$ as otherwise the trivial tree decomposition on a tree with a single vertex r and $\beta(r) = V(G)$ and $A_r = V(G)$ would satisfy our claim. Moreover, from our minimality assumptions it follows that $|Z| = 3\alpha$ since otherwise we could add another vertex to Z and find a smaller counter example.

Claim 1. For all separations (X_1, X_2) of order less than α we have $|X_i \cap Z| \leq |X_1 \cap X_2|$ for exactly one $i \in [2]$.

Proof of Claim 1. Let (X_1, X_2) be a separation of order less than α . If $|X_i \cap Z| \leq |X_1 \cap X_2| < \alpha$ we would have $|Z| < 2\alpha$ contradicting our assumption that $|Z| = 3\alpha$.

Hence, it suffices to show that $|X_i \cap Z| > |X_1 \cap X_2|$ cannot hold for both $i \in [2]$. Towards a contradiction let us assume that the inequality holds for both $i \in [2]$. Note that we must have that (X_1, X_2) is non-trivial. For each $i \in [2]$ let $Z_i := (Z \cap X_i) \cup (X_1 \cap X_2)$. By minimality the theorem holds for the two graphs $G_i := G[X_i]$ with the set Z_i respectively. So for each $i \in [2]$ there exists a tree decomposition (T_i, β_i) with root r_i such that for every $d \in V(T_i)$ the torso $G_{i,d}$ of G_i at d has an apex set $A_{i,d}$ of size at most 4α such that $G_{i,d} - A_{i,d}$ has Euler-genus at most γ . Let T be the tree formed by introducing a new vertex r to $T_1 \cup T_2$ and joining it with edges to the vertices r_1 and r_2 . Let $\beta(r) := Z \cup (X_1 \cap X_2)$ and $\beta(d) := \beta_i(d)$ if $d \in V(T_i)$ for all $d \in V(T_1) \cup V(T_2)$. Note that $Z_i \subseteq \beta(r_i)$ by assumption. As $|X_1 \cap X_2| < \alpha$ and thus $|\beta(r)| < 4\alpha$ it follows that (T, β) is a tree decomposition for G of adhesion at most α such that the torso at every vertex of T has

Euler-genus at most γ after deleting at most 4α vertices. This means that G could not have been a counter example to our assertion in the first place and we have obtained a contradiction. ■

We may now consider the collection \mathcal{T} of all separations (X, Y) of order less than α in G such that $|Z \cap Y| > \alpha$. Recall that \mathcal{S}_α is the family of all separations of order less than α in G . It follows from Claim 1 that \mathcal{T} is in fact an orientation of \mathcal{S}_α . Moreover, given three separations $(X_1, Y_1), (X_2, Y_2), (X_3, Y_3) \in \mathcal{T}$, as each Y_i contains at most $\alpha - 1$ vertices of Z and $|Z| = 3\alpha$ we have $V(G) \neq Y_1 \cup Y_2 \cup Y_3$ and thus \mathcal{T} is a tangle of order α .

As H is a minor of the shallow vortex grid of order t , which itself has $2t^2$ vertices and thus is contained in K_{2t^2} , G cannot have a K_{2t^2} -minor. Hence, by Proposition 43 there exists a surface Σ of Euler-genus at most $2t^2(2t^2 + 1) = \gamma$, a subset $A' \subseteq V(G)$ with $|A'| \leq 5(2t^2)^2 \cdot (2t^2)^{10(2t^2)^7(2t^2)^{26}}$, and a Σ -decomposition ρ of $G - A'$ of breadth at most $2(2t^2)^2 = 8t^4$ and depth at most $(2t^2)^{10^7(2t^2)^{26}} \leq t^{10^8 t^{78}}$ which is $\mathcal{T} - A'$ -central. Moreover, by applying Proposition 24 we may further assume that for each vortex cell $c \in C(\rho)$ there exists a nest \mathcal{C}_c in ρ around the unique disk $\Delta \in \mathcal{D}$ corresponding to c of order $10^{21}(2t^2)^{100}$. Also, for each vortex cell $c \in C(\rho)$, if $\Delta_c \subseteq \Sigma$ is the disk bounded by the trace of $C_{10^{21}(2t^2)^{100}} \in \mathcal{C}_c$, then for each pair of distinct vortex cells $c, c' \in C(\rho)$ we have that $\Delta_c \cap \Delta_{c'} = \emptyset$. If we use Theorem 27 we get, in addition to the above properties, that for each vortex cell $c \in C(\rho)$ the rendition $(\rho_{\mathcal{C}_c}, H_{\mathcal{C}_c}, \Omega_{\mathcal{C}_c})$ of the nest \mathcal{C}_c is $(2t^2)^{10^7(2t^2)^{26}} \leq t^{10^8 t^{78}}$ -tight.

We now break G into subgraphs based on ρ . By the minimality of G , each subgraph has a tree decomposition with the desired properties that can be attached to the part of G that is properly drawn on Σ . The only difference between this proof and the one from [23] is the way we treat vortices. In the end, all decompositions we obtain for the subgraphs will be combined into one tree decomposition for the entirety of G , resulting in a contradiction to G being a counter example.

Let us start by treating the vortices. Let $c \in C(\rho)$ be an arbitrary vortex cell. Then there exists a nest $\mathcal{C}_c = \{C_1, \dots, C_{10^{21}(2t^2)^{100}}\}$ of order $10^{21}(2t^2)^{100}$ and a disk $\Delta_{\mathcal{C}_c}$ bound by the trace of $C_{10^{21}(2t^2)^{100}}$ such that the rendition $(\rho_{\mathcal{C}_c}, H_{\mathcal{C}_c}, \Omega_{\mathcal{C}_c})$ is $t^{10^8 t^{78}}$ -tight. By Lemma 31 we either find a shallow vortex grid of order t as a minor, and thus we obtain an H minor in G which is impossible, or there exists a separation (A_c, B_c) of order at most $12 \cdot t^{10^8 t^{78}}(t - 1)$ with $V(\Omega_{\mathcal{C}_c}) \subseteq A$ such that $(H_{\mathcal{C}_c}[A], \Omega'_{\mathcal{C}_c})$ has a vortex-free rendition in the disk. Let $S_c := A_c \cap B_c$. Since G is H -minor-free we can find the separation $(A_{c'}, B_{c'})$ for every vortex in ρ . Moreover, $B_{c'} \cap B_{c''} = \emptyset$ for all distinct vortices $c', c'' \in C(\rho)$. Hence, by removing the union of all $B_{c'}$ from G we obtain a vortex-free Σ -decomposition for the resulting graph G' . Let Z be the union of all $S_{c'}$ and B be the union of all $B_{c'}$ over all vortex cells of ρ . Since there are at most $8t^4$ vortices we have that

$$|Z| \leq 8t^4 \cdot 12 \cdot t^{10^8 t^{78}}(t - 1) \leq 96 \cdot t^{10^8 t^{78} + 5}.$$

Now let $G' := G - (B \setminus S)$, moreover, let $A := A' \cup S$. Notice that $|A| \leq \alpha - 1$. Then ρ contains a vortex-free Σ -decomposition ρ' of $G' - A$.

It follows that for every vortex cell c we have that $|B_c \cap Z| \leq \alpha - 1$. Let H_c be the subgraph of G induced by $B_c \cup A \cup (B_c \cap Z)$. By the minimality of G , there exists a tree decomposition (T_c, β_c) of adhesion at most α such that the torso of every bag has Euler-genus at most γ after the deletion of an apex set of size at most 4α . Moreover, T_c has a root r_c such that $A \cup (B_c \cap Z) \in \beta_c(r_c)$ and it is a subset of the apex set of the torso of H_c at r_c .

For every $o \in C(\rho')$ that is not a vertex let H_o be the subgraph of G induced by $V(\sigma(o)) \cup A$. Moreover, let B_o be the collection of vertices of G' that are drawn on the boundary of o . Notice

that $|B_o| \leq 3$. As before $A \cup B_o$ is a separator of order at most $\alpha - 1$ separating the vertices of H_o from the rest of G . Hence, H_o may contain at most $\alpha - 1$ vertices of Z . So by the minimality of G , there exists a tree decomposition (T_o, β_o) of adhesion at most α such that the torso of every bag has Euler-genus at most γ after the deletion of an apex set of size at most 4α . Moreover, T_o has a root r_o such that $A \cup (V(H_o) \cap Z) \in \beta_o(r_o)$ and it is a subset of the apex set of the torso of H_o at r_o .

Now let G'' be the graph obtained from G' by deleting the vertex set $V(H_o) \setminus (A \cup Z \cup B_o)$ for every $o \in C(\rho')$ that is not a vertex and adding an edge between every pair of vertices of $A \cup Z$ as well as between every pair of vertices of B_o for every o .

Let T be a tree obtained by introducing a new vertex r and joining it with an edge to every r_o for every $o \in C(\rho')$ that is not a vertex and every r_c where $c \in C(\rho)$ is a vortex. Let $\beta(r) := V(G'')$, $\beta(d) = \beta_o(d)$ if $d \in V(T_o)$ for some o , and $\beta(d) = \beta_c(d)$ if $d \in V(T_c)$ for some c . Then (T, β) is a tree decomposition for G with adhesion at most α . For all $d \neq r$ the torso of G at d has Euler-genus at most γ after deleting an apex set of size at most 4α , we have $Z \subseteq \beta(r)$, and G'' has Euler-genus at most γ after removing the vertices of $A \cup Z$ with $|A \cup Z| \leq 4\alpha$. Hence, (T, β) is the desired tree decomposition for G and our proof is complete. \square

In [23] it is mentioned that, in time $\mathcal{O}(f(|V(H)|) \cdot |G|^3)$, where $f: \mathbb{N} \rightarrow \mathbb{N}$ is some computable function, one either finds the Σ -decomposition of Proposition 24 or a $K_{|V(H)|}$ minor for any graph H . Since all other results and constructions used in the proof of Theorem 44 can be obtained in cubic time, we obtain the following corollary.

Corollary 45. *There exists a computable function $f: \mathbb{N} \rightarrow \mathbb{N}$ such that for every positive integer t , every graph H which is a minor of the shallow vortex grid of order t , and every graph G , we can find in time $\mathcal{O}(f(t) \cdot |G|^3)$ either a minor model of H in G , or a tree decomposition (T, β) with a root $r \in V(T)$ and adhesion at most α such that for every $d \in V(T)$, the torso G_d of G at d has a set $A_d \subseteq V(G_d)$ of size at most 4α for which the graph $G_d - A_d$ has Euler-genus at most γ , where γ and α are defined as in Theorem 44. Moreover, for every $(d_1, d_2) \in E(T)$ we have $|(\beta(d_1) \setminus A_{d_1}) \cap (\beta(d_2) \setminus A_{d_2})| \leq 3$, and if $|(\beta(d_1) \setminus A_{d_1}) \cap (\beta(d_2) \setminus A_{d_2})| = 3$ and $\beta(d_1)$ is larger than 4α , then $\beta(d_1) \setminus A_{d_1} \cap (\beta(d_2) \setminus A_{d_2})$ induces a triangle in $G_{d_1} - A_{d_1}$ which bounds a face.*

Recall that, from Theorem 44, in Corollary 45, $\alpha(x) = \text{poly}(x)$ and $\gamma(x) = 2^{\text{poly}(x)}$.

3.4 Proof of the combinatorial lower bound

In this subsection we establish that Theorem 7 is tight in the sense that the exclusion of a graph H which is not a shallow vortex minor can never guarantee the absence of vortices while, at the same time, provide a global bound on the Euler-genus of all torsos. This shows that the class of shallow vortex minors is exactly the class of all graphs whose exclusion as a minor allows for a version of the Graph Minors Structure Theorem without vortices.

Before we proceed to the proof we give some more definitions and make some observations.

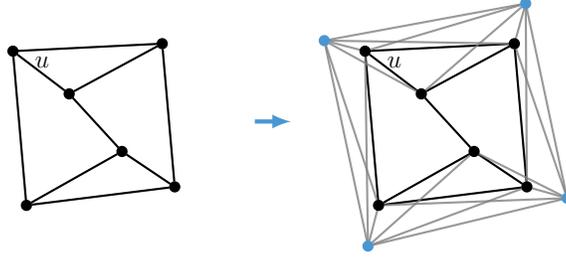


Figure 13: The ring blowup of some cross-free drawing on a disk of a planar graph.

Definition 46 (Ring Blowup Graphs). Let γ is the cross-free drawing on a disk Δ of some (planar) graph G . The face of γ whose closure contains the boundary of Δ is called *external face* of γ . Let Q be the vertices of G that are incident to this (unique) external face.

The *ring blowup* of γ is the graph G' obtained from G by introducing for each $u \in Q$ a new vertex v_u together with the edge uv_u , the edges $\{wv_u \mid w \in N_G(u)\}$, and the edges $\{v_uv_w \mid w \in N_G(u) \cap Q\}$. A graph G is called a *ring blowup graph* if it is a subgraph of some graph isomorphic to the ring blowup of a cross-free drawing on a disk Δ of some planar graph. For instance K_7 is the ring blowup of any cross-free drawing of K_4 on a disk Δ (see Figure 13 for another example).

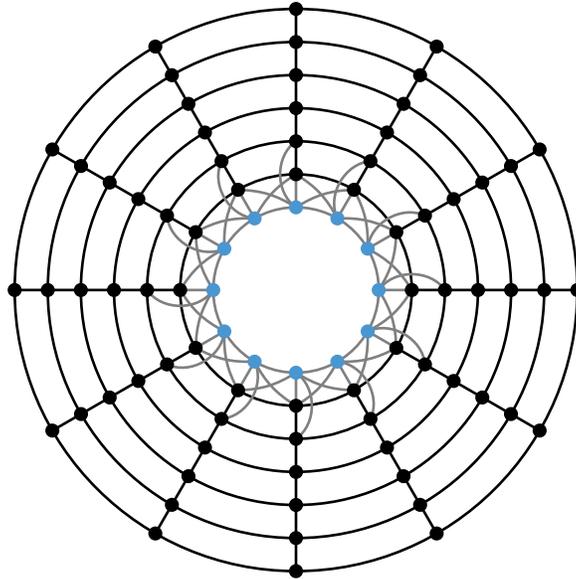


Figure 14: A (6×12) -cylindrical grid ring blowup.

Recall that, given $t \geq 1$ and $s \geq 3$, we define the $(t \times s)$ -cylindrical grid as the graph obtained if we take a $(t \times s)$ -grid and then add the edge $(j, 1)(j, s)$, for every $j \in [t]$. The *extremal cycles* of a $(t \times s)$ -cylindrical grid are the two cycles that have all vertices of degree three (in Figure 2, these cycles are depicted in blue).

A *standard cross-free drawing on a disk Δ* of the $(t \times s)$ -cylindrical grid is one where of the extremal cycles is drawn on the boundary of Δ . We say that a graph is a $(t \times s)$ -cylindrical grid ring

blowup if it is the ring blow up of a standard cross-free drawing on a disk Δ of the $(t \times s)$ -cylindrical grid.

Lemma 47. *There is a function $g : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that every $(t \times s)$ -cylindrical grid ring blowup is a minor of $\mathcal{S}_{g(t,s)}$.*

Proof. The proof is depicted in Figure 15. □

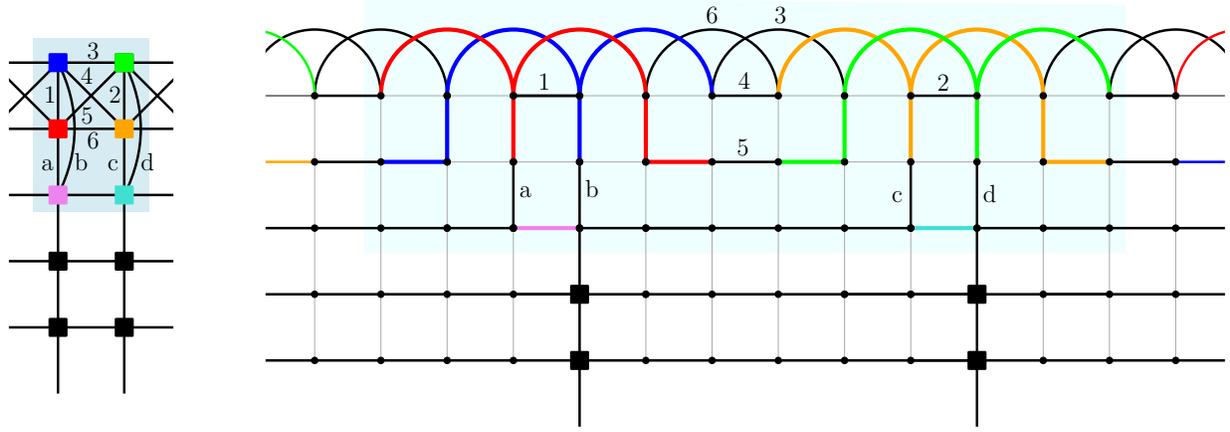


Figure 15: A visualization of the proof of Lemma 47. The leftmost picture is seen as a portion of the blown up part of the $(t \times s)$ -cylindrical grid ring blowup (see Figure 14) and the right part is how this portion is routed through some big enough shallow vortex grid.

Definition 48 (The graph $Q_{s,r}$). Given $s, r \geq 1$ we define the graph $Q_{s,r}$ as follows: We first consider the $(s \times s \cdot r)$ -grid where we denote by $x_1, \dots, x_{s \cdot r}$ the vertices of some of the paths of length $s \cdot r$ where all internal vertices have degree three and the endpoints have degree two. We then introduce r pairs of vertices $\{t_1, t'_1, \dots, t_r, t'_r\}$ (we call them *terminals*) and for every $i \in [r]$ we make t_i and t'_i adjacent with all the vertices in $\{x_{(i-1)+1}, \dots, x_{(i-1)+s}\}$. For example, the graph $Q_{5,4}$ is depicted in Figure 16, where the terminals of $Q_{s,r}$ are the squares vertices.

Using Figure 14 and Figure 16, one may easily verify the following:

Lemma 49. *For every $s, r \geq 1$, $Q_{s,r}$ is a minor of an $((s+1) \times (r \cdot s))$ -cylindrical grid ring blowup.*

We are now ready to prove the main result of this section.

Lemma 50. *There exists a function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that for every $k \in \mathbb{N}$ and every graph G , if $\mathfrak{p}(G) \geq f(k)$, then $\mathfrak{p}_{\text{-ga}}(G) \geq k$. In other words we have $\mathfrak{p} \preceq \mathfrak{p}_{\text{-ga}}$.*

Proof. Let $g : \mathbb{N} \rightarrow \mathbb{N}$ be a function such that $\text{genus}(K_{g(k)}) > k$ and $\text{genus}(K_{3,g(k)}) > k$ (such a function exist because of the standard estimations on the Euler-genus of K_k and $K_{3,k}$, see e.g., [22]). We consider the graph $Q_{s,r}$ where $s = g(k) + k$ and $r = k + 1$. By Lemma 47 and Lemma 49, there is a function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that $Q_{s,r}$ is a minor of $\mathcal{S}_{f(k)}$. Also, it is easy to verify that $\mathfrak{p}_{\text{-ga}}$ is minor-closed. By the definition of \mathfrak{p} , if, for some graph G , $\mathfrak{p}(G) \geq f(k)$, then $\mathcal{S}_{f(k)}$ is a minor of G , which implies that $\mathfrak{p}_{\text{-ga}}(G) \geq \mathfrak{p}_{\text{-ga}}(\mathcal{S}_{f(k)}) \geq \mathfrak{p}_{\text{-ga}}(Q_{s,r})$. It now remains to prove that $\mathfrak{p}_{\text{-ga}}(Q_{s,r}) > k$.

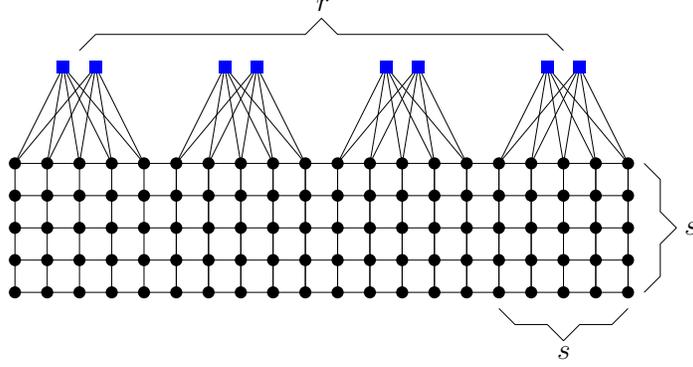


Figure 16: The graph $Q_{s,r}$ (in particular, here $s = 5$ and $t = 4$).

Suppose to the contrary that $Q_{s,r}$ has a tree decomposition (T, β) where every torso G_t contains an apex set A where $|A| \leq k$ and $\text{genus}(G_t - A) \leq k$.

We first rule out the possibility for two terminals x, x' of $Q_{s,r}$ not to belong in the same torso of (T, β) . Indeed, suppose to the contrary that there exist distinct $t, t' \in V(T)$ such that $x \in \beta(t) \setminus \beta(t')$, $x' \in \beta(t') \setminus \beta(t)$. Notice that there are $s = g(k) + k$ internally disjoint paths in $Q_{s,r}$ between x and x' , which implies that $\beta(t) \cap \beta(t') \geq g(k) + k$. As (T, β) is a tree decomposition, this implies that $\beta(t) \cap \beta(t')$ is a subset of some adhesion set of G_t that induces a clique of size $\geq g(k) + k$ in G_t . Let A be the apex set of G_t . The fact that $|A| \leq k$ implies that $G_t - A$ contains a clique of size $\geq g(k)$, therefore $\text{genus}(G_t - A) \geq \text{genus}(K_{g(k)}) > k$, a contradiction. We conclude that there is some $t \in V(T)$ such that all $2r$ terminals of $Q_{s,r}$ belong in G_t .

Notice now that the property that every two terminals are connected by s internally disjoint paths in $Q_{s,r}$ holds for G_t as well because their subpaths that do not belong to G_t can be replaced by edges of cliques induced by the adhesion sets of G_t . As $|A| \leq k$ and $r = k + 1$, at least one pair x, x' of twin terminals is still present in $G_t - A$ as well as a third terminal x'' . Notice that x'' is connected with x (as well as with x') with $s = g(k) + k$ internally disjoint paths of G_t and, among them, at least $g(k)$ also belong to $G_t - A$. As x and x' are twins, these two collections of $g(k)$ paths each differ only on the edges incident to x and x' . Therefore, we may use them in order to find $K_{3,g(k)}$ as a minor of $G_t - A$, therefore $\text{genus}(G_t - A) \geq \text{genus}(K_{3,g(k)}) > k$, a contradiction. \square

4 Algorithmic consequences

With our structural results in place we are now able to discuss how to exploit the decomposition guaranteed by the absence of a large shallow vortex grid as a minor. This section is divided into two parts. First, in [Subsection 4.1](#), we design a dynamic programming algorithm on the decomposition provided by [Theorem 44](#). Then, in [Subsection 4.2](#), we show that, for proper minor-closed classes, this result can be seen as best possible.

4.1 Dynamic programming for the generating function

We present the dynamic programming necessary for the algorithm of [Theorem 8](#) in iterative steps. The algorithm will be performed in a bottom up fashion along the decomposition of [Theorem 44](#).

First we describe the tables we compute in each step for some tuple (G, \mathbf{p}, X) , where G is a graph, \mathbf{p} is a particular labelling of the edges of G , and X is some specified set of “boundary” vertices. Then we provide a subroutine that computes the table of (G, \mathbf{p}, X) from the table of some (H, \mathbf{p}, Y) where the set $(V(G) \setminus V(H)) \cup X \cup Y$ is of bounded size. The next step is the introduction for a subroutine that produces a table for some graph of bounded Euler-genus which is then extended to the case where we encounter a bag of unbounded size in our decomposition, but have already computed all necessary tables for the subtrees below it. The final piece of the algorithm will then be a procedure that merges the tables of several subtrees joined at a common adhesion set.

The generating function of perfect matchings Our goal is to not only count the perfect matchings of a graph, but to also differentiate between perfect matchings of different weight. This will be captured by the generating function of perfect matchings as explained below.

For our purposes it will be convenient to allow for a more general type of edge weighting where edges are labelled by fractions of integer polynomials. This will allow us to encode parts of the tables of our dynamic programming directly on the edges. For this purpose let $\mathbb{Z}[x]$ be the set of all polynomials with integer coefficients and let

$$\mathbb{Z}(x) := \left\{ \frac{p}{q} \mid p, q \in \mathbb{Z}[x], q \neq 0 \right\}.$$

Now let $\mathbf{p}: E(G) \rightarrow \mathbb{Z}(x)$, we call (G, \mathbf{p}) an *labelled graph* and \mathbf{p} is its *labeling*. If (G, \mathbf{w}) is an edge-weighted graph, we derive a labelled graph $(G, \mathbf{p}_\mathbf{w})$ from (G, \mathbf{w}) by setting, for every $e \in E(G)$,

$$\mathbf{p}_\mathbf{w}(e) := x^{\mathbf{w}(e)}.$$

We will proceed to give general definitions in terms of labelings, but using the definition of $\mathbf{p}_\mathbf{w}$ one can easily derive the more conventional versions of these concepts.

Now assume we are given a labelled graph (G, \mathbf{p}) with a perfect matching M . We express the total weight of M under \mathbf{p} as the monomial

$$\mathbf{p}(M) := \prod_{e \in M} \mathbf{p}(e).$$

Definition 51 ((labelled) Generating Function of Perfect Matchings). Let (G, \mathbf{p}) be a labelled graph. The *labelled generating function of the perfect matchings of G* , usually abbreviated as the *\mathbf{p} -generating function* or the *generating function*, is defined to be the polynomial

$$\text{PerfMatch}(G, \mathbf{p}) := \sum_{M \in \mathcal{M}(G)} \mathbf{p}(M).$$

Let (G, \mathbf{w}) be an edge-weighted graph, then the *generating function of all weighted perfect matchings in G* is the polynomial

$$\text{PerfMatch}(G, \mathbf{w}) := \sum_{M \in \mathcal{M}(G)} \mathbf{p}_\mathbf{w}(M).$$

Our aim is to eventually compute the generating function of all weighted perfect matchings of an edge-weighted graph (G, \mathbf{w}) . For this we will compute $\text{PerfMatch}(G, \mathbf{w})$ from a series of partial generating functions of labelled subgraphs of G with some specified set of “boundary” vertices, where the labelings are derived from \mathbf{w} .

Let G be a graph with a perfect matching and $F \subseteq E(G)$ be some (not necessarily perfect) matching in G . We say that a vertex $v \in V(G)$ is *covered* by F if v is an endpoint of some edge in F , the vertex v is said to be *exposed* by F if it is not covered by it. We say that F is *extendable* if there exists a perfect matching $M \in \mathcal{M}(G)$ such that $F \subseteq M$. In other words, F is extendable if it can be extended to a perfect matching of the entire graph.

Given a pair (G, X) where $X \subseteq V(G)$ we say that a matching F of G is *internally extendible* (or, simply, *extendible*) in (G, X) if F is an extendable matching in the graph $G - (X \setminus V(F))$. We also say that F is *aligned* with (G, X) if it is internally extendible in (G, X) and every edge of F contains some endpoint in X .

Roughly speaking, F being aligned with (G, X) means that F classifies exactly the matchings of G which are extensions of F , leave the vertices of $X \setminus V(F)$ exposed, but cover every other vertex $G - (X \setminus V(F))$. Another way of explaining the idea behind aligned matchings is the following. Let G be a graph with a perfect matching, let (A, B) be a separation of G and let $H := G[B]$. Consider the pair $(H, A \cap B)$. Now for every perfect matching M of G we can find the set $F_M \subseteq M$ of exactly those edges in M with both endpoints in H and at least one endpoint in $A \cap B$. This allows us to define an equivalence relation on $\mathcal{M}(G)$ by saying that two perfect matchings M and M' are equivalent if $F_M = F_{M'}$. Hence, each of these equivalence classes is uniquely defined by some set $F \subseteq E(H)$ which is aligned with $(H, A \cap B)$.

Let us denote by $\text{Aligned}(G, X)$ the collection of all matchings that are aligned with (G, X) .

Please note that testing whether a given matching F is extendable can be done in polynomial time by simply checking if $G - V(F)$ has a perfect matching using known polynomial time algorithms for maximum matching [15]. It follows that also checking for the extendability of some pair (F, X) as above is possible in polynomial time.

A pair (G, X, \mathbf{p}) where (G, \mathbf{p}) is a labelled graph and $X \subseteq V(G)$ is a set of vertices is called a (*labelled*) *boundary graph*. We say that X is the *boundary* of (G, X, \mathbf{p}) . Let $H \subseteq G$ be an induced subgraph of G , and $Y \subseteq V(H)$ be a set of vertices in H such that in $G - Y$ there is no edge with one endpoint in $H - Y$ and the other in $G - H$. We call (H, Y, \mathbf{p}) an (*labelled*) *boundary subgraph* of (G, X, \mathbf{p}) , with *boundary* Y . In a slight abuse of notation we use \mathbf{p} in (H, Y, \mathbf{p}) , but actually mean the restriction $\mathbf{p}|_{E(H)}$ of \mathbf{p} to the edges of H . Note that (H, Y, \mathbf{p}) is itself also a boundary graph.

Definition 52 (Partial Generating Function). Let (G, X, \mathbf{p}) be a labelled boundary graph and $F \in \text{Aligned}(G, X)$. The *partial generating function of (G, X) under F* is defined to be

$$\text{PerfMatch}(G, X, \mathbf{p}, F) := \left(\prod_{e \in F} \mathbf{p}(e) \right) \cdot \text{PerfMatch}(G - V(F) - X, \mathbf{p}).$$

We define the set

$$\text{Gen}_{(G, X, \mathbf{p})} := \{ \text{PerfMatch}(G, X, \mathbf{p}, F) \mid F \in \text{Aligned}(H, X) \}$$

to be the collection of all partial generating functions of (G, X, \mathbf{p}) under its aligned matchings. For every $F \in \text{Aligned}(H, X)$ we use $\text{Gen}_{(G, X, \mathbf{p})}(F)$ to denote $\text{PerfMatch}(G, X, \mathbf{p}, F)$.

Note that Gen is technically a mapping itself. We write its argument in the subscript to better distinguish $\text{Gen}_{(G,X,\mathbf{p})}$ and $\text{Gen}_{(G,X,\mathbf{p})}(F)$.

The definitions above correspond a set $\text{Gen}_{(G,X,\mathbf{p})}$ to every boundary graph (G, X, \mathbf{p}) . This collection of partial generating functions will serve as the table of our dynamic programming for (G, X, \mathbf{p}) .

Bags of Bounded Size We start the discussion of our algorithm with a way to compute the table entries for our dynamic programming in the case where we are concerned with a single bag of bounded size and we are given the unified table for all of its children. The lemma presented here could as well be regarded as a special case of the way we treat bags of unbounded size by simply assuming that the apex set equals the entire bag. However, we deem this special case to be a nice illustration of how we compute our tables.

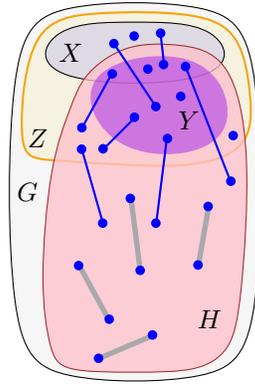


Figure 17: A visualization of the bounded graphs (G, X) and (H, Y) and of the set Z in the statement of [Lemma 53](#). Also the matching F (in turquoise) is aligned with (G, Z) .

Lemma 53. *Let $k \geq 1$ be a positive integer. Let (G, X, \mathbf{p}) be a labelled boundary graph and (H, Y, \mathbf{p}) be a labelled boundary subgraph of (G, X, \mathbf{p}) such that there exists a set $Z \subseteq V(G)$ with $X, Y \subseteq Z$, $G - Z = H - Y$, and $|Z| \leq k$. There exists an algorithm that, given the set $\text{Gen}_{(H,Y,\mathbf{p})}$, computes the set $\text{Gen}_{(G,X,\mathbf{p})}$ in time $|G|^{\mathcal{O}(k)}$.*

Proof. Notice that (G, Z, \mathbf{p}) is also a labelled boundary graph. Let \mathcal{F} be the collection of all matchings F in G such that $F \in \text{Aligned}(G, Z)$ and observe that, with $|Z| \leq k$, we have $|\text{Aligned}(G, Z)| \in \mathcal{O}(|G|^{2k})$. Moreover, it takes time $\mathcal{O}(|G|^c)$ to check for any such matching F whether it is extendable in $G - (Z \setminus V(F))$, for some constant c . Hence, we can find \mathcal{F} , and $\text{Aligned}(G, Z)$ as a consequence, in time $\mathcal{O}(|G|^{2k+c})$. Using the same argument and the fact that $X, Y \subseteq Z$, we are able to find the sets $\text{Aligned}(G, X)$ and $\text{Aligned}(H, Y)$. (See [Figure 17](#) for a visualization of the bounded graphs (G, X) and (H, Y) .)

Now each $F \in \mathcal{F}$ can be covered by the sets F_1, F_2 , and F_3 as follows.

- Let F_1 be the set of edges in F with at least one endpoint in X ,
- let F_2 be the set of edges in F with both endpoints in $V(H)$, and

- let F_3 be the set of edges in F with at most one endpoint in Y .

It follows that, for every $F \in \mathcal{F}$, $F_1 \in \text{Aligned}(G, X)$, $F_2 \in \text{Aligned}(H, Y)$, and $F_3 = F \setminus F_2$.

For each $R \in \text{Aligned}(G, X)$ let $\mathcal{F}_R \subseteq \mathcal{F}$ be the collection of matchings $F \in \mathcal{F}$ such that $R = F_1$. Notice that $\{\mathcal{F}_R \mid R \in \text{Aligned}(G, X)\}$ is a partition of \mathcal{F} . Moreover, let $F' \in \text{Aligned}(G, X)$. Then $G - (X \setminus V(F'))$ has a perfect matching, say M , with $F' \subseteq M$ and we may set F'' to be the set of edges from M with at least one endpoint in Z . It follows that $F' \subseteq F'' \in \text{Aligned}(G, Z) = \mathcal{F}$. Thus, for each $R \in \text{Aligned}(G, X)$ there is at least one $F \in \mathcal{F}$ such that $R = F_1$.

By definition, for each $R \in \text{Aligned}(G, X)$, $\text{Gen}_{(G, X, \mathbf{p})}(R)$ is the generating function for all perfect matchings of $G - (X \setminus V(R))$ that contain R . Let M be such a perfect matching of $G - (X \setminus V(R))$, then M can be partitioned in to four sets as follows

- $M_1 := R$,
- M_2 , which is the set of all edges in $M \setminus R$ with at most one endpoint in H ,
- M_3 , which is the set of all edges of $M \setminus R$ with both endpoints in H and at least one endpoint in Y , and
- $M_4 := M \setminus (M_1 \cup M_2 \cup M_3)$.

We may further split R into the sets $R_1 := R \cap E(H)$ and $R_2 := R \setminus E(H)$. Observe the following:

1. By the definition of boundary subgraphs and the choice of Z we have that every edge of M with exactly one endpoint in H has an endpoint in Y .
2. $M_1 \cup M_2 \cup M_3$ is aligned with (G, Z) .
3. $R_1 \cap M_3$ is aligned with (H, Y) and the vertices in $Y \setminus V(R_1 \cap M_3)$ are exactly those vertices of H which either belong to $X \setminus V(R)$ or are covered by $M_1 \cup M_2$.

Now consider some $F \in \mathcal{F}_R$ and recall the bipartition of F into F_2 and F_3 as above. As $R \subseteq F$ it follows that $R_1 = F_2 \cap R$ and $R_2 = F_3 \cap R$ form a bipartition of the set R itself. By definition we have that

$$\begin{aligned} \text{Gen}_{(G, Z, \mathbf{p})}(F) &= \left(\prod_{e \in F} \mathbf{p}(e) \right) \cdot \frac{\text{Gen}_{(H, Y, \mathbf{p})}(F_2)}{\prod_{e' \in F_2} \mathbf{p}(e')} \\ &= \left(\prod_{e \in F_3} \mathbf{p}(e) \right) \cdot \text{Gen}_{(H, Y, \mathbf{p})}(F_2). \end{aligned}$$

It follows that

$$\begin{aligned} \text{Gen}_{(G, X, \mathbf{p})}(R) &= \left(\prod_{e \in R} \mathbf{p}(e) \right) \cdot \sum_{F \in \mathcal{F}_R} \frac{\text{Gen}_{(G, Z, \mathbf{p})}(F)}{\prod_{e \in R} \mathbf{p}(e)} \\ &= \left(\prod_{e \in R} \mathbf{p}(e) \right) \cdot \sum_{F \in \mathcal{F}_R} \frac{\left(\prod_{e \in F_3} \mathbf{p}(e) \right) \cdot \text{Gen}_{(H, Y, \mathbf{p})}(F_2)}{\prod_{e \in R} \mathbf{p}(e)} \end{aligned}$$

$$\begin{aligned}
&= \left(\prod_{e \in R} \mathbf{p}(e) \right) \cdot \sum_{F \in \mathcal{F}_R} \left(\prod_{e \in F_3 \setminus R} \mathbf{p}(e) \right) \cdot \frac{\text{Gen}_{(H,Y,\mathbf{p})}(F_2)}{\prod_{e' \in F_2 \cap R} \mathbf{p}(e')} \\
&= \sum_{F \in \mathcal{F}_R} \left(\prod_{e \in F_3} \mathbf{p}(e) \right) \cdot \text{Gen}_{(H,Y,\mathbf{p})}(F_2).
\end{aligned}$$

Given the sets \mathcal{F} , $\text{Aligned}(G, X)$, and $\text{Gen}_{(H,Y)}$ it follows that $\text{Gen}_{(G,X)}(R)$ can be found in time $|G|^{\mathcal{O}(k)}$ and, since $|\text{Aligned}(G, X)| = \mathcal{O}(|G|^k)$, our claim follows. \square

Generating Functions, Permanents, and Pfaffian Orientations. The generating function of planar graphs and, in general, of graphs of bounded Euler-genus, is usually computed using Pfaffian orientations (in the case of planar graphs) or a linear combination of many different orientations derived from Pfaffian orientations (in the case of general graphs of bounded Euler-genus). Before we go on we introduce this notion and explain how the generating function of an edge-weighted graph of bounded Euler-genus can be found using this concept. We then elaborate on this a bit more and explain how we can, essentially, replace the monomial of an edge by a more complicated polynomial in order to encode entire generating functions linked to this edge. This last step is of particular importance as it provides us with a way to handle the (possibly) unboundedly many different sets of size at most three onto which subtrees below a bag of unbounded size might attach.

An *orientation* of a graph G is a digraph \vec{G} with vertex set $V(G)$ whose edge set is obtained by introducing for every edge $uv \in E(G)$ exactly one of the edges (u, v) or (v, u) . Let C be an even cycle of G and let \vec{G} be an orientation of G . C is said to be *oddly oriented* by \vec{G} if it has an odd number of directed edges in agreement with the clockwise traversal of C . Notice that, since C is even, if it is oddly oriented it must also have an odd number of directed edges in agreement with the counterclockwise traversal, so the property of being oddly oriented does not depend on the direction of traversal after all. Similarly, we say that a cycle C of G is *evenly oriented* by \vec{G} if it is not oddly oriented.

Let G be a graph with a perfect matching and C be an even cycle. We call C a *conformal cycle* if $G - V(C)$ has a perfect matching. Notice here that, since C is even, it has two disjoint perfect matchings and thus every perfect matching of $G - V(C)$ can be completed to a perfect matching of G by choosing one of these two matchings. If M is a perfect matching of G and C is a conformal cycle in G such that M contains a perfect matching of C we say that C is *M -conformal*.

Definition 54 (Sign Polynomials of Matchings and Orientations). Let (G, \mathbf{p}) be a labelled graph with a perfect matching M and let \vec{G} be an orientation of G . For every perfect matching N of G let $\text{sgn}(\vec{G}, M, N) := (-1)^n$, where n is the number of M -conformal cycles of G which are also N -conformal and evenly oriented by \vec{G} . We define the *M -polynomial* of \vec{G} as follows

$$\mathbf{p}_{\text{sgn}}(\vec{G}, M, \mathbf{p}) := \sum_{N \in \mathcal{M}(G)} \text{sgn}(\vec{G}, M, N) \cdot \mathbf{p}(N).$$

Definition 55 (Pfaffians of Skew-Symmetric Matrices). Let (G, \mathbf{p}) be a labelled graph and \vec{G} be an orientation of G such that $|G| = 2n$ for some $n \in \mathbb{N}$. We denote by $\mathbf{A}(\vec{G}, \mathbf{p})$ the skew-symmetric

matrix with rows and columns indexed by $V(G)$, where $a_{uv} = \mathbf{p}(uv)$ in case $(u, v) \in E(\vec{G})$, $a_{uv} = -\mathbf{p}(uv)$ if $(v, u) \in E(\vec{G})$, and $a_{uv} = 0$ otherwise. The *Pfaffian* of $\mathbf{A}(\vec{G}, \mathbf{p})$ is defined as

$$\text{pf}(\mathbf{A}(\vec{G}, \mathbf{p})) := \sum_{\pi} \mathbf{s}(\pi) \cdot a_{i_1, j_1} \cdots a_{i_n, j_n}$$

where $\pi = \{\{i_1, j_1\}, \dots, \{i_n, j_n\}\}$ is a partition of the set $[2n]$ into pairs $i_k < j_k$ for every $k \in [n]$, and $\mathbf{s}(\pi)$ equals the sign of the permutation $i_1 j_1 \dots i_n j_n$ of $12 \dots (2n)$.

Each non-zero term of the expression of the Pfaffian of $\mathbf{A}(\vec{G}, \mathbf{p})$ equals $\mathbf{p}(M)$ or $-\mathbf{p}(M)$ where M is a perfect matching of G . We denote by $\mathbf{s}(\vec{G}, M)$ the sign of the term $\mathbf{p}(M)$ in said expression of the Pfaffian of $\mathbf{A}(\vec{G}, \mathbf{p})$. Notice that $\mathbf{s}(\vec{G}, M)$ is independent of \mathbf{p} as it only depends on M , the encoding the permutation, and \vec{G} determining whether $\mathbf{p}(uv)$ appears with its original sign or the opposing one. It has been observed (see e.g., [17]) that

$$\text{pf}(\mathbf{A}(\vec{G}, \mathbf{p})) = \sum_{M \in \mathcal{M}(G)} \mathbf{s}(\vec{G}, M) \cdot \mathbf{p}(M).$$

The following theorem was proven by Kasteleyn [26] for the case where $\mathbf{p}(e) = x^{w(e)}$, $w(e) \in \mathbb{Z}$. It still holds for arbitrary labelings \mathbf{p} and we state it here in this more general form.

Proposition 56 ([26]). *Let (G, \mathbf{p}) be a labelled graph with two perfect matchings M, N , and an orientation \vec{G} . Then $\mathbf{s}(\vec{G}, N) = \mathbf{s}(\vec{G}, M) \cdot \text{sgn}(\vec{G}, M, N)$. Hence, for every perfect matching N ,*

$$\begin{aligned} \text{pf}(\mathbf{A}(\vec{G}, \mathbf{p})) &= \sum_{N \in \mathcal{M}(G)} \mathbf{s}(\vec{G}, N) \cdot \mathbf{p}(N) \\ &= \mathbf{s}(\vec{G}, M) \sum_{N \in \mathcal{M}(G)} \text{sgn}(\vec{G}, M, N) \cdot \mathbf{p}(N) = \mathbf{s}(\vec{G}, M) \cdot \mathbf{p}_{\text{sgn}}(\vec{G}, M, \mathbf{p}). \end{aligned}$$

While computing the permanent, and thus computing the generating function for perfect matchings is generally a $\#\text{P}$ -hard problem [57], the Pfaffian of a skew-symmetric matrix can be expressed through its determinant and is thereby efficiently computable.

Proposition 57 ([6, 27]). *Let (G, \mathbf{p}) be a labelled graph and an orientation \vec{G} , then*

$$\text{pf}(\mathbf{A}(\vec{G}, \mathbf{p}))^2 = \det(\mathbf{A}(\vec{G}, \mathbf{p})).$$

In case a graph G has an orientation \vec{G} such that the signs produced by $\mathbf{s}(\vec{G}, \cdot)$ and $\text{sgn}(\vec{G}, \cdot, \cdot)$ would be independent of the matchings, this technique could be used to find the generating functions for the perfect matchings of G efficiently. This fact lead Kasteleyn to the definition of the so called ‘‘Pfaffian orientations’’, which are orientations that ensure $\text{sgn}(\vec{G}, M, N) = 1$ for all choices of M and N .

Definition 58 (Pfaffian Orientation). Let G be a graph with a perfect matching M . An orientation \vec{G} is called a *Pfaffian orientation* if every M -conformal cycle of G is oddly oriented by \vec{G} . If G has a Pfaffian orientation, G is called *Pfaffian*.

As a first application of this idea, Kasteleyn showed that every planar graph is Pfaffian [26]. He generalized this idea and stated that the generating function of a graph G of orientable genus

g could be expressed as a linear combination of 4^g Pfaffians of different orientations of G . This was later turned into a theorem by Galluccio and LoebL [17] and adapted as a time $\mathcal{O}_k(|G|^{\mathcal{O}(1)})$ algorithm for edge-weighted graphs whose weights are bounded in size by some polynomial in the size of the graph by Galluccio, LoebL, and Vondrák [18]. An extension of this idea to graphs of bounded Euler-genus was found by [52]. He uses 2^g many orientations, where g now is the Euler-genus of G . Both approaches were unified into a framework that uses planarizing gadgets instead of orientations by Curticapean and Xia in [9].

These algorithms are the centerpiece of our own algorithm as they allow us to find the generating functions for the torsos of unbounded size in our decomposition. We slightly adapt the formulation of the results of Galluccio, LoebL, Vondrák, and Tesler to match our more general setting of labelled graphs. Since the edge weights only come into play in the computation of the 2^g Pfaffians and their linear combination, the more general version of their result still holds.

Let (G, \mathbf{p}) be a labelled graph. We say that \mathbf{p} is *polynomially bounded* if there exists a polynomial \mathbf{p} such that the degrees of the polynomials in

$$\left\{ p, q \in \mathbb{Z}[x] \mid \text{there exists } e \in E(G) \text{ s.t. } \mathbf{p}(e) = \frac{p}{q} \text{ is fully reduced} \right\}$$

are bounded by $\mathbf{p}(|G|)$.

Proposition 59 ([18, 52]). *Let $g \in \mathbb{N}$ be an integer and (G, \mathbf{b}) be a labelled graph whose labeling is polynomially bounded. There exists an algorithm that computes the labelled generating function of all perfect matchings of (G, \mathbf{p}) in time $\mathcal{O}_k(|G|^{\mathcal{O}(1)})$, where $k = \text{genus}(G)$.*

There are two important observations to take away from this quick introduction to Pfaffians and, in particular, from Proposition 59. The first is, that the generating function of (labelled) perfect matchings can be expressed as a linear combination of Pfaffians of the skew-symmetric matrices corresponding to some orientations of a bounded-genus graph. The second is that the Pfaffian can be expressed as a function of the determinant of a skew-symmetric matrix, which is polynomially computable.

A polynomial algorithm for apex-bounded-genus graphs. The next subroutine for our algorithm will be a way to produce the generating function for the labelled perfect matchings of graphs that have Euler-genus at most t after the deletion of at most t vertices. The lemma we prove here can be seen as a slight generalization of Lemma 53 which now also incorporates Proposition 59.

Lemma 60. *Let k be a positive integer. Let (G, X, \mathbf{p}) be a labelled boundary graph with $|X| \leq k$ and assume there exists a set $A \subseteq V(G)$ with $|A| \leq k$ such that the Euler-genus of $G - A$ is at most k . There exists an algorithm that computes in time $|G|^{\mathcal{O}(k)}$ the set $\text{Gen}_{(G, X, \mathbf{p})}$.*

Proof. Notice that $(G, A \cup X, \mathbf{p})$ is a boundary graph with a boundary of size at most $2k$. Let \mathcal{F} be the collection of all matchings F such that $F \in \text{Aligned}(G, A \cup X)$ and every vertex of $A \setminus X$ is covered by an edge of F . As discussed in the proof of Lemma 53 we have $|\mathcal{F}| \in |G|^{\mathcal{O}(k)}$ and \mathcal{F} , as well as the set $\text{Aligned}(G, X)$, can be found in time $|G|^{\mathcal{O}(k)}$. Each set $F \in \mathcal{F}$ can be partitioned into two sets as follows:

- Let F_1 be the set of edges in F with at least one endpoint in X , and

- let $F_2 := F \setminus F_1$.

Observe that for each $F \in \mathcal{F}$ we have $F_1 \in \text{Aligned}(G, X)$, and $G - (X \cup V(F)) \subseteq G - A$ has Euler-genus at most k . For every $R \in \text{Aligned}(G, X)$ let $\mathcal{F}_R \subseteq \mathcal{F}$ be the collection of all sets $F \in \mathcal{F}$ with $F_1 = R$. It follows that

$$\begin{aligned} \text{Gen}_{(G, X, \mathbf{p})}(R) &= \left(\prod_{e \in R} \mathbf{p}(e) \right) \cdot \text{PerfMatch}(G - V(R) - X, \mathbf{p}) \\ &= \left(\prod_{e \in R} \mathbf{p}(e) \right) \cdot \sum_{F \in \mathcal{F}_R} \left(\prod_{e \in F \setminus R} \mathbf{p}(e) \right) \cdot \text{PerfMatch}(G - V(F) - X, \mathbf{p}) \\ &= \sum_{F \in \mathcal{F}_R} \left(\prod_{e \in F} \mathbf{p}(e) \right) \cdot \text{PerfMatch}(G - V(F) - X, \mathbf{p}). \end{aligned}$$

Since $G - A$ has Euler-genus at most k and there are at most $|G|^{\mathcal{O}(k)}$ many sets in \mathcal{F}_R , by calling the algorithm from [Proposition 59](#) $|G|^{\mathcal{O}(k)}$ many times we can compute $\text{PerfMatch}(G - (V(F) \cup X), \mathbf{p})$ for every $R \in \text{Aligned}(G, X)$ and every $F \in \mathcal{F}_R$ in time $|G|^{\mathcal{O}(k)}$ and thus our claim follows. \square

Bags of unbounded size. We are now ready to generalize [Lemma 60](#) to the setting where the boundary graph (H, X) arises as the torso of some bag in the structural decomposition provided by [Theorem 44](#). For a detailed introduction into the theory of graph minors we refer the reader back to [Subsection 2.1](#). The definition of a Σ -decomposition is given in [Definition 12](#) and [Theorem 44](#) is the relevant structure theorem for this part.

As a quick reminder, [Theorem 44](#) provides, for every graph G which excludes the shallow vortex grid as a minor, a tree decomposition where the torso of the bag at any vertex t is a graph of bounded Euler-genus after the deletion of a small set of vertices called the *apex set*. Moreover, each node of this tree decomposition is equipped with such an apex set such that the adhesion set of any edge dt , that is the intersection of the bags of t and d , in the decomposition tree avoids the apex set of t in at most three vertices. Additionally, the torso of G at t , after deleting the corresponding apex set, is embedded in a surface of bounded Euler-genus in such a way that the, at most three, vertices in the adhesion of dt sit on a common face.

More formally, this means that we have to address the following issue. Let G be a graph excluding some shallow vortex minor and let (T, β) be the decomposition of G provided by [Theorem 44](#) where $r \in V(T)$ is the root of T . Now let $t \in V(T)$ be some vertex with $|\beta(t)| > 4\alpha(t)$ and let d_1, \dots, d_ℓ be the children of t . For each $i \in [\ell]$ the intersection $\beta(t) \cap \beta(d_i)$ may contain up to three vertices which do not belong to the apex set of the torso G_t of G at t . Moreover, the number ℓ is unbounded. Hence, if we were to generalize our approach from [Lemma 60](#) directly we would need to produce a set \mathcal{F} of matchings not only covering the boundary $\beta(t) \cap \beta(t')$ to the ancestor of t , but also covering $\beta(t) \cap \beta(d_i)$ for all $i \in [\ell]$. As a result, we would be unable to bound the size of \mathcal{F} .

To get a better grip on the situation let us introduce some more definitions.

Definition 61 (Branching of a Graph). Let $k \in \mathbb{N}$ be some integer. Let (G, X, \mathbf{p}) be a labelled boundary graph with an *apex set* $A \subseteq V(G)$. Let $(B_1, Y_1, \mathbf{p}), \dots, (B_\ell, Y_\ell, \mathbf{p})$ be boundary subgraphs of (G, X, \mathbf{p}) . We call $\mathcal{B} = ((G, X), A, \mathbf{p}, (B_1, Y_1), \dots, (B_\ell, Y_\ell))$ a *k-branching* (of (G, X, \mathbf{p})) if

- For every $i \in [\ell]$, $(V(B_i) \setminus Y_i) \cap A = \emptyset$.

- $|X \setminus A| \leq 3$ and, for all $i \in [\ell]$, $|Y_i \setminus A| \leq 3$,
- $|X|, |A| \leq k$ and, for all $i \in [\ell]$, $|Y_i| \leq k$,
- if $i \neq j \in [\ell]$ then $V(B_i) \cap V(B_j) \subseteq Y_i \cap Y_j$,
- if $i \neq j \in [\ell]$ then neither $Y_i \setminus A \subseteq Y_j \setminus A$ nor $Y_j \setminus A \subseteq Y_i \setminus A$ holds,
- if $G_{\mathcal{B}}$ is the subgraph obtained from $G - (\bigcup_{i \in [\ell]} (V(B_i) \setminus Y_i))$ by turning every set Y_i , $i \in [\ell]$, and the set X into cliques, then $G_{\mathcal{B}} - A$ has a drawing γ , without crossings, on some surface Σ , without boundary, of Euler-genus at most k ,
- the vertices of X are incident to a face of γ and, for every $i \in [\ell]$, the vertices of $Y_i \setminus A$ are incident to a face of γ .

See [Figure 18](#) for a visualisation of [Definition 61](#).

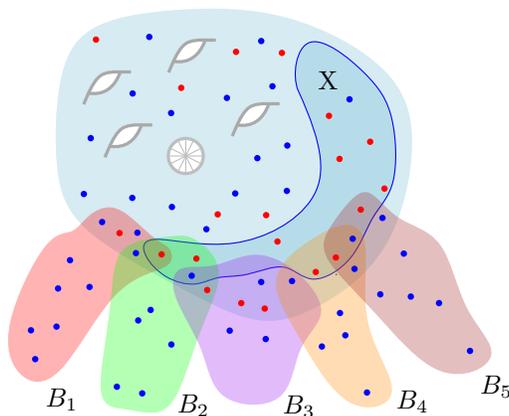


Figure 18: A visualization of [Definition 61](#). The **red** vertices depict the vertices in A . For each B_i , the boundary Y_i consists of the vertices in the intersection of the **light blue** territory and the territory corresponding in B_i .

The (B_i, Y_i, \mathbf{p}) are called the *branches* of \mathcal{B} . Let (G, X, \mathbf{p}) be a labelled boundary graph and let $\mathcal{B} = ((G, X), \mathbf{p}, A, (B_1, Y_1), \dots, (B_\ell, Y_\ell))$ be a k -branching of (G, X, \mathbf{p}) . We call a function $\mathbf{sp}: [\ell] \rightarrow 2^{\mathbb{Z}^{[x]}}$ the *sign post* of \mathcal{B} if for every $i \in [\ell]$ we have $\mathbf{sp}(i) = \text{Gen}_{(B_i, Y_i, \mathbf{p})}$.

Our goal is to compute the family of all generating functions $\text{Gen}_{(G, X, \mathbf{p})}$ for the labelled boundary graph (G, X, \mathbf{p}) in the case where we are given a branching \mathcal{B} together with its sign post \mathbf{sp} .

Matchgates. To manage the problem of the unbounded number of non-trivial residual boundaries in a branching, we employ the technique from [\[50\]](#) which is similar to Valiant’s “matchgates” [\[58\]](#). matchgates were applied for the problem of counting perfect matchings in single crossing minor-free graphs in [\[7\]](#). However, for our purpose of producing the entire generating function the gadgets from [\[50\]](#) appear to be more suited.

If (G, X, \mathbf{p}) is some labelled boundary graph and $\mathcal{B} = ((G, X), A, \mathbf{p}, (B_1, Y_1), \dots, (B_\ell, Y_\ell))$ is a k -branching of (G, X, \mathbf{p}) , then what remains of any B_i in $G_{\mathcal{B}} - A$ is at most a triangle and this triangle bounds a face of γ . Let F be matching in G such that

- every edge in F covers some vertex in $A \cup X$,
- every vertex of $A \setminus X$ is covered by an edge of F , and
- there exists a perfect matching M of $G - V(F) - X$.

Now define the F -reduced branching $\mathcal{B}_F := ((G_F, \emptyset), \emptyset, \mathbf{p}, (B_{1,F}, Y_{1,F}), \dots, (B_{\ell,F}, Y_{\ell,F}))$, where

$$\begin{aligned} G_F &:= G - V(F) - X, \\ B_{i,F} &:= B_i - V(F) - X \text{ for all } i \in [\ell], \text{ and} \\ Y_{i,F} &:= Y_i \setminus (V(F) \cup X) \text{ for all } i \in [\ell]. \end{aligned}$$

Note that every $(B_{i,F}, Y_{i,F}, \mathbf{p})$ still is a labelled boundary subgraph of $(G - V(F) - X, \emptyset, \mathbf{p})$. Suppose we are given, for every $i \in [\ell]$, the set $\text{Gen}_{(B_{i,F}, Y_{i,F}, \mathbf{p})}$. Let us consider some $i \in [\ell]$. There are eight possible cases that might arise. In each of these cases, we replace the entire graph $B_{i,F}$ by a certain graph $J_{i,F}$ together with a labelling $\lambda_{i,F}$ of the edges of $J_{i,F}$ with elements from $\text{Quot}(\mathbb{Z}[x])$. These polynomial fractions will then be used to encode the generating functions of $(B_{i,F}, Y_{i,F}, \mathbf{p})$ for a skew-symmetric matrix. The Pfaffian of said matrix will be used (via [Proposition 59](#)) to compute the generating function for the perfect matchings in $(G_F, \emptyset, \mathbf{p})$. Formally, we will replace the subgraph $B_{i,F}$ by the graph $J_{i,F}$, thereby producing some labelled boundary graph $(G', \emptyset, \mathbf{p}')$, and adjust the function \mathbf{p}' by setting $\mathbf{p}'(e) := \lambda_{i,F}(e)$ for all $e \in E(J_{i,F})$, while \mathbf{p}' equals \mathbf{p} on all other edges of G' .

We now describe the *matchgates* that we will be using. In the non-trivial cases, we will attribute the respective matchgate with a property called “representativeness”. In this property we subsume the different ways in which perfect matchings of G' can interact with the matchgate. The role of the representativeness is to ensure that $(J_{i,F}, \lambda_{i,F})$ correctly encodes the generating function of the boundary graph it represents.

Case 1: $Y_{i,F} = \emptyset$. In this case $J_{i,F}$ is just the empty graph without any vertices and $\lambda_{i,F}$ is empty as well.

Case 2: $Y_{i,F} = \{a\}$. In this case $J_{i,F}$ consists exactly of the vertex a and $\lambda_{i,F}$ is empty.

Case 3: $Y_{i,F} = \{a, b\}$ and $|V(B_{i,F})|$ is even. Here we define $J_{i,F}$ to be the graph with vertex set $\{a, b, u, v\}$ together with the edges $\{au, uv, vb\}$, where u and v are newly introduced vertices. Note that, since $|V(B_{i,F})|$ is even, every perfect matching of G_F either covers both a and b with edges of $B_{i,F}$, or none of them. Moreover, after replacing $B_{i,F}$ with $J_{i,F}$, every perfect matching of the resulting graph G'_F either contains the edges au and bv , or the edge uv , let us call this fact the *representativeness* of $J_{i,F}$. The labels $\lambda_{i,F}$ should now express two possible states; the contribution of the edges au and bv to the Pfaffian of the skew-symmetric matrix we want to construct should equal the labelled generating function of all perfect matchings of $B_{i,F}$, let us call this function \mathbf{p}_\emptyset since no vertex from $\{a, b\}$ is matched outside of $B_{i,F}$, while the contribution of the edge uv should equal the labelled generating function of all perfect

matchings of $B_{i,F} - a - b$, we denote this function by \mathfrak{p}_{ab} . Finally, we set $\lambda_{i,F}(au) := \mathfrak{p}_\emptyset$, $\lambda_{i,F}(bv) := 1$, and $\lambda_{i,F} := \mathfrak{p}_{ab}$.

Case 4: $Y_{i,F} = \{a, b\}$ and $|V(B_{i,F})|$ is odd. This case is similar to Case 3 with the difference that every perfect matching of G_F must match exactly one of the two vertices a and b within $B_{i,F}$ while the other one cannot be matched within $B_{i,F}$ at the same time. To model this with our matchgate $J_{i,F}$ we define its vertex set to be $\{a, b, u\}$, where u is a newly introduced vertex. The edge set is defined to be $E(J_{i,F}) := \{au, bu\}$. For each $x \in \{a, b\}$ let \mathfrak{p}_x be the labelled generating function for all perfect matchings of $B_{i,F} - x$ and let H'_F be the graph obtained from H_F by replacing $B_{i,F}$ with $J_{i,F}$. In this case the *representativeness* of $J_{i,F}$ is the fact that every perfect matching of H'_F must contain exactly one of the edges of $J_{i,F}$.

Case 5: $Y_{i,F} = \{a, b, c\}$, $|V(B_{i,F})|$ is even, and $B_{i,F}$ has a perfect matching. For each $S \subseteq \{a, b, c\}$, let \mathfrak{p}_S be the labelled generating function of all perfect matchings of $B_{i,F} - S$. Since $|V(B_{i,F})|$ is even, any matching F' with $F' \in \text{Aligned}(B_{i,F}, Y_{i,F})$ must be of odd size and thus $\mathfrak{p}_S \neq 0$ is possible if and only if $|S|$ is even. In particular, we have that $\mathfrak{p}_\emptyset \neq 0$ since $B_{i,F}$ has a perfect matching. We introduce three new vertices u, v , and w and define $J_{i,F}$ and $\lambda_{i,F}$ as depicted in **Case 5** of [Figure 19](#). Now let G'_F be the graph obtained from G_F by replacing $B_{i,F}$ with $J_{i,F}$. In this case the *representativeness* of $J_{i,F}$ is slightly more complicated than in previous cases. The new vertices u, v , and w must be covered by every perfect matching of H'_F . Hence, every perfect matching M of G'_F contains exactly two edges of $J_{i,F}$ or three. This means there always exists a set $S \in \{\emptyset, \{a, b\}, \{a, c\}, \{b, c\}\}$ such that M contains a perfect matching of $J_{i,F} - S$. Moreover, in case $S = \emptyset$, M must contain the edges au, bv , and cw whose labels under $\lambda_{i,F}$, in total, multiply to \mathfrak{p}_\emptyset . If $S = \{a, b\}$ then M contains uv and cw , evaluating to \mathfrak{p}_{ab} , similarly $S = \{a, c\}$ forces the labels of the edges of M within $J_{i,F}$ to multiply to \mathfrak{p}_{ac} . Finally, M contains the edge vw if and only if M contains the edge au and thus, since $\mathfrak{p}_\emptyset \neq 0$, we obtain $\lambda_{i,F}(au) \cdot \lambda_{i,F}(vw) = \mathfrak{p}_{bc}$.

Case 6: $Y_{i,F} = \{a, b, c\}$, $|V(B_{i,F})|$ is even, and $B_{i,F}$ has no perfect matching. As for the above case, or each $S \subseteq \{a, b, c\}$ let \mathfrak{p}_S be the labelled generating function of all perfect matchings of $B_{i,F} - S$. Since $|V(B_{i,F})|$ is even, any matching F' with $F' \in \text{Aligned}(B_{i,F}, Y_{i,F})$ must be of odd size and thus $\mathfrak{p}_S \neq 0$ is possible if and only if $|S|$ is even. Moreover, since $B_{i,F}$ has no perfect matching we know $\mathfrak{p}_\emptyset = 0$. We introduce a single new vertex w and define $J_{i,F}$ and $\lambda_{i,F}$ as depicted in **Case 6** of [Figure 19](#). Now let G'_F be the graph obtained from G_F by replacing $B_{i,F}$ with $J_{i,F}$. The *representativeness* of this case is the fact that the new vertex w must be covered by every perfect matching of G'_F . Hence, every such perfect matching must contain exactly one of the edges of $J_{i,F}$. Since for every $x \in \{a, b, c\}$ we have $\lambda_{i,F}(wx) = \mathfrak{p}_{\{a,b,c\} \setminus \{x\}}$, the perfect matchings of $B_{i,F} - (\{a, b, c\} \setminus \{x\})$ are correctly represented by $(J_{i,F}, \lambda_{i,F})$.

Case 7: $Y_{i,F} = \{a, b, c\}$, $|V(B_{i,F})|$ is odd, and $B_{i,F} - a$ has a perfect matching. In this case for every $S \subseteq \{a, b, c\}$, let \mathfrak{p}_S be the labelled generating function of all perfect matchings of $B_{i,F} - S$. Since $|V(B_{i,F})|$ is odd, any matching F' with $F' \in \text{Aligned}(B_{i,F}, Y_{i,F})$ must be of even size and thus $\mathfrak{p}_S \neq 0$ is possible if and only if $|S|$ is odd. In particular, we have $\mathfrak{p}_a \neq 0$ since $B_{i,F} - a$ has a perfect matching. We introduce two new vertices, v and w , and define $J_{i,F}$ and $\lambda_{i,F}$ as depicted in **Case 7** of [Figure 19](#). Next let G'_F be the graph obtained from G_F by replacing $B_{i,F}$ with $J_{i,F}$. Observe that every perfect matching of G'_F must cover the

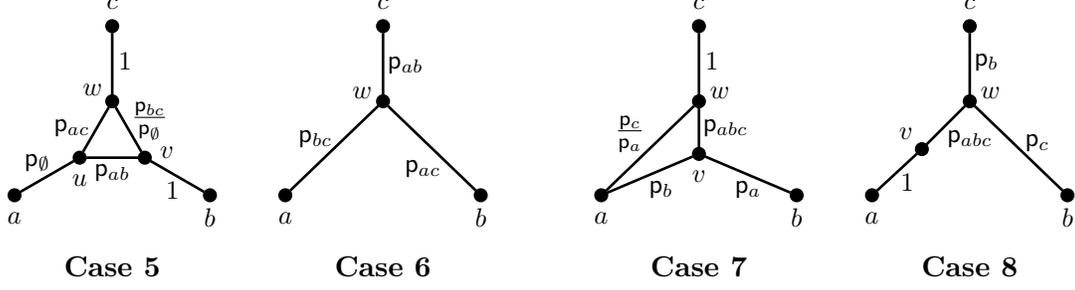


Figure 19: We assume that (G, X, \mathbf{p}) is some labelled boundary graph and, moreover, let $\mathcal{B} = ((G, X), A, \mathbf{p}, (B_1, Y_1), \dots, (B_\ell, Y_\ell))$ be a k -branching of (G, X, \mathbf{p}) . Let F be an extendable matching such that every vertex of $A \setminus X$ is covered by an edge of F , every edge in F has an endpoint in $X \cup A$, and $G - X - V(F)$ has a perfect matching. Let $i \in [\ell]$. The figure shows the four matchgates for the case where $Y_i \setminus (V(F) \cup X) = \{a, b, c\}$ for some branch (B_i, Y_i, \mathbf{p}) of \mathcal{B} . In case $B_i - (V(F) \cup X)$ has an even number of vertices (**Cases 5 and 6**) and an odd number of vertices (**Cases 7 and 8**). If $|B_i - (V(F) \cup X)|$ is even we use **Case 5** except in the situation where $\mathbf{p}_\emptyset = 0$, then we use **Case 6** instead. Similarly, in case $|B_i - (V(F) \cup X)|$ is odd, we use **Case 7** except if $\mathbf{p}_a = 0$, then we use **Case 8** instead.

vertices v and w , for this the two cases are possible; Let M be a perfect matching of G'_F , then either $vw \in M$, or there exist $x, y \in \{a, b, c\}$ such that $vx, wy \in M$. In the first case all three vertices of $\{a, b, c\}$ are matched by M with vertices outside of $J_{i,F}$, while in the second case only the single remaining vertex of $\{a, b, c\} \setminus \{x, y\}$ is matched to a vertex outside of $J_{i,F}$. If M contains the edge aw , then M must match v within $J_{i,F}$ and the only way to do so is via the edge bv . Hence, we obtain $\lambda_{i,F}(aw) \cdot \lambda_{i,F}(bv) = \mathbf{p}_c$, correctly representing the case where exactly c is matched by M with a vertex not in $J_{i,F}$. In case $av \in M$ we must have $cw \in M$ and $\lambda_{i,F}(av) \cdot \lambda_{i,F}(cw) = \mathbf{p}_b$, and if $bv, cw \in M$ we have $\lambda_{i,F}(bv) \cdot \lambda_{i,F}(cw) = \mathbf{p}_a$. This leaves only the case where $vw \in M$ which means all three vertices a, b , and c must be matched outside of $J_{i,F}$ and thus the contribution of $M \cap E(J_{i,F})$ in this case is exactly $\lambda_{i,F}(vw) = \mathbf{p}_{abc}$. We refer to these observations as the *representativeness* of $(J_{i,F}, \lambda_{i,F})$ in this case.

Case 8: $Y_{i,F} = \{a, b, c\}$, $|V(B_{i,F})|$ is odd, and $B_{i,F} - a$ has no perfect matching. As in the previous cases, for each $S \subseteq \{a, b, c\}$ let \mathbf{p}_S be the labelled generating function of all perfect matchings of $B_{i,F} - S$. Since $|V(B_{i,F})|$ is odd, any matching F' with $F' \in \text{Aligned}(B_{i,F}, Y_{i,F})$ must be of even size and thus $\mathbf{p}_S \neq 0$ is possible if and only if $|S|$ is odd. Moreover, we have $\mathbf{p}_a = 0$ since $B_{i,F} - a$ has no perfect matching. We introduce two new vertices, v and w , and define $J_{i,F}$ and $\lambda_{i,F}$ as depicted in **Case 8** of **Figure 19**. Let G'_F be the graph obtained from G_F by replacing $B_{i,F}$ with $J_{i,F}$. We complete the introduction of the matchgates by discussing the *representativeness* of $(J_{i,F}, \lambda_{i,F})$ in this case. Notice that for every perfect matching M of G'_F , the vertex v must either be matched with a or with w . In case $vw \in M$ no other edge of $J_{i,F}$ can be contained in M and thus all three vertices, a, b , and c must be matched outside of $J_{i,F}$. With $\lambda_{i,F}(vw) = \mathbf{p}_{abc}$ this is correctly represented. Hence, we may assume $av \in M$. This means we must either have bw or cw in M . Since $\lambda_{i,F}(a) = 1$, the total contribution of the two edges of $J_{i,F}$ in M is exactly as intended in both cases.

We can now employ the matchgates to produce the labelled generating functions for a boundary graph with a k -branching, given the corresponding sign post.

Lemma 62. *Let $k \in \mathbb{N}$ be some integer. Let (G, X, \mathbf{p}) be a labelled boundary graph with an additional set $A \subseteq V(G)$ such that $|X|, |A| \leq k$. If we are given a k -branching $\mathcal{B} = ((G, X), A, \mathbf{p}, (B_1, Y_1), \dots, (B_\ell, Y_\ell))$ together with the sign post sp of \mathcal{B} , then the set $\text{Gen}_{(G, X, \mathbf{p})}$ can be computed in time $|G|^{\mathcal{O}(k)}$.*

Proof. Note that $(G, A \cup X, \mathbf{p})$ is also a labelled boundary graph. Let \mathcal{F} be the collection of all matchings F in G such that $F \in \text{Aligned}(G, A \cup X)$ and $A \subseteq V(F)$. Since $|A \cup X| = \mathcal{O}(k)$ we have $|\mathcal{F}| \in |G|^{\mathcal{O}(k)}$ and \mathcal{F} can be found in time $|G|^{\mathcal{O}(k)}$. Similarly, we can find the set $\text{Aligned}(G, X)$. Now every set $F \in \mathcal{F}$ can be covered by $\ell + 2$ sets F_i as follows.

- i) Let F_1 be the set of all edges in F with at least one endpoint in X ,
- ii) let F_2 be the set of all edges in F with at most one endpoint in $\bigcup_{j \in [\ell]} Y_j$, and
- iii) for every $i \in [3, \ell + 2]$ let F_i be the set of all edges in F with both endpoints in B_{i-2} .

It follows that for every F we have $F_1 \in \text{Aligned}(G, X)$ and $F_2 = F \setminus (\bigcup_{i \in [3, \ell + 2]} F_i)$.

We now fix some $F \in \mathcal{F}$. Let also $i \in [\ell]$. Then $|Y_i \setminus V(F)| \leq 3$, but it is not necessary equal to zero. Hence, there may be $\mathcal{O}(|G|^3)$ matchings W such that $F_{i+2} \subseteq W$ and $W \in \text{Aligned}(B_i, Y_i)$. For every $i \in [\ell]$ let $\mathcal{W}_{i,F}$ be the collection of all matchings W such that $F_{i+2} \subseteq W$ and $W \in \text{Aligned}(B_i, Y_i)$. Moreover, for every $S \subseteq Y_i \setminus V(F)$ we denote by $\mathcal{W}_{i,F,S}$ the subset of $\mathcal{W}_{i,F}$ such that for all $W \in \mathcal{W}_{i,F,S}$ we have $Y_i \setminus V(W) = S$ and $V(W) \cap X = V(F) \cap X$. This means that the members of $\mathcal{W}_{i,F,S}$ are exactly the matchings of the members of $\text{Aligned}(B_i, Y_i)$ that contain F_{i+2} and expose the set $S \cup (X \setminus V(F))$. We define

$$\mathfrak{p}_S^{i,F} := \sum_{W \in \mathcal{W}_{i,F,S}} \left(\prod_{\substack{e \in F_{i+2} \\ \mathbf{p}(e) \neq 0}} \mathbf{p}(e) \right)^{-1} \cdot \text{Gen}_{(B_i, Y_i, \mathbf{p})}(W).$$

Note that, since we are given the sign post of \mathbf{B} and $|Y_i \setminus V(F)| \leq 3$, we compute all $\mathfrak{p}_S^{i,F}$ for all i and S in time $|VG|^{\mathcal{O}(k)}$. Moreover, $\mathfrak{p}_S^{i,F}$ is exactly the labelled generating function of all perfect matchings in the labelled graph $(B_i - V(F) - X - S, \mathbf{p})$. Note that, for fixed $i \in [\ell]$ and $F \in \mathcal{F}$, the $\mathfrak{p}_S^{i,F}$ can be used to replace the functions \mathfrak{p}_S used in the construction of the matchgates and we maintain their representativeness. We will make use of this observation to replace each (B_i, Y_i) by some matchgate in the graph $G - V(F) - X$ in order to produce a labelled graph of bounded Euler-genus.

The next steps of this proof are as follows: We first formally describe the construction and discuss its validity. Then we use [Proposition 59](#) to produce the labelled generating function of the resulting graph with its adjusted labeling. Finally, we show that the resulting labelled generating functions can be used to correctly produce $\text{Gen}_{(G, X, \mathbf{p})}(R)$ for all $R \in \text{Aligned}(G, X)$.

For some fixed $F \in \mathcal{F}$, recall according to the definition of the F -reduced branching $\mathcal{B}_F := ((G_F, \emptyset), \emptyset, \mathbf{p}, (B_{1,F}, Y_{1,F}), \dots, (B_{\ell,F}, Y_{\ell,F}))$. Where

$$\begin{aligned} G_F &:= G - V(F) - X, \\ B_{i,F} &:= B_i - V(F) - X \text{ for all } i \in [\ell], \text{ and} \\ Y_{i,F} &:= Y_i \setminus (V(F) \cup X) \text{ for all } i \in [\ell]. \end{aligned}$$

Observe that, with $A \subseteq V(F)$ and because of the definition of k -branchings, the graph G_F admits a drawing γ_F without crossings in a surface Σ (without boundary) of Euler-genus at most k , and for every $i \in [\ell]$, the vertices of $Y_{i,F}$ lie on a common face of γ_F . Now let G'_F be obtained from G_F by replacing, for every $i \in [\ell]$, the subgraph $G_F[Y_{i,F}]$ with the appropriate matchgate while using the functions $\mathbf{p}_S^{i,F}$ for the functions necessarily to produce the labellings of the matchgates. Let \mathbf{p}' be the resulting labeling of G'_F . Note that, since none of the $Y_{i,F}$ is contained in some other $Y_{j,F}$ by definition, and the fact that the matchgates only introduce new vertices and never delete the vertices of $Y_{i,F}$, the different $Y_{i,F}$'s do not interfere with the introduction of matchgates for other $Y_{j,F}$'s and thus G'_F is well-defined. Moreover, as the vertices of $Y_{i,F}$ lie on a common face of γ_F and all matchgates are planar with the vertices of $Y_{i,F}$ lying on the outer face, G'_F also admits a drawing in Σ without crossings, thus has Euler-genus at most k .

Let $\text{PerfMatch}(G'_F, \mathbf{p}')$ be the labelled generating function of the labelled graph (G'_F, \mathbf{p}') . We claim that $\text{PerfMatch}(G'_F, \mathbf{p}') = \text{PerfMatch}(G_F, \mathbf{p})$. If this holds, we may use [Proposition 59](#) on (G'_F, \mathbf{p}) to produce $\text{PerfMatch}(G_F, \mathbf{p})$ in time $\mathcal{O}_k(|G|^{\mathcal{O}(1)})$. For each $i \in [\ell]$ and every $M \in \mathcal{M}(G'_F)$, let us denote by $S_{i,F,M}$ the set of all vertices of $Y_{i,F}$ which are **not** covered by an edge of the matchgate $J_{i,F}$. Recall that, by the representativeness of our matchgates, we have

$$\prod_{e \in M \cap E(J_{i,F})} \mathbf{p}'(e) = \mathbf{p}_{S_{i,F,M}}^{i,F}.$$

Finally, let us denote by M^- the set $M \setminus (\cup_{i \in [\ell]} E(J_{i,F}))$. Then

$$\begin{aligned} \sum_{M \in \mathcal{M}(G'_F)} \mathbf{p}'(M) &= \sum_{M \in \mathcal{M}(G'_F)} \prod_{e \in M} \mathbf{p}'(e) \\ &= \sum_{M \in \mathcal{M}(G'_F)} \left(\prod_{e \in M^-} \mathbf{p}'(e) \right) \left(\prod_{i \in [\ell]} \prod_{e \in M \cap E(J_{i,F})} \mathbf{p}'(e) \right) \\ &= \sum_{M \in \mathcal{M}(G'_F)} \left(\prod_{e \in M^-} \mathbf{p}(e) \right) \left(\prod_{i \in [\ell]} \prod_{e \in M \cap E(J_{i,F})} \mathbf{p}'(e) \right) \\ &= \sum_{M \in \mathcal{M}(G'_F)} \left(\prod_{e \in M^-} \mathbf{p}(e) \right) \left(\prod_{i \in [\ell]} \mathbf{p}_{S_{i,F,M}}^{i,F} \right) \\ &= \sum_{M \in \mathcal{M}(G'_F)} \left(\prod_{e \in M^-} \mathbf{p}(e) \right) \left(\prod_{i \in [\ell]} \text{PerfMatch}(B_{i,F} - S_{i,F,M}, \mathbf{p}) \right) \\ &= \sum_{M \in \mathcal{M}(G'_F)} \left(\prod_{e \in M^-} \mathbf{p}(e) \right) \left(\prod_{i \in [\ell]} \sum_{N \in \mathcal{M}(B_{i,F} - S_{i,F,M})} \mathbf{p}(N) \right) \end{aligned}$$

$$\begin{aligned}
&= \sum_{M \in \mathcal{M}(G'_F)} \left(\prod_{e \in M^-} \mathbf{p}(e) \right) \left(\prod_{i \in [\ell]} \sum_{N \in \mathcal{M}(B_{i,F} - S_{i,F,M})} \prod_{e \in N} \mathbf{p}(e) \right) \\
&= \sum_{M \in \mathcal{M}(G'_F)} \left(\prod_{e \in M^-} \mathbf{p}(e) \right) \left(\sum_{N \in \mathcal{M}(\bigcup_{i \in [\ell]} (B_{i,F} - S_{i,F,M}))} \prod_{e \in N} \mathbf{p}(e) \right) \\
&= \sum_{M \in \mathcal{M}(G'_F)} \sum_{N \in \mathcal{M}(\bigcup_{i \in [\ell]} (B_{i,F} - S_{i,F,M}))} \left(\prod_{e \in M^-} \mathbf{p}(e) \right) \left(\prod_{e \in N} \mathbf{p}(e) \right) \\
&= \sum_{M \in \mathcal{M}(G_F)} \prod_{e \in M} \mathbf{p}(e) \\
&= \text{PerfMatch}(G_F, \mathbf{p})
\end{aligned}$$

Now fix some $R \in \text{Aligned}(G, X)$ and let $\mathcal{F}_R \subseteq \mathcal{F}$ be the set of all $F \in \mathcal{F}$ with $R = F_1$. It follows that

$$\begin{aligned}
\text{Gen}_{(G, X, \mathbf{p})}(R) &= \left(\prod_{e \in R} \mathbf{p}(e) \right) \cdot \text{PerfMatch}(G - V(R) - X, \mathbf{p}) \\
&= \left(\prod_{e \in R} \mathbf{p}(e) \right) \cdot \sum_{F \in \mathcal{F}_R} \left(\prod_{e \in F \setminus R} \mathbf{p}(e) \right) \cdot \text{PerfMatch}(G - V(F) - X, \mathbf{p}) \\
&= \sum_{F \in \mathcal{F}_R} \left(\prod_{e \in F} \mathbf{p}(e) \right) \cdot \text{PerfMatch}(G - V(F) - X, \mathbf{p}).
\end{aligned}$$

Since $|\mathcal{F}_R| \in |G|^{\mathcal{O}(k)}$ and, as we have seen in the discussion above, $\text{PerfMatch}(G - V(F) - X, \mathbf{p})$ can be computed in time $\mathcal{O}_k(|G|^{\mathcal{O}(1)})$ using [Proposition 59](#) on the labelled graph (G'_F, \mathbf{p}') , $\text{Gen}_{(G, X, \mathbf{p})}(R)$ can be computed in time $|G|^{\mathcal{O}(k)}$. This, together with the bounded size of $\text{Aligned}(G, X)$ completes our proof. \square

Merging bags. The remaining piece for our algorithm is to merge the tables of subtrees that attach to a common boundary. For the sake of simplicity, in all algorithms above we have always assumed that the boundaries of boundary subgraph are always distinct and never nested. As the merging of subtrees can be performed iteratively on two subtrees, it suffices to consider the case where we have two boundary graphs with nested boundaries. Towards this, we will prove the following stronger version that contains the nested assumption as a special case. While being more general, it enjoys some notational symmetry on the role of the sets X_1 and X_2 , which makes more easy the presentation of the proof.

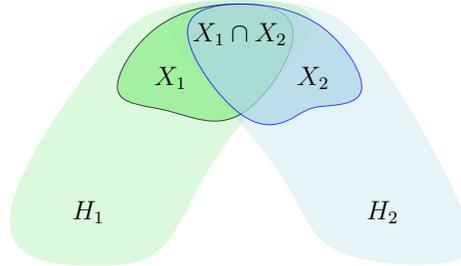


Figure 20: A visualization of [Lemma 63](#).

Lemma 63. *Let $k \in \mathbb{N}$ be some integer. Let (H_1, X_1, \mathbf{p}) and (H_2, X_2, \mathbf{p}) be two labelled boundary graphs with $V(H_1) \cap V(H_2) \subseteq X_1 \cap X_2$ and let $X := X_1 \cap X_2$ where $|X_1|, |X_2| \leq k$. Suppose we are given $\text{Gen}_{(H_1, X_1, \mathbf{p})}$ and $\text{Gen}_{(H_2, X_2, \mathbf{p})}$. Then $\text{Gen}_{(H_1 \cup H_2, X, \mathbf{p})}$ can be computed in time $|V(H_1 \cup H_2)|^{\mathcal{O}(k)}$.*

Proof. Similarly to the arguments before, we start by computing the set $\text{Aligned}(H_1 \cup H_2, X)$ and extract the set \mathcal{F} of all matchings F with $F \in \text{Aligned}(H_1 \cup H_2, X)$ in time $|V(H_1 \cup H_2)|^{\mathcal{O}(k)}$. Now fix some $F \in \mathcal{F}$ and observe that for every $i \in [2]$, there might be several extendable pairs $F_i \in \text{Aligned}(H_i, X_i)$ such that $F \cap E(H_i) \subseteq F_i$ and $X \setminus V(F) \subseteq X_i \setminus V(F_i \cup F)$. Moreover, note that for every such F_i we have $V(F_i) \setminus V(F) \subseteq V(H_i) \setminus V(H_{3-i})$ by the definition of X . Finally, any extendable pair $F_i \in \text{Aligned}(H_i, F_i)$ such that $X \setminus V(F) \subsetneq X_i \setminus V(F_i \cup F)$ will not make any contribution to $\text{PerfMatch}(H_1 \cup H_2, X, \mathbf{p}, F)$, since $\text{PerfMatch}(H_i, X_i, \mathbf{p}, F_i)$ counts all perfect matchings of $H_i - (X_i \setminus V(F_i))$ and therefore includes matchings that expose vertices of $X_i \setminus X$. (See [Figure 20](#) for a visualization of (H_1, X_1, \mathbf{p}) and (H_2, X_2, \mathbf{p}) .)

Thus, for each $i \in [2]$ and every $F \in \mathcal{F}$ let $\mathcal{F}_{i,F}$ be the collection of all F_i such that $F_i \in \text{Aligned}(H_i, X_i)$, $F \cap E(H_i) \subseteq F_i$, and $X_i \setminus V(F_i \cup F) = X \setminus V(F)$. Since $|X_i| \leq k$ we can be sure that $|\mathcal{F}_{i,F}| \in |V(H_i)|^{\mathcal{O}(k)}$. We define

$$\mathbf{p}_{i,F} := \sum_{F_i \in \mathcal{F}_{i,F}} \left(\prod_{e \in F \cap E(H_i)} \mathbf{p}(e) \right)^{-1} \cdot \text{Gen}_{H_i, X_i, \mathbf{p}}(F_i).$$

Hence, $\mathbf{p}_{i,F}$ is the labelled generating function for the perfect matchings of $(H_i - (V(F \cap E(H_i)) \cup X), \mathbf{p})$. From this observation we obtain that

$$\text{PerfMatch}(H_1 \cup H_2, X, \mathbf{p}, F) = \mathbf{p}_{1,F} \cdot \mathbf{p}_{2,F} \cdot \prod_{e \in F} \mathbf{p}(e),$$

which can be computed in time $|V(H_1) \cup V(H_2)|^{\mathcal{O}(k)}$ from $\text{Gen}_{(H_1, X_1, \mathbf{p})}$ and $\text{Gen}_{(H_2, X_2, \mathbf{p})}$ and thus our claim follows. \square

Proof of [Theorem 8](#). We are given a weighted graph (G, \mathbf{w}) as input. In the following we will be working with the labelled graph $(G, \mathbf{p}_{\mathbf{w}})$ instead.

Let α and γ be defined as in [Theorem 44](#). Then by [Corollary 45](#) we can find a tree decomposition (T, β) , where we see T as rooted on a root $r \in V(T)$. Also, adhesion at most α such that for every $d \in V(T)$, the torso G_d of G at d has a set $A_d \subseteq V(G_d)$ of size at most 4α for which the graph $G_d - A_d$ has Euler-genus at most γ in time $\mathcal{O}_t(|G|^3)$. Moreover, for every $(d_1, d_2) \in E(T)$ we have $|(\beta(d_1) \setminus A_{d_1}) \cap (\beta(d_2) \setminus A_{d_2})| \leq 3$, and if $|(\beta(d_1) \setminus A_{d_1}) \cap (\beta(d_2) \setminus A_{d_2})| = 3$ and $\beta(d_1)$ is larger than 4α , then $(\beta(d_1) \setminus A_{d_1}) \cap (\beta(d_2) \setminus A_{d_2})$ induces a triangle in the corresponding drawing of $G_{d_1} - A_{d_1}$ which bounds a face. Observe that $|V(T)| \in \mathcal{O}(|G|)$.

Set $k = \max\{4\alpha, \gamma\}$ and let $t \in V(T)$ be an arbitrary leaf. If $|\beta(t)| \leq 4\alpha$ we may use [Lemma 53](#) to initialize the tables at t , otherwise we use [Lemma 60](#). Now let $t \in V(T)$ be some internal vertex such that the tables for the subtrees rooted at the children of t have already been computed. We proceed by iteratively applying [Lemma 53](#) and [Lemma 63](#) as long as we deal with torsos of size at most 4α . Then we find a k -branching $\mathcal{B} = ((H, X), A_t, \mathbf{p}_{\mathbf{w}}, (B_1, Y_1), \dots, (B_\ell, Y_\ell))$, where the (B_i, Y_i) represent the graphs induced by the union of the bags in the subtrees rooted at the children of t which were merged into the adhesion set $Y_i \subseteq \beta(t)$, and (H, X) is the boundary graph where H is the subgraph of G induced by the union of all bags in the subtree of T rooted at t and X is

the adhesion of $\beta(t)$ to the bag of the parent node of t . In case $t = r$ we have $X = \emptyset$. Using the procedure from [Lemma 62](#) allows us to correctly compute the table for \mathcal{B} .

Hence, by induction, we are able to compute $\text{Gen}_{(G, \emptyset, \mathbf{p}_w)}$ in time $|G|^{\mathcal{O}(\alpha)}$. Notice that the only member of $\text{Gen}_{(G, \emptyset, \mathbf{p}_w)}$ is exactly $\text{PerfMatch}(G, \mathbf{w})$ and our proof is complete. \square

4.2 The complexity lower bound

In the previous section we have seen how to compute the generating function for all perfect matchings efficiently on a graph that excludes some shallow vortex minor. Once the generating function is known, the number of perfect matchings can be found by evaluating the function for $x = 1$. Hence, if \mathcal{G} is a proper minor-closed class of graphs excluding some shallow vortex minor, computing the number of perfect matchings is tractable on \mathcal{G} . Towards completing the proof of our main result (that is [Theorem 2](#)), it remains to prove [Theorem 9](#), i.e., show that on any proper minor-closed class of graphs which contains all shallow vortex minors, the problem of counting perfect matchings remains $\#P$ -hard.

To show this we make use of a recent complexity result from [\[11\]](#) which was used by Curticapean and Xia to prove that counting perfect matchings is $\#P$ -hard on K_8 -minor free graphs.

Proposition 64 ([\[11\]](#)). *Counting perfect matchings is $\#P$ -hard on the class of ring blowup graphs.*

The next observation follows by straightforward graph drawing arguments (see e.g., [\[2, Lemma 5.5\]](#)). Recall that we defined a *standard cross-free drawing on a disk* Δ of the $(t \times s)$ -cylindrical grid one where of the extremal cycles is drawn on the boundary of Δ

Lemma 65. *There exists a function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that the following holds: Suppose that γ is the embedding of some 2-connected planar graph G on n vertices in a closed disk Δ and let C be the cycle of G whose drawing is the boundary of Δ . Let also γ' be a standard cross-free drawing on a disk of the $(f(n) \times f(n))$ -cylindrical grid G' in a closed disk Δ' and let C' be the cycle of G' defined by the boundary of the external face of γ' . Then there is a minor model $\{X_v \mid v \in V(H)\}$ of G to G' such that for every $v \in V(C)$, X_v is a subpath of C' . Moreover, f is a quadratic function.*

It is interesting to observe that, according to [Lemma 47](#) and [Lemma 65](#), the class \mathcal{S} of shallow vortex minors is *exactly* the class of graphs that are minors of ring blowups.

We now have what we need for the proof of [Theorem 9](#).

Proof of Theorem 9. From [Proposition 64](#), it is enough to prove that the class of all ring blowup graphs is a subset of \mathcal{S} . In the special case where G is not a 2-connected graph, we may create a new graph G^+ by adding in G a minimum number of edges that can make it 2-connected, while maintaining its planarity and the property of having a cross-free drawing γ^+ on a disk where Q is the set of vertices incident in its external face. Notice that the ring blowup of γ is a subgraph of the ring blowup of γ^+ . This implies that we may assume that G is a 2-connected planar graph and prove that the ring blowup \hat{G} of γ of G is a minor of some, big enough, shallow vortex grid. The 2-connectivity of G permits us to assume that the vertices that are incident with the external face of γ are the vertices of some cycle C of G . By [Lemma 65](#), there is a $(f(n) \times f(n))$ -cylindrical grid G' whose standard cross-free drawing on a disk is γ' and, given that C' is the cycle defined by the boundary of the external face of γ' , there exists some minor model $\{X_v \mid v \in V(H)\}$ of G to G'

such that for every $v \in V(C)$, X_v is a subpath of C' . It is now easy to observe that, because of this last property, \hat{G} is a minor of the ring blowup \hat{G}' of γ' . The result follows as \hat{G}' is a $(f(n) \times f(n))$ -cylindrical grid ring blowup, which from [Lemma 47](#) is the minor of $\mathcal{S}_{g(f(n))} \in \mathcal{S}$ (where g is the function of [Lemma 47](#)). \square

5 Conclusion

Notice that, by definition, \mathcal{S} is a minor-closed graph class. Let $\mathcal{Q} = \text{obs}(\mathcal{S})$ be its minor-obstruction set, that is the set of all minor-minimal graphs not contained in \mathcal{S} . We know, from Robertson and Seymour’s theorem that \mathcal{Q} is a finite set [\[46\]](#). This set provides a “finite” version of the characterization in [Theorem 2](#) as follows.

Corollary 66. *Let \mathcal{F} be a finite set of graphs. $\#\text{PERFECT MATCHING}(\mathcal{F})$ is polynomial-time solvable if \mathcal{F} contains some \mathcal{Q} -minor free graph; otherwise it is $\#\text{P}$ -complete.*

The exact identification of the finite set \mathcal{Q} seems to be an interesting, however hard, combinatorial problem.

As already mentioned by Curticapean and Xia in [\[10\]](#), it is interesting to investigate the complexity dichotomies of other families of counting problems in the realm of minor-closed graph classes. Such a framework is the one of *Holant problems*: $\#\text{PERFECT MATCHING}$ is a typical example of this family of problems [\[4, 5\]](#).

Another direction on the study of $\#\text{PERFECT MATCHING}$ is to look for classes of bipartite graphs, ordered by the *matching-minor* relation. The most general result in this direction is the one of McCuaig, Robertson, Seymour, and Thomas [\[34, 38\]](#), implying that such an algorithm exists for Pfaffian bipartite graphs, that, according to Little [\[30\]](#) are exactly the graphs excluding $K_{3,3}$ as a *matching minor*. Moreover, the number of perfect matchings can be found efficiently on bipartite graphs excluding a planar bipartite matching minor [\[19\]](#). This induces an alternative line of research, asking for more general matching-minor-closed bipartite graph classes where $\#\text{PERFECT MATCHING}$ is polynomially solvable. This line of research may be of particular importance as the complexity of the permanent is directly linked to the complexity of counting perfect matchings on bipartite graphs.

Acknowledgements. The first author wishes to thank [Ioannis Mourtos](#) for old discussions that offered some early inspiration on this project.

We wish to thank the anonymous referees for helping to greatly improve the presentation of the paper.

References

- [1] Manindra Agrawal. Determinant versus permanent. In *In Proceedings of the 25th International Congress of Mathematicians, ICM 2006*, volume 3, page 985–997, 2006. [3](#)

- [2] Julien Baste, Ignasi Sau, and Dimitrios M. Thilikos. Hitting minors on bounded treewidth graphs. IV. an optimal algorithm. *SIAM J. Comput.*, 52(4):865–912, 2023. doi:10.1137/21m140482x. 60
- [3] Jacek Blazewicz, Piotr Formanowicz, Marta Kasprzak, Petra Schuurman, and Gerhard J. Woeginger. A polynomial time equivalence between DNA sequencing and the exact perfect matching problem. *Discret. Optim.*, 4(2):154–162, 2007. doi:10.1016/j.disopt.2006.07.004. 5
- [4] Jin-yi Cai, Pinyan Lu, and Mingji Xia. Holant problems and counting CSP. In Michael Mitzenmacher, editor, *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 715–724. ACM, 2009. doi:10.1145/1536414.1536511. 61
- [5] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Holographic algorithms. In *Encyclopedia of Algorithms*, pages 921–926. Springer, Berlin, 2016. doi:10.1007/978-1-4939-2864-4_746. 61
- [6] Arthur Cayley. Sur les Déterminants Gauches. In *Journal für die reine und angewandte Mathematik Band 38*, pages 93–96. De Gruyter, 1847. 49
- [7] Radu Curticapean. Counting perfect matchings in graphs that exclude a single-crossing minor. *CoRR*, abs/1406.4056, 2014. URL: <http://arxiv.org/abs/1406.4056>, arXiv:1406.4056. 4, 5, 9, 10, 52
- [8] Radu Curticapean. Counting problems in parameterized complexity. In *13th International Symposium on Parameterized and Exact Computation (IPEC 2018)*, volume 115 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 1:1–1:18, Dagstuhl, Germany, 2019. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. URL: <http://drops.dagstuhl.de/opus/volltexte/2019/10202>, doi:10.4230/LIPIcs.IPEC.2018.1. 9
- [9] Radu Curticapean and Mingji Xia. Parameterizing the permanent: genus, apices, minors, evaluation mod $2k$. In Venkatesan Guruswami, editor, *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 994–1009. IEEE Computer Society, 2015. doi:10.1109/FOCS.2015.65. 4, 9, 10, 50
- [10] Radu Curticapean and Mingji Xia. Parameterizing the permanent: hardness for K_8 -minor-free graphs. *CoRR*, abs/2108.12879, 2021. URL: <https://arxiv.org/abs/2108.12879>, arXiv:2108.12879. 4, 5, 6, 61
- [11] Radu Curticapean and Mingji Xia. Parameterizing the Permanent: Hardness for Fixed Excluded Minors. In *Symposium on Simplicity in Algorithms (SOSA)*, pages 297–307. SIAM, 2022. 3, 10, 60
- [12] Reinhard Diestel. *Graph Theory*, volume 173. Springer-Verlag, 5th edition, 2017. doi:10.1007/978-3-662-53622-3. 11
- [13] Reinhard Diestel, Ken-ichi Kawarabayashi, Theodor Müller, and Paul Wollan. On the excluded minor structure theorem for graphs of large tree-width. *J. Comb. Theory, Ser. B*, 102(6):1189–1210, 2012. 11, 19, 21, 22, 37

- [14] Walther Dyck. Beiträge zur Analysis situs I. *Math. Ann.*, 32(4):457–512, 1888. doi:10.1007/BF01443580. 12
- [15] Jack Edmonds. Paths, Trees, and Flowers. *Canad. J. of Mathematics*, 17:449–467, 1965. 45
- [16] David Eppstein and Vijay V. Vazirani. NC algorithms for computing a perfect matching, the number of perfect matchings, and a maximum flow in one-crossing-minor-free graphs. In Christian Scheideler and Petra Berenbrink, editors, *The 31st ACM on Symposium on Parallelism in Algorithms and Architectures, SPAA 2019, Phoenix, AZ, USA, June 22-24, 2019*, pages 23–30. ACM, 2019. doi:10.1145/3323165.3323206. 4, 5, 8, 9, 10
- [17] Anna Galluccio and Martin Loeb. On the theory of pfaffian orientations. I. perfect matchings and permanents. *Electron. J. Comb.*, 6, 1999. URL: http://www.combinatorics.org/Volume_6/Abstracts/v6i1r6.html. 4, 5, 9, 10, 49, 50
- [18] Anna Galluccio, Martin Loeb, and Jan Vondrák. Optimization via Enumeration: A New Algorithm for the Max Cut Problem. *Math. Programming*, 90(2):273–290, 2001. 9, 10, 50
- [19] Archontia C Giannopoulou, Stephan Kreutzer, and Sebastian Wiederrecht. Excluding a Planar Matching Minor in Bipartite Graphs. *arXiv preprint arXiv:2106.00703*, 2021. 61
- [20] Petr A. Golovach, Giannos Stamoulis, and Dimitrios M. Thilikos. Hitting topological minor models in planar graphs is fixed parameter tractable. In Shuchi Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 931–950. SIAM, 2020. doi:10.1137/1.9781611975994.56. 32
- [21] Rohit Gurjar, Arpita Korwar, Jochen Messner, and Thomas Thierauf. Exact perfect matching in complete graphs. *ACM Trans. Comput. Theory*, 9(2):8:1–8:20, 2017. doi:10.1145/3041402. 5
- [22] Frank Harary. *Graph theory*. Addison-Wesley, 1991. 42
- [23] Ken ichi Kawarabayashi, Robin Thomas, and Paul Wollan. Quickly excluding a non-planar graph, 2021. arXiv:2010.12397. 6, 8, 11, 12, 13, 15, 16, 17, 18, 19, 20, 21, 24, 35, 37, 38, 39, 40
- [24] Heinz A Jung. Eine Verallgemeinerung des n -fachen Zusammenhangs für Graphen. *Mathematische Annalen*, 187(2):95–103, 1970. 15
- [25] P. W. Kasteleyn. Dimer statistics and phase transitions. *Journal of Mathematical Physics*, 4(2):287–293, 1963. arXiv:<https://doi.org/10.1063/1.1703953>, doi:10.1063/1.1703953. 3
- [26] P.M. Kasteleyn. The statistics of dimers on a lattice: I. The number of dimer arrangements on a quadratic lattice. *Physica*, 27:1209–1225, 1961. 3, 4, 5, 10, 49
- [27] P.M. Kasteleyn. Graph theory and crystal physics. In *In Graph Theory and Theoretical Physics*, pages 43–110. Academic Press, 1967. 3, 4, 10, 49

- [28] Ken-ichi Kawarabayashi, Robin Thomas, and Paul Wollan. A new proof of the flat wall theorem. *J. Comb. Theory, Ser. B*, 129:204–238, 2018. doi:[10.1016/j.jctb.2017.09.006](https://doi.org/10.1016/j.jctb.2017.09.006). [6](#), [11](#), [13](#), [15](#), [16](#), [17](#)
- [29] Ken-ichi Kawarabayashi and Paul Wollan. A shorter proof of the graph minor algorithm: the unique linkage theorem. In Leonard J. Schulman, editor, *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 687–694. ACM, 2010. doi:[10.1145/1806689.1806784](https://doi.org/10.1145/1806689.1806784). [32](#)
- [30] Charles Little. *An Extension of kasteleyn’s method of enumerating the 1-factors of planar graphs*, volume 403, pages 63–72. Springer, Berlin, 08 2006. doi:[10.1007/BFb0057377](https://doi.org/10.1007/BFb0057377). [5](#), [61](#)
- [31] Tonglai Liu, Jigang Wu, Jiaying Li, Jingyi Li, and Zikai Zhang. Efficient algorithms for storage load balancing of outsourced data in blockchain network. *The Computer Journal*, 2021. [5](#)
- [32] C. R. F. Maunder. *Algebraic topology*. Dover Publications, Inc., Mineola, NY, 1996. Reprint of the 1980 edition. [12](#)
- [33] William McCuaig. Pólya’s permanent problem. *Electron. J. Comb.*, 11(1), 2004. URL: http://www.combinatorics.org/Volume_11/Abstracts/v11i1r79.html. [4](#)
- [34] William McCuaig, Neil Robertson, Paul D. Seymour, and Robin Thomas. Permanents, pfaffian orientations, and even directed circuits (extended abstract). In Frank Thomson Leighton and Peter W. Shor, editors, *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, El Paso, Texas, USA, May 4-6, 1997*, pages 402–405. ACM, 1997. doi:[10.1145/258533.258625](https://doi.org/10.1145/258533.258625). [4](#), [5](#), [61](#)
- [35] Christos H. Papadimitriou and Mihalis Yannakakis. The complexity of restricted spanning tree problems. *J. ACM*, 29(2):285–309, 1982. doi:[10.1145/322307.322309](https://doi.org/10.1145/322307.322309). [5](#)
- [36] Christophe Paul, Evangelos Protopapas, and Dimitrios M. Thilikos. Universal obstructions of graph parameters, 2023. [arXiv:2304.14121](https://arxiv.org/abs/2304.14121). [5](#)
- [37] György Pólya. Aufgabe 424. *Arch. Math. Phys.*, 1913. [3](#), [5](#)
- [38] Neil Robertson, P. D. Seymour, and Robin Thomas. Permanents, pfaffian orientations, and even directed circuits. *Annals of Mathematics*, 150(3):929–975, 1999. URL: <http://www.jstor.org/stable/121059>. [4](#), [5](#), [61](#)
- [39] Neil Robertson and Paul Seymour. Graph Minors. XXI. Graphs with unique linkages. *J. Comb. Theory, Ser. B*, 99(3):583–616, 2009. [32](#)
- [40] Neil Robertson and Paul D Seymour. Graph Minors: IX. Disjoint crossed paths. *J. Comb. Theory, Ser. B*, 49(1):40–77, 1990. [15](#)
- [41] Neil Robertson and Paul D. Seymour. Graph minors. x. obstructions to tree-decomposition. *J. Comb. Theory, Ser. B*, 52(2):153–190, 1991. doi:[10.1016/0095-8956\(91\)90061-N](https://doi.org/10.1016/0095-8956(91)90061-N). [7](#), [19](#), [38](#)
- [42] Neil Robertson and Paul D Seymour. Graph Minors: X. Obstructions to tree-decomposition. *J. Comb. Theory, Ser. B*, 52(2):153–190, 1991. [37](#)

- [43] Neil Robertson and Paul D. Seymour. Excluding a graph with one crossing. In *Graph structure theory (Seattle, WA, 1991)*, pages 669–675. Amer. Math. Soc., Providence, RI, 1993. [8](#), [38](#)
- [44] Neil Robertson and Paul D Seymour. Graph Minors: XIII. The disjoint paths problem. *J. Comb. Theory, Ser. B*, 63(1):65–110, 1995. [15](#), [17](#)
- [45] Neil Robertson and Paul D. Seymour. Graph minors. XVI. Excluding a non-planar graph. *J. Comb. Theory, Ser. B*, 89(1):43–76, 2003. [doi:10.1016/S0095-8956\(03\)00042-X](#). [6](#), [8](#), [19](#), [21](#)
- [46] Neil Robertson and Paul D. Seymour. Graph minors. XX. Wagner’s conjecture. *J. Comb. Theory, Ser. B*, 92(2):325–357, 2004. [doi:10.1016/j.jctb.2004.08.001](#). [4](#), [61](#)
- [47] Paul D Seymour. Disjoint Paths in Graphs. *Discr. Mathematics*, 29(3):293–309, 1980. [15](#)
- [48] Paul D Seymour and Robin Thomas. Graph searching and a min-max theorem for tree-width. *J. Comb. Theory, Ser. B*, 58(1):22–33, 1993. [34](#)
- [49] Yossi Shiloach. A polynomial solution to the undirected two paths problem. *J. of the ACM (JACM)*, 27(3):445–456, 1980. [15](#)
- [50] Simon Straub, Thomas Thierauf, and Fabian Wagner. Counting the number of perfect matchings in K_5 -free graphs. In *IEEE 29th Conference on Computational Complexity, CCC 2014, Vancouver, BC, Canada, June 11-13, 2014*, pages 66–77. IEEE Computer Society, 2014. [doi:10.1109/CCC.2014.15](#). [4](#), [5](#), [52](#)
- [51] H. N. V. Temperley and Michael E. Fisher. Dimer problem in statistical mechanics-an exact result. *The Philosophical Magazine: A Journal of Theoretical Experimental and Applied Physics*, 6(68):1061–1063, 1961. [arXiv:https://doi.org/10.1080/14786436108243366](#), [doi:10.1080/14786436108243366](#). [3](#), [4](#), [5](#), [10](#)
- [52] Glenn Tesler. Matchings in graphs on non-orientable surfaces. *J. Comb. Theory, Ser. B*, 78(2):198–231, 2000. [doi:10.1006/jctb.1999.1941](#). [4](#), [5](#), [9](#), [10](#), [50](#)
- [53] Dimitrios M. Thilikos and Sebastian Wiederrecht. Killing a vortex. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 - November 3, 2022*, pages 1069–1080. IEEE, 2022. [doi:10.1109/FOCS54457.2022.00104](#). [1](#)
- [54] Dimitrios M. Thilikos and Sebastian Wiederrecht. Approximating branchwidth on parametric extensions of planarity, 2023. [arXiv:2304.04517](#). [8](#)
- [55] Dimitrios M. Thilikos and Sebastian Wiederrecht. Excluding surfaces as minors in graphs, 2023. [arXiv:2306.01724](#). [8](#)
- [56] Carsten Thomassen. 2-linked graphs. *Europ. J. of Combinatorics.*, 1(4):371–378, 1980. [15](#)
- [57] Leslie G. Valiant. The complexity of computing the permanent. *Theor. Comput. Sci.*, 8:189–201, 1979. [doi:10.1016/0304-3975\(79\)90044-6](#). [4](#), [5](#), [49](#)
- [58] Leslie G Valiant. Holographic algorithms. *SIAM J. Comput.*, 37(5):1565–1594, 2008. [52](#)

- [59] Vijay V. Vazirani. NC algorithms for computing the number of perfect matchings in $K_{3,3}$ -free graphs and related problems. *Inf. Comput.*, 80(2):152–164, 1989. doi:[10.1016/0890-5401\(89\)90017-5](https://doi.org/10.1016/0890-5401(89)90017-5). 4, 5
- [60] Guohun Zhu, Xiangyu Luo, and Yuqing Miao. Exact weight perfect matching of bipartite graph is NP-complete. In *Proceedings of the World Congress on Engineering*, volume 2, pages 1–7, 2008. 5