

Hitting Minors on Bounded Treewidth Graphs. IV. An Optimal Algorithm

Julien Baste, Ignasi Sau, Dimitrios M. Thilikos

► To cite this version:

Julien Baste, Ignasi Sau, Dimitrios M. Thilikos. Hitting Minors on Bounded Treewidth Graphs. IV. An Optimal Algorithm. SIAM Journal on Computing, 2023, 52 (4), pp.865-912. 10.1137/21M140482X . lirmm-04171495

HAL Id: lirmm-04171495 https://hal-lirmm.ccsd.cnrs.fr/lirmm-04171495

Submitted on 26 Jul 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.

Hitting minors on bounded treewidth graphs. IV. An optimal algorithm 1

Julien Baste^{2,3}

Ignasi Sau³

Dimitrios M. Thilikos³

Abstract

For a fixed finite collection of graphs \mathcal{F} , the \mathcal{F} -M-DELETION problem asks, given an *n*-vertex input graph G, for the minimum number of vertices that intersect all minor models in G of the graphs in \mathcal{F} . by Courcelle's Theorem, this problem can be solved in time $f_{\mathcal{F}}(\mathsf{tw}) \cdot n^{\mathcal{O}(1)}$, where tw is the treewidth of G, for some function $f_{\mathcal{F}}$ depending on \mathcal{F} . In a recent series of articles, we have initiated the programme of optimizing asymptotically the function $f_{\mathcal{F}}$. Here we provide an algorithm showing that $f_{\mathcal{F}}(\mathsf{tw}) = 2^{\mathcal{O}(\mathsf{tw} \cdot \log \mathsf{tw})}$ for every collection \mathcal{F} . Prior to this work, the best known function $f_{\mathcal{F}}$ was double-exponential in tw. In particular, our algorithm vastly extends the results of Jansen et al. [SODA 2014] for the particular case $\mathcal{F} = \{K_5, K_{3,3}\}$ and of Kociumaka and Pilipczuk [Algorithmica 2019] for graphs of bounded genus, and answers an open problem posed by Cygan et al. [Inf Comput 2017]. We combine several ingredients such as the machinery of boundaried graphs in dynamic programming via representatives, the Flat Wall Theorem, Bidimensionality, the irrelevant vertex technique, treewidth modulators, and protrusion replacement. Together with our previous results providing single-exponential algorithms for particular collections F [Theor Comput Sci 2020] and general lower bounds [J Comput Syst Sci 2020], our algorithm yields the following complexity dichotomy when $\mathcal{F} = \{H\}$ contains a single connected graph H, assuming the Exponential Time Hypothesis: $f_H(\mathsf{tw}) = 2^{\Theta(\mathsf{tw})}$ if H is a contraction of the chair or the banner, and $f_H(\mathsf{tw}) = 2^{\Theta(\mathsf{tw} \cdot \log \mathsf{tw})}$ otherwise.

Keywords: parameterized complexity; graph minors; treewidth; hitting minors; Flat Wall Theorem; irrelevant vertex; dynamic programming; complexity dichotomy.

¹An extended abstract of this article appeared in the *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 951-970, Salt Lake City, Utah, U.S., January 2020.* The first author was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - 388217545. The two last authors were supported by the ANR projects DEMOGRAPH (ANR-16-CE40-0028), ESIGMA (ANR-17-CE23-0010), ELIT (ANR-20-CE48-0008), the French-German Collaboration ANR/DFG Project UTMA (ANR-20-CE92-0027), and the French Ministry of Europe and Foreign Affairs, via the Franco-Norwegian project PHC AURORA.

²Univ. Lille, CNRS, Centrale Lille, UMR 9189 - CRIStAL - Centre de Recherche en Informatique, Signal et Automatique de Lille, F-59000 Lille, France. Email: julien.baste@univ-lille.fr.

³LIRMM, Univ. Montpellier, CNRS, Montpellier, France. Emails: ignasi.sau@lirmm.fr, sedthilk@thilikos.info.

Contents

1	Introduction	1
2	Overview of the algorithm	3
3	Preliminaries	6
	 3.1 Basic definitions	6 7 8
4	Flat walls 4.1 Walls and subwalls	9 10 11 13 14 17 18
5	Finding an irrelevant vertex 5.1 A lemma for model taming 5.2 Model rerouting in partially disk-embedded graphs 5.3 Levelings and well-aligned flatness pairs 5.4 Rerouting minors of small intrusion	19 19 22 27 29
6	Bounding the size of the representatives6.1 Finding a treewidth modulator of linear size6.2 Finding a linear protrusion decomposition and reducing protrusions	36 36 38
7	Further research	41
A	Illustration of the complexity dichotomy	46
в	An estimation of the constants depending on \mathcal{F} in our algorithm B.1 An improved version of Lemma 27 B.2 Upper bounds on the constants depending on the excluded minors	46 47 48

1 Introduction

Let \mathcal{F} be a finite non-empty collection of non-empty graphs. In the \mathcal{F} -M-DELETION problem, we are given a graph G and an integer k, and the objective is to decide whether there exists a set $S \subseteq V(G)$ with $|S| \leq k$ such that $G \setminus S$ does not contain any of the graphs in \mathcal{F} as a minor. This problem belongs to the family of graph modification problems and has a big expressive power, as instantiations of it correspond, for instance, to VERTEX COVER ($\mathcal{F} = \{K_2\}$), FEEDBACK VERTEX SET ($\mathcal{F} = \{K_3\}$), and VERTEX PLANARIZATION ($\mathcal{F} = \{K_5, K_{3,3}\}$). Note that if \mathcal{F} contains a graph with at least one edge, then \mathcal{F} -M-DELETION is NPhard [37].

We study the parameterized complexity of \mathcal{F} -M-DELETION in terms of the treewidth of the input graph (while the size of k may be unbounded). Since the property of containing a graph as a minor can be expressed in monadic second-order logic [32], by Courcelle's Theorem [13], \mathcal{F} -M-DELETION can be solved in time¹ $\mathcal{O}^*(f_{\mathcal{F}}(\mathsf{tw}))$ on graphs with treewidth at most tw, where $f_{\mathcal{F}}$ is some computable function depending on \mathcal{F} . As the function $f_{\mathcal{F}}(\mathsf{tw})$ given by Courcelle's Theorem is typically enormous, our goal is to determine, for a fixed collection \mathcal{F} , what is the best possible such function $f_{\mathcal{F}}$ that one can (asymptotically) hope for, subject to reasonable complexity assumptions. Besides being an interesting objective in its own, optimizing the running time of algorithms parameterized by treewidth has usually side effects. Indeed, black-box subroutines parameterized by treewidth are nowadays ubiquitous in parameterized [14], exact [19], and approximation [50] algorithms. Note that, for the sake of the running times discussed below, we may assume that we are given a tree decomposition of the input graph whose width is within a constant-factor of its treewidth, by using for instance the recent result of Korhonen [34].

Previous work. This line of research has attracted considerable attention in the parameterized complexity community during the last years. For instance, VERTEX COVER is easily solvable in time $\mathcal{O}^*(2^{\mathcal{O}(\mathsf{tw})})$, called *single-exponential*, by standard dynamic programming techniques, and no algorithm with running time $\mathcal{O}^*(2^{\mathcal{O}(\mathsf{tw})})$ exists, unless the Exponential Time Hypothesis (ETH) fails [26]. (The ETH implies that 3-SAT on *n* variables cannot be solved in time $2^{o(n)}$; see [26] for more details.) For FEEDBACK VERTEX SET, standard dynamic programming techniques give a running time of $\mathcal{O}^*(2^{\mathcal{O}(\mathsf{tw} \cdot \log \mathsf{tw})})$, while the lower bound under the ETH [26] is again $\mathcal{O}^*(2^{o(\mathsf{tw})})$. This gap remained open for a while, until Cygan et al. [16] presented an optimal (randomized) algorithm running in time $\mathcal{O}^*(2^{\mathcal{O}(\mathsf{tw})})$, introducing the celebrated *Cut & Count* technique. This article triggered several other (deterministic) techniques to obtain single-exponential algorithms for so-called *connectivity problems* on graphs of bounded treewidth, mostly based on algebraic tools [10, 20].

Concerning VERTEX PLANARIZATION, Jansen et al. [27] presented an algorithm running in time $\mathcal{O}^*(2^{\mathcal{O}(\mathsf{tw} \cdot \log \mathsf{tw})})$ as a crucial subroutine in an algorithm running in time $\mathcal{O}^*(2^{\mathcal{O}(k \cdot \log k)})$ where k is the solution size. Marcin Pilipczuk [40] proved afterwards that this running time is optimal under the ETH, by using the framework introduced by Lokshtanov et al. [38] for proving superexponential lower bounds.

Generalizing the above algorithm, the main technical contribution of the recent paper of Kociumaka and Pilipczuk [33] is an algorithm running in time $\mathcal{O}^*(2^{\mathcal{O}((\mathsf{tw}+g)\cdot \log(\mathsf{tw}+g))})$ to solve the GENUS VERTEX DELETION problem, which consists in deleting the minimum number of vertices from an input graph in order to obtain a graph embeddable on a surface of Euler genus at most g.

Cygan et al. [15] studied the problem of hitting subgraphs (instead of minors), and proved that the problem of hitting all copies of a fixed *path* as a subgraph can be solved in time $\mathcal{O}^*(2^{\mathcal{O}(\mathsf{tw} \cdot \log \mathsf{tw})})$. As a path occurs as a subgraph if and only if it occurs as a minor, their result implies that $\{P_h\}$ -M-DELETION can be solved in time $\mathcal{O}^*(2^{\mathcal{O}(\mathsf{tw} \cdot \log \mathsf{tw})})$ for every fixed integer $h \geq 2$, where P_h is the path on h vertices. They left as an open problem whether the algorithm in time $\mathcal{O}^*(2^{\mathcal{O}(\mathsf{tw} \cdot \log \mathsf{tw})})$ of Jansen et al. [27] for VERTEX

¹The notation $\mathcal{O}^*(\cdot)$ suppresses polynomial factors depending on the size of the input graph.

PLANARIZATION could be generalized to more minor-closed graph classes, other than planar graphs.

In a recent series of three papers [6–8], we initiated a systematic study of the complexity of \mathcal{F} -M-DELETION, parameterized by treewidth². Before stating these results, we say that a collection \mathcal{F} is *connected* if it contains only connected graphs.

In [6] we showed that, for every fixed collection \mathcal{F} , \mathcal{F} -M-DELETION can be solved in time $2^{2^{\mathcal{O}(\mathsf{tw} \cdot \log \mathsf{tw})}} \cdot n$ by a natural dynamic programming algorithm, and that if \mathcal{F} contains a planar graph, the running time can be improved³ to $2^{\mathcal{O}(\mathsf{tw} \cdot \log \mathsf{tw})} \cdot n$. If the input graph G is planar or, more generally, embedded in a surface of bounded genus, then we showed that the running time can be further improved to $2^{\mathcal{O}(\mathsf{tw})} \cdot n$.

In [7] we provided single-exponential algorithms for $\{H\}$ -M-DELETION when H is either P_4 , C_4 , the claw, the paw, the chair, or the banner; see Figure 17 in Appendix A for an illustration of these graphs.

In [8] we focused on lower bounds under the ETH. We proved that for any connected \mathcal{F} containing graphs on at least two vertices, \mathcal{F} -M-DELETION cannot be solved in time $\mathcal{O}^*(2^{o(\mathsf{tw})})$, even if the input graph G is planar. More notably, we proved that \mathcal{F} -M-DELETION cannot be solved in time $\mathcal{O}^*(2^{o(\mathsf{tw} \cdot \log \mathsf{tw})})$ for collections \mathcal{F} satisfying some generic conditions. In particular, these conditions apply when \mathcal{F} contains a single connected graph H that is not a contraction of the chair or the banner. Note that the connected graphs H with $|V(H)| \ge 2$ that are a contraction of the chair or the banner are those on the left in Figure 17, and for each of them $\{H\}$ -M-DELETION can be solved in (optimal) single-exponential time [7,8].

Our results. In this article we present an algorithm to solve \mathcal{F} -M-DELETION in time $2^{\mathcal{O}(\mathsf{tw} \cdot \log \mathsf{tw})} \cdot n$ for every collection \mathcal{F} , thus making a significant step towards a complete classification of the complexity of the \mathcal{F} -M-DELETION problem parameterized by treewidth. That is, we drop the condition that \mathcal{F} contains a planar graph, which was critically needed in the algorithm presented in [6] in order to bound the treewidth of an \mathcal{F} -minor-free graph. Our algorithm can be interpreted as an exponential "collapse" of the natural dynamic programming algorithm running in time $2^{2^{\mathcal{O}(\mathsf{tw} \cdot \log \mathsf{tw})}} \cdot n$ given in [6]. Besides largely improving our previous results [6], this algorithm generalizes the ones for $\mathcal{F} = \{K_5, K_{3,3}\}$ given by Jansen et al. [27] and for the GENUS VERTEX DELETION problem given by Kociumaka and Pilipczuk [33], which are based on embeddings, and answers positively the open problem of Cygan et al. [15] for every minor-closed graph class. Since the algorithm is quite involved, we provide an overview of it in Section 2.

This algorithm in time $2^{\mathcal{O}(\text{tw} \cdot \log \text{tw})} \cdot n$ for every collection \mathcal{F} , together with the lower bounds under the ETH given in [8], the single-exponential algorithms given in [7], and the known cases $\mathcal{F} = \{P_2\}$ [14, 26], $\mathcal{F} = \{P_3\}$ [4, 49], and $\mathcal{F} = \{C_3\}$ [10, 16], imply the following complexity dichotomy when \mathcal{F} consists of a single connected graph H, which we suppose to have at least one edge.

Theorem 1. Let H be a connected graph. Under the ETH, $\{H\}$ -M-DELETION is solvable in time⁴

- $2^{\Theta(tw)} \cdot n^{\mathcal{O}(1)}$, if H is a contraction of the chair or the banner, and
- $2^{\Theta(\mathsf{tw} \cdot \log \mathsf{tw})} \cdot n^{\mathcal{O}(1)}$, otherwise.

This dichotomy is depicted in Figure 17, containing all connected graphs H with $2 \leq |V(H)| \leq 5$; note that if $|V(H)| \geq 6$, then H is not a contraction of the chair or the banner, and therefore the second item above applies. Note also that K_4 and the diamond are the only graphs on at most four vertices for which the problem is solvable in time $2^{\Theta(\text{tw} \cdot \log \text{tw})} \cdot n^{\mathcal{O}(1)}$ and that the chair and the banner are the only graphs on at least five vertices for which the problem is solvable in time $2^{\Theta(\text{tw} \cdot \log \text{tw})} \cdot n^{\mathcal{O}(1)}$ and that the chair and the banner are the only graphs on at least five vertices for which the problem is solvable in time $2^{\Theta(\text{tw})} \cdot n^{\mathcal{O}(1)}$.

The crucial role played by the chair and the banner in the complexity dichotomy may seem surprising at first sight. In fact, we realized a posteriori that the "easy" cases can be succinctly described in terms of

²In these papers [6–8], in some results we also considered the version of the problem where the graphs in \mathcal{F} are forbidden as topological minors; in the current paper we focus exclusively on the minor version.

³In the conference version of [6] we additionally required \mathcal{F} to be connected; in the journal version we proved this result without this assumption.

 $^{^{4}}$ We use n and tw for the number of vertices and the treewidth of the input graph, respectively.

the chair and the banner by taking a look at Figure 17. Note that the "easy" graphs can be equivalently characterized as those that are minors of the banner, with the exception of P_5 . Nevertheless, there is some intuitive reason for which excluding the chair or the banner constitutes the horizon on the existence of single-exponential algorithms. Namely, focusing on the banner, every connected component (with at least five vertices) of a graph that excludes the banner as a minor is either a cycle (of any length) or a tree in which some vertices have been replaced by triangles; both such types of components can be maintained by a dynamic programming algorithm in single-exponential time [7]. A similar situation occurs when excluding the chair. It appears that if the characterization of the allowed connected components is enriched in some way, such as restricting the length of the allowed cycles or forbidding certain degrees, the problem becomes inherently more difficult, inducing a transition from time $\mathcal{O}^*(2^{\Theta(tw)})$ to $\mathcal{O}^*(2^{\Theta(tw \cdot \log tw)})$.

Organization of the paper. In Section 2 we provide a high-level overview of the algorithm running in time $\mathcal{O}^*(2^{\mathcal{O}(\mathsf{tw} \cdot \log \mathsf{tw})})$. In Section 3 we give some preliminaries. In Section 4 we deal with flat walls, in Section 5 we apply the irrelevant vertex technique in the context of boundaried graphs, and in Section 6 we use this in order to bound the size of the dynamic programming tables. We conclude the article in Section 7 with some open questions for further research. In Appendix B we present an estimation of the constants depending on the (fixed) collection \mathcal{F} in our algorithm (cf. Corollary 36).

2 Overview of the algorithm

In order to obtain our algorithm of time $2^{\mathcal{O}(\mathsf{tw} \cdot \log \mathsf{tw})} \cdot n$ for every collection \mathcal{F} , our approach can be streamlined as follows. We use the machinery of boundaried graphs, equivalence relations, and representatives originating in the seminal work of Bodlaender et al. [11] and subsequently used, for instance, in [6, 21, 22, 32]. Let h be a constant depending only on the collection \mathcal{F} (to be defined in the formal description of the algorithm) and let t be a positive integer that is at most the treewidth of the input graph plus one. Skipping several technical details, a t-boundaried graph is a graph with a distinguished set of vertices -its boundary-labeled bijectively with integers from the set [t]. We say that two t-boundaried graphs are h-equivalent if for any other t-boundaried graph that we can "glue" to each of them, resulting in graphs G_1 and G_2 , and every graph H on at most h vertices, H is a minor of G_1 if and only if it is a minor of G_2 (cf. Section 3 for the precise definitions). Let $\mathcal{R}_{h}^{(t)}$ be a set of minimum-sized representatives of this equivalence relation. Since h-equivalent (boundaried) graphs have the same behavior in terms of eventual occurrences of minors of size up to h, there is a generic dynamic programming algorithm (already used in [6]) to solve \mathcal{F} -M-DELETION on a rooted tree decomposition of the input graph, via a typical bottom-up approach: at every bag B of the tree decomposition, naturally associated with a t-boundaried graph G_B , and for every representative $R \in \mathcal{R}_h^{(t)}$, store the minimum size of a set $S \subseteq V(G_B)$ such that the graph $G_B \setminus S$ is h-equivalent to R (cf. Subsection 6.2) for some more details⁵). This yields an algorithm running in time $\mathcal{O}^*(|\mathcal{R}_h^{(t)}|^2)$, and therefore it suffices to prove that $|\mathcal{R}_h^{(t)}| = 2^{\mathcal{O}_h(t \cdot \log t)}$, where the notation " $\mathcal{O}_h(\cdot)$ " means that the hidden constants depend only on h. Since we may assume that the graphs in $\mathcal{R}_h^{(t)}$ exclude some graph on at most h vertices as a minor (as all those that do not are h-equivalent), hence they have a linear number of edges [39], it is enough to prove that, for every $R \in \mathcal{R}_h^{(t)}$, it holds that

$$|V(R)| = \mathcal{O}_h(t). \tag{1}$$

Note that this is indeed sufficient in order to obtain an algorithm within the claimed running time, as there are at most $\binom{|V(R)|^2}{|E(R)|} = 2^{\mathcal{O}_h(|V(R)| \cdot \log |V(R)||)}$ representatives, and $t! = 2^{\mathcal{O}(t \cdot \log t)}$ possible labelings for

⁵This step was the only reason for which in the conference version of this article we required the collection \mathcal{F} to be connected. As mentioned in Section 1, in the full version of [6] we dropped the connectivity assumption, which implies that in the current article we can drop it as well.

the vertices in the boundary. In order to prove Equation 1, we combine a number of different techniques, which we proceed to discuss informally, and that are schematically summarized in Figure 1:



Figure 1: Diagram of the algorithm in time $2^{\mathcal{O}(\mathsf{tw} \cdot \log \mathsf{tw})} \cdot n$ for any collection \mathcal{F} .

▶ We use the Flat Wall Theorem of Robertson and Seymour [43], in particular a version (Theorem 5) that has been recently proved in [46] (and is based on the framework of Kawarabayashi et al. [30]), which incorporates the so-called regularity property. In a nutshell, this theorem says that every K_h -minor-free graph G has a set of vertices $A \subseteq V(G)$ –called *apices*– with $|A| = \mathcal{O}_h(1)$ such that $G \setminus A$ contains a *flat wall* of height $\Omega_h(\mathsf{tw}(G))$. Here, the definition of "flat wall" is quite involved and is detailed in Section 4; it essentially means a subgraph that has a bidimensional grid-like structure, separated from the rest of the graph by its perimeter, and that is "close" to being planar, in the sense that it can be embedded in the plane in a way that its potentially non-planar pieces, called *flaps*, have a well-defined structure along larger pieces called *bricks*.

▶ A subwall of a flat wall is *h*-homogeneous if for every brick of the subwall, the flaps within that brick have the same variety of *h*-folios, that is, the same sets of "boundaried" minors of detail at most h (the detail of

a boundaried graph is the maximum between its number of edges and its number of non-boundary vertices). This notion is inspired (but is not the same) by the one defined by Robertson and Seymour in [43]. Using standard "zooming" arguments, we can prove that, given a flat wall, we can find a large h-homogeneous subwall inside it (Proposition 9). Homogeneous subwalls are very useful because, as we explain below, they permit the application of the irrelevant vertex technique adapted to our purposes.

▶ We say that a vertex set *S* affects a flat wall if some vertex within the wall has a neighbor in *S* that is not an apex. With these definitions at hand, we define a parameter, denoted by $\mathbf{p}_{h,r}$ in this informal description, mapping every graph *G* to the smallest size of a vertex set that affects all *h*-homogeneous flat walls with at most *h* apices and height at least *r* in *G*. It is not hard to prove that the parameter $\mathbf{p}_{h,r}$ has a "bidimensional" behavior [17, 18, 21], in the sense that its value on a flat wall depends quadratically on the height of the wall (Lemma 10) and it is separable [11, 18, 21] (Lemma 11).

▶ The most complicated step towards proving Equation 1 is to find an "irrelevant" vertex inside a sufficiently large (in terms of h) flat wall of a boundaried graph that is *not* affected by its boundary (Theorem 23). Informally, here "irrelevant" means a non-boundary vertex of R that can be avoided by any minor model of a graph on at most h vertices and edges that traverses the boundary of R, *no matter* the graph that may be glued to it and *no matter how* this model traverses the boundary of R; see Section 5 for the precise definition. The *irrelevant vertex technique* originated in the seminal work of Robertson and Seymour [43, 44] and has become a very useful tool used in various kinds of linkage and cut problems [1, 27, 33, 35, 41]. Nevertheless, given the nature of our setting, it is critical that the size of the flat wall where the irrelevant vertex appears does *not depend* on the boundary size. To the best of our knowledge, this property is not guaranteed by the existing results on the irrelevant vertex technique (such as [43, (10.2)] and its subsequent proof in [44]). To achieve it and, moreover, in order to make an estimation of the parametric dependencies, we develop a self-reliant theoretical framework that uses the following ingredients:

- With a flat wall W we associate a bipartite graph W, which we call its *leveling* as defined in [46]; cf. Subsection 5.3 for the precise definition. In particular, this graph has a vertex for every flap of the flat wall, and can be embedded in a disk in a planar way.
- It turns out to be more convenient to work with *topological* minor models instead of minor models; we can afford it since for every graph H there are at most f(H) different topological minor minimal graphs that contain H as a minor (Observation 3). The reason for this is that it is easier to deal with the branch vertices of a topological minor model in the analysis. Given a topological minor model, we say that a flap of a wall is *dirty* if it contains a branch vertex of the model, or there is an edge from the flap to an apex vertex of the wall. We also define the leveling of a topological minor model, and we equip its dirty flaps with colors that encode their *h*-folios. We now proceed to explain how to reroute the colored leveling of a topological minor model.
- In order to reroute (colored levelings of) topological minor models, it will be helpful to use *railed annuli*, a structure introduced in [28] that occurs as a subgraph inside a flat wall (Proposition 12) and that has the following nice property, recently proved in [25] (Proposition 13): if a railed annulus is large enough compared to h, every topological minor model of a graph on at most h vertices traversing it can be rerouted so that the branch vertices are preserved and such that, more importantly, the intersection of the new model with a large *prescribed* part of the railed annulus is *confined*, in the sense that it is only allowed to use a well-defined set of paths in that part, which does not depend on the original model.
- We also need a technical result with a graph drawing flavor (Lemma 16) guaranteeing that large enough railed annuli contain topological minor models of every graph of maximum degree three with the property, in particular, that certain vertices are pairwise far apart in the embedding. Using this result and the one proved in [25] mentioned above, we can finally prove (Theorem 17) that every

topological minor model of a graph H inside a graph with a large flat wall W can be "collapsed" inside the wall, in the following sense: G contains another topological minor model of a graph H', such that H is a minor of H', and such that the new model avoids the central part of the annulus; here is where the irrelevant vertex will be found.

• To conclude, it just remains to "lift" the constructed embedding of the colored leveling of the topological minor to an embedding of the "original" minor in the flat wall (Theorem 23). For that, we exploit the fact that we have rerouted the model inside an h-homogeneous subwall not affected by the boundary, which allows to mimic the behavior of the original minor inside the flaps of the wall, using that all bricks have the same variety of h-folios.

The above arguments, incorporated in the proof of Theorem 23, imply that if $R \in \mathcal{R}_h^{(t)}$ is a minimumsized representative, then its boundary affects all large enough *h*-homogeneous flat walls, as otherwise we could remove an irrelevant vertex and find a smaller equivalent representative. In particular, it follows that, for every $R \in \mathcal{R}_h^{(t)}$, we have $\mathbf{p}_{h,r}(R) \leq t$ (Corollary 25).

▶ Combining that the parameter $\mathbf{p}_{h,r}$ is "bidimensional" and separable along with the fact that $\mathbf{p}_{h,r}(R) \leq t$ for every $R \in \mathcal{R}_h^{(t)}$, we prove in Lemma 27 that every representative $R \in \mathcal{R}_h^{(t)}$ has a vertex subset S containing its boundary, with $|S| = \mathcal{O}_h(t)$, whose removal leaves a graph of treewidth bounded by a function of h; such a set is called a *treewidth modulator*. (In Appendix B we provide an improved version of Lemma 27, namely Lemma 34, by adapting the proof of [21, Lemma 3.6].)

▶ Once we have a treewidth modulator of size $\mathcal{O}_h(t)$ of a representative R, all that remains is to pipeline it with known techniques to compute an appropriate protrusion decomposition [32] (Lemma 29) and to reduce protrusions to smaller equivalent ones of size bounded by a function of h –we use the version given in [6] adapted to the \mathcal{F} -M-DELETION problem– (Lemma 30), implying that $|V(R)| = \mathcal{O}_h(t)$ for every every $R \in \mathcal{R}_h^{(t)}$ and concluding the proof of Equation 1.

It should be noted that all the items above do not need to be converted into an algorithm, they are just used in the analysis: the conclusion is that if $R \in \mathcal{R}_h^{(t)}$ is a minimum-sized representative, then $|V(R)| = \mathcal{O}_h(t)$, as otherwise some reduction rule could be applied to it (either by removing an irrelevant vertex or by protrusion replacement), thus obtaining an equivalent representative of smaller size and contradicting its minimality. Our main result can be formally stated as follows.

Theorem 2. Let \mathcal{F} be a finite non-empty collection of non-empty graphs. There exists a constant $c_{\mathcal{F}}$ such that the \mathcal{F} -M-DELETION problem is solvable in time $c_{\mathcal{F}}^{\mathsf{tw} \cdot \log \mathsf{tw}} \cdot n$ on n-vertex graphs of treewidth at most tw .

In Appendix B we provide an estimation of the constant $c_{\mathcal{F}}$ in the above theorem based on the parametric dependencies of the Unique Linkage Theorem [31, 44].

3 Preliminaries

3.1 Basic definitions

Sets and integers. We denote by \mathbb{N} the set of non-negative integers. Given two integers p, q, where $p \leq q$, we denote by [p,q] the set $\{p, \ldots, q\}$. For an integer $p \geq 1$, we set [p] = [1,p] and $\mathbb{N}_{\geq p} = \mathbb{N} \setminus [0, p-1]$. Given a non-negative integer p, we denote by $\mathsf{odd}(p)$ the minimum odd number that is not smaller than p. For a set S, we denote by 2^S the set of all subsets of S and by $\binom{S}{2}$ the set of all subsets of S of size two. If S is a collection of objects where the operation \cup is defined, then we denote $\bigcup S = \bigcup_{X \in S} X$.

Basic concepts on graphs. As a graph G we denote any pair (V, E) where V is a finite set and $E \subseteq \binom{V}{2}$, that is, all graphs of this paper are undirected, finite, and without loops or multiple edges. We also define V(G) = V and E(G) = E. Unless stated otherwise, we denote by n and m the number of vertices and edges, respectively, of the graph under consideration. We say that a pair $(L, R) \in 2^{V(G)} \times 2^{V(G)}$ is a separation of G if $L \cup R = V(G)$ and there is no edge in G between $L \setminus R$ and $R \setminus L$. The order of a separation (L, R) is $|L \cap R|$. Given a vertex $v \in V(G)$, we denote by $N_G(v)$ the set of vertices of G that are adjacent to v in G. Also, given a set $S \subseteq V(G)$, we set $N_G(S) = (\bigcup_{v \in S} N_G(v)) \setminus S$. A vertex $v \in V(G)$ is isolated if $N_G(v) = \emptyset$. For $S \subseteq V(G)$, we set $G[S] = (S, E \cap \binom{S}{2})$ and use $G \setminus S$ to denote $G[V(G) \setminus S]$. Given an edge $e = \{u, v\} \in E(G)$, we define the subdivision of e to be the operation of deleting e, adding a new vertex w, and making it adjacent to u and v. Given two graphs H, G, we say that H is a subdivision of G if H can be obtained from G by subdividing edges. The contraction of an edge $e = \{u, v\}$ of a simple graph G results in a simple graph G' obtained from $G \setminus \{u, v\}$ by adding a new vertex uv adjacent to all the vertices in the set $N_G(u) \cup N_G(v) \setminus \{u, v\}$. A graph G' is a minor of a graph G if G' can be obtained from a subgraph of G after a series of edge contractions. The distance between two vertices x and y of a graph G is the number of edges of a shortest path between x and y in G.

Treewidth. Let G = (V, E) be a graph. A tree decomposition of G is a pair $(T, \mathcal{X} = \{X_t\}_{t \in V(T)})$ where T is a tree and \mathcal{X} is a collection of subsets of V such that

- $\bigcup_{t \in V(T)} X_t = V$,
- $\forall e = \{u, v\} \in E, \exists t \in V(T) : \{u, v\} \subseteq X_t$, and
- $\forall v \in V, T[\{t \mid v \in X_t\}]$ is connected.

We call the vertices of T nodes and the sets in \mathcal{X} bags of the tree decomposition (T, \mathcal{X}) . The width of (T, \mathcal{X}) is equal to max $\{|X_t| - 1 \mid t \in V(T)\}$ and the treewidth of G is the minimum width over all tree decompositions of G. We denote the treewidth of a graph G by tw(G).

For $t \in \mathbb{N}$, we say that a set $S \subseteq V(G)$ is a *t*-treewidth modulator of G if $\mathsf{tw}(G \setminus S) \leq t$.

3.2 Formal definition of the problem

Let \mathcal{F} be a finite non-empty collection of non-empty graphs; we call such a collection *proper*. We extend the minor relation to \mathcal{F} such that, given a graph G, $\mathcal{F} \preceq_{\mathsf{m}} G$ if and only if there exists a graph $H \in \mathcal{F}$ such that $H \preceq_{\mathsf{m}} G$. We also denote $\mathsf{exc}_{\mathsf{m}}(\mathcal{F}) = \{G \mid \mathcal{F} \not\preceq_{\mathsf{m}} G\}$, i.e., $\mathsf{exc}_{\mathsf{m}}(\mathcal{F})$ is the class of graphs that do *not* contain any graph in \mathcal{F} as a minor.

Let \mathcal{F} be a proper collection. We define the graph parameter $\mathbf{m}_{\mathcal{F}}$ as the function that maps graphs to non-negative integers as follows:

$$\mathbf{m}_{\mathcal{F}}(G) = \min\{|S| \mid S \subseteq V(G) \land G \setminus S \in \mathsf{exc}_{\mathsf{m}}(\mathcal{F})\}.$$

The main objective of this paper is to study the problem of computing the parameter $\mathbf{m}_{\mathcal{F}}$ for graphs of bounded treewidth. The corresponding decision problem is formally defined as follows.

 \mathcal{F} -M-DELETION **Input:** A graph G and an integer $k \in \mathbb{N}$. **Parameter:** The treewidth of G. **Output:** Is $\mathbf{m}_{\mathcal{F}}(G) \leq k$?

3.3 Boundaried graphs, folios, and representatives

Boundaried graphs. Let $t \in \mathbb{N}$. A *t*-boundaried graph is a triple $\mathbf{G} = (G, B, \rho)$ where G is a graph, $B \subseteq V(G), |B| = t$, and $\rho : B \to [t]$ is a bijection. We say that two *t*-boundaried graphs $\mathbf{G}_1 = (G_1, B_1, \rho_1)$ and $\mathbf{G}_2 = (G_2, B_2, \rho_2)$ are *isomorphic* if there is an isomorphism from G_1 to G_2 that extends the bijection $\rho_2^{-1} \circ \rho_1$. The triple (G, B, ρ) is a boundaried graph if it is a *t*-boundaried graph for some $t \in \mathbb{N}$. As in [43], we define the detail of a boundaried graph $\mathbf{G} = (G, B, \rho)$ as detail(\mathbf{G}) := max{ $|E(G)|, |V(G) \setminus B|$ }. We denote by $\mathcal{B}^{(t)}$ the set of all (pairwise non-isomorphic) *t*-boundaried graphs and by $\mathcal{B}^{(t)}_h$ the set of all (pairwise non-isomorphic) *t*-boundaried graphs with detail at most *h*. We also set $\mathcal{B} = \bigcup_{t \in \mathbb{N}} \mathcal{B}^{(t)}$.

Minors and topological minors of boundaried graphs. We say that a t-boundaried graph $\mathbf{G}_1 = (G_1, B_1, \rho_1)$ is a minor of a t-boundaried graph $\mathbf{G}_2 = (G_2, B_2, \rho_2)$, denoted by $\mathbf{G}_1 \leq_{\mathsf{m}} \mathbf{G}_2$, if there is a sequence of removals of non-boundary vertices, edge removals, and edge contractions in G_2 , disallowing contractions of edges with both endpoints in B_2 , that transforms \mathbf{G}_2 to a boundaried graph that is isomorphic to \mathbf{G}_1 (during edge contractions, boundary vertices prevail). Note that this extends the usual definition of minors in graphs without boundary.

We say that (M, T) is a tm-pair if M is a graph, $T \subseteq V(M)$, and all vertices in $V(M) \setminus T$ have degree two. We denote by diss(M, T) the graph obtained from M by dissolving all vertices in $V(M) \setminus T$, that is, for every vertex $v \in V(M) \setminus T$, with neighbors u and w, we delete v and, if u and w are not adjacent, we add the edge $\{u, w\}$. A tm-pair of a graph G is a tm-pair (M, T) where M is a subgraph of G.

Given two graphs H and G, we say that a tm-pair (M,T) of G is a topological minor model of H in Gif H is isomorphic to diss(M,T). We denote this isomorphism by $\sigma_{M,T}: V(H) \to T$. We call the vertices in T branch vertices of (M,T). We call each path in M between two distinct branch vertices and with no internal branch vertices a subdivision path of (M,T) and the internal vertices of such paths, i.e., the vertices of $V(M) \setminus T$, are the subdivision vertices of (M,T). We also extend $\sigma_{M,T}$ so to also map each $e = \{x,y\} \in E(H)$ to the subdivision path of M with endpoints $\sigma_{M,T}(x)$ and $\sigma_{M,T}(y)$. Furthermore, we extend $\sigma_{M,T}$ so to also map each subgraph H' of H to the subgraph of M consisting of the vertices of $\sigma_{M,T}(T)$ and the paths in $\sigma_{M,T}(e), e \in E(H')$.

If $\mathbf{M} = (M, B, \rho) \in \mathcal{B}$ and $T \subseteq V(M)$ with $B \subseteq T$ and such that all vertices in $V(M) \setminus T$ have degree two, we call (\mathbf{M}, T) a btm-pair and we define diss $(\mathbf{M}, T) = (\text{diss}(M, T), B, \rho)$. Note that we do not permit dissolution of boundary vertices, as we consider all of them to be branch vertices. If $\mathbf{G} = (G, B, \rho)$ is a boundaried graph and (M, T) is a tm-pair of G where $B \subseteq T$, then we say that (\mathbf{M}, T) , where $\mathbf{M} = (M, B, \rho)$, is a btm-pair of $\mathbf{G} = (G, B, \rho)$. Let $\mathbf{G}_i = (G_i, B_i, \rho_i), i \in [2]$. We say that \mathbf{G}_1 is a *topological minor* of \mathbf{G}_2 , denoted by $\mathbf{G}_1 \preceq_{tm} \mathbf{G}_2$, if \mathbf{G}_2 has a btm-pair (\mathbf{M}, T) such that diss (\mathbf{M}, T) is isomorphic to \mathbf{G}_1 .

Given a $\mathbf{G} = (G, B, \rho) \in \mathcal{B}$, we define $ext(\mathbf{G})$ as the set containing every topological-minor-minimal boundaried graph $\mathbf{G}' = (G', B, \rho)$ among those that contain \mathbf{G} as a minor. Notice that we insist that Band ρ are the same for all graphs in $ext(\mathbf{G})$. Moreover, we do not consider isomorphic boundaried graphs in $ext(\mathbf{G})$ as different boundaried graphs. The set $ext(\mathbf{G})$ helps us to express the minor relation in terms of the topological minor relation because of the following simple observation. Note that this definition extends naturally to graphs, seen as boundaried graphs with empty boundary.

Observation 3. If $\mathbf{G}_1, \mathbf{G}_2 \in \mathcal{B}$, then $\mathbf{G}_1 \preceq_{\mathsf{m}} \mathbf{G}_2 \iff \exists \mathbf{G} \in \mathsf{ext}(\mathbf{G}_2) : \mathbf{G}_1 \preceq_{\mathsf{tm}} \mathbf{G}$. Moreover, if \mathbf{G} is a boundaried graph with detail h, then every graph in $\mathsf{ext}(\mathbf{G})$ has detail at most 3h.

Folios. We define the *h*-folio of $\mathbf{G} = (G, B, \rho) \in \mathcal{B}$ as

h-folio(**G**) = {**G**' $\in \mathcal{B} \mid \mathbf{G}' \preceq_{\mathsf{tm}} \mathbf{G}$ and **G**' has detail at most h}.

Using the fact that an h-folio is a collection of K_{h+1} -minor-free boundaried graphs, it follows that the

h-folio of a *t*-boundaried graph has at most $2^{\mathcal{O}((h+t)\cdot\log(h+t))}$ elements. Therefore, the number of distinct *h*-folios of *t*-boundaried graphs is given by the following lemma (also observed in [6]).

Lemma 4. There exists a function $f_1 : \mathbb{N}^2 \to \mathbb{N}$ such that for every $t, h \in \mathbb{N}$, $|\{h\text{-folio}(\mathbf{G}) \mid \mathbf{G} \in \mathcal{B}_h^{(t)}\}| \leq f_1(t,h)$. Moreover, $f_1(t,h) = 2^{2^{\mathcal{O}((h+t) \cdot \log(h+t))}}$.

Equivalent boundaried graphs and representatives. We say that two boundaried graphs $\mathbf{G}_1 = (G_1, B_1, \rho_1)$ and $\mathbf{G}_2 = (G_2, B_2, \rho_2)$ are compatible if $\rho_2^{-1} \circ \rho_1$ is an isomorphism from $G_1[B_1]$ to $G_2[B_2]$. Given two compatible boundaried graphs $\mathbf{G}_1 = (G_1, B_1, \rho_1)$ and $\mathbf{G}_2 = (G_2, B_2, \rho_2)$, we define $\mathbf{G}_1 \oplus \mathbf{G}_2$ as the graph obtained if we take the disjoint union of G_1 and G_2 and, for every $i \in [|B_1|]$, we identify vertices $\rho_1^{-1}(i)$ and $\rho_2^{-1}(i)$.

Given $h \in \mathbb{N}$, we say that two boundaried graphs \mathbf{G}_1 and \mathbf{G}_2 are *h*-equivalent, denoted by $\mathbf{G}_1 \equiv_h \mathbf{G}_2$, if they are compatible and, for every graph H on at most h vertices and h edges and every boundaried graph \mathbf{F} that is compatible with \mathbf{G}_1 (hence, with \mathbf{G}_2 as well), it holds that

$$H \preceq_{\mathsf{m}} \mathbf{F} \oplus \mathbf{G}_1 \iff H \preceq_{\mathsf{m}} \mathbf{F} \oplus \mathbf{G}_2.$$
⁽²⁾

Note that \equiv_h is an equivalence relation on \mathcal{B} . A minimum-sized (in terms of number of vertices) element of an equivalence class of \equiv_h is called a *representative* of \equiv_h . For $t \in \mathbb{N}$, a set of t-representatives for \equiv_h is a collection containing a minimum-sized representative for each equivalence class of \equiv_h restricted to $\mathcal{B}^{(t)}$. Given $t, h \in \mathbb{N}$, we denote by $\mathcal{R}_h^{(t)}$ a set of t-representatives for \equiv_h .

At this point, we wish to stress that the folio-equivalence defined in Equation 2 is related but is not the same as the one defined by "having the same h-folio". Indeed, observe first that if \mathbf{G}_1 and \mathbf{G}_2 are compatible t-boundaried graphs and h-folio(\mathbf{G}_1) = h-folio(\mathbf{G}_2) then $\mathbf{G}_1 \equiv_h \mathbf{G}_2$, therefore the folio-equivalence is a refinement of \equiv_h . In fact, a dynamic programming procedure for solving \mathcal{F} -M-DELETION can also be based on the folio-equivalence, and this has already been done in the general algorithm in [6], which has a double-exponential parametric dependence due to the bound of Lemma 4. In this paper we build our dynamic programming on the equivalence \equiv_h and we essentially prove that \equiv_h is "coarse enough" so to reduce the double-exponential parametric dependence of the dynamic programming to a single-exponential one. In fact, this has already been done in [6] for the case where \mathcal{F} contains some planar graph, as this structural restriction directly implies an upper bound on the treewidth of the representatives. To deal with the general case, the only structural restriction for the (non-trivial) representatives is the exclusion of H as a minor. All the combinatorial machinery that we introduce in the next two sections is intended to deal with the structure of this general and (more entangled) setting.

4 Flat walls

In this section we deal with flat walls. More precisely, in Subsections 4.1, 4.2, and 4.3 we give the definition of a *flat wall* in the form of a *flatness pair*. In Subsection 4.4 we define the notion of *regular* flatness pair and we give a version of the Flat Wall Theorem of Robertson and Seymour [43] that has been recently proved in [46]. This version (Theorem 5) incorporates the regularity property and is based on the recent results and the terminology of Kawarabayashi et al. [30]. In Subsection 4.5 we define a notion of homogeneity of flat walls, also introduced in [46], that along with Theorem 5 will be the combinatorial framework for the proofs of Section 5. We stress that the notion of homogeneity that we use is different from that defined by Robertson and Seymour in [43] and can serve as an alternative for further applications based on the technology of flat walls (see e.g. [45, 47, 48]). In Subsection 4.6 we define a graph parameter related to flat walls and show that it enjoys a series of properties related to Bidimensionality (as introduced in [17] and further developed in [21]).

4.1 Walls and subwalls

We first introduce some basic concepts such as partially disk-embedded graphs, walls, subwalls, tilts, and layers (for an example of all the concepts defined in this subsection, see Figure 2).

Partially disk-embedded graphs. A closed (resp. open) disk is a set homeomorphic to the set $\{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 \leq 1\}$ (resp. $\{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 < 1\}$). Let Δ be an open or closed disk. We use $\mathsf{bd}(\Delta)$ to denote the boundary of Δ and, if Δ is closed, we use $\mathsf{int}(\Delta)$ to denote the open disk $\Delta \setminus \mathsf{bd}(\Delta)$. Also, if Δ is an open disk, we use $\overline{\Delta} = \Delta \cup \mathsf{bd}(\Delta)$ for the closure of Δ . When we embed a graph G in the plane or in a disk, we treat G (both its vertex and edge sets) as a set of points. This permits us to make set operations between graphs and sets of points.

If Δ is a closed disk, we say that a graph G is Δ -embedded if G is embedded in Δ without crossings such that the intersection of $bd(\Delta)$ and G (seen as a set of points of Δ) is a subset of V(G). We say that a graph G is partially disk-embedded in some closed disk Δ , if there is some Δ -embedded subgraph, say K, of G such that $G \cap \Delta = K$ and $(V(G) \cap \Delta, V(G) \setminus int(\Delta))$ is a separation of G. From now on, we use the term partially Δ -embedded graph G to denote that a graph G is partially disk-embedded in some closed disk Δ .

A circle of Δ is any set homeomorphic to $\{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 = 1\}$. Given two distinct points $x, y \in \Delta$, an (x, y)-arc of Δ is any subset of Δ that is homeomorphic to the closed interval [0, 1].



Figure 2: An 13-wall W along with a choice of pegs and corners. The six layers of W are colored alternatively in red and green and the central 5-subwall of W appears in grey. The pegs are the squared vertices while, among them, those that are black are the corners. The original vertices that are not pegs are turquoise circles while the subdivision vertices are the small yellow circles. The central vertices are the two 3-branch vertices that are surrounded by white squares. Notice that W has 144 bricks and, among them, 100 are internal.

Walls. Let $k, r \in \mathbb{N}$. The $(k \times r)$ -grid is the graph whose vertex set is $[k] \times [r]$ and two vertices (i, j)and (i', j') are adjacent if |i - i'| + |j - j'| = 1. An elementary r-wall, for some odd integer $r \ge 3$, is the graph obtained from a $(2r \times r)$ -grid with vertices $(x, y) \in [2r] \times [r]$, after the removal of the "vertical" edges $\{(x, y), (x, y + 1)\}$ for odd x + y, and then the removal of all vertices of degree one. Notice that, as $r \ge 3$, an elementary r-wall is a planar graph that has a unique (up to topological isomorphism) embedding in the plane \mathbb{R}^2 such that all its finite faces are incident to exactly six edges. The perimeter of an elementary r-wall is the cycle bounding its infinite face, while the cycles bounding its finite faces are called *bricks*. Also, the vertices in the perimeter of an elementary *r*-wall that have degree two are called *pegs*, while the vertices (1,1), (2,r), (2r-1,1), and (2r,r) are called *corners* (notice that the corners are also pegs).

An *r*-wall is any graph W obtained from an elementary *r*-wall W after subdividing edges. A graph W is a wall if it is an *r*-wall for some odd $r \ge 3$ and we refer to *r* as the *height* of *W*. Given a graph *G*, a wall of *G* is a subgraph of *G* that is a wall. We insist that, for every *r*-wall, the number *r* is always odd.

We call the vertices of degree three of a wall W 3-branch vertices. The vertices that are created by subdivisions are called *subdivision vertices* while the rest are called *original* vertices of W. A cycle of W is a *brick* (resp. the *perimeter*) of W if its 3-branch vertices are the vertices of a brick (resp. the perimeter) of \overline{W} . We use D(W) in order to denote the perimeter of the wall W. A brick of W is *internal* if it is disjoint from D(W). Note that every wall W has a unique (up to homeomorphism) embedding in the plane whose infinite face is bounded by the perimeter D(W) of the wall. Each time we consider a plane-embedded wall, we consider this embedding.

Given two vertices x and y of a plane graph G, we define their *face-distance* in G as the smallest integer i such that there exists an arc of the plane (i.e., a subset homeomorphic to the interval [0, 1]) between x and y that does not cross the infinite face of the embedding, crosses no vertices of G, and intersects at most i faces of G. Note that two distinct vertices of a plane wall W are within face-distance one if and only if they belong to the same brick. Given two vertex sets X, Y of a plane graph G, we define the face-distance between X and Y as the minimum face-distance between a vertex in X and a vertex in Y.

Subwalls. Given an elementary r-wall \overline{W} , some $i \in \{1, 3, \ldots, 2r-1\}$, and i' = (i+1)/2, the *i'*-th vertical path of \overline{W} is the one whose vertices, in order of appearance, are $(i, 1), (i, 2), (i+1, 2), (i+1, 3), (i, 3), (i, 4), (i+1, 4), (i+1, 5), (i, 5), \ldots, (i, r-2), (i, r-1), (i+1, r-1), (i+1, r)$. Also, given some $j \in [2, r-1]$ the *j*-th horizontal path of \overline{W} is the one whose vertices, in order of appearance, are $(1, j), (2, j), \ldots, (2r, j)$.

A vertical (resp. horizontal) path of W is one that is a subdivision of a vertical (resp. horizontal) path of \overline{W} . Notice that the perimeter of an r-wall W is uniquely defined regardless of the choice of the elementary r-wall \overline{W} . A subwall of W is any subgraph W' of W that is an r'-wall, with $r' \leq r$ and such the vertical (resp. horizontal) paths of W' are subpaths of the vertical (resp. horizontal) paths of W.

Tilts. The *interior* of a wall W is the graph obtained from W if we remove from it all edges of D(W) and all vertices of D(W) that have degree two in W. Given two walls W and \tilde{W} of a graph G, we say that \tilde{W} is a *tilt* of W if \tilde{W} and W have identical interiors.

Layers. The *layers* of an *r*-wall *W* are recursively defined as follows. The first layer of *W* is its perimeter. For i = 2, ..., (r-1)/2, the *i*-th layer of *W* is the (i-1)-th layer of the subwall *W'* obtained from *W* after removing from *W* its perimeter and removing recursively all occurring vertices of degree one. The *central vertices* of an *r*-wall are its two 3-branch vertices that do not belong to any of its layers. See Figure 2 for an illustration of the notions defined above. Given an *r*-wall *W* and an odd integer *q*, where $3 \le q \le r$, the *central q-subwall of W* is the subwall of *W* of height *q* whose central vertices are the central vertices of *W*.

4.2 Paintings and renditions

Before defining flat walls, we need to introduce paintings and renditions. Here we closely follow the terminology of [30].

Paintings. Let Δ be a closed disk. A Δ -painting is a pair $\Gamma = (U, N)$ where

- N is a finite set of points of Δ ,
- $N \subseteq U \subseteq \Delta$, and
- $U \setminus N$ has finitely many arcwise-connected components, called *cells*, where, for every cell c,
 - the closure \bar{c} of c is a closed disk and
 - $\circ |\tilde{c}| \leq 3$, where $\tilde{c} := \mathsf{bd}(c) \cap N$.

We use the notation $U(\Gamma) := U$, $N(\Gamma) := N$ and denote the set of cells of Γ by $C(\Gamma)$. For convenience, we may assume that each cell of Γ is an open disk of Δ . See Figure 3 for an example of a Δ -painting.



Figure 3: A Δ -painting $\Gamma = (U, N)$. The red circles are the points of N that are points of the boundary of Δ (whose complement is drawn in grey) and the blue circles are those that lie in the interior of Δ . The set $U \setminus N$ is depicted in green.

Notice that, given a Δ -painting Γ , the pair $(N(\Gamma), \{\tilde{c} \mid c \in C(\Gamma)\})$ is a hypergraph whose hyperedges have cardinality at most three and Γ can be seen as a plane embedding of this hypergraph in Δ .

Renditions. Let G be a graph, and let Ω be a cyclic permutation of a subset of V(G) that we denote by $V(\Omega)$. By an Ω -rendition of G we mean a triple (Γ, σ, π) , where

- (a) Γ is a Δ -painting for some closed disk Δ ,
- (b) $\pi: N(\Gamma) \to V(G)$ is an injection, and
- (c) σ assigns to each cell $c \in C(\Gamma)$ a subgraph $\sigma(c)$ of G, such that
 - (1) $G = \bigcup_{c \in C(\Gamma)} \sigma(c),$
 - (2) for distinct $c, c' \in C(\Gamma)$, $\sigma(c)$ and $\sigma(c')$ are edge-disjoint,
 - (3) for every cell $c \in C(\Gamma)$, $\pi(\tilde{c}) \subseteq V(\sigma(c))$,

- (4) for every cell $c \in C(\Gamma)$, $V(\sigma(c)) \cap \bigcup_{c' \in C(\Gamma) \setminus \{c\}} V(\sigma(c')) \subseteq \pi(\tilde{c})$, and
- (5) $\pi(N(\Gamma) \cap \mathsf{bd}(\Delta)) = V(\Omega)$, such that the points in $N(\Gamma) \cap \mathsf{bd}(\Delta)$ appear in $\mathsf{bd}(\Delta)$ in the same ordering as their images, via π , in Ω .

Given an Ω -rendition (Γ, σ, π) of a graph G, we call a cell c of Γ trivial if $\pi(\tilde{c}) = V(\sigma(c))$.

Tight renditions. We say that an Ω -rendition (Γ, σ, π) of a graph G is *tight* if the following conditions are satisfied:

- (i) If there are two points x, y of $N(\Gamma)$ such that $e = \{\pi(x), \pi(y)\} \in E(G)$, then there is a cell $c \in C(\Gamma)$ such that $\sigma(c)$ is the two-vertex connected graph $(e, \{e\})$,
- (ii) for every $c \in C(\Gamma)$, every two vertices in $\pi(\tilde{c})$ belong to some path of $\sigma(c)$,
- (iii) for every $c \in C(\Gamma)$ and every connected component C of the graph $\sigma(c) \setminus \pi(\tilde{c})$, if $N_{\sigma(c)}(V(C)) \neq \emptyset$, then $N_{\sigma(c)}(V(C)) = \pi(\tilde{c})$,
- (iv) there are no two distinct non-trivial cells c_1 and c_2 such that $\pi(\tilde{c_1}) = \pi(\tilde{c_2})$, and
- (v) for every $c \in C(\Gamma)$ there are $|\tilde{c}|$ vertex-disjoint paths in G from $\pi(\tilde{c})$ to the set $V(\Omega)$.

As proved in [46], it is possible to transform any Ω -rendition to a tight one. For this reason, in this paper, we always assume that Ω -renditions are tight.

4.3 Flat walls and flatness pairs

We are now in position to define the notion of a flat wall. We further encode it into the concept of a *flatness* pair of a graph.

Flat walls. Let G be a graph and let W be an r-wall of G, for some odd integer $r \ge 3$. We say that a pair $(P,C) \subseteq D(W) \times D(W)$ is a choice of pegs and corners for W if W is the subdivision of an elementary r-wall \overline{W} where P and C are the pegs and the corners of \overline{W} , respectively (clearly, $C \subseteq P$). To get more intuition, notice that a wall W can occur in several ways from the elementary wall \overline{W} , depending on the way the vertices in the perimeter of \overline{W} are subdivided. Each of them gives a different selection (P,C) of pegs and corners of W (see Figure 2 for an example of a choice of pegs and corners (P,C) in a 13-wall W).

We say that W is a *flat* r-wall of G if there is a separation (X, Y) of G and a choice (P, C) of pegs and corners for W such that:

- $V(W) \subseteq Y$,
- $P \subseteq X \cap Y \subseteq V(D(W))$, and
- if Ω is the cyclic ordering of the vertices $X \cap Y$ as they appear in D(W), then there exists an Ω -rendition (Γ, σ, π) of G[Y].



Figure 4: A graph G and a flatness pair (W, \mathfrak{R}) of G where W is a 5-wall and $\mathfrak{R} = (X, Y, P, C, \Gamma, \sigma, \pi)$ is a 7-tuple certifying the flatness of W in G. The edges of W are drawn in orange. In the corresponding separation (X, Y), the vertices of X are green and yellow while the vertices in Y are all the non-green vertices. Consequently, the yellow vertices are the vertices in $X \cap Y$. The pegs and the corners are the squared vertices where the pegs that are not corners are purple and the the corners are black. Untidy cells are marked by green stars. The Δ -painting of the Ω -rendition (Γ, σ, π) of G' = G[Y] is the one depicted in Figure 3.

Flatness pairs. Given the above, we say that the choice of the 7-tuple $\mathfrak{R} = (X, Y, P, C, \Gamma, \sigma, \pi)$ certifies that W is a flat wall of G. We call the pair (W, \mathfrak{R}) a flatness pair of G and define the height of the pair (W, \mathfrak{R}) to be the height of W. We use the term cell of \mathfrak{R} in order to refer to the cells of Γ (see Figure 4 for an example of a flatness pair (W, \mathfrak{R}) of a graph).

We call the graph G[Y] the \mathfrak{R} -compass of W in G, denoted by $\mathsf{compass}_{\mathfrak{R}}(W)$ (see Figure 5 for the \mathfrak{R} compass of W, corresponding to the flatness pair (W, \mathfrak{R}) of Figure 4). We define the flaps of the wall W in \mathfrak{R} as $\mathsf{flaps}_{\mathfrak{R}}(W) := \{\sigma(c) \mid c \in C(\Gamma)\}$. Given a flap $F \in \mathsf{flaps}_{\mathfrak{R}}(W)$, we define its base as $\partial F := V(F) \cap \pi(N(\Gamma))$. A flap $F \in \mathsf{flaps}_{\mathfrak{R}}(W)$ is trivial if $|\partial F| = 2$ and F consists of one edge between the two vertices in ∂F . We call the edges of the trivial flaps short edges of $\mathsf{compass}_{\mathfrak{R}}(W)$. A cell c of \mathfrak{R} is untidy if $\pi(\tilde{c})$ contains a vertex x of W such that two of the edges of W that are incident to x are edges of $\sigma(c)$. Notice that if c is untidy then $|\tilde{c}| = 3$. A cell is tidy if it is not untidy (in Figure 4 untidy cells are marked by green stars).

4.4 Influence, regularity, and tilts of flatness pairs

We now introduce a classification of the cells of a flatness pair (W, \mathfrak{R}) . This classification will be used in order to define the concepts of regularity and W'-tilts of flatness pairs that will be important for our proofs.

Cell classification. Given a cycle C of $\operatorname{compass}_{\mathfrak{R}}(W)$, we say that C is \mathfrak{R} -normal if it is not a subgraph of a flap $F \in \operatorname{flaps}_{\mathfrak{R}}(W)$. Given an \mathfrak{R} -normal cycle C of $\operatorname{compass}_{\mathfrak{R}}(W)$, we call a cell c of \mathfrak{R} C-perimetric if $\sigma(c)$ contains some edge of C. Notice that if c is C-perimetric, then $\pi(\tilde{c})$ contains two points $p, q \in N(\Gamma)$ such that $\pi(p)$ and $\pi(q)$ are vertices of C where one, say P_c^{in} , of the two $(\pi(p), \pi(q))$ -subpaths of C is a subgraph of $\sigma(c)$ and the other, denoted by P_c^{out} , $(\pi(p), \pi(q))$ -subpath contains at most one internal vertex of $\sigma(c)$,



Figure 5: The \mathfrak{R} -compass of the 5-wall W for the flatness pair (W, \mathfrak{R}) depicted in Figure 4, and a subwall W' of W whose edges are depicted in bold. The red curve is the curve $K_{W'}$. The W'-internal cells are depicted in grey while the W'-perimetric cells are depicted in green. W'-marginal cells are marked with orange stars. The set $\operatorname{influence}_{\mathfrak{R}}(W')$ contains all the flaps that are drawn inside the grey or the green cells (the W'-external cells are not depicted).

which should be the (unique) vertex z in $\partial \sigma(c) \setminus \{\pi(p), \pi(q)\}$. We pick a (p, q)-arc A_c in $\hat{c} := c \cup \tilde{c}$ such that $\pi^{-1}(z) \in A_c$ if and only if P_c^{in} contains the vertex z as an internal vertex.

We consider the circle $K_C = \bigcup \{A_c \mid c \text{ is a } C\text{-perimetric cell of } \mathfrak{R}\}\$ and we denote by Δ_C the closed disk bounded by K_C that is contained in Δ . A cell c of \mathfrak{R} is called *C*-internal if $c \subseteq \Delta_C$ and is called *C*-external if $\Delta_C \cap c = \emptyset$. Notice that the cells of \mathfrak{R} are partitioned into *C*-internal, *C*-perimetric, and *C*-external cells.

Let c be a tidy C-perimetric cell of \mathfrak{R} where $|\tilde{c}| = 3$. Notice that $c \setminus A_c$ has two arcwise-connected components and one of them is an open disk D_c that is a subset of Δ_C . If the closure \overline{D}_c of D_c contains only two points of \tilde{c} then we call the cell c C-marginal.

Influence. For every \mathfrak{R} -normal cycle C of $\operatorname{compass}_{\mathfrak{R}}(W)$ we define the set

influence_{\Re}(C) = { $\sigma(c) \mid c$ is a cell of \Re that is not C-external}.

A wall W' of $\operatorname{compass}_{\mathfrak{R}}(W)$ is \mathfrak{R} -normal if D(W') is \mathfrak{R} -normal. Notice that every wall of W (and hence every subwall of W) is an \mathfrak{R} -normal wall of $\operatorname{compass}_{\mathfrak{R}}(W)$. We denote by $S_{\mathfrak{R}}(W)$ the set of all \mathfrak{R} -normal walls of $\operatorname{compass}_{\mathfrak{R}}(W)$. Given a $W' \in S_{\mathfrak{R}}(W)$ and a cell c of \mathfrak{R} we say that c is W'perimetric/internal/external/marginal if c is D(W')-perimetric/internal/external/marginal (see Figure 5 for an example). We also use $K_{W'}$, $\Delta_{W'}$, influence $\mathfrak{R}(W')$ as shortcuts for $K_{D(W')}$, $\Delta_{D(W')}$, influence $\mathfrak{R}(D(W'))$.

Regularity. Let (W, \mathfrak{R}) be a flatness pair of a graph G. We call a flatness pair (W, \mathfrak{R}) of a graph G regular if none of its cells is W-external, W-marginal, or untidy. Notice that the flatness pair of Figure 4 is not regular (for an example of a regular flatness pair of a graph that is a modification of the one in Figure 4, see Figure 9). The notion of regularity has been defined in [46] and will be useful in Subsection 5.3. In fact,

regularity permits the definition of a "well-allinged" Δ -embedded representation of the \Re -compass that will be valuable in the proofs of Subsection 5.4. The precise definition of the notion of well-allinged flatness pairs is given in Subsection 5.3.

The next result has been proved in [46]. It can be seen as a version of the Flat Wall Theorem incorporating the concept of regularity, which is necessary for our proofs.

Theorem 5. There exist two functions $f_2, f_3 : \mathbb{N} \to \mathbb{N}$ such that for every graph G, every odd integer $r \ge 3$, and every $q \in \mathbb{N}_{>1}$, one of the following is true:

- K_q is a minor of G,
- $\mathsf{tw}(G) \leq f_2(q) \cdot r, \ or$
- there exist a set $A \subseteq V(G)$, where $|A| \leq f_3(q)$, and a regular flatness pair (W, \mathfrak{R}) of $G \setminus A$ of height r.

Moreover, $f_2(q) = 2^{O(q^2 \log q)}$ and $f_3(q) = O(q^{24})$.

Tilts of flatness pairs. Let (W, \mathfrak{R}) and $(\tilde{W}', \tilde{\mathfrak{R}}')$ be two flatness pairs of a graph G and let $W' \in S_{\mathfrak{R}}(W)$. We also assume that $\mathfrak{R} = (X, Y, P, C, \Gamma, \sigma, \pi)$ and $\tilde{\mathfrak{R}}' = (X', Y', P', C', \Gamma', \sigma', \pi')$. We say that $(\tilde{W}', \tilde{\mathfrak{R}}')$ is a W'-tilt of (W, \mathfrak{R}) if

- $\tilde{\mathfrak{R}}'$ does not have \tilde{W}' -external cells,
- \tilde{W}' is a tilt of W',
- the set of \tilde{W}' -internal cells of $\tilde{\mathfrak{R}}'$ is the same as the set of W'-internal cells of \mathfrak{R} and their images via σ' and σ are also the same,
- compass_{\mathfrak{H}'} (\tilde{W}') is a subgraph of \bigcup influence_{\mathfrak{H}} (W'), and
- if c is a cell in $C(\Gamma') \setminus C(\Gamma)$, then $|\tilde{c}| \leq 2$.

The next observation follows from the definitions of regular flatness pairs and tilts.

Observation 6. If (W, \mathfrak{R}) is a regular flatness pair, then for every $W' \in S_{\mathfrak{R}}(W)$, every W'-tilt of (W, \mathfrak{R}) is also regular.

We need one more observation, which follows from the third item above and the fact that the cells corresponding to flaps containing a central vertex of W' are all internal (recall that the height of a wall is always at least three).

Observation 7. The central vertices of W' belong to every W'-tilt of (W, \mathfrak{R}) .

The need to define W'-tilts of flatness pairs emerges from the fact that not every subwall W' of a flat wall W is necessarily flat, recently observed in [46]. The next proposition, proved in [46], suggests that there is always a slight modification of W' in the \Re -compass of W that is indeed a flat wall. This "tilt" preserves the internal cells, and therefore the "essential" part of the influence of W'. That way, it permits us to define a notion of compass relative to a subwall of a flat wall.

Proposition 8. For every flatness pair (W, \mathfrak{R}) of a graph G and every $W' \in S_{\mathfrak{R}}(W)$, there exists a flatness pair $(\tilde{W}', \tilde{\mathfrak{R}}')$ of G that is a W'-tilt of (W, \mathfrak{R}) .

4.5 Homogeneous walls

Homogeneous walls were a basic ingredient of the seminal algorithm of Robertson and Seymour for the DISJOINT PATHS problem in [43]. This algorithm introduced the *Irrelevant Vertex Technique* that consisted in the identification of a vertex in an instance of the DISJOINT PATHS problem that is *irrelevant* in the sense that its removal does not change the YES/NO-status of the instance. The notion of wall homogeneity was given in [43] and was based on the concept of the *vision* of an "internal" flap of a flat wall. It was proved in [44] that the central vertices of a sufficiently big homogenous flat wall are indeed *irrelevant* with respect to the DISJOINT PATHS problem and therefore they could be safely discarded. Our results are following the same technique. However, we need an alternative notion of homogeneity that we introduce in this subsection.

Let G be a graph, let $A \subseteq V(G)$, and let (W, \mathfrak{R}) be a flatness pair of $G \setminus A$, where $\mathfrak{R} = (X, Y, P, C, \Gamma, \sigma, \pi)$ and (Γ, σ, π) is an Ω -rendition of G[Y]. Recall that $\Gamma = (U, N)$ is a Δ -painting of the closed disk Δ .

Augmented flaps. For each flap $F \in \text{flaps}_{\mathfrak{R}}(W)$ we consider a labeling $\ell_F : \partial F \to \{1, 2, 3\}$ such that the set of labels assigned by ℓ_F to ∂F is one of $\{1\}$, $\{1, 2\}$, $\{1, 2, 3\}$. We also consider a bijection $\rho_A : A \to [a]$, where a = |A|. The labelings in $\mathcal{L} = \{\ell_F \mid F \in \text{flaps}_{\mathfrak{R}}(W)\}$ and the labeling ρ_A will be useful for defining a set of boundaried graphs that we will call augmented flaps. We first need some more definitions.

Given a flap $F \in \mathsf{flaps}_{\mathfrak{R}}(W)$, we define an ordering $\Omega(F) = (x_1, \ldots, x_q)$, with $q \leq 3$, of the vertices of ∂F so that

• (x_1, \ldots, x_q) is a counter-clockwise cyclic ordering of the vertices of ∂F as they appear in the corresponding cell of $C(\Gamma)$. Notice that this cyclic ordering is significant only when $|\partial F| = 3$, in the sense that (x_1, x_2, x_3) remains invariant under shifting, i.e., (x_1, x_2, x_3) is the same as (x_2, x_3, x_1) but not under inversion, i.e., (x_1, x_2, x_3) is not the same as (x_3, x_2, x_1) , and

• for
$$i \in [q], \ell_F(x_i) = i$$
.

Notice that the second condition is necessary for completing the definition of the ordering $\Omega(F)$, and this is the reason why we set up the labelings in \mathcal{L} .

For each $F \in \mathsf{flaps}_{\mathfrak{R}}(W)$ with $t_F = |\partial F|$, we fix $\rho_F : \partial F \to [a+1, a+t_F]$ such that $(\rho_F^{-1}(a+1), \ldots, \rho_F^{-1}(a+t_F)) = \Omega(F)$. Also, we define the boundaried graph

$$\mathbf{F}^{A} = (G[A \cup F], A \cup \partial F, \rho_{A} \cup \rho_{F})$$
(3)

and we denote by F^A the underlying graph of \mathbf{F}^A . We call \mathbf{F}^A an *augmented flap* of the flatness pair (W, \mathfrak{R}) of $G \setminus A$ in G.

Palettes and homogeneity. For each cycle C of W, we define (A, ℓ) -palette $(C) = \{\ell \text{-} \mathsf{folio}(\mathbf{F}^A) \mid F \in \mathsf{influence}_{\mathfrak{R}}(C)\}$. We say that the flatness pair (W, \mathfrak{R}) of $G \setminus A$ is ℓ -homogeneous with respect to the pair (G, A) if every internal brick B of W (seen as a cycle of W) has the same (A, ℓ) -palette.

Apex-wall triples. Let G be a graph, let $A \subseteq V(G)$ with $|A| \leq a$, and let (W, \mathfrak{R}) be a regular flatness pair of $G \setminus A$ such that W has height r and is ℓ -homogenous with respect to (G, A) for some $\ell \in \mathbb{N}$. We call such a triple (A, W, \mathfrak{R}) an (a, r, ℓ) -apex-wall triple of G.

The next proposition, proved in [46], implies that it is possible to find an ℓ -homogeneous flat wall inside the compass of a sufficiently big flat wall.

Proposition 9. There exists a function $f_4 : \mathbb{N}^2 \to \mathbb{N}$ such that if $r \in \mathbb{N}_{\geq 3}$, G is a graph, $A \subseteq V(G)$, and (W, \mathfrak{R}) is a flatness pair of $G \setminus A$ of height $f_4(r, w)$, where $w = f_1(|A| + 3, \ell)$, then W contains some subwall W' of height r such that every W'-tilt of (W, \mathfrak{R}) is ℓ -homogeneous with respect to (G, A). Moreover, $f_4(r, w) = \mathcal{O}(r^w)$.

4.6 A parameter for affecting flat walls

We proceed to define a graph parameter that will be useful for our proofs. We prove that it satisfies some properties related to Bidimensionality theory [17, 18, 21] that will be used later in Subsection 6.2.

Let G be a graph and let (A, W, \mathfrak{R}) be an (a, r, ℓ) -apex-wall triple of G. We say that S affects (A, W, \mathfrak{R}) if $N_G[V(\mathsf{compass}_{\mathfrak{R}}(W))] \cap (S \setminus A) \neq \emptyset$. For $a, r, \ell \in \mathbb{N}$, we define

 $\mathbf{p}_{a,r,\ell}(G) = \min\{k \mid \exists S \subseteq V(G) : |S| \le k \land S \text{ affects every } (a,r,\ell)\text{-apex-wall triple of } G\}.$

Using Theorem 5, Proposition 8, and Proposition 9, we prove that the above parameter grows quadratically with its treewidth.

Lemma 10. There is a function $f_5 : \mathbb{N}^3 \to \mathbb{N}$ such that if $q, r, \ell \in \mathbb{N}_{>0}$, and G is a K_q -minor-free graph, then $\mathsf{tw}(G) \leq f_5(q, r, \ell) \cdot \max\left\{1, \sqrt{\mathbf{p}_{f_3(q), r, \ell}(G)}\right\}$. In particular, one may choose $f_5(q, r, \ell) = \mathcal{O}(f_2(q) \cdot r^w)$, where $w = f_1(f_3(q) + 3, \ell) = 2^{2^{\mathcal{O}((q^{24} + \ell) \cdot \log(q^{24} + \ell))}}$.

Proof. The lemma follows easily if we prove that, for every positive integer p, tw $(G) > f_2(q) \cdot ((r^w + 2) \cdot [\sqrt{p+1}] + 2)$ implies that $\mathbf{p}_{f_3(q),r,\ell}(G) > p$.

By Theorem 5, it follows that, for any $q, r, \ell \in \mathbb{N}_{>0}$, if G is K_q -minor-free and $\mathsf{tw}(G) > f_2(q) \cdot ((r^w + 2) \cdot \lceil \sqrt{p+1} \rceil + 2)$, then G contains some vertex set A such that $|A| \leq f_3(q)$ and $G \setminus A$ has a regular flatness pair (W, \mathfrak{R}) of height $(r^w + 2) \cdot \lceil \sqrt{p+1} \rceil + 2$. Let $\hat{W}_1, \ldots, \hat{W}_{p+1} \in S_{\mathfrak{R}}(W)$ be a collection of p+1 pairwise disjoint subwalls of W, each of height r^w , such that there are no two vertices w_1, w_2 in W of face-distance at most one such that w_1 belongs to some \hat{W}_i and w_2 belongs either in D(W) or in some other $\hat{W}_{i'}, i' \neq i$.

By Proposition 8, for every $i \in [p + 1]$, there is a \hat{W}_i -tilt of (W, \mathfrak{R}) that we denote by (W_i, \mathfrak{R}_i) . Since (W, \mathfrak{R}) is regular, Observation 6 implies that (W_i, \mathfrak{R}_i) is also regular for every $i \in [p + 1]$. Moreover, by the definition of a \hat{W}_i -tilt, it follows that, for every two distinct $i, i' \in [p + 1]$, $N_G[V(\text{compass}_{\mathfrak{R}_i}(W_i))] \cap N_G[V(\text{compass}_{\mathfrak{R}_i}(W_i))] \subseteq A$. Note that this is correct because the face-distance demand leaves a "buffer" among the flat walls and the perimeter of W to guarantee that the neighborhoods of their compasses do not intersect, except possibly at apex vertices. Let $w = f_1(|A| + 3, \ell)$. From Proposition 9, for each regular flatness pair (W_i, \mathfrak{R}_i) of $G \setminus A$ there is a subwall \hat{W}'_i of height r such that every regular \hat{W}'_i -tilt of (W, \mathfrak{R}) is ℓ -homogeneous with respect to (G, A). We denote this \hat{W}'_i -tilt by (W'_i, \mathfrak{R}'_i) and we conclude that $(A, W'_i, \mathfrak{R}'_i)$ is an (a, r, ℓ) -apex-wall triple of G, for every $i \in [p + 1]$. As before, $N_G[V(\text{compass}_{\mathfrak{R}'_i}(W'_i))] \cap N_G[V(\text{compass}_{\mathfrak{R}'_i}(W'_i))] \subseteq A$ for every $i, i' \in [p + 1], i \neq i'$. Therefore, every set $S \subseteq V(G)$ affecting every $(f_3(q), r, \ell)$ -apex-wall triple of G needs to contain at least one vertex from each of the sets $\{N_G[V(\text{compass}_{\mathfrak{R}}(W_i))] \mid i \in [p]\}$, implying that $\mathbf{p}_{f_3(q), r, \ell}(G) > p$.

We now prove that the parameter $\mathbf{p}_{a,r,\ell}$ is *separable*, that is, that when considering a separation of a graph, the value of the parameter is "evenly" split along both sides of the separation, possibly with an offset bounded by the order of the separation.

Lemma 11. Let $a, r, \ell \in \mathbb{N}$, let G be a graph, and let $S \subseteq V(G)$ such that S affects every (a, r, ℓ) -apex-wall triple of G. Then, for every separation (L, R) of S in G, the set $L \cap (R \cup S)$ affects every (a, r, ℓ) -apex-wall triple of G[L].

Proof. Suppose for contradiction that (A, W, \mathfrak{R}) is an (a, r, ℓ) -apex-wall triple of G[L] that is not affected by $L \cap (R \cup S)$. In particular, it holds that $V(\mathsf{compass}_{\mathfrak{R}}(W)) \subseteq L \setminus R$. Since by assumption (A, W, \mathfrak{R}) is affected by S but not by $L \cap (R \cup S)$, there should exist a vertex $v \in S \cap (R \setminus L)$ with a neighbor in $V(\mathsf{compass}_{\mathfrak{R}}(W)) \subseteq L \setminus R$, contracting the hypothesis that (L, R) is a separation of G. \Box

5 Finding an irrelevant vertex

In this section we show how to find inside a sufficiently large flat wall of a boundaried graph $\mathbf{G} = (G, B, \rho)$ a flat subwall whose compass is "irrelevant" with respect to the presence of a graph H in \mathcal{F} as a minor. Here the term "irrelevant" is not only related to G but to every graph $\mathbf{K} \oplus \mathbf{G}$ that can be obtained by gluing \mathbf{G} with another boundaried graph \mathbf{K} . For this we need a stronger notion of irrelevancy, defined in Subsection 5.4, that takes into account only the "essential part" of a topological minor model (M, T) of Hthat is "invading" \mathbf{G} .

We start in Subsection 5.1, by detecting in every wall a railed annulus. This structure, introduced in [28] and reused later in [24, 25], turns out to be quite handy in order to guarantee a "taming property" of topological minor models (cf. Proposition 13). In Subsection 5.2 we first use graph drawing tools to prove that we can assume that our model is embedded "nicely" inside a railed annulus, in the sense that certain vertices are sufficiently pairwise far apart (cf. Lemma 16); this will be helpful in order to reroute the model of every "invading" topological minor model (M, T) of H. With the help of Proposition 13, we prove (cf. Theorem 17) that, given a partially disk-embedded graph that contains a railed annulus, the topological minor model (M, T) of a graph H can be rerouted so to obtain another topological minor model that can be contracted back to H and such that a "large enough" central region of the railed annulus is avoided. The rerouting of (M, T) will be done so that a prescribed subset of degree-3 vertices of the original model will not be affected by contractions.

Once we have all the above ingredients, we consider in Subsection 5.4 a boundaried graph $\mathbf{G} = (G, B, \rho)$ and an apex wall triple (A, W, \mathfrak{R}) that is not affected by B and we show, in Theorem 23, how every topological minor model (M, T) of a graph H in G can be rerouted away from the compass of the central subwall W' of W. This will permit us later to declare the whole compass of W' irrelevant and rule out the possibility that W has size exceeding some function depending on the "intrusion" of H in \mathbf{G} . The proof of Theorem 23 is the most technical part of this paper. For this, we define an appropriate "flat" representation of the \mathfrak{R} -compass of W, called its *leveling*, and a representation of the wall W in the leveling that is "well-aligned". This wellalignment property, defined in Subsection 5.3, emerges from the regularity of the flatness pair (W, \mathfrak{R}) and permits the representation of (M, T) by a topological minor model (\tilde{M}, \tilde{T}) of the leveling, accompanied with a suitable encoding of the parts of (M, T) that have been suppressed by the leveling. This will permit us to obtain, using Theorem 17, a rerouting (\hat{M}, \hat{T}) of (\tilde{M}, \tilde{T}) inside the leveling. Finally, using the homogeneity property, we will translate back (\hat{M}, \hat{T}) to a rerouting of (M, T) that will avoid the compass of the central subwall W'.

5.1 A lemma for model taming

We introduce the concept of a railed annulus and present the main combinatorial result of [25].

Railed annuli. Let G be a partially Δ -embedded graph and let $\mathcal{C} = [C_1, \ldots, C_r], r \geq 2$, be a collection of vertex-disjoint cycles of the compass of G. We say that the sequence \mathcal{C} is a Δ -nested sequence of cycles of G if every C_i is the boundary of an open disk D_i of Δ such that $\Delta \supseteq D_1 \supseteq \cdots \supseteq D_r$. From now on, each Δ -nested sequence $\mathcal{C} = [C_1, \ldots, C_r]$ will be accompanied with the sequence $[D_1, \ldots, D_r]$ of the corresponding open disks as well as the sequence $[\overline{D}_1, \ldots, \overline{D}_r]$ of their closures. Given $x, y \in [r]$ with $x \leq y$, we call the set



Figure 6: An example of a (5, 8)-railed annulus and its inner disk D_5 .

 $D_x \setminus D_y(x, y)$ -annulus of \mathcal{C} and we denote it by $\operatorname{ann}(\mathcal{C}, x, y)$. Finally we say that $\operatorname{ann}(\mathcal{C}, 1, r)$ is the annulus of \mathcal{C} and we denote it by $\operatorname{ann}(\mathcal{C})$.

Let $r \in \mathbb{N}_{\geq 3}$ and $q \in \mathbb{N}_{\geq 3}$ with r odd. An (r,q)-railed annulus of a Δ -partially-embedded graph G is a pair $\mathcal{A} = (\mathcal{C}, \mathcal{P})$ where $\mathcal{C} = [C_1, \ldots, C_r]$ is a Δ -nested collection of cycles of G and $\mathcal{P} = [P_1, \ldots, P_q]$ is a collection of pairwise vertex-disjoint paths in G, called *rails*, such that

- for every $j \in [q], P_j \subseteq \operatorname{ann}(\mathcal{C})$, and
- for every $(i, j) \in [r] \times [q], C_i \cap P_j$ is a non-empty path that we denote by $P_{i,j}$.

See Figure 6 for an example of a (5, 8)-railed annulus. The following proposition states that large railed annuli can be found inside a modestly larger wall and will be used in the next section. A similar (but less precise) statement can be found in [28].

Proposition 12. If $x, z \ge 3$ are odd integers, $y \ge 1$, and W is an $odd(2x + max\{z, \frac{y}{4} - 1\})$ -wall, then

- there is a collection \mathcal{P} of y paths in W such that if C is the collection of the first x layers of W, then $(\mathcal{C}, \mathcal{P})$ is an (x, y)-railed annulus of W where the first cycle of C is the perimeter of W, and
- the open disk defined by the x-th cycle of C contains the vertices of the compass of the central z-subwall of W.

Proof. Let $C = \{C_1, \ldots, C_x\}$ be the collection of the first x layers of W. Notice that the open disk defined by C_x contains the vertices of the compass of the central w'-subwall, where $w' := w - 2x = \operatorname{odd}(\max\{z, \frac{y}{4} - 1\}) \ge z$. On the other hand, W contains a collection \mathcal{P} of at least 2w' + 2(w' + 2) pairwise vertex-disjoint paths from C_1 to C_x . Since $2w' + 2(w' + 2) = 4w' + 4 \ge y$, the pair $(\mathcal{C}, \mathcal{P})$ is an (x, y)-railed annulus of W where the first cycle of \mathcal{C} is the perimeter of W (see Figure 7).

We define the annulus of $\mathcal{A} = (\mathcal{C}, \mathcal{P})$ as the annulus of \mathcal{C} . We call C_1 and C_r the outer and the inner cycle of \mathcal{A} , respectively. Also, if $(i, i') \in [r]^2$ with i < i' then we define $\mathcal{A}_{i,i'} = ([C_i, \cdots, C_{i'}], \mathcal{P} \cap \operatorname{ann}(\mathcal{C}, i, i'))$.



Figure 7: A visualization of the proof of Proposition 12 in the 33-wall where x = 9, y = 64, and z = 15. For simplicity, the paths in $\{P_{i,j} \mid (i,j) \in [33]^2\}$ are depicted as vertices.

The union-graph of an (r, q)-railed annulus $\mathcal{A} = (\mathcal{C}, \mathcal{P})$ is defined as $G(\mathcal{A}) := (\bigcup_{i \in [r]} C_i) \cup (\bigcup_{i \in [q]} P_i)$. Clearly, $G(\mathcal{A})$ is a planar graph and we always assume that its infinite face is the one whose boundary is the first cycle of \mathcal{C} .

Let \mathcal{A} be a (r, q)-railed annulus of a partially Δ -embedded graph G. Let r = 2t + 1, for some $t \geq 0$. Let also $s \in [r]$ where s = 2t' + 1, for some $0 \leq t' \leq t$. Given some $I \subseteq [q]$, we say that a subgraph M of G is (s, I)-confined in \mathcal{A} if

$$M \cap \operatorname{ann}(\mathcal{C}, t - t', t + t') \subseteq \bigcup_{i \in I} P_i.$$

The following proposition has been recently proved by Golovach et al. [25, Theorem 2.1], where it has been dubbed as the "Model Taming Lemma".

Proposition 13. There exist two functions $f_6, f_7 : \mathbb{N}_{\geq 0} \to \mathbb{N}_{\geq 0}$ such that the images of f_7 are even and such that if

- s is a positive odd integer,
- *H* is a graph on at most ℓ edges,
- G is a Δ -partially-embedded graph,
- $\mathcal{A} = (\mathcal{C}, \mathcal{P})$ is an (r, q)-railed annulus of G, where $r = f_7(\ell) + 2 + s$ and $q \ge 5/2 \cdot f_6(\ell)$,
- (M,T) is a topological minor model of H in G such that $T \cap \operatorname{ann}(\mathcal{A}) = \emptyset$, and
- $I \subseteq [q]$ where $|I| > f_6(\ell)$,

then G contains a topological minor model (\tilde{M}, \tilde{T}) of H in G such that

1.
$$T = T$$

- 2. \tilde{M} is (s, I)-confined in A, and
- 3. $\tilde{M} \setminus \operatorname{ann}(\mathcal{A}) \subseteq M \setminus \operatorname{ann}(\mathcal{A}).$

Moreover $f_7(\ell) = \mathcal{O}((f_6(\ell))^2).$

Remark 14. It is worth mentioning here that the function $f_6(\ell)$ of Proposition 13 depends on the constants involved in the Unique Linkage Theorem [31,44]; see Subsection B.2 for a more detailed discussion. At this point we just remark that, according to the results of Adler and Krause [2], we have that $f_6(\ell) = 2^{\Omega(\ell)}$. This permits us to henceforth make the (generous) assumption that $\ell = \mathcal{O}(f_6(\ell))$.

5.2 Model rerouting in partially disk-embedded graphs

Using classic results on how to optimally draw planar graphs of maximum degree three into grids (see e.g., [29]) one may easily derive the following.

Proposition 15. There exists a function $f_8 : \mathbb{N} \to \mathbb{N}$ such that for every ℓ -vertex planar graph H with maximum degree three there is a tm-pair (M,T) of the $(f_8(\ell) \times f_8(\ell))$ -grid, denoted by Γ , that is a topological minor model of H in Γ . Moreover, it holds that $f_8(\ell) = \mathcal{O}(\ell)$.

Let Γ be an $(r \times r)$ -grid for some $r \geq 3$. We see a Γ -grid as the union of r horizontal paths and r vertical paths. Given an $i \in \lfloor \frac{r}{2} \rfloor$, we define the *i*-th layer of Γ recursively as follows: the first layer of Γ is its perimeter, while, if $i \geq 2$, the *i*-th layer of Γ is the perimeter of the $(r - 2(i - 1) \times r - 2(i - 1))$ -grid created if we remove from Γ its i - 1 first layers. When we deal with a $(r \times r)$ -grid Γ , we always consider its embedding where the infinite face is bounded by the first layer of Γ .

Safely arranged models. Let G be a plane graph. Given two subgraphs of G, we define their *face-distance* as the minimum face-distance between two of their vertices. We denote by $\mathbf{F}_{G}^{(i)}(x)$ the set of all vertices of G that are within face-distance at most i from vertex x.

Given a $c \ge 0$ and a tm-pair (M,T) of G, we say that (M,T) is safely c-dispersed in G if

- every two distinct vertices $t, t' \in T$ are within face-distance at least 2c + 1 in G, and
- for every $t \in T$ of degree d in M, the graph $M[\mathbf{F}_G^{(c)}(t) \cap V(M)]$ consists of d paths with t as a unique common endpoint.

With Proposition 15 at hand, we can prove the following useful lemma.

Lemma 16. There exists a function $f_9 : \mathbb{N}^3 \to \mathbb{N}$ such that the following holds. Let $c, r, r', \ell \in \mathbb{N}, r' \leq r, H$ be a *D*-embedded $(\ell + r')$ -vertex graph, and $Z := \{z_1, \ldots, z_{r'}\} \subseteq V(H)$ such that

- the vertices of H have degree at most three,
- Z is an independent set of H,
- all vertices of Z have degree one in H,
- $bd(D) \cap H = Z$, and
- $(z_1, \ldots, z_{r'})$ is the cyclic ordering of the vertices of Z as they appear in the boundary of D.

Let also G be a Δ -embedded graph, $\mathcal{A} = (\mathcal{C}, \mathcal{P})$ be an (x, y)-railed annulus of G, where x, y are integers such that $\min\{x, y\} \geq f_9(c, r, \ell) \geq r$, and where $\mathcal{C} = [C_1, \ldots, C_x]$ and $\mathcal{P} = [P_1, \ldots, P_y]$, w_i be the endpoint of P_i that is contained in C_1 , for $i \in [r]$, and $I := \{i_1, \ldots, i_{r'}\} \subseteq [r]$. Then the union-graph $G(\mathcal{A})$ of \mathcal{A} contains a tm-pair (M, T) that is a topological minor model of H in $G(\mathcal{A})$ such that

- for each $j \in [r']$, $\sigma_{M,T}(z_j) = w_{i_j}$,
- the tm-pair (M,T) is safely c-dispersed in the union graph $G(\mathcal{A})$, and

• none of the vertices of $T \setminus \{w_{i_1}, \ldots, w_{i_{r'}}\}$ is within face-distance less than c from some vertex in C_1 or in C_r .

Moreover, it holds that $f_9(c, r, \ell) = \mathcal{O}(cr(\ell + r)).$

Proof. Using H, we construct a new graph H' as follows: consider a copy \tilde{H} of H where the copy of z_i in \tilde{H} is denoted by \tilde{z}_i , for each $i \in [r']$. We take the disjoint union of H and \tilde{H} , add the edges $\{z_1, z_2\}, \ldots, \{z_{r'}, z_1\}$, forming a cycle C, subdivide the edges in C, add the edges $\{\tilde{z}_1, \tilde{z}_2\}, \ldots, \{\tilde{z}_r, \tilde{z}_1\}$, forming a cycle \tilde{C} , subdivide the edges in C, add the edges $\{\tilde{z}_1, \tilde{z}_2\}, \ldots, \{\tilde{z}_r, \tilde{z}_1\}$, forming a cycle \tilde{C} , subdivide the edges in \tilde{C} , and, given that, for $i \in [r']$, x_i (resp. \tilde{x}_i) is the vertex created after the subdivision of $\{z_i, z_{i+1}\}$ (resp. $\{\tilde{z}_i, \tilde{z}_{i+1}\}$) (here r' + 1 is interpreted as 1), add the edges $\{x_1, \tilde{x}_1\}, \ldots, \{x_r, \tilde{x}_{r'}\}$. The resulting graph H' has $2(\ell + r') \leq 2(\ell + r)$ vertices, is planar, and has maximum degree three. Let $s = f_8(2(\ell + r))$. By Proposition 15, there is a tm-pair (M', T') of the $(s \times s)$ -grid Γ that is a topological minor model of H' in Γ . We now subdivide $\overline{r} = 2(c+1)r$ times each of the edges of Γ and see the resulting graph Γ' as a subgraph of a $((\overline{r}+1)(s+2) \times (\overline{r}+1)(s+2))$ -grid Γ'' in a way that none of the \overline{r} first layers of Γ'' intersects Γ' . By also subdividing \overline{r} times each of the edges of M' we construct a tm-pair (M'', T') of Γ'' that is a topological minor model of H' in Γ'' .

Let w_1, \ldots, w_r be the first r vertices of the lower path of Γ'' . Recall that H is a subgraph of H', therefore we can define $M = \sigma_{M',T'}(H)$. Let also $T = \sigma_{M',T'}(V(H))$. Notice that (M,T) is a tm-pair of Γ'' that is a topological minor model of H. Let $\hat{z}_i = \sigma_{M,T}(z_i), i \in [r']$. We make two observations about the position of these vertices in Γ'' . The first is that, because of the construction of $H', \hat{z}_1, \ldots, \hat{z}_{r'}$ appear, in this ordering, on a cycle of Γ' bounding a closed disk, say Δ , that contains the whole M. The second is that, as M is a subgraph of Γ' , each pair $\hat{z}_i, \hat{z}_j, i \neq j$, is at distance at least $\overline{r} + 1$ in the graph $\Gamma''' := \Gamma'' \setminus \operatorname{int}(\Delta)$. It is now easy to observe that the two previous observations permit to find in Γ''' pairwise disjoint paths joining \hat{z}_j with w_{ij} , for $j \in [r']$. By adding these paths in M and including in T the set $\{w_{i_1}, \ldots, w_{i_{r'}}\}$, we construct a tm-pair (M,T) that is a topological minor model of H in Γ' such that for each $j \in [r']$, the function $\sigma_{M,T}$ maps the vertex z_j to w'_{ij} and the intersection of V(M) and the upper path of Γ'' is empty. Moreover, as we applied at least $\overline{r} = (c+2)r$ subdivisions, it also holds that the set T is safely 2c-dispersed in Γ'' . Moreover, it is easy to observe that none of the vertices of $T \setminus \{w_{i_1}, \ldots, w_{i_{r'}}\}$ is within face-distance less than c from some vertex in the perimeter of Γ'' .

Consider now an (x, y)-railed annulus $(\mathcal{C}, \mathcal{P})$ of some Δ -embedded graph G, with $\min\{x, y\} \geq q$, where $q = (\overline{r} + 1)(s + 2)$. Let $\mathcal{C} = [C_1, \ldots, C_x]$ and $\mathcal{P} = [P_1, \ldots, P_y]$ as in the statement of the lemma. Let also $\tilde{\Gamma} = G(\mathcal{A})$. For every $i \in [q]$, we define $F_{\mathcal{A}}^{(i)}$ as the edge set of the unique path in C_i with one endpoint in $P_{i,q}$ and the other in $P_{i,1}$, that does not contain internal vertices of the paths $P_{i,q}$ or $P_{i,1}$, and does not contain any vertex from P_2 . We denote by $F_{\mathcal{A}}^e$ (resp. $F_{\mathcal{A}}^\vee$) the set of all edges (resp. internal vertices) of the paths $F_{\mathcal{A}}^{(i)}$, $i \in [q]$. Notice that the grid Γ'' occurs from $(\tilde{\Gamma} \setminus F_{\mathcal{A}}^e) \setminus F_{\mathcal{A}}^\vee$ if, for every $(i, j) \in [q]^2$, we contract the path $P_{i,j}$ defined by the intersection of the *i*-th horizontal path and the *j*-th vertical path of W. It is easy to see that if in $\tilde{\Gamma}$ we uncontract each vertex, say (i, j) of M to the path $P_{i,j}$, one can transform (M, T) to a tm-pair of W that is a topological minor model of H in W and additionally, for each $i \in [r']$, the function $\sigma_{M,T}$ maps the vertex z_j to w_{ij} . This implies the first condition of the lemma. The second condition follows directly from the fact that the pair (M, T) was already safely 2*c*-dispersed before applying the uncontractions and such uncontractions cannot reduce the distance to more than half of it. The third condition is also an obvious consequence of the uncontraction procedure. Therefore, the lemma holds if we set $f_9(c, r, \ell) := q = \mathcal{O}(cr(\ell + r))$.

Let G be a partially Δ -embedded graph and let $\mathcal{C} = [C_1, \ldots, C_r]$ be a Δ -nested sequence of cycles of G and let $[D_1, \ldots, D_r]$ (resp. $[\overline{D}_1, \ldots, \overline{D}_r]$) be the sequences of the corresponding open (resp. closed) disks.

Let also (M, T) be a tm-pair of G and $p \in [r]$. We define the *p*-crop of (M, T) in \mathcal{C} , denoted by $(M, T) \cap \overline{D}_p$, as the tm-pair (M', T') where $M' = M \cap \overline{D}_p$ and $T' = (T \cap \overline{D}_p) \cup (V(C_p \cap M))$.

Given a graph H a set $Q \subseteq V(H)$ and a graph G, we say that $\phi : V(H) \to 2^{V(G)}$ is a *Q*-respecting contraction-mapping of H to G if

- $\bigcup_{x \in V(H)} \phi(x) = V(G),$
- for every $x, y \in V(H)$, if $x \neq y$ then $\phi(x) \cap \phi(y) = \emptyset$,
- for every $x \in V(H)$, $G[\phi(x)]$ is connected,
- for every $\{x, y\} \in E(H)$, $G[\phi(x) \cup \phi(y)]$ is connected, and
- for every $x \in Q$, $|\phi(x)| = 1$.

The critical point in the above definitions is that vertices in Q are not "uncontracted" when transforming H to G.

Intrusion of a topological minor model. Let G be a graph, let $S \subseteq V(G)$, and let (M,T) be a tm-pair of G. We define the S-intrusion of (M,T) in G as the maximum value between $|S \cap T|$ and the number of subdivision paths of (M,T) that contain vertices of S. It is important to notice that S can intersect many times a subdivision path of (M,T), however the value of the S-intrusion counts each such a path only once.

Using Proposition 12, Proposition 13, and Lemma 16 we prove the following.

Theorem 17. There exist three functions $f_{10} : \mathbb{N}^2 \to \mathbb{N}$, $f_{11} : \mathbb{N}^2 \to \mathbb{N}$, and $f_{12} : \mathbb{N}^3 \to \mathbb{N}$ such that the following holds. Let $c, \ell \in \mathbb{N}$, $z \geq 3$ be an odd integer, and G be a partially Δ -embedded graph, whose compass contains an $f_{12}(c, z, \ell)$ -wall W with $bd(\Delta)$ as perimeter. Let also $\mathcal{C} = [C_1, \ldots, C_{f_{10}(c,\ell)}]$ be the first $f_{10}(c, \ell)$ -layers of W and $D_1, \ldots, D_{f_{10}(c,\ell)}$ be the open disks of Δ that they define. If (M, T) is a tm-pair of G whose $\Delta \cap V(G)$ -intrusion in G is at most ℓ and Q is a subset of T containing vertices of degree at most three in M, then there is a tm-pair (\hat{M}, \hat{T}) of G and an integer $b \in [f_{10}(c, \ell)]$ such that

- 1. $\hat{M} \setminus D_b$ is a subgraph of $M \setminus D_b$,
- 2. $\operatorname{ann}(\mathcal{C}, b, b + f_{11}(c, \ell) 1) \cap (T \cup \hat{T}) = \emptyset$,
- 3. $(\hat{M}, \hat{T}) \cap \overline{D}_{b+f_{11}(c,\ell)}$ is a tm-pair of W that is safely c-dispersed in W and none of the vertices of $\hat{T} \cap \overline{D}_{b+f_{11}(c,\ell)}$ is within face-distance less than c in W from some vertex of $C_{b+f_{11}(c,\ell)} \cup C_{f_{10}(c,\ell)}$,
- $4. \ \hat{M} \cap D_{f_{10}(c,\ell)} = \emptyset,$
- 5. the compass of the central z-subwall of W is a subset of $D_{f_{10}(c,\ell)}$, and
- 6. there is a Q-respecting contraction-mapping of diss(M, T) to diss (\hat{M}, \hat{T}) .

Moreover, it holds that $f_{10}(c,\ell) = \mathcal{O}(c \cdot (f_6(\ell))^3), f_{11}(c,\ell) = \mathcal{O}(c \cdot (f_6(\ell))^2), and f_{12}(c,z,\ell) = \mathcal{O}(c \cdot (f_6(\ell))^3 + z).$

See Figure 8 for an illustration of the conditions guaranteed by Theorem 17.

Proof. Let $r = f_6(\ell) + 1$, $s = \operatorname{odd}(f_9(c, r, 3\ell + r))$, $x' = \operatorname{odd}(f_7(\ell) + 2 + s)$, $y = \max\{s, \lceil 5/2 \cdot f_6(\ell) \rceil\}$, $x = \operatorname{odd}((\ell + 1) \cdot x')$, and $w = \operatorname{odd}(2x + \max\{z, \frac{y}{4} - 1\})$. We will prove the theorem for $f_{10}(c, \ell) = x$, $f_{11}(\ell) = \frac{x'-s}{2}$, and $f_{12}(c, z, \ell) = w$. Let G be a partially Δ -embedded graph, whose compass contains a w-wall W with $\operatorname{bd}(\Delta)$ as perimeter. Let also $\mathcal{C} = [C_1, \ldots, C_x]$ be the first x layers of W and let $[D_1, \ldots, D_x]$ (resp. $[\overline{D}_1, \ldots, \overline{D}_x]$) be the sequences of the corresponding open (resp. closed) disks of Δ bounded by the cycles in \mathcal{C} . From Proposition 12 there is a collection $\mathcal{P} = \{P_1, \ldots, P_y\}$ of paths in W such that $\mathcal{A} = (\mathcal{C}, \mathcal{P})$



Figure 8: A visualization of how a tm-pair (M, T) is rearranged to a new tm-pair (\hat{M}, \hat{T}) as in Theorem 17. The figure depicts in red the part of the tm-pair (\hat{M}, \hat{T}) that intersects the disk Δ . The cycles correspond to the first $f_{10}(c, \ell)$ layers of W. The black vertices are the vertices in Q, while the circled vertices inside the turquoise area are the "new" branch vertices of \tilde{T} that are vertices of W. The "green clouds" are the non-singleton images of the Q-respecting contraction-mapping of diss(M, T) to diss (\hat{M}, \hat{T}) . We stress that in this picture, the way the model enters the turquoise area does not reflect the fact that \tilde{M} is (s, I)-confined in \mathcal{A}' , as it is argued in the proof. We opted not to reflect this fact in the figure as we prioritized the visualization of other, more important, aspects of the proof.

is an (x, y)-railed annulus of W where the outer cycle of C is the perimeter of W and such that the vertices of the compass of the central z-subwall of W belong to D_x , and Property 5 follows.

Let \check{M} be the union of all subdivision paths of (M, T) that intersect $\Delta \cap V(G)$ and let \check{T} be the endpoints of these paths. Moreover, we denote $\check{H} = \operatorname{diss}(\check{M}, \check{T})$ and observe that \check{H} is a subgraph of H. Intuitively, \check{H} is the subgraph of H whose topological minor model (\check{M}, \check{T}) is the part of (M, T) that intersects the closed disk Δ . As the $\Delta \cap V(G)$ -intrusion of (M, T) in G is at most ℓ , the same bound applies to the $\Delta \cap V(G)$ -intrusion of (\check{M}, \check{T}) in G. This in turn implies that $|\check{T} \cap \Delta| \leq \ell$ and that $|E(\check{H})| \leq \ell$.

Since, $x = (\ell + 1) \cdot x'$, there is a $b \leq \ell \cdot x' + 1 \leq x$ such that $\mathbb{A} := \operatorname{ann}(\mathcal{C}, b, b + x' - 1)$ does not contain any vertex of T. We define $T^{\operatorname{out}} = \check{T} \setminus \overline{D}_b$ and $T^{\operatorname{in}} = \check{T} \cap D_{b+x'-1}$. Clearly, $\{T^{\operatorname{out}}, T^{\operatorname{in}}\}$ is a partition of \check{T} .

We set $\mathcal{A}' = ([C_b, \ldots, C_{b+x'-1}], \mathcal{P} \cap \mathbb{A})$. By applying Proposition 13 on $s, \check{H}, g := \ell$, the Δ -boundaried graph G, the (x', y)-railed annulus \mathcal{A}' , the tm-pair (\check{M}, \check{T}) , and the set I = [r], we have that G contains a

topological minor model (\tilde{M}, \tilde{T}) of \check{H} in G such that \tilde{M} is (s, I)-confined in \mathcal{A}' and $\check{M}\setminus \operatorname{ann}(\mathcal{A}') \subseteq \check{M}\setminus \operatorname{ann}(\mathcal{A}')$. We enhance \tilde{M} by adding to it all subdivision paths of (M, T) that are not intersecting Δ . That way, we have that (\tilde{M}, T) is a topological minor model of H in G such that \tilde{M} is (s, I)-confined in \mathcal{A}' and $\tilde{M} \setminus \operatorname{ann}(\mathcal{A}') \subseteq M \setminus \operatorname{ann}(\mathcal{A}')$.

Let $p = b + \frac{x'-s}{2}$ and $q = b + \frac{x'+s}{2} - 1$ and notice that $q \leq x$. We set $\mathbb{A}' := \operatorname{ann}(\mathcal{C}, p, q)$ and we define $\mathcal{A}'' := ([C_p, \ldots C_q], \mathcal{P}')$ where $\mathcal{P}' = \mathcal{P} \cap \mathbb{A}'$. Let $\mathcal{P}' = \{P'_1, \ldots, P'_y\}$. Observe that, from the second property of Proposition 13, the connected components of $\tilde{M} \cap \mathbb{A}'$ are some of the first r paths in \mathcal{P}' . This means that there is a subset of indices $\{i_1, \ldots, i_{r'}\} \subseteq I$ such that $\tilde{M} \cap \mathbb{A}' = P'_{i_1} \cup \cdots \cup P'_{i_{r'}}$. Let $Z = \{z_{i_1}, \ldots, z_{i_{r'}}\}$ be the set of endpoints of the paths $P'_{i_1}, \ldots, P'_{i_{r'}}$ that are contained in C_p .

Let $\tilde{M}^{\text{in}} = \tilde{M} \cap \overline{D}_p$, $\tilde{M}^{\text{out}} = (\tilde{M} \setminus D_p) \setminus E(C_p)$, and observe that $\tilde{M} = \tilde{M}^{\text{in}} \cup \tilde{M}^{\text{out}}$ and that $Z = V(\tilde{M}^{\text{in}}) \cap V(\tilde{M}^{\text{out}})$. Moreover, all vertices of Z have degree one in both \tilde{M}^{in} and \tilde{M}^{out} . Let \tilde{H}^{in} (resp. \tilde{H}^{out}) be the graph obtained from \tilde{M}^{in} (resp. \tilde{M}^{out}) by dissolving all vertices except from those in $T^{\text{in}} \cup Z$ (resp. $T^{\text{out}} \cup Z$). Note that $(\tilde{M}^{\text{in}}, T^{\text{in}} \cup Z)$ (resp. $(\tilde{M}^{\text{out}}, T^{\text{out}} \cup Z)$) is a topological minor model of \tilde{H}^{in} (resp. \tilde{H}^{out}).

Notice that \tilde{H}^{in} has vertex set $T^{\text{in}} \cup Z$ and can be seen as a *D*-embedded graph, for some closed disk *D*, on at most $\ell + r$ edges where $\mathsf{bd}(D) \cap H = Z$ and $(z_{i_1}, \ldots, z_{i_{r'}})$ is the ordering of the vertices of *Z* as they appear in C_p . Observe now that \tilde{H}^{in} can be seen as the contraction of another *D*-embedded graph \hat{H}^{in} with detail at most $3\ell + r$ that has maximum degree at most three. Moreover, we can assume that the vertices of \tilde{H}^{in} that have degree at most three are also vertices of \hat{H}^{in} that are not affected by the contractions while transforming \hat{H}^{out} to \tilde{H}^{out} . This implies that there is a *Q*-respecting contraction-mapping of \tilde{H}^{out} to \hat{H}^{out} . Again, in the embedding of \hat{H}^{in} in *D*, $(z_{i_1} \ldots, z_{i_{r'}})$ is the ordering of the vertices of *Z* as they appear in $\mathsf{bd}(D)$.

Keep in mind that $\tilde{H}^+ := \tilde{H}^{\text{out}} \cup \tilde{H}^{\text{in}}$ is a minor of $\hat{H}^+ := \tilde{H}^{\text{out}} \cup \hat{H}^{\text{in}}$ and that, if we dissolve in \tilde{H}^+ all the vertices in Z, we obtain H. Also let \hat{H} be the graph obtained if we dissolve in \hat{H}^+ all the vertices in Z. Clearly, \hat{H} is a minor of H.

We now apply Lemma 16 for $c, r, r', 3\ell$, the *D*-embedded graph \hat{H}^{in} , the set *Z*, and the (s, y)-railed annulus \mathcal{A}'' of the \overline{D}_p -disk embedded graph $G \cap \overline{D}_p$ and obtain a tm-pair $(\hat{M}^{\text{in}}, \hat{T}^{\text{in}})$ of $G(\mathcal{A}'')$ that is a topological minor model of \hat{H}^{in} and such that for each $j \in [r']$, the function $\sigma_{\hat{M}^{\text{in}}, \hat{T}^{\text{in}}}$ maps vertex z_{i_j} to itself. Notice that $G(\mathcal{A}'')$ is a subgraph of $W \cap \operatorname{ann}(\mathcal{C}, p, q)$. From the second property of Lemma 16, $(\hat{M}^{\text{in}}, \hat{T}^{\text{in}})$ is safely c-dispersed in $W \cap \operatorname{ann}(\mathcal{C}, p, q)$. From the third property of Lemma 16, it follows that none of the vertices of $\hat{T}^{\text{in}} \setminus \{w_{i_1}, \ldots, w_{i_{r'}}\}$ is within face-distance less than c from some vertex of $C_p \cup C_q$ in $W \cap \operatorname{ann}(\mathcal{C}, p, q)$.

We now consider the graph $\hat{M} = \hat{M}^{\text{in}} \cup \tilde{M}^{\text{out}}$. Property 3 follows by the conclusions of the previous paragraph. Moreover, \hat{M} does not intersect D_q and, as $q \leq x$, it neither intersects D_x , hence Property 4 holds. Notice also that $\tilde{M} \setminus \text{ann}(\mathcal{A}') \subseteq M \setminus \text{ann}(\mathcal{A}')$ implies $\tilde{M} \setminus \overline{D}_b \subseteq M \setminus \overline{D}_b$. This along with the fact that $\hat{M} \setminus \overline{D}_b = \tilde{M} \setminus \overline{D}_b$, yield Property 1.

Observe that $(\hat{M}, \hat{T}^{in} \cup T^{out} \cup Z)$ is a topological minor model of \hat{H}^+ , which in turn implies that $(\hat{M}, \hat{T}^{in} \cup T^{out})$ is a topological minor model of \hat{H} . We now set $\hat{T} = \hat{T}^{in} \cup T^{out}$. As there is a Q-respecting contractionmapping of \hat{H}^{out} to \hat{H}^{out} , we also have that there is a Q-respecting contraction-mapping of H = diss(M, T)to $\hat{H} = \text{diss}(\hat{M}, \hat{T})$ and Property 6 holds. As $\hat{T}^{in} \subseteq \text{int}(\text{ann}(\mathcal{A}'') \subseteq D_p = D_{b+f_{11}(\ell)}$ and $T^{out} \subseteq G \setminus \overline{D}_b$, we deduce that $\hat{T} \in G \setminus \text{ann}(\mathcal{C}, b, b + f_{11}(\ell) - 1)$, which together with the fact that $\text{ann}(\mathcal{C}, b, b + x' - 1)$) does not contain any vertex of T, yield Property 2.



Figure 9: A graph G and a regular flatness pair (W, \mathfrak{R}) of G.

To conclude the proof, let us provide upper bounds on the claimed functions. By definition, it holds that

$$\begin{aligned} f_{10}(c,\ell) &= \mathcal{O}(\ell \cdot (f_7(\ell) + f_9(c,f_6(\ell) + 1, 3\ell + f_6(\ell) + 1))), \\ f_{11}(c,\ell) &= \mathcal{O}(f_7(\ell) + f_9(c,f_6(\ell) + 1, 3\ell + f_6(\ell) + 1)), \text{ and} \\ f_{12}(c,z,\ell) &= \mathcal{O}(\ell \cdot (f_7(\ell) + f_9(c,f_6(\ell) + 1, 3\ell + f_6(\ell) + 1)) + \\ &\qquad z + f_6(\ell) + f_9(c,f_6(\ell) + 1, 3\ell + f_6(\ell) + 1)). \end{aligned}$$

Since by Proposition 13 we have that $f_7(\ell) = \mathcal{O}((f_6(\ell))^2)$, by Lemma 16 we have that $f_9(c, r, \ell) = \mathcal{O}(cr(\ell+r))$, and by Remark 14 we may assume that $\ell = \mathcal{O}(f_6(\ell))$, the above can be simplified to

$$f_{10}(c,\ell) = \mathcal{O}(\ell \cdot \left((f_6(\ell))^2 + c \cdot (f_6(\ell))^2\right)) \\ = \mathcal{O}(c \cdot (f_6(\ell))^3), \\ f_{11}(c,\ell) = \mathcal{O}(c \cdot (f_6(\ell))^2), \text{ and} \\ f_{12}(c,z,\ell) = \mathcal{O}(\ell \cdot \left((f_6(\ell))^2 + c \cdot (f_6(\ell))^2\right) + z) \\ = \mathcal{O}(c \cdot (f_6(\ell))^3 + z), \end{cases}$$

and the theorem follows.

5.3 Levelings and well-aligned flatness pairs

Let G be a graph and let (W, \mathfrak{R}) be a flatness pair of G. Let also $\mathfrak{R} = (X, Y, P, C, \Gamma, \sigma, \pi)$, where (Γ, σ, π) is an Ω -rendition of G[Y] and $\Gamma = (U, N)$ is a Δ -painting. The ground set of W in \mathfrak{R} is ground_{\mathfrak{R}} $(W) := \pi(N(\Gamma))$ and we refer to the vertices of this set as the ground vertices of the \mathfrak{R} -compass of W in G. Notice that ground_{\mathfrak{R}}(W) may contain vertices of compass_{\mathfrak{R}}(W) that are not necessarily vertices of W.

In the flatness pairs of Figure 9 and Figure 4 the ground vertices are the vertices on the boundaries of the green cells. (Notice also that the flatness pair in Figure 9 is regular, while the one in Figure 4 is not.)



Figure 10: The leveling $W_{\mathfrak{R}}$ corresponding to the regular flatness pair (W, \mathfrak{R}) in Figure 9. The groundvertices of $W_{\mathfrak{R}}$ are the circled vertices while the flap-vertices are the rhombic vertices. The representation R_W of W in $W_{\mathfrak{R}}$ is obtained from $W_{\mathfrak{R}}$ after removing the black squared vertices. The ground vertices in $\mathsf{bd}(\Delta) \cap W_{\mathfrak{R}} = \mathsf{bd}(\Delta) \cap W_{\mathfrak{R}}$ are depicted in red.

Levelings. We define the \Re -leveling of W in G, denoted by W_{\Re} , as the bipartite graph where one part is the ground set of W in \Re , the other part is a set $\mathsf{vflaps}_{\Re}(W) = \{v_F \mid F \in \mathsf{flaps}_{\Re}(W)\}$ containing one new vertex v_F for each flap F of W in \Re , and, given a pair $(x, F) \in \mathsf{ground}_{\Re}(W) \times \mathsf{flaps}_{\Re}(W)$, the set $\{x, v_F\}$ is an edge of W_{\Re} if and only if $x \in \partial F$. We call the vertices of $\mathsf{ground}_{\Re}(W)$ (resp. $\mathsf{vflaps}_{\Re}(W)$) ground-vertices (resp. flap-vertices) of W_{\Re} . Notice that the incidence graph of the plane hypergraph $(N(\Gamma), \{\tilde{c} \mid c \in C(\Gamma)\})$ is isomorphic to W_{\Re} via an isomorphism that extends π and, moreover, bijectively corresponds cells to flap-vertices. This permits us to treat W_{\Re} as a Δ -embedded graph where $\mathsf{bd}(\Delta) \cap W_{\Re}$ is the set $X \cap Y$. As an example, see Figure 10 for the \Re -leveling corresponding to the flatness pair (W, \Re) in Figure 9.

The following observation is a consequence of the definition of leveling and condition (v) of the tightness property of a rendition.

Observation 18. Let G be a graph, let (W, \mathfrak{R}) be a flatness pair of G, and let $W_{\mathfrak{R}}$ be the leveling of W in G. For every $v_F \in \mathsf{vflaps}_{\mathfrak{R}}(W)$ of degree r in $W_{\mathfrak{R}}$, there exist r internally vertex-disjoint paths in $W_{\mathfrak{R}}$ from v_F to r distinct ground-vertices of $W_{\mathfrak{R}}$ that belong to the perimeter of W.

Well-aligned flatness pairs. We denote by W^{\bullet} the graph obtained from W if we subdivide *once* every edge of W that is short in $\text{compass}_{\mathfrak{R}}(W)$. The graph W^{\bullet} is a "slightly richer variant" of W that is necessary for our definitions and proofs, namely to be able to associate every flap-vertex of an appropriate subgraph of $W_{\mathfrak{R}}$ (that we will denote by R_W) with a non-empty path of W^{\bullet} , as we proceed to formalize. We say that (W, \mathfrak{R}) is *well-aligned* if the following holds:

 $W_{\mathfrak{R}}$ contains as a subgraph an *r*-wall R_W where $D(R_W) = D(W_{\mathfrak{R}})$ and W^{\bullet} is isomorphic to some subdivision of R_W via an isomorphism that maps each ground vertex to itself.

Suppose now that the flatness pair (W, \mathfrak{R}) is well-aligned. We call the wall R_W in the above condition a *representation* of W in $W_{\mathfrak{R}}$. Note that, as R_W is a subgraph of $W_{\mathfrak{R}}$, it is bipartite as well. The above property

gives us a way to represent a flat wall by a wall of its leveling in a way that ground vertices are not altered. The following proposition, proved in [46], indicates that such a representation is yielded by regularity.

Proposition 19. Every regular flatness pair (W, \mathfrak{R}) of a graph G is well-aligned.

Notice that both $W_{\mathfrak{R}}$ and its subgraph R_W can be seen as Δ -embedded graphs where $\mathsf{bd}(\Delta) \cap W_{\mathfrak{R}} = \mathsf{bd}(\Delta) \cap R_W \subseteq V(D(W_{\mathfrak{R}})) = V(D(R_W))$. This establishes a bijection from the set of cycles of W to the set of cycles of R_W , which allows to reinterpret the homogeneity property of a regular flatness pair in terms of its representation, as stated in the following observation. This translation will be used in the proof of Theorem 23. Given the Δ -embedded graph R_W , we define, for every brick B of R_W , vflaps_{R_W}(B) as the flap-vertices of the leveling $W_{\mathfrak{R}}$ that belong to the closed disk of the plane bounded by B disjoint from the infinite face. (Recall Equation 3 for the definition of the augmented flap \mathbf{F}^A corresponding to a flap-vertex v_F of the leveling $W_{\mathfrak{R}}$.)

Observation 20. If (A, W, \mathfrak{R}) is an (a, r, ℓ) -apex-wall triple of a graph G and R_W is the representation of W in $W_{\mathfrak{R}}$, then for any two internal bricks B, B' of R_W , it holds that

$$\{\ell\text{-}\mathsf{folio}(\mathbf{F}^A) \mid v_F \in \mathsf{vflaps}_{R_W}(B)\} = \{\ell\text{-}\mathsf{folio}(\mathbf{F}^A) \mid v_F \in \mathsf{vflaps}_{R_W}(B')\}.$$

Note that, in the above equation, a flap F is notationally associated with both \mathbf{F}^A and v_F .

5.4 Rerouting minors of small intrusion

Let W be a plane-embedded r-wall and $c \ge 1$. We call a cycle C of W c-internal if V(C) and V(D(W)) are within face-distance at least c. Given a 1-internal cycle C of W, we define its internal pegs (resp. external pegs) as its vertices that are incident to edges of W that belong to the interior (resp. exterior) of C with respect to the embedding of the wall (we see edges as open sets). Notice that each vertex of C is either an internal or an external peg.

Observation 21. Let W be an r-wall and let C_1 and C_2 be two cycles of W within face-distance at least four and such that C_2 is a subset of the closed disk bounded by C_1 . Let $y \in [3]$, let p_1, \ldots, p_y be internal pegs of C_1 and $\bar{p}_1, \ldots, \bar{p}_y$ be external pegs of C_2 , assuming that both these sets of vertices are ordered as they appear in their corresponding cycles in counter-clockwise order. Then there are y pairwise vertex-disjoint paths $\hat{P}_1, \ldots, \hat{P}_y$ such that, for $i \in [y]$, \hat{P}_i joins p_i with $\bar{p}_i, V(\hat{P}_i) \cap V(C_1) = \{p_i\}$, and $V(\hat{P}_i) \cap V(C_2) = \{\bar{p}_i\}$.

Given a 1-internal brick B of W, one can see the union of all bricks of W that have a common vertex with B, as a subdivision of the graph in the left part of Figure 11. We call this subgraph X of W the *brick-neighborhood* of B in W. The *perimeter* of a brick-neighborhood is defined in the obvious way.

The next lemma is based on Observation 18.

Lemma 22. Let (W, \mathfrak{R}) be a well-aligned flatness pair of a graph G and let R_W be its representation in the leveling $W_{\mathfrak{R}}$ of W. For every 2-internal brick B of R_W and every flap vertex $v_F \in vflaps_{\mathfrak{R}}(B)$, R_W contains $|\partial F|$ internally vertex-disjoint paths from v_F to the external pegs of the perimeter of the brick-neighborhood of B in R_W . Moreover, these paths belong to the closed disk bounded by the perimeter of the brick-neighborhood of B in R_W .

Proof. Let $r = |\partial F|$, X be the brick-neighborhood of B in R_W , and P be the perimeter of X in R_W . We call *frontier-path* of X a subpath of P that joins two external pegs and does not contain any other external peg. Notice that P is the union of the frontier-paths of X and that there are exactly 12 such paths. Notice that for every frontier-path Q of X there is a path \hat{Q} of X such that

• its endpoints are in P but not in Q,



Figure 11: On the left: the base graph for the definition of a brick-neighborhood – the external pegs of the perimeter of X are the black round vertices. On the right we depict a 2-internal brick B of an r-wall W, $r \ge 6$, contained as a subgraph in a plane graph G along with three internally vertex-disjoint paths from a vertex F of G to the external pegs of the perimeter of the brick-neighborhood of B.

- it does not contain any internal vertex in P, and
- every path in $W_{\mathfrak{R}}$ from v_F to a vertex of Q intersects some vertex of \hat{Q} .

See Figure 12 for two indicative examples of the above correspondence.



Figure 12: Two indicative examples for the choice of \hat{Q} (in blue), depending on the choice of Q (in green).

From Observation 18, there are r internally vertex-disjoint paths P_1, \ldots, P_r in W_{\Re} starting from v_F and finishing in vertices v_1, \ldots, v_r of P. Moreover, we can assume that $(P_1 \cup \cdots \cup P_r) \cap P = \{v_1, \ldots, v_r\}$. The lemma is trivial in case r < 3 as the paths P_1, \ldots, P_r can easily be extended so to finish in external pegs of P. Moreover, the lemma also follows easily if r = 3 and v_1, v_2, v_3 do not all belong to some of the frontier-paths of X. It remains to examine the case where v_1, v_2, v_3 are vertices of some frontier-path Q of X. Let \hat{Q} be as defined above and let \hat{q} be one of its endpoints. Let z_Q be the first vertex of $P_1 \cup P_2 \cup P_3$ that is met while following \hat{Q} starting from \hat{q} and moving towards its other endpoint. The vertex z_Q exists because of the definition of \hat{Q} . W.l.o.g., we assume that z_Q is an internal vertex of P_1 . We now define P'_1 by first removing from P_1 all vertices of its subpath from z_Q to v_i , except z_Q and then taking the union of the resulting path with the subpath of \hat{Q} between \hat{q} and z_Q . It is now easy to see that P_1 can be extended to some external peg that is different from the endpoints of P_2 and P_3 , while P_2 and P_3 can be extended towards the external pegs that are endpoints of \hat{Q} . This completes the proof of the lemma.

Irrelevant vertices in boundaried graphs. Let G be a graph, H be a minor of G, and $S \subseteq V(G)$. We define the S-minor-intrusion of H in G as the minimum S-intrusion in G over all tm-pairs (M, T) of G such that (M, T) is a topological minor model in G where $diss(M, T) \in ext(H)$.

Let $\mathbf{Z} = (Z, B, \rho)$ be a *t*-boundaried graph and let $\ell \in \mathbb{N}$. We say that a vertex set $S \subseteq V(Z) \setminus B$ is ℓ -*irrelevant* if for every boundaried graph $\mathbf{K} = (K, B, \rho)$ that is compatible with \mathbf{Z} , every minor of $\mathbf{K} \oplus \mathbf{Z}$ with $(V(Z) \setminus B)$ -minor-intrusion at most ℓ , is also a minor of $\mathbf{K} \oplus (Z \setminus S, B, \rho)$. Informally, an ℓ -irrelevant

set of vertices can be removed without affecting the occurrences of any minor of minor-intrusion at most ℓ , where the intrusion is defined without taking into account the terminal vertices in the boundary.

Using Theorem 17, we can finally prove the main result of this section.

Theorem 23. There exist two functions $f_{13} : \mathbb{N}^3 \to \mathbb{N}$ and $f_{14} : \mathbb{N}^2 \to \mathbb{N}$ such that, for every $a, z, \ell \in \mathbb{N}$ and every boundaried graph $\mathbf{Z} = (Z, B, \rho)$, if (A, W, \mathfrak{R}) is an $(a, f_{13}(a, z, \ell), f_{14}(a, \ell))$ -apex-wall triple of Z that is not affected by B, then the vertex set of the compass of every W'-tilt of (W, \mathfrak{R}) , where W' is the central z-subwall of W, is ℓ -irrelevant. Moreover, it holds that $f_{13}(a, z, \ell) = \mathcal{O}((f_6(16a + 12\ell))^3 + z)$ and $f_{14}(a, \ell) = a + 3 + \ell$.

Proof. Let $\tilde{\ell} = 16a + 12\ell$, $\hat{\ell} = a + 3 + \ell$, $r = f_{12}(6, z, \tilde{\ell})$, and $\tilde{r} = f_{10}(6, \tilde{\ell})$. We prove the theorem for $f_{13}(a, z, \ell) = r$ and $f_{14}(a, \ell) = \hat{\ell}$. Note that, by Theorem 17, $f_{13}(a, z, \ell) = \mathcal{O}((f_6(\tilde{\ell}))^3 + z)$. Let $\mathbf{Z} = (Z, B, \rho)$ and let (A, W, \mathfrak{R}) be an $(a, r, \hat{\ell})$ -apex-wall triple of the graph Z that is not affected by B. Let W' be the central z-subwall of W, and let $(\tilde{W}', \tilde{\mathfrak{R}}')$ be a flatness pair of $Z \setminus A$ that is a W'-tilt of (W, \mathfrak{R}) . Our objective is to prove that $V(\text{compass}_{\mathfrak{R}'}(\tilde{W}'))$ is ℓ -irrelevant.

Let $\mathbf{K} = (K, B, \rho)$ be a boundaried graph compatible with \mathbf{Z} . As (A, W, \mathfrak{R}) is not affected by B, we have that (A, W, \mathfrak{R}) is an $(a, r, \hat{\ell})$ -apex-wall triple of the graph $G := \mathbf{K} \oplus \mathbf{Z}$ as well. Let now H be a minor of G whose $(V(Z) \setminus B)$ -minor-intrusion in G is at most ℓ . Let also (M, T) be a tm-pair in G that is a topological minor model of a graph $H' \in \text{ext}(H)$, such that $|T \setminus V(K)| = |T \cap (V(Z) \setminus B)| \leq \ell$ and with at most ℓ subdivision paths intersecting $V(Z) \setminus B$. Our purpose is to find a tm-pair $(\overline{M}, \overline{T})$ of G where $V(\text{compass}_{\mathfrak{K}'}(\tilde{W}')) \cap V(\overline{M}) = \emptyset$ and such that H' = diss(M, T) is a minor of $\text{diss}(\overline{M}, \overline{T})$, which in turn implies that H is a minor of $G \setminus V(\text{compass}_{\mathfrak{K}'}(\tilde{W}')) = \mathbf{K} \oplus (Z \setminus V(\text{compass}_{\mathfrak{K}'}(\tilde{W}')), B, \rho)$, as required.



Figure 13: An illustration of the graph $G = \mathbf{K} \oplus \mathbf{Z}$ and a tm-model (M, T) in it. The apex set A is cyan. We also draw a single flap F and its magnification.

We proceed with the definition of a series of auxiliary graphs that will permit to work on a partially planarized version of G. Suppose that $\mathfrak{R} = (X, Y, P, C, \Gamma, \sigma, \pi)$. We first define the graph G_{out} as the graph obtained from G if we remove all the vertices of $\text{compass}_{\mathfrak{R}}(W)$, except from those in $X \cap Y$. We then define $\tilde{G} := G_{\text{out}} \cup W_{\mathfrak{R}}$ where $W_{\mathfrak{R}}$ is the leveling of W. Since (W, \mathfrak{R}) is regular, it is also well-aligned by Proposition 19, hence $W_{\mathfrak{R}}$ contains a representation R_W of W, which is a subgraph of $W_{\mathfrak{R}}$ that is isomorphic to some subdivision of W^{\bullet} via an isomorphism that maps each ground vertex to itself (recall that W^{\bullet} is the graph obtained from W if we subdivide once every short edge in W). Notice that \hat{G} is a partially Δ -embedded graph whose compass is W_{\Re} that, in turn, is the compass in \tilde{G} of the *d*-wall R_W . Recall that each flap F of W corresponds to a flap-vertex v_F of W_{\Re} .

We enhance (M,T) by defining another tm-pair (M_A, T_A) , where $M_A = (V(M) \cup A, E(M))$ and $T_A = T \cup A$, i.e., (M_A, T_A) is obtained from (M,T) by including all the apices in A. We set $H_A = \text{diss}(M_A, T_A)$ and observe that H is a (topological) minor of H_A . We set $T_{\text{in}} = A \cup (T_A \cap \text{compass}_{\mathfrak{R}}(W))$ and $T_{\text{out}} = A \cup (T_A \setminus (\text{compass}_{\mathfrak{R}}(W) \setminus (X \cap Y)))$. As (A, W, \mathfrak{R}) is not affected by B, it follows that $T_{\text{in}} \cap \text{compass}_{\mathfrak{R}}(W) \subseteq T \setminus (V(K) \cup B)$, hence $|T_{\text{in}} \cap \text{compass}_{\mathfrak{R}}(W)| \leq \ell$.

We call a vertex $v \in \mathsf{compass}_{\mathfrak{R}}(W)$ an *apex-jump* vertex if there exists an edge $\{v, w\}$ of M_A with $w \in A$. Notice that there are at most $a + \ell$ apex-jump vertices. We define a set of (M_A, T_A) -dirty flaps of the flatness pair (W, \mathfrak{R}) by applying the following definition: first we declare as (M_A, T_A) -dirty every flap F such that $V(F) \setminus \partial F$ contains an apex-jump vertex or a vertex in T_{in} . Second, for every apex-jump vertex $v \in \mathsf{ground}(W)$ that is not in the boundary of some (M_A, T_A) -dirty flap, we arbitrarily pick a flap F with $v \in \partial F$ and we declare it (M_A, T_A) -dirty. Observe that (W, \mathfrak{R}) has at most $a + \ell + |T_{in} \cap \mathsf{compass}_{\mathfrak{R}}(W)| \leq a + 2\ell (M_A, T_A)$ dirty flaps.



Figure 14: An example of how part of the model (M_A, T_A) (in red edges) traverses the flaps of W along with the model $(\tilde{M}_{in}, \tilde{T}_{in})$ of \tilde{W} (i.e., the leveling of (M_A, T_A)). The dashed red edges are edges pointing to apices. The orange flaps are the (M_A, T_A) -dirty flaps. The black vertices are the ground vertices that do not belong to M.

We define M_{in} as the subgraph of the leveling W_{\Re} induced by the vertices in ground $(W) \cap V(M_A)$ and all the flap-vertices v_F of W_{\Re} such that F is either (M_A, T_A) -dirty or contains an edge of M (in Figure 14, these latter flaps are turquoise). Note that the induced edges may increase the degree of some, say w, of the degree-2 ground vertices in $V(M_A) \setminus T_A$ because of some edge $e = \{v_F, w\}$ between a flap-vertex v_F of W_{\Re} and w; we then remove from \tilde{M}_{in} all such edges (in Figure 14, these edges are the black dotted edges). Actually this last modification could be avoided, however it facilitates the presentation of the last part of the proof.

Notice that if a flap F is not (M_A, T_A) -dirty and contains an edge e of M, then this edge should belong to a subpath P of a subdivision path of M such that the endpoints of P are two of the vertices of ∂F . We call such flap-vertices of \tilde{M}_{in} subdivision flap-vertices of \tilde{M}_{in} . If a flap-vertex of \tilde{M}_{in} is not a subdivision flap-vertex, then we call it branch flap-vertex of \tilde{M}_{in} . We denote by Q the set of all the branch flap-vertices of \tilde{M}_{in} , and as each branch flap-vertex corresponds to a (M_A, T_A) -dirty flap, we have that $|Q| \leq a + 2\ell$. Notice that all vertices of Q have degree at most three in W_{\Re} (and therefore in \tilde{M}_{in} as well). The ground-vertices of $M_{\rm in}$ that are apex-jump vertices or belong to $T_{\rm in} \setminus A$ are called *branch ground-vertices* of $\tilde{M}_{\rm in}$, while the rest of the ground vertices of $\tilde{M}_{\rm in}$ are called *subdivision ground-vertices* of $\tilde{M}_{\rm in}$. Notice that there are at most $a + 2\ell$ branch ground-vertices in $\tilde{M}_{\rm in}$. Let $M_{\rm out} = M_A \cap G_{\rm out}$ and $\tilde{M} = M_{\rm out} \cup \tilde{M}_{\rm in}$. We define $\tilde{T}_{\rm in}$ as the union of the set of branch ground-vertices of $\tilde{M}_{\rm in}$ and the set Q of the branch flap-vertices of $\tilde{M}_{\rm in}$. Observe that $|\tilde{T}_{\rm in}| \leq 2a + 4\ell$. Moreover, we set $T_{\rm out} = T_A \cap V(M_{\rm out})$ and $\tilde{T} = T_{\rm out} \cup \tilde{T}_{\rm in}$.

Notice that (\tilde{M}, \tilde{T}) is a tm-pair of \tilde{G} ; we refer to it as the *leveling of the* tm-*pair* (M_A, T_A) of G with respect to (A, W, \mathfrak{R}) . Clearly, \tilde{M} is a partially Δ -embedded graph whose compass is \tilde{M}_{in} , i.e., $\tilde{M}_{in} = \Delta \cap \tilde{M}$. Also, as $|\tilde{T}_{in}| \leq 2a + 4\ell$, the $\Delta \cap V(G)$ -intrusion of (\tilde{M}, \tilde{T}) in \tilde{G} can be bounded, using the fact that \tilde{M}_{in} is a subgraph of the planar graph \tilde{G}_{in} , by $3(2a + 4\ell) = \tilde{\ell}$.

Let us now give some intuition on the definition of (\tilde{M}, \tilde{T}) . We see \tilde{G} and its tm-pair (\tilde{M}, \tilde{T}) as a "projection" of the graph G and its tm-pair (M, T), respectively, in what concerns the leveling W_{\Re} of W. This projection is loosing some of the information of (\tilde{M}, \tilde{T}) , however it will be valuable as now both \tilde{G} and \tilde{M} are partially Δ -embedded graphs. The lost information is encoded, for every branch flap-vertex v_F in Q, by ℓ -folio(\mathbf{F}^A). Recall that $\mathbf{F}^A := (G[A \cup F], A \cup \partial F, \rho_A \cup \rho_F)$. In what follows, we will use Theorem 17 in order to draw in \tilde{G} a modification (\hat{M}, \hat{T}) of (\tilde{M}, \tilde{T}) so that (\hat{M}, \hat{T}) does not go "too deeply" in the representation R_W of W in W_{\Re} , and the parts of (\hat{M}, \hat{T}) that are different from (\tilde{M}, \tilde{T}) are routed through R_W in a "dispersed enough" way. Moreover, again from Theorem 17, (\tilde{M}, \tilde{T}) can be seen as a contraction of (\hat{M}, \hat{T}) that does not identify any of the vertices of Q. Maintaining the vertices in Q intact, while performing contractions in (\hat{M}, \hat{T}) , is important as we need to keep the information given by $\hat{\ell}$ -folio(\mathbf{F}^A) for each branch flap-vertex $v_F \in Q$.

For each branch flap-vertex v_F of Q we define $\mathbf{M}_F = (G[A \cup V(F)] \cap M_A, A \cup \partial F, \rho_A \cup \rho_F)$, where the functions ρ_A and ρ_F are defined as explained in the beginning of Subsection 4.5, and $T_F = A \cup \partial F \cup (T_A \cap F)$. Notice that (\mathbf{M}_F, T_F) is a btm-pair of \mathbf{F}^A such that $\operatorname{diss}(\mathbf{M}_F, T_F) \in \hat{\ell}$ -folio (\mathbf{F}^A) , given that $\hat{\ell} = a + 3 + \ell$. This means that if $v_{\bar{F}}$ is another flap-vertex of $W_{\mathfrak{R}}$ with $\hat{\ell}$ -folio $(\mathbf{F}^A) = \hat{\ell}$ -folio $(\bar{\mathbf{F}}^A)$, then there is a btm-pair $(\mathbf{M}_{\bar{F}}, T_{\bar{F}})$ of $\bar{\mathbf{F}}^A$ such that $\operatorname{diss}(\mathbf{M}_F, T_F) = \operatorname{diss}(\mathbf{M}_{\bar{F}}, T_{\bar{F}})$.

Recall that \tilde{G} is a partially Δ -embedded graph whose compass is $\tilde{G} \cap \Delta = W_{\mathfrak{R}}$ and that \tilde{M} can be seen as a partially Δ -embedded graph whose compass is $\tilde{M} \cap \Delta$. Recall also that the $\Delta \cap V(G)$ -intrusion of (\tilde{M}, \tilde{T}) in \tilde{G} is at most $\tilde{\ell}$.

Let $C_1, \ldots, C_{\tilde{r}}$ be the first \tilde{r} layers of R_W and let $[D_1, \ldots, D_{\tilde{r}}]$ (resp. $[\overline{D}_1, \ldots, \overline{D}_{\tilde{r}}]$) be the sequences of the corresponding open (resp. closed) disks. We can now apply Theorem 17 with input 6, $\tilde{\ell}$, z, \tilde{G} , the *r*-wall R_W , the tm-pair (\tilde{M}, \tilde{T}) of \tilde{G} , and the set Q of branch flap-vertices defined above. Let $d = f_{11}(6, \tilde{\ell})$. Theorem 17 guarantees the existence of a tm-pair (\hat{M}, \hat{T}) of \tilde{G} and a $b \in [\tilde{r}]$ such that

- 1. $\hat{M} \setminus D_b$ is a subgraph of $\hat{M} \setminus D_b$,
- 2. $\operatorname{ann}(\mathcal{C}_{b,b+d-1}) \cap (\tilde{T} \cup \hat{T}) = \emptyset$,
- 3. $(\hat{M}, \hat{T}) \cap \overline{D}_{b+d}$ is a tm-pair of R_W that is safely 6-dispersed in R_W and none of the vertices of $\hat{T} \cap \overline{D}_{b+d}$ is within face-distance less than six in R_W from some vertex of $C_{b+d} \cup C_{\tilde{r}}$,
- 4. $\hat{M} \cap D_{\tilde{r}} = \emptyset$,
- 5. the compass of the central z-subwall of W is a subset of $D_{\tilde{r}}$, and
- 6. there is a Q-respecting contraction-mapping ϕ of diss (\tilde{M}, \tilde{T}) to diss (\hat{M}, \hat{T}) .

Our next step is to modify the tm-pair (\hat{M}, \hat{T}) to obtain another tm-pair (\hat{M}^+, \hat{T}^+) so that, in addition to Property 1 and Property 4 above (where \hat{M} is replaced by \hat{M}^+), (\hat{M}^+, \hat{T}^+) satisfies the following stronger version of Property 6:

6⁺. there is a Q-respecting contraction-mapping ϕ^+ of diss (\tilde{M}, \tilde{T}) to diss (\hat{M}^+, \hat{T}^+) such that for every branch flap-vertex $v_F \in Q$, if $\phi^+(v_F) = \{v_{\hat{F}}\}$, then $\hat{\ell}$ -folio $(\mathbf{F}^A) = \hat{\ell}$ -folio $(\hat{\mathbf{F}}^A)$.

To force Property 6^+ , we will make strong use of the fact that the flatness pair (W, \mathfrak{R}) is $\hat{\ell}$ -homogeneous with respect to (G, A). This will permit us to modify the subdivision paths of (\hat{M}, \hat{T}) by pointing to images of flaps in Q that have been displaced after the application of Theorem 17.

We proceed as follows: let $v_F \in Q$ be some flap-vertex, and let $\phi(v_F) = \{w\}$ be its image given by the function ϕ guaranteed by Property 6. From Property 2, w cannot belong to $\operatorname{ann}(\mathcal{C}_{b,b+d-1})$. Also, from Property 4, w cannot belong to $D_{\tilde{r}}$. If $w \in \operatorname{ann}(\mathcal{C}_{1,b})$, it follows from Property 1 that $w = v_F$ and in this case Property 6⁺ already holds trivially. The only remaining case is when $w \in \operatorname{ann}(\mathcal{C}_{b+d,\tilde{r}})$. In this case, from Property 3, w is a vertex of $R_W \cap \operatorname{ann}(\mathcal{C}_{b+d,\tilde{r}})$ that is within face-distance at least six from both C_{b+d} and $C_{\tilde{r}}$.



Figure 15: Some part of $R_W \cap \operatorname{ann}(\mathcal{C}_{b+d,\tilde{r}})$ around a vertex w of \hat{T} . The interior of the brick \hat{R} is depicted in yellow. The cycle C_w is depicted in blue.

Let R and \hat{R} be bricks of R_W such that $v_F \in \mathsf{vflaps}_{R_W}(R)$ and $w \in \mathsf{vflaps}_{R_W}(\hat{R})$, respectively. Certainly, both R and \hat{R} are internal bricks of R_W . Moreover, w is a vertex of the cycle \hat{R} (bounding the disk depicted in bright yellow in Figure 15). By Observation 20, there exists a vertex $v_{\hat{F}}$ in \hat{R} such that $\hat{\ell}$ -folio(\mathbf{F}^A) = $\hat{\ell}$ -folio($\hat{\mathbf{F}}^A$). Let y be the degree of w in \hat{M} (which equals the degree of v_F in \tilde{M}) and let P_1, \ldots, P_y be the subdivided paths of the model (\hat{M}, \hat{T}) that have w as an endpoint, depicted as fat green lines in Figure 15. Note than $y \leq 3$.

Let C_w be the cycle of R_W induced by the vertices that are within face-distance exactly six from w (in Figure 15, this cycle is depicted with thick blue edges). From Property 3, C_w is within face-distance at least $(2 \cdot 6 + 1) - 6 = 7$ from any other vertex in $\hat{T} \cap \overline{D}_{b+d}$ in R_W . For $i \in [y]$, we define P'_i as the unique path, among the paths obtained from P_i by the removal the open disk bounded by C_w , that has only one endpoint in the cycle C_w . For $i \in [y]$, we denote by p_i the endpoint of P'_i that is not an endpoint of P_i (depicted in purple in Figure 15). Notice that p_1, \ldots, p_y are internal pegs of C ordered as they appear in the perimeter in counter-clockwise order.

Consider now the brick-neighborhood, say X, of \hat{R} . Notice that the face-distance between the perimeter of X and C_w is at least four. From Lemma 22, there are y internally vertex-disjoint paths $\bar{P}_1, \ldots, \bar{P}_y$ from $v_{\hat{F}}$ to the external pegs of the perimeter C_X of X (depicted in red lines in Figure 15). Let these pegs be $\bar{p}_1, \ldots, \bar{p}_y$ ordered as they appear in the perimeter in counter-clockwise order. We are now in position to apply Observation 21 to C_w , C_X , p_1, \ldots, p_y , and $\bar{p}_1, \ldots, \bar{p}_y$ and find y pairwise vertex-disjoint paths $\hat{P}_1, \ldots, \hat{P}_y$ each joining p_i with \bar{p}_i (depicted in orange lines in Figure 15). We now remove from \hat{M} the open disk bounded by C_w and we add the graph $\bar{P}_1 \cup \cdots \cup \bar{P}_y \cup \hat{P}_1 \cup \cdots \cup \hat{P}_y$. We also modify \hat{T} by substituting wby $v_{\hat{F}}$. By repeating this procedure for every flap-vertex of Q, and defining (\hat{M}^+, \hat{T}^+) as the updated tm-pair obtained when this algorithm terminates, we enforce Property 6⁺. Namely the function ϕ^+ is defined from modifying ϕ so that, whenever $\phi(v_F) \neq \{v_F\}$, we set $\phi^+(v_F) = \{v_{\hat{F}}\}$ as above. Note that (\hat{M}^+, \hat{T}^+) indeed satisfies, besides Property 6⁺, Properties 1 and 4.

Our next (and last) step is to further modify the tm-pair (\hat{M}^+, \hat{T}^+) of \tilde{G} so to obtain a tm-pair $(\overline{M}, \overline{T})$ of the original graph G so that diss(M, T) is a minor of $diss(\overline{M}, \overline{T})$. Notice that each edge of \hat{M}^+ that is an edge of R_W (resp. $W_{\mathfrak{R}}$) has one endpoint that is a ground-vertex of R_W (resp. $W_{\mathfrak{R}}$) and another one that is a flap-vertex of R_W (resp. $W_{\mathfrak{R}}$).

Let $v_{\hat{F}}$ be a flap-vertex of \hat{M}^+ . We modify (\hat{T}^+, \hat{M}^+) by distinguishing the following cases:

- $v_{\hat{F}} \notin \hat{T}^+$. Then the neighbors of $v_{\hat{F}}$ in R_W are two ground-vertices g and g'. From tightness property (ii) of a rendition, there is a path $P_{\hat{F}}$ in \hat{F} with g and g' as endpoints. We substitute the edges $\{g, \hat{F}\}$ and $\{\hat{F}, g'\}$ of R_W with $P_{\hat{F}}$. Notice that $P_{\hat{F}}$ is a path of $\text{compass}_{\Re}(W)$.
- $v_{\hat{F}} \in \hat{T}^+$ and $v_{\hat{F}} \in \phi^+(x)$ for some $x \notin Q$. In this case, from tightness property (iii) of a rendition, \hat{F} has $y \in [3]$ neighbors v_1, \ldots, v_y , and a vertex $z_{\hat{F}}$ that is connected to v_1, \ldots, v_y via y internally vertexdisjoint paths. We substitute the edges $\{z_{\hat{F}}, v_1\}, \ldots, \{z_{\hat{F}}, v_y\}$ and the vertex $z_{\hat{F}}$ of R_W with the union of these y internally vertex-disjoint paths. Notice that these paths are also paths of compass_{\mathfrak{R}}(W). We also update $\hat{T}^+ := \hat{T}^+ \setminus \{v_{\hat{F}}\} \cup \{z_{\hat{F}}\}$ and we call the vertex $z_{\hat{F}}$ replacement of $v_{\hat{F}}$.
- $v_{\hat{F}} \in \hat{T}^+$ and $\{v_{\hat{F}}\} = \phi^+(v_F)$ for some $v_F \in Q$. Let $\mathbf{M}_F = (F \cap M, \partial F, \rho_A \cup \rho_F)$ and let $T_F = T \cap F$. From Property 6⁺, $\hat{\ell}$ -folio($\hat{\mathbf{F}}^A$) = $\hat{\ell}$ -folio(\mathbf{F}^A), which implies that there is a btm-pair ($\mathbf{M}_{\hat{F}}, T_{\hat{F}}$) of $\hat{\mathbf{F}}^A$ such that diss($\mathbf{M}_{\hat{F}}, T_{\hat{F}}$) = diss(\mathbf{M}_F, T_F). We denote by $M_{\hat{F}}$ the underlying graph of $\mathbf{M}_{\hat{F}}$ and we substitute the vertex $v_{\hat{F}}$ and its incident edges in $W_{\mathfrak{R}}$ with the graph $M_{\hat{F}}$ (that is a subgraph of compass_{\mathfrak{R}}(W)), by identifying the boundary vertices according to the functions $\rho_A \cup \rho_F$. We also update $\hat{T}^+ := \hat{T}^+ \setminus \{v_{\hat{F}}\} \cup T_{\hat{F}}$.

The above operations create a tm-pair of G that we denote henceforth by $(\overline{M},\overline{T})$. Since $(\overline{M},\overline{T})$ satisfies Property 4 and $\tilde{r} \geq z$, $\overline{M} \cap \bigcup$ influence_{\Re} $(W') = \emptyset$, and therefore $(\overline{M},\overline{T})$ is a tm-pair of $G \setminus V(\text{compass}_{\Re}(\tilde{W}'))$. It is worth mentioning that Property 1 implies the strong property that, throughout the rerouting procedure, the part of the topological minor model outside of the \Re -compass of W in $G \setminus A$ can only be reduced; formally, $\overline{M} \cap X \subseteq M \cap X$ (see Remark 24 after the end of this proof).

As H is a minor of $H' = \operatorname{diss}(M, T)$, hence a minor of $\operatorname{diss}(M_A, T_A)$ as well, and $G \setminus V(\operatorname{compass}_{\mathfrak{R}}(W')) = \mathbf{K} \oplus (Z \setminus V(\operatorname{compass}_{\mathfrak{R}}(\tilde{W}')), B, \rho)$, it remains to verify that $\operatorname{diss}(\overline{M}, \overline{T})$ contains $\operatorname{diss}(M_A, T_A)$ as a minor. For this, consider every $x \in \tilde{T} \setminus Q$ and construct the set T_x by taking $T_x := \phi^+(x)$ and substituting each branch flap-vertex $v_{\hat{F}} \in T_x$ with its replacement $z_{\hat{F}}$ defined as in the second case of the above case analysis. Notice that each set T_x is a subset of \overline{T} . We now construct (M^*, T^*) as follows: \tilde{M}^* is obtained by contracting, for each $x \in \tilde{T} \setminus Q$, all subdivision paths in \overline{M} that have as endpoints two vertices in T_x to a single vertex, which we again call x. This operation identifies, for each $x \in \tilde{T} \setminus Q$, all vertices of T_x into x, thus we set $T^* = (\overline{T} \setminus (\bigcup_{x \in \tilde{T} \setminus Q} T_x)) \cup (\tilde{T} \setminus Q)$. That way (M^*, T^*) can be seen as the tm-pair (\tilde{M}, \tilde{T}) where each of its branch flap-vertices has been substituted as in the third case of the above case analysis. This permits us to



Figure 16: The tm-pairs considered in the proof of Theorem 23. The cyan-shadowed pairs are tm-pairs of G and the pink-shadowed pairs are tm-pairs of \tilde{G} .

verify that $diss(M_A, T_A)$ and $diss(M^*, T^*)$ are isomorphic (see Figure 16 for the relation between the models that we have defined). As $diss(M^*, T^*)$ is a minor of $diss(\overline{M}, \overline{T})$, the theorem follows.

Remark 24. In the above proof, the rerouted topological minor model $(\overline{M}, \overline{T})$ obtained from the original model (M, T) satisfies the following property: the part of the topological minor model outside of the \mathfrak{R} compass of W in $G \setminus A$ can only be reduced with respect to the original one; formally, $\overline{M} \cap X \subseteq M \cap X$,
where X is the "external" set corresponding to the separation of the considered apex-wall triple (A, W, \mathfrak{R}) .

By definition of the set $\mathcal{R}_h^{(t)}$, its elements are of minimum size, and therefore a boundaried graph $\mathbf{G} = (G, B, \rho) \in \mathcal{R}_h^{(t)}$ does not contain any 3*h*-irrelevant vertex. To see this, recall that in Equation 2 the equivalence is defined in terms of graphs H with detail at most h (i.e., with at most h vertices and at most h edges), and that by Observation 3 every graph in ext(H) has detail at most 3h. On the other hand, Observation 7 implies that, in the setting of Theorem 23, there is a vertex belonging to the compass of every W'-tilt of (W, \mathfrak{R}) , where W' is the central z-subwall of W. Thus, from Theorem 23 for the particular case z = 3 and $\ell = 3h$, B should affect every $(a, f_{13}(a, 3, 3h), f_{14}(a, 3h))$ -apex-wall triple of G, for every value of a. As |B| = t, we conclude the following.

Corollary 25. If $t, h, a \in \mathbb{N}$ and $\mathbf{G} = (G, B, \rho)$ is a boundaried graph in $\mathcal{R}_h^{(t)}$, then B affects every $(a, f_{13}(a, 3, 3h), f_{14}(a, 3h))$ -apex-wall triple of G, in particular, $\mathbf{p}_{a, f_{13}(a, 3, 3h), f_{14}(a, 3h)}(G) \leq t$.

6 Bounding the size of the representatives

In this section we use the results obtained in the previous sections to prove that every representative in $\mathcal{R}_{h}^{(t)}$ has size linear in t. For this, we first prove in Subsection 6.1 that every representative in $\mathcal{R}_{h}^{(t)}$ has a set of $\mathcal{O}_{h}(t)$ vertices containing its boundary whose removal leaves a graph with treewidth bounded by a constant depending only on the collection \mathcal{F} ; such a set is called a *treewidth modulator*.

Once we have the treewidth modulator, we can use known results from the protrusion machinery to achieve our goal. Namely, in Subsection 6.2 we show how to obtain a linear protrusion decomposition of a representative, and we reduce each of the linearly many protrusions in the decomposition to an equivalent protrusion of constant size. Once we have this, a dynamic programming algorithm similar to that of [6] yields Theorem 33.

6.1 Finding a treewidth modulator of linear size

Given a graph G and a set $S \subseteq V(G)$, we say that a separation (L, R) of G is a 2/3-balanced separation of S in G if $|(L \setminus R) \cap S|, |(R \setminus L) \cap S| \leq \frac{2}{3}|S|$. We need the following well-known property of graphs of bounded treewidth (see e.g. [9,14]).

Lemma 26. Let G be a graph and let $S \subseteq V(G)$. There is a 2/3-balanced separation (L, R) of S in G of order at most tw(G) + 1.

Using Lemma 10, Lemma 11, Corollary 25, and Lemma 26 we prove the following result, whose proof uses Akra-Bazzi Theorem [3], in particular its extension provided by Leighton [36]. We stress that \mathbf{p} is not a bidimensional parameter in the precise way that is defined in [17, 18, 21], therefore Lemma 27 cannot be derived by directly applying the results of [21].

Lemma 27. There exist two functions $f_{15}, f_{16} : \mathbb{N}^2 \to \mathbb{N}$ such that if $t, q, h \in \mathbb{N}$ and $\mathbf{G} = (G, B, \rho)$ is a K_q -minor-free boundaried graph in $\mathcal{R}_h^{(t)}$, then G contains an $f_{15}(q, h)$ -treewidth modulator that contains B and has at most $f_{16}(q, h) \cdot t$ vertices. Moreover, it holds that $f_{15}(q, h) = \mathcal{O}((f_5(q, f_{13}(q), 3, 3h), f_{14}(f_3(q), 3h)))^2).$

Proof. Let $q, h \in \mathbb{N}$. We use q as a shortcut for the triple $(f_3(q), f_{13}(f_3(q), 3, 3h), f_{14}(f_3(q), 3h))$ and we set $s = f_5(q, f_{13}(f_3(q), 3, 3h), f_{14}(f_3(q), 3h))$. Let $t_0 = \max\{\min\{t' \mid s \cdot \sqrt{t'} + s + 1 \le t' / \log^2 t'\}, 42534179953\}$ and let $x = s \cdot \sqrt{t_0} + s$. We define the function $\mathbf{z} : \mathbb{N} \to \mathbb{N}$ where

 $\mathbf{z}(t) = \min\{z \mid \forall G \; \forall B \subseteq V(G) \text{ if } G \text{ is } K_q \text{-minor-free, } |B| \leq t, \text{ and } B \text{ affects every } q \text{-apex-wall triple of } G, \\ \text{then } \exists Z \subseteq V(G) : |Z| \leq z, B \subseteq Z, \mathsf{tw}(G \setminus Z) \leq x\}.$

Let G be a K_q -minor-free graph and let $B \subseteq V(G)$ such $|B| \leq t$ and B affects every q-apex-wall triple of G. From Lemma 10, tw $(G) \leq s \cdot \max\{1, \sqrt{|B|}\} \leq s \cdot \max\{1, \sqrt{t}\} \leq s \cdot \sqrt{t} + s$. From Lemma 26, G has a 2/3balanced separator (L, R) of B where $|L \cap R| \leq s \cdot \sqrt{t} + s + 1$. This means that $|(L \setminus R) \cap B|, |(R \setminus L) \cap B| \leq \frac{2}{3}|B|$. We set $G_L = G[L], G_R = G[R], B_L = L \cap (R \cup B) \subseteq V(G_L)$, and $B_R = R \cap (L \cup B) \subseteq V(G_R)$ and observe that $B \subseteq B_L \cup B_R$ and that both G_L and G_R are K_q -minor-free. Notice that $B_L = ((L \setminus R) \cap B) \cup (L \cap R)$ and $B_R = ((R \setminus L) \cap B) \cup (R \cap L)$, therefore there is some $\alpha \in [\frac{1}{2}, \frac{2}{3}]$, such that $|B_L| \leq \alpha \cdot t + s \cdot \sqrt{t} + s + 1$ and $|B_R| \leq (1 - \alpha) \cdot t + s \cdot \sqrt{t} + s + 1$. From Lemma 11, B_L affects every q-apex-wall triple of G_L and B_R affects every q-apex-wall triple of G_R . This means that there exists some $Z_L \subseteq V(G_L)$ such that $|Z_L| \leq \mathbf{z}(\alpha \cdot t + s \cdot \sqrt{t} + s + 1), B_L \subseteq Z_L$, and tw $(G_L \setminus Z_L) \leq x$. Also, there exists some $Z_R \subseteq V(G_R)$ such that $|Z_R| \leq \mathbf{z}(\alpha \cdot t + s \cdot \sqrt{t} + s + 1) + \mathbf{z}((1 - \alpha) \cdot t + s \cdot \sqrt{t} + s + 1)$. Moreover, $B \subseteq B_L \cup B_R \subseteq Z_L \cup Z_R = Z$ and, since $L \cap R \subseteq Z$, it holds that tw $(G \setminus Z) \leq x$. We obtain that

$$\mathbf{z}(t) \leq \mathbf{z}(\alpha \cdot t + s \cdot \sqrt{t} + s + 1) + \mathbf{z}((1 - \alpha) \cdot t + s \cdot \sqrt{t} + s + 1).$$
(4)

Let now $f: \mathbb{N}_{\geq 1} \to \mathbb{N}$ be the solution of the following recurrence:

$$f(t) = \begin{cases} f\left(\frac{t}{3} + s \cdot \sqrt{t} + s + 1\right) + f\left(\frac{2t}{3} + s \cdot \sqrt{t} + s + 1\right) & \text{if } t > t_0\\ t_0 & \text{if } 1 \le t \le t_0 \end{cases}$$

By the choice of t_0 , it holds that $t_0 = \Theta_s(1)$. Also, the choice of t_0 is made so that the conditions for applying the extended version of Akra-Bazzi Theorem [3] provided by Leighton [36] are satisfied⁶. Consequently, the solution of the above recurrence is $f(t) = \Theta_s(t^{\alpha})$ where α is the unique solution of the equation $(1/3)^{\alpha} + (2/3)^{\alpha} = 1$. Therefore $f(t) = \Theta_s(t)$.

Note that, if $1 \le t \le t_0$, then from Lemma 10 we have that $\mathbf{z}(t) \le t \le t_0 = f(t)$. On the other hand, by convexity, the right part of (4) is upper-bounded by f(t), so for all $t > t_0$ we have that $\mathbf{z}(t) \le f(t)$. Summarizing, we have that $\mathbf{z}(t) \le f(t)$ for all $t \ge 1$.

Let now $\mathbf{G} = (G, B, \rho)$ be a K_q -minor-free boundaried graph in $\mathcal{R}_h^{(t)}$. Applying Corollary 25 with $a = f_3(q)$, we obtain that B affects every q-apex-wall triple of G. Therefore G contains an x-treewidth modulator

⁶We verified these conditions using an elementary MATLAB program, from which the number 42534179953 was generated.

that contains B and has $\mathbf{z}(t) \leq f(t) = \mathcal{O}_s(t)$ vertices, as required. Therefore, $\mathbf{z}(t) \leq f_{16}(q,h) \cdot t$ for some function $f_{16} : \mathbb{N}^2 \to \mathbb{N}$. Observe that $t_0 = \mathcal{O}(s^2)$, therefore $x = \mathcal{O}(s^2)$ as well. The lemma follows with $f_{15}(q,h) := x = \mathcal{O}(s^2)$.

Note that the above proof does not give any estimation on the function $f_{16}(q, h)$. In the Appendix (Subsection B.1) we provide an improved version of Lemma 27, namely Lemma 34, with $f_{16}(q, h) = 2$. This will permit us to make a better estimation of the contribution of h in the running time of our algorithm (cf. Subsection B.2). The proof of Lemma 34 is an adaptation to our setting of the one of [21, Lemma 3.6].

6.2 Finding a linear protrusion decomposition and reducing protrusions

Equipped with Lemma 27, the next step is to construct an appropriate protrusion decomposition of a representative. We first need to define protrusions and protrusion decompositions of graphs and boundaried graphs.

Protrusion decompositions of unboundaried graphs. Given a graph G, a set $X \subseteq V(G)$ is a β protrusion of G if $|\partial(X)| \leq \beta$ and tw $(G[X]) \leq \beta - 1$. Given $\alpha, t \in \mathbb{N}$, an (α, β) -protrusion decomposition of G is a sequence $\mathcal{P} = \langle R_0, R_1, \ldots, R_\ell \rangle$ of pairwise disjoint subsets of V(G) such that

- $\bigcup_{i \in [\ell]} = V(G),$
- $\max\{\ell, |R_0|\} \le \alpha$,
- for $i \in [\ell]$, $N[R_i]$ is a β -protrusion of G, and
- for $i \in [\ell]$, $N(R_i) \subseteq R_0$.

We call the sets $N[R_i]$ $i \in [\ell]$, the protrusions of \mathcal{P} and the set R_0 the core of \mathcal{P} .

The above notions can be naturally generalized to boundaried graphs, just by requiring that both boundaries –of the host graph and of the protrusion– behave as one should expect, namely that the intersection of the protrusion with the boundary of the considered graph is a subset of the boundary of the protrusion.

Protrusions and protrusion decompositions of boundaried graphs. Given a boundaried graph $\mathbf{G} = (G, B, \rho)$, a tree decomposition of \mathbf{G} is any tree decomposition of G with a bag containing B. The treewidth of a boundaried graph \mathbf{G} , denoted by tw(\mathbf{G}), is the minimum width of a tree decomposition of \mathbf{G} . A boundaried graph $\mathbf{G}' = (G', B', \rho')$ is a β -protrusion of \mathbf{G} if

- V(G') is a β -protrusion of G,
- $\mathsf{tw}(\mathbf{G}') \leq \beta 1$,
- $\partial(V(G')) \subseteq B'$, and
- $B \cap V(G') \subseteq B'$.

Given a boundaried graph $\mathbf{G} = (G, B, \rho)$ and $\alpha, t \in \mathbb{N}$, an (α, β) -protrusion decomposition of \mathbf{G} is a sequence $\mathcal{P} = \langle R_0, R_1, \ldots, R_\ell \rangle$ of pairwise disjoint subsets of V(G) such that

- $\bigcup_{i \in [\ell]} = V(G),$
- $\max\{\ell, |R_0|\} \le \alpha$,
- $B \subseteq R_0$,

- for $i \in [\ell]$, $(G(N[R_i]), \partial(N[R_i]), \rho|_{\partial(N[R_i])})$ is a β -protrusion of **G**, and
- for $i \in [\ell]$, $N(R_i) \subseteq R_0$.

As in the unboundaried case, we call the sets $N[R_i]$ $i \in [\ell]$, the protrusions of \mathcal{P} and the set R_0 the core of \mathcal{P} .

The following theorem is a reformulation using our notation of one of the main results of Kim et al. [32], which is stronger than what we need, in the sense that also applies to *topological*-minor-graphs. It is worth mentioning that, for H-minor-free-graphs, an appropriate protrusion decomposition can also be found using the results in [21, Lemma 3.10].

Theorem 28. Let c, β, t be positive integers, let H be a q-vertex graph, and let G be an n-vertex H-topologicalminor-free graph. If we are given a set $M \subseteq V(G)$ with $|M| \leq c \cdot t$ such that $\operatorname{tw}(G - M) \leq \beta$, then we can compute in time $\mathcal{O}(n)$ an $((\alpha_H \cdot \beta \cdot c) \cdot t, 2\beta + q)$ -protrusion decomposition \mathcal{P} of G with M contained in the core of \mathcal{P} , where α_H is a constant depending only on H such that $\alpha_H \leq 40q^2 2^{5q \log q}$.

Having stated the above definitions, the following lemma is an easy consequence of Lemma 34 and Theorem 28.

Lemma 29. There exists a function $f_{17} : \mathbb{N}^2 \to \mathbb{N}$ such that if $t, q, h \in \mathbb{N}$ and $\mathbf{G} = (G, B, \rho)$ is a K_q -minor-free boundaried graph in $\mathcal{R}_h^{(t)}$, then \mathbf{G} admits a $(f_{17}(q,h) \cdot t, f_{17}(q,h))$ -protrusion decomposition. Moreover, it holds that $f_{17}(q,h) = f_{15}(q,h) \cdot f_{16}(q,h) \cdot 2^{\mathcal{O}(q \log q)}$.

Proof. By Lemma 34, G contains an $f_{21}(q,h)$ -treewidth modulator M that contains B and has at most 2t vertices. We can now apply Theorem 28 to G and M with $H = K_q$, $c = f_{16}(q,h)$, and $\beta = f_{15}(q,h)$, obtaining a $(f_{17}(q,h) \cdot t, f_{17}(q,h))$ -protrusion decomposition \mathcal{P} of G with M contained in the core of \mathcal{P} and $f_{17}(q,h) := f_{15}(q,h) \cdot f_{16}(q,h) \cdot 40q^2 2^{5q \log q}$. Since $B \subseteq M$ and M contained in the core of \mathcal{P} , it can be easily checked that \mathcal{P} is also a $(f_{17}(q,h) \cdot t, f_{17}(q,h))$ -protrusion decomposition of \mathbf{G} .

Once we have the protrusion decomposition given by Lemma 29, all that remains is to replace the protrusions by equivalent ones of size depending only on the collection \mathcal{F} . The protrusion replacement technique, which is nowadays part of the basic toolbox of parameterized complexity, originated in the meta-theorem of Bodlaender et al. [11], whose objective was to produce linear kernels for a wide family of problems on graphs of bounded genus. This technique was later extended to graphs excluding a fixed minor by Fomin et al. [21] and then to graphs excluding a fixed topological minor by Kim et al. [32]. We could directly apply the results of Fomin et al. [21] to the protrusion decomposition of a representative given by Lemma 29, hence reducing each protrusion to an equivalent one of size $\mathcal{O}_{\mathcal{F}}(1)$, yielding an equivalent representative of size $\mathcal{O}_{\mathcal{F}}(t)$. However, the drawback of the results in [21] (and also in [11,32]) is that they do not provide explicit bounds on the hidden constants. In order to be able to do so (cf. Subsection B.2), we apply the protrusion replacement used by Baste et al. [6], which is suited for the \mathcal{F} -M-DELETION problem. This yields explicit constants because it uses ideas similar to the ones presented by Garnero et al. [22] (later generalized in [23]) for obtaining kernels with explicit constants.

Given a function $\xi : \mathbb{N}^2 \to \mathbb{N}$ and a *t*-boundaried graph **G**, we say that **G** is ξ -protrusion-bounded if, for every $t' \in \mathbb{N}$, all β -protrusions of **G** have at most $\xi(\beta)$ vertices. The following lemma is again a reformulation using our notation of one of the results of Baste et al. [6]. Namely, it is a consequence of the proof⁷ of [6, Lemma 7.2].

⁷In the statement of [6, Lemma 7.2] it is required that the family \mathcal{F} contains a planar graph, an assumption that is not true anymore in our case. However, in the proof this fact is only used to guarantee that the considered protrusion has treewidth bounded by a function depending only on \mathcal{F} . Thanks to Lemma 29, we can assume that this also holds in our setting.

Lemma 30. There exists a function $f_{18} : \mathbb{N}^2 \to \mathbb{N}$ such that if $t, q, h \in \mathbb{N}$ and $\mathbf{G} = (G, B, \rho)$ is a K_q -minor-free boundaried graph in $\mathcal{R}_h^{(t)}$, then \mathbf{G} is $f_{18}(q, h)$ -protrusion-bounded. Moreover, $f_{18}(q, h) = 2^{2^{\mathcal{O}(f_{17}(q,h) \cdot \log f_{17}(q,h))}}$.

Using Lemma 29 and Lemma 30, we can easily prove Theorem 31, that is the main result on which the algorithm of Theorem 2 is based (cf. Section 2). In particular, it implies Equation 1.

Theorem 31. There exists a function $f_{19} : \mathbb{N}^2 \to \mathbb{N}$ such that, for every $t \in \mathbb{N}$ and $q, h \in \mathbb{N}_{\geq 1}$, if $\mathbf{G} = (G, B, \rho)$ is a K_q -minor-free boundaried graph in $\mathcal{R}_h^{(t)}$, then $|V(G)| \leq f_{19}(q, h) \cdot t$. Moreover, it holds that $f_{19}(q, h) \leq f_{17}(q, h) \cdot (f_{18}(q, h) + 1)$.

Proof. By Lemma 29, **G** admits an $(f_{17}(q, h) \cdot t, f_{17}(q, h))$ -protrusion decomposition \mathcal{P} . By Lemma 30, each of the protrusions of \mathcal{P} has at most $f_{18}(q, h)$ vertices. Therefore,

$$|V(G)| \le f_{17}(q,h) \cdot t + f_{17}(q,h) \cdot f_{18}(q,h) \cdot t,$$

and the theorem follows with $f_{19}(q,h) := f_{17}(q,h) \cdot (f_{18}(q,h)+1)$.

Let $h := \max_{F \in \mathcal{F}} \{ \max_{H \in \mathsf{ext}(F)} \mathsf{detail}(H) \}$. The following corollary is an immediate consequence of Theorem 31, by using the fact that all t-representatives in $\mathcal{R}_h^{(t)}$, except one, are K_h -minor-free, hence they have $\mathcal{O}(f_{19}(h,h) \cdot h\sqrt{\log h}) \cdot t$ edges; see for instance [39]. Note that are at most $\binom{n^2}{m} = 2^{\mathcal{O}(n \log m)}$ different graphs on n vertices and m edges and that, if $(G, B, \rho) \in \mathcal{R}_h^{(t)}$, then Theorem 31 implies that $|V(G)| \leq f_{19}(h,h) \cdot t$. Note also that there are $\binom{|V(G)|}{t} = 2^{\mathcal{O}(t \log |V(G)|)}$ choices for B, and $t! = 2^{\mathcal{O}(t \log t)}$ choices for ρ . Therefore, $|\mathcal{R}_h^{(t)}| = 2^{\mathcal{O}\left(f_{19}(h,h) \cdot t \cdot \log(f_{19}(h,h) \cdot h\sqrt{\log h} \cdot t) + f_{19}(h,h) \cdot t \log(f_{19}(h,h) \cdot t + \log t)\right)}$ and we can conclude the following.

Corollary 32. There exists a function $f_{20} : \mathbb{N} \to \mathbb{N}$ such that for every $t \in \mathbb{N}_{\geq 1}$, $|\mathcal{R}_h^{(t)}| \leq 2^{f_{20}(h) \cdot t \cdot \log t}$. In particular, the relation \equiv_h partitions $\mathcal{B}^{(t)}$ into at most $2^{f_{20}(h) \cdot t \cdot \log t}$ equivalence classes. Moreover, it holds that $f_{20}(h) = \mathcal{O}(f_{19}(h,h) \cdot \log(f_{19}(h,h) \cdot h\sqrt{\log h}))$.

The dynamic programming algorithm. Having proved Corollary 32, we can apply [6, Theorem 9] to compute the parameter $\mathbf{m}_{\mathcal{F}}(G)$ within the claimed running time.

For the sake of completeness, let us comment some details of this dynamic programming algorithm, whose details can be found in [6, Section 8]. First of all, to run the algorithm we need to have the set $\mathcal{R}_{h}^{(t)}$ of representatives at hand. This can be done easily relying on Theorem 31, by generating all t-boundaried graphs on at most $f_{19}(h,h) \cdot t$ vertices and $\mathcal{O}(f_{19}(h,h) \cdot h\sqrt{\log h}) \cdot t$ edges, partitioning them into equivalence classes according to \equiv_h , and picking an element of minimum size in each of them; see [6, proof of Lemma 7.1] for more details. To simplify the description of the dynamic programming update operations, the main algorithm in [6] is written in terms of *branchwidth* instead of treewidth. Without defining branchwidth here, it is enough to say that it is linearly equivalent to treewidth, in the sense that both parameters differ by a constant factor and whose corresponding decompositions can be easily transformed from one to the other [42]. Also, the main algorithm in [6] is written in terms of topological minors, that is, given a finite graph class \mathcal{F}' and a graph G, it computes $\operatorname{tm}_{\mathcal{F}'}(G)$, that is the minimum-size set of vertices $S \subseteq V(G)$ whose removal leaves a graph without any of the graphs in a fixed collection \mathcal{F}' as a topological minor. This works for our purposes because of the translation of the question on minors to one on topological minors, provided by Observation 3. The dynamic algorithm computes, in a typical bottom-up manner, at every bag separator Bof the branch decomposition associated with a t-boundaried graph \mathbf{G}_B and for every representative $\mathbf{R} \in \mathcal{R}_h^{(t)}$ the minimum size of a set $S \subseteq V(\mathbf{G}_B)$ such that $\mathbf{G}_B \setminus S \equiv_h \mathbf{R}$. These values can be computed in a standard way by combining the values associated with the children of a given node; cf. [6, Theorem 9]. The overall running time is bounded by $\mathcal{O}(|\mathcal{R}_h^{(t)}|^2 \cdot |E(G)|)$, and taking into account that $|E(G)| \leq \mathsf{tw}(G) \cdot |V(G)|$, from Corollary 32 we obtain the following theorem, which is a more precise reformulation of Theorem 2.

Theorem 33. Let $t, h \in \mathbb{N}$, \mathcal{F} be a proper collection of size at most h, and G be an n-vertex graph of treewidth at most t. Then $\mathbf{m}_{\mathcal{F}}(G)$ can be computed by an algorithm that runs in $2^{\mathcal{O}(f_{20}(h) \cdot t \log t)} \cdot n$ steps.

In Appendix B we give upper bounds on the constants depending on the collection \mathcal{F} involved in our algorithm. These upper bounds depend explicitly on the parametric dependencies of the Unique Linkage Theorem [31,44].

7 Further research

We presented an algorithm for solving the \mathcal{F} -M-DELETION problem in time $2^{\mathcal{O}(\mathsf{tw} \cdot \log \mathsf{tw})} \cdot n$ for every collection \mathcal{F} . This algorithm together with the single-exponential algorithms and lower bounds presented in previous papers of this series [7, 8] yield a complete classification of the asymptotic complexity of \mathcal{F} -M-DELETION parameterized by treewidth assuming the ETH, when $\mathcal{F} = \{H\}$ and H is connected (Theorem 1). However, we do not have a complete classification when $|\mathcal{F}| > 2$, even for connected \mathcal{F} . To ease the presentation, let us call a connected graph H easy (resp. hard) if $\{H\}$ -M-DELETION is solvable in time $\mathcal{O}^*(2^{\Theta(\mathsf{tw})})$ (resp. $\mathcal{O}^*(2^{\Theta(\mathsf{tw} \cdot \log \mathsf{tw})}))$. Suppose that $\mathcal{F} = \{H_1, H_2\}$ with both H_1 and H_2 being connected. Using the recent results of Baste [5], it is possible to prove that if both H_1 and H_2 are easy, then \mathcal{F} is easy as well (easiness of graph collections is defined in the obvious way). However, if both H_1 and H_2 are hard, then strange things may happen. For instance, Bodlaender et al. [12] presented an algorithm running in time $\mathcal{O}^*(2^{\mathcal{O}(\mathsf{tw})})$ for PSEUDOFOREST DELETION, which consists in, given a graph G and an integer k, deciding whether one can delete at most k vertices from G to obtain a *pseudoforest*, i.e., a graph where each connected component contains at most one cycle. Note that PSEUDOFOREST DELETION is equivalent to {diamond, butterfly}-M-DELETION. While both the diamond and the butterfly are hard graphs (cf. Figure 17), {diamond, butterfly} is an easy collection. The cases where H_1 is easy and H_2 is hard seem even trickier. Obtaining (tight) lower bounds when \mathcal{F} may contain disconnected graphs is another challenging avenue for further research.

It is also interesting to consider the version of the problem where the graphs in \mathcal{F} are forbidden as topological minors; we call this problem \mathcal{F} -TM-DELETION. While most of the lower bounds that we presented in [8] also hold for \mathcal{F} -TM-DELETION, the algorithm in time $2^{\mathcal{O}(\mathsf{tw} \cdot \log \mathsf{tw})} \cdot n$ of this paper does not work for topological minors. In this direction, the algorithm in time $2^{\mathcal{O}(\mathsf{tw} \cdot \log \mathsf{tw})} \cdot n$ for \mathcal{F} -M-DELETION when \mathcal{F} contains a planar graph given in [6] also works for \mathcal{F} -TM-DELETION, if we additionally require \mathcal{F} to contain a subcubic planar graph (in order to bound the treewidth of the representatives). The main obstacle for applying our approach in order to achieve a time $\mathcal{O}^*(2^{\mathcal{O}(\mathsf{tw} \cdot \log \mathsf{tw})})$ for every collection \mathcal{F} , is that topological-minor-free graphs do not enjoy the flat wall structure that is omnipresent in our proofs. Another reason is that in our rerouting procedure, in order to find an irrelevant vertex (Theorem 17), we may find a different topological minor model that corresponds to the same minor. Nevertheless, we think that this latter difficulty can be overcome for planar graphs –or even minor-free graphs– by making use of the rerouting potential of Proposition 13, as this is done in [25] for planar graphs.

Finally, it is worth mentioning that the algorithm presented in this paper, as well as the main combinatorial result (Theorem 23), have been used in [45] (see [47] for the full version) to obtain a fixed-parameter algorithm for the \mathcal{F} -M-DELETION problem parameterized by k. Theorem 23 has also been used in [48] in order to provide explicit upper bounds on the size of the minor-obstructions of the set of yes-instances of the \mathcal{F} -M-DELETION problem, as a function of \mathcal{F} and k.

Acknowledgement. We would like to thank the reviewers for helpful remarks that improved the presentation of the article. Special thanks to Giannos Stamoulis for his suggestions and viewpoints on several aspects of this paper.

References

- Isolde Adler, Stavros G. Kolliopoulos, Philipp Klaus Krause, Daniel Lokshtanov, Saket Saurabh, and Dimitrios M. Thilikos. Irrelevant vertices for the planar Disjoint Paths Problem. *Journal of Combinatorial Theory, Series B*, 122:815–843, 2017. doi:10.1016/j.jctb.2016.10.001.
- [2] Isolde Adler and Philipp Klaus Krause. A lower bound on the tree-width of graphs with irrelevant vertices. Journal of Combinatorial Theory, Series B, 137:126–136, 2019. doi:10.1016/j.jctb.2018. 12.008.
- [3] Mohamad A. Akra and Louay Bazzi. On the solution of linear recurrence equations. Computational Optimization and Applications, 10(2):195–210, 1998. doi:10.1023/A:1018373005182.
- [4] Zongwen Bai, Jianhua Tu, and Yongtang Shi. An improved algorithm for the vertex cover P₃ problem on graphs of bounded treewidth. *Discrete Mathematics & Theoretical Computer Science*, 21(4), 2019. doi:10.23638/DMTCS-21-4-17.
- [5] Julien Baste. Composing dynamic programming tree-decomposition-based algorithms. Manuscript, 2019. arXiv:1904.12500.
- [6] Julien Baste, Ignasi Sau, and Dimitrios M. Thilikos. Hitting minors on bounded treewidth graphs. I. General upper bounds. SIAM Journal on Discrete Mathematics, 34(3):1623–1648, 2020. doi:10.1137/19M1287146.
- Julien Baste, Ignasi Sau, and Dimitrios M. Thilikos. Hitting minors on bounded treewidth graphs. II. Single-exponential algorithms. *Theoretical Computer Science*, 814:135–152, 2020. doi:10.1016/j.tcs. 2020.01.026.
- [8] Julien Baste, Ignasi Sau, and Dimitrios M. Thilikos. Hitting minors on bounded treewidth graphs. III. Lower bounds. Journal of Computer and System Sciences, 109:56-77, 2020. doi:10.1016/j.jcss. 2019.11.002.
- Hans L. Bodlaender. A Partial k-Arboretum of Graphs with Bounded Treewidth. Theoretical Computer Science, 209(1-2):1–45, 1998. doi:10.1016/S0304-3975(97)00228-4.
- [10] Hans L. Bodlaender, Marek Cygan, Stefan Kratsch, and Jesper Nederlof. Deterministic single exponential time algorithms for connectivity problems parameterized by treewidth. *Information and Computation*, 243:86–111, 2015. doi:10.1016/j.ic.2014.12.008.
- [11] Hans L. Bodlaender, Fedor V. Fomin, Daniel Lokshtanov, Eelko Penninkx, Saket Saurabh, and Dimitrios M. Thilikos. (Meta) Kernelization. *Journal of the ACM*, 63(5):44:1–44:69, 2016. doi: 10.1145/2973749.
- [12] Hans L. Bodlaender, Hirotaka Ono, and Yota Otachi. A faster parameterized algorithm for pseudoforest deletion. Discrete Applied Mathematics, 236:42–56, 2018. doi:10.1016/j.dam.2017.10.018.
- [13] Bruno Courcelle. The Monadic Second-Order Logic of Graphs. I. Recognizable Sets of Finite Graphs. Information and Computation, 85(1):12–75, 1990. doi:10.1016/0890-5401(90)90043-H.
- [14] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015. doi:10.1007/ 978-3-319-21275-3.

- [15] Marek Cygan, Dániel Marx, Marcin Pilipczuk, and Michal Pilipczuk. Hitting forbidden subgraphs in graphs of bounded treewidth. *Information and Computation*, 256:62–82, 2017. doi:10.1016/j.ic. 2017.04.009.
- [16] Marek Cygan, Jesper Nederlof, Marcin Pilipczuk, Michal Pilipczuk, Johan M. M. van Rooij, and Jakub Onufry Wojtaszczyk. Solving Connectivity Problems Parameterized by Treewidth in Single Exponential Time. In Proc. of the 52nd IEEE Symposium on Foundations of Computer Science (FOCS), pages 150–159, 2011. doi:10.1109/FOCS.2011.23.
- [17] Erik D. Demaine, Fedor V. Fomin, Mohammadtaghi Hajiaghayi, and Dimitrios M. Thilikos. Subexponential parameterized algorithms on bounded-genus graphs and *H*-minor-free graphs. *Journal of the* ACM, 52(6):866-893, 2005. doi:10.1145/1101821.1101823.
- [18] Fedor V. Fomin, Erik D. Demaine, Mohammad Taghi Hajiaghayi, and Dimitrios M. Thilikos. Bidimensionality. In *Encyclopedia of Algorithms*, pages 203–207. Springer, 2016. doi:10.1007/ 978-1-4939-2864-4_47.
- [19] Fedor V. Fomin and Dieter Kratsch. Exact Exponential Algorithms. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2010. doi:10.1007/978-3-642-16533-7.
- [20] Fedor V. Fomin, Daniel Lokshtanov, Fahad Panolan, and Saket Saurabh. Efficient computation of representative families with applications in parameterized and exact algorithms. *Journal of the ACM*, 63(4):29:1–29:60, 2016. doi:10.1145/2886094.
- [21] Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, and Dimitrios M. Thilikos. Bidimensionality and kernels. SIAM Journal on Computing, 49(6):1397–1422, 2020. doi:10.1137/16M1080264.
- [22] Valentin Garnero, Christophe Paul, Ignasi Sau, and Dimitrios M. Thilikos. Explicit linear kernels via dynamic programming. SIAM Journal on Discrete Mathematics, 29(4):1864–1894, 2015. doi: 10.1137/140968975.
- [23] Valentin Garnero, Christophe Paul, Ignasi Sau, and Dimitrios M. Thilikos. Explicit linear kernels for packing problems. Algorithmica, 81(4):1615–1656, 2019. doi:10.1007/s00453-018-0495-5.
- [24] Petr A. Golovach, Marcin Kaminski, Spyridon Maniatis, and Dimitrios M. Thilikos. The parameterized complexity of graph cyclability. SIAM Journal on Discrete Mathematics, 31(1):511–541, 2017. doi: 10.1137/141000014.
- [25] Petr A. Golovach, Giannos Stamoulis, and Dimitrios M. Thilikos. Hitting topological minor models in planar graphs is fixed parameter tractable. In Proc. of the 31st ACM-SIAM Symposium on Discrete Algorithms, (SODA), pages 931–950, 2020. doi:10.1137/1.9781611975994.56.
- [26] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? Journal of Computer and System Sciences, 63(4):512–530, 2001. doi:10.1006/jcss. 2001.1774.
- [27] Bart M. P. Jansen, Daniel Lokshtanov, and Saket Saurabh. A near-optimal planarization algorithm. In Proc. of the 25th ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 1802–1811, 2014. doi:10.1137/1.9781611973402.130.
- [28] Marcin Kaminski and Dimitrios M. Thilikos. Contraction checking in graphs on surfaces. In Proc. of the 29th International Symposium on Theoretical Aspects of Computer Science, (STACS), volume 14 of LIPIcs, pages 182–193, 2012. doi:10.4230/LIPIcs.STACS.2012.182.

- [29] Goossen Kant. Algorithm for Drawing Planar Graphs. PhD thesis, Department of Information and Computing Sciences - Utrecht University, 1993. URL: https://dspace.library.uu.nl/bitstream/ handle/1874/842/full.pdf.
- [30] Ken-ichi Kawarabayashi, Robin Thomas, and Paul Wollan. A new proof of the flat wall theorem. Journal of Combinatorial Theory, Series B, 129:204–238, 2018. doi:10.1016/j.jctb.2017.09.006.
- [31] Ken-ichi Kawarabayashi and Paul Wollan. A shorter proof of the graph minor algorithm: the unique linkage theorem. In Proc. of the 42nd ACM Symposium on Theory of Computing (STOC), pages 687– 694, 2010. doi:10.1145/1806689.1806784.
- [32] Eun Jung Kim, Alexander Langer, Christophe Paul, Felix Reidl, Peter Rossmanith, Ignasi Sau, and Somnath Sikdar. Linear kernels and single-exponential algorithms via protrusion decompositions. ACM Transactions on Algorithms, 12(2):21:1–21:41, 2016. doi:10.1145/2797140.
- [33] Tomasz Kociumaka and Marcin Pilipczuk. Deleting vertices to graphs of bounded genus. Algorithmica, 81(9):3655–3691, 2019. doi:10.1007/s00453-019-00592-7.
- [34] Tuukka Korhonen. A single-exponential time 2-approximation algorithm for treewidth. In Proc. of the 62nd IEEE Annual Symposium on Foundations of Computer Science (FOCS), pages 184–192, 2021. doi:10.1109/F0CS52979.2021.00026.
- [35] Stefan Kratsch and Magnus Wahlström. Representative sets and irrelevant vertices: New tools for kernelization. Journal of the ACM, 67(3):16:1–16:50, 2020. doi:10.1145/3390887.
- [36] Tom Leighton. Notes on Better Master Theorems for Divide-and-Conquer Recurrences, 1996. URL: https://courses.csail.mit.edu/6.046/spring04/handouts/akrabazzi.pdf.
- [37] John M. Lewis and Mihalis Yannakakis. The node-deletion problem for hereditary properties is NP-complete. Journal of Computer and System Sciences, 20(2):219–230, 1980. doi:10.1016/ 0022-0000(80)90060-4.
- [38] Daniel Lokshtanov, Dániel Marx, and Saket Saurabh. Slightly superexponential parameterized problems. SIAM Journal on Computing, 47(3):675–702, 2018. doi:10.1137/16M1104834.
- [39] Joseph Samuel Myers and Andrew Thomason. The extremal function for noncomplete minors. *Combinatorica*, 25(6):725–753, 2005. doi:10.1007/s00493-005-0044-0.
- [40] Marcin Pilipczuk. A tight lower bound for Vertex Planarization on graphs of bounded treewidth. Discrete Applied Mathematics, 231:211-216, 2017. doi:10.1016/j.dam.2016.05.019.
- [41] Bruce A. Reed. Rooted routing in the plane. Discrete Applied Mathematics, 57(2-3):213-227, 1995.
 doi:10.1016/0166-218X(94)00104-L.
- [42] Neil Robertson and Paul D. Seymour. Graph Minors. X. Obstructions to Tree decomposition. Journal of Combinatorial Theory, Series B, 52(2):153–190, 1991. doi:10.1016/0095-8956(91)90061-N.
- [43] Neil Robertson and Paul D. Seymour. Graph Minors. XIII. The Disjoint Paths Problem. Journal of Combinatorial Theory, Series B, 63(1):65-110, 1995. doi:10.1006/jctb.1995.1006.
- [44] Neil Robertson and Paul D. Seymour. Graph Minors. XXII. Irrelevant Vertices in Linkage Problems. Journal of Combinatorial Theory, Series B, 102(2):530–563, 2012. doi:10.1016/j.jctb.2007.12.007.

- [45] Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. An FPT-algorithm for recognizing k-apices of minor-closed graph classes. In Proc. of the 47th International Colloquium on Automata, Languages, and Programming (ICALP), volume 168 of LIPIcs, pages 95:1–95:20, 2020. doi:10.4230/LIPIcs.ICALP. 2020.95.
- [46] Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. A more accurate view of the Flat Wall Theorem. Manuscript, 2021. arXiv:2102.06463.
- [47] Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. k-apices of minor-closed graph classes. II. parameterized algorithms. ACM Transactions on Algorithms, 18(3):21:1–21:30, 2022. doi:10.1145/ 3519028.
- [48] Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. k-apices of minor-closed graph classes.
 I. Bounding the obstructions. Journal of Combinatorial Theory, Series B, 161:180-227, 2023. doi: 10.1016/j.jctb.2023.02.012.
- [49] Jianhua Tu, Lidong Wu, Jing Yuan, and Lei Cui. On the vertex cover P₃ problem parameterized by treewidth. Journal of Combinatorial Optimization, 34(2):414–425, 2017. doi:10.1007/ s10878-016-9999-6.
- [50] David P. Williamson and David B. Shmoys. *The Design of Approximation Algorithms*. Cambridge University Press, 2011. URL: https://www.designofapproxalgs.com.

A Illustration of the complexity dichotomy



Figure 17: Classification of the complexity of $\{H\}$ -M-DELETION for all connected simple graphs H with $2 \leq |V(H)| \leq 5$, according to our results: for the nine graphs on the left, the problem is solvable in time $2^{\Theta(\mathsf{tw})} \cdot n^{\mathcal{O}(1)}$ under the ETH. For the 21 graphs on the right and for all the connected graphs on at least six vertices, the problem is solvable in time $2^{\Theta(\mathsf{tw} \cdot \log \mathsf{tw})} \cdot n^{\mathcal{O}(1)}$ under the ETH.

B An estimation of the constants depending on \mathcal{F} in our algorithm

The main result of this paper is that the \mathcal{F} -M-DELETION problem can be solved in time $2^{\mathcal{O}(f(h) \cdot tw \cdot \log tw)} \cdot n$ on *n*-vertex graphs of treewidth at most tw, for some computable function $f : \mathbb{N} \to \mathbb{N}$, where here *h* is an upper bound on the size of the graphs in \mathcal{F} . This appendix is dedicated to an estimation of an upper bound on the function *f*. Notice that almost all the statements of the results in this paper are accompanied with specific bounds on the involved functions, usually in terms of functions defined in previous statements. However, there are two exceptions. The first one is the function f_6 of Proposition 35, which we discuss in Subsection B.2. The second one is Lemma 27, where no explicit bound for $f_{16}(q, h)$ is given. This is because the existence of $f_{16}(q, h)$ follows by applying Akra-Bazzi Theorem [3] as a black box and this does not provide any estimation of f_{16} . To circumvent this issue, in Subsection B.1 we provide an improved version of Lemma 27, namely Lemma 34, whose proof uses a direct induction, without invoking the Akra-Bazzi Theorem [3, 36]. This alternative proof is strongly based on the proof of [21, Lemma 3.6]. Finally, in Subsection B.2 we provide an upper bound on the constants involving \mathcal{F} in our algorithm. For this, we will use the stronger version of Lemma 27 given in Subsection B.1.

B.1 An improved version of Lemma 27

In this section we provide an improved version of Lemma 27, whose proof is an adaptation of the proof of [21, Lemma 3.6].

Lemma 34. There exists a function $f_{21} : \mathbb{N}^2 \to \mathbb{N}$ such that if $t, q, h \in \mathbb{N}$ and $\mathbf{G} = (G, B, \rho)$ is a K_q -minor-free boundaried graph in $\mathcal{R}_h^{(t)}$, then G contains an $f_{21}(q, h)$ -treewidth modulator that contains B and has at most 2t vertices. Moreover, it holds that $f_{21}(q, h) = \mathcal{O}((f_5(q, f_{13}(f_3(q), 3, 3h), f_{14}(f_3(q), 3h)))^2)$.

Proof. For simplicity, we use q as a shortcut for the triple $(f_3(q), f_{13}(f_3(q), 3, 3h), f_{14}(f_3(q), 3h))$. We define the constants $s = f_5(q, f_{13}(f_3(q), 3, 3h), f_{14}(f_3(q), 3h)), t_0 = 256s^2$, and $c = s \cdot \sqrt{t_0}$. We define the relation \leq_{\circ} so that $a \leq_{\circ} b$ means that $a \leq \max\{0, b\}$. We first prove, by induction on t, the following statement.

Claim: For every non-negative integer t and every K_q -minor-free graph G, if $\mathbf{p}_q(G) \leq t$ then G has a c-treewidth modulator Z with $|Z| \leq_{\circ} t - 16s \cdot \sqrt{t}$.

Proof of claim: In the base case we consider any t with $0 \le t \le t_0$. As $\mathbf{p}_q(G) \le t$, Lemma 10 implies that $\mathsf{tw}(G) \le s \cdot \max\{\sqrt{t}, 1\} \le s \cdot \sqrt{t_0}$. Thus G has a c-treewidth modulator of size $0 \le_{\circ} t - 16s \cdot \sqrt{t}$, and the claim follows.

For the inductive step, let $t > t_0$ and suppose that the claim is true for every t' with $0 \le t' \le t - 1$. We prove that the claim holds also for t. Consider a graph G with $\mathbf{p}_q(G) \le t$ and let S be a set of at most t vertices affecting every q-apex-wall triple of G. Because of Lemma 10, $\mathbf{p}_q(G) \le t$ implies that $\mathsf{tw}(G) \le s \cdot \max\{\sqrt{t}, 1\} = s \cdot \sqrt{t}$.

By applying Lemma 26 to G and S, there is a 2/3-balanced separation (L, R) of S in G such that $|L \cap R| \leq \mathsf{tw}(G) + 1 \leq s \cdot \sqrt{t} + 1$ and there exists some $\alpha \in [\frac{1}{3}, \frac{2}{3}]$ such that $|(L \setminus R) \cap S| \leq \alpha \cdot |S| \leq \alpha \cdot t$ and $|(R \setminus L) \cap S| \leq (1 - \alpha) \cdot |S| \leq (1 - \alpha) \cdot t$.

Since S affects every q-apex-wall triple of G, Lemma 11 gives that the set $L \cap (R \cup S)$ affects every q-apex-wall triple of G[L]. This implies

$$\mathbf{p}_{\mathsf{q}}(G[L]) \le |L \cap (R \cup S)| = |(L \setminus R) \cap S| + |L \cap R| \le \alpha \cdot t + (s \cdot \sqrt{t} + 1) \le \alpha \cdot t + 2s \cdot \sqrt{t}.$$

Here the last inequality follows from the assumption that $t \ge t_0 \ge 1$.

In order to apply the inductive hypothesis, note that $t' := \alpha \cdot t + 2s \cdot \sqrt{t} \le t - 1$ for $t \ge t_0$. This can be verified by using the fact that $s \ge 1$, $\alpha \le \frac{2}{3}$, and checking that $\frac{2}{3}t + 2s \cdot \sqrt{t} \le t - 1$ for $t \ge t_0$. Indeed, the inequality holds whenever $\sqrt{t} \ge 3s + \frac{3}{2}\sqrt{(2s)^2 + \frac{4}{3}}$ and this is the case as $\sqrt{t} \ge \sqrt{t_0} = 16s \ge 3 \cdot s + \frac{3}{2}\sqrt{(2s)^2 + \frac{4}{3}}$.

Therefore we can apply the induction hypothesis to G[L] and t' and obtain a *c*-treewidth modulator Z_L of G[L], such that

$$|Z_L| \leq_{\circ} t' - 16s \cdot \sqrt{t'} \leq_{\circ} (\alpha \cdot t + 2s \cdot \sqrt{t}) - 16s \cdot \sqrt{\alpha \cdot t} + 2s \cdot \sqrt{t} \leq_{\circ} (\alpha \cdot t + 2s \cdot \sqrt{t}) - 16s \cdot \sqrt{\alpha \cdot t}.$$

A symmetric argument applied to G[R] yields a treewidth modulator Z_R of G[R], such that

$$|Z_R| \leq_{\circ} \left((1-\alpha) \cdot t + 2s \cdot \sqrt{t} \right) - 16s \cdot \sqrt{(1-\alpha) \cdot t}.$$

We now construct a c-treewidth modulator Z of G as follows by setting $Z := Z_L \cup (L \cap R) \cup Z_R$. The set Z is a c-treewidth-modulator of G because every connected component of G - Z is a subset of either $(L \setminus (L \cap R)) \setminus Z_L$ or $(R \setminus (L \cap R)) \setminus Z_R$, and Z_L and Z_R are c-treewidth modulators for G[L] and G[R] respectively. Finally we bound the size of Z.

$$\begin{aligned} |Z| &\leq |Z_L| + |Z_R| + |S| \\ &\leq_{\circ} (\alpha \cdot t + 2s \cdot \sqrt{t}) - 16s \cdot \sqrt{t \cdot \alpha} + \left((1 - \alpha) \cdot t + 2s \cdot \sqrt{t} \right) - 16s \cdot \sqrt{t \cdot (1 - \alpha)} + s \cdot \sqrt{t} + 1 \\ &\leq_{\circ} (\alpha \cdot t + 2s \cdot \sqrt{t}) - 16s \cdot \sqrt{t \cdot \alpha} + \left((1 - \alpha) \cdot t + 2s \cdot \sqrt{t} \right) - 16s \cdot \sqrt{t \cdot (1 - \alpha)} + 2s \cdot \sqrt{t} \\ &\leq_{\circ} t - \left(16 \left(\sqrt{\alpha} + \sqrt{1 - \alpha} \right) - 6 \right) \cdot s \cdot \sqrt{t} \\ &\leq_{\circ} t - 16s \cdot \sqrt{t}. \end{aligned}$$

The last inequality uses the fact that $16\left(\sqrt{\alpha} + \sqrt{1-\alpha}\right) - 6 \ge 16$, for every $\alpha \in \left[\frac{1}{3}, \frac{2}{3}\right]$. The claim follows.

Suppose now that $\mathbf{G} = (G, B, \rho)$ is a K_q -minor-free boundaried graph in $\mathcal{R}_h^{(t)}$. From Corollary 25, $\mathbf{p}_q(G) \leq t$ and, because of the above claim, G contains a c-treewidth modulator Z where $|Z| \leq_o t - 16s \cdot \sqrt{t} \leq t$. This, in turn, implies that $B \cup Z$ is a c-treewidth modulator of G that contains B and has size 2t. Also observe that $c = 16s^2$. Therefore, the lemma holds for $f_{21}(q, h) = \mathcal{O}(s^2)$.

We stress that, as it is done in the proof of [21, Lemma 3.6], it is possible to find a modulator of size at most $(1 + \varepsilon) \cdot t$, for every positive real $\varepsilon > 0$. Nevertheless, we have provided the proof of Lemma 34 for the particular case $\varepsilon = 2$, which is enough for our purposes.

B.2 Upper bounds on the constants depending on the excluded minors

In this section we provide an estimation on the function f_{20} in Theorem 33, or equivalently on the constant c_F in Theorem 2. We first provide some definitions in order to introduce the Unique Linkage Theorem [31, 44].

A linkage in a graph G is a subgraph L of G whose connected components are all non-trivial paths. The paths of a linkage L are its connected components and we denote them by $\mathcal{P}(L)$. The size of L is the number of its paths and is denoted by |L|. The terminals of a linkage L, denoted by T(L), are the endpoints of the paths in $\mathcal{P}(L)$, and the pattern of L is the set

 $\{\{s,t\} \mid \mathcal{P}(L) \text{ contains some } (s,t)\text{-path}\}.$

Two linkages L_1, L_2 of G are *equivalent* if they have the same pattern and we denote this fact by $L_1 \equiv L_2$. We say that a linkage L in a graph G is *unique* if for every linkage L' that is equivalent to L it holds that V(L') = V(L).

According to the proof of Proposition 13 in [25], the function f_6 emerges from the following result, known as the Unique Linkage Theorem.

Proposition 35 ([31,44]). There exists a function $f_6 : \mathbb{N}_{\geq 0} \to \mathbb{N}_{\geq 0}$ such that if G is a graph and L is a unique linkage of G, then $\mathsf{tw}(G) \leq f_6(|L|)$.

It is worth mentioning that [31,44] do not provide the precise number of exponentiations involved in the function f_6 , and therefore we will express our upper bounds in terms of this function. Namely, in order to provide an upper bound on $f_{20}(h)$, we backtrack the functions involved in the intermediate results of this paper as follows:

- According to Corollary 32, $f_{20}(h) = \mathcal{O}(f_{19}(h,h) \cdot \log(f_{19}(h,h) \cdot h\sqrt{\log h})).$
- By Theorem 31, $f_{19}(h,h) = \mathcal{O}(f_{17}(h,h) \cdot f_{18}(h,h)).$
- By Lemma 30, $f_{18}(h,h) = 2^{2^{2^{\mathcal{O}(f_{17}(h,h) \cdot \log f_{17}(h,h))}}}$

- By Lemma 29, $f_{17}(h,h) = f_{15}(h,h) \cdot f_{16}(h,h) \cdot 2^{\mathcal{O}(h \log h)}$.
- By Lemma 34, we can take $f_{16}(h, h) = 2$.
- By Lemma 27 (and Lemma 34 as well), $f_{15}(h,h) = \mathcal{O}((f_5(h, f_{13}(h), 3, 3h), f_{14}(f_3(h), 3h)))^2)$.
- By Theorem 5, $f_3(h) = \mathcal{O}(h^{24}) = h^{\mathcal{O}(1)}$.
- By Theorem 23, $f_{14}(f_3(h), 3h) = \mathcal{O}(h^{24}) = h^{\mathcal{O}(1)}$.
- By Theorem 23, $f_{13}(f_3(h), 3, 3h) = \mathcal{O}((f_6(h^{\mathcal{O}(1)}))^3).$
- By Lemma 10, $f_5(h, r, \hat{\ell}) = f_2(h) \cdot r^{2^{\mathcal{O}((h^{24} + \hat{\ell}) \cdot \log(h^{24} + \hat{\ell}))}}$, which is $r^{2^{2^{(h+\hat{\ell})}\mathcal{O}(1)}}$ by Theorem 5. If we set $r = f_{13}(f_3(h), 3, 3h) = \mathcal{O}((f_6(h^{\mathcal{O}(1)}))^3)$ and $\hat{\ell} = f_{14}(f_3(h), 3h) = h^{\mathcal{O}(1)}$, we have that

$$f_{15}(h,h) = \mathcal{O}((f_5(h,r,\hat{\ell}))^2) = \left((f_6(h^{\mathcal{O}(1)}))^3\right)^{2^{2^{(h+h^{\mathcal{O}(1)})^{\mathcal{O}(1)}}} = \left(f_6(h^{\mathcal{O}(1)})\right)^{2^{2^{h^{\mathcal{O}(1)}}}}$$

• We set $\lambda = f_6(h^{\mathcal{O}(1)})$. Given that $f_{17}(h,h) = f_{15}(h,h) \cdot f_{16}(h,h) \cdot 2^{\mathcal{O}(h \log h)}$ and that $f_{16}(h,h) = 2$, we obtain that $f_{17}(h,h) = \lambda^{2^{2^{h^{\mathcal{O}(1)}}}}$. We now have that $f_{18}(h,h) = 2^{2^{2^{\lambda^{2^{2^{h^{\mathcal{O}(1)}}}}}}$, which implies that $f_{19}(h,h) = 2^{2^{2^{\lambda^{2^{2^{h^{\mathcal{O}(1)}}}}}}$ and thus $f_{20}(h) = 2^{2^{2^{\lambda^{2^{h^{\mathcal{O}(1)}}}}}$.

From Theorem 33 and the above discussion, we conclude the following corollary, which gives an explicit upper bound on the contribution of the maximum size of the graphs in \mathcal{F} in the complexity of our algorithm, depending on the function f_6 given by Proposition 35.

Corollary 36. Let \mathcal{F} be a collection of graphs each of size at most h, and let G be a graph. Then the parameter $\mathbf{m}_{\mathcal{F}}(G)$ can be computed in time

$$2^{\left(2^{2^{2^{\lambda^{2^{e^{\lambda^{\mathcal{O}(1)}}}}}}\right)\cdot \mathsf{tw}(G)\cdot \log(\mathsf{tw}(G))}} \cdot |V(G)|, \text{ where } \lambda = f_6(h^{\mathcal{O}(1)})$$