



HAL
open science

A complexity dichotomy for hitting connected minors on bounded treewidth graphs: the chair and the banner draw the boundary

Julien Baste, Ignasi Sau, Dimitrios M. Thilikos

► **To cite this version:**

Julien Baste, Ignasi Sau, Dimitrios M. Thilikos. A complexity dichotomy for hitting connected minors on bounded treewidth graphs: the chair and the banner draw the boundary. SODA 2020 - 31st Annual ACM-SIAM Symposium on Discrete Algorithms, Jan 2020, Salt Lake City, UT, United States. pp.951-970, 10.1137/1.9781611975994.57 . lirmm-02991215

HAL Id: lirmm-02991215

<https://hal-lirmm.ccsd.cnrs.fr/lirmm-02991215>

Submitted on 5 Nov 2020

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.



Distributed under a Creative Commons Attribution 4.0 International License

A complexity dichotomy for hitting connected minors on bounded treewidth graphs: the chair and the banner draw the boundary

Julien Baste*

Ignasi Sau†

Dimitrios M. Thilikos†,‡

Abstract

For a fixed connected graph H , the $\{H\}$ -M-DELETION problem asks, given a graph G , for the minimum number of vertices that intersect all minor models of H in G . It is known that this problem can be solved in time $f(\text{tw}) \cdot n^{\mathcal{O}(1)}$, where tw is the treewidth of G . We determine the asymptotically optimal function $f(\text{tw})$, for each possible choice of H . Namely, we prove that, under the ETH, $f(\text{tw}) = 2^{\Theta(\text{tw})}$ if H is a contraction of the chair or the banner, and $f(\text{tw}) = 2^{\Theta(\text{tw} \cdot \log \text{tw})}$ otherwise. Prior to this work, such a complete characterization was only known when H is a planar graph with at most five vertices. For the upper bounds, we present an algorithm in time $2^{\Theta(\text{tw} \cdot \log \text{tw})} \cdot n^{\mathcal{O}(1)}$ for the more general problem where all minor models of connected graphs in a finite family \mathcal{F} need to be hit. We combine several ingredients such as the machinery of bounded treewidth graphs in dynamic programming via representatives, the Flat Wall Theorem, Bidimensionality, the irrelevant vertex technique, treewidth modulators, and protrusion replacement. In particular, this algorithm vastly generalizes a result of Jansen et al. [SODA 2014] for the particular case $\mathcal{F} = \{K_5, K_{3,3}\}$. For the lower bounds, our reductions are based on a generic construction building on the one given by the authors in [IPEC 2018], which uses the framework introduced by Lokshtanov et al. [SODA 2011] to obtain superexponential lower bounds.

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 NP-hard [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 is not bounded). Since the property of containing a graph as a minor can be expressed in Monadic Second Order logic [34], by Courcelle's theorem [13], \mathcal{F} -M-DELETION can be solved in time $\mathcal{O}^*(f(\text{tw}))$ on graphs with treewidth at most tw , where f is some computable function¹. As the function $f(\text{tw})$ given by Courcelle's theorem is typically enormous, our goal is to determine, for a fixed collection \mathcal{F} , which is the *best possible* such function 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 [47] algorithms.

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 $\mathcal{O}^*(2^{\mathcal{O}(\text{tw})})$, called *single-exponential*, by standard dynamic programming techniques, and no algorithm with running time $\mathcal{O}^*(2^{o(\text{tw})})$ exists, unless the Exponential Time Hypothesis (ETH)² fails [27]. For FEEDBACK VERTEX SET, standard dynamic programming techniques give a running time of $\mathcal{O}^*(2^{\mathcal{O}(\text{tw} \cdot \log \text{tw})})$, while the

*Institute of Optimization and Operations Research, Ulm University, Ulm, Germany. julien.baste@uni-ulm.de.

†CNRS, LIRMM, Université de Montpellier, Montpellier, France. Supported by projects DEMOGRAPH (ANR-16-CE40-0028) and ESIGMA (ANR-17-CE23-0010). Emails: ignasi.sau@lirmm.fr, sedthilk@thilikos.info.

‡Supported by the Research Council of Norway and the French Ministry of Europe and Foreign Affairs, via the Franco-Norwegian project PHC AURORA 2019.

The full version of this article is permanently available at <https://arxiv.org/abs/1907.04442>.

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

²The ETH states that 3-SAT on n variables cannot be solved in time $2^{o(n)}$; see [27] for more details.

lower bound under the ETH [27] is again $\mathcal{O}^*(2^{o(\text{tw})})$. This gap remained open for a while, until Cygan et al. [15] presented an optimal (randomized) algorithm running in time $\mathcal{O}^*(2^{\mathcal{O}(\text{tw})})$, 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 [8, 20].

Concerning VERTEX PLANARIZATION, Jansen et al. [28] presented an algorithm running in time $\mathcal{O}^*(2^{\mathcal{O}(\text{tw} \cdot \log \text{tw})})$ as a crucial subroutine in an algorithm running in time $\mathcal{O}^*(2^{\mathcal{O}(k \cdot \log k)})$. Marcin Pilipczuk [41] proved afterwards that this running time is optimal under the ETH, by using the framework introduced by Lokshтанov et al. [38] for proving superexponential lower bounds.

Generalizing the above algorithm, the main technical contribution of the recent paper of Kociumaka and Pilipczuk [35] is an algorithm running in time $\mathcal{O}^*(2^{\mathcal{O}((\text{tw}+g) \cdot \log(\text{tw}+g))})$ 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 .

In a recent pair of papers [5, 6], we initiated a systematic study of the complexity of \mathcal{F} -M-DELETION, parameterized by treewidth³. Before stating these results, we say that a collection \mathcal{F} is *connected* if it contains only connected graphs. In [5] we showed that, for every \mathcal{F} , \mathcal{F} -M-DELETION can be solved in time $\mathcal{O}^*(2^{2^{\mathcal{O}(\text{tw} \cdot \log \text{tw})}})$, and that if \mathcal{F} is connected and contains a planar graph, the running time can be improved to $\mathcal{O}^*(2^{\mathcal{O}(\text{tw} \cdot \log \text{tw})})$. If the input graph G is planar or, more generally, embedded in a surface of bounded genus, and \mathcal{F} is connected, then the running time can be further improved to $\mathcal{O}^*(2^{\mathcal{O}(\text{tw})})$. We also provided single-exponential algorithms for the cases where $\mathcal{F} \in \{\{P_3\}, \{P_4\}, \{C_4\}\}$. Concerning lower bounds under the ETH, we proved that for any connected \mathcal{F} , \mathcal{F} -M-DELETION cannot be solved in time $\mathcal{O}^*(2^{o(\text{tw})})$, even if the input graph G is planar. Inspired by the reduction of Pilipczuk [41], we proved that the problem cannot be solved in time $\mathcal{O}^*(2^{o(\text{tw} \cdot \log \text{tw})})$ for some families of collections \mathcal{F} , for example, when all graphs in \mathcal{F} are planar and 3-connected. In the subsequent paper [6], we focused on small planar graphs. Namely, we classified the optimal asymptotic complexity of $\{H\}$ -M-DELETION when H




Figure 1: The chair and the banner.

is a connected *planar* graph on at most five vertices. To achieve that, we provided single-exponential algorithms for a number of small patterns not considered in [5] and superexponential lower bounds for the remaining cases, this time inspired by a reduction of Bonnet et al. [11] for generalized feedback vertex set problems. Full proofs of the results in [5, 6] are available at [4].

Our results. In this article we make significant steps towards a complete classification of the complexity of the \mathcal{F} -M-DELETION problem parameterized by treewidth, by improving both the known upper and lower bounds. Namely, we prove the following results:

► Our main contribution is an algorithm to solve \mathcal{F} -M-DELETION in time $\mathcal{O}^*(2^{\mathcal{O}(\text{tw} \cdot \log \text{tw})})$ for every connected collection \mathcal{F} , hence dropping the condition that it contains a planar graph, which was critically needed in the algorithm presented in [5] in order to bound the treewidth of an \mathcal{F} -minor-free graph. Besides largely improving our previous results [5, 6], this algorithm also generalizes the one for $\mathcal{F} = \{K_5, K_{3,3}\}$ given by Jansen et al. [28], which is based on embeddings. It can be interpreted as an exponential “collapse” of the natural dynamic programming algorithm running in time $\mathcal{O}^*(2^{2^{\mathcal{O}(\text{tw} \cdot \log \text{tw})}})$ given in [5]. The algorithm is quite involved, and we provide an overview of it in §2.

► Concerning lower bounds, we vastly improve all previous super-exponential lower bounds [5, 6, 41] for \mathcal{F} -M-DELETION by proving that for every connected graph H that is *not* a contraction of the chair or the banner, depicted in Figure 1, $\{H\}$ -M-DELETION cannot be solved in time $\mathcal{O}^*(2^{o(\text{tw} \cdot \log \text{tw})})$ under the ETH. We also prove a lower bound of $\mathcal{O}^*(2^{o(\text{tw} \cdot \log \text{tw})})$ for \mathcal{F} -M-DELETION when \mathcal{F} is any finite non-empty subset of all connected graphs that contain a block with at least five edges. In particular, the former result applies to K_5 and all the connected graphs with at least six vertices.

Our reductions are based on a generic framework that generalizes the one given in [6], which was inspired by a reduction of Bonnet et al. [11]. These lower bounds also subsume the ones in [5], which were proved using a different reduction inspired by the one of Pilipczuk [41]. More precisely, in [6] we proved subexponential lower bounds for P_5 , $K_{1,i}$ with $i \geq 4$, $K_{2,i}$ and θ_i for $i \geq 3$ (θ_i is the graph consisting of two vertices and i parallel edges), and the following graphs (see the full version

³In these papers [5, 6], we also considered the version of the problem where the graphs in \mathcal{F} are forbidden as *topological* minors; in the current paper we will focus exclusively on the minor version.

for a figure containing these graphs): the px, the kite, the dart, the bull, the butterfly, the cricket, and the co-banner. All these reductions were based on a general construction, which is a less general version of the construction that we present here, and then we needed particular small gadgets to deal with each of the graphs. Here we present a more general version of this approach that has the following advantage: in order to prove lower bounds, we need to distinguish cases not depending on *particular* instantiations of the collection \mathcal{F} , but on *general structural properties* of the graphs in \mathcal{F} , like containing a block with at least five edges or the number and relative position of cut vertices and cycles.

The above results, together with the lower and upper bounds for planar graphs on at most five vertices given in [5, 6] and the known cases $\mathcal{F} = \{P_2\}$ [14, 27], $\mathcal{F} = \{P_3\}$ [2, 46], and $\mathcal{F} = \{C_3\}$ [8, 15] 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 1.1. *Let H be a connected graph. Under the ETH, $\{H\}$ -M-DELETION is solvable in time⁴*

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

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 $\mathcal{O}^*(2^{\Theta(\mathbf{tw} \cdot \log \mathbf{tw})})$ and that the chair and the banner are the only graphs on at least five vertices for which the problem is solvable in time $\mathcal{O}^*(2^{\Theta(\mathbf{tw})})$.

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 the chair and the banner. 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 [4].

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

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 $\mathcal{O}^*(2^{\Theta(\mathbf{tw})})$ to $\mathcal{O}^*(2^{\Theta(\mathbf{tw} \cdot \log \mathbf{tw})})$.

Organization of the paper. In §2 we provide a high-level overview of the algorithm running in time $\mathcal{O}^*(2^{\Theta(\mathbf{tw} \cdot \log \mathbf{tw})})$. In §3 we give some preliminaries. In §4 we deal with flat walls, in §5 we apply the irrelevant vertex technique in the context of boundaried graphs, and in §6 we use this in order to bound the size of the dynamic programming tables. The lower bounds can be found in the full version. We conclude the article in §7. Due to space limitations, the proofs of all the results marked with ‘(★)’ can be found in the full version.

2 Overview of the algorithm

In order to obtain our algorithm of time $\mathcal{O}^*(2^{\Theta(\mathbf{tw} \cdot \log \mathbf{tw})})$ for every connected 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. [9] and subsequently used, for instance, in [5, 21, 23, 34]. 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 (see §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 [5]) 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 details). 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 ’ 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, it is enough to prove that, for every $R \in \mathcal{R}_h^{(t)}$, it holds that

$$(2.1) \quad |V(R)| = \mathcal{O}_h(t).$$

Note that this is indeed sufficient as there are at most $\binom{|V(R)|}{|E(R)|} = 2^{\mathcal{O}_h(|V(R)| \log |V(R)|)}$ representatives. In order to prove (2.1), we combine a number of different techniques, which we proceed to discuss informally, and that are schematically summarized in Figure 2:




Figure 2: Diagram of the algorithm in time $\mathcal{O}^*(2^{\mathcal{O}(tw \cdot \log tw)})$ for connected \mathcal{F} .

► We use the Flat Wall Theorem of Robertson and Seymour [44], in particular the recent optimized versions by Kawarabayashi et al. [32] and by Chuzhoy [12]. 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 $\Omega_h(\text{tw}(G))$. Here, the definition of “flat wall” is quite involved and is detailed in §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*.

► 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 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 [16, 18], in the sense that its value on a flat wall depends quadratically on the height of the wall (Lemma 4.1) and *separable* [9, 18, 21] (Lemma 4.2).

► 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 “bordered” minors of *detail* at most h (the detail of a bordered 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 [44]. Using standard “zooming” arguments, we can prove that, given a flat wall, we can find a large h -homogeneous subwall inside it (Lemma 4.3). Homogeneous subwalls are very useful because, as we explain below, they permit the application of the irrelevant vertex technique adapted to our purposes.

► The most complicated step towards proving (2.1) is to find an “irrelevant” vertex inside a sufficiently large (in terms of h) flat wall of a bordered graph that is *not* affected by its boundary (Theorem 5.2). 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 §5 for the precise definition. The *irrelevant vertex technique* originated in the seminal work of Robertson and Seymour [44, 45] and has become a very useful tool used in various kinds of linkage and cut problems [1, 28, 35, 36, 42]. 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 [44, (10.2)] and its subsequent proof in [45]). 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 \tilde{W} , which we call its *leveling*; cf. Subsection 4.3 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 1). 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 [29] that occurs as a subgraph inside a flat wall (Proposition 5.1) and that has the following nice property, recently proved in [26] (Proposition 5.2): 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 5.1) 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 in [26] mentioned above, we can finally prove (Theorem 5.1) 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 5.2). 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 5.2, imply that if $R \in \mathcal{R}_h^{(t)}$ is a minimum-sized representative, then its boundary affects all large enough 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 5.1).

► 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 6.2 (whose proof is an adaptation of [22, Lemma 3.6] – see also [21]) that every representative $R \in \mathcal{R}_h^{(t)}$ has a vertex subset S containing its boundary, with $|S| \leq 2t$, whose removal leaves a graph of treewidth bounded by a function of h ; such a set is called a *treewidth modulator*.

► Once we have a treewidth modulator of size $\mathcal{O}(t)$ of a representative R , all that remains is to pipeline it with known techniques to compute an appropriate *protrusion decomposition* [34] (Lemma 6.3) and to *reduce protrusions* to smaller equivalent ones of size bounded by a function of h – we use the version given in [5] adapted to the \mathcal{F} -M-DELETION problem – (Lemma 6.4), implying that $|V(R)| = \mathcal{O}_h(t)$ for every every $R \in \mathcal{R}_h^{(t)}$ and concluding the proof of (2.1).

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 2.1. *Let \mathcal{F} be a finite non-empty collection of non-empty connected graphs. There exists a constant*

$c_{\mathcal{F}}$ such that the \mathcal{F} -M-DELETION problem is solvable in time $c_{\mathcal{F}}^{\text{tw} \cdot \log \text{tw}} \cdot n$ on n -vertex graphs of treewidth at most tw .

In the full version we provide an estimation of the constant $c_{\mathcal{F}}$ in Theorem 2.1 based on the parametric dependencies of the Unique Linkage Theorem [33, 45].

3 Preliminaries

Sets and integers. We denote by \mathbb{N} the set of non-negative integers and we set $\mathbb{N}^+ = \mathbb{N} \setminus \{0\}$. Given two integers p and q , the set $[p, q]$ refers to the set of every integer r such that $p \leq r \leq q$. For an integer $p \geq 1$, we set $[p] = [1, p]$ and $\mathbb{N}_{\geq p} = \mathbb{N} \setminus [0, p-1]$. In the set $[1, k] \times [1, k]$, a *row* is a set $\{i\} \times [1, k]$ and a *column* is a set $[1, k] \times \{i\}$ for some $i \in [1, k]$. For a set S , we denote by 2^S the set of all the subsets of S .

Graphs. All the graphs that we consider in this paper are undirected, finite, and without loops or multiple edges. We use standard graph-theoretic notation, and we refer the reader to [17] for any undefined terminology. Let G be a graph. 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 the value $|L \cap R|$. Given a graph G and a set $S \subseteq V(G)$, we denote by $\partial(S)$ the set of vertices in S that have a neighbor in $V(G) \setminus S$. For $S \subseteq V(G)$, we use the shortcut $G \setminus S$ to denote $G[V(G) \setminus S]$. 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 $N_G(u) \cup N_G(v) \setminus \{u, v\}$, where $N_G(u)$ denotes the set of neighbors of u in G . A graph G' is a *minor* of a graph G if G' can be obtained from G by a sequence of vertex removals, edge removals, and edge contractions. If only edge contractions are allowed, we say that G' is a *contraction* of 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 $\text{tw}(G)$.

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

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 $\rho : B \rightarrow [t]$ is a bijection. We say that $\mathbf{G}_1 = (G_1, B_1, \rho_1)$ 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 [44], we define the *detail* of a boundaried graph $\mathbf{G} = (G, B, \rho)$ as $\text{detail}(\mathbf{G}) := \max\{|E(G)|, |V(G) \setminus B|\}$. We denote by $\mathcal{B}^{(t)}$ the set of all (pairwise non-isomorphic) t -boundaried graphs and by $\mathcal{B}_h^{(t)}$ 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 $\mathbf{G}_1 \preceq_m \mathbf{G}_2$, 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 $\text{diss}(M, T)$ 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 if 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 G if H is isomorphic to $\text{diss}(M, T)$. We denote this isomorphism by $\sigma_{M, T} : V(H) \rightarrow T$. 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 $\sigma_{M, T}(H')$ 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 $B \subseteq T$, we call (\mathbf{M}, T) a *btm-pair* and we define $\text{diss}(\mathbf{M}, T) = (\text{diss}(M, T), B, \rho)$. 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 $\text{diss}(\mathbf{M}, T)$ is isomorphic to \mathbf{G}_1 .

Given a $\mathbf{G} = (G, B, \rho) \in \mathcal{B}$, we define $\text{ext}(\mathbf{G})$ 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 B and ρ are the same for all graphs in $\text{ext}(\mathbf{G})$. Moreover, we do not consider isomorphic boundaried graphs in $\text{ext}(\mathbf{G})$ as different graphs. The set $\text{ext}(\mathbf{G})$ helps us to express the minor relation in terms of the topological minor relation because of the following simple observation.

OBSERVATION 1. *If $\mathbf{G}_1, \mathbf{G}_2 \in \mathcal{B}$, then $\mathbf{G}_1 \preceq_m \mathbf{G}_2 \iff \exists \mathbf{G} \in \text{ext}(\mathbf{G}_2) : \mathbf{G}_1 \preceq_{tm} \mathbf{G}$. Moreover, there is a function $f_1 : \mathbb{N} \rightarrow \mathbb{N}$ such that if \mathbf{G} is a boundaried graph with detail h , then every graph in $\text{ext}(\mathbf{G})$ has detail at most $f_1(h)$.*

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

$$h\text{-folio}(\mathbf{G}) = \{\mathbf{G}' \in \mathcal{B} \mid \mathbf{G}' \preceq_{tm} \mathbf{G} \wedge \mathbf{G}' \text{ has detail } \leq h\}.$$

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 [5]).

LEMMA 3.1. *For every $t, h \in \mathbb{N}$, $|\{h\text{-folio}(\mathbf{G}) \mid \mathbf{G} \in \mathcal{B}_h^{(t)}\}| = 2^{2^{\mathcal{O}((h+t) \cdot \log(h+t))}}$.*

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

$$(3.2) \quad H \preceq_m \mathbf{F} \oplus \mathbf{G}_1 \iff H \preceq_m \mathbf{F} \oplus \mathbf{G}_2.$$

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 *representative* of \equiv_h . A *set of t-representatives* for \equiv_h is a collection

containing a minimum-sized representative for each equivalence class of \equiv_h . Given $t, h \in \mathbb{N}$, we denote by $\mathcal{R}_h^{(t)}$ a set of *t-representatives* for \equiv_h .

Let \mathcal{F} be a finite non-empty collection of non-empty graphs; we call such a collection *proper*. A proper collection is called *connected* if all its graphs are connected. We extend the minor relation to \mathcal{F} such that, given a graph G , $\mathcal{F} \preceq_m G$ if and only if there exists a graph $H \in \mathcal{F}$ such that $H \preceq_m G$. We also denote $\text{ex}_m(\mathcal{F}) = \{G \mid \mathcal{F} \not\preceq_m G\}$, i.e., $\text{ex}_m(\mathcal{F})$ is the class of graphs that do not contain any graph in \mathcal{F} as a minor.

Definition of the problem. 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:

$$\mathbf{m}_{\mathcal{F}}(G) = \min\{|S| \mid S \subseteq V(G) \wedge G \setminus S \in \text{ex}_m(\mathcal{F})\}.$$

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 $\mathbf{m}_{\mathcal{F}}(G) \leq k$?

At this point, we wish to stress that the folio-equivalence defined in (3.2) 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 $h\text{-folio}(\mathbf{G}_1) = h\text{-folio}(\mathbf{G}_2)$ 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 [5], which has a double-exponential parametric dependence due to the bound of Lemma 3.1. 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 [5] 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 Subsection 4.1 we define them and we state the Flat Wall Theorem of Robertson and Seymour [44], using the terminology of [32]. In Subsection 4.2 we define a graph parameter related to flat walls and prove that it behaves bidimensionally and is separable. Finally, in Subsection 4.3 we define homogeneous subwalls and prove that a flat wall contains a large homogeneous subwall.

4.1 The Flat Wall Theorem Before defining flat walls, we need to introduce walls and renditions, following the recent framework of [32].

Walls. Let $k, r \in \mathbb{N}$. The $(k \times r)$ -grid is the Cartesian product of two paths on k and r vertices, respectively. An *elementary r -wall*, for some odd $r \geq 3$, 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 $r \geq 3$, 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 $(1, 1), (2, r), (2r - 1, 1), (2r, r)$ are called *corners* (notice that the corners are also pegs).

Given an elementary r -wall \overline{W} , some $i \in \{1, 3, \dots, 2r - 1\}$, and $i' = (i + 1)/2$, the i' -th *vertical path* of \overline{W} is the one whose vertices, in order of appearance, are $(i, 1), (i, 2), (i + 1, 2), (i + 1, 3), (i, 3), (i, 4), (i + 1, 4), (i + 1, 5), (i, 5), \dots, (i, r - 2), (i, r - 1), (i + 1, r - 1), (i + 1, r)$. 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 $(1, j), (2, j), \dots, (2r, j)$. The first horizontal path is the one containing the vertices $(1, 1), (2, 1), \dots, (2r - 1, 1)$ while the r -th horizontal path is the one containing the vertices $(1, r), (2, r), \dots, (2r, r)$. We call these two last paths the *lowest* and the *highest* paths of W , respectively.

An r -wall is any graph W obtained from an elementary r -wall \overline{W} after subdividing edges. The following theorem of Kawarabayashi and Kobayashi [31] provides a *linear* relation between the treewidth and the height of a largest wall in a minor-free graph.

THEOREM 4.1. *There is a function $f_2 : \mathbb{N} \rightarrow \mathbb{N}$ such that, for every $q, r \in \mathbb{N}$ and every K_q -minor-free graph G , if $\text{tw}(G) \geq f_2(q) \cdot r$, then G contains an r -wall. In*

particular, one may choose $f_2(q) = 2^{\mathcal{O}(q^2 \log q)}$.

We call the vertices of an r -wall W that were added after the subdivision operations *subdivision vertices*, while we call the rest of the vertices (i.e., those of \overline{W}) *branch vertices*. A cycle of W is a *brick* (resp. the *perimeter*) of W if its branch vertices are the vertices of a brick (resp. the perimeter) of \overline{W} . We denote by $\mathcal{C}(W)$ the set of all cycles of W , by $\text{bricks}(W)$ the set of all the bricks of W , and we use $D(W)$ in order to denote the perimeter of the wall W .

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

Given an r -wall W , 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, $C \subseteq P$). A subgraph W of a graph G is called a *wall* of G if W is an r -wall for some odd $r \geq 3$ and we refer to r as the *height* of the wall W .

The *layers* of an r -wall W are recursively defined as follows. The first layer of W is its perimeter. For $i = 2, \dots, (r - 1)/2$, 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 branch vertices that do not belong to any layer. See Figure 3 for an illustration of these notions.




Figure 3: An 11-wall and its five layers.

Renditions. Let Δ be a closed disk. Given a subset X of Δ , we denote its closure by \overline{X} and its boundary by $\text{bd}(X)$. A Δ -*painting* is a pair $\Gamma = (U, N)$ where N is a finite set of points of Δ , $N \subseteq U \subseteq \Delta$, $U \setminus N$ has finitely many arcwise-connected components, called *cells*, such that, for every cell c , \bar{c} is a closed disk,

$\text{bd}(c) \cap \Delta \subseteq N$, and $|\text{bd}(c) \cap N| \leq 3$. We use the notation $U(\Gamma) := U$, $N(\Gamma) := N$, and denote the set of cells of Γ by $C(\Gamma)$.

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

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

- Γ is a Δ -painting for some closed disk Δ ,
- $\pi : N(\Gamma) \rightarrow V(G)$ is an injection, and
- 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 $c \in C(\Gamma)$, $\pi(c \cap N) \subseteq V(\sigma(c))$,
 - (4) for every cell $c \in C(\Gamma)$, $V(\sigma(c)) \cap \bigcup_{c' \in C(\Gamma) \setminus \{c\}} V(\sigma(c')) \subseteq \pi(c \cap N)$, and
 - (5) $\pi(N(\Gamma) \cap \text{bd}(\Delta)) = V(\Omega)$, such that the points in $N(\Gamma) \cap \text{bd}(\Delta)$ appear in $\text{bd}(\Delta)$ in the same ordering as their images, via π , in Ω .

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

1. for every $c \in C(\Gamma)$ there is a path in $\sigma(c)$ between any two vertices in $\pi(c \cap N)$ and
2. there is no other Ω -rendition of G satisfying Condition 1 with smaller number of cells.

Flat walls. Let G be a graph and let W be a wall a of G . We say that W is a *flat 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 $G[Y]$.

Given a flat wall W of a graph G as above, we call $G[Y]$ the *compass* of W in G , denoted by $\text{compass}(W)$. We call $X \cap Y$ the *frontier* of W . We call the set $\text{ground}(W) := \pi(N(\Gamma))$ *ground set* of W . We clarify that $\text{ground}(W)$ consists of vertices of the compass of W that are not necessarily vertices of W . We also call the graphs

in $\text{flaps}(W) := \{\sigma(c) \mid c \in C(\Gamma)\}$ *flaps* of the wall W . For each flap $F \in \text{flaps}(W)$ we define its *base*⁵ as the set $\partial F := V(F) \cap \text{ground}(W)$. We also refer to the triple (Γ, σ, π) as a *rendition* of the compass of W in G . We always assume that this rendition is a tight one. Based on this assumption and by using Menger's theorem, it is easy to prove the following.

OBSERVATION 2. *Let W be a flat wall of a graph G and let K be the compass of W . For every flap F of W , there exists $|\partial F|$ pairwise vertex-disjoint paths in K from ∂F to the frontier of W . Moreover, any two vertices in ∂F are connected by a path in F .*

Let G be a graph. We say that a pair (A, W) is an (a, r) -*apex-wall pair* of G if $A \subseteq V(G)$, $|A| \leq a$ and W is a flat r -wall of $G \setminus A$. We are now ready to state the Flat Wall Theorem, first proved by Robertson and Seymour [44] and then reproved by Kawarabayashi et al. [32] and Chuzhoy [12]. The version we state here is by Chuzhoy [12].

THEOREM 4.2. *There is a constant $c_1 \in \mathbb{N}$ such that, for every odd $r \geq 3$ and every $q \in \mathbb{N}$, every graph G that*

- is K_q -minor-free and
- contains a z -wall where $z = c_1 \cdot (q \cdot (r + q))$,

contains a $(q - 5, r)$ -apex-wall pair (A, W) .

4.2 Affecting flat walls We proceed to define a graph parameter and then prove that it is bidimensional and separable.

Let G be a graph, (A, W) be an (a, r) -apex-wall pair of G , and $S \subseteq V(G)$. We say that S *affects* (A, W) if $N_G[V(\text{compass}(W))] \cap (S \setminus A) \neq \emptyset$. For $a, r \in \mathbb{N}$, we define

$$\mathbf{p}_{a,r}(G) = \min\{k \mid \exists S \subseteq V(G) : |S| \leq k \wedge S \text{ affects every } (a, r)\text{-apex-wall pair } (A, W) \text{ of } G\}.$$

Using Theorem 4.1 and Theorem 4.2, we prove that the above parameter grows *quadratically* with the height of a largest wall, or equivalently, by Theorem 4.1, with its treewidth.

LEMMA 4.1. (*) *There is a function $f_3 : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that if $q, r \in \mathbb{N}_{>0}$, and G is a K_q -minor-free graph, then $\mathbf{tw}(G) \leq f_3(q, r) \cdot \max\{1, \sqrt{\mathbf{p}_{q-5,r}(G)}\}$. In particular, one may choose $f_3(q, r) = r \cdot 2^{\mathcal{O}(q^2 \log q)}$.*

⁵We point out that this definition differs from the notation $\partial(S)$ for a set of vertices defined in §3.

We now prove that the parameter $\mathbf{p}_{a,r}$ 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 4.2. *Let $a, r \in \mathbb{N}$, let G be a graph, and let $S \subseteq V(G)$ such that S affects every (a, r) -apex-wall pair of G . Then, for every separation (L, R) of S in G , the set $L \cap (R \cup S)$ affects every (a, r) -apex-wall pair of $G[L]$.*

Proof. Suppose for contradiction that (A, W) is an (a, r) -apex-wall pair of $G[L]$ that is not affected by $L \cap (R \cup S)$. In particular, it holds that $V(\text{compass}(W)) \subseteq L \setminus R$. Since by assumption (A, W) 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 $V(\text{compass}(W)) \subseteq L \setminus R$, contradicting the hypothesis that (L, R) is a separation of G . \square

4.3 Homogeneous subwalls Before defining homogeneous subwalls, we need to define partially disk-embedded graphs and introduce the concept of a *leveling* of a flat wall. This concept can be seen as a way to capture the “plane structure” of a flat wall in terms of a graph embeddable in a disk and might be useful in further applications of the Flat Wall Theorem.

Partially disk-embedded graphs. A *closed disk* (resp. *open disk*) Δ is a set homeomorphic to the set $\{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 \leq 1\}$ (resp. $\{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 < 1\}$). A *disk* of Δ is a closed or an open disk that is a subset of Δ . We say that a graph G is *partially disk-embedded in some closed disk Δ* , if there is some subgraph K of G that is embedded in Δ such that $(V(G) \cap \Delta, V(G) \setminus \text{int}(\Delta))$ is a separation of G , where int is used to denote the interior of a subset of the plane. 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 . We also call the graph $K = G \cap \Delta$ *compass* of the Δ -embedded graph G and we always assume that G is accompanied by an embedding of its compass in Δ , that is the set $G \cap \Delta$. We say that G is a Δ -embedded graph if it is partially Δ -embedded graph and $G \subseteq \Delta$ (the whole G is embedded in Δ).

Levelings. Let W be a flat wall of a graph G . We define the *leveling* of W in G , denoted by \tilde{W} , as the bipartite graph where one part is the ground set of W , the other part is the set of flaps of W , and, given a pair $(v, F) \in \text{ground}(W) \times \text{flaps}(W)$, the set $\{v, F\}$ is an edge of \tilde{W} if and only if $v \in \partial F$. Again, keep in mind that \tilde{W} may contain (many) vertices that are not in W . Notice that the incidence graph of the plane hypergraph $(N(\Gamma), \{c \cap N \mid c \in C(\Gamma)\})$ is isomorphic to \tilde{W} via an

isomorphism that extends π and, moreover, bijectively corresponds cells to flaps. This permits us to treat \tilde{W} as a Δ -embedded graph where $\text{bd}(\Delta) \cap \tilde{W}$ is the frontier of W . We call the vertices of $\text{ground}(\tilde{W})$ (resp. $\text{flaps}(\tilde{W})$) *ground-vertices* (resp. *flap-vertices*) of \tilde{W} .

Recall that each edge of $\text{compass}(W)$ belongs to exactly one flap of W . If both of the endpoints of this edge are in the boundary of this flap, then we say that this edge is a *short edge* of $\text{compass}(W)$. We define the graph W^\bullet as the graph obtained from W if we subdivide once every short edge in W .

The next observation, which is used in the proof of Theorem 5.2, is a consequence of the following three facts: flap-vertices of \tilde{W} have degree at most three, all the vertices of a wall have degree at most three, and every separation (A, B) of order at most three of a wall is trivial.

OBSERVATION 3. *If W is a flat wall of a graph G , then the leveling \tilde{W} of W in G contains a subgraph W' that is isomorphic to some subdivision of W^\bullet via an isomorphism that maps each ground vertex to itself.*

We call the graph W' as in Observation 3 *representation* of the flat wall W in \tilde{W} and we see it as a Δ -embedded subgraph of \tilde{W} . Notice that the above observation permits to bijectively map each cycle of W to a cycle of W' that is also a cycle of \tilde{W} . That way, each cycle C of W corresponds to a cycle C' of W' denoted by C' and we call C' as the *representation* of C in \tilde{W} . From now on, we reserve the “ $\tilde{\cdot}$ ”-notation to denote the correspondence between W (resp. C) and W' (resp. C'). We define the function $\text{flaps} : \mathcal{C}(W) \rightarrow 2^{\text{flaps}(W)}$ so that, for each cycle C of W , $\text{flaps}(C)$ contains each flap F of W that, when seen as a flap-vertex of the Δ -embedded graph \tilde{W} , belongs to the closed disk bounded by C' .

Homogeneous subwalls. Let G be a graph and W be a flat wall of G . Let also (Γ, σ, π) be a rendition of the compass of W in G . Recall that $\Gamma = (U, N)$ is a Δ -painting for some closed disk Δ . Given a flap F , we denote by $\Omega(F)$ the counter-clockwise ordering of the vertices of ∂F as they appear in the corresponding cell of $C(\Gamma)$. Notice that as $|\partial F| \leq 3$, 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., $(v_1, v_2, v_3) \equiv (v_2, v_3, v_1)$ but not under inversion, i.e., $(v_1, v_2, v_3) \not\equiv (v_3, v_2, v_1)$.

Let G be a graph and let (A, W) be an (a, r) -apex-wall pair of G . For each cell $F \in \text{flaps}(W)$ with $t_F = |\partial F|$, we fix $\rho_F : \partial F \rightarrow [a+1, a+t_F]$ such that $(\rho_F^{-1}(a+1), \dots, \rho_F^{-1}(a+t_F)) \equiv \Omega(c)$. We also fix a bijection $\rho_A : A \rightarrow [a]$. For each flap $F \in \text{flaps}(W)$ we define the *boundaried graph $\mathbf{F}^A := (G[A \cup F], A \cup \partial F, \rho_A \cup \rho_F)$* and

we denote by F^A the underlying graph of \mathbf{F}^A . Notice that $G[V(\text{compass}(W)) \cup A] = \bigcup_{F \in \text{flaps}(W)} F^A$.

Given some $\ell \in \mathbb{N}$, we say that two flaps $F_1, F_2 \in \text{flaps}(W)$ are (A, ℓ) -equivalent, denoted by $F_1 \sim_{A, \ell} F_2$, if

$$\ell\text{-folio}(\mathbf{F}_1^A) = \ell\text{-folio}(\mathbf{F}_2^A).$$

For each $F \in \text{flaps}(W)$, we define the (ℓ, a) -color of F , denoted by (ℓ, a) -color(F), as the equivalence class of $\sim_{A, \ell}$ to which \mathbf{F}^A belongs.

Let \tilde{W} be the leveling of W in $G \setminus A$ and let W' be the representation of W in \tilde{W} . Recall that \tilde{W} is a Δ -embedded graph. For each cycle C of W , we define the (ℓ, a) -palette of C , denoted by (ℓ, a) -palette(C), as the set of all the (ℓ, a) -colors of the flaps that appear as vertices of \tilde{W} in the closed disk bounded by C' in Δ (recall that by C' we denote the representation of C in W). Let \tilde{W} be a subwall of W . We say that \tilde{W} is an (ℓ, a) -homogeneous subwall of W if every brick B of W has the same (ℓ, a) -palette (seen as a cycle of W).

We would like to stress that, according to our definition, we do not consider a homogeneous subwall of a flat wall as a flat subwall itself. This permits us to avoid to define, in particular, the pegs and the rendition that would be associated with that flat subwall.

In the next lemma we prove that a sufficiently large flat wall contains a large enough homogeneous subwall.

LEMMA 4.3. (\star) *There is a function $f_4 : \mathbb{N}^3 \rightarrow \mathbb{N}$ such that if $\ell, r, a \in \mathbb{N}$, where r is an odd number, G is a graph, and (A, W) is an $(a, f_4(\ell, r, a))$ -apex-wall pair of G , then W has an (ℓ, a) -homogeneous subwall of height r . Moreover, it holds that $f_4(\ell, r, a) = r^{2^{2^{\mathcal{O}((a+\ell) \cdot \log(a+\ell))}}}$.*

5 Finding an irrelevant vertex

In this section we show how to find an irrelevant vertex inside a sufficiently large (in terms of the collection \mathcal{F}) flat wall of a boundaried graph that is *not* affected by its boundary. We start by defining in Subsection 5.1 *railed annuli*, a structure introduced in [29], and reused later in [25, 26], which occurs as a subgraph inside a flat wall and that turns out to be quite handy to guarantee a “confinement” property of topological minor models (cf. Proposition 5.2). 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 5.1), which will be helpful in order to reroute the model. We then prove (cf. Theorem 5.1) that the model can be rerouted—possibly, by obtaining another topological minor model that can be contracted to a dissolution of the original model—so that, in particular, a large enough central

region of the railed annulus is avoided by the model, and such that a prescribed subset of degree-3 vertices of the original model cannot be affected by contractions in order to obtain the desired model. Once we have all these ingredients, we finally show in Subsection 5.3 how the desired irrelevant vertex can be found inside a large enough flat wall that is not affected by the boundary (cf. Theorem 5.2). In order to prove this result, the notion of homogeneous subwall is crucial. We also need, in particular, to define an appropriate “flat” representation of a topological minor model, called its *leveling*, which we equip with colors encoding the “critical” zones in which we will have to be very careful to do the rerouting.

5.1 A lemma for model taming Before stating Proposition 5.2, we need some definitions.

Railed annuli. Let G be a partially Δ -embedded graph and let $\mathcal{C} = [C_1, \dots, C_r]$, $r \geq 2$, 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 $\Delta \supseteq D_1 \supseteq \dots \supseteq D_r$. From now on, each Δ -nested sequence $\mathcal{C} = [C_1, \dots, C_r]$ will be accompanied with the sequence $[D_1, \dots, D_r]$ of the corresponding open disks as well as the sequence $[\overline{D}_1, \dots, \overline{D}_r]$ of their closures. Given $x, y \in [r]$ with $x \leq y$, we call the set $\overline{D}_x \setminus D_y$ (x, y) -annulus of \mathcal{C} and we denote it by $\text{ann}(\mathcal{C}, x, y)$. Finally we say that $\text{ann}(\mathcal{C}, 1, r)$ is the *annulus* of \mathcal{C} and we denote it by $\text{ann}(\mathcal{C})$.

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 $\mathcal{C} = [C_1, \dots, C_r]$ is a Δ -nested collection of cycles of G and $\mathcal{P} = [P_1, \dots, P_q]$ is a collection of pairwise vertex-disjoint paths in G , called *rails*, such that

- for every $j \in [q]$, $P_j \subseteq \text{ann}(\mathcal{C})$, 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 4 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. The proof is easy and can be found, for instance, in [29].

PROPOSITION 5.1. *For every odd $x \in \mathbb{N}$, if W is a $(\lceil \frac{9}{2} \cdot x \rceil \times \lceil \frac{9}{2} \cdot x \rceil)$ -wall, then there is a collection \mathcal{P} of x paths in W such that if \mathcal{C} is the collection of the first x layers of W , then $(\mathcal{C}, \mathcal{P})$ is an (x, x) -railed annulus of W where the first cycle of \mathcal{C} is the perimeter of W . Moreover, the open disk defined by the x -th cycle of \mathcal{C} contains the central vertices of W .*




Figure 4: An example of a $(5, 8)$ -railed annulus and its inner disk D_5 .

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 $\mathcal{A}_{i,i'} = ([C_i, \dots, C_{i'}], \mathcal{P} \cap \text{ann}(\mathcal{C}, i, i'))$.

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

$$M \cap \text{ann}(\mathcal{C}, t + 1 - t', t + 1 + t') \subseteq \bigcup_{i \in I} P_i.$$

The following proposition has been recently proved by Golovach et al. [26, Theorem 3].

PROPOSITION 5.2. *There exist two functions $f_5, f_6 : \mathbb{Z}_{\geq 0} \rightarrow \mathbb{Z}_{\geq 0}$ such that if*

- s is a positive odd integer,
- H is a graph on g edges,
- G is a Δ -partially-embedded graph,
- $\mathcal{A} = (\mathcal{C}, \mathcal{P})$ is an (r, q) -railed annulus of G , where $r \geq f_6(g) + 2 + s$ and $q \geq 5/2 \cdot f_5(g)$,

- (M, T) is a topological minor model of H in G such that $T \cap \text{ann}(\mathcal{A}) = \emptyset$, and
- $I \subseteq [q]$ where $|I| > f_5(g)$,

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

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

Moreover $f_6(g) = O((f_5(g))^2)$.

5.2 Model rerouting in disk-embedded graphs

Using classic results on how to optimally draw planar graphs of maximum degree three into grids (see e.g., [30]) one may easily derive the following.

PROPOSITION 5.3. *There is a function $f_7 : \mathbb{N} \rightarrow \mathbb{N}$ such that for every ℓ -vertex planar graph H with maximum degree three there is a tm-pair (M, T) of the $(f_7(\ell) \times f_7(\ell))$ -grid, denoted by Γ , that is a topological minor model of H in Γ . Moreover, it holds that $f_7(\ell) = \mathcal{O}(\ell)$.*

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 $i \in [\frac{r}{2}]$, 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 vertices x and y of 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 crosses at most i edges of G . Given two subgraphs of W , 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 $c \geq 0$ 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 5.3 at hand, we can prove the following useful lemma.

LEMMA 5.1. (\star) *There is a (polynomial) function $f_8 : \mathbb{N}^3 \rightarrow \mathbb{N}$ 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 $Z := \{z_1, \dots, z_{r'}\} \subseteq V(H)$ 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 ,*
- *$\text{bd}(D) \cap H = Z$, and*
- *$(z_1, \dots, z_{r'})$ 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 a $(f_8(c, r, \ell), f_8(c, r, \ell))$ -railed annulus of G , where $\mathcal{C} = [C_1, \dots, C_{f_8(c, r, \ell)}]$, w_i be the endpoint of P_i that is contained in C_1 , for $i \in [r]$, and $I := \{i_1, \dots, i_{r'}\} \subseteq [r]$. Then the union-graph $\tilde{\Gamma} := G(\mathcal{A})$ of \mathcal{A} contains a tm -pair (M, T) that is a topological minor model of H in $\tilde{\Gamma}$ 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 $\tilde{\Gamma}$, and*
- *none of the vertices of $T \setminus \{w_{i_1}, \dots, w_{i_{r'}}\}$ is within face-distance less than c from some vertex in C_1 or in C_r .*

Moreover, it holds that $f_8(c, r, \ell) = \mathcal{O}(cr(\ell + r))$.

Let G be a partially Δ -embedded graph and let $\mathcal{C} = [C_1, \dots, C_r]$ be a Δ -nested sequence of cycles of G and let $[D_1, \dots, D_r]$ (resp. $[\overline{D}_1, \dots, \overline{D}_r]$) 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 $(M, T) \pitchfork \overline{D}_p$, 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 $\phi : V(G) \rightarrow 2^{V(H)}$ is a Q -respecting contraction-mapping of H to G if

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

The critical point in the above definitions is that vertices in Q are not “uncontracted” when transforming H to G . Given a non-negative integer x , we denote by $\text{odd}(x)$ the minimum odd number that is not smaller than x .

Intrusion of a topological minor model. Let G be graph, let $S \subseteq V(G)$, and let (M, T) be a tp -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 . Notice that S can intersect many times a subdivision path of (M, T) , however the value of the S -intrusion counts each such path only once.

Using Proposition 5.1, Proposition 5.2, and Lemma 5.1 we prove the following.

THEOREM 5.1. *There are functions $f_9 : \mathbb{N}^2 \rightarrow \mathbb{N}$ and $f_{10} : \mathbb{N} \rightarrow \mathbb{N}$ such that the following holds. Let $c, \ell \in \mathbb{N}$ and let G be a partially Δ -embedded graph, whose compass contains a $\lceil \frac{9}{2} \cdot f_9(c, \ell) \rceil$ -wall W with $\text{bd}(\Delta)$ as perimeter. Let also $C_1, \dots, C_{f_9(c, \ell)}$ be the first $f_9(c, \ell)$ -layers of W and $D_1, \dots, D_{f_9(c, \ell)}$ 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 $b \in [f_9(c, \ell)]$ such that*

1. $\hat{M} \setminus D_b$ is a subgraph of $M \setminus D_b$,
2. $\text{ann}(C_{b, b+f_{10}(c, \ell)-1}) \cap \hat{T} = \emptyset$,
3. $(\hat{M}, \hat{T}) \pitchfork \overline{D}_{b+f_{10}(c, \ell)}$ is a tm -pair of W that is safely c -dispersed in W and none of the vertices of $\hat{T} \cap \overline{D}_{b+f_{10}(c, \ell)}$ is within face-distance less than c from some vertex of $C_{b+f_{10}(c, \ell)} \cup C_{f_9(c, \ell)}$ in W .
4. $\hat{M} \cap \text{int}(D_{f_9(c, \ell)}) = \emptyset$,
5. there is a Q -respecting contraction-mapping of $\text{diss}(M, T)$ to $\text{diss}(\hat{M}, \hat{T})$, and
6. \hat{M} does not intersect the two centers of W .

Moreover, it holds that $f_9(c, \ell) = \mathcal{O}(c \cdot (f_5(\ell^2))^2)$.

See Figure 5 for an illustration of the conditions guaranteed by Theorem 5.1.

Proof. Let $g = \binom{\ell}{2}$, $r = f_5(g) + 1$, $s = \text{odd}(f_8(c, r, 3\ell + r))$, $x = \text{odd}(\max\{f_6(\ell) + 2 + s, \lceil 5/2 \cdot f_5(h) \rceil\})$, $l = (\ell + 1) \cdot x$, and $y = \lceil \frac{9}{2} l \rceil$. We will prove the theorem for $f_9(\ell) = l$ and $f_{10}(\ell) = \frac{x-s}{2}$. Let G be a partially Δ -embedded graph, whose compass contains a y -wall W with $\text{bd}(\Delta)$ as perimeter. Let also $\mathcal{C} = [C_1, \dots, C_l]$ be the first l layers of W and let $[D_1, \dots, D_r]$ (resp. $[\overline{D}_1, \dots, \overline{D}_r]$) be the sequences of the corresponding open (resp. closed) disks of Δ bounded by the cycles in \mathcal{C} . From Proposition 5.1 there is a collection $\mathcal{P} = \{P_1, \dots, P_x\}$ of paths in W such that $\mathcal{A} = (\mathcal{C}, \mathcal{P})$ is an




Figure 5: A visualization of how a tm-pair (M, T) is rearranged to a new tm-pair (\hat{M}, \hat{T}) as in Theorem 5.1. 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 $f_9(c, \ell)$ 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 \hat{T} that are vertices of W . The “green clouds” are the non-singleton images of the Q -respecting contraction-mapping of $\text{diss}(M, T)$ to $\text{diss}(\hat{M}, \hat{T})$.

(l, l) -railed annulus of W where the outer cycle of C is the perimeter of W and such that the two central vertices of W belong to the interior of \bar{D}_l .

Let \tilde{M} the union of all subdivision paths of (M, T) that intersect $\Delta \cap V(G)$ and let \tilde{T} be the endpoints of these paths. Moreover, we denote $\tilde{H} = \text{diss}(\tilde{M}, \tilde{T})$ and observe that \tilde{H} is a subgraph of H . Intuitively, \tilde{H} is the subgraph of H whose topological minor model (\tilde{M}, \tilde{T}) is the part of (M, T) intersects the 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 (\tilde{M}, \tilde{T}) in G . This in turn implies that $|\tilde{T} \cap \Delta| \leq \ell$ and that $|E(\tilde{H})| \leq \ell$.

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

We set $\mathcal{A}' = (\mathcal{C}_{b, b+x-1}, \mathcal{P} \cap \mathbb{A})$. By applying Propo-

sition 5.2 on s, H, g , the Δ -boundaried graph G , the (x, x) -railed annulus \mathcal{A}' , the tm-pair (\tilde{M}, \tilde{T}) , and the set $I = [r]$, we have that G contains a topological minor model (\tilde{M}, \tilde{T}) of \tilde{H} in G such that \tilde{M} is (s, I) -confined in \mathcal{A}' and $\tilde{M} \setminus \text{ann}(\mathcal{A}') \subseteq \tilde{M} \setminus \text{ann}(\mathcal{A}')$. 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 $\tilde{M} \setminus \text{ann}(\mathcal{A}') \subseteq \tilde{M} \setminus \text{ann}(\mathcal{A}')$.

Let $p = b + \frac{x-s}{2}$ and $q = b + \frac{x+s}{2} - 1$ and notice that $q \leq l$. We set $\mathbb{A}' := \text{ann}(\mathcal{C}_{p,q})$ and we define $\mathcal{A}'' := (\mathcal{C}_{p,q}, \mathcal{P}')$ where $\mathcal{P}' = \mathcal{P} \cap \mathbb{A}'$. Let $\mathcal{P}' = \{P'_1, \dots, P'_{r'}\}$. Observe that, from the second property of Proposition 5.2, the connected components of $\tilde{M} \cap \mathbb{A}'$ are some of the paths in \mathcal{P}' . This means that there is a subset of indices $\{i_1, \dots, i_{r'}\} \subseteq I$ such that $\tilde{M} \cap \mathbb{A}' = P'_{i_1} \cup \dots \cup P'_{i_{r'}}$. Let $Z = \{z_{i_1}, \dots, z_{i_{r'}}\}$ be the set of endpoints of the paths $P'_{i_1}, \dots, P'_{i_{r'}}$ that are contained in \mathcal{C}_p .

Let $\tilde{M}^{\text{in}} = \tilde{M} \cap \bar{D}_p$, $\tilde{M}^{\text{out}} = (G \setminus D_p) \setminus E(\mathcal{C}_p)$, 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 in T^{in} (resp. T^{out}) except from those in Z . Also $(\tilde{M}^{\text{in}}, T^{\text{in}} \cup Z)$ (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 on at most $3\ell + r$ edges where $\text{bd}(D) \cap H = Z$ and $(z_{i_1}, \dots, z_{i_{r'}})$ is the ordering of the vertices of Z as they appear in \mathcal{C}_p . Observe now that \tilde{H}^{in} can be seen as the contraction of another D -embedded graph \hat{H}^{in} on at most $3\ell + r$ vertices that has maximum degree at most three. Moreover, we can assume that the vertices of \hat{H}^{in} that have degree at most three are also vertices of \tilde{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 , $(z_{i_1}, \dots, z_{i_{r'}})$ is the ordering of the vertices of Z as they appear in $\text{bd}(D)$.

Keep in mind that $\hat{H}^+ = \hat{H}^{\text{out}} \cup \hat{H}^{\text{in}}$ is a minor of $\hat{H}^+ := \tilde{H}^{\text{out}} \cup \hat{H}^{\text{in}}$ and that if we dissolve in \hat{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 5.1 for c, r, r', ℓ , the D -embedded graph \hat{H}^{in} , the set Z , and the (s, s) -railed annulus \mathcal{A}'' of the \bar{D}_p -disk embedded graph $G \cap \bar{D}_p$ 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 $W \cap \text{ann}(\mathcal{C}_{p,q})$. From

the second property of Lemma 5.1, $(\hat{M}^{\text{in}}, \hat{T}^{\text{in}})$ is safely c -dispersed in $W \cap \text{ann}(\mathcal{C}_{p,q})$. From the third property of Lemma 5.1, it follows that none of the vertices of $\hat{T}^{\text{in}} \setminus \{w_{i_1}, \dots, w_{i_r}\}$ is within face-distance less than c from some vertex in $C_p \cup C_q$ in $W \cap \text{ann}(\mathcal{C}_{p,q})$.

We now consider the graph $\hat{M} = \hat{M}^{\text{in}} \cup \tilde{M}^{\text{out}}$. Properties 3 and 6 follow by the conclusions of the previous paragraph. Moreover, \hat{M} does not intersect \overline{D}_{q+1} and, as $q \leq l$, it neither intersects \overline{D}_{l+1} and 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 $\tilde{M} \setminus \overline{D}_b = \tilde{M} \setminus \overline{D}_b$, yield Property 1.

Observe that $(\hat{M}, \hat{T}^{\text{in}} \cup T^{\text{out}} \cup Z)$ is a topological minor model of \hat{H}^+ , which in turn implies that $(\hat{M}, \hat{T}^{\text{in}} \cup T^{\text{out}})$ is a topological minor model of \hat{H} . We now set $\hat{T} = \hat{T}^{\text{in}} \cup T^{\text{out}}$. As there is a Q -respecting contraction-mapping of \tilde{H}^{out} to \hat{H}^{out} , we also have that there is a Q -respecting contraction-mapping of $H = \text{diss}(M, T)$ to $\hat{H} = \text{diss}(\hat{M}, \hat{T})$ and Property 5 holds. As $\hat{T}^{\text{in}} \subseteq \text{int}(\text{ann}(\mathcal{A}'')) \subseteq D_p = D_{b+f_{10}(\ell)-1}$ and $T^{\text{out}} \subseteq G \setminus \overline{D}_b$, we deduce that $\hat{T} \in G \setminus \text{ann}(\mathcal{C}_{b, b+f_{10}(\ell)-1})$, which yields Property 2. \square

5.3 Rerouting minors of small intrusion Let W be an r -wall and $c \geq 0$. We call a cycle C of W c -internal if it is within face-distance at least c from the perimeter of W . Given a 0-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 (we see edges as open sets). Notice that each vertex of C is either an internal or an external peg. Given two subgraphs H_1 and H_2 of a graph H we define the *distance* in H between H_1 and H_2 as the minimum distance between a vertex in H_1 and a vertex in H_2 .

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 Figure 6. We call this subgraph X of W the *brick-neighborhood* of B . The perimeter P of a brick-neighborhood is defined in the obvious way. The next lemma is based on Observation 2.




Figure 6: The base graph for the definition of a brick-neighborhood – the external pegs of the perimeter of X are the black round vertices.

LEMMA 5.2. (\star) *Let W be a flat wall of a graph G and let W' be its representation wall in the leveling \tilde{W} of W . For every 1-internal brick B of W' and every flap $F \in \text{flaps}(B)$, the brick-neighborhood X of B contains $|\partial F|$ internally vertex-disjoint paths of W' from F to the external pegs of the perimeter of X .*

Vertices irrelevant to minors. 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 (T, M) of G such that (T, M) is a topological minor model of G and $\text{diss}(T, M) \in \text{ext}(H)$.

Let $\mathbf{Z} = (Z, B, \rho)$ be a t -boundaried graph and $\ell \in \mathbb{N}$. We say that a vertex $v \in V(Z) \setminus B$ is an ℓ -irrelevant vertex of \mathbf{Z} if for every boundaried graph $\mathbf{C} = (C, B, \rho)$ that is compatible with \mathbf{Z} , every minor of $\mathbf{C} \oplus \mathbf{Z}$ with $(V(Z) \setminus B)$ -minor-intrusion in G at most ℓ , is also a minor of $\mathbf{C} \oplus (Z \setminus v, B, \rho)$. Informally, such an irrelevant vertex can be removed without affecting the occurrences of any minor of small minor-intrusion, where the intrusion is defined without taking into account the branching vertices in the boundary.

Using Theorem 5.1 we can finally prove the main result of this section.

THEOREM 5.2. (\star) *There is a function $f_{11} : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that, for every $a, \ell \in \mathbb{N}$ and every boundaried graph $\mathbf{Z} = (Z, B, \rho)$, if (A, W) is an $(a, f_{11}(a, \ell))$ -apex-wall pair of Z that is not affected by B , then \mathbf{Z} contains an ℓ -irrelevant vertex. Moreover, it holds that $f_{11}(a, \ell) = (f_5(\ell^2))^{2^{2^{(a+\ell^2) \cdot \log(a+\ell^2)}}}$.*

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 $2h$ -irrelevant vertex. To see this, recall that in (3.2) the equivalence is defined in terms of graphs H on at most h vertices and at most h edges, and that every topological minor-minimal in $\text{ext}(H)$ has at most $2h$ vertices and at most $2h$ edges. Thus, from Theorem 5.2, B should affect every $(a, f_{11}(a, 2h))$ -apex-wall pair of G , for every value of a . We conclude the following.

COROLLARY 5.1. *If $t, h, a \in \mathbb{N}$ and $\mathbf{G} = (G, B, \rho)$ is a boundaried graph in $\mathcal{R}_h^{(t)}$, then it holds that $\mathbf{P}_{a, f_{11}(a, 2h)}(\mathbf{G}) \leq t$.*

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 at most $2t$ 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. In the full version 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 [33, 45].

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. [7, 14]).

LEMMA 6.1. *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 $\mathbf{tw}(G) + 1$.*

Using Lemma 4.1, Lemma 4.2, Corollary 5.1, and Lemma 6.1 we prove the following result, whose proof is an adaptation to our setting of the one of [22, Lemma 3.6] (see also [21]). We stress that \mathbf{p} is not a bidimensional parameter in the way this is defined in [16, 18, 21, 22], therefore Lemma 6.2 cannot be derived by directly applying the results of [22].

LEMMA 6.2. (\star) *There is a function $f_{12} : \mathbb{N} \rightarrow \mathbb{N}$ such that if $t, h, q \in \mathbb{N}$ and $\mathbf{G} = (G, B, \rho)$ is a K_q -minor-free boundaried graph in $\mathcal{R}_h^{(t)}$, then G contains an $f_{12}(q, h)$ -treewidth modulator that contains B and has at most $2t$ vertices. Moreover, it holds that $f_{12}(q, h) = (f_5(h^2))^{2^{2^{\mathcal{O}((q+h^2) \cdot \log(q+h^2))}}}$.*

6.2 Finding a linear protrusion decomposition and reducing protrusions Equipped with Lemma 6.2, 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 $\mathbf{tw}(G[X]) \leq \beta - 1$. Given $\alpha, t \in \mathbb{N}$, an (α, β) -protrusion decomposition of G is a sequence $\mathcal{P} = \langle R_0, R_1, \dots, R_\ell \rangle$ of pairwise disjoint subsets of $V(G)$ such that

- $\bigcup_{i \in [\ell]} = V(G)$,
- $\max\{\ell, |R_0|\} \leq \alpha$,
- 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 $\mathbf{tw}(\mathbf{G})$, 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 ,
- $\mathbf{tw}(\mathbf{G}') \leq \beta - 1$,
- $\partial(V(G')) \subseteq B'$, 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 $\mathcal{P} = \langle R_0, R_1, \dots, R_\ell \rangle$ of pairwise disjoint subsets of $V(G)$ such that

- $\bigcup_{i \in [\ell]} = V(G)$,
- $\max\{\ell, |R_0|\} \leq \alpha$,
- $B \subseteq R_0$,
- for $i \in [\ell]$, $(G(N[R_i]), \partial(N[R_i]), \rho|_{\partial(N[R_i])})$ is a β -protrusion of \mathbf{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 *core* of \mathcal{P} .

The following theorem is a reformulation using our notation of one of the main results of Kim et al. [34], which is stronger than what we need in the sense that also applies to graphs excluding a *topological* minor.

THEOREM 6.1. *Let c, β, t be positive integers, let H be an q -vertex graph, and let G be an n -vertex H -topological-minor-free graph. If we are given a set $M \subseteq V(G)$ with $|M| \leq c \cdot t$ such that $\mathbf{tw}(G - M) \leq \beta$, 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 $\alpha_H \leq 40q^{2 \cdot 2^{5q \log q}}$.*

Having stated the above definitions, the following lemma is an easy consequence of Lemma 6.2 and Theorem 6.1.

LEMMA 6.3. *There is a function $f_{13} : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that if $t, h, q \in \mathbb{N}$ and $\mathbf{G} = (G, B, \rho)$ is a K_q -minor-free boundaried graph in $\mathcal{R}_h^{(t)}$, then \mathbf{G} admits a $(f_{13}(q, h) \cdot t, f_{13}(q, h))$ -protrusion decomposition. Moreover, it holds that $f_{13}(q, h) = (f_5(h^2))^{2^{2^{\mathcal{O}((q+h^2) \cdot \log(q+h^2))}}}$.*

Proof. By Lemma 6.2, G contains an $f_{12}(q, h)$ -treewidth modulator M that contains B and has at most $2t$ vertices. We can now apply Theorem 6.1 to G and M with $H = K_q$, $c = 2$, and $\beta = f_{12}(q, h)$, obtaining a $(f_{13}(q, h) \cdot t, f_{13}(q, h))$ -protrusion decomposition \mathcal{P} of G with M contained in the core of \mathcal{P} and $f_{13}(q, h) := 2 \cdot f_{12}(q, h) \cdot 40q^2 2^{5q \log q}$. Since $B \subseteq M$ and M contained in the core of \mathcal{P} , it can be easily seen that \mathcal{P} is also a $(f_{13}(q, h) \cdot t, f_{13}(q, h))$ -protrusion decomposition of \mathbf{G} . \square

Once we have the protrusion decomposition given by Lemma 6.3, 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. [9], 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. [34]. We could directly apply the results of Fomin et al. [21] to the protrusion decomposition of a representative given by Lemma 6.3, 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 [9, 34]) is that they do not provide explicit bounds on the hidden constants. In order to be able to do so, we apply the protrusion replacement used by Baste et al. [5], 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. [23] (later generalized in [24]) for obtaining kernels with explicit constants.

Given a function $\xi : \mathbb{N}^2 \rightarrow \mathbb{N}$ and a t -boundaried graph \mathbf{G} , we say that \mathbf{G} is ξ -protrusion-bounded if, for every $t' \in \mathbb{N}$, all β -protrusions of \mathbf{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. [5]. Namely, it is a consequence of the proof⁶ of [4, Lemma 15].

⁶In the statement of [4, Lemma 15] it is required that the

LEMMA 6.4. *There is a function $f_{14} : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that if $t, h, q \in \mathbb{N}$ and $\mathbf{G} = (G, B, \rho)$ is a K_q -minor-free boundaried graph in $\mathcal{R}_h^{(t)}$, then \mathbf{G} is $f_{14}(q, h)$ -protrusion-bounded. Moreover, $f_{14}(q, h) = 2^{2^{2^{\mathcal{O}(f_{13}(q, h)) \cdot \log f_{13}(q, h)}}}$.*

Using Lemma 6.3 and Lemma 6.4, we can easily prove Theorem 6.2, that is the main result on which the algorithm of Theorem 2.1 is based (cf. §2). In particular, it implies (2.1).

THEOREM 6.2. *There is a function $f_{15} : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that, for every $t, h \in \mathbb{N}_{\geq 1}$, if $\mathbf{G} = (G, B, \rho)$ is a K_q -minor-free boundaried graph in $\mathcal{R}_h^{(t)}$, then $|V(G)| \leq f_{15}(q, h) \cdot t$. Moreover, it holds that $f_{15}(q, h) \leq f_{13}(q, h) \cdot (f_{14}(q, h) + 1)$.*

Proof. By Lemma 6.3, \mathbf{G} admits a $(f_{13}(q, h) \cdot t, f_{13}(q, h))$ -protrusion decomposition \mathcal{P} . By Lemma 6.4, each of the protrusions of \mathcal{P} has at most $f_{14}(q, h)$ vertices. Therefore,

$$|V(G)| \leq f_{13}(q, h) \cdot t + f_{13}(q, h) \cdot f_{14}(q, h) \cdot t,$$

and the theorem follows with $f_{15}(q, h) := f_{13}(q, h) \cdot (f_{14}(q, h) + 1)$. \square

Let $h := \max_{F \in \mathcal{F}} \{\max_{H \in \text{ext}(F)} |V(H)|\}$. The following corollary is an immediate consequence of Theorem 6.2, by using the fact that all t -representatives in $\mathcal{R}_h^{(t)}$, except one, are K_h -minor-free, hence they have $\mathcal{O}(f_{15}(h, h) \cdot h \sqrt{\log h}) \cdot t$ edges; see for instance [40].

COROLLARY 6.1. *There is a function $f_{16} : \mathbb{N} \rightarrow \mathbb{N}$ such that for every $t \in \mathbb{N}_{\geq 1}$, $|\mathcal{R}_h^{(t)}| \leq 2^{f_{16}(h) \cdot t \cdot \log t}$. In particular, the relation \equiv_h partitions $\mathcal{B}^{(t)}$ into $2^{f_{16}(h) \cdot t \cdot \log t}$ equivalence classes. Moreover, it holds that $f_{16}(h) = \mathcal{O}(f_{13}(h, h) \cdot f_{14}(h, h)) = \mathcal{O}(f_{14}(h, h))$.*

The dynamic programming algorithm. Having proved Corollary 6.1, we can just reuse the dynamic programming algorithm given in [5] to compute the parameter $\mathbf{m}_{\mathcal{F}}(G)$ in the claimed running time. For the sake of completeness, let us comment some details of this algorithm, whose details can be found in [4, proof of Theorem 3]. 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 6.2, by generating all t -boundaried graphs on at most $f_{15}(h, h) \cdot t$

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 6.3, we can assume that this also holds in our setting.

vertices and $\mathcal{O}(f_{15}(h, h) \cdot h\sqrt{\log h}) \cdot t$ edges, partitioning them into equivalence classes according to \equiv_h , and picking an element of minimum size in each of them; see [4, proof of Lemma 14] for more details. To simplify the description of the dynamic programming update operations, the algorithm in [5] is written in terms of *branchwidth* instead of *treewidth*. Without defining branchwidth here, it is enough to say that it is parametrically 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 [43]. Also, the algorithm in [5] is written in terms of *topological minors*, that is, it computes a 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*; we denote $|S| =: \mathbf{tm}_{\mathcal{F}}(G)$. It is easy to see that computing this parameter suffices for computing $\mathbf{m}_{\mathcal{F}}(G)$, since, as observed in [4, Lemma 4], for every proper collection \mathcal{F} and every graph G , it holds that $\mathbf{m}_{\mathcal{F}}(G) = \mathbf{tm}_{\mathcal{F}'}(G)$, where \mathcal{F}' is the family containing every topological minor minimal graph among those that contain some graph in \mathcal{F} as a minor; note that \mathcal{F}' has size bounded by a small function of \mathcal{F} (see Observation 1).

The algorithm then computes, in a typical bottom-up manner, at every bag separator B of 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. [4, proof of Theorem 3]. The overall running time is bounded by $|\mathcal{R}_h^{(t)}|^2 \cdot |E(G)|$, and taking into account that $|E(G)| \leq \mathbf{tw}(G) \cdot |V(G)|$, from Corollary 6.1 we obtain the following theorem, which is a more precise reformulation of Theorem 2.1.

THEOREM 6.3. *Let $t, h \in \mathbb{N}$, \mathcal{F} be a proper connected 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 $2^{(2 \cdot f_{16}(h) \cdot t + 1) \cdot \log t} \cdot n$ steps.*

7 Further research

Our main algorithmic result is an algorithm solving \mathcal{F} -M-DELETION in time $\mathcal{O}^*(2^{\mathcal{O}(\mathbf{tw} \cdot \log \mathbf{tw})})$ for every connected collection \mathcal{F} . One may wonder why the connectivity of \mathcal{F} is necessary. In fact, in the whole algorithm (see Figure 2) the connectivity of the graphs in \mathcal{F} is only used at the very end, when we apply the dynamic programming algorithm of [5] based on representatives. This algorithm uses the connectivity of \mathcal{F} in the “base case”, namely to guarantee that the representative of a graph G without boundary is the empty boundaried

graph if and only if G does not contain any of the graphs in \mathcal{F} as a minor (see [4, Lemma 7]). We think that this is a technical hurdle, rather than an intrinsic one, and we believe that the condition on the connectivity of \mathcal{F} can be dropped, that is, that there exists an algorithm to solve \mathcal{F} -M-DELETION in time $\mathcal{O}^*(2^{\mathcal{O}(\mathbf{tw} \cdot \log \mathbf{tw})})$ for every collection \mathcal{F} . As an evidence towards this, note that the minor obstructions for being embeddable on a surface of Euler genus at most g contain disconnected graphs if $g \geq 2$ (for instance, the disjoint union of two K_5 's is an obstruction for being embeddable on the torus [39]), and that Kociumaka and Pilipczuk [35] presented an algorithm running in time $\mathcal{O}^*(2^{\mathcal{O}((\mathbf{tw}+g) \cdot \log(\mathbf{tw}+g))})$ for deleting a minimum number of vertices to obtain a graph embeddable on a surface of Euler genus at most g .

We also presented a framework to obtain lower bounds for ruling out algorithms in time $\mathcal{O}^*(2^{\mathcal{O}(\mathbf{tw} \cdot \log \mathbf{tw})})$ under the ETH. In particular, when $\mathcal{F} = \{H\}$ and H is connected, it settles completely the asymptotic complexity of $\{H\}$ -M-DELETION (Theorem 1.1). However, we do not have a complete classification when $|\mathcal{F}| \geq 2$, 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 $\mathcal{O}^*(2^{\Theta(\mathbf{tw})})$ (resp. $\mathcal{O}^*(2^{\Theta(\mathbf{tw} \cdot \log \mathbf{tw})})$). Suppose that $\mathcal{F} = \{H_1, H_2\}$ with both H_1 and H_2 being connected. Using the recent results of Baste [3], 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. [10] presented an algorithm running in time $\mathcal{O}^*(2^{\mathcal{O}(\mathbf{tw})})$ for PSEUDOFORREST 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 PSEUDOFORREST DELETION is equivalent to $\{\mathbf{diamond}, \mathbf{butterfly}\}$ -M-DELETION. While both the *diamond* and the *butterfly* are hard graphs, $\{\mathbf{diamond}, \mathbf{butterfly}\}$ 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 a 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 the lower bounds that we presented in this article also hold for \mathcal{F} -TM-DELETION (with the exception of $K_{1,i}$ for $i \geq 4$; see [4]), the algorithm in time $\mathcal{O}^*(2^{\mathcal{O}(\mathbf{tw} \cdot \log \mathbf{tw})})$ for every connected collection \mathcal{F} does not work for topological minors. In this direction, the algorithm in time $\mathcal{O}^*(2^{\mathcal{O}(\mathbf{tw} \cdot \log \mathbf{tw})})$ in [5] for \mathcal{F} -M-

DELETION (when \mathcal{F} is connected and contains a *planar* graph) 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 $\mathcal{O}^*(2^{\mathcal{O}(\text{tw} \cdot \log \text{tw})})$ for every connected 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 5.1), 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 5.2, as this is done in [26] for planar graphs.

References

- [1] I. ADLER, S. G. KOLLIPOULOS, P. K. KRAUSE, D. LOKSHANOV, S. SAURABH, AND D. M. THILIKOS, *Irrelevant vertices for the planar Disjoint Paths Problem*, Journal of Combinatorial Theory, Series B, 122 (2017), pp. 815–843.
- [2] Z. BAI, J. TU, AND Y. SHI, *An improved algorithm for the vertex cover P_3 problem on graphs of bounded treewidth*, CoRR, abs/1603.09448 (2016).
- [3] J. BASTE, *Composing dynamic programming tree-decomposition-based algorithms*, CoRR, abs/1904.12500 (2019).
- [4] J. BASTE, I. SAU, AND D. M. THILIKOS, *Hitting minors on bounded treewidth graphs*, CoRR, abs/1704.07284 (2017).
- [5] ———, *Optimal Algorithms for Hitting (Topological) Minors on Graphs of Bounded Treewidth*, in Proc. of the 12th International Symposium on Parameterized and Exact Computation (IPEC), vol. 89 of LIPIcs, 2017, pp. 4:1–4:12.
- [6] ———, *A Complexity Dichotomy for Hitting Small Planar Minors Parameterized by Treewidth*, in Proc. of the 13th International Symposium on Parameterized and Exact Computation (IPEC), vol. 115 of LIPIcs, 2018, pp. 2:1–2:13.
- [7] H. L. BODLAENDER, *A Partial k -Arboretum of Graphs with Bounded Treewidth*, Theoretical Computer Science, 209 (1998), pp. 1–45.
- [8] H. L. BODLAENDER, M. CYGAN, S. KRATSCH, AND J. NEDERLOF, *Deterministic single exponential time algorithms for connectivity problems parameterized by treewidth*, Information and Computation, 243 (2015), pp. 86–111.
- [9] H. L. BODLAENDER, F. V. FOMIN, D. LOKSHANOV, E. PENNINKX, S. SAURABH, AND D. M. THILIKOS, *(Meta) Kernelization*, Journal of the ACM, 63 (2016), pp. 44:1–44:69.
- [10] H. L. BODLAENDER, H. ONO, AND Y. OTACHI, *A faster parameterized algorithm for pseudoforest deletion*, Discrete Applied Mathematics, 236 (2018), pp. 42–56.
- [11] É. BONNET, N. BRETTELL, O. KWON, AND D. MARX, *Generalized Feedback Vertex Set Problems on Bounded-Treewidth Graphs: Chordality Is the Key to Single-Exponential Parameterized Algorithms*, in Proc. of the 12th International Symposium on Parameterized and Exact Computation (IPEC), vol. 89 of LIPIcs, 2017, pp. 7:1–7:13.
- [12] J. CHUZHUY, *Improved bounds for the flat wall theorem*, in Proc. of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2015, pp. 256–275.
- [13] B. COURCELLE, *The Monadic Second-Order Logic of Graphs. I. Recognizable Sets of Finite Graphs*, Information and Computation, 85 (1990), pp. 12–75.
- [14] M. CYGAN, F. V. FOMIN, L. KOWALIK, D. LOKSHANOV, D. MARX, M. PILIPCZUK, M. PILIPCZUK, AND S. SAURABH, *Parameterized Algorithms*, Springer, 2015.
- [15] M. CYGAN, J. NEDERLOF, M. PILIPCZUK, M. PILIPCZUK, J. M. M. VAN ROOIJ, AND J. O. WOJTASZCZYK, *Solving Connectivity Problems Parameterized by Treewidth in Single Exponential Time*, in Proc. of the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2011, pp. 150–159.
- [16] E. D. DEMAINE, F. V. FOMIN, M. HAJIAGHAYI, AND D. M. THILIKOS, *Subexponential parameterized algorithms on graphs of bounded genus and H -minor-free graphs*, Journal of the ACM, 52 (2005), pp. 866–893.
- [17] R. DIESTEL, *Graph Theory*, vol. 173, Springer-Verlag, 4th edition, 2010.
- [18] F. V. FOMIN, E. D. DEMAINE, M. T. HAJIAGHAYI, AND D. M. THILIKOS, *Bidimensionality*, in Encyclopedia of Algorithms, publisher-springer, 2016, pp. 203–207.
- [19] F. V. FOMIN AND D. KRATSCH, *Exact Exponential Algorithms*, Texts in Theoretical Computer Science. An EATCS Series, Springer, 2010.
- [20] F. V. FOMIN, D. LOKSHANOV, F. PANOLAN, AND S. SAURABH, *Efficient computation of representative families with applications in parameterized and exact algorithms*, Journal of the ACM, 63 (2016), pp. 29:1–29:60.
- [21] F. V. FOMIN, D. LOKSHANOV, S. SAURABH, AND D. M. THILIKOS, *Bidimensionality and ker-*

- nels, in Proc. of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2010, pp. 503–510.
- [22] —, *Bidimensionality and kernels*, CoRR, abs/1606.05689 (2016).
- [23] V. GARNERO, C. PAUL, I. SAU, AND D. M. THILIKOS, *Explicit Linear Kernels via Dynamic Programming*, SIAM Journal on Discrete Mathematics, 29 (2015), pp. 1864–1894.
- [24] —, *Explicit linear kernels for packing problems*, Algorithmica, 81 (2019), pp. 1615–1656.
- [25] P. A. GOLOVACH, M. KAMINSKI, S. MANIATIS, AND D. M. THILIKOS, *The parameterized complexity of graph cyclability*, SIAM Journal on Discrete Mathematics, 31 (2017), pp. 511–541.
- [26] P. A. GOLOVACH, G. STAMOULIS, AND D. M. THILIKOS, *Hitting Topological Minor Models in Planar Graphs is Fixed Parameter Tractable*, CoRR, abs/1907.02919 (2019). To appear in the Proc. of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2019).
- [27] R. IMPAGLIAZZO, R. PATURI, AND F. ZANE, *Which problems have strongly exponential complexity?*, Journal of Computer and System Sciences, 63 (2001), pp. 512–530.
- [28] B. M. P. JANSEN, D. LOKSHTANOV, AND S. SAURABH, *A near-optimal planarization algorithm*, in Proc. of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2014, pp. 1802–1811.
- [29] M. KAMINSKI AND D. M. THILIKOS, *Contraction checking in graphs on surfaces*, in Proc. of the 29th International Symposium on Theoretical Aspects of Computer Science, (STACS), vol. 14 of LIPIcs, 2012, pp. 182–193.
- [30] G. KANT, *Algorithms for Drawing Planar Graphs*, PhD thesis, Department of Information and Computing Sciences - Utrecht University, 1993.
- [31] K. KAWARABAYASHI AND Y. KOBAYASHI, *Linear min-max relation between the treewidth of H -minor-free graphs and its largest grid*, in Proc. of the 29th International Symposium on Theoretical Aspects of Computer Science (STACS), vol. 14 of LIPIcs, 2012, pp. 278–289.
- [32] K. KAWARABAYASHI, R. THOMAS, AND P. WOLLAN, *A new proof of the flat wall theorem*, Journal of Combinatorial Theory, Series B, 129 (2018), pp. 204–238.
- [33] K. KAWARABAYASHI AND P. 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), 2010, pp. 687–694.
- [34] E. J. KIM, A. LANGER, C. PAUL, F. REIDL, P. ROSSMANITH, I. SAU, AND S. SIKDAR, *Linear kernels and single-exponential algorithms via protrusion decompositions*, ACM Transactions on Algorithms, 12 (2016), pp. 21:1–21:41.
- [35] T. KOCIUMAKA AND M. PILIPCZUK, *Deleting vertices to graphs of bounded genus*, Algorithmica, 81 (2019), pp. 3655–3691.
- [36] S. KRATSCHE AND M. WAHLSTRÖM, *Representative Sets and Irrelevant Vertices: New Tools for Kernelization*, in Proc. of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2012, pp. 450–459.
- [37] J. M. LEWIS AND M. YANNAKAKIS, *The node-deletion problem for hereditary properties is NP-complete*, Journal of Computer and System Sciences, 20 (1980), pp. 219–230.
- [38] D. LOKSHTANOV, D. MARX, AND S. SAURABH, *Slightly superexponential parameterized problems*, in Proc. of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2011, pp. 760–776.
- [39] B. MOHAR AND C. THOMASSEN, *Graphs on surfaces*, Johns Hopkins University Press, 2001.
- [40] J. S. MYERS AND A. THOMASON, *The extremal function for noncomplete minors*, Combinatorica, 25 (2005), pp. 725–753.
- [41] M. PILIPCZUK, *A tight lower bound for Vertex Planarization on graphs of bounded treewidth*, Discrete Applied Mathematics, 231 (2017), pp. 211–216.
- [42] B. A. REED, *Rooted routing in the plane*, Discrete Applied Mathematics, 57 (1995), pp. 213–227.
- [43] N. ROBERTSON AND P. D. SEYMOUR, *Graph Minors. X. Obstructions to Tree decomposition*, Journal of Combinatorial Theory, Series B, 52 (1991), pp. 153–190.
- [44] N. ROBERTSON AND P. D. SEYMOUR, *Graph Minors. XIII. The Disjoint Paths Problem*, Journal of Combinatorial Theory, Series B, 63 (1995), pp. 65–110.
- [45] —, *Graph Minors. XXII. Irrelevant vertices in linkage problems*, Journal of Combinatorial Theory, Series B, 102 (2012), pp. 530–563.
- [46] J. TU, L. WU, J. YUAN, AND L. CUI, *On the vertex cover P_3 problem parameterized by treewidth*, Journal of Combinatorial Optimization, 34 (2017), pp. 414–425.
- [47] D. P. WILLIAMSON AND D. B. SHMOYS, *The Design of Approximation Algorithms*, Cambridge University Press, 2011.