

Turing kernelization for finding long paths in graph classes excluding a topological minor

Citation for published version (APA):

Jansen, B. M. P., Pilipczuk, M., & Wrochna, M. (2017). Turing kernelization for finding long paths in graph classes excluding a topological minor. *arXiv*, 1-22. [1707.01797].

Document status and date:

Published: 01/01/2017

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Turing Kernelization for Finding Long Paths in Graph Classes Excluding a Topological Minor*

Bart M. P. Jansen[†], Marcin Pilipczuk[‡] and Marcin Wrochna[‡]

Abstract

The notion of Turing kernelization investigates whether a polynomial-time algorithm can solve an NP-hard problem, when it is aided by an oracle that can be queried for the answers to bounded-size subproblems. One of the main open problems in this direction is whether k -PATH admits a polynomial Turing kernel: can a polynomial-time algorithm determine whether an undirected graph has a simple path of length k , using an oracle that answers queries of size $k^{\mathcal{O}(1)}$?

We show this can be done when the input graph avoids a fixed graph H as a topological minor, thereby significantly generalizing an earlier result for bounded-degree and $K_{3,t}$ -minor-free graphs. Moreover, we show that k -PATH even admits a polynomial Turing kernel when the input graph is not H -topological-minor-free itself, but contains a known vertex modulator of size bounded polynomially in the parameter, whose deletion makes it so. To obtain our results, we build on the graph minors decomposition to show that any H -topological-minor-free graph that does not contain a k -path, has a separation that can safely be reduced after communication with the oracle.

1 Introduction

Suppose that Alice is a polynomial-time agent faced with an input to an NP-hard problem that she wishes to solve exactly. To facilitate her in this process, she can ask questions to an all-knowing oracle. These will be answered truthfully and instantly, but the oracle is memory-less and will not take previous questions into account when answering the next one. How large do these questions have to be, to allow Alice to find the answer to her problem? Clearly, the answer can be established by sending the entire input to the oracle, who determines the answer and sends it to Alice. Could there be a more clever strategy? Alice can attempt to isolate a small but meaningful question about the behavior of her input, such that after learning its answer, she can reduce to a smaller input without changing the outcome. Iterating this process solves her problem: when it has become sufficiently small, it can be posed to the oracle in its entirety.

Such problem-solving strategies can be rigorously analyzed using the notion of *Turing kernelization* that originated in parameterized algorithmics. The parameter makes it possible to express how the size of the questions that Alice asks, depends on properties of the input that she is given. (See Section 3 for a formal definition.)

Understanding the power of Turing kernelization is one of the main open research horizons in parameterized algorithmics. There is a handful of problems for which a nontrivial Turing

*This work was supported by the Netherlands Organization for Scientific Research (NWO) Veni grant 639.021.437 “Frontiers in Parameterized Preprocessing” and Gravitation grant 024.002.003 “Networks”. Marcin Pilipczuk is supported by the “Recent trends in kernelization: theory and experimental evaluation” project, carried out within the Homing programme of the Foundation for Polish Science co-financed by the European Union under the European Regional Development Fund. Marcin Wrochna is supported by the National Science Centre of Poland grant number 2013/11/D/ST6/03073 and by the Foundation for Polish Science (FNP) via the START stipend programme.

[†]Eindhoven University of Technology, The Netherlands

[‡]University of Warsaw, Poland

kernelization is known [1, 3, 4, 6, 14, 17, 19, 20, 26, 28]. On the other hand, there is a hierarchy of parameterized complexity classes which are conjectured not to admit polynomial Turing kernels [16]. Arguably, the main open problem (cf. [5, 4, 16]) in this direction is to determine whether the k -PATH problem (determine whether an undirected graph has a simple path of length k) has a polynomial Turing kernel. In earlier work [18], the first author showed that k -PATH indeed admits polynomial Turing kernels on several graph classes. In this work, we develop Turing kernels for k -PATH in a much more general setting.

Our results Our algorithmic contributions are twofold. First of all, we extend the Turing kernelization for k -PATH to much broader families of sparse graphs. Whereas the earlier work could only deal with $K_{3,t}$ -minor-free graphs, claw-free graphs, and bounded-degree graphs, we show that a Turing kernelization exists on H -minor-free graphs for all fixed graphs H . We even lift the kernelization to H -topological-minor-free graphs, thereby capturing a common generalization of the bounded-degree and $K_{3,t}$ -minor-free cases.

Theorem 1.1. *For every fixed graph H , the k -PATH problem, restricted to graphs excluding H as a topological minor, admits a polynomial Turing kernel. Furthermore, the kernel runs in time $k^{\mathcal{O}_H(1)}n^2m$ and invokes $k^{\mathcal{O}_H(1)} \cdot n$ calls to the oracle.*

Our second contribution is the following theorem. By a novel algorithmic approach, we obtain a Turing kernelization even when the input graph does not belong to the desired restricted graph class itself, but contains a small known vertex modulator whose deletion places the graph in such a graph class.

Theorem 1.2. *For every fixed graph H , the k -PATH problem, on instances consisting of a graph G , integer k , and a modulator $M \subseteq V(G)$ such that $G - M$ is H -topological-minor-free, admits a polynomial Turing kernel, when parameterized by k and $|M|$.*

Techniques To explain our approach, we briefly recall the idea behind the Turing kernelization for k -PATH on planar graphs. At the core lies a win/win: there is a polynomial-time algorithm that either (i) establishes that a planar graph G has a k -path (a simple path on k vertices), or (ii) finds a separation (A, B) in G with the following property: the size of A is polynomially bounded in k , but large enough that after marking a witness structure for each reasonable way in which a k -path might intersect A , some vertex remains unmarked. Using bounded-size oracle queries to mark the witness structures, this allows the problem to be simplified by removing an unmarked vertex from A without changing the answer.

Theorem 1.1 is established by lifting this win/win approach to H -(topological)-minor-free graphs. This requires an adaptation of the decomposition theorems of Robertson and Seymour [25] (for minors) and of Grohe and Marx [15] (for topological minors), to obtain the following. Every H -free graph that does not have a k -path, has a tree decomposition of constant adhesion and width $\text{poly}(k)$. A reducible separation can be found by inspecting this tree decomposition. To establish this result, we exploit known theorems stating that triconnected n -vertex graphs that exclude $K_{3,t}$ as a minor for some t [8], contain paths of length $\Omega(n^\varepsilon)$ for some $\varepsilon > 0$. Roughly speaking, this allows us to infer the existence of a k -path if there is a large embedded part in the nearly-embeddable graph corresponding to a bag of the graph minors decomposition, since graphs embeddable in a fixed surface are $K_{3,t}$ -minor-free for some t . We use lower bounds on the circumference of graphs of bounded degree [7, 27] to achieve a similar conclusion from the existence of a large bounded-degree bag in the topological-minor-free decomposition. Several technical steps are needed to translate this into the desired win/win, due to the existence of vortices, virtual edges, and the lack of a direct polynomial-time algorithm to compute the decomposition.

To prove Theorem 1.2, we introduce a new algorithmic tool for finding irrelevant vertices for the k -PATH problem in the presence of a modulator M in the input graph G . Since Theorem 1.1

can be applied to find a k -path in $G - M$ if one exists, the challenge is to detect a k -path in G that jumps between M and $G - M$ several times. The absence of a k -path in $G - M$ implies it has a tree decomposition of width $\text{poly}(k)$ and constant adhesion. Using Theorem 1.1 as a subroutine, along with a packing argument, we can compute a vertex set X of size polynomial in $k + |M|$ with the following guarantee. If there is a k -path, then there is a *guarded* k -path P in which each successive pair of vertices in $M \cap P$ are connected by a subpath through $G - M$ that intersects X . Using the tree decomposition of $G - M$, the standard ancestor-marking technique allows us to identify a vertex subset C of $G - (M \cup X)$ that is adjacent to constantly many vertices from X . Unless G is already small, we can find such a set C that is sufficiently large to be reducible but small enough that we may invoke the oracle for questions about it. We can then reduce the graph without losing the existence of a guarded k -path, by marking a witness for each sensible way in which a constant-size subset from M can connect to prescribed vertices in X through C . The fact that C only has constantly many neighbors in X implies that there are only polynomially many relevant choices. We may then safely remove the unmarked vertices.

Organization After preliminaries in Section 2, we give a generic Turing-style reduction rule for k -PATH in Section 3. In Section 4 we show that an H -minor-free graph either has a k -path or a separation that is suitable for reduction. In Section 5 we extend this to topological minors. Finally, in Section 6 we present a Turing kernel applicable when the input graph has a small modulator to a suitable graph class.

2 Preliminaries

Notation. All graphs we consider are finite, simple, and undirected. A *separation* of a graph G is a pair (A, B) , $A, B \subseteq V(G)$ such that $A \cup B = V(G)$ and there are no edges between $A \setminus B$ and $B \setminus A$. The *order* of the separation (A, B) is $|A \cap B|$. A graph is *triconnected* if it is connected and cannot be disconnected by deleting fewer than three vertices. When referring to the *size* of a graph in our statements, we mean the number of vertices.

A *tree decomposition* of a graph G is a pair (T, \mathcal{X}) where T is a rooted tree and \mathcal{X} is a function that assigns to every node $t \in V(T)$ a subset $\mathcal{X}(t)$ of $V(G)$ called a *bag* such that:

- $\bigcup_{t \in V(T)} \mathcal{X}(t) = V(G)$;
- for each edge $uv \in E(G)$, there is a node $t \in V(T)$ with $u, v \in \mathcal{X}(t)$;
- for each $v \in V(G)$, the nodes $\{t \mid v \in \mathcal{X}(t)\}$ induce a (connected) subtree of T .

The *width* of (T, \mathcal{X}) is $\max_{t \in V(T)} |\mathcal{X}(t)| - 1$. Its *adhesion* is $\max_{tt' \in E(T)} |\mathcal{X}(t) \cap \mathcal{X}(t')|$. We also call the set $\mathcal{X}(t) \cap \mathcal{X}(t')$ the *adhesion of tt'* , for every edge tt' of T . For a decomposition (T, \mathcal{X}) of G and a node $t \in V(T)$, the *torso*, denoted $\text{TORSO}(G, \mathcal{X}(t))$, is the graph obtained from $G[\mathcal{X}(t)]$ by adding an edge between each pair of vertices in $\mathcal{X}(t) \cap \mathcal{X}(t')$, for every neighbor t' of t in T (so each adhesion induces a clique in the torso). Added edges not present in G are called *virtual edges*. For a subtree $T' \subseteq T$ we write $\mathcal{X}(T')$ for the union $\bigcup_{t \in V(T')} \mathcal{X}(t)$ of bags in T' .

For an edge $t_1 t_2 \in E(T)$, let T_i be the connected component of $T - \{t_1 t_2\}$ that contains t_i . Let $V_i = \mathcal{X}(T_i)$. Observe that the properties of a tree decomposition imply that (V_1, V_2) is a separation with $V_1 \cap V_2 = \mathcal{X}(t_1) \cap \mathcal{X}(t_2)$.

A decomposition (T, \mathcal{X}) is *connected* if for every $t \in V(T)$ and its child t' , if $T_{t'}$ is the subtree of T rooted at t' , we have (i) that $G[\mathcal{X}(T_{t'}) \setminus \mathcal{X}(t)]$ is connected, and (ii) that $\mathcal{X}(T_{t'}) \setminus \mathcal{X}(t)$ has edges to every vertex of the adhesion $\mathcal{X}(t) \cap \mathcal{X}(t')$. It is straightforward to turn any decomposition into a connected one without increasing its width nor adhesion. For (i), as long as there exists a pair (t, t') violating the requirement, make a distinct copy $T_{t'}^C$ of $T_{t'}$ for each connected component C of $G[\mathcal{X}(T_{t'}) \setminus \mathcal{X}(t)]$, restrict the bags of $T_{t'}^C$ to the vertices of $C \cup \mathcal{X}(t)$ only, and attach $T_{t'}^C$ as a subtree with the root being a child of t . For (ii), while there is a vertex $v \in \mathcal{X}(t) \cap \mathcal{X}(t')$ that has no neighbors in $\mathcal{X}(T_{t'}) \setminus \mathcal{X}(t)$, simply remove v from all bags in $T_{t'}$.

We will also need the following non-standard complexity measure of a tree decomposition (T, \mathcal{X}) . For every $t \in V(T)$, the number of distinct adhesions $\mathcal{X}(t) \cap \mathcal{X}(t')$ for $t' \in N_T(t)$ is called the *adhesion degree* of t . The maximum adhesion degree over all nodes t is the *adhesion degree* of the decomposition (T, \mathcal{X}) . Observe that if a tree decomposition (T, \mathcal{X}) has width less than ℓ and adhesions of size at most h , then its adhesion degree is at most

$$\sum_{i=0}^h \binom{\ell}{i} \leq (1 + \ell)^h.$$

However, in sparse graph classes we can prove a much better bound on the adhesion degree due to linear bounds on the number of cliques in such graphs; cf. Lemma 4.3.

A *path decomposition* is a tree decomposition where T is a path; we will denote it simply as a sequence of bags $\mathcal{Z}_1, \dots, \mathcal{Z}_{|V(T)|}$.

For an integer θ , a *tangle of order θ* in a graph G is a family \mathcal{T} of separations of order at most θ such that, for every separation (X, Y) of order at most θ in G , exactly one of the separations (X, Y) and (Y, X) belongs to \mathcal{T} . If $(X, Y) \in \mathcal{T}$ we call X the *small side* of the separation and Y the *large side*. Furthermore, we require that for every three separations (X_1, Y_1) , (X_2, Y_2) , (X_3, Y_3) in \mathcal{T} , we have $G[X_1] \cup G[X_2] \cup G[X_3] \neq G$.

3 Turing kernels

In this section we introduce a general toolbox and notation for proving our Turing kernel bounds.

3.1 Definitions and the auxiliary problem

For a parameterized problem Π and a computable function f , a *Turing kernel of size f* is an algorithm that solves an input instance (x, k) of Π in polynomial time, given access to an oracle that solves instances (x', k') of Π with $|x'|, k' \leq f(k)$. A Turing kernel is a *polynomial* one if f is a polynomial.

If we are only interested in distinguishing between NP-complete problems admitting a polynomial Turing kernel from the ones that do not admit such a kernel, we can assume that the oracle solves an arbitrary problem in NP, not necessarily the k -PATH problem. Indeed, note that by the definition of NP-completeness, an oracle to a problem in NP can be implemented with an oracle to k -PATH with only polynomial blow-up in the size of the passed instances.

In our work, it will be convenient to reduce to the AUXILIARY LINKAGE problem, defined as follows. The input consists of an undirected graph G' , an integer k' , a set of terminals $S \subseteq V(G')$, and a number of requests R_1, R_2, \dots, R_r ; a request is a set of at most two terminals. A path P_i in G is said to *satisfy a request R_i* if $V(P_i) \cap S = R_i$ and every vertex of $V(P_i) \cap S$ is an endpoint of P_i . With such an input, the AUXILIARY LINKAGE problem asks for a sequence of r paths P_1, P_2, \dots, P_r such that P_i satisfies R_i for every $1 \leq i \leq r$, $|\bigcup_{i=1}^r V(P_i)| = k'$, and every vertex of $V(G) \setminus S$ is contained in at most one path P_i (i.e., the paths P_i are vertex-disjoint, except that they may share an endpoint, but only if the requests ask them to do so).

We remark that AUXILIARY LINKAGE is a more general problem than k -PATH: an instance with $G' = G$, $k' = k$, $S = \emptyset$, $r = 1$, and $R_1 = \emptyset$ asks precisely for a k -path in G .

Clearly, the decision version of the AUXILIARY LINKAGE problem belongs to the class NP. By using its self-reducibility (cf. [18, Lemma 2]), we assume that the oracle returns a sequence of paths $(P_i)_{i=1}^r$ in case of a positive answer. That is, in all subsequent bounds on the number of AUXILIARY LINKAGE oracle calls, the bound adheres to the number of calls to an oracle that returns the actual paths P_i ; if one wants to use a decision oracle, one should increase the bound by the blow-up implied by the self-reducibility application (i.e., at most $|E(H)|$ for calls on a graph H).




Figure 1: A set A (blue) and a path with three A -traverses (bold). The path is guarded w.r.t. $Z \subseteq N(A)$ (the red set), since each A -traverse has an endpoint in it. (A path fully contained in A (with one A -traverse) or disjoint from A (with no A -traverses) would also be guarded.)

3.2 Generic reduction rule

We now show a generic reduction rule for the k -PATH problem. We start with a few definitions.

Definition 3.1. For a graph G , a subset $A \subseteq V(G)$, and a simple path P in G , an A -traverse of P is a maximal subpath of P that contains at least one vertex of A and has all its internal vertices in A .

Note that if Q is an A -traverse of P , then every endpoint of Q is either an endpoint of P or lies in $N_G(A)$. See Figure 1.

Definition 3.2. Let G be a graph, $A \subseteq V(G)$, and let k be an integer. A set $Z \subseteq N(A)$ is called a k -guard of A if the following implication holds: if G admits a k -path, then there exists a k -path P in G that is either contained in A or such that every A -traverse of P has at least one endpoint in Z .

Given a graph G , a set $A \subseteq V(G)$, and a k -guard $Z \subseteq N(A)$ of A , a k -path P satisfying properties as in the above definition is called *guarded* (w.r.t. k , A , and Z). If the integer k and the set A are clear from the context, we call such a set Z simply a guard.

Observe that $Z = N(A)$ is always a guard, but sometimes we will be able to find smaller ones. Of particular interest will be guards of constant size, as our kernel sizes will depend exponentially on the guard size. To describe our single reduction rule, we show how solutions to AUXILIARY LINKAGE can be used to preserve the existence of guarded k -paths.

Assume we are given a graph G , a set $A \subseteq V(G)$, an integer k , and a k -guard $Z \subseteq N(A)$ of A . Let $h = |Z|$ and $\ell = |N(A)|$. Furthermore, assume that G admits a k -path, and let P be a guarded one w.r.t. A and Z . Let (Q_1, Q_2, \dots, Q_r) be the A -traverses of P , let $R_i = V(Q_i) \setminus A = V(Q_i) \cap N(A)$ for $1 \leq i \leq r$, let $G' = G[N[A]]$, $S = N(A)$, and let $k' = |\bigcup_{i=1}^r V(Q_i)|$. Observe that (Q_1, Q_2, \dots, Q_r) is a feasible solution to the AUXILIARY LINKAGE instance $\mathcal{I}_P := (G[N[A]], k', S, (R_i)_{i=1}^r)$; the instance \mathcal{I}_P is henceforth called *induced by P and A* . Furthermore, it is easy to see that if $(Q'_1, Q'_2, \dots, Q'_r)$ is a different feasible solution to \mathcal{I}_P , then a path P' obtained from P by replacing every subpath Q_i with Q'_i is also a guarded k -path in G .

The crucial observation is that a small guard limits the number of A -traverses.

Lemma 3.3. *The number r of traverses of the guarded k -path P is bounded by $\max(1, 2|Z|)$.*

Proof. Every vertex of Z can be an endpoint of at most two traverses. If $r > 1$, then none of the traverses Q_i are contained in $G[A]$, and thus every traverse has at least one endpoint in the guard Z . \square

Lemma 3.3 in turn limits the number of possible instances \mathcal{I} that can be induced by a guarded k -path, for a fixed set A and guard Z . Note that we have $0 \leq k' \leq k$ and $0 \leq r \leq \max(1, 2|Z|)$. Furthermore, unless $r = 1$ and $R_1 = \emptyset$, we have $R_i \subseteq N(A)$, $|R_i| \in \{1, 2\}$, and every set R_i needs to have at least one element of Z ; there are at most $|Z| \cdot (|N(A)| + 1) = h(\ell + 1)$ choices for such a set R_i . Consequently, the number of possibilities for the instance \mathcal{I} is at most

$$(k + 1) \cdot \left(1 + \sum_{r=0}^{2h} h^r (\ell + 1)^r \right) \leq (k + 1) \cdot (h(\ell + 1))^{2h+1} =: \mathfrak{p}(k, \ell, h). \quad (1)$$

Reduction rule. If $|A| > k \cdot \mathfrak{p}(k, \ell, h)$, then we can apply the following reduction rule. For each AUXILIARY LINKAGE instance \mathcal{I} out of at most $\mathfrak{p}(k, \ell, h)$ reasonable instances for A -traverses of a guarded k -path in G , we invoke an oracle on the instance \mathcal{I} , and mark the vertices of the solution if the oracle finds one. The whole process will mark at most $k \cdot \mathfrak{p}(k, \ell, h) < |A|$ vertices, thus at least one vertex of $|A|$ will remain unmarked. We delete any such vertices.

The observation that on a guarded k -path P one can replace a solution to the instance \mathcal{I}_P induced by P and A by a different solution provides safeness of this reduction. Finally, note that the reduction invokes at most $\mathfrak{p}(k, \ell, h)$ calls to the oracle; each call operates on a subgraph of the graph $G[N[A]]$ with $k' \leq k$ and $r \leq 2|Z|$.

We shall apply the Reduction Rule for a medium-sized set A and a guard set Z of constant size formed from adhesions of a tree decomposition. For most of the paper we will use $Z = N(A)$ with $\ell = h = |Z|$ a constant (depending on the excluded (topological) minor, in the results of Sections 4 and 5). Only in Section 6, when dealing with a modulator M such that $G - M$ has an appropriate structure, it will be important to consider $N(A)$ potentially containing all of M , with a guard set Z of constant size disjoint from M .

3.3 Separation oracles

The natural way of using our reduction rule is to find in a graph a large (but not too large) part of the graph with a small (preferably, constant) boundary. Let us first make an abstract definition of an algorithm finding such a separation.

Definition 3.4. For a graph class \mathcal{G} , a constant h , and a computable coordinate-wise nondecreasing function $q : \mathbb{Z}_{\geq 0} \times \mathbb{Z}_{\geq 0} \rightarrow \mathbb{Z}_{\geq 0}$, an algorithm \mathcal{S} is called a $(h, q, T_{\mathcal{S}})$ -separation oracle if, given a graph $G \in \mathcal{G}$ and integers k and p , in time $T_{\mathcal{S}}(|G|, k, p)$ it finds a separation (A, B) in G of order at most h with $p < |A| \leq q(k, p)$, or correctly concludes that G contains a k -path.

For all considered graph classes, we will be able to provide a separation oracle with q being a polynomial. This, in turn, allows the following generic Turing kernel.

Lemma 3.5. *Let \mathcal{S} be a $(T_{\mathcal{S}}, h, q)$ -separation oracle for a hereditary graph class \mathcal{G} . Take $\hat{h} := (2h)^{4h+3}$. Then, the k -PATH problem restricted to graphs from \mathcal{G} can be solved:*

- in time $\mathcal{O}(T_{\mathcal{S}}(|G|, k, k^2\hat{h}) \cdot |V(G)| + k\hat{h} \cdot |V(G)| \cdot |E(G)|)$,
- using at most $k\hat{h} \cdot |V(G)|$ calls to AUXILIARY LINKAGE
- each call on an induced subgraph of the input graph of size at most $q(k, k^2\hat{h})$.

Proof. Let

$$p = k \cdot \mathfrak{p}(k, h, h) \leq k(k + 1)(h(h + 1))^{2h+1} \leq 2k^2(2h)^{4h+2} \leq k^2\hat{h}.$$

As long as $|V(G)| > p$, we proceed as follows. Invoke algorithm \mathcal{S} on G . If \mathcal{S} claims that G admits a k -PATH, we simply output the answer yes. Otherwise, let (A', B') be the separation output by \mathcal{S} . Apply the Reduction Rule for k , $A := A' \setminus B'$, and $Z = N(A) \subseteq A' \cap B'$. Note that

as $|Z| \leq h$, the Reduction Rule deletes at least one vertex of A . Furthermore, the Reduction Rule invokes at most

$$p(k, h, h) \leq k\hat{h}$$

calls to the oracle, each call on an induced subgraph of G of size at most

$$|A'| \leq q(k, p) = q(k, k^2\hat{h}).$$

Once we obtain $|V(G)| \leq p$, we solve the instance using a single call to AUXILIARY LINKAGE with $k' = k$, $r = 1$, and $R_1 = \emptyset$. The bounds follow, as there are at most $|V(G)|$ applications of the Reduction Rule, and each call to the oracle takes $\mathcal{O}(|E(G)|)$ time to prepare the instance and parse the output. \square

Note that for any graph class where separations as in Definition 3.4 exist, there exists a trivial separation oracle which finds them, running in time $n^{h+\mathcal{O}(1)}$: one iterates over every candidate for $A \cap B$ and, for fixed set $A \cap B$, a straightforward knapsack-type dynamic programming algorithm checks if one can assemble $A \setminus B$ of the desired size from the connected components of $G - (A \cap B)$.

However, this running time bound is unsatisfactory, as it greatly exceeds the number of used oracle calls. For all considered graph classes we prove a much stronger property than just merely the prerequisites of Lemma 3.5, in particular providing a more efficient separation oracle. We provide necessary definitions in the next section.

3.4 Decomposable graph classes

The following definition captures the key concept of this section.

Definition 3.6. For a constant h and a computable nondecreasing function $w : \mathbb{Z}_{\geq 0} \rightarrow \mathbb{Z}_{\geq 0}$, a graph class \mathcal{G} is called (w, h) -decomposable if for every positive integer k and every $G \in \mathcal{G}$ that does not admit a k -path, the graph G admits a tree decomposition of width less than $w(k)$ and adhesions of size at most h .

A standard argument shows that in a decomposable graph class, given the decomposition with appropriate parameters, it is easy to provide a separation oracle.

Lemma 3.7. *Assume we are given a graph G and a tree decomposition (T, \mathcal{X}) of G of width less than w , adhesion at most h , and adhesion degree at most $a \geq 2$. Then, given an integer p such that $|V(G)| > p$, one can in time $h^{\mathcal{O}(1)} \cdot (|V(G)| + |E(G)| + |V(T)| + \sum_{t \in V(T)} |\mathcal{X}(t)|)$ find a separation (A, B) of order at most h such that*

$$p < |A| \leq w + p \cdot a.$$

Proof. Root the tree T in an arbitrary node, and for $t \in V(T)$ let T_t be the subtree of T rooted in t . Let t_0 be the lowest node of T such that $|\mathcal{X}(V(T_{t_0}))| > p$; such a node can be computed in linear time in the size of G and (T, \mathcal{X}) .

Group the children t' of t according to their adhesions $\mathcal{X}(t') \cap \mathcal{X}(t)$. Due to the bound on the adhesion degree, there are at most a groups. For every adhesion S , let $X_{t,S}$ be the set of the children of t with $S = \mathcal{X}(t') \cap \mathcal{X}(t)$. Define

$$V_S = \bigcup_{t' \in X_{t,S}} \mathcal{X}(T_{t'}).$$

We consider now two cases. First, assume that $|V_S| \leq p$ for every adhesion S . Then, by the adhesion degree bound, we have

$$|\mathcal{X}(T_t)| \leq |\mathcal{X}(t)| + ap \leq w + ap.$$

Consequently, we can return the separation (A, B) with $A = \mathcal{X}(T_t)$ and $B = \mathcal{X}(T - V(T_t))$.

In the other case, there exists an adhesion S with $V_S > p$. We greedily take a minimal subset $Y_{t,S} \subseteq X_{t,S}$ such that

$$V'_S := \bigcup_{t' \in Y_{t,S}} \mathcal{X}(T_{t'})$$

is of size greater than p . By the minimality of t , for every $t' \in X_{t,S}$ we have $|\mathcal{X}(T_{t'})| \leq p$ and, consequently $|V'_S| \leq 2p$. Thus, we can return the separation (A, B) for $A = V'_S$ and $B = N_G[V(G) \setminus V'_S]$, as then $A \cap B \subseteq S$. \square

A critical insight is that the decomposition used by Cygan et al. [9] to solve the MINIMUM BISECTION problem in fact provides an approximate decomposition in a decomposable graph class. Let us first recall the main technical result of [9].

Definition 3.8. A vertex set $X \subseteq V(G)$ of a graph G is called (q, h) -unbreakable if every separation (A, B) of order at most h satisfies $|(A \setminus B) \cap X| \leq q$ or $|(B \setminus A) \cap X| \leq q$.

Theorem 3.9 ([9]). *There is an algorithm that given a graph G and integer h runs in time $2^{\mathcal{O}(h^2)}|V(G)|^2|E(G)|$ and outputs a connected tree decomposition (T, \mathcal{Y}) of G such that: (i) for each $t \in V(T)$, the bag $\mathcal{Y}(t)$ is $(2^{\mathcal{O}(h)}, h)$ -unbreakable in G , and (ii) for each $tt' \in E(T)$ the adhesion $\mathcal{Y}(t) \cap \mathcal{Y}(t')$ has at most $2^{\mathcal{O}(h)}$ vertices and is $(2h, h)$ -unbreakable in G .*

Lemma 3.10. *Let G be a graph and suppose there exists a decomposition (T, \mathcal{X}) of G of width less than w , adhesion h , and adhesion degree a . Let (T', \mathcal{Y}) be a tree decomposition of G such that for each $t \in V(T')$, the bag $\mathcal{Y}(t)$ is $(2^{\mathcal{O}(h)}, h)$ -unbreakable in G . Then $|\mathcal{Y}(t)| \leq w + a \cdot 2^{\mathcal{O}(h)}$.*

Proof. Consider a bag $Y = \mathcal{Y}(t)$ of the second decomposition, or any $(2^{\mathcal{O}(h)}, h)$ -unbreakable set $Y \subseteq V(G)$. For every edge $tt' \in E(T)$ of the h -decomposition, removing it partitions T into subtrees $T_{t't}$ and $T_{tt'}$ containing t' and t , respectively. The induced separation $(\mathcal{X}(T_{t't}), \mathcal{X}(T_{tt'}))$ on G has order at most h , so one of the sides, say $\mathcal{X}(T_{t't}) \setminus \mathcal{X}(t)$, contains at most $2^{\mathcal{O}(h)}$ vertices of Y , by definition of unbreakability. Let us orient the edge tt' away from t' (the ‘smaller’ side). By orienting each edge of T this way, we find a single node t_0 of T such that all incident edges point to it. That is, for every neighbor t' of t_0 , we have $|Y \cap \mathcal{X}(T_{t't_0}) \setminus \mathcal{X}(t_0)| \leq 2^{\mathcal{O}(h)}$.

While t_0 can have many neighbors, we can group those neighbors t' according to the adhesion $\mathcal{X}(t') \cap \mathcal{X}(t_0)$ to which they correspond. By the adhesion degree bound, there are at most a such groups. If for any such group $U \subseteq N(t_0) \subseteq V(T)$, the union $\bigcup_{t' \in U} \mathcal{X}(T_{t't_0}) \setminus \mathcal{X}(t_0)$ contained more than $3 \cdot 2^{\mathcal{O}(h)}$ vertices of Y , then the group can be partitioned into two parts with more than $2^{\mathcal{O}(h)}$ vertices of Y each. This would give a separation of order h of G with too many vertices of Y on both sides, contradicting its unbreakability. Therefore, there are at most a groups, each containing at most $3 \cdot 2^{\mathcal{O}(h)}$ vertices of Y , thus the size of Y is bounded by $|\mathcal{X}(t_0)| + a \cdot 3 \cdot 2^{\mathcal{O}(h)}$ \square

Thus, the decomposition computed by the algorithm of Theorem 3.9 approximates the desired decomposition of a decomposable graph class.

Corollary 3.11. *Let \mathcal{G} be a (w, h) -decomposable graph class. Then, for every $G \in \mathcal{G}$ and every integer k , one can in $2^{\mathcal{O}(h^2)}|V(G)|^2|E(G)|$ time either correctly conclude that G admits a k -path, or find a tree decomposition of G of width at most $(w(k) + 1)^{\mathcal{O}(h)}$ and adhesion at most $2^{\mathcal{O}(h)}$.*

Let us now combine all the above. That is, given an integer k and a graph G from a hereditary (w, h) -decomposable graph class \mathcal{G} , we start by computing the tree decomposition of Corollary 3.11 (or conclude there is a k -path). In general this approximated decomposition has adhesion degree $(w(k) + 1)^{2^{\mathcal{O}(h)}}$. We use this decomposition to find separations of any induced subgraphs of G using the algorithm of Lemma 3.7 in time $2^{\mathcal{O}(h)}$ times linear in the size of G and the computed decomposition. This gives a (h, q, T) -separation oracle with $q(k, p) = p \cdot (w(k) + 1)^{2^{\mathcal{O}(h)}}$ and $T(n, k, p) = 2^{\mathcal{O}(h)} \cdot n \cdot (w(k) + 1)^{\mathcal{O}(h)}$, for any hereditary (w, h) -decomposable graph class. By plugging it into Lemma 3.5, we obtain the following.

Corollary 3.12. *Let \mathcal{G} be a hereditary (w, h) -decomposable graph class. Then, the k -PATH problem, restricted to graphs from \mathcal{G} , can be solved in time $2^{2^{\mathcal{O}(h)}}|V(G)|^2|E(G)|$ using $2^{2^{\mathcal{O}(h)}}kn$ calls to AUXILIARY LINKAGE on induced subgraphs of the input graph of size $k^2(1 + w(k))^{2^{\mathcal{O}(h)}}$.*

In the next section, we prove that H -minor-free graphs are $(\mathcal{O}_H(k^{22}), \mathcal{O}_H(1))$ -decomposable by analyzing the Global Structure Theorem of H -minor-free graphs due to Robertson and Seymour. A subsequent section provides an analogous result for graphs excluding a fixed topological minor. In both cases we also get better bounds on the adhesion degree of the approximate decomposition outputted by Theorem 3.9, improving the bounds in the final kernel.

We would like to remark that we do not want to claim in this paper the idea that, in the context of H -(topological)-minor-free graphs, the decomposition of Theorem 3.9 should be related to the decomposition of the Global Structure Theorem via an argument as in the proof of Lemma 3.10. In particular, this observation appeared previously in a work of the second author with Daniel Lokshtanov, Michał Pilipczuk, and Saket Saurabh [21].

4 Excluding a minor

In this section we tackle proper minor-closed graph classes, that is, we prove Theorem 1.1 for graph classes excluding a fixed minor, by proving the following.

Theorem 4.1. *For every fixed graph H , the k -PATH problem restricted to H -minor-free graphs can be solved in time $\mathcal{O}_H(n^2m)$ using $\mathcal{O}_H(kn)$ calls to AUXILIARY LINKAGE on instances being induced subgraphs of the input graph of size $\mathcal{O}_H(k^{24})$.*

Our main technical result is the following:

Theorem 4.2. *For every graph H , the class of H -minor-free graphs is (w, h) -decomposable for $w(k) = \mathcal{O}_H(k^{22})$ and $h = \mathcal{O}_H(1)$.*

By plugging the above into Corollary 3.12, we obtain the desired polynomial Turing kernel, but with worse bounds than promised by Theorem 4.1. To obtain better bounds, we need to recall the folklore bound on the adhesion degree in sparse graph classes; for completeness, we provide a full proof in Appendix A.1.

Lemma 4.3. *Let G be a graph not containing H as a topological minor, and let (T, \mathcal{X}) be a connected tree decomposition of G of width less than ℓ and adhesion h . Then the adhesion degree of (T, \mathcal{X}) is bounded by $f(h, H) \cdot \ell$ for some integer $f(h, H)$ depending only on h and H .*

This way, we conclude that H -minor-free graphs without k -paths have tree decompositions of width $\mathcal{O}_H(k^{22})$, adhesion $\mathcal{O}_H(1)$ and adhesion degree $\mathcal{O}_H(k^{22})$. We can use the algorithm of Theorem 3.9 for h to find (by Lemma 3.10) an approximate decomposition of width $\mathcal{O}_H(k^{22})$, adhesion $\mathcal{O}_H(1)$, and thus, again using Lemma 4.3, of adhesion degree $\mathcal{O}_H(k^{22})$. Theorem 4.1 follows from Lemma 3.5 if we find separations using the algorithm of Lemma 3.7 applied to this decomposition.

Thus, it remains to prove Theorem 4.2. For the proof, we use the graph minors structure theorem, decomposing an H -minor-free graph G into parts ‘nearly embeddable’ in surfaces (precise definitions are given in the next subsection). By carefully analyzing details of the structure, we either find a large triconnected embedded part, which must contain a long path by the following theorem of Chen et al. [8], or we tighten the graph structure to give a tree decomposition where all parts are small (polynomial in k) and adhesions (‘boundaries’) between them are of constant size.

Theorem 4.4 ([8]). *There is a constant $\varepsilon > 0$ such that for every integer t , every triconnected graph on $n \geq 3$ vertices embeddable in a surface of (Euler) genus g contains a cycle of length at least $n^\varepsilon / 2^{(2g+3)^2}$.*

We note that Chen et al. phrase the theorem (more generally) for $K_{3,t}$ -minor-free graphs, but a folklore edge-counting argument shows that graphs embeddable on a surface of genus g are $K_{3,t}$ -minor-free for $t = 2g + 3$ (see e.g. [2]).

Two intertwined problems that arise with this approach is that torsos of decompositions are not necessarily triconnected, and long paths in them do not necessarily imply long paths in the original graph, because of virtual edges added in torsos. Torsos can be made triconnected if their near-embeddings include cycles or paths around each vortex, but these may use virtual edges in essential ways. On the other hand, the decomposition can be modified so that virtual edges can be replaced with paths in the original graph, but this requires changes that remove virtual edges, hence potentially removing paths around vortices and destroying triconnectedness.

Because of that, we need to go a little deeper and use a local, strong version of the structure theorem from Graph Minors XVII [25]. For the same reason we cannot use existing algorithms for finding the graph minors decompositions. Instead, we only prove the *existence* of a tree decomposition of bounded adhesion, small width, and with nearly embeddable bags.

Global and local graph minor structure theorems

We now define near-embeddability and the graph minors α -decomposition.

Definition 4.5. For an integer α , an α -near embedding of G consists of:

- (i) a set A of at most α vertices (called the *apex* set);
- (ii) a family $G_0 \cup G_1 \cup \dots \cup G_{\alpha'} = G \setminus A$ of edge-disjoint subgraphs of $G \setminus A$, where:
 - G_0 is called the *embedded part*,
 - $\mathcal{V} = \{G_1, \dots, G_{\alpha'}\}$ for some $\alpha' \leq \alpha$ are called (*large*) *vortices*,
 - the intersection $\Omega_i := V(G_0) \cap V(G_i)$ is called the *society* of vortex G_i ,
 - vortices are pairwise vertex-disjoint,
- (iii) an embedding of G_0 in a surface of genus at most α such that for $i \in [\alpha']$, the society Ω_i is embedded on the boundary of a disk whose interior is empty (i.e., does not intersect the embedding or other disks), called the disk *accomodating* G_i ;
- (iv) a linear ordering $w_1, \dots, w_{|\Omega_i|}$ of each vortex society Ω_i , corresponding to its natural ordering around its disk (for some choice of direction and starting point);
- (v) for each large vortex $G_i \in \mathcal{V}$, a path decomposition $\mathcal{Z}_1, \dots, \mathcal{Z}_{|\Omega_i|}$ of G_i of width at most α such that $w_j \in \mathcal{Z}_j$, for $j \in [|\Omega_i|]$.

We denote such an α -near embedding as (A, G_0, \mathcal{V}) , with the embedding and path decompositions only implicit in the notation.

Definition 4.6. A (*graph minors*) α -decomposition of a graph G consists of:

- a rooted tree decomposition (T, \mathcal{Y}) of G of adhesion at most α ;
- for each $t \in V(T)$, an α -near embedding $(A_t, G_{t_0}, \mathcal{V}_t)$ of $\text{TORSO}(G, \mathcal{X}(t))$.

In this subsection we prove the following variant of the Global Structure Theorem which implies Theorem 4.2.

Theorem 4.7. *For every graph H , there is a constant $\alpha = \alpha(H)$ such that the following holds, for any integer k : any graph G excluding H as a minor and without a k -path has an α -decomposition of width at most $\alpha \cdot k^{22}$.*

We deduce Theorem 4.7 from a similarly modified variant of the Local Structure Theorem. An α -near local embedding of a graph G is defined similarly to an α -near embedding, but we allow an arbitrary number of ‘small vortices’ \mathcal{W} and we allow the path decompositions of all vortices to have only bounded adhesion instead of bounded width; thus arbitrarily complicated graphs can hide behind vortices. We additionally require each large vortex to be surrounded by a certain path. Formally:

Definition 4.8. A *comb* is a union of a path P with some mutually vertex-disjoint (possibly trivial) paths that have their first vertex and no other vertex on P . The last vertices of those paths are called the *teeth* of the comb, they are naturally ordered by P .

Definition 4.9. For an integer α , an α -near local embedding of G consists of:

- (i) a set A of at most α vertices (called the *apex* set);
- (ii) a family $G_0 \cup G_1 \cup \dots \cup G_n = G \setminus A$ of edge-disjoint subgraphs of $G \setminus A$, where:
 - G_0 is called the *embedded part*,
 - $\mathcal{V} = \{G_1, \dots, G_{\alpha'}\}$ for some $\alpha' \leq \alpha$ are called *large vortices*,
 - $\mathcal{W} = \{G_{\alpha'+1}, \dots, G_n\}$ are called *small vortices*,
 - vortices intersect only in G_0 : $V(G_i) \cap V(G_j) \subseteq V(G_0)$ for $i \neq j \in [n]$,
 - the intersection $\Omega_i := V(G_0) \cap V(G_i)$ is called the *society* of vortex G_i ,
 - large vortices are pairwise vertex-disjoint,
 - small vortices have societies of size ≤ 3 ;
- (iii) an embedding of G_0 in a surface of genus at most α such that for $i \in [n]$, the society Ω_i is embedded on the boundary of a disk whose interior is empty (i.e., does not intersect the embedding or other disks), called the disk *accomodating* G_i ;
- (iv) a linear ordering $w_1, \dots, w_{|\Omega_i|}$ of each society Ω_i , corresponding to its natural ordering around its disk (for some choice of direction and starting point);
- (v) for each large vortex $G_i \in \mathcal{V}$, a path decomposition $\mathcal{Z}_1, \dots, \mathcal{Z}_{|\Omega_i|}$ of G_i of adhesion at most α such that $\mathcal{Z}_j \cap \Omega_i = \{w_{j-1}, w_j\}$ for $1 < j \leq |\Omega_i|$ and $\mathcal{Z}_1 \cap \Omega_i = \{w_1\}$.
- (vi) for each large vortex $G_i \in \mathcal{V}$, a comb C_i in $G_i \cup G_0 \cup \bigcup \mathcal{W}$ whose teeth are vertices $w_1, \dots, w_{|\Omega_i|}$ of Ω_i , in the same order.

We denote such an α -near local embedding as $(A, G_0, \mathcal{V}, \mathcal{W})$, with the embedding and path decompositions only implicit in the notation. The near embedding is said to *respect* a tangle \mathcal{T} if the large side of a separation in $\mathcal{T} \setminus A$ is never contained in a vortex $W \in \mathcal{W}$ or in a bag of the decomposition of a vortex $V \in \mathcal{V}$.

Definition 4.10. In a graph G , a tangle \mathcal{T} *controls an H -minor* if there is a minor model in G (defined by branch sets $(B_h)_{h \in V(H)}$ inducing vertex-disjoint connected subgraphs in G and with an edge between B_h and $B_{h'}$ whenever $hh' \in E(H)$) such that no branch set is fully contained in a small side of a separation in \mathcal{T} .

The following Local Structure Theorem follows from [25], as explained in [12] (we note that while [12] assumes that G is H -minor-free, the original statement in [25] only requires that \mathcal{T} controls no H -minor – we will need this stronger version when dealing with topological minors). Since this is crucial for our approach, we stress that the comb C_i contains a path of length at least $|\Omega_i|$ (a scrupulous reading of [25] implies even stronger statements, but this will suffice).

Theorem 4.11 ([12]). *For every graph H there exist integers α, θ such that: for every graph G and every tangle \mathcal{T} in G of order $\geq \theta$ that controls no H -minor, there is an α -near local embedding of G which respects \mathcal{T} .*

We now improve the statement of the Local Structure Theorem 4.11 in steps: first requiring small vortices to be ‘well attached’, then making the torsos triconnected, and finally deducing bounds on width in terms of longest path length.

Lemma 4.12. *In Theorem 4.11, we can additionally assume that:*

- (vii) *for each small vortex $G_i \in \mathcal{W}$ and every $v_1, v_2 \in \Omega_i$, there is a path in G_i between v_1 and v_2 , with no internal vertices in Ω_i .*

Proof. Property (vii) is guaranteed by (9.1) in [25]. (As explained in [12], the border cells of [25] are translated into bags of large vortex decompositions, so small vortices arise only from internal cells; for an internal cell c_0 , $\alpha(c_0)^*$ is translated into a small vortex G_i , while $\gamma(\tilde{c}_0)$ is translated into Ω_i , giving exactly the statement we want. We also note that the proof of (9.1) simply partitions a vortex with no v_1 - v_2 path into two vortices with at most two vertices in their societies). \square

Definition 4.13. For an α -near local embedding $(A, G_0, \mathcal{V}, \mathcal{W})$ of G , define G_0^* to be the graph obtained from G_0 by adding an edge between every two consecutive vertices in each society Ω_i , and a new vertex for each society of a large vortex, with edges to every vertex of the society. The embedding of G_0 is naturally extended to an embedding of G_0^* (the new vertices and edges embedded in place of the accommodating disks).

Lemma 4.14. *In Lemma 4.12, we can additionally assume that:*

(viii) G_0^* is triconnected.

Proof. Suppose G_0^* has a separation (X, Y) of order ≤ 2 . Since societies of small vortices induce cliques in G_0^* , they are contained in X or Y . Similarly for each large vortex $G_i \in \mathcal{V}$, its society Ω_i , together with the new vertex in G_0^* with Ω_i as neighborhood, induces a wheel in G_0^* , which is triconnected, hence contained in X or Y . In the language of [25], this translates back into a partition I, J of cells, whose intersection corresponds to at most two embedded vertices v, v' (that is, $\gamma(I) \cap \gamma(J) = \{v, v'\}$). But this contradicts (11.1) in [25], which states that such an intersection has size at least 3. \square

Lemma 4.15. *In Lemma 4.14, we can additionally assume that there is a constant α' depending on H only such that:*

(ix) for any integer k , if G does not contain a path of length $\geq k$, then $|V(G_0)| < \alpha' \cdot k^{22}$.

Proof. For a given graph H , let α be the constant given by Theorem 4.11. By Theorem 4.4, there is a universal constant ε such that any triconnected graph G_0 embeddable in a surface of genus at most α contains a cycle of length at least $|G_0|^\varepsilon / 2^{(2\alpha+2)^2}$. Let $c := \frac{2}{\varepsilon}$ and $\alpha' := (\alpha \cdot 2^{(2\alpha+2)^2})^{1/\varepsilon}$. For an H -minor-free graph G , suppose the α -near local embedding given by Lemma 4.14 has $|V(G_0)| \geq \alpha' \cdot k^c$. Then G_0^* is a triconnected graph (by property (viii)) embedded in a surface of genus at most α (by Definition 4.13) with at least as many vertices, so G_0^* must contain a cycle C of length at least $(\alpha' k^c)^\varepsilon / 2^{(2\alpha+2)^2} = \alpha \cdot k^2$.

If the society Ω_i of any vortex of the α -near local embedding of G has at least k vertices, then G contains a path of at least that length (in the comb from property (vi)), in which case the lemma follows. Otherwise, there are at most α large vortices, hence at most $\alpha \cdot (k-1)$ vertices in their societies, and at most α virtual vertices that were inserted as the centers of the wheels. Let P be the longest subpath of C between any two such vortex-related vertices (or any subpath not visiting any such vertices, if there are less than two). Then P is a path of length at least $\frac{\alpha \cdot k^2}{\alpha \cdot k} = k$ in G_0^* whose edges in $E(G_0^*) \setminus E(G_0)$ came only from small vortices, by the definition of G_0^* (Definition 4.13). By property (vii), these edges can be replaced with paths (each of length at least 2) in the corresponding small vortices, giving a path in G . Only one or two consecutive edges can come from the same small vortex (since their societies have at most 3 vertices), so the resulting path has length at least k .

The constant from Theorem 4.4 given by Chen et al. [8] is $\varepsilon = \log_{1729} 2$, hence the universal constant here is $c = \frac{2}{\varepsilon} = 2 \log_2(1729) < 22$. \square

Finally, we deduce Theorem 4.7 (the Global Structure Theorem) from Theorem 4.11 (the Local Structure Theorem) by a standard induction, exactly as done by Diestel et al. [10, Theorem 4]. (Note that in case the excluded graph H is planar, we can already conclude the theorem

trivially from the fact that there is a tree-decomposition of width $\mathcal{O}_H(1)$ [23]). The only difference is that in the Local Structure Theorem we add the bound on the size of the embedded part from Lemma 4.15. In the proof from [10], every bag of the created decomposition is either constructed as a set of size bounded as $\mathcal{O}_H(1)$, or is constructed from an α -near local embedding by taking the vertices of: the embedded part G_0 , the set of apices A , and for each large vortex, the intersection of every two consecutive bags of its path decomposition, each of size at most $\alpha = \mathcal{O}_H(1)$. The number of bags in a path decomposition of a large vortex is equal to the size of its society, and large vortex societies are disjoint subsets of $V(G_0)$. Therefore every tree decomposition bag constructed in the proof has size at most $|V(G_0)| + \alpha + \alpha \cdot |V(G_0)| = \mathcal{O}_H(k^{22})$, by Lemma 4.15. This proves the additional condition we require in Theorem 4.7.

5 Excluding a topological minor

In this section we tackle graph classes excluding a topological minor, that is, we prove Theorem 1.1 by proving the following.

Theorem 5.1. *For every fixed graph H , the k -PATH problem restricted to H -topological-minor-free graphs can be solved in time $\mathcal{O}_H(n^2m)$ using $\mathcal{O}_H(kn)$ calls to AUXILIARY LINKAGE on instances being induced subgraphs of the input graph of size $k^{\mathcal{O}_H(1)}$.*

This follows as before from the following decomposability theorem. Note the exponent in the polynomial bound on width (bag size) now depends on H .

Theorem 5.2. *For every graph H , the class of H -topological-minor-free graphs is (w, h) -decomposable for $w(k) = k^{\mathcal{O}_H(1)}$ and $h = \mathcal{O}_H(1)$.*

We proceed with the proof of the above theorem. Grohe and Marx [15] proved that when excluding a topological minor, graphs admit a similar structure as for excluding a minor, but apart from nearly embeddable parts, one needs to consider parts that have bounded degree except for a bounded number of vertices:

Definition 5.3. A graph G has (a, d) -bounded degree (for $a, d \in \mathbb{N}$) if $\leq a$ vertices of G have degree $> d$ in G .

The main ‘Global Structure Theorem’ of [15] is that graphs excluding a topological minor H admit a tree decomposition of bounded adhesion, where the torso of every bag is either nearly embeddable (as when excluding a minor), or has (h, d) -bounded degree, for $d = \mathcal{O}_H(1)$. We still need the slightly stronger notion of near embeddability (with combs) for our proofs, which is why we will work with the following definition of ‘Local Structure’, analogous to α -near local embeddings.

A *star decomposition* of a graph is a (rooted) tree decomposition $(\mathcal{T}, \mathcal{X})$ where T is a star: a root node t_{root} with only leaf nodes attached. We call the bag $\mathcal{X}(t_{\text{root}})$ the *root bag* and the remaining ones *leaf bags*. The *root torso* is $\text{TORSO}(G, \mathcal{X}(t_{\text{root}}))$. A star decomposition *respects* a tangle \mathcal{T} if the large side of any separation in \mathcal{T} is not fully contained in any single leaf bag.

Grohe and Marx [15, Lemma 6.12] proved the following Local Structure Theorem.¹

Theorem 5.4 ([15]). *For every integer k there exist integers θ, d, k' such that: for every graph G excluding K_k as a topological minor and every tangle \mathcal{T} of order $\geq \theta$ that controls a $K_{k'}$ -minor, there is a star decomposition of G that respects \mathcal{T} , has adhesion $< \theta$, and such that the root torso has (k, d) -bounded degree.*

¹Grohe and Marx [15] use a slightly stronger definition of *respecting a tangle*, which is however equivalent to the conjunction of this one and the assertion that all adhesions in the star decomposition are strictly smaller than the order of \mathcal{T} , which we prefer to state separately. They also assume a slightly weaker condition than that of a tangle controlling a minor (which they call *not being removed by the tangle*), but it is trivially implied.

Corollary 5.5. *For every graph H there exist integers θ, α, d such that: for every graph G excluding H as a topological minor and every tangle \mathcal{T} in G of order $\geq \theta$, G has one of the following:*

- *a star decomposition that respects \mathcal{T} , has adhesion α , and such that the root torso has $(|V(H)|, d)$ -bounded degree, or*
- *an α -near local embedding that respects \mathcal{T} , satisfying properties (vii)-(ix).*

Proof. For $k = |V(H)|$ let the constants given by Theorem 5.4 be θ', d', k' . Let α'', θ'' be the constants given by Theorem 4.11 for graphs excluding $K_{k'}$ as a minor. We prove the claim for $\theta = \max(\theta', \theta'')$, $\alpha = \max(\theta', \alpha'')$ and $d = d'$. If the tangle \mathcal{T} of G controls no $K_{k'}$ -minor, then the claim follows from Theorem 4.11. Otherwise \mathcal{T} does control an $K_{k'}$ -minor and Theorem 5.4 applies. \square

The main theoretical tool we use, besides the Grohe and Marx decomposition in the form of Corollary 5.5, is the following lower bound on cycles in graphs of bounded degree by Shan [27] (we note the constant $\Delta - 1$ is asymptotically optimal; for small Δ , previous results by Chen et al. [7] give a slightly better bound).

Theorem 5.6 ([27]). *If G is a triconnected graph with maximum degree at most $\Delta \geq 425$, then G has a cycle of length at least $n^{1/\log_2(\Delta-1)}/4 + 2$.*

With this in hand, we can proceed to bounding the bags of almost bounded degree.

Lemma 5.7. *In Corollary 5.5, we can additionally assume that in the first outcome, for any integer k , if G has no k -path, then the root bag has size $k^{\mathcal{O}_H(1)}$.*

Proof. Without loss of generality assume that $\theta \geq \alpha(|V(H)|+2)$, in the statement of Corollary 5.5 (otherwise increase θ). Let us first ensure that the star decomposition (T, \mathcal{X}) is connected.

Claim 5.8. *We can additionally assume that*

$$\text{For every leaf node } t \in V(T), \text{ the set } \mathcal{X}(t) \setminus \mathcal{X}(t_{\text{root}}) \text{ is connected, and} \quad (\text{C1})$$

$$\mathcal{X}(t) \setminus \mathcal{X}(t_{\text{root}}) \text{ has edges to each vertex in the adhesion } \mathcal{X}(t) \cap \mathcal{X}(t_{\text{root}}) \quad (\text{C2})$$

Proof. If (C1) fails we can partition the leaf bag into many leaf bags. If (C2) fails we can remove the adhesion vertex in question from the leaf bag. \lrcorner

We need to define a variant of the torso that is similarly connected as the root torso, but has no virtual edges (so that long paths in it imply long paths in G). Construct G_{root}^* from G as follows: contract $\mathcal{X}(t) \setminus \mathcal{X}(t_{\text{root}})$ to a single vertex x_t for each leaf $t \in T$ and merge any such vertices $x_t, x_{t'}$ that have equal neighborhoods (corresponding to the same adhesion) into one. Let X be the set of all those new vertices.

Claim 5.9. *If G has no k -path, then neither does G_{root}^* .*

Proof. For each leaf node $t \in V(T)$, the set $\mathcal{X}(t) \setminus \mathcal{X}(t_{\text{root}})$ is connected by (C1) and can therefore be contracted to a single vertex x_t by repeated edge contractions. Since G_{root}^* is a subgraph of the graph obtained by these edge contractions, it follows that G_{root}^* is a minor of G . As minor operations do not increase the length of a longest path, the claim follows. \lrcorner

Claim 5.10. *G_{root}^* has $(|V(H)|, \max(\alpha, d + 2^d))$ -bounded degree.*

Proof. Observe that G_{root}^* can be obtained from the root torso of (T, \mathcal{X}) by removing virtual edges and adding the vertices of X . Each vertex in X has degree equal to the size of some corresponding adhesion, which is $\leq \alpha$. Each vertex v in the torso gains as neighbors only vertices in X , which correspond to different (as vertex subsets) adhesions containing v ; since these adhesions induced cliques containing v in the torso (which had $(|V(H)|, d)$ -bounded degree), their number is bounded by 2^d (unless v was already one of the $\leq |V(H)|$ vertices of high degree). \lrcorner

Claim 5.11. *If G_{root}^* does not contain a path of length at least k , then G_{root}^* has a tree decomposition of adhesion $\leq |V(H)| + 2$ and width $k^{\mathcal{O}_H(1)}$.*

Proof. Assume that G_{root}^* does not contain a path of length at least k . By the previous claim, in G_{root}^* one can delete $\leq |V(H)|$ vertices to obtain a subgraph of degree bounded by $\Delta := \max(\alpha, d + 2^d)$; name this subgraph G_{root}^\dagger . Consider the Tutte decomposition [29] of G_{root}^\dagger . In modern terms (cf. [18, Theorem 1]), it is a tree decomposition of G_{root}^\dagger of adhesion at most two, whose torsos are triconnected topological minors of G_{root}^\dagger . Consequently, each torso again has degree bounded by Δ and cannot contain a path of length at least k . By Theorem 5.6, if n denotes the size of such a bag, then $n^\varepsilon/4 + 2 < k$, where $\varepsilon = 1/\log_2(\max(425, \alpha, d + 2^d) - 1)$. Thus the size of each bag is $n < (4(k - 2))^{1/\varepsilon} = k^{\mathcal{O}(d + \log_2 \alpha)} = k^{\mathcal{O}_H(1)}$. Adding the $\leq |V(H)|$ deleted vertices back to every bag, we obtain a tree decomposition of G_{root}^* as claimed. \lrcorner

Denote the root torso of (T, \mathcal{X}) as G_{root} .

Claim 5.12. *For every tree decomposition of G_{root}^* , one can find a tree decomposition of G_{root} with adhesion and maximum bag size at most α times larger than the adhesion and maximum bag size of the original tree decomposition.*

Proof. Let (S, \mathcal{Y}) be a tree decomposition of G_{root}^* . Define $\mathcal{Y}'(s)$ as $\mathcal{Y}(s)$ with vertices $x_t \in X$ replaced by all of $N_{G_{\text{root}}^*}(x_t)$. We claim (S, \mathcal{Y}') is a tree decomposition of G_{root} . Every edge of G_{root} is either an edge of $G_{\text{root}}^* - X$, still contained in the same bag, or a virtual edge between two vertices of an adhesion, hence contained in $N_{G_{\text{root}}^*}(x_t)$ for some $x_t \in X$ and thus in a bag $\mathcal{Y}'(s)$ (for an s such that $\mathcal{Y}(s)$ contained x_t). To show that the subset $\mathcal{Y}'^{-1}(v) := \{s \in S : v \in \mathcal{Y}'(s)\}$ is connected, for any given vertex $v \in V(G_{\text{root}})$, observe that it is the union of $\mathcal{Y}^{-1}(v)$ with $\mathcal{Y}^{-1}(x_t)$ for all $x_t \in X$ adjacent to v in G_{root}^* ; since such sets $\mathcal{Y}^{-1}(v)$ and $\mathcal{Y}^{-1}(x_t)$ are connected in S and intersect (as x_t is adjacent to v), their union is connected. Each vertex in each bag has been replaced by at most $|N_{G_{\text{root}}^*}(x_t)|$ vertices, which is the size of an adhesion in (T, \mathcal{X}) and hence at most α . It follows that the maximum bag size increases by a factor at most α .

To bound any adhesion, observe similarly that for $s_1, s_2 \in V(S)$, the adhesion equals

$$\mathcal{Y}'(s_1) \cap \mathcal{Y}'(s_2) = ((\mathcal{Y}(s_1) \cap \mathcal{Y}(s_2)) \setminus X) \cup \bigcup_{x_{t_1} \in \mathcal{Y}(s_1), x_{t_2} \in \mathcal{Y}(s_2)} N(x_{t_1}) \cap N(x_{t_2}).$$

For any vertex $v \in N(x_{t_1}) \cap N(x_{t_2})$ (for some $x_{t_i} \in \mathcal{Y}(s_i)$), the set $\mathcal{Y}^{-1}(v)$ is connected and must intersect both $\mathcal{Y}^{-1}(x_{t_1})$ and $\mathcal{Y}^{-1}(x_{t_2})$. Hence either $\mathcal{Y}^{-1}(x_{t_1})$ contains both s_1 and s_2 , or $\mathcal{Y}^{-1}(x_{t_2})$ does, or $\mathcal{Y}^{-1}(v)$ does. This means x_{t_1}, x_{t_2} or v is contained in $\mathcal{Y}(s_1) \cap \mathcal{Y}(s_2)$. Therefore

$$\mathcal{Y}'(s_1) \cap \mathcal{Y}'(s_2) \subseteq ((\mathcal{Y}(s_1) \cap \mathcal{Y}(s_2)) \setminus X) \cup \bigcup_{x_t \in \mathcal{Y}(s_1) \cap \mathcal{Y}(s_2)} N(x_t).$$

This again implies each adhesion increases at most α times. \lrcorner

By the last three claims, assuming G has no k -path, G_{root} has a tree decomposition (S, \mathcal{Y}) of adhesion $\alpha' := \alpha(|V(H)| + 2)$ and width $k^{\mathcal{O}_H(1)}$. We focus this decomposition on the unique bag that, in some sense, respects the tangle \mathcal{T} . Let us first see how \mathcal{T} orients separations of G_{root} .

For a separation (A, B) of G_{root} , every adhesion $\mathcal{X}(t) \cap \mathcal{X}(t_{\text{root}})$ of (T, \mathcal{X}) induces a clique in G_{root} , hence it is fully contained in A or in B ; thus (A, B) can be extended to a separation of G by adding all of $\mathcal{X}(t)$ to A or B , correspondingly, with an arbitrary choice for $\mathcal{X}(t)$ such that $\mathcal{X}(t) \cap \mathcal{X}(t_{\text{root}}) \subseteq A \cap B$. We call such separations of G *extensions* of the separation (A, B) of G_{root} . For a separation (A, B) of order $\leq \theta$ of G_{root} , observe that either all extensions of (A, B) and no extensions of (B, A) belong to \mathcal{T} , or vice-versa: this follows from tangle axioms and the fact that (T, \mathcal{X}) respects \mathcal{T} (that is, each $\mathcal{X}(t)$ such that $\mathcal{X}(t) \cap \mathcal{X}(t_{\text{root}}) \subseteq A \cap B$ is small w.r.t. \mathcal{T} , so putting it one side of an extension (A', B') instead of the other cannot change whether A' is small, as otherwise the two small sides together with $\mathcal{X}(t)$ would cover all of G).

We now orient edges of the decomposition (S, \mathcal{Y}) according to \mathcal{T} , to find a (unique) bag we can focus on. For each edge $s_1 s_2 \in E(S)$, if S_1, S_2 denote the two connected components of $S - s_1 s_2$ containing s_1, s_2 , respectively, then $(\mathcal{Y}(S_1), \mathcal{Y}(S_2))$ is a separation with $|\mathcal{Y}(S_1) \cap \mathcal{Y}(S_2)| = |\mathcal{Y}(s_1) \cap \mathcal{Y}(s_2)| \leq \alpha' \leq \theta$, hence either all of its extensions or all extensions of $(\mathcal{Y}(S_2), \mathcal{Y}(S_1))$ belongs to the tangle \mathcal{T} . We direct the edge $s_1 s_2$ to either s_2 or s_1 accordingly (towards the side with large extensions). After repeating this for every edge of S , we find a node s_0 with no outgoing edges. This means that if we root S at s_0 and S_1, S_2, \dots, S_ℓ are the subtrees of $S - s_0$, then all extensions of $(\mathcal{Y}(S_i), \mathcal{Y}(S - S_i))$ belong to \mathcal{T} , for $i \in [\ell]$.

We can now forget about some details of the tree decomposition and define the following star decomposition (S', \mathcal{Y}') of G_{root} : its nodes are $s_{\text{root}} := s_0$ and its direct children s_1, \dots, s_ℓ in S , while its bags are $\mathcal{Y}'(s_{\text{root}}) := \mathcal{Y}(s_0)$ and $\mathcal{Y}'(s_i) := \mathcal{Y}(S_i)$ for $i \in [\ell]$. By construction (S', \mathcal{Y}') is a star decomposition of G_{root} of adhesion α' , with a root bag of size $k^{\mathcal{O}_H(1)}$, which satisfies the following:

$$\text{all extensions of } (\mathcal{Y}'(s_i), \mathcal{Y}'(S - s_i)) \text{ are in } \mathcal{T}, \text{ for } i \in [\ell]. \quad (2)$$

The following claim thus concludes the proof of the lemma, by extending this star decomposition of G_{root} to a star decomposition of G respecting \mathcal{T} , with a root torso of almost bounded degree and size $k^{\mathcal{O}_H(1)}$.

Claim 5.13. *If G_{root} has a star decomposition (S, \mathcal{Y}) of adhesion $\alpha' \leq \theta$ satisfying (2), then G has a star decomposition of adhesion $\max(\alpha, \alpha')$ respecting \mathcal{T} where the root torso has $(|V(H)|, d + \alpha' \cdot d)$ -bounded degree and its vertex set is equal to the root bag of (S, \mathcal{Y}) .*

Proof. Let (S, \mathcal{Y}) be a star decomposition of G_{root} satisfying (2). Without loss of generality assume that for every leaf $s \in V(S)$, every vertex of the adhesion $\mathcal{Y}(s) \cap \mathcal{Y}(s_{\text{root}})$ has a neighbor in $\mathcal{Y}(s) \setminus \mathcal{Y}(s_{\text{root}})$ (otherwise we can remove it from $\mathcal{Y}(s)$). Every adhesion of the (larger) star decomposition (T, \mathcal{X}) of G induces a clique in G_{root} by definition, hence we can choose for each $t \in V(T)$ a node $s(t) \in V(S)$ such that $\mathcal{X}(t) \cap \mathcal{X}(t_{\text{root}}) \subseteq \mathcal{Y}(s(t))$. We construct a new star decomposition (T', \mathcal{X}') of G from (T, \mathcal{X}) as follows: the root bag becomes $\mathcal{Y}(s_{\text{root}})$, leaves t with $s(t) = s_{\text{root}}$ and their bags remain unchanged, the remaining leaves are deleted, and each leaf s of (S, \mathcal{Y}) is added as a new leaf of T' , with $\mathcal{X}'(s) := \mathcal{Y}(s) \cup \bigcup_{t: s(t)=s} \mathcal{X}(t)$.

It is straightforward to check that (T', \mathcal{X}') is a star decomposition of G of adhesion $\max(\alpha, \alpha')$, as follows. Every edge is covered by a bag in \mathcal{X}' , because it was covered by a bag in \mathcal{X} , which is either a subset of a bag in \mathcal{X}' , or in $\mathcal{X}(t_{\text{root}})$, and hence in a bag of \mathcal{Y} , which are again subsets of bags in \mathcal{X}' . Every vertex occurs in bags \mathcal{X}' of a connected subtree of T' , because if a vertex occurs in two different leaf bags of \mathcal{X}' , then either it occurs in two different leaf bags of \mathcal{Y} and hence in $\mathcal{Y}(s_{\text{root}}) = \mathcal{X}'(t_{\text{root}})$, or in two leaf bags $\mathcal{X}(t)$ and $\mathcal{X}(t')$ with $s(t) \neq s(t')$ and hence in $\mathcal{X}(t_{\text{root}})$, which means it occurs in the adhesions of $\mathcal{X}(t)$ and $\mathcal{X}(t')$; these adhesions are contained in $\mathcal{Y}(s(t))$ and $\mathcal{Y}(s(t'))$ respectively, so the vertex must be contained in $\mathcal{Y}(s_{\text{root}}) = \mathcal{X}'(t_{\text{root}})$ as well. Finally to check the sizes of adhesions, observe that each adhesion of (T, \mathcal{X}') is either (for leaf nodes t coming from T) a subset of an adhesion in \mathcal{X} , namely $\mathcal{X}'(t) \cap \mathcal{X}'(t_{\text{root}}) \subseteq \mathcal{X}(t) \cap \mathcal{X}(t_{\text{root}})$, or

(for leaf nodes s coming from S) equal to an adhesion of (S, \mathcal{Y}) :

$$\mathcal{X}'(s) \cap \mathcal{X}'(t_{\text{root}}) = \left(\mathcal{Y}(s) \cup \bigcup_{t: s(t)=s} \mathcal{X}(t) \right) \cap \mathcal{Y}(s_{\text{root}}) = \mathcal{Y}(s) \cap \mathcal{Y}(s_{\text{root}}),$$

since every vertex in $\mathcal{X}(t) \cap \mathcal{Y}(s_{\text{root}}) \subseteq \mathcal{X}(t) \cap \mathcal{X}'(t_{\text{root}})$ is contained in $\mathcal{Y}(s(t))$.

To show that the new decomposition respects \mathcal{T} , suppose to the contrary that some leaf bag contains the large side B of a separation $(A, B) \in \mathcal{T}$. Then either this bag is $\mathcal{X}'(t) = \mathcal{X}(t)$ for $t \in V(T)$, contradicting that (T, \mathcal{X}) respects \mathcal{T} , or $\mathcal{X}'(s) = \mathcal{Y}(s) \cup \bigcup_{t: s(t)=s} \mathcal{X}(t)$ for some $s \in V(S)$. Then since $(\mathcal{X}'(s), \bigcup_{s' \neq s \in V(S)} \mathcal{X}'(s'))$ is a separation induced by the adhesion at $\mathcal{X}'(s)$, of order $\leq \alpha'$ (as shown in the previous paragraph), $\mathcal{X}'(s)$ (which contains B) must be the large side of this separation in \mathcal{T} . But this is an extension of the separation $(\mathcal{Y}(s), \bigcup_{s' \neq s \in V(S)} \mathcal{Y}(s'))$, contradicting (2).

The part that is not entirely straightforward is bounding the degrees in the root torso of the new star decomposition (T', \mathcal{X}') . Observe that this new torso can be obtained from the root torso G_{root} of (T, \mathcal{X}) by taking the subgraph induced by $\mathcal{X}'(t_{\text{root}}) = \mathcal{Y}(s_{\text{root}})$, removing certain virtual edges (from adhesions corresponding to deleted leaves) and adding virtual edges between every two vertices of each new adhesion $\mathcal{X}'(s) \cap \mathcal{X}'(t_{\text{root}})$, $s \in V(S)$. We need to bound the number of such edges added to a vertex $v \in \mathcal{Y}(s_{\text{root}})$, assuming it was not already one of the $|V(H)|$ high-degree vertices of the root torso of (T, \mathcal{X}) . Thus v had at most d neighbors in G_{root} . Since for each new adhesion $\mathcal{X}'(s) \cap \mathcal{X}'(t_{\text{root}}) = \mathcal{Y}(s) \cap \mathcal{Y}(s_{\text{root}})$ (as shown above) which contains v we assumed that v has a neighbor in $\mathcal{Y}(s) \setminus \mathcal{Y}(s_{\text{root}})$ (and these are disjoint sets, for distinct s), it follows that v received new edges from at most d distinct new adhesions, at most α' edges from each. Therefore, the degree of all but $|V(H)|$ vertices in the root torso of (T', \mathcal{X}') is bounded by $d + \alpha' \cdot d$, that is, Corollary 5.5 holds with this constant in place of d . \lrcorner

This concludes the proof of Lemma 5.7. \square

As before, the global decomposition stated in Theorem 5.2 follows from the Local Structure Theorem in Corollary 5.5 improved by the bound on bag size in Lemma 5.7. The construction only changes in that it can take root bags from the local structure (the star decomposition) given by the first outcome of Corollary 5.5 and use them as bags of the global decomposition, without any modifications. We note this can also be seen as an instance of a general theorem (11.1) by Robertson and Seymour [24] which allows to turn ‘local decompositions’ (with respect to a tangle) into global ones; however, giving a precise translation would involve too much notational overhead.

6 Adding a modulator

In this section we prove Theorem 1.2 in a more general setting of Section 3. More precisely, Theorem 1.2 follows directly from the following theorem via Theorems 4.2 and 5.2.

Theorem 6.1. *One can solve in polynomial time a given k -PATH instance (G, k) , given access to a set $M \subseteq V(G)$ such that $G - M$ admits a tree decomposition of width less than w and adhesion $h = \mathcal{O}(1)$, and an oracle that solves the AUXILIARY LINKAGE problem for instances $(G', k', S, (R_i)_{i=1}^r)$ with G' being a subgraph of G , $r, k' \leq k$, $|S| \leq |M| + \mathcal{O}(1)$, and $|V(G')|$ being bounded polynomially in k, w , and $|M|$.*

Contrary to Section 3, in this section we are not precise with the polynomial dependency on the parameters k, w , and $|M|$, as well as the dependency on the adhesion $h = \mathcal{O}(1)$ of the decomposition. This is due to the fact that the latter dependency on h is substantial, most notably in the exponent of the polynomial bound on the size of the oracle calls. Therefore, the

result of this section is a purely theoretical result classifying the aforementioned parameterization as admitting a polynomial Turing kernel.

Decomposition. We apply the algorithm of Theorem 3.9 to the graph $G - M$ and separation size h . Since $h = \mathcal{O}(1)$, the algorithm runs in polynomial time. By Lemma 3.10, the output decomposition (T_0, \mathcal{X}_0) has width $w^{\mathcal{O}(1)}$ and adhesion $\mathcal{O}(1)$.

We modify the decomposition (T_0, \mathcal{X}_0) as follows. For every node $t \in V(T)$ we group the children t' of t according to the adhesions $\mathcal{X}(t) \cap \mathcal{X}(t')$. For every group $U \subseteq N_T(t)$, we add a new child t_U of t with $\mathcal{X}(t_U) = \mathcal{X}(t) \cap \mathcal{X}(t')$, $t' \in U$, and reattach the nodes of U as children of t_U . Then, we duplicate t_U into $\max(1, |U| - 1)$ copies, arranged in a binary tree, with the nodes of U being different leaves of the binary tree. Let (T, \mathcal{X}) be the final decomposition. In this manner, we maintain the property that (T, \mathcal{X}) has width $w^{\mathcal{O}(1)}$ and adhesion $\mathcal{O}(1)$, while every node has a polynomial in w number of children.

Marking. We continue the proof of Theorem 6.1 with a somewhat standard marking process. For an integer $0 \leq k' \leq k - 2$ and two vertices $u, v \in M$, define a (u, v, k') -path in G as a (simple) path with $k' + 2$ vertices: two endpoints u and v and exactly k' internal vertices, all in $G - M$. Note that a $(u, v, 0)$ -path is a path consisting of an edge uv . Similarly, for $0 \leq k' \leq k - 1$ and a vertex $u \in M$, a (u, k') -path in G is a path with $k' + 1$ vertices: one endpoint being u and all other vertices in $G - M$.

For every $0 \leq k' \leq k - 2$ and every $u, v \in M$, we compute an inclusion-wise maximal family $\mathcal{P}_{u,v,k'}^0$ of internally vertex-disjoint (u, v, k') -paths. Similarly, for every $0 \leq k' \leq k - 1$ and $u \in M$ we compute an inclusion-wise maximal family $\mathcal{P}_{u,k'}^0$ of internally vertex-disjoint (u, k') -paths. The computation can be done via the previously established kernel for decomposable graph classes.

Lemma 6.2. *The families $\mathcal{P}_{u,v,k'}^0$ and $\mathcal{P}_{u,k'}^0$ can be computed in polynomial time with the access to the assumed AUXILIARY LINKAGE oracle.*

Proof. We focus on $\mathcal{P}_{u,v,k'}^0$; the proof for the second family is analogous. Since the required family is inclusion-wise maximal, it suffices to find its paths one-by-one. That is, we show how to find, in a given induced subgraph of $G - (M \setminus \{u, v\})$ obtained by removing the internal vertices of the previously-found paths, a next path of the desired length $k' \leq k$ between u and v (or conclude that none exists and the packing is maximal). For that, we use a slight modification of the previously established kernel in Section 3 for decomposable graph classes (Corollary 3.12).

More precisely, first observe that the Reduction Rule is still correct when we are looking for a k' -path between fixed endpoints u and v (as opposed to any k' -path) as long as the endpoints do not lie in the set A . Thus, we can (just as in Corollary 3.12) iteratively use the decomposition (T, \mathcal{X}) of $G - M$, reintroduce u and v to every bag, and find a separation (A', B') by Lemma 3.7 such that the Reduction Rule applies, and iterate this until we find the k' -path between u and v or conclude there is none, as in Lemma 3.5. \square

We define $\mathcal{P}_{u,v,k'}$ to be a subfamily of $k + 1$ arbitrarily chosen paths from $\mathcal{P}_{u,v,k'}^0$, or let $\mathcal{P}_{u,v,k'} = \mathcal{P}_{u,v,k'}^0$ if $|\mathcal{P}_{u,v,k'}^0| < k + 1$. We define $\mathcal{P}_{u,k'}$ similarly.

Let $|M| = s$. Let X_1 be the set of all vertices of $G - M$ that lie on some path in any of the sets $\mathcal{P}_{u,v,k'}$ or $\mathcal{P}_{u,k'}$ ($u, v \in M$, $k' \leq k$). Note that, as we picked at most $k + 1$ paths to the families $\mathcal{P}_{u,v,k'}$ and $\mathcal{P}_{u,k'}$, each consisting of at most $k - 1$ vertices from $G - M$, for each of at most k choices of the length $k' \leq k$, and $\binom{s}{2} + s \leq 2s^2$ choices for the endpoint(s), we have

$$|X_1| < (k + 1) \cdot (k - 1) \cdot k \cdot 2s^2.$$

For every $x \in X_1$, pick one node $t(x) \in V(T)$ such that $x \in \mathcal{X}(t(x))$. Let $B_1 = \{t(x) \mid x \in X_1\}$. Let $B_2 \subseteq V(T)$ consist of all of B_1 , the root of T , and the lowest common ancestor of every pair

$(t_1, t_2) \in B_1 \times B_1$. Let $X_2 = \mathcal{X}(B_2)$. It is folklore that $|B_2| \leq 2|B_1|$, hence:

$$\begin{aligned} |B_2| &\leq 2(k+1) \cdot (k-1) \cdot k \cdot 2s^2, \\ |X_2| &= w^{\mathcal{O}(1)} k^3 s^2. \end{aligned}$$

Let \mathcal{C} be the set of connected components of $T - B_2$. As B_2 is closed under taking lowest common ancestors of nodes, we have that $|N_T(C)| \leq 2$ for any $C \in \mathcal{C}$: one of the elements of $N_T(C)$ is a parent of the root of C (if it exists), and additionally one vertex of C may have a child in $N_T(C)$.

Reducing a large component. Our main goal is to show that if the bags of any component C contain in total more than $m = \text{poly}(k, w, s)$ vertices of G , that is, $|\mathcal{X}(C)| > m$, then one can use the Reduction Rule from Section 3 to reduce the graph G by at least one vertex. Fix such a component C .

Let $t_0 \in V(C)$ be the lowest node in C such that, if D is the set of nodes of C in the subtree of T rooted in t_0 , then $|\mathcal{X}(D)| > m$.

Let $Z = \mathcal{X}(D) \cap X_2$. Recall that $|N_T(C)| \leq 2$ due to the fact that B_2 is closed under taking lowest common ancestors. By the properties of the tree decomposition, every vertex of Z lies in $\mathcal{N}(t)$ for some $t \in N_T(C)$. Since the adhesions are of constant size, we have that $|Z| = \mathcal{O}(1)$.

Let $A = \mathcal{X}(D) \setminus Z$. Note that $A \cap X_2 = \emptyset$. Since in (T, \mathcal{X}) every node has a number of children bounded polynomially in w , the choice of t_0 ensures that $|\mathcal{X}(D)| \leq \text{poly}(w) \cdot m$. Furthermore, note that $N_G(A) \subseteq Z \cup M$; in particular, $|N_G(A)| \leq |M| + \mathcal{O}(1)$.

The marking process ensures the following crucial property.

Lemma 6.3. *The set A is guarded by Z .*

Proof. Let P be a k -path in G that contains a maximum possible number of vertices from X_1 ; we claim that P is guarded by Z in A . Assume the contrary, let Q be a traverse of P through A for which neither endpoint belongs to Z . Since P is not completely contained in $G[A]$, and $N_G(A) \subseteq Z \cup M$, we have that Q is either a (u, v, k') -path for some $u, v \in M$ and $1 \leq k' \leq k-2$ or a (u, k') -path for some $u \in M$ and $1 \leq k' \leq k-1$. Assume the first case; the proof for the second case is analogous.

Since A does not contain any vertex of $X_2 \supseteq X_1$, from the maximality of $\mathcal{P}_{u,v,k'}^0$ we infer that $\mathcal{P}_{u,v,k'} \subsetneq \mathcal{P}_{u,v,k'}^0$ (since it contains Q) and, consequently, $|\mathcal{P}_{u,v,k'}| = k+1$. Thus, at least one (u, v, k') -path $Q' \in \mathcal{P}_{u,v,k'}$ has no internal vertices on the k -path P . Consequently, by replacing Q with Q' on P , we obtain a simple k -path P' with strictly more vertices of X_1 than P . This is a contradiction. \square

Thus, if $m = \text{poly}(k, w, s)$ is large enough, we can apply the Reduction Rule for the set A and the guard Z ; recall that $|Z| = \mathcal{O}(1)$ while $m - \mathcal{O}(1) < |A| \leq \text{poly}(w) \cdot m$. The Reduction Rule deletes at least one vertex after invoking a number of calls to the AUXILIARY LINKAGE oracle on the subgraph $G[N[A]]$; note that $N[A] \subseteq A \cup M \cup Z$.

Wrap up. Recall that we have shown that one can apply the Reduction Rule if there exists a component $C \in \mathcal{C}$ with $|\mathcal{X}(C)| > m$ for some m bounded polynomially in k, w , and s . However, since $|B_2| \leq 2(k+1) \cdot (k-1) \cdot k \cdot 2s^2$ and a node of T has $\text{poly}(w)$ children, there are at most $\text{poly}(w)k^3s^2$ components. Consequently, if the reduction is not applicable, we have $|V(G)|$ bounded polynomially in k, w , and s . Such an instance can be solved with a single call to AUXILIARY LINKAGE with $k' = k, S = \emptyset, r = 1$, and $R_1 = \emptyset$. This finishes the proof of Theorem 6.1 and, consequently, of Theorem 1.2.

7 Conclusions

We significantly extended the graph classes on which k -PATH has a polynomial Turing kernel. In addition, we showed that even an instance that does not belong to such a class, but has a small vertex modulator whose deletion makes it so, can be solved efficiently using small queries to an oracle. A subdivision-based argument (cf. [13]) shows that we cannot generalize much beyond H -topological-minor-free graphs without settling the problem in general. In particular, the existence of a polynomial Turing kernel for graphs of bounded expansion implies its existence in general graphs.

While our narrative focused on k -PATH, after small modifications our techniques can also be applied to prove analogues of Theorems 1.1 and 1.2 for the k -CYCLE problem of detecting a simple cycle of length *at least* k . The main difficulty in adapting our arguments to k -CYCLE is the fact that, a priori, the only cycles of length at least k may be arbitrarily much larger than k . However, this issue can easily be resolved in the following way. Since a cycle is contained within a single biconnected component, a Turing kernelization can decompose its input into biconnected components and solve the problem independently in each of them. We then start by testing for the existence of a *path* with k^2 vertices using the algorithms developed in the paper. If there is a path of length k^2 in a biconnected component, then by a classic theorem of Dirac [11] there is a cycle of length at least k , and we are done. If no such path exists, then the longest cycle in G has length less than $2k$, and we can continue under the guarantee that the cycle we are looking for has length at least k and less than $2k$. In this setting, our arguments can be easily adapted. In particular, the absence of a path of length k^2 implies the existence of suitable tree decompositions from which reducible separations can be extracted.

A significant portion of the technical work in this paper was devoted to modifying the graph minors decomposition to obtain the win/win that either answers the problem or finds a reducible separation. In this way, the algorithmic question has driven a challenging graph-theoretic project. It would be interesting to find more problems amenable to such an approach. We conclude with some concrete open problems. Does k -PATH have a polynomial Turing kernel on chordal graphs? How about INDUCED k -PATH or DIRECTED k -PATH, on planar graphs?

References

- [1] A. M. Ambalath, R. Balasundaram, C. R. H., V. Koppula, N. Misra, G. Philip, and M. S. Ramanujan. On the kernelization complexity of colorful motifs. In *Proc. 5th IPEC*, pages 14–25, 2010. doi:10.1007/978-3-642-17493-3_4.
- [2] D. Archdeacon. Topological graph theory. a survey. *Congressus Numerantium*, 115(5-54):18, 1996.
- [3] F. Barbero, C. Paul, and M. Pilipczuk. Exploring the complexity of layout parameters in tournaments and semi-complete digraphs. In *Proc. 44th ICALP*, 2017. In press. arXiv:1706.00617.
- [4] D. Binkele-Raible, H. Fernau, F. V. Fomin, D. Lokshantov, S. Saurabh, and Y. Villanger. Kernel(s) for problems with no kernel: On out-trees with many leaves. *ACM Trans. Algorithms*, 8(4):38, 2012. doi:10.1145/2344422.2344428.
- [5] H. L. Bodlaender, E. D. Demaine, M. R. Fellows, J. Guo, D. Hermelin, D. Lokshantov, M. Müller, V. Raman, J. v. Rooij, and F. A. Rosamond. Open problems in parameterized and exact computation - IWPEC 2008. Technical Report UU-CS-2008-017, Utrecht University, 2008.
- [6] H. L. Bodlaender, B. M. P. Jansen, and S. Kratsch. Kernelization lower bounds by cross-composition. *SIAM J. Discrete Math.*, 28(1):277–305, 2014. doi:10.1137/120880240.
- [7] G. Chen, Z. Gao, X. Yu, and W. Zang. Approximating longest cycles in graphs with bounded degrees. *SIAM J. Comput.*, 36(3):635–656, 2006. doi:10.1137/050633263.
- [8] G. Chen, X. Yu, and W. Zang. The circumference of a graph with no $K_{3,t}$ -minor, II. *J. Comb. Theory, Ser. B*, 102(6):1211–1240, 2012. doi:10.1016/j.jctb.2012.07.003.
- [9] M. Cygan, D. Lokshantov, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. Minimum bisection is fixed parameter tractable. In *Proc. STOC 2014*, pages 323–332. ACM, 2014. doi:10.1145/2591796.2591852.

- [10] R. Diestel, K. Kawarabayashi, T. Müller, and P. Wollan. On the excluded minor structure theorem for graphs of large tree-width. *J. Comb. Theory, Ser. B*, 102(6):1189–1210, 2012. doi:10.1016/j.jctb.2012.07.001.
- [11] G. A. Dirac. Some theorems on abstract graphs. *Proceedings of the London Mathematical Society*, s3-2(1):69–81, 1952. doi:10.1112/plms/s3-2.1.69.
- [12] J. Fröhlich and T. Müller. Linear connectivity forces large complete bipartite minors: An alternative approach. *J. Comb. Theory, Ser. B*, 101(6):502–508, 2011. doi:10.1016/j.jctb.2011.02.002.
- [13] J. Gajarský, P. Hliněný, J. Obdržálek, S. Ordyniak, F. Reidl, P. Rossmanith, F. S. Villaamil, and S. Sikdar. Kernelization using structural parameters on sparse graph classes. *J. Comput. Syst. Sci.*, 84:219–242, 2017. doi:10.1016/j.jcss.2016.09.002.
- [14] V. Garnero and M. Weller. Parameterized certificate dispersal and its variants. *Theor. Comput. Sci.*, 622:66–78, 2016. doi:10.1016/j.tcs.2016.02.001.
- [15] M. Grohe and D. Marx. Structure theorem and isomorphism test for graphs with excluded topological subgraphs. *SIAM J. Comput.*, 44(1):114–159, 2015. doi:10.1137/120892234.
- [16] D. Hermelin, S. Kratsch, K. Soltys, M. Wahlström, and X. Wu. A completeness theory for polynomial (Turing) kernelization. *Algorithmica*, 71(3):702–730, 2015. doi:10.1007/s00453-014-9910-8.
- [17] F. Hüffner, C. Komusiewicz, and M. Sorge. Finding highly connected subgraphs. In *Proc. 41st SOFSEM*, pages 254–265, 2015. doi:10.1007/978-3-662-46078-8_21.
- [18] B. M. P. Jansen. Turing kernelization for finding long paths and cycles in restricted graph classes. *J. Comput. Syst. Sci.*, 85:18–37, 2017. doi:10.1016/j.jcss.2016.10.008.
- [19] B. M. P. Jansen and D. Marx. Characterizing the easy-to-find subgraphs from the viewpoint of polynomial-time algorithms, kernels, and Turing kernels. In *Proc. 26th SODA*, pages 616–629, 2015. doi:10.1137/1.9781611973730.42.
- [20] S. Kolay and F. Panolan. Parameterized algorithms for deletion to (r, ℓ) -graphs. In *Proc. 35th FSTTCS*, pages 420–433, 2015. doi:10.4230/LIPIcs.FSTTCS.2015.420.
- [21] D. Lokshitanov, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. Manuscript, 2017.
- [22] J. Nešetřil and P. Ossona de Mendez. *Sparsity: Graphs, Structures, and Algorithms*, volume 28 of *Algorithms and Combinatorics*. Springer, 2012. doi:10.1007/978-3-642-27875-4.
- [23] N. Robertson and P. D. Seymour. Graph minors. V. Excluding a planar graph. *J. Comb. Theory, Ser. B*, 41(1):92–114, 1986. doi:10.1016/0095-8956(86)90030-4.
- [24] N. Robertson and P. D. Seymour. Graph minors. X. Obstructions to tree-decomposition. *J. Comb. Theory, Ser. B*, 52(2):153–190, 1991. doi:10.1016/0095-8956(91)90061-N.
- [25] N. Robertson and P. D. Seymour. Graph minors: XVII. Taming a vortex. *J. Comb. Theory, Ser. B*, 77(1):162–210, 1999. doi:10.1006/jctb.1999.1919.
- [26] A. Schäfer, C. Komusiewicz, H. Moser, and R. Niedermeier. Parameterized computational complexity of finding small-diameter subgraphs. *Optim. Lett.*, 6(5):883–891, 2012. doi:10.1007/s11590-011-0311-5.
- [27] S. Shan. *Homeomorphically Irreducible Spanning Trees, Halin Graphs, and Long Cycles in 3-connected Graphs with Bounded Maximum Degrees*. PhD thesis, Georgia State University, 2015. URL: http://scholarworks.gsu.edu/math_diss/23/.
- [28] S. Thomassé, N. Trotignon, and K. Vuskovic. A polynomial Turing-kernel for weighted independent set in bull-free graphs. In *Proc. 40th WG*, pages 408–419. Springer, 2014. doi:10.1007/978-3-319-12340-0_34.
- [29] W. T. Tutte. *Connectivity in graphs*. Mathematical expositions. University of Toronto Press, 1966.

A Omitted proofs

A.1 Proof of Lemma 4.3

Consider a node $t \in V(T)$, let T_t be the subtree of T rooted in t , and let $G_t = G[\mathcal{X}(T_t)]$. Note that (T_t, \mathcal{X}) is a tree decomposition of G_t with t being a root; with regards to this decomposition, we consider the torso $H_t := \text{TORSO}(G_t, \mathcal{X}(t))$. Note that the connectivity of (T, \mathcal{X}) implies that H_t can be obtained from $G_t[\mathcal{X}(t)]$ by turning the neighborhood $N_{G_t}(C)$ into a clique for every connected component C of $G_t - \mathcal{X}(t)$. Hence every adhesion to a child of t corresponds to a clique in H_t , and we may bound the adhesion degree by bounding the number of distinct cliques in H_t .

To show the bound on the number of cliques, we show a longer, more elaborate process that turns G_t into H_t . Let \mathcal{C} be the family of connected components of $G_t - \mathcal{X}(t)$. For every $C \in \mathcal{C}$ and every vertex $v \in N(C) \subseteq \mathcal{X}(t)$, pick one neighbor $x_{v,C} \in N(v) \cap C$. Let S_C be a minimal connected subgraph of C that contains every vertex $x_{v,C}$ for $v \in N(C)$. Note that by minimality, S_C is a tree with at most $|N(C)| \leq h$ leaves and thus less than h vertices of degree at least 3. Delete from G_t all vertices and edges of $G[C]$ that do not lie in S_C , and suppress every node of degree 2 in S_C that is not of the form $x_{v,C}$ for $v \in N(C)$. By performing this operation for every component $C \in \mathcal{C}$, we obtain a topological minor G_t^1 of G_t with the following property: the graph $G_t[\mathcal{X}(t)]$ remains untouched in G_t^1 , while every $C \in \mathcal{C}$ is turned into its topological minor C^1 with $N(C) = N(C^1)$ and $|C^1| < 2|N(C)| \leq 2h$. Since G does not contain H as a topological minor, neither does G_t^1 .

We now switch to the theory of graphs of bounded expansion; for an introduction and more notation we refer to the textbook [22]. A graph class \mathcal{G} is of *bounded expansion* if there exists a function $\nabla : \mathbb{Z}_{\geq 0} \rightarrow \mathbb{Z}_{\geq 0}$ such that for every radius r , every graph $H \in \mathcal{G}$, and every r -shallow minor H' of H , we have $|E(H')|/|V(H')| \leq \nabla(r)$.

Let G_t^2 be a graph obtained from G_t^1 by contracting every component C^1 for $C \in \mathcal{C}$ into a single vertex x_C . Then, as $|C^1| < 2h$, G_t^2 is an $2h$ -shallow minor of G_t^1 . Let G_t^3 be a graph obtained from G_t^2 by replacing every vertex x_C for $C \in \mathcal{C}$ with a clique K_C of $|N(C)|$ copies of x_C . Finally, note that H_t can be obtained from G_t^3 by contracting every vertex of K_C onto a distinct vertex of $N(C)$, that is, H_t is a 1-shallow minor of G_t^3 .

Since taking an $\mathcal{O}(1)$ -shallow-minor or replacing every vertex with a clique of constant size in a graph from a class of bounded expansion gives a graph also from a class of bounded expansion (but possibly worse function ∇) [22], we have that there exists a graph class \mathcal{G} of bounded expansion (with the function ∇ depending on H and h) such that $H_t \in \mathcal{G}$. Since graphs of bounded expansion have bounded degeneracy and therefore have linearly many cliques [22, Lemma 3.1], we have that H_t contains at most $f'(h, H) \cdot |\mathcal{X}(t)|$ cliques for some constant $f'(h, H)$ depending on h and H . Thus, in Lemma 4.3 we can take $f(h, H) = 1 + f'(h, H)$ to accommodate for the additional adhesion corresponding to the parent of t .