Polynomial bounds for centered colorings on proper minor-closed graph classes*

Michał Pilipczuk† Sebastian Siebertz‡

July 10, 2018

Abstract

For $p \in \mathbb{N}$, a coloring λ of the vertices of a graph G is *p-centered* if for every connected subgraph H of G, either H receives more than p colors under λ or there is a color that appears exactly once in H. Centered colorings play an important role in the theory of sparse graph classes introduced by Nešetřil and Ossona de Mendez [27], as they structurally characterize classes of *bounded expansion* — one of the key sparsity notions in this theory. More precisely, a class of graphs $\mathscr{C}$ has bounded expansion if and only if there is a function $f\colon \mathbb{N} \to \mathbb{N}$ such that every graph $G \in \mathscr{C}$ for every $p \in \mathbb{N}$ admits a p-centered coloring with at most $f(p)$ colors. Unfortunately, known proofs of the existence of such colorings yield large upper bounds on the function f governing the number of colors needed, even for as simple classes as planar graphs.

In this paper, we prove that every K_t-minor-free graph admits a p-centered coloring with $\mathcal{O}(p^{g(t)})$ colors for some function g. In the special case that the graph is embeddable in a fixed surface Σ we show that it admits a p-centered coloring with $\mathcal{O}(p^{19})$ colors, with the degree of the polynomial independent of the genus of Σ. This provides the first polynomial upper bounds on the number of colors needed in p-centered colorings of graphs drawn from proper minor-closed classes, which answers an open problem posed by Dvořák [1].

As an algorithmic application, we use our main result to prove that if $\mathscr{C}$ is a fixed proper minor-closed class of graphs, then given graphs H and G, on p and n vertices, respectively, where $G \in \mathscr{C}$, it can be decided whether H is a subgraph of G in time $2^{\mathcal{O}(p \log p)} \cdot n^{\mathcal{O}(1)}$ and space $n^{\mathcal{O}(1)}$.

*The work of M. Pilipczuk and S. Siebertz is supported by the National Science Centre of Poland via POLONEZ grant agreement UMO-2015/19/P/ST6/03998, which has received funding from the European Union's Horizon 2020 research and innovation programme (Marie Skłodowska-Curie grant agreement No. 665778).

†Institute of Informatics, University of Warsaw, Poland, `michal.pilipczuk@mimuw.edu.pl`.

‡Institute of Informatics, University of Warsaw, Poland, `siebertz@mimuw.edu.pl`.

1 Introduction

Structural graph theory provides a wealth of tools that can be used in the design of efficient algorithms for generally hard graph problems. In particular, the algorithmic properties of classes of graphs of bounded treewidth, of planar graphs, and more generally, of classes which exclude a fixed minor have been studied extensively in the literature. The celebrated structure theory developed by Robertson and Seymour for graphs with excluded minors had an immense influence on the design of efficient algorithms. Nešetřil and Ossona de Mendez introduced the even more general concepts of *bounded expansion* [25] and *nowhere denseness* [26], which offer abstract and robust notions of sparseness in graphs, and which also lead to a rich algorithmic theory. Bounded expansion and nowhere dense graph classes were originally defined by restricting the edge densities of bounded depth minors that may occur in these classes; in particular, every class that excludes a fixed topological minor has bounded expansion. In this work we are going to study *p-centered colorings*, which may be used to give a structural characterization of bounded expansion and nowhere dense classes, and which are particularly useful in the algorithmic context.

Definition 1 ([24]). Let G be a graph, $p \in \mathbb{N}$, and let C be a set of colors. A coloring $\lambda\colon V(G) \to C$ of the vertices of G is called *p-centered* if for every connected subgraph H of G, either H receives more than p colors or there is a color that appears exactly once in H under λ.

Definition 2. For a function $f : \mathbb{N} \to \mathbb{N}$, we say that a graph class $\mathscr{C}$ admits *p-centered colorings* with $f(p)$ colors if for every $p \in \mathbb{N}$, every graph $G \in \mathscr{C}$ admits a p-centered coloring using at most $f(p)$ colors. If for the class $\mathscr{C}$ we can choose f to be a polynomial, say of degree d, then we say that $\mathscr{C}$ admits *polynomial centered colorings* of degree d.

Nešetřil and Ossona de Mendez [25] proved that classes of bounded expansion can be characterized by admitting centered colorings with a bounded number of colors, as explained below.

Theorem 1 ([25]). *A class $\mathscr{C}$ of graphs has bounded expansion if and only if there exists a function $f : \mathbb{N} \to \mathbb{N}$ such $\mathscr{C}$ admits p-centered colorings with $f(p)$ colors.*

A similar characterization is known for nowhere dense classes as well, but this notion will not be directly relevant to the purpose of this work. Note that 1-centered colorings are exactly proper colorings of a graph, thus centered colorings are a generalization of proper colorings. On the other hand, every p-centered coloring of a graph G is also a *treedepth-p coloring* of G, in the sense that the union of every i color classes, $i \leqslant p$, induces a subgraph of G of treedepth at most i; see [24]. Here, the *treedepth* of a graph is the minimum height of a rooted forest whose ancestor-descendant closure contains the graph; this parameter is never smaller than the treewidth. Hence, a p-centered coloring of a graph G can be understood as a decomposition of $V(G)$ into disjoint pieces, so that any subgraph induced by at most p pieces is strongly structured — it has treedepth at most p, so also treewidth at most p.

The inspiration of Nešetřil and Ossona de Mendez for introducing low treedepth colorings in [25] was a long line of research on *low treewidth colorings* in proper minor-closed classes (i.e., minor-closed classes excluding at least one minor). It is a standard observation, underlying the classic Baker's approach, that if in a connected planar graph G we fix a vertex u and we color all the vertices according to the residue of their distance from u modulo $p+1$, then the obtained coloring with $p+1$ colors has the following property: the union of any p color classes induces a graph of treewidth $\mathcal{O}(p)$. As proved by Demaine et al. [6] and by DeVos et al. [9], such colorings with $p+1$ colors can be found for any proper minor-closed class of graphs. Decompositions of this kind, together with similar statements for colorings of edges, are central in the design of approximation and parameterized algorithms in proper minor-closed graph classes, see e.g. [6, 7, 8, 9] and the discussion therein.

Thus, low treedepth colorings offer a somewhat different view compared to low treewidth colorings: we obtain a stronger structure — bounded treedepth instead of bounded treewidth — at the cost of having significantly more colors — some function of p instead of just $p + 1$. While admittedly not that useful for approximation algorithms, low treedepth colorings are a central algorithmic tool in the design of parameterized algorithms in classes of bounded expansion. For instance, as observed in [25], using low treedepth colorings one can give a simple fpt algorithm for testing subgraph containment on classes of bounded expansion: to check whether a graph H on p vertices is a subgraph of a large graph G, we compute a treedepth-p coloring of G, say with $f(p)$ colors, and for every p-tuple of color classes we use dynamic programming to verify whether H is a subgraph of the graph induced by those color classes. A much more involved generalization of this idea led to an fpt algorithm for testing any first-order definable property in any class of bounded expansion, first given by Dvořák, Král' and Thomas [12] using different tools. We remark that proofs of this result using low treedepth colorings [16, 19, 30] crucially use the fact that any p-tuple of color classes induce a graph of bounded treedepth, and not just bounded treewidth.

The running times of algorithms based on p-centered colorings strongly depend on the number $f(p)$ of colors used. Unfortunately, the known approaches to constructing centered colorings produce a very large number of colors, typically exponential in p. As shown in recent experimental works [28], this is actually one of major bottlenecks for applicability of these techniques in practice.

The original proof of Theorem 1 in [25] gives a bound for $f(p)$ that is at least doubly exponential in p for general classes of bounded expansion. Somewhat better bounds for proper minor-closed classes can be established via a connection to yet another family of parameters, namely the *weak coloring numbers*, introduced by Kierstead and Yang [22]. We refrain from giving formal definitions, as they are not directly relevant to our purposes here, but intuitively the weak p-coloring number of a graph G measures reachability properties up to distance p in a linear vertex ordering of the graph G. It was shown by Zhu [34] that the number of colors needed for a p-centered coloring of a graph is bounded by its weak 2^{p-2}-coloring number. The weak r-coloring number of a graph G is bounded by $\mathcal{O}(r^3)$ if G is planar and by $\mathcal{O}(r^{t-1})$ if G excludes K_t as a minor [33]. Combining the two results gives a bound of $\mathcal{O}(2^{3p})$ colors needed for a p-centered coloring on planar graphs and $\mathcal{O}(2^{(t-1)p})$ on graphs which exclude K_t as a minor. To the best of authors' knowledge, so far no bounds polynomial in p were known even for the case of planar graphs.

Motivated by this state of the art, Dvořák [1] asked whether one could obtain a polynomial bound on the number of colors needed for p-centered colorings on proper minor-closed graph classes.

Out results. We answer the question of Dvořák in affirmative by proving the following theorems.

Theorem 2. *Every proper minor-closed class admits polynomial time computable polynomial centered colorings, of some degree depending on the class.*

Theorem 3. *For every surface Σ, the class of graphs embeddable in Σ admits polynomial time computable polynomial centered colorings of degree 19. More precisely, if the Euler genus of Σ is g, then the obtained p-centered coloring uses $\mathcal{O}(g^2p^3 + p^{19})$ colors.*

Observe that in case of surface-embedded graphs we obtain degree independent of the genus, however for general proper minor-closed classes the degree depends on the class.

Our techniques. Our proof proceeds by establishing the result for larger and larger graph classes.

We first focus on graphs of bounded treewidth, where we prove that the class of graphs of treewidth at most k admits polynomial centered colorings of degree k, i.e. with $\mathcal{O}(p^k)$ colors for a p-centered coloring. The key to this result is the combinatorial core of the proof of Grohe et al. [20] that every graph of treewidth at most k has weak p-coloring number $\mathcal{O}(p^k)$.

We next move to the case of planar graphs, where for every planar graph G we construct a p-centered coloring of G using $\mathcal{O}(p^{19})$ colors. The idea is to first prove a structure theorem for planar graphs, which may be of independent interest. To state it, we first need a few definitions

A path P in a graph G is called a *geodesic* if it is a shortest path between its endpoints. A *partition* of a graph G is any family $\mathcal{P}$ of induced subgraphs of G such that every vertex of G is in exactly one of subgraphs from $\mathcal{P}$. For a partition $\mathcal{P}$ of G, we define the *quotient graph* $G/\mathcal{P}$ as follows: it has $\mathcal{P}$ as the vertex set and two parts $X, Y \in \mathcal{P}$ are adjacent in $G/\mathcal{P}$ if and only if there exist $x \in X$ and $y \in Y$ that are adjacent in G. The structure theorem then can be stated as follows.

Theorem 4. *For every planar graph G there exists a partition $\mathcal{P}$ of G such that $\mathcal{P}$ is a family of geodesics in G and $G/\mathcal{P}$ has treewidth at most 8. Moreover, such a partition $\mathcal{P}$ of G together with a tree decomposition of $G/\mathcal{P}$ of width at most 8 can be computed in time $\mathcal{O}(n^2)$.*

The idea of using separators that consist of a constant number of geodesics is not new. A classic result of Lipton and Tarjan [23] states that in every n-vertex planar graph one can find two geodesics whose removal leaves components of size at most $2n/3$. By recursively applying this result, one obtains a decomposition of logarithmic depth along geodesic separators, which has found many algorithmic applications, see e.g. the notion of k-path separable graphs of Abraham and Gavoille [2]. However, there is a subtle difference between this decomposition and the decomposition given by Theorem 4: in Theorem 4 all the paths are geodesics in the *whole* graph G, while in the decomposition obtained as above the paths are ordered as $P_1, \ldots, P_\ell$ so that each P_i is only geodesic in the graph $G - \bigcup_{j<i} P_j$. This difference turns out to be crucial in our proof.

Let us come back to the issue of finding a p-centered coloring of a planar graph G. By applying the layering technique, we may assume that that G has radius bounded by $2p$. Hence, every geodesic in the partition $\mathcal{P}$ given by Theorem 4 has at most $4p + 1$ vertices. By the already established case of graphs of bounded treewidth, the quotient graph $G/\mathcal{P}$ admits a p-centered coloring κ with $\mathcal{O}(p^8)$ colors (this is later blown up to $\mathcal{O}(p^{19})$ by layering). We can now assign every vertex a color consisting of the color under κ of the geodesic that contains it, and its distance from a fixed end of the geodesic. This resolves the planar case.

We next lift the result to graphs embeddable in a fixed surface. Here, the idea is to cut the surface along a short cut-graph that can be decomposed into $\mathcal{O}(g)$ geodesics; a construction of such a cut-graph was given by Erickson and Har-Peled [14]. Then the case of embeddable graphs is generalized to nearly embeddable graphs using a technical construction inspired by the work of Grohe [17]. Finally, we lift the case of nearly embeddable graphs to graphs from a fixed proper minor-closed class using the structure theorem of Robertson and Seymour [32]. Here, we observe that the (already proved) bounded treewidth case can be lifted to a proof that p-centered colorings can be conveniently combined along tree decompositions with small adhesions, that is, where every two adjacent bags intersect only at a constant number of vertices.

Applications. Finally, we show a concrete algorithmic application of our main result. There is one aspect where having a treedepth decomposition of small height is more useful than having a tree decomposition of small width, namely space complexity. Dynamic programming algorithms on tree decompositions typically use space exponential in the width of the decomposition, and there are complexity-theorerical reasons to believe that without significant loss on time complexity, this cannot be avoided. On the other hand, on treedepth decompositions one can design algorithms with polynomial space usage. We invite the reader to the work of Pilipczuk and Wrochna [31] for an in-depth study of this phenomenon. This feature of treedepth can be used to prove the following.

Theorem 5. *Let $\mathscr{C}$ be a proper minor-closed class. Then given graphs H and G, on p and n vertices, respectively, where $G \in \mathscr{C}$, it can be decided whether H is a subgraph of G in time $2^{\mathcal{O}(p \log p)} \cdot n^{\mathcal{O}(1)}$ and space $n^{\mathcal{O}(1)}$.*

The proof of Theorem 5 follows the same strategy as before: having computed a p-centered coloring with $p^{\mathcal{O}(1)}$ colors, we iterate over all the p-tuples of color classes, of which there are $2^{\mathcal{O}(p\log p)}$, and for each p-tuple we use an algorithm that solves the problem on graphs of treedepth at most p. This algorithm can be implemented to work in time $2^{\mathcal{O}(p\log p)}\cdot n^{\mathcal{O}(1)}$ and use polynomial space. We remark that this is not a straightforward dynamic programming, in particular we use the color-coding technique of Alon et al. [3] to ensure the injectivity of the constructed subgraph embedding.

The subgraph containment problem in proper minor-closed classes has a large literature. By applying the same technique on a treewidth-p coloring with $p+1$ colors and using a standard dynamic programming algorithm for graphs of treewidth p one can obtain an algorithm with time and space complexity $2^{\mathcal{O}(p\log p)}\cdot n^{\mathcal{O}(1)}$ working on any proper minor-closed class. This was first observed for planar graphs by Eppstein [13], and the running time in the planar case was subsequently improved by Dorn [11] to $2^{\mathcal{O}(p)}\cdot n$. More generally, the running time can be improved to $2^{\mathcal{O}(p/\log p)}\cdot n^{\mathcal{O}(1)}$ for apex-minor-free classes and connected pattern graphs H, and even to $2^{\mathcal{O}(\sqrt{p}\log^2 p)}\cdot n^{\mathcal{O}(1)}$ under the additional assumption that H has constant maximum degree [15]. All the abovementioned algorithms use space exponential in p. See also [5] for lower bounds under ETH. Thus, Theorem 5 offers a reduction of space complexity to polynomial at the cost of having a moderately worse time complexity than the best known.

2 Lifting constructions

Preliminaries. All graphs in this paper are finite and simple, that is, without loops at vertices or multiple edges connecting the same pair of vertices. They are also undirected unless explicitly stated. We use the notation of Diestel's textbook [10] and refer to it for all undefined notation.

When we say that some object in a graph G is *polynomial-time computable*, or *ptime computable* for brevity, we mean that there is a polynomial-time algorithm that given a graph G computes such an object in polynomial time. When we say that G admits a ptime computable p-centered coloring, we mean that the algorithm computing the coloring takes p on input and runs in time $c\cdot n^c$ for some constant c, independent of p. This definition carries over to classes of graphs: class $\mathscr{C}$ admits ptime computable p-centered colorings if there is an algorithm as above working on every graph from $\mathscr{C}$. Note that in particular, the constant c may depend on $\mathscr{C}$, but may not depend on p given on input.

We now give three constructions that enable us to lift the existence of polynomial centered colorings from simpler to more complicated graph classes. The first one is based on the layering technique, and essentially states that to construct p-centered colorings with $p^{\mathcal{O}(1)}$ colors in a graph class it suffices to focus on connected graphs of radius at most $2p$. The second states that having a partition of the graph into small pieces, we can lift centered colorings from the quotient graph to the original graph. The third allows lifting centered colorings through tree decompositions with small adhesion (maximum size of an intersection of two adjacent bags). All proofs of claims from this section are deferred to Section 8 and replaced by sketches.

2.1 Lifting through layering

We first show that using classic layering one can reduce the problem of finding p-centered colorings to connected graphs of radius at most $2p$.

Lemma 6. *Let $\mathscr{C}$ be a minor-closed class of graphs. Suppose that for some function $f\colon\mathbb{N}\to\mathbb{N}$ the following condition holds: for every $p\in\mathbb{N}$ and connected graph $G\in\mathscr{C}$ of radius at most $2p$, the graph G has a ptime computable p-centered coloring with $f(p)$ colors. Then $\mathscr{C}$ admits a ptime computable p-centered colorings with $(p+1)\cdot f(p)^2$ colors.*

PROOF (SKETCH). Assuming w.l.o.g. that G is connected, run a breadth-first search from any fixed vertex u and partition $V(G)$ into layers according to the distance from u. Consider blocks of $2p$ consecutive layers starting at layers with indices divisible by p; thus every layer is contained in at most two such blocks. For every block we may find a p-centered coloring by removing all layers above it, contracting all layers below it onto u, and applying the assumed result to the obtained graph of radius at most $2p$. It then suffices to overlay the obtained colorings for blocks together with a coloring that colors each vertex with the index of its layer modulo $p+1$. □

Note that in Lemma 6, if $f(p)$ is a polynomial of degree d, then $\mathscr{C}$ admits ptime computable polynomial centered colorings of degree $2d+1$.

2.2 Lifting through partitions

The following lemma will be useful for lifting the existence of centered colorings through partitions and quotient graphs.

Lemma 7. *Let p, q be positive integers. Suppose a graph G has a partition $\mathcal{P}$, computable in ptime for p given on input, so that*

- *$|V(A)| \leqslant q$ for each $A \in \mathcal{P}$, and*
- *the graph $G/\mathcal{P}$ admits a ptime computable p-centered coloring with $f(p)$ colors, for some $f\colon \mathbb{N} \to \mathbb{N}$.*

Then the graph G has a ptime computable p-centered coloring with $q \cdot f(p)$ colors.

PROOF (SKETCH). Having found a p-centered coloring κ of $G/\mathcal{P}$, color every vertex v of G with a pair consisting of the color in κ of the part $A \in \mathcal{P}$ containing v and a number from $\{1, \ldots, q\}$, so that vertices within each part of $\mathcal{P}$ receive different second coordinates. □

We will often use the following combination of Lemma 6 and Lemma 7, where each part of the partition is a geodesic.

Corollary 8. *Suppose a minor-closed class of graphs $\mathscr{C}$ has the following property: for every graph $G \in \mathscr{C}$ there exists a ptime computable partition $\mathcal{P}_G$ of G into geodesics in G so that the class*

$$\mathscr{D} = \{G/\mathcal{P}_G \colon G \in \mathscr{C}\}$$

admits ptime computable p-centered colorings with $f(p)$ colors, for some function $f\colon \mathbb{N} \to \mathbb{N}$. Then $\mathscr{C}$ admits ptime computable p-centered colorings with $(p+1)(4p+1)^2 \cdot f(p)^2$ colors.

2.3 Lifting through tree decompositions

In this paper, it will be convenient to work with rooted tree decompositions. That is, the shape of a tree decomposition will be a *directed tree* T: an acyclic directed graph with one *root* node having out-degree 0 and all other nodes having out-degree 1. This imposes standard parent/child relation in T, where the parent of a non-root node is its unique out-neighbor.

Definition 3. A *tree decomposition* of a graph G is a pair $\mathcal{T} = (T, \beta)$, where T is a directed tree and $\beta\colon V(T) \to 2^{V(G)}$ is a mapping that assigns each node x of T its *bag* $\beta(x) \subseteq V(G)$ so that the following conditions are satisfied:

(T1) For each $u \in V(G)$, the set $\{x \colon u \in \beta(x)\}$ is non-empty and induces a connected subtree of T.

(T2) For every edge $uv \in E(G)$, there is a node $x \in V(T)$ such that $\{u, v\} \subseteq \beta(x)$.

Let $\mathcal{T} = (T, \beta)$ be a tree decomposition of G. The *width* of $\mathcal{T}$ is the maximum bag size minus 1, i.e., $\max_{x \in V(T)} |\beta(x)| - 1$. The *treewidth* of a graph G is the minimum possible width of a tree decomposition of G. For a non-root node x with parent y, we define the *adhesion set* of x as $\alpha(x) = \beta(x) \cap \beta(y)$. If x is the root, then we set $\alpha(x) = \emptyset$ by convention. The *adhesion* of the tree decomposition $\mathcal{T} = (T, \beta)$ is the maximum size of an adhesion set in $\mathcal{T}$, i.e., $\max_{x \in V(T)} |\alpha(x)|$.

The following lemma expresses how centered colorings can be combined along tree decompositions with small adhesion.

Lemma 9. *Let k be a fixed integer and $\mathscr{C}$ be a class of graphs that admits ptime computable polynomial centered colorings of degree d. Suppose a class of graphs $\mathscr{D}$ has the following property: every graph $G \in \mathscr{D}$ admits a ptime computable tree decomposition over $\mathscr{C}$ with adhesion at most k. Then $\mathscr{D}$ has ptime computable polynomial centered colorings of degree $d + k$.*

The proof of Lemma 9 is entirely relegated to Section 8.3, but we remark that the main idea is to reuse the combinatorial core of the proof of Grohe et al. [20] that graphs of treewidth k have weak p-coloring number $\mathcal{O}(p^k)$. For reference, we now state the immediate corollary for graphs of bounded treewidth.

Corollary 10. *For every $k \in \mathbb{N}$, the class of graphs of treewidth at most k has ptime computable polynomial centered colorings of degree k.*

Proof. Clearly, the class $\mathscr{C}$ of graphs with at most $k+1$ vertices has ptime computable polynomial centered colorings of degree 0. As shown by Bodlaender [4], for every $k \in \mathbb{N}$, given a graph G of treewidth at most k we may compute a tree decomposition $\mathcal{T}$ of G of width at most k in linear time. We may further assume that in $\mathcal{T}$ there are no two adjacent nodes with equal bags, as such two nodes can be contracted to one node with the same bag. Then $\mathcal{T}$ has adhesion at most k and we may conclude by Lemma 9. □

3 Planar graphs

In this section we establish the result for planar graphs. We first prove Theorem 4, which we repeat for convenience.

Theorem 4. *For every planar graph G there exists a partition $\mathcal{P}$ of G such that $\mathcal{P}$ is a family of geodesics in G and $G/\mathcal{P}$ has treewidth at most 8. Moreover, such a partition $\mathcal{P}$ of G together with a tree decomposition of $G/\mathcal{P}$ of width at most 8 can be computed in time $\mathcal{O}(n^2)$.*

Proof. We provide a proof of the existential statement and at the end we briefly discuss how it can be turned into a suitable algorithm with quadratic time complexity.

We may assume that G is connected, for otherwise we may apply the claim to every connected component of G separately and take the union of the obtained partitions. Let us fix any plane embedding of G. We also fix any triangulation G^+ of G. That is, G^+ is a plane supergraph of G with $V(G^+) = V(G)$ whose embedding extends that of G, and every face of G^+ is a triangle. Let f_{out} be the outer face of G^+.

In the following, by a *cycle* C is G^+ we mean a simple cycle, i.e., a subgraph of G consisting of a sequence of pairwise different vertices $(v_1, \ldots, v_k)$ and edges $v_1v_2, v_2v_3, \ldots, v_{k-1}v_k, v_kv_1$ connecting them in order. The embedding of such a cycle C splits the plane into two regions, one bounded and one unbounded. By the subgraph *enclosed* by C, denoted $\mathrm{enc}(C)$, we mean the subgraph of G^+ consisting of all vertices and edges embedded into the closure of the bounded region. Note that C itself is a subgraph of $\mathrm{enc}(C)$. Moreover, $\mathrm{enc}(C)$ is a plane graph whose outer face is C and every non-outer face is a triangle.

A cycle C in G^+ shall be called *tight* if C can be partitioned into paths $P_1, P_2, \ldots, P_k$ for some $k \leqslant 6$ so that each P_i ($i \in \{1, \ldots, k\}$) is a geodesic in G; in particular, all edges of P_i belong to $E(G)$. Note that thus a tight cycle can contain at most 6 edges from $E(G^+) - E(G)$, as each such edge must connect endpoints of two cyclically consecutive paths among $P_1, \ldots, P_k$. The crux of the proof lies in the following claim; see Figure 1 for a visualization.

Claim 1. *Let C be a tight cycle, let $P_1, \ldots, P_k$ be a partition of C witnessing its tightness, and let $H = \mathrm{enc}(C)$. Then there exists a partition $\mathcal{Q}$ of H such that:*

(S1) $\mathcal{Q}$ is a family of geodesics in G containing $P_1, \ldots, P_k$; and

(S2) $H/\mathcal{Q}$ admits a rooted tree decomposition of width at most 8 in which $P_1, \ldots, P_k$ belong to the root bag.

Theorem 4 then follows by applying Claim 1 to the outer face f_{out}, regarded as a tight cycle enclosing the whole graph G^+. Indeed, f_{out} is a triangle, so partitioning it into three single-vertex geodesics witnesses that it is tight.

We prove Claim 1 by induction with respect to the number of bounded faces of $H = \mathrm{enc}(C)$. For the base of the induction, if H has one bounded face, then $H = C$ is in fact a triangular, bounded face of G^+ and we may set $\mathcal{Q} = \{P_1, \ldots, P_k\}$.

Suppose then that H has more than one face. Since C is a simple cycle in G^+, it has length at least 3. Therefore, we may assume that $k \geqslant 3$, for otherwise we may arbitrarily split one of geodesics P_i into two or three subpaths so that C is partitioned into three geodesics, apply the reasoning, and at the end merge back the split geodesic in the obtained partition of H. Now, as $3 \leqslant k \leqslant 6$, we may partition C into three paths Q_1, Q_2, Q_3 so that each Q_j is either equal to some P_i, or is equal to the concatenation of some P_i and P_{i+1} together with the edge of C connecting them. Note that paths Q_j are not necessarily geodesics.

For $j \in \{1, 2, 3\}$, let $A_j = V(Q_j)$. Since G is connected, for every vertex v of H there is some path in G connecting v with $V(C)$. If we take the shortest such path, then it is entirely contained in the graph H. For $v \in V(H)$, let $\pi(v)$ be the vertex of $V(C)$ that is the closest to v in G; in case of ties, prefer a vertex belonging to A_j with a smaller index j, and among one A_j break ties arbitrarily. Further, for each $v \in V(H)$ fix $\Pi(v)$ to be any shortest path in G connecting v with $\pi(v)$; note that $\Pi(v)$ is a geodesic in G. For $j \in \{1, 2, 3\}$, let $B_j \subseteq V(H)$ be the set of those vertices v of H for which $\pi(v) \in A_j$; clearly, $\{B_1, B_2, B_3\}$ is a partition of $V(H)$. Observe that $A_j \subseteq B_j$ and for every vertex $v \in B_j$, the path $\Pi(v)$ connects v with A_j, all its vertices belong to B_j, and all its vertices apart from the endpoint $\pi(v)$ do not lie on C.

Thus, we have partitioned the vertices of the disk-embedded graph H into three parts B_1, B_2, B_3 so that C — the boundary of the disk into which H is embedded — is split into three nonempty segments: one contained in B_1, one contained in B_2, and one contained in B_3. All bounded faces of H are triangles. Hence, we may apply Sperner's Lemma to H to infer that there is a bounded face f of H with vertices v_1, v_2, v_3 such that $v_j \in B_j$ for all $j \in \{1, 2, 3\}$.

Let $K_j = \Pi(v_j)$, for $j \in \{1, 2, 3\}$. Observe that paths K_1, K_2, K_3 are geodesics in G and they are pairwise vertex-disjoint, as each K_j is entirely contained in B_j. Furthermore, the only vertex of K_j that lies on C is $\pi(v_j)$, and moreover we have $\pi(v_j) \in A_j$. For $j \in \{1, 2, 3\}$ define C_j as the concatenation of: path K_j, edge $v_j v_{j+1}$ of the face f, path K_{j+1}, and the subpath of C between $\pi(v_j)$ and $\pi(v_{j+1})$ that is disjoint from A_{j+2}; here, indices behave cyclically. From the asserted properties of K_1, K_2, K_3 it follows that C_j is a simple cycle, unless it degenerates to a single edge $v_j v_{j+1}$ traversed there and back in case f shares this edge with C. In the following we shall assume for simplicity that all of C_1, C_2, C_3 are simple cycles; in case one of them degenerates, it should be simply ignored in the analysis. Observe that the disks bounded by C_1, C_2, C_3 are pairwise disjoint and if we denote $H_j = \mathrm{enc}(C_j)$ for $j \in \{1, 2, 3\}$, then graphs H_j and H_{j+1} share only the path K_{j+1}. Moreover, each graph H_j has strictly fewer bounded faces than H, since f is not a face of any H_j.

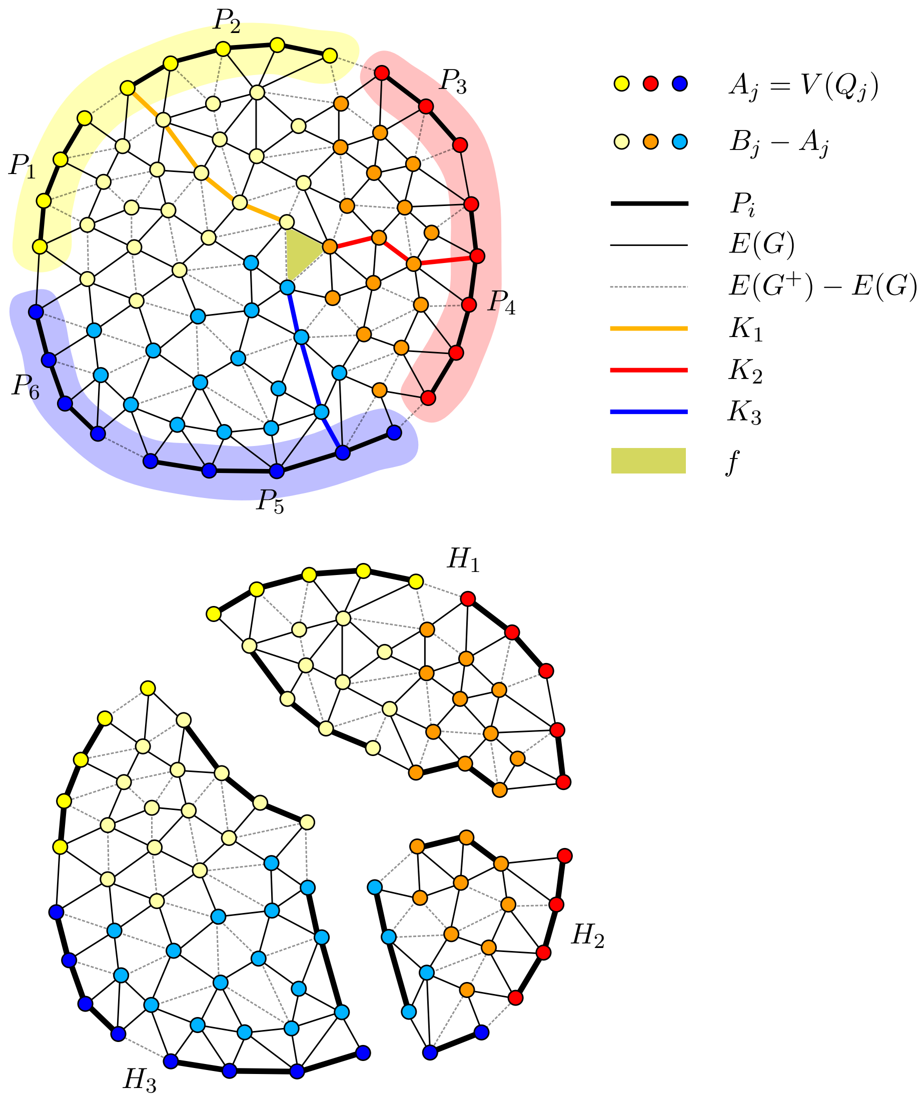

Figure 1: An example situation in Claim 1 and its proof. The top panel depicts the construction of sets A_j and B_j, face f, and paths K_j. Note that Q_1 is the concatenation of P_1 and P_2, Q_2 is the concatenation of P_3 and P_4, and Q_3 is the concatenation of P_5 and P_6. The bottom panel depicts the graphs H_j, with cycles C_j enclosing them, to which the claim is applied inductively. Note that a partition of each cycle C_j into at most 6 geodesics is depicted, which witnesses that C_j is tight.

Denote $L_j = K_j - \pi(v_j)$ for $j \in \{1, 2, 3\}$; in other words, L_j is the path obtained from K_j by removing the endpoint lying on C. Observe that L_j is a geodesic, unless it is empty in case $v_j = \pi(v_j)$ lies on C. We now observe that each cycle C_j, for $j \in \{1, 2, 3\}$, is tight. Indeed, C_j can be partitioned into geodesics L_j and L_{j+1} (provided they are not empty), a subpath of Q_j, and a subgraph Q_{j+1}. By construction, path Q_j can be partitioned into one or two geodesics, so the same holds also for any its subpath; similarly for any subpath of Q_{j+1}. We conclude that C_j can be partitioned into at most six geodesics: L_j and L_{j+1} (provided they are not empty), one or two contained in Q_j, and one or two contained in Q_{j+1}. This witnesses the tightness of C_j.

We now apply the induction hypothesis to each C_j with the partition witnessing its tightness as described above. This yields a suitable partition $\mathcal{Q}_j$ of H_j and tree decomposition $\mathcal{T}_j$ of $H_j/\mathcal{Q}_j$ of width at most 8. Obtain a family $\mathcal{Q}'_j$ from the partition $\mathcal{Q}_j$ as follows: for every path $R \in \mathcal{Q}_j$ that is contained in some geodesic P_i, for $i \in \{1, \dots, k\}$, replace R with P_i. Note here that such paths R have to be in the partition of C_j witnessing the tightness of C_j and there can be at most 4 of them. Now let

$$\mathcal{Q} = \mathcal{Q}'_1 \cup \mathcal{Q}'_2 \cup \mathcal{Q}'_3.$$

It follows readily from the construction that $\mathcal{Q}$ is a partition of H into geodesics that contains all paths P_i for $i \in \{1, \dots, k\}$; this yields condition (S1).

For condition (S2), consider a rooted tree decomposition $\mathcal{T}$ of $H/\mathcal{Q}$ obtained as follows. Construct a root node x with bag $\beta(x) = \{P_1, \dots, P_k\} \cup \{L_1, L_2, L_3\}$. Further, for each $j \in \{1, 2, 3\}$ obtain $\mathcal{T}'_j$ from $\mathcal{T}_j$ by performing the same replacement as before in every bag of $\mathcal{T}_j$: for every $R \in \mathcal{Q}_j$ contained in a geodesic P_i, replace R with P_i. Finally, attach tree decompositions $\mathcal{T}'_j$ for $j \in \{1, 2, 3\}$ below x by making their roots into children of x. It can be easily seen that $\mathcal{T}$ obtained in this manner is indeed a tree decomposition of $H/\mathcal{Q}$. Moreover, we have $|\beta(x)| \leqslant 9$ and each decomposition $\mathcal{T}'_j$ has width at most 8 by the induction hypothesis, so $\mathcal{T}$ has width at most 8 as well.

This finishes the proof of Claim 1 and of the existential statement of Theorem 4. The algorithmic statement follows by turning the inductive proof into a recursive algorithm with time complexity $\mathcal{O}(n^2)$ in a straightforward way. Indeed, it is easy to see that given C as in Claim 1, the cycles C_1, C_2, C_3 can be computed in linear time, and in the recursion we investigate a linear number of recursive calls. ☐

From Theorem 4, Corollary 10, and Corollary 8 we infer the result for planar graphs.

Theorem 11. *The class of planar graphs admits ptime computable polynomial centered colorings of degree* 19.

4 Bounded genus graphs

In this section we lift the result to surface-embedded graphs. By a *surface* we mean a compact, connected 2-dimensional manifold Σ without boundary. An *embedding* of a graph G in Σ maps vertices of G to distinct points in Σ and edges of G to pairwise non-crossing curves on Σ connecting respective endpoints. When we talk about a Σ-embedded graph, we implicitly identify the graph with its embedding in Σ. For a Σ-embedded graph G, every connected component of $\Sigma - G$ is called a *face*. The set of faces of G is denoted by $F(G)$. The embedding is *proper* if every face is homeomorphic to an open disk.

Recall that every surface Σ has its *Euler genus* $g = g(\Sigma)$, which is an invariant for which the following holds: for every properly Σ-embedded connected graph G, we have

$$|V(G)| - |E(G)| + |F(G)| = 2 - 2g.$$

A subgraph K of a properly Σ-embedded graph G is called a *cut-graph* of G if the topological space $\Sigma - K$ is homeomorphic to a disk. The lift of our results from planar graphs to Σ-embeddable graphs is based on the following lemma, which was essentially proved by Erickson and Har-Peled in [14, Lemma 5.7]. Since we need a slightly different phrasing, which puts focus on some properties of the construction that are implicit in [14], we provide our proof in Section 9.

Lemma 12. *Let G be a connected graph properly embedded in a surface Σ of Euler genus $g = g(\Sigma)$. Then G contains a ptime computable cut-graph K such that K can be partitioned into at most $4g$ geodesics in G.*

Lemma 12 can be now used to lift Theorem 4 to graphs embeddable into a surface of fixed genus. The proof is a technical lift of the proof of Theorem 4, precomposed with cutting the surface using the cut graph provided by Lemma 12. We provide here a sketch that contains all the necessary ideas, while the full proof is in Section 9.

Theorem 13. *Let Σ be a surface of Euler genus g. Then for every graph G that can be embedded into Σ, there is a ptime computable partition $\mathcal{P}$ of G and a subset $\mathcal{Q} \subseteq \mathcal{P}$ with $|\mathcal{Q}| \leqslant 16g$ such that $\mathcal{P}$ is a family of geodesics in G and $(G/\mathcal{P}) - \mathcal{Q}$ has treewidth at most 8.*

Proof (Sketch). By Lemma 12, G has a ptime computable cut-graph K which admits a partition $\mathcal{Q}$ into at most $4g$ geodesics in G. We may cut Σ along the cut-graph K, thus turning G into a disc-embedded graph $\widehat{G}$, obtained by duplicating every edge of K, copying every vertex of K the number of times equal to its degree in K, and "opening" the graph in the expected way; see Figure 2. There is a natural projection π from $\widehat{G}$ to G that maps every vertex and edge of $\widehat{G}$ to its origin in G. Since K could be partitioned into at most $4g$ geodesics, it is not hard to see that the boundary of $\widehat{G}$ is a simple cycle that can be partitioned into a set $\mathcal{R}$ of at most $16g$ paths that map in π to subpaths of paths from $\mathcal{Q}$. Now apply a reasoning along the lines of the proof of Claim 1 in Theorem 4 to $\widehat{G}$, where we redefine the notion of a tight cycle: a cycle is now tight if it can be partitioned into an arbitrary number of paths, out of which all but at most 6 are contained in a path from $\mathcal{R}$, and the remaining ones are geodesics in G. It is not hard to see that this notion of tightness can be pushed through the inductive proof of Claim 1, yielding a decomposition of $\widehat{G}$ into $\mathcal{R}$ plus a family of geodesics in G. It now remains to set $\mathcal{P}$ to be those geodesics plus $\mathcal{Q}$. ☐

We note that to prove Theorem 13, one cannot just remove the cut-graph K, apply the planar case (Theorem 4), and take the union of the resulting partition and the partition of K into $\mathcal{O}(g)$ geodesics. This is because the obtained geodesics would be geodesics in $G - V(K)$, and not in G.

From Theorem 13 we may infer the result for graphs embedded into a fixed surface.

Theorem 3. *For every surface Σ, the class of graphs embeddable in Σ admits polynomial time computable polynomial centered colorings of degree 19. More precisely, if the Euler genus of Σ is g, then the obtained p-centered coloring uses $\mathcal{O}(g^2p^3 + p^{19})$ colors.*

Proof. By Lemma 6, it suffices to prove that every connected graph G embeddable in Σ of radius at most $2p$ admits a p-centered coloring with $\mathcal{O}(gp + p^9)$ colors. Apply Theorem 13 to G, yielding a partition $\mathcal{P}$ of G into geodesics and $\mathcal{Q} \subseteq \mathcal{P}$ with $|\mathcal{Q}| \leqslant 16g$ such that $(G/\mathcal{P}) - \mathcal{Q}$ has treewidth at most 8. Since a geodesic in a graph of radius at most $2p$ contains at most $4p+1$ vertices, we have that each geodesic in $\mathcal{P}$ involves at most $4p+1$ vertices and in particular the total number of vertices involved in geodesics from $\mathcal{Q}$ is at most $16g(4p+1) = \mathcal{O}(gp)$. By Corollary 10, the graph $(G/\mathcal{P}) - \mathcal{Q}$ admits a ptime computable p-centered coloring with $\mathcal{O}(p^8)$ colors. Applying Lemma 7 to the graph $G' = G - \bigcup_{Q \in \mathcal{Q}} V(Q)$ and its partition $\mathcal{P} - \mathcal{Q}$, we obtain a ptime computable p-centered coloring of G' with $\mathcal{O}(p^{19})$ colors. It now remains to extend this coloring to G by assigning each of the $\mathcal{O}(gp)$ vertices of $\bigcup_{Q \in \mathcal{Q}} V(Q)$ a fresh, individual color. ☐

5 Nearly embeddable graphs

We now move to nearly embeddable graphs. Roughly saying, a graph G is (a, q, w, g)-*nearly embeddable* if it is embeddable into a surface of Euler genus g modulo at most a *apices* and at most q *vortices* of width at most w each. This is formalized next, following the definitional layer of Grohe [17].

For two graphs G and H, by $G \cup H$ we denote the graph with vertex set $V(G) \cup V(H)$ and edge set $E(G) \cup E(H)$; note that this makes sense also when G and H share vertices or edges. A *path decomposition* is a tree decomposition where the underlying tree is a path. A *boundaried surface* Σ is a 2-dimensional compact manifold with boundary homeomorphic to q copies of S^1, for some $q \in \mathbb{N}$, which shall be called the *boundary cycles*. The *Euler genus* of such a boundaried surface Σ is the Euler genus of the closed surface obtained from Σ by gluing a disc along each boundary cycle.

Definition 4. A graph G is (a, q, w, g)-*nearly embeddable* if there is a vertex subset A with $|A| \leqslant a$, (possibly empty) subgraphs $G_0, \ldots, G_q$ of G, and a boundaried surface Σ of genus g with q boundary cycles $C^1, \ldots, C^q$ such that the following conditions hold:

- We have $G - A = G_0 \cup G_1 \cup \ldots \cup G_q$.
- Graphs $G_0, G_1, \ldots, G_q$ have disjoint edge sets, and graphs $G_1, \ldots, G_q$ are pairwise disjoint.
- Graph G_0 has an embedding into Σ such that all vertices of $V(G_0) \cap V(G_i)$ are embedded on C^i, for $1 \leqslant i \leqslant q$.
- For $1 \leqslant i \leqslant q$, let $m_i = |V(G_0) \cap V(G_i)|$ and let $u^i_1, u^i_2, \ldots, u^i_{m_i}$ be the vertices of $V(G_0) \cap V(G_i)$ in the order of appearance on the cycle C^i. Then G_i has a path decomposition $\mathcal{T}_i = (T_i, \beta_i)$ of width at most w, where T_i is a path $(x^i_1, \ldots, x^i_{m_i})$ and $u^i_j \in \beta(x^i_j)$ for all $1 \leqslant j \leqslant m_i$.

The vertices of A in the above definition are called *apices*, the subgraphs $G_1, \ldots, G_q$ are called *vortices*, and G_0 is called the *skeleton graph*. We now lift the results to nearly embeddable graphs.

Theorem 14. *For any fixed $a, q, w, g \in \mathbb{N}$, the class of (a, q, w, g)-nearly embeddable graphs admits polynomial centered colorings of degree $\mathcal{O}(gq \cdot w^q)$. Moreover, these centered colorings are ptime computable assuming the input graph is given together with a decomposition as in Definition 4.*

Proof. Fix p for which we need to find a p-centered coloring of the input graph G. Since q is a fixed constant, we may assume that $p \geqslant 4q$, for otherwise we replace p with $\max(p, 4q)$ in the reasoning. As in Definition 4, let A be the apex set, G_0 be the skeleton graph, $G_1, \ldots, G_q$ be vortices, and Σ be the target surface of the near-embedding. Note that all of the above is given on input. Let $G' = G - A = G_0 \cup G_1 \cup \ldots \cup G_q$.

Assign to each apex $w \in A$ an individual color c_w that will not be used for any other vertex. Thus, by using at most a additional colors we may focus on finding a p-centered coloring of G'.

We first argue that without loss of generality we may assume that G' is connected and has radius bounded linearly in p. In previous sections we used Lemma 6 for such purposes, but this time there is a technical issue: the class of (a, q, w, g)-nearly embeddable graphs is not necessarily minor-closed. However, the proof of Lemma 6 can be amended, as explained next.

Claim 1. *Without loss of generality we may assume that G' is connected and has radius at most $2p + 2q$.*

Proof. As we shall amend the proof of the Lemma 6, we assume that the reader is familiar with it.

Clearly, we may assume that G' is connected, because we may treat every connected component separately and take the union of the obtained colorings. Also, we argue that we may assume that each vortex G_i, for $1 \leqslant i \leqslant q$, is connected and has radius at most 1. This can be easily achieved by adding edges to G_i so that one of its vertices, say v_i, becomes universal, i.e., is adjacent to all the other vertices of G_i. Note that v_i can be added to every bag of the assumed path decomposition $\mathcal{T}_i$ of G_i, so the width of every vortex grows to $w + 1$ at most. Thus, after modification the obtained graph G' is $(0, q, w + 1, g)$-nearly embeddable, but whether the width of vortices is w or $w + 1$ has no impact on the claimed asymptotic bound of $\mathcal{O}(gq \cdot w^q)$ on the degree of polynomial centered colorings.

We now follow the steps of the proof of Lemma 6. Construct the same layering structure: pick any vertex u and partition the vertex set of G' into layers $L_0, L_1, L_2, \ldots$ according to distances from u; let k be such that L_k is the largest nonempty layer. Since we assumed that each vortex G_i has radius at most 1, it is entirely contained in 3 consecutive layers. Next, for every $j \in \{0, 1, \ldots, k\}$ divisible by p we considered the graph G'_j obtained from G' by contracting all layers L_t for $t < j$ onto u and removing all layers L_t for $t \geqslant j + 2p$. Let $I_j = \{j, j+1, \ldots, j+2p-1\}$ be the interval of indices of layers that are preserved by this construction.

We now extend the interval I_j to achieve the following condition: every vortex G_i is either entirely contained or entirely disjoint with $\bigcup_{t \in I_j} L_t$. This can be achieved by adding one or two indices to I_j, either from the lower or the higher end, as long as there exists a vortex that intersects some layer L_t with $t \in I_j$ and some other layer $L_{t'}$ with $t' \notin I_j$. Note that here we use the assumption that every vortex is contained in at most 3 consecutive layers. Since there are q vortices in total, after this operation the interval I_j consists of at most $2p + 2q$ consecutive layers and

$$I_j \subseteq \{j - 2q, \ldots, j + 2p - 1 + 2q\} \subseteq \{j - p/2, \ldots, j + 2p - 1 + p/2\},$$

where the last containment follows from the assumption $p \geqslant 4q$. Observe that this means that every index t is contained in at most 3 intervals I_j, for $j \in \{0, 1, \ldots, k\}$ divisible by p.

After this modification we proceed as in the proof of Lemma 6. Namely, let G'_j be the graph obtained from G' by contracting all layers below the lower end of I_j onto u and removing all layers above the higher end of I_j. Observe that G'_j is still $(0, q, w+1, g)$-nearly embeddable, as every vortex either got entirely contracted onto u, or got entirely removed, or is entirely preserved intact. Moreover, as in Lemma 6 we have that G'_j is connected and has radius at most $2p + 2q$.

We may apply the assumption to compute a p-centered coloring λ_j of G'_j, for each relevant j, and superimpose the colorings λ_j just as in Lemma 6. Note that now, in the obtained coloring λ the color of each vertex v is a 4-tuple instead of a 3-tuple, because it consists of the index of v's layer modulo $p+1$ and the colors of v under λ_j for those indices j for which $v \in I_j$, and there are at most three such js. Thus, if each λ_j uses $p^{\mathcal{O}(gq \cdot w^q)}$ colors, then λ uses $(p+1) \cdot p^{3 \cdot \mathcal{O}(gq \cdot w^q)} = p^{\mathcal{O}(gq \cdot w^q)}$ colors. It is straightforward to see that the remainder of the reasoning of the proof of Lemma 6 goes through without changes, yielding that λ is p-centered. ⌟

By Claim 1, from now on we assume that G' is connected and has radius at most $2p + 2q$. We may also assume that each vortex G_i for $1 \leqslant i \leqslant q$ is non-empty, otherwise we apply the reasoning for smaller q. Note that by the connectedness this implies $m_i \geqslant 1$. We now follow the construction of Proposition 3.8 of Grohe [17], who showed that nearly embeddable graphs without apices have bounded local treewidth.

Construct a graph $\widehat{G}$ from G_0 as follows. For each $i \in \{1, \ldots, q\}$, introduce a new vertex z^i and add edges: $z^i u^i_j$ for all $1 \leqslant j \leqslant m_i$, and $u^i_j u^i_{j+1}$ for all $1 \leqslant j \leqslant m_i$, where $u^i_{m_i+1} = u^i_1$. Let $\widehat{G}_i$ be the subgraph of $\widehat{G}$ consisting of vertices $\{z^i\} \cup \{u^i_j \colon 1 \leqslant j \leqslant m_i\}$ and edges added above; note that $\widehat{G}_i$ has diameter 2. Let $\widehat{\Sigma}$ be the closed surface of Euler genus g obtained from Σ by gluing a disk D^i along the boundary cycle C^i, for each $1 \leqslant i \leqslant q$. Then $\widehat{G}$ is embeddable into $\widehat{\Sigma}$, because each subgraph $\widehat{G}_i$ can be embdedded into D^i. Also, $\widehat{G}$ has not much larger diameter than G'.

Claim 2. *The graph $\widehat{G}$ is connected and has diameter at most $8(p+q)+2$.*

Proof. Take any two vertices a and b of $\widehat{G}$, and suppose for a moment that $a, b \in V(G_0)$. Then a and b can be connected by a path P in G' of length at most $4(p+q)$, because the radius of G' is at most $2(p+q)$. Observe that every maximal infix of P that traverses the edges of some vortex G_i, where $1 \leqslant i \leqslant q$, can be replaced by a path of length 2 in H through the vertex z^i. By performing such replacement for every such infix we obtain a path P' in H with the same endpoints and length at most $8(p+q)$. To resolve the case when a or b is among the new vertices $z^1, \ldots, z^q$, it suffices to observe that each of them is adjacent to some vertex of G_0. ⌟

Now that $\widehat{G}$ is embeddable into $\widehat{\Sigma}$, we may apply Theorem 13 to construct a partition $\widehat{\mathcal{P}}$ of $\widehat{G}$ such that $\widehat{\mathcal{P}}$ is a family of geodesics in $\widehat{G}$ and $\widehat{G}/\widehat{\mathcal{P}}$ has treewidth $\mathcal{O}(g)$. Note that since $\widehat{G}$ has diameter at most $8(p+q)+2$ by Claim 2, each geodesic in $\widehat{\mathcal{P}}$ has at most $8(p+q)+3$ vertices.

We now observe that because graphs $\widehat{G}_i$ have diameter at most 2, geodesics in $\widehat{P}$ have only small interaction with them.

Claim 3. *For each $1 \leqslant i \leqslant q$, every path $P \in \widehat{\mathcal{P}}$ contains at most 3 vertices of $\widehat{G}_i$.*

Proof. If P contained more than 3 vertices in $\widehat{G}_i$, then two of them would be at distance more than 2 on P, but $\widehat{G}_i$ has diameter at most 2. This would contradict the assumption that P is a geodesic in $\widehat{G}$. ⌟

Construct a partition $\mathcal{P}$ of G' into paths as follows. First, each vertex $v \in V(G_1) \cup \ldots \cup V(G_q)$, i.e. participating in any vortex, gets assigned to a single-vertex path consisting only of v. The remaining vertices, those of $V(G_0) - (V(G_1) \cup \ldots \cup V(G_q))$, are partitioned into inclusion-wise maximal paths contained in paths in $\widehat{P}$. In other words, for every path $P \in \widehat{\mathcal{P}}$ we remove all vertices of $V(\widehat{G}_1) \cup \ldots \cup V(\widehat{G}_q)$, thus splitting P into several paths, and put all those paths into $\mathcal{P}$. Note that since each part of $\widehat{\mathcal{P}}$ contains at most $8(p+q)+3$ vertices, the same holds also for $\mathcal{P}$, even though paths in $\mathcal{P}$ are not necessarily geodesics in G'.

For the next, crucial step we shall need the following technical lemma of Grohe [17], which enables gluing tree decompositions along common interfaces.

Lemma 15 (Lemma 2.2 of [17]). *Let G, H be graphs and let (T, β) be a path decomposition of H of width at most k. Assume that T is a path $(x_1, \ldots, x_m)$ for some $m \in \mathbb{N}$. Let $v_1, \ldots, v_m$ be a path in G such that $v_i \in \beta(x_i)$ for $1 \leqslant i \leqslant m$ and $V(G) \cap V(H) = \{v_1, \ldots, v_m\}$. Then $\mathrm{tw}(G \cup H) \leqslant (\mathrm{tw}(G)+1)(k+1)-1$.*

We now claim the following.

Claim 4. *The graph $G'/\mathcal{P}$ has treewidth $\mathcal{O}(gq \cdot w^q)$.*

Proof. Let $\widehat{\mathcal{P}}'$ be a partition of $\widehat{G}$ obtained as follows. Examine every path $P \in \widehat{\mathcal{P}}$ and partition it into subpaths: a single-vertex path for each $v \in V(\widehat{G}_1) \cup \ldots \cup V(\widehat{G}_q)$ traversed by P, and maximal infixes of P consisting of vertices not in $V(\widehat{G}_1) \cup \ldots \cup V(\widehat{G}_q)$. Add all the obtained subpaths to $\widehat{\mathcal{P}}'$, thus eventually obtaining a partition of $\widehat{G}$. Note that by Claim 3, we add at most $6q+1$ subpaths of P for each $P \in \widehat{\mathcal{P}}$.

As $\widehat{G}/\widehat{\mathcal{P}}$ has treewidth $\mathcal{O}(g)$, it is easy to see that $\widehat{G}/\widehat{\mathcal{P}}'$ has treewidth $\mathcal{O}(gq)$. Indeed, we may take a tree decomposition of $\widehat{G}/\widehat{\mathcal{P}}$ of width $\mathcal{O}(g)$ and replace every path $P \in \widehat{\mathcal{P}}$ with all its subpaths added to $\widehat{\mathcal{P}}'$ in every bag, thus obtaining a tree decomposition of $\widehat{G}/\widehat{\mathcal{P}}'$ of width $\mathcal{O}(gq)$.

Note that since each subgraph $\widehat{G}_i$ of $\widehat{G}$ contains the cycle $(u^i_1, \ldots, u^i_{m_i})$, in $\widehat{G}/\widehat{\mathcal{P}}'$ the single-vertex paths consisting of vertices $u^i_1, \ldots, u^i_{m_i}$ form a cycle in the same way.

Now it remains to observe that $G'/\mathcal{P}$ is a subgraph of a graph that can be obtained from $\widehat{G}/\widehat{\mathcal{P}'}$ by iteratively adding vortices $G_1, \ldots, G_q$, with path decompositions $\mathcal{T}_1, \ldots, \mathcal{T}_q$ of width at most w (here, we implicitly identify vertices of vortices with single-vertex paths consisting of them). Noting that the prerequisites of Lemma 15 are satisfied, every such addition increases the treewidth from the current value, say t, to at most $(t+1)(w+1)-1$. Since the treewidth of $\widehat{G}/\widehat{\mathcal{P}'}$ is $\mathcal{O}(gq)$, it follows that the treewidth of $G'/\mathcal{P}$ is $\mathcal{O}(gq \cdot w^q)$. ⌟

Since $G'/\mathcal{P}$ has treewidth $\mathcal{O}(gq \cdot w^q)$, by Corollary 10 it admits a ptime computable p-centered coloring with $p^{\mathcal{O}(qg \cdot w^q)}$ colors. As each part of $\mathcal{P}$ has at most $8(p+q)+3$ vertices, we may conclude by Lemma 7. □

6 Proper minor-closed classes

Our main result now follows easily by combining the structure theorem of Robertson and Seymour with the already prepared tools.

Theorem 16 (Robertson and Seymour [32]). *For every $t \in \mathbb{N}$ there exist $a, q, w, g, k \in \mathbb{N}$ such that every graph G excluding K_t as a minor admits a tree decomposition of adhesion at most k over the class of (a, q, w, g)-nearly embeddable graphs.*

Furthermore, it is known that a tree decomposition as stated in Theorem 16, together with decompositions of torsos witnessing their (a, q, w, g)-near embeddability, can be computed in polynomial time, see [6, 18, 21]. Now Theorem 2 is an immediate consequence of this result combined with Lemma 9 and Theorem 14, as every proper minor-closed class excludes some clique K_t as a minor.

7 Subgraph Isomorphism for graphs of bounded treedepth

In this section we prove Theorem 5, but the main technical contribution is the proof of the following lemma.

Lemma 17. *Suppose we are given a graph H on p vertices and a graph G on n, together with a treedepth decomposition of G of depth d. Then it can be decided whether H is a subgraph of G in time $2^{\mathcal{O}((p+d)\log p)} \cdot n^{\mathcal{O}(1)}$ and space $n^{\mathcal{O}(1)}$.*

We then apply the following connection of p-centered colorings with low-treedepth colorings.

Definition 5. Let F be a rooted forest, i.e., a graph whose connected components are rooted trees. The *closure* of F, denoted $\mathrm{clos}(F)$ has as its vertex set the set $V(F)$ and it contains every edge uv such that u, v are vertices of a tree T of F and $u \leqslant_T v$. The *height* of a tree T is the maximal number of vertices on a root-leaf path of T. The *treedepth* of a graph G is the minimum height of a forest F such that $G \subseteq \mathrm{clos}(F)$. Such a forest F is called a *treedepth decomposition* of G.

Proposition 18 (Nešetřil and Ossona de Mendez [24]). *Every p-centered coloring $\lambda\colon V(G) \to C$ of a graph G is also a treedepth-p coloring of G in the following sense: for any color subset $X \subseteq C$ with $|X| \leqslant p$, the graph $G[\lambda^{-1}(X)]$ has treedepth at most $|X|$. Furthermore, a treedepth decomposition of $G[\lambda^{-1}(X)]$ of depth at most $|X|$ can be computed in linear time.*

Lemma 17 combined with the above can be now used to give a space-efficient fixed-parameter algorithm for Subgraph Isomorphism on proper minor-closed classes, as explained in Theorem 5.

Theorem 5. *Let $\mathscr{C}$ be a proper minor-closed class. Then given graphs H and G, on p and n vertices, respectively, where $G \in \mathscr{C}$, it can be decided whether H is a subgraph of G in time $2^{\mathcal{O}(p\log p)} \cdot n^{\mathcal{O}(1)}$ and space $n^{\mathcal{O}(1)}$.*

Proof. By Theorem 2, in polynomial time we can compute a p-centered coloring λ of G that uses at most $c \cdot p^c$ colors, where c is a constant depending only on $\mathscr{C}$. Iterate through all color subsets of consisting p colors and for each such subset X consider the graph $G_X = G[\lambda^{-1}(X)]$. Observe that since H has p vertices, H is subgraph of G if and only if H is a subgraph of G_X for any such color subset X. By Proposition 18, G_X has treedepth at most p and a treedepth decomposition of G_X of depth at most p can be computed in linear time. Hence, we may apply Lemma 17 to verify whether H is a subgraph of G_X in time $2^{\mathcal{O}(p\log p)} \cdot n^{\mathcal{O}(1)}$ and space $n^{\mathcal{O}(1)}$. Since there are $(c \cdot p^c)^p = 2^{\mathcal{O}(p\log p)}$ color subsets X to consider, and for each we apply an algorithm with time complexity $2^{\mathcal{O}(p\log p)} \cdot n^{\mathcal{O}(1)}$ and space complexity $n^{\mathcal{O}(1)}$, the claimed complexity bounds follow. □

In the remainder of this section we give a polynomial-space fixed-parameter algorithm for the Subgraph Isomorphism problem on graphs of bounded treedepth, i.e. we prove Lemma 17. Recall that we are given graphs H and G, where H has p vertices and G has n vertices, and moreover we are given a treedepth decomposition F of G of depth at most d. The goal is to check whether H is a subgraph of G; that is, whether there exists a *subgraph embedding* from H to G, which is an injective mapping from $V(H)$ to $V(G)$ such that $uv \in E(H)$ entails $\eta(u)\eta(v) \in E(G)$.

We first use the color coding technique of Alon et al. [3] to reduce the problem to the colored variant, where in addition vertices of G are labeled with vertices of H and the sought subgraph embedding has to respect these labels. More precisely, suppose we are given a mapping $\alpha\colon V(G) \to V(H)$. We say that a subgraph embedding η from H to G is *compliant* with α if $\alpha(\eta(u)) = u$ for each $u \in V(H)$. The following lemma encapsulates the application of color coding to our problem.

Lemma 19. *There exists a family $\mathcal{F}$ consisting of $2^{\mathcal{O}(p \log p)} \cdot \log n$ functions from $V(G)$ to $V(H)$ so that the following condition holds: for each injective function $\eta\colon V(H) \to V(G)$ there exists at least one function $\alpha \in \mathcal{F}$ such that $\alpha(\eta(u)) = u$ for each $u \in V(H)$. Moreover, $\mathcal{F}$ can be enumerated with polynomial delay and using polynomial working space.*

Proof. For positive integers $p \leqslant q$, a family $\mathcal{S}$ of functions from $V(G)$ to $\{1, \ldots, q\}$ is called *p-perfect* if for every subset $W \subseteq V(G)$ of size p there exists a function $f \in \mathcal{S}$ that is injective on W. Alon et al. [3] proved that for $q = p^2$ there exists a p-perfect family $\mathcal{S}$ of size $p^{\mathcal{O}(1)} \log n$ that can be computed in polynomial time. Define the family $\mathcal{F}$ as follows:

$$\mathcal{F} = \{g \circ f \;:\; f \in \mathcal{S} \text{ and } g \in \{1, \ldots, p^2\}^{V(H)}\}.$$

In other words, $\mathcal{F}$ consists of all functions constructed by composing a function from $\mathcal{S}$ with any function from $\{1, \ldots, p^2\}$ to the vertex set of H. Note that $|\mathcal{S}| = p^{\mathcal{O}(1)} \cdot \log n$ and there are $p^{2p} = 2^{\mathcal{O}(p \log p)}$ functions from $\{1, \ldots, p^2\}$ to $V(H)$, so indeed $|\mathcal{F}| \leqslant 2^{\mathcal{O}(p \log p)} \cdot \log n$. Also, clearly $\mathcal{F}$ can be enumerated with polynomial delay and using polynomial working space.

Finally, we verify that $\mathcal{F}$ satisfies the promised condition. Consider $W = \eta(V(H))$; by the properties of $\mathcal{S}$, there exists $f \in \mathcal{S}$ that is injective on W. Since f is injective and η is injective on the image of f, we may construct a function $g\colon \{1, \ldots, p^2\} \to V(H)$ as follows: if x belongs to the image of $f \circ \eta$ then $g(x)$ is the unique vertex u of H such that $f(\eta(u)) = x$, and otherwise $g(x)$ is set arbitrarily. Let $\alpha = g \circ f$. Clearly α belongs to $\mathcal{F}$ and we have $\alpha(\eta(u)) = u$ for each $u \in V(H)$. $\square$

By applying Lemma 19, to prove Lemma 17 we may focus on the variant where we are additionally given a mapping $\alpha\colon V(G) \to V(H)$ and we seek a subgraph embedding that is compliant with α. Indeed, if we give an algorithm with the promised time and space complexity for this variant, then we may apply it for every function α from the family $\mathcal{F}$ enumerated using Lemma 19. This adds a multiplicative factor of $2^{\mathcal{O}(p \log p)} \cdot n^{\mathcal{O}(1)}$ to the time complexity and an additive factor of $n^{\mathcal{O}(1)}$ to the space complexity, which is fine for the claimed complexity bounds.

Before we proceed to the algorithm, we introduce some notation. Recall that F is the given treedepth decomposition G; that is, F is a rooted forest of depth at most d on the same vertex as G such that every edge of G connects a vertex with its ancestor in F. For $u \in V(G)$, we introduce the following notation:

- $\mathsf{Chld}(u)$ is the set of children of u in F;
- $\mathsf{Tail}(u)$ is the set of all strict ancestors of u in F (i.e., excluding u itself);
- G_u is the subgraph of G induced by the ancestors and descendants of u, including u itself.

A pair (X, D) of disjoint subsets of vertices of H is a *chunk* if

- X is either empty or it induces a connected subgraph of H; and
- in H there is no edge with one endpoint in X and second in $V(H) - (X \cup D)$.

Note that the second condition is equivalent to saying that that D is contained in $N_H(X)$. A *subproblem* is a quadruple (u, X, D, γ), where

- u is a vertex of G,
- (X, D) is a chunk, and
- γ is an injective function from D to $\mathsf{Tail}(u)$.

Note that the number of different subproblems is at most $3^p \cdot p^d \cdot n = 2^{\mathcal{O}(p+d \log p)} \cdot n$. The *value* of the subproblem (u, X, D, γ), denoted $\mathsf{Val}(u, X, D, \gamma)$, is the boolean value of the following assertion:

There exists a subgraph embedding η from the graph $H[X \cup D]$ to the graph G_u such that $\eta(X) \cap \mathsf{Tail}(u) = \emptyset$ and η restricted to D is equal to γ.

An embedding η satisfying the above will be called a *solution* to the subproblem (u, X, D, γ). Our algorithm will compute the values of subproblems in a recursive manner using the formula presented in the following lemma. Here, $\gamma[w \to u]$ denotes γ extended by mapping w to u, whereas for $Y \subseteq V(H)$ by $\mathsf{CC}(Y)$ we denote the family of vertex sets of the connected components of $H[Y]$.

Lemma 20. *Suppose (u, X, D, γ) is a subproblem. Then the following assertions hold:*

(i) If u is a leaf of F, then $\mathsf{Val}(u, X, D, \gamma)$ is true if and only if either $X = \emptyset$ and γ is a subgraph embedding from $H[D]$ to G_u, or $X = \{w\}$ with $w = \alpha(u)$ and $\gamma[w \to u]$ is a subgraph embedding from $H[D \cup \{w\}]$ to G_u.

(ii) If u is not a leaf of F and $\alpha^{-1}(u) \notin X$, then

$$\mathsf{Val}(u, X, D, \gamma) = \bigvee_{v \in \mathsf{Chld}(u)} \mathsf{Val}(v, X, D, \gamma).$$

(iii) If u is not a leaf of F and $w = \alpha^{-1}(u) \in X$, then

$$\mathsf{Val}(u, X, D, \gamma) = \bigvee_{v \in \mathsf{Chld}(u)} \mathsf{Val}(v, X, D, \gamma) \vee \bigwedge_{Z \in \mathsf{CC}(X - \{w\})} \bigvee_{v \in \mathsf{Chld}(u)} \mathsf{Val}(v, Z, D \cup \{w\}, \gamma[w \to u]).$$

Proof. Assertion (i) is straightforward.

For assertion (ii), observe that a solution η to the subproblem (u, X, D, γ) cannot map any vertex of X to u, because only the vertex $w = \alpha^{-1}(u)$ can be mapped to u, and w does not belong to X by assumption. Moreover, since $H[X]$ is connected (due to (X, D) being a chunk), $\eta(X)$ has to be entirely contained in one subtree of F rooted at a child of u. It follows that every solution to the subproblem (u, X, D, γ) is a solution to one of the subproblems $\mathsf{Val}(v, X, D, \gamma)$ for v ranging over the children of u in F, and conversely every solution to any of these subproblems is trivially also a solution to (u, X, D, γ). The formula follows.

For assertion (iii), observe that every solution η to the subproblem (u, X, D, γ) either maps $w = \alpha^{-1}(u)$ to u, or does not map any vertex to u. In the latter case, the same reasoning as for assertion (ii) yields that η

is also a solution to one of subproblems (v, X, D, γ) for v ranging over the children of u; this corresponds to the first part of the formula. Consider now the former case, that is, suppose that indeed $\eta(w) = u$. Denote $D' = D \cup \{w\}$ and $\gamma' = \gamma[w \to u]$ for brevity. Then, for every connected component $Z \in \mathsf{CC}(X - \{w\})$, η has to map Z entirely into one subtree of F rooted at a child of u. Hence, for at least one child v of u we have that η restricted to $Z \cup D'$ witnesses that $\mathsf{Val}(v, Z, D', \gamma')$ is true. Since this holds for every $Z \in \mathsf{CC}(X - \{w\})$, we conclude that $\bigwedge_{Z \in \mathsf{CC}(X-\{w\})} \bigvee_{v \in \mathsf{Chld}(u)} \mathsf{Val}(v, Z, D', \gamma')$ is true. Conversely, supposing that $\bigwedge_{Z \in \mathsf{CC}(X-\{w\})} \bigvee_{v \in \mathsf{Chld}(u)} \mathsf{Val}(v, Z, D', \gamma')$ is true, for each $Z \in \mathsf{CC}(X - \{w\})$ we may find $v_Z \in \mathsf{Chld}(u)$ and a solution η_Z to the subproblem $\mathsf{Val}(v_Z, Z, D', \gamma')$. Observe that solutions η_Z match γ' on D', hence we may consider their union; call it η. It is straightforward to see that η is a subgraph embedding from $H[D \cup X]$ to G_u that is compliant with α and extends γ, i.e., it is a solution to (u, X, D, γ). Here, the only non-trivial condition is injectivity, but this is ensured by the fact that each solution η_Z is compliant with α: each vertex $x \in X - \{w\}$, say $x \in Z$, is mapped by η_Z to a vertex belonging to $\alpha^{-1}(x)$, so no two vertices lying in different connected components $Z, Z' \in \mathsf{CC}(X - \{w\})$ can be mapped by η_Z and $\eta_{Z'}$ to the same vertex of G_u. The formula follows. $\square$

Lemma 20 suggests the following algorithm for our problem. First, define a recursive procedure $\mathsf{ComputeVal}$ that given a subproblem (u, X, D, γ) computes its value using the recursive formula provided by Lemma 20. Then the algorithm proceeds as follows: for each connected component of H, say with vertex set X, verify whether there exists a root r of F for which $\mathsf{Val}(r, X, \emptyset, \emptyset)$ is true; this is done by invoking $\mathsf{ComputeVal}(r, X, \emptyset, \emptyset)$ for each root r of F. Finally, conclude that there is a subgraph embedding from H to G compliant with α if and only if for each connected component of H this verification was positive. The same reasoning as for Lemma 20, assertion (iii), proves that this algorithm is correct. Further, throghout the algorithm we store only a stack consisting of at most d frames of recursive calls to $\mathsf{ComputeVal}$, each of polynomial size, so the overall space complexity is polynomial in n. It remains to argue that the time complexity is $2^{\mathcal{O}((p+d)\log p)} \cdot n^{\mathcal{O}(1)}$.

To this end, we claim that $\mathsf{ComputeVal}$ is invoked on every subproblem (u, X, D, γ) at most once. As there are $2^{\mathcal{O}(p+d\log p)} \cdot n$ subproblems in total and the internal computation of $\mathsf{ComputeVal}$ for each of them takes polynomial time, the promised running time follows from this claim. To see the claim, observe that if $\mathsf{ComputeVal}$ is invoked on a subproblem (u, X, D, γ), then one of the following assertions holds:

- u is a root of F and $\mathsf{ComputeVal}(u, X, D, \gamma)$ is invoked directly in the main algorithm;
- u has a parent v and the subproblem (u, X, D, γ) uniquely defines the call to $\mathsf{ComputeVal}$ where $\mathsf{ComputeVal}(u, X, D, \gamma)$ was invoked: this call was to subproblem (v, X', D', γ') where
  - $D' = D - \gamma^{-1}(v)$ if $v \in \gamma(D)$ and $D' = D$ otherwise,
  - $\gamma' = \gamma|_{D'}$, and
  - X' is the vertex set of the unique connected component of $G - D'$ that contains $X \cup \gamma^{-1}(v)$, or $X' = \emptyset$ when $X \cup \gamma^{-1}(v) = \emptyset$.

With the above observation, the claim follows immediately: the subproblems on which $\mathsf{ComputeVal}$ is invoked in the main algorithm are pairwise different, while every other subproblem on which $\mathsf{ComputeVal}$ is invoked has a uniquely defined parent in the recursion tree. This means that every subproblem is solved at most once and we are done.

8 Omitted proofs from Section 2

8.1 Lifting through layering

Proof (of Lemma 6). Fix $p \in \mathbb{N}$. For any graph $G \in \mathscr{C}$ we shall construct a p-centered coloring of G using $(p+1) \cdot f(p)^2$ colors. We may assume that G is connected, as otherwise we treat each connected component of G separately and take the union of the obtained colorings. Note here that each connected component of G belongs to $\mathscr{C}$, because $\mathscr{C}$ is minor-closed.

Fix any vertex u of G and partition $V(G)$ into layers $L_0, L_1, L_2, \ldots \subseteq V(G)$ according to the distance from u: layer L_i comprises vertices exactly at distance i from u. Thus, $L_0 = \{u\}$, $\{L_0, L_1, L_2, \ldots\}$ forms a partition of $V(G)$, and every edge of G connects two vertices from same or adjacent layers. Let k be the largest integer such that layer L_k is non-empty.

For every $j \in \{0, 1, \ldots, k\}$ divisible by p, consider the graph G_j defined as follows: take the subgraph of G induced by $L_0 \cup L_1 \cup \ldots \cup L_{j+2p-1}$ and, provided $j > 0$, contract all vertices of $L_0 \cup L_1 \cup \ldots \cup L_{j-1}$ onto u; note that this is possible since $L_0 \cup L_1 \cup \ldots \cup L_{j-1}$ induces a connected subgraph of G. Note that G_j is obtained from G by vertex removals and edge contractions, so G_j is a minor of G; since $\mathscr{C}$ is minor-closed, we have $G_j \in \mathscr{C}$. Moreover, G_j is connected and has radius at most $2p$: this is straightforward for $j = 0$, while for $j > 0$ it can be easily seen that every vertex of G_j is at distance at most $2p$ from the vertex resulting from contracting $L_0 \cup L_1 \cup \ldots \cup L_{j-1}$. Finally, the vertex set of G_j contains the $2p$ consecutive layers $L_j, L_{j+1}, \ldots, L_{j+2p-1}$, plus one more vertex when $j > 0$. Thus, for every $i \in \{0, 1, \ldots, k\}$ and vertex $v \in L_i$ we have that $v \in V(G_{j-1})$ and $v \in V(G_j)$, where $j = p \cdot \lfloor i/p \rfloor$ is the largest integer divisible by p not larger than i. Here, G_{j-1} should be ignored if $j = 0$. Clearly, the layers and the graphs G_j are polynomial time computable.

Since each G_j is a graph from $\mathscr{C}$ that is connected and has radius at most $2p$, we may apply the assumed property of $\mathscr{C}$ to G_j in order to compute in polynomial time a p-centered coloring λ_j of G_j using $f(p)$ colors. We may assume that all colorings λ_j use the color set $\{1, \ldots, f(p)\}$. Now, define a coloring λ of G as follows: for $i \in \{0, 1, \ldots, k\}$ with $j = p \cdot \lfloor i/p \rfloor$, to each vertex $v \in L_i$ assign a color $\lambda(v)$ consisting of the following three of numbers:

$$i \bmod (p+1) \qquad ; \qquad \lambda_j(v) \qquad ; \qquad \lambda_{j-1}(v) \text{ if } j > 0 \text{, and } 1 \text{ otherwise.}$$

These three numbers are arranged into an ordered triple as follows: $i \bmod (p+1)$ is always the first coordinate, while $\lambda_j(v)$ is on the second coordinate if j is even and on the third coordinate if j is odd. The value $\lambda_{j-1}(v)$ (or 1 if $j = 0$) is put on the remaining coordinate. The ordered triple defined in this manner is set as the color $\lambda(v)$. Observe that thus, λ is a coloring of G using the color set $\{0, 1, \ldots, p\} \times \{1, \ldots, f(p)\} \times \{1, \ldots, f(p)\}$, which consists of $(p+1) \cdot f(p)^2$ colors. Clearly, λ is polynomial time computable from the layers and the colorings λ_i. So it remains to prove that λ is a p-centered coloring of G.

To this end, fix any connected subgraph H of G. Let $I \subseteq \{0, 1, \ldots, k\}$ be the set of those indices i, for which $V(H) \cap L_i \neq \emptyset$. Since H is connected, we have that I is an interval, i.e., $I = \{a, a+1, \ldots, b\}$ for some $0 \leqslant a \leqslant b \leqslant k$.

Suppose first that $b - a > p$. Then for each residue $r \in \{0, 1, \ldots, p\}$ there is $i \equiv r \bmod p$ such that $i \in I$, hence there is a vertex of H whose color under λ has r on the first coordinate. We infer that vertices of H receive more than p different colors under λ.

Suppose now that $b - a \leqslant p$, which means that $V(H) \subseteq L_a \cup L_{a+1} \cup \ldots \cup L_{a+p-1}$. Let $j = p \cdot \lfloor a/p \rfloor$ be the largest integer divisible by p not larger than a. Then $a - j < p$, hence $V(H) \subseteq L_j \cup L_{j+1} \cup \ldots \cup L_{j+2p-1} \subseteq V(G_j)$ and H is an induced subgraph of G_j. Since λ_j is a p-centered coloring of G_j and H is a connected subgraph of G_j, we infer that either H receives more than p colors under λ_j, or some color in λ_j appears exactly once among vertices of H. Moving to the coloring λ, observe that for every vertex v of H, the color $\lambda_j(v)$ appears either on the second or on the third coordinate of the color $\lambda(v)$, depending on whether j is even or odd. Consequently, for any two vertices $v, v' \in V(H)$ we have that $\lambda_j(v) \neq \lambda_j(v')$ implies $\lambda(v) \neq \lambda(v')$, and the above mentioned property of H under the coloring λ_j carries over to H under the coloring λ. $\square$

8.2 Lifting through partitions

Proof (of Lemma 7). Since each part of $\mathcal{P}$ has at most q vertices, we can compute a coloring $\kappa\colon V(G) \to C$ for a color set C of size q so that within each part of $\mathcal{P}$ all vertices receive pairwise different colors. By assumption, we can also compute in polynomial time a p-centered coloring $\lambda_0\colon \mathcal{P} \to D$ of $G/\mathcal{P}$ for a color set D of size $f(p)$. Let $\lambda\colon V(G) \to D$ be a natural lift of λ_0 to G: for each $u \in V(G)$ we put $\lambda(u) = \lambda_0(A)$, where $A \in \mathcal{P}$ is such that $u \in V(A)$. We now construct the product coloring $\rho\colon V(G) \to C \times D$ defined as

$$\rho(u) = (\kappa(u), \lambda(u)) \qquad \text{for each } u \in V(G).$$

Since ρ uses $q \cdot f(p)$ colors, it suffices to verify that ρ is p-centered.

Let $G' = G/\mathcal{P}$. Take any connected subgraph H of G. Let $\mathcal{X} \subseteq \mathcal{P}$ be the set of those parts of $\mathcal{P}$ that intersect H. Since H is connected, the graph $G'[\mathcal{X}]$ is connected as well. We infer that either parts from $\mathcal{X}$ receive more than p different colors in λ_0, or there is a part $A \in \mathcal{X}$ whose color is unique in $\mathcal{X}$ under λ_0. In the first case, it follows immediately that H receives more than p different colors in ρ, as there are already p different second coordinates of the colors of vertices of H. In the second case, each vertex of A receives a different color under λ, and no other vertex of H can share this color, because A is colored uniquely among $\mathcal{X}$. It follows that every vertex of $V(A) \cap V(H)$ has a unique color under ρ among vertices of H; since this intersection is non-empty, the claim follows. $\square$

Proof (of Corollary 8). By Lemma 6, it suffices to show that for every $p \in \mathbb{N}$, every connected graph $G \in \mathscr{C}$ of radius at most $2p$ has a polynomial time computable p-centered coloring with $(4p+1) \cdot f(p)$ colors. By assumption, there is a ptime computable partition $\mathcal{P}_G$ of G such that every $P \in \mathcal{P}_G$ is a geodesic in G and the graph $H = G/\mathcal{P}_G$ admits a ptime computable p-centered coloring with $f(p)$ colors. Observe that any geodesic in a graph of radius at most $2p$ has length at most $4p$, hence each geodesic $P \in \mathcal{P}_G$ contains at most $4p+1$ vertices. The claim follows by Lemma 7. $\square$

8.3 Lifting through tree decompositions

Before proving Lemma 9, we collect several properties of tree decompositions. Let $\mathcal{T} = (T, \beta)$ be a tree decomposition of G. We use the following notation whenever $\mathcal{T}$ is clear from the context.

1. We have a natural ancestor/descendant relation in T: a node is a descendant of all the nodes that appear on the unique path leading from it to the root. Note that every node of T is also its own ancestor and descendant. We write $x \leqslant_T y$ if x is an ancestor of y. Then $\leqslant_T$ is a partial order on the nodes of T with the root being the unique $\leqslant_T$-minimal element.

2. The *margin* of a node x is the set $\mu(x) = \beta(x) - \alpha(x)$. Recall here that $\alpha(x)$ is the adhesion set of x.

3. For every vertex u of G, let $x(u)$ be the unique $\leqslant_T$-minimal node of T with $u \in \beta(x)$. Note that this node is unique due to condition (T1). We define a quasi-order $\leqslant_{\mathcal{T}}$ on the vertex set of G as follows: $u \leqslant_{\mathcal{T}} v$ if and only if $x(u) \leqslant_T x(v)$.

4. The *torso* of a node x is the graph $\Gamma(x)$ on vertex set $\beta(x)$ where two vertices $u, v \in \beta(x)$ are adjacent if and only if $uv \in E(G)$ or if there exists $y \neq x$ such that $u, v \in \beta(y)$. Equivalently, $\Gamma(x)$ is obtained from $G[\beta(x)]$ by turning the adhesion sets of x and of all children of x into cliques.

5. We call $\mathcal{T}$ a tree-decomposition *over* a class $\mathscr{C}$ of graphs if $\Gamma(x) \in \mathscr{C}$ for every node x of T.

6. The *skeleton* of G over $\mathcal{T}$ is the directed graph S with vertex set $V(S) = V(G)$ and arc set defined as follows: for each $x \in V(T)$, $u \in \mu(x)$, and $v \in \alpha(x)$, we put the arc (u, v) into the arc set of S.

Note that if (u, v) is an arc in the skeleton S, then in particular $v <_{\mathcal{T}} u$, equivalently $x(v) <_T x(u)$. This implies that the skeleton is always acyclic (i.e. it is a DAG).

The following lemmas express well-known properties of tree decompositions.

Lemma 21. *If $\mathcal{T} = (T, \beta)$ is a tree decomposition of a graph G and uv is an edge in G, then $x(u)$ is an ancestor of $x(v)$ or vice versa. Consequently, $u \leqslant_{\mathcal{T}} v$ or $v \leqslant_{\mathcal{T}} u$.*

Proof. Otherwise the sets of nodes whose bags contain u and v, respectively, would be disjoint, which would be a contradiction with the existence of the edge uv by condition (T2). □

Lemma 22. *Let $\mathcal{T} = (T, \beta)$ be a tree decomposition of a graph G. For every vertex u of G, the node $x(u)$ is the unique node of T whose margin contains u. Consequently, $\{\mu(x)\}_{x \in V(T)}$ is a partition of the vertex set of G.*

Proof. Vertex u belongs to $\mu(x)$ for some node x if and only if $u \in \beta(x)$ and either x is the root, or the parent y of x satisfies $u \notin \beta(y)$. By condition (T1), among nodes x with $u \in \beta(x)$ there is exactly one satisfying the second condition, being $x(u)$. □

Note that by Lemma 22, the margins of nodes of T are exactly the classes of equivalence in the quasi-order $\leqslant_{\mathcal{T}}$ on $V(G)$.

We start the proof of Lemma 9 by observing some properties of the skeleton graph. Fix $p \in \mathbb{N}$, a graph G, a tree decomposition $\mathcal{T} = (T, \beta)$ of G with adhesion at most k, and let S be the skeleton of G over $\mathcal{T}$.

First, we show that restricted reachability in G implies reachability in the skeleton.

Lemma 23. *Let u, v be vertices of G with $v <_{\mathcal{T}} u$ and let P be a path in G with endpoints u and v such that every vertex w of P apart from v satisfies $v <_{\mathcal{T}} w$. Then there exists a directed path Q in S leading from u to v and satisfying $V(Q) \subseteq V(P)$.*

Proof. We proceed by induction on the length of path P. Let w be the first (closest to u) vertex on P satisfying $w <_{\mathcal{T}} u$; such w exists because v satisfies the condition.

Claim 1. *(u, w) is an arc in S.*

Proof. Let w' be the predecessor of w on P. We argue by induction that every vertex t on the prefix of P between u and w' satisfies $u \leqslant_{\mathcal{T}} t$. This holds trivially for $t = u$. Supposing it holds for some vertex t, we argue that it holds also for the successor t' of t on the prefix. Indeed, we have $t \leqslant_{\mathcal{T}} t'$ or $t' \leqslant_{\mathcal{T}} t$ by Lemma 21, implying that either $u \leqslant_{\mathcal{T}} t'$ or $u >_{\mathcal{T}} t'$, but the latter case is excluded by the choice of w, proving the induction step. In particular, we infer that $u \leqslant_{\mathcal{T}} w'$, implying $w <_{\mathcal{T}} u \leqslant_{\mathcal{T}} w'$, which means that $x(w) <_T x(u) \leqslant_T x(w')$.

Since ww' is an edge in G, assertion $x(w) <_T x(w')$ together with conditions (T1) and (T2) imply that w is contained in all the bags of nodes on the unique path in T between $x(w)$ and $x(w')$. Since

$x(w) <_T x(u) \leqslant_T x(w')$, we infer that w is contained in the bag of both $x(u)$ and of its parent, which means that $w \in \alpha(x(u))$. As $u \in \mu(x(u))$ by Lemma 22, (u, w) is an arc in S, as claimed. ⌟

Since w lies on P, by assumption we have either $w = v$ or $v <_{\mathcal{T}} w$. In the former case we are immediately done, as we take Q to be the path consisting only of the arc $(u, w) = (u, v)$. In the latter case we may apply the induction assumption to w and v connected by the suffix of P from w to v. This yields a path Q', which may be extended to a suitable path Q by adding the edge (u, v) at the front. □

Lemma 23 suggests studying reachability in the skeleton. We next show that if the considered tree decomposition has small adhesion, then every vertex reaches only a small number of vertices via short paths in the skeleton. The argument essentially boils down to the combinatorial core of the proof that graphs of treewidth k have weak p-coloring number $\mathcal{O}(p^k)$ [20].

Lemma 24. *Consider any vertex u of G and let $p \in \mathbb{N}$. Then there exist at most $\binom{p+k}{k}$ vertices of G that are reachable by a directed path of length at most p from u in S.*

Proof. We proceed by induction on $p + k$, with base case $p = 1$ or $k = 1$. When $p = 1$, we observe that the out-degrees in the skeleton are bounded by k, so there can be at most $k + 1 = \binom{k+1}{k}$ vertices reachable from u by a path of length at most 1. When $k = 1$, the skeleton S is a directed forest and every vertex reaches at most $p + 1 = \binom{p+1}{1}$ vertices by paths of length at most p.

We proceed to the induction step. Let $x = x(u)$. By definition, the set of out-neighbors of u in S is exactly the adhesion set $\alpha(x)$. If $\alpha(x) = \emptyset$, then u has no out-neighbors and there is nothing to prove, so assume otherwise. For each $w \in \alpha(x)$ we have $x(w) <_T x$, hence nodes $\{x(w) \colon w \in \alpha(x)\}$ are pairwise comparable in the order $\leqslant_T$. Let y be the $\leqslant_T$-minimal element of $\{x(w) \colon w \in \alpha(x)\}$. Note that $y <_T x$. Fix any $a \in \alpha(x)$ with $y = x(a)$.

Let R be the set of vertices reachable from u by a directed path of length at most p in S; we need to prove that $|R| \leqslant \binom{p+k}{k}$. Note that whenever some $v \in R$ can be reached from u by a directed path P of length at most p in S, then all vertices on P are contained in R; this is witnessed by the prefixes of P. Since arcs in the skeleton S point always to a vertex that is strictly smaller in the quasi-order $\leqslant_{\mathcal{T}}$, we have that $v <_{\mathcal{T}} u$ for each $v \in R - \{u\}$; this in particular implies $x(v) \leqslant_T x$ for all $v \in R$. As $y <_T x$, we may partition R into two subsets R_1 and R_2 as follows:

$$R_1 = (R \cap \{v \colon x(v) <_T y\}) \cup \{a\} \qquad \text{and} \qquad R_2 = (R \cap \{v \colon y \leqslant_T x(v) \leqslant_T x\}) - \{a\}.$$

Note that in particular $a \in R_1$ and $u \in R_2$, and by the choice of y we have that a is the only out-neighbor of u in R_1. We now analyze the interaction between R_1 and R_2.

Claim 1. *In S there is no arc with tail in R_1 and head in R_2.*

Proof. Arcs in S always point to a vertex that is strictly smaller in the quasi-order $\leqslant_{\mathcal{T}}$, but for each $s \in R_2$ and $t \in R_1$ we have $x(t) \leqslant_T y \leqslant_T x(s)$. ⌟

Claim 2. *For every arc (s, t) of S with $s \in R_2$ and $t \in R_1$, either $t = a$ or (a, t) is an arc in S.*

Proof. Supposing $t \neq a$ we have $x(t) <_T y$. As (s, t) is an arc of S, we have $t \in \alpha(x(s))$, which implies $t \in \beta(x(s))$. As $s \in R_2$, we have $y \leqslant_T x(s)$, hence $x(t) <_T y \leqslant_T x(s)$. By condition (T1), $t \in \beta(x(s))$ entails that t belongs to every bag on the unique path in T between $x(t)$ and $x(s)$, so in particular $t \in \alpha(y)$. As $x(a) = y$, we find that $a \in \mu(y)$ and $t \in \alpha(y)$, hence (a, t) is an arc in S. ⌟

Claim 1 and Claim 2 together imply the following.

Claim 3. *Every vertex $v \in R_2$ is reachable from u in S by a directed path of length at most p whose vertices are all contained in R_2. Every vertex $v \in R_1$ is reachable from a in S by a directed path of length at most $p-1$ whose vertices are all contained in R_1.*

Proof. For the first statement, by Claim 1 every path P witnessing $v \in R$ has to be entirely contained in R_2. For the second statement, suppose P is a path witnessing $v \in R$. Let s be the last vertex on P that belongs to R_1 and let t be its successor on P; these vertices exist due to $u \in R_2$ and $v \in R_1$. If $t = a$, then we may take the suffix of P from t to v. Otherwise, by Claim 2 we have that (a, t) is an arc in S. It then suffices to augment the suffix of P from t to v by adding the arc (a, t) at the front. Observe that the path obtained in this manner has length at most $p - 1$, because $t \in R_1 - \{a\}$ entails that t is not an out-neighbor of u, which means that the prefix of P from u to t has length at least 2. ⌟

By combining the induction assumption with the second statement of Claim 3 we infer that

$$|R_1| \leqslant \binom{p-1+k}{k}.$$

Observe now that for the graph $G[R_2]$ we may construct a tree decomposition $\mathcal{T}'$ from $\mathcal{T}$ by removing all vertices outside of R_2 from all the bags, and moreover removing all nodes not lying on the path from x to y in T; this is because $x(v)$ lies on this path for each $v \in R_2$. It is easy to see that the skeleton of $G[R_2]$ over $\mathcal{T}'$ is equal to $S[R_2]$. Further, observe that since (u, a) is an arc of S, we have $a \in \alpha(x) \subseteq \beta(x)$, which together with $a \in \beta(y)$ implies that in $\mathcal{T}$, the vertex a is contained in the bags of all the nodes z satisfying $y \leqslant_T z \leqslant_T x$. Since $a \notin R_2$, this implies that the adhesion of $\mathcal{T}'$ is at most $k - 1$. Then combining the induction assumption with the first statement of Claim 3 yields

$$|R_2| \leqslant \binom{p+k-1}{k-1}.$$

All in all, we have

$$|R| = |R_1| + |R_2| \leqslant \binom{p+k-1}{k} + \binom{p+k-1}{k-1} = \binom{p+k}{k};$$

this concludes the induction step. □

We may now use Lemma 24 to find a coloring of the skeleton using a small number of colors so that all pairs of vertices connected by short paths receive different colors.

Lemma 25. *There is a coloring $\kappa\colon V(S) \to C$ with a color set C of size $\binom{p+k}{k}$ such that:*

- *if there exists a directed path of length at most p from u to v in S, then $\kappa(u) \neq \kappa(v)$; and*
- *for every node $x \in V(T)$, all vertices of $\mu(x)$ receive the same color under κ.*

Moreover, given $G, \mathcal{T}, p$ on input, such a coloring can be computed in polynomial time.

Proof. As observed before, the skeleton S can be computed in polynomial time. We first compute the p-transitive closure S^p of S, which is a directed graph on the same vertex set as S where we put an arc (u, v) whenever there is a path in S from u to v of length at most p. By Lemma 24, every vertex has out-degree at most $\binom{p+k}{k} - 1$ in S^p; the -1 summand comes from counting the vertex itself among reachable ones in Lemma 24. As S is acyclic, so is S^p as well. Hence, we can compute any topological ordering of S^p and iterate along the ordering while coloring the vertices greedily using the color set $\{1, \ldots, \binom{p+k}{k}\}$.

Specifically, every vertex u receives the smallest color that is not present among the out-neighbors of u, which were all colored in the previous iterations. This yields a proper coloring κ of the undirected graph underlying S^p, which is also coloring of S with the sought properties. Note here that all vertices residing in the same margin $\mu(x)$ for some $x \in V(T)$ have exactly the same out-neighbors in S, so they will receive the same color in the procedure. □

We are ready to prove Lemma 9.

Proof (of Lemma 9). Fix any graph $G \in \mathscr{D}$ and let $\mathcal{T}$ be a tree decomposition of G over $\mathscr{C}$ with adhesion at most k; by assumption, such $\mathcal{T}$ can be computed in polynomial time. Let S be the skeleton of G over $\mathcal{T}$ and let $\kappa\colon V(G) \to C$ be the coloring of S provided by Lemma 25. As asserted, $|C| = \binom{p+k}{k} = \mathcal{O}(p^k)$ and κ can be computed in polynomial time.

We now define a coloring λ of G as follows. By assumption, for every node x of T, we may compute in polynomial time a p-centered coloring λ_x of the torso $\Gamma(x)$, where each coloring λ_x uses the same color set D of size $\mathcal{O}(p^d)$. Then $\lambda\colon V(G) \to D$ is defined as

$$\lambda(u) = \lambda_{x(u)}(u) \qquad \text{for each } u \in V(G).$$

In other words, we restrict each coloring λ_x to $\mu(x)$ and λ is the union of all those restrictions.

We finally define a coloring $\rho\colon V(G) \to C \times D$ as the product of κ and λ, that is

$$\rho(u) = (\kappa(u), \lambda(u)) \qquad \text{for each } u \in V(G).$$

Note that ρ uses $\mathcal{O}(p^{d+k})$ colors, hence it suffices to prove that ρ is a p-centered coloring of G.

To this end, fix any connected subgraph H of G. Since H is connected, by conditions (T1) and (T2) we infer that the set $\{x\colon \beta(x) \cap V(H) \neq \emptyset\}$ is connected in T. Consequently, this set contains a unique $\leqslant_T$-minimal node; call it z. Note that $z \leqslant_T x(u)$ for each $u \in V(H)$. Since the bag of the parent of z (provided it exists) is disjoint with $V(H)$, we infer that $V(H) \cap \beta(z) = V(H) \cap \mu(z)$. Fix any vertex $v \in V(H) \cap \mu(z)$ and let $\kappa(v) = c$.

Suppose first that there exists a vertex $u \in V(H)$ with $x(u) \neq z$ and $\kappa(u) = c$. Let P be any path in H connecting u and v and let v' be the first (closest to u) vertex on P satisfying $x(v') = z$. Then we have $v, v' \in \mu(z)$, so by the properties of κ asserted by Lemma 25 we have that $\kappa(v') = \kappa(v) = c$. As $x(u) \neq z$, we have $u \neq v'$.

Let P' be the prefix of P from u to v'. By the choice of z and of v' we have that $v' <_{\mathcal{T}} w$ for each vertex w on P' different from v'. Hence, we may apply Lemma 23 to u and v' to infer that in S there is a directed path Q leading from u to v' and satisfying $V(Q) \subseteq V(P') \subseteq V(H)$. By the properties of κ asserted by Lemma 25, every $p+1$ consecutive vertices on Q receive pairwise different colors under κ. Since $u \neq v'$ and $\kappa(u) = \kappa(v') = c$, we conclude that Q has length at least $p+2$ and among the first $p+1$ vertices of Q there are $p+1$ different colors present. But $V(Q) \subseteq V(H)$, so H receives more than p different colors under κ, hence it also receives more than p different colors under ρ and we are done.

We are left with the case when the vertices of H that receive color c under κ are exactly the vertices of $\mu(z)$. Let H' be the subgraph of $\Gamma(z)$ induced by $V(H) \cap \mu(z)$. We claim that H' is connected. For this, take any $a, b \in V(H) \cap \mu(z)$ and let R be a path in H connecting a and b. By the properties of tree decompositions, for every maximal infix of R lying outside of $V(H) \cap \mu(z)$, the two vertices immediately preceding and immediately succeeding this infix on R have to belong to the same adhesion set $\alpha(z')$ for some child z' of z. As $\alpha(z')$ is turned into a clique in $\Gamma(z)$, we may shortcut this infix by using the edge connecting the two vertices. By performing this operation for infix of R as above, we turn R into a path R' in H' connecting a and b.

Since H' is connected, the vertices of $V(H') = V(H) \cap \mu(z)$ either receive more than p different colors in λ_z, or some color appears exactly once in H' under λ_z. Recall that λ and λ_z coincide on $\mu(z)$, so the above alternative holds for λ as well. In the former case, H' receives more than p different colors in λ, implying the same for H and ρ. In the latter case, let v_0 be a vertex whose color under λ is unique in H'. Then the color of v_0 under ρ is unique in H: the colors of all other vertices of $V(H) \cap \mu(z)$ differ on the second coordinate, while the colors of all vertices of $V(H) - \mu(z)$ differ on the first coordinate. Since H was chosen arbitrarily, this concludes the proof. □

9 Omitted proofs from Section 4

We first give the proof of Lemma 12.

Proof (of Lemma 12). As mentioned, this result was essentially proved in [14, Lemma 5.7] and we follow closely the reasoning described there. Our presentation is based on that of [29, Lemma 10.1], which was another rephrasing of the same combinatorial fact.

In the proof we will use the dual multigraph $G^\star$. Recall that the vertex set of $G^\star$ is the face set of G, and for every edge e of G we put in $G^\star$ the *dual edge* $e^\star$ connecting the two faces of G incident to e. We may define a proper Σ-embedding of $G^\star$ by mapping every face $f \in F(G)$ to any fixed point inside it, and every dual edge $e^\star = ff'$ to a suitable chosen curve connecting the points assigned to f, f' that is contained in $f \cup f' \cup e$ and crosses e once. Also, clearly $G^\star$ is connected.

Fix any vertex $u \in V(G)$. Let T be the tree of a breadth-first search from u in G. Since G is connected, T is a tree on the same vertex set as G. Moreover, for every $v \in V(G)$ the unique path in T between u and v is a shortest u-v path in G; in particular, it is a geodesic in G. Let $M = \{e^\star \colon e \in E(T)\} \subseteq E(G^\star)$ be the set of edges dual to the edges of T.

Consider now the graph $G^\star - M$ and observe that it is connected, because T is acyclic. Let S be any spanning tree of $G^\star - M$. Denote $N = E(S)$. Clearly, M and N are disjoint subsets of $E(G^\star)$ and we have $|M| = |V(G)| - 1$ and $|N| = |V(G^\star)| - 1 = |F(G)| - 1$. Consequently, if we define $X^\star = E(G^\star) - (M \cup N)$, then

$$|X^\star| = |E(G^\star)| - |M| - |N| = |E(G)| - |V(G)| - |F(G)| + 2 = 2g,$$

where g is the Euler genus of Σ. Denote $X = \{e \colon e^\star \in X^\star\}$.

Define a subgraph K of G by taking the union of the edges of X and the u-v paths in T, for all vertices v that are endpoints of edges in X. Since $|X| = 2g$, there are at most $4g$ such vertices v, hence K is the union of the $2g$ edges of X and at most $4g$ paths $P_1, \ldots, P_k$ in T, where each P_i is a geodesic in G. By considering paths $Q_i = P_i - (V(P_1) \cup \ldots \cup V(P_{i-1}))$ and removing all Q_is that turn out to be empty, we find that K can be partitioned into at most $4g$ geodesics in G. Therefore, it remains to argue that K is a cut-graph of G.

Consider first the graph $\widetilde{K}$ obtained from T by adding all edges of X. We argue that $\widetilde{K}$ is a cut-graph of G. To see this, fix any face $f \in F(G)$ and consider S as a tree rooted at f. Note that $\Sigma - \widetilde{K}$ can be obtained from f by iteratively gluing faces of $F(G) - \{f\}$ along edges dual to the edges of S, in a top-down manner on S. Each face of f is homeomorphic to a disk and gluing two disks along a common segment of their boundaries yields a disk. Thus, throughout the above process we maintain the invariant that the topological space glued so far is homeomorphic to a disk, yielding at the end that $\Sigma - \widetilde{K}$ is homeomorphic to a disk.

Next, observe that K can be obtained from $\widetilde{K}$ by iteratively removing vertices of degree 1. Note the following claim: if H is a cut-graph of G and w is a vertex of degree 1 in H, then $H' = H - w$ is also a cut-graph of G. Indeed, $\Sigma - H'$ is obtained from the disk $\Sigma - H$ by gluing together the two segments of its boundary corresponding to the two sides of the unique edge of H incident to w. These two segments share an endpoint and are glued in opposite directions, so $\Sigma - H'$ remains homeomorphic to a disk. By applying this claim iteratively starting with $\widetilde{K}$, we infer that K is indeed a cut-graph of G. ☐

Next, we give a full proof of Theorem 13. Since our reasoning uses parts of the proof of Theorem 4, we assume familiarity with it.

Proof (of Theorem 13). We assume that G can be properly embedded in Σ, since otherwise G is embeddable in a surface of smaller genus and we can apply the reasoning for this surface.

By Lemma 12, we can compute a cut-graph K of G and its partition $\mathcal{Q}$ into at most $4g$ geodesics in G. Consider a graph $\widehat{G}$ obtained by *cutting G along K* as follows; see Figure 2 for reference. Starting from G, first duplicate every edge e of K and embed the two copies of e in a small neighborhood of the original embedding, thus creating a very thin face of length 2 incident to both copies. Next, for every vertex u of K scan the edges of K incident to u in the cyclic order around u in the embedding; we note that we do not assume here that Σ is orientable, as reversing the order yields the same construction. For every pair of consecutive edges e, e' in the order, create a copy of u and make it incident to one copy of e, one copy of e', and all the edges of G lying between e and e' in the cyclic order, as in Figure 2. The original vertex u is removed and copies are shifted within a small neighborhood of the original placement of u as in Figure 2.

Finally, remove from Σ the space that is not contained in the (slightly shifted) faces of G; this space is depicted in grey in the last panel of Figure 2. Since K was a cut-graph, the obtained topological space is homeomorphic to a closed disc and the graph $\widehat{G}$ is embedded into it. Moreover, the boundary of $\widehat{G}$ is a simple cycle C_0 that contains two copies of every edge of K and as many copies of every vertex of K as its degree in K. Let $\pi\colon V(\widehat{G}) \cup E(\widehat{G}) \to V(G) \cup E(G)$ be the mapping sending every vertex and edge of $\widehat{G}$ to its origin in G.

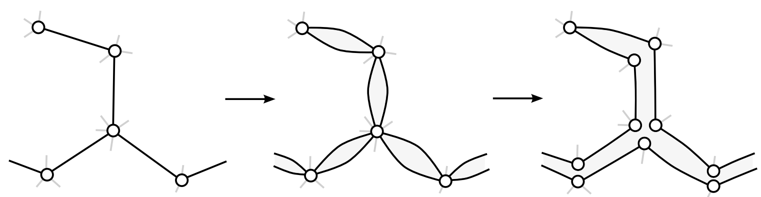

Figure 2: Cutting G along K. The edges of K are depicted in black, the remaining edges of G are in grey. The grey area in the last panel is the part of the surface that gets removed.

The graph $\widehat{G}$ is often called the *polygonal schema* of Σ, where Σ is treated as a 2-dimensional cell complex consisting of faces of G glued along the edges of G. We remark that in [14], Erickson and Har-Peled explain a different, equivalent construction of $\widehat{G}$: take the faces of G, considered as closed discs, and instead of gluing them along all the edges of G to obtain surface Σ, we glue them along the edges of $E(G) - E(K)$ to obtain the disc $\Sigma - K$. The graph obtained in this manner is $\widehat{G}$.

Call an edge e of K *delimiting* if e is not contained in any geodesic from $\mathcal{Q}$. In the proof of Lemma 12 we have argued that K is a tree with $2g$ edges added, thus $|E(K)| = |V(K)| - 1 + 2g$. Since $|\mathcal{Q}| \leqslant 4g$ and every geodesic $Q \in \mathcal{Q}$ satisfies $|E(Q)| = |V(Q)| - 1$, we have

$$|E(K)| - \sum_{Q\in\mathcal{Q}} |E(Q)| = |E(K)| - \sum_{Q\in\mathcal{Q}} |V(Q)| + |\mathcal{Q}| = |V(K)| - |E(K)| + |\mathcal{Q}| = 2g - 1 + |\mathcal{Q}| \leqslant 6g.$$

Hence, in K there is at most $6g$ delimiting edges, so on C_0 there are at most $12g$ copies of delimiting edges.

Let $\mathcal{R}$ be the partition of C_0 into paths obtained by removing from C_0 all copies of all delimiting edges and taking all the obtained connected components as parts of $\mathcal{R}$. Then $|\mathcal{R}| \leqslant 12g$ and every path in $\mathcal{R}$ is mapped under π to a subpath of a path in $\mathcal{Q}$.

We now apply the reasoning leading to Claim 1 from the proof of Theorem 4. First, we triangulate $\widehat{G}$, obtaining a triangulated, disc-embedded graph $\widehat{G}^+$ with boundary being a simple cycle C_0. Next, we redefine the notion of a tight cycle as follows: a cycle C in $\widehat{G}^+$ is *tight* if it admits a partition into paths in $\widehat{G}$, out of which all but at most 6 are subpaths of different paths from $\mathcal{R}$, and the remaining at most 6 paths are geodesics in G. One can then readily verify that with the definition of tightness amended in this way, the inductive proof of Claim 1 works just as before. Here are the main differences:

- When defining paths Q_1, Q_2, Q_3, instead of requiring that they consist of at most two geodesics on C, they are now a concatenation of an arbitrary number of paths from the partition witnessing tightness of C, however all but at most 2 of these paths have to be subpaths of paths from $\mathcal{R}$.

- When splitting the graph along paths K_1, K_2, K_3, we observe that L_1, L_2, L_3 (paths K_1, K_2, K_3 trimmed by removing the vertex lying on C) are geodesics in G, as K_1, K_2, K_3 are defined as shortest paths in $\widehat{G}$ from the face f to the cycle C.

We leave verifying the straighforward details to the reader.

All in all, the above reasoning yields a partition $\widehat{\mathcal{P}}$ of $\widehat{G}$ into paths such that $\mathcal{R} \subseteq \widehat{\mathcal{P}}$, all paths in $\widehat{\mathcal{P}} - \mathcal{R}$ are geodesics in G, and $\widehat{G}/\widehat{\mathcal{P}}$ admits a tree decomposition where every bag contains at most 9 paths from $\widehat{\mathcal{P}} - \mathcal{R}$. Note now that since $\bigcup_{R \in \mathcal{R}} \pi(V(R)) = V(K)$, we have that $\widehat{\mathcal{P}} - \mathcal{R}$ is a partition of $G - V(K)$ such that $(G - V(K))/(\widehat{\mathcal{P}} - \mathcal{R})$ has treewidth at most 8. Since $\mathcal{Q}$ consists of geodesics in G and $\bigcup_{Q \in \mathcal{Q}} V(P) = V(K)$, we may output $\mathcal{Q}$ and the partition $\mathcal{P} = (\widehat{\mathcal{P}} - \mathcal{R}) \cup \mathcal{Q}$ of G. $\square$

10 Conclusions

In this paper we gave the first polynomial upper bounds on the number of colors needed for p-centered colorings on proper minor-closed graph classes, including the first such upper bounds for planar graphs. Admittedly, the obtained $\mathcal{O}(p^{19})$ upper bound for planar graphs does not look very practical. We see that with a deeper technical analysis of our construction one can reduce the degree of the polynomial to below 10, but we decided to keep the presentation clean and settle for a slightly higher, yet still polynomial bound. A real challenge would be to obtain tight bounds for planar graphs: can the number of colors be, say, quadratic or even near-linear in p?

Perhaps more importantly, so far we are lacking any tools for proving lower bounds on the number of colors needed for p-centered colorings. It is known that there exist graphs of treewidth k that have weak p-coloring number $\Omega(p^k)$ [20]. We conjecture that the same holds also for p-centered colorings: the degree of the polynomial needs to increase with k. Similarly, we conjecture that the number of colors needed for p-centered colorings in the class of graphs of maximum degree 3 is exponential in p. Note that this class admits every graph as a minor, but has bounded expansion.

Acknowledgements. The authors thank Zdenek Dvořák and Patrice Ossona de Mendez for insightful discussions on the problem.

References

[1] *Question posed at the Workshop on Algorithms, Logic and Structure, 12-14 December, 2016, University of Warwick.* https://warwick.ac.uk/fac/sci/maths/people/staff/ daniel_kral/alglogstr/openproblems.pdf.

[2] I. Abraham and C. Gavoille. Object location using path separators. In *PODC 2006*, pages 188–197. ACM, 2006.

[3] N. Alon, R. Yuster, and U. Zwick. Color-coding. *J. ACM*, 42(4):844–856, 1995.

[4] H. L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM Journal on Computing*, 25(6):1305–1317, 1996.

[5] H. L. Bodlaender, J. Nederlof, and T. C. van der Zanden. Subexponential time algorithms for embedding H-minor free graphs. In *ICALP 2016*, volume 55 of *LIPIcs*, pages 9:1–9:14. Schloss Dagstuhl — Leibniz-Zentrum für Informatik, 2016.

[6] E. Demaine, M. Hajiaghayi, and K. Kawarabayashi. Algorithmic graph minor theory: Decomposition, approximation, and coloring. In *FOCS 2005*, pages 637–646. IEEE, 2005.

[7] E. D. Demaine, M. Hajiaghayi, and K. Kawarabayashi. Contraction decomposition in H-minor-free graphs and algorithmic applications. In *STOC 2011*, pages 441–450. ACM, 2011.

[8] E. D. Demaine, M. Hajiaghayi, and B. Mohar. Approximation algorithms via contraction decomposition. *Combinatorica*, 30(5):533–552, 2010.

[9] M. DeVos, G. Ding, B. Oporowski, D. P. Sanders, B. Reed, P. Seymour, and D. Vertigan. Excluding any graph as a minor allows a low tree-width 2-coloring. *Journal of Combinatorial Theory, Series B*, 91(1):25–41, 2004.

[10] R. Diestel. *Graph Theory, 4th Edition*, volume 173 of *Graduate Texts in Mathematics*. Springer, 2012.

[11] F. Dorn. Planar Subgraph Isomorphism revisited. In *STACS 2010*, volume 5 of *LIPIcs*, pages 263–274. Schloss Dagstuhl — Leibniz-Zentrum für Informatik, 2010.

[12] Z. Dvořák, D. Král', and R. Thomas. Testing first-order properties for subclasses of sparse graphs. *J. ACM*, 60(5):36, 2013.

[13] D. Eppstein. Subgraph isomorphism in planar graphs and related problems. *J. Graph Algorithms Appl.*, 3(3), 1999.

[14] J. Erickson and S. Har-Peled. Optimally cutting a surface into a disk. *Discrete & Computational Geometry*, 31(1):37–59, 2004.

[15] F. V. Fomin, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. Subexponential parameterized algorithms for planar and apex-minor-free graphs via low treewidth pattern covering. In *FOCS 2016*, pages 515–524. IEEE Computer Society, 2016.

[16] J. Gajarský, S. Kreutzer, J. Nešetřil, P. Ossona de Mendez, M. Pilipczuk, S. Siebertz, and S. Toruńczyk. First-order interpretations of bounded expansion classes. In *ICALP 2018*, pages 126:1–126:14, 2018.

[17] M. Grohe. Local tree-width, excluded minors, and approximation algorithms. *Combinatorica*, 23(4):613–632, 2003.

[18] M. Grohe, K. Kawarabayashi, and B. A. Reed. A simple algorithm for the graph minor decomposition — logic meets structural graph theory. In *SODA 2013*, pages 414–431. SIAM, 2013.

[19] M. Grohe and S. Kreutzer. Methods for algorithmic meta theorems. *Model Theoretic Methods in Finite Combinatorics*, 558:181–206, 2011.

[20] M. Grohe, S. Kreutzer, R. Rabinovich, S. Siebertz, and K. Stavropoulos. Colouring and covering nowhere dense graphs. In *WG 2015*, pages 325–338, 2015.

[21] K. Kawarabayashi and P. Wollan. A simpler algorithm and shorter proof for the graph minor decomposition. In *STOC 2011*, pages 451–458. ACM, 2011.

[22] H. A. Kierstead, D. Yang. Orderings on graphs and game coloring number. *Order*, 20(3):255–264, 2003.

[23] R. J. Lipton and R. E. Tarjan. A separator theorem for planar graphs. *SIAM Journal on Applied Mathematics*, 36(2):177–189, 1979.

[24] J. Nešetřil and P. Ossona de Mendez. Tree-depth, subgraph coloring and homomorphism bounds. *European Journal of Combinatorics*, 27(6):1022–1041, 2006.

[25] J. Nešetřil and P. Ossona de Mendez. Grad and classes with bounded expansion I. Decompositions. *European Journal of Combinatorics*, 29(3):760–776, 2008.

[26] J. Nešetřil and P. Ossona de Mendez. On nowhere dense graphs. *European Journal of Combinatorics*, 32(4):600–617, 2011.

[27] J. Nešetřil and P. Ossona de Mendez. *Sparsity — Graphs, Structures, and Algorithms*, volume 28 of *Algorithms and combinatorics*. Springer, 2012.

[28] M. P. O'Brien and B. D. Sullivan. Experimental evaluation of counting subgraph isomorphisms in classes of bounded expansion. *CoRR*, abs/1712.06690, 2017.

[29] M. Pilipczuk, M. Pilipczuk, P. Sankowski, and E. J. van Leeuwen. Network sparsification for Steiner problems on planar and bounded-genus graphs. *CoRR*, abs/1306.6593, 2013. Version 1.

[30] M. Pilipczuk, S. Siebertz, and S. Toruńczyk. Parameterized circuit complexity of model-checking on sparse structures. In *LICS 2018*, pages 789–798, 2018.

[31] M. Pilipczuk and M. Wrochna. On space efficiency of algorithms working on structural decompositions of graphs. *ACM Trans. on Computation Theory (TOCT)*, 9(4):18:1–18:36, 2018.

[32] N. Robertson and P. D. Seymour. Graph minors. XVI. Excluding a non-planar graph. *Journal of Combinatorial Theory, Series B*, 89(1):43–76, 2003.

[33] J. van den Heuvel, P. Ossona de Mendez, D. Quiroz, R. Rabinovich, S. Siebertz. On the generalised colouring numbers of graphs that exclude a fixed minor. *European Journal of Combinatorics*, 66:129–144, 2017.

[34] X. Zhu. Colouring graphs with bounded generalized colouring number. *Discrete Mathematics*, 309(18):5562–5568, 2009.