Computing Excluded Minors

Isolde Adler Martin Grohe
Institut für Informatik,
Humboldt Universität zu Berlin
{adler,grohe}@informatik.hu-berlin.de

Stephan Kreutzer
Oxford University
Computing Laboratory
kreutzer@comlab.ox.ac.uk

Abstract

By Robertson and Seymour's graph minor theorem, every minor ideal can be characterised by a finite family of excluded minors. (A *minor ideal* is a class of graphs closed under taking minors.) We study algorithms for computing excluded minor characterisations of minor ideals. We propose a general method for obtaining such algorithms, which is based on definability in monadic second-order logic and the decidability of the monadic second-order theory of trees. A straightforward application of our method yields algorithms that, for a given k, compute excluded minor characterisations for the minor ideal $\mathcal{T}_k$ of all graphs of tree width at most k, the minor ideal $\mathcal{B}_k$ of all graphs of branch width at most k, and the minor ideal $\mathcal{G}_k$ of all graphs of genus at most k.

Our main results are concerned with constructions of new minor ideals from given ones. Answering a question that goes back to Fellows and Langston [11], we prove that there is an algorithm that, given excluded minor characterisations of two minor ideals $\mathcal{C}$ and $\mathcal{D}$, computes such a characterisation for the ideal $\mathcal{C} \cup \mathcal{D}$. Furthermore, we obtain an algorithm for computing an excluded minor characterisation for the class of all apex graphs over a minor ideal $\mathcal{C}$, given an excluded minor characterisation for $\mathcal{C}$. (An apex graph over $\mathcal{C}$ is a graph G from which one vertex can be removed to obtain a graph in $\mathcal{C}$.) A corollary of this result is a uniform ftp-algorithm for the "distance k from planarity" problem.

1 Introduction

One of the most stunning consequences of Robertson and Seymour's graph minor theory is the algorithmic result stating that every minor ideal has a cubic time membership test. Here a *minor ideal* is a class of graphs closed under taking minors. A puzzling feature of this result is that it merely states the existence of an algorithm deciding membership in a minor ideal, without giving an explicit algorithm. For example, the class of all graphs that can knotlessly[1] be embedded into 3-space is a minor ideal and hence has a cubic time membership test by Robertson and Seymour's theorem, but to the best of our knowledge still no explicit membership test for this minor ideal is known. Let us analyse how this comes about. The cubic time membership test for minor ideals is the direct consequence of two other results due to Robertson and Seymour. The first is the well-known graph minor theorem [22], stating that every minor ideal is characterised by a finite family of excluded minors, that is, for every minor ideal $\mathcal{C}$ there are an $n \geq 1$ and graphs $H_1, \ldots, H_n$ such that a graph G belongs to $\mathcal{C}$ if and only if none of the graphs $H_1, \ldots, H_n$ is a minor of G. We call $\{H_1, \ldots, H_n\}$ an *excluded minor characterisation* of $\mathcal{C}$.[2] The second result states that for every graph H there is a cubic time algorithm deciding whether a given graph G has H as a minor [21]. Let us call such an algorithm an H*-minor test*. The algorithmic result is perfectly constructive; for every H, Robertson and Seymour construct an explicit cubic H-minor test. However, the graph minor theorem is not constructive; in general, we do not know how to obtain an excluded minor characterisation of a given minor ideal. For example, we do not know such a characterisation for the minor ideal of knotlessly embeddable graphs. Fellows and Langston [12] observed that there is no algorithm that, given a Turing machine which is a membership test for a minor ideal, computes an excluded minor characterisation of the ideal. Later, Courcelle et al. [5] showed that there is no algorithm that, given a sentence of monadic second-order logic which defines a minor ideal, computes an excluded minor characterisation of the ideal.

Fellows and Langston [9, 10, 11, 12] were the first to study these algorithmic issues related to the graph minor theorem (also see [15, 6]). They devised several strategies to overcome the "non-constructiveness" of algorithms based on the minor theorem. The most obvious, of course, is to compute excluded minor char-

[1] A *knotless embedding* of a graph into 3-space is an embedding where no cycle of the graph is non-trivially knotted. The class of knotlessly embeddable graphs only serves as examples in this introduction and will play no further role in the paper.

[2] Instead of *excluded minor characterisation*, the term *obstruction set* is often used in the literature.

acterisations of the classes in question. In [11], Fellows and Langston proposed a general method for computing excluded minor characterisations. It is based on a generalisation of the Myhill-Nerode theorem of formal language theory to certain "graph languages". Here, we propose a similar method that, instead of on Myhill-Nerode congruences, is based on definability in monadic second order logic and generalisations of the Büchi-Elgot-Trakhtenbrot theorem to graphs [4, 23]. Using our method, we can easily show that there are algorithms that, for a given k, compute excluded minor characterisations of the minor ideal $\mathcal{T}_k$ of all graphs of tree width at most k, the minor ideal $\mathcal{B}_k$ of all graphs of branch width at most k, and the minor ideal $\mathcal{G}_k$ of all graphs of genus at most k. These results are not new: The computability of excluded minor characterisations of the families $\mathcal{T}_k$ of bounded tree width follows easily from [16], and for the families $\mathcal{B}_k$ of bounded branch width it follows from [13] (and is also mentioned, without proof, in [2]). For the families $\mathcal{G}_k$ of bounded genus, it follows from [24] or a combination of [11] and [25] (cf. [18], p.222).

Our main results are concerned with constructions of new minor ideals from given ones. Answering a question that goes back to Fellows and Langston [11] (and later was also asked by Courcelle et al. [5]), we prove that there is an algorithm that, given excluded minor characterisations for two minor ideals $\mathcal{C}$ and $\mathcal{D}$, computes an excluded minor characterisation of the minor ideal $\mathcal{C} \cup \mathcal{D}$. (Note that it is trivial to obtain such a characterisation for the minor ideal $\mathcal{C} \cap \mathcal{D}$.) Fellows and Langston [11] observed that an algorithm computing an excluded minor characterisation of $\mathcal{C} \cup \mathcal{D}$ exists if an effective bound on the tree width of the excluded minors for $\mathcal{C} \cup \mathcal{D}$ could be obtained, but left open the problem of obtaining such a bound. Using a deep result due to Robertson and Seymour [21, 19], we give such a bound. Cattell et al. [3] had earlier established the special case of our result where one of the ideals $\mathcal{C}, \mathcal{D}$ is of bounded pathwidth.

We say that a graph G is an *apex graph* over a class $\mathcal{C}$ of graphs if there is a vertex $v \in V(G)$ such that the graph $G \setminus \{v\}$ obtained from G by deleting v belongs to $\mathcal{C}$. Let us denote the class of all apex graphs over $\mathcal{C}$ by $\mathcal{C}^{\text{apex}}$. Observe that if $\mathcal{C}$ is a minor ideal then so is $\mathcal{C}^{\text{apex}}$. We show that there is an algorithm that, given an excluded minor characterisation of a minor ideal $\mathcal{C}$, computes an excluded minor characterisation of $\mathcal{C}^{\text{apex}}$. Iterating the algorithm, of course we can also compute excluded minor characterisations of the classes of all graphs from which we can delete a fixed number k of vertices to obtain a graph in $\mathcal{C}$.

As a corollary, we obtain a uniform fixed-parameter tractable algorithm for the problem of deciding whether a graph has "distance k" from the class of planar graphs or from any other minor ideal. More precisely, for every minor ideal $\mathcal{C}$ we obtain an algorithm that, given a graph G and a nonnegative integer k, decides in time $f(k) \cdot n^3$ if there is a set X of at most k vertices of G such that $G \setminus X \in \mathcal{C}$. Here f is some computable function on the integers. Note that it is an immediate consequence of the graph minor theorem that for every fixed k there is a cubic time algorithm that decides if a given graph has distance at most k from planarity; the point here is the *uniformity*, which means that we have one algorithm which works for all k. As was pointed out to us by Mike Fellows, the result can also be derived by a technique from [12] called *effectivization by self reduction* (cf. [6], Sec. 7.9.1). For planar graphs, it has independently been obtained by Marx and Schlotter [17].

2 Preliminaries

In this section we briefly recall basic notions from graph minors and logic and prove some basic lemmas.

2.1 Graphs, Minors, and Minor Ideals

Graphs in this paper are finite, simple, and undirected. We denote the vertex set of a graph G by $V(G)$ and its edge set by $E(G)$. We always assume that $V(G) \cap E(G) = \varnothing$. The degree of a vertex v in a graph G is denoted by $\deg_G(v)$, and the maximum degree of G is denoted by $\Delta(G)$.

A *tree decomposition* of a graph G is a pair $\mathcal{T} := (T, B)$, where T is a tree and B is a mapping that associates with every node $t \in V(T)$ a set $B_t \subseteq V(G)$ such that $G = \bigcup_{t \in V(T)} G[B_t]$, and for every $v \in V(G)$ the set $B^{-1}(v) = \{t \in V(T) : v \in B(t)\}$ is connected in T. Here, $G[B_t]$ denotes the subgraph of G induced by B_t. The sets B_t, for $t \in V(T)$, are called the *bags* of the decomposition $\mathcal{T}$. The *width* width(D) of a tree decomposition D is defined as $\max\{|B_t| - 1 : t \in V(T)\}$ and the *tree width* tw(G) of a graph G is defined as $\min\{\text{width}(\mathcal{T}) : \mathcal{T} \text{ tree decomposition of } G\}$. Finally, a class $\mathcal{C}$ of graphs has *bounded tree width*, if there is a constant $k \in \mathbb{N}$ such that tw$(G) \leq k$ for all $G \in \mathcal{C}$. We denote the class of graphs of tree width $\leq k$ by $\mathcal{T}_k$.

For the rest of this paper we assume that the tree T underlying a tree decomposition (T, D) is sub-cubic, i.e. a tree of maximum degree at most 3. This is w.l.o.g. as any tree decomposition can easily be converted into one where the underlying tree is sub-cubic.

A graph H is a *minor* of G, in terms $H \preccurlyeq G$, if for each $v \in V(H)$ there is a non-empty connected subgraph $T_v \subseteq G$ such that for all $u, v \in V(H)$, if $u \neq v$ then $V(T_v) \cap V(T_u) = \varnothing$ and if there is an edge $e \in E(H)$ between u and v then there is an edge

$\mathbf{i}_e \in \mathbf{E}(\mathbf{G})$ with one endpoint in $\mathbf{T}_u$ and one endpoint in $\mathbf{T}_v$. The subgraph $\bigcup_{v \in V(H)} \mathbf{T}_v \cup \bigcup_{e \in E(H)} \mathbf{i}_e$ of $\mathbf{G}$ is called a *model* of $\mathbf{H}$ in $\mathbf{G}$ and the function taking $\mathbf{v} \in \mathbf{V}(\mathbf{H})$ to $\mathbf{T}_v$ and $\mathbf{e} \in \mathbf{E}(\mathbf{H})$ to $\mathbf{i}_e \in \mathbf{E}(\mathbf{G})$ is called the corresponding *minor map*. By a *minimal model* of $\mathbf{H}$ in $\mathbf{G}$ we mean a model where all $\mathbf{T}_v$ are trees with $\deg_H(\mathbf{v})$ many leaves. Clearly, if $\mathbf{H} \preccurlyeq \mathbf{G}$, then there is a minimal model of $\mathbf{H}$ in $\mathbf{G}$.

A *minor ideal* is a class $\mathcal{C}$ of graphs that is downward closed under the minor relation $\preccurlyeq$, that is, if $\mathbf{G} \in \mathcal{C}$ and $\mathbf{H} \preccurlyeq \mathbf{G}$, then $\mathbf{H} \in \mathcal{C}$.

Definition 2.1. *Let $\mathcal{C}$ be a minor ideal. A set $\mathcal{F}$ of finite graphs is a set of excluded minors for $\mathcal{C}$ if*

- $\mathbf{H} \not\preccurlyeq \mathbf{G}$ *for all* $\mathbf{G} \in \mathcal{C}$ *and* $\mathbf{H} \in \mathcal{F}$
- *for every* $\mathbf{G} \notin \mathcal{C}$ *there is a* $\mathbf{H} \in \mathcal{F}$ *with* $\mathbf{H} \preccurlyeq \mathbf{G}$

We define an ordering $\leq_{\mathcal{C}}$ among the sets of excluded minors for $\mathcal{C}$ where $\mathcal{F}_1 := \{\mathbf{F}_1, \dots, \mathbf{F}_k\} \leq_{\mathcal{C}} \mathcal{F}_2 := \{\mathbf{H}_1, \dots, \mathbf{H}_l\}$ if either $\mathbf{f}_1 := \sum_{i=1}^{k} |\mathbf{V}(\mathbf{F}_i)| < \mathbf{f}_2 := \sum_{i=1}^{l} |\mathbf{V}(\mathbf{H}_i)|$ or $\mathbf{f}_1 = \mathbf{f}_2$ and $\sum_{i=1}^{k} |\mathbf{E}(\mathbf{F}_i)| < \sum_{i=1}^{l} |\mathbf{E}(\mathbf{H}_i)|$.

A set $\mathcal{F}$ of excluded minors for $\mathcal{C}$ is minimal*, if it is minimal with respect to $\leq_{\mathcal{C}}$.*

Theorem 2.1. (Graph Minor Theorem [22]) *Every minor ideal has a finite set of excluded minors.*

Lemma 2.1. *Up to isomorphism, there is only one minimal set of excluded minors for any minor ideal $\mathcal{C}$.*

We omit the straightforward proof.

For a minor ideal $\mathcal{C}$ we denote the minimal set of excluded minors by $\mathcal{F}(\mathcal{C})$.

Lemma 2.2. *Let $\mathbf{k} > 0$, and let $\mathbf{H} \preccurlyeq \mathbf{G}$ with $|\mathbf{V}(\mathbf{H})| = \mathbf{k}$. Then any minimal model $\mathbf{G}'$ of $\mathbf{H}$ in $\mathbf{G}$ has $\operatorname{tw}(\mathbf{G}') \leq \mathbf{k}^2 + 1$.*

We omit the straightforward proof.

Corollary 2.2. *For all graphs $\mathbf{H}, \mathbf{G}$, if $\mathbf{H} \preccurlyeq \mathbf{G}$ then there is a subgraph $\mathbf{G}'$ of $\mathbf{G}$ with $\mathbf{H} \preccurlyeq \mathbf{G}'$ satisfying $\operatorname{tw}(\mathbf{G}') \leq |\mathbf{V}(\mathbf{H})|^2 + 1$.*

Observe that by the graph minor theorem and Corollary 2.2, for every minor ideal $\mathcal{C}$ there is a $\mathbf{k}$ such that every graph $\mathbf{G} \notin \mathcal{C}$ has a subgraph $\mathbf{G}' \subseteq \mathbf{G}$ such that $\mathbf{G}' \notin \mathcal{C}$ and $\operatorname{tw}(\mathbf{G}') \leq \mathbf{k}$.

Definition 2.2. *Let $\mathcal{C}$ be a minor ideal. The least $\mathbf{k}$ such that for all $\mathbf{G} \notin \mathcal{C}$ there is a $\mathbf{G}' \subseteq \mathbf{G}$ with $\mathbf{G}' \notin \mathcal{C}$ and $\operatorname{tw}(\mathbf{G}') \leq \mathbf{k}$ is called the* width *of $\mathcal{C}$ and denoted by* $\operatorname{width}(\mathcal{C})$.

Although every minor ideal has bounded width, it is not clear how to compute an upper bound for the width of a class. An upper bound follows from Lemma 2.2 if we are given the excluded minors. Otherwise, it is far from clear how to compute the width and indeed much of the present paper will be spent on proving upper bounds for the width of various minor ideals.

2.2 Monadic Second-Order Logic We briefly recall monadic second-order logic (see e.g. [7]). A *signature* is a finite set $\boldsymbol{\tau} := \{\mathbf{R}_1, \dots, \mathbf{R}_k\}$ of relation symbols $\mathbf{R}_i$, each equipped with an arity $ar(\mathbf{R}_i)$. A $\boldsymbol{\tau}$-structure is tuple $\mathfrak{A} := (\mathbf{U}, (\mathbf{R})^{\mathfrak{A}}_{R \in \tau})$ consisting of a finite universe $\mathbf{U}$ and for each $\mathbf{R} \in \boldsymbol{\tau}$ a relation $\mathbf{R}^{\mathfrak{A}} \subseteq \mathbf{U}^{ar(R)}$. With every graph $\mathbf{G}$ we associate a relational structure $\mathfrak{G} := (\mathbf{V}(\mathbf{G}) \mathbin{\dot\cup} \mathbf{E}(\mathbf{G}), \mathbf{V}^{\mathfrak{G}}, \mathbf{E}^{\mathfrak{G}}, \mathbf{I}^{\mathfrak{G}})$, over the signature $\boldsymbol{\tau}_{\text{graph}} := \{\mathbf{V}, \mathbf{E}, \mathbf{I}\}$, where $\mathbf{V}^{\mathfrak{G}} := \mathbf{V}(\mathbf{G})$, $\mathbf{E}^{\mathfrak{G}} := \mathbf{E}(\mathbf{G})$, and $\mathbf{I}^{\mathfrak{G}} := \{(\mathbf{v}, \mathbf{e}) : \mathbf{e} \in \mathbf{E}(\mathbf{G}), \mathbf{v} \in \mathbf{e} \cap \mathbf{V}(\mathbf{G})\}$. In the sequel, we do not distinguish notationally between a graph and the associated relational structure.

Monadic Second-Order Logic (MSO) is the extension of first-order logic by quantification over sets. E.g. formulas of MSO over $\boldsymbol{\tau}_{\text{graph}}$ are built up inductively from atoms $\mathbf{x} \in \mathbf{V}, \mathbf{e} \in \mathbf{E}$, and $(\mathbf{x}, \mathbf{e}) \in \mathbf{I}$ using Boolean connectives, first-order quantifiers $\exists \mathbf{x}, \forall \mathbf{x}$ and second-order quantifiers $\exists \mathbf{X}, \forall \mathbf{X}$. Here, $\mathbf{x}, \mathbf{y}, \mathbf{e}, \dots$ are first-order variables interpreted by single elements form $\mathbf{V}(\mathbf{G}) \cup \mathbf{E}(\mathbf{G})$ and $\mathbf{X}, \mathbf{Y}, \dots$ are second-order variables interpreted by sets of elements. As the universe of structures $\mathfrak{G}$ considered here contains edges and vertices, we can use set quantification over sets of edges and sets of vertices. We illustrate the definition by an example.

Example. Let $\mathbf{H}$ be a finite graph with $\mathbf{V}(\mathbf{H}) := \{\mathbf{v}_1, \dots, \mathbf{v}_k\}$. Let $\boldsymbol{\phi}_H$ be the MSO-sentence $\exists \mathbf{X}_1 \dots \exists \mathbf{X}_k \boldsymbol{\psi}$ where $\boldsymbol{\psi}$ is defined as

$$\left(\begin{array}{l} \bigwedge_{i \neq j} \mathbf{X}_i \cap \mathbf{X}_j = \varnothing \wedge \bigwedge_i (\varnothing \subsetneq \mathbf{X}_i \subseteq \mathbf{V} \wedge \\ \qquad\qquad\qquad\qquad \text{"}\mathbf{X}_i \text{ is connected"}) \wedge \\ \bigwedge_{\{v_i, v_j\} \in E(H)} \exists \mathbf{x}_i \in \mathbf{X}_i \exists \mathbf{x}_j \in \mathbf{X}_j \exists \mathbf{z} ((\mathbf{x}_i, \mathbf{z}) \in \mathbf{I} \wedge \\ \qquad\qquad\qquad (\mathbf{x}_j, \mathbf{z}) \in \mathbf{I}), \end{array} \right)$$

where "$\mathbf{X}_i$ is connected" can easily be expressed in MSO. The formula is true in a graph $\mathbf{G}$ if there is a minor map $\boldsymbol{\phi}$ from $\mathbf{H}$ into $\mathbf{G}$, where the sets $\mathbf{X}_i$ represent the images $\boldsymbol{\phi}(\mathbf{v}_i)$.

Hence, for every graph $\mathbf{G}$, $\mathbf{G} \models \boldsymbol{\phi}_H$ if, and only if, $\mathbf{H} \preccurlyeq \mathbf{G}$. If $\mathcal{H}$ is a finite set of graphs then $\boldsymbol{\phi}_{\mathcal{H}} := \bigvee_{H \in \mathcal{H}} \boldsymbol{\phi}_H$ is true in a graph $\mathbf{G}$ if, and only if, at least one graph of $\mathcal{H}$ is a minor of $\mathbf{G}$. Hence, $\neg \boldsymbol{\phi}_{\mathcal{H}}$ defines the minor ideal $\mathcal{F}(\mathcal{H})$.

THEOREM 2.3. (SEESE [23]) *For each* $\mathbf{k} \in \mathbb{N}$*, it is decidable for a given* MSO*-formula whether it is true in a graph of tree width at most* $\mathbf{k}$.

2.3 Extending Structures by Tree Decompositions To define certain properties of graphs in MSO, it is sometimes more convenient if the MSO-formulas can explicitly refer to tree decompositions of graphs. We will therefore equip graphs with a tree decomposition, i.e. consider structures consisting of a graph and a tree decomposition of the graph. Recall that we only consider tree decompositions in this paper where the underlying graph is sub-cubic.

DEFINITION 2.3. (TREE-DEC EXPANSIONS) *Let* $\mathbf{G}$ *be a graph and* $\mathcal{T} := (\mathbf{T}, \mathbf{B})$ *be a tree decomposition of* $\mathbf{G}$. *Let* $\tau_{\text{ex}} := \{\mathbf{V}, \mathbf{E}, \mathbf{I}, \mathbf{T}, \mathbf{E}_T, \mathbf{I}_T, \mathbf{B}\}$*, where* $\mathbf{V}, \mathbf{E}, \mathbf{T}, \mathbf{E}_T$ *are unary relation symbols and* $\mathbf{I}, \mathbf{I}_T, \mathbf{B}$ *are binary.*

The tree-dec expansion $\text{TreeExp}(\mathbf{G}, \mathcal{T})$ *of* $\mathbf{G}$ *and* $\mathcal{T}$ *is the* τ_{ex}*-structure*

$$(\mathbf{V}(\mathbf{G}) \dot\cup \mathbf{E}(\mathbf{G}) \dot\cup \mathbf{V}(\mathbf{T}) \dot\cup \mathbf{E}(\mathbf{T}), \mathbf{V}^{\mathfrak{G}}, \mathbf{E}^{\mathfrak{G}}, \mathbf{I}^{\mathfrak{G}}, \mathbf{T}^{\mathfrak{G}}, \mathbf{E}_T^{\mathfrak{G}}, \mathbf{I}_T^{\mathfrak{G}}, \mathbf{B}^{\mathfrak{G}}),$$

where $\mathbf{V}^{\mathfrak{G}}, \mathbf{E}^{\mathfrak{G}}, \mathbf{I}^{\mathfrak{G}}$ *are as above and* $\mathbf{T}^{\mathfrak{G}} := \mathbf{V}(\mathbf{T})$*,* $\mathbf{E}_T^{\mathfrak{G}} := \mathbf{E}(\mathbf{T})$*,* $\mathbf{I}_T^{\mathfrak{G}} := \{(\mathbf{t}, \mathbf{e}) : \mathbf{e} \in \mathbf{E}(\mathbf{T}), \mathbf{t} \in \mathbf{e} \cap \mathbf{V}(\mathbf{T})\}$ *and* $\mathbf{B}^{\mathfrak{G}} := \{(\mathbf{t}, \mathbf{u}) : \mathbf{t} \in \mathbf{V}(\mathbf{T}), \mathbf{u} \in \mathbf{B}_t \cap \mathbf{V}(\mathbf{G})\}$.

Intuitively, we think of tree-dec expansions $\mathfrak{G}$ as of two-sorted structures consisting of the graph $\mathbf{G}$ and the tree decomposition $\mathcal{T}$. The proof of the following lemma is straightforward.

LEMMA 2.3. *(i) If* $\mathbf{G}$ *is a graph and* $\mathcal{T}$ *a tree decomposition of* $\mathbf{G}$ *of width* $\mathbf{k}$*, then the tree width of* $\text{TreeExp}(\mathbf{G}, \mathcal{T})$ *is at most* $\mathbf{k} + 2$.

(ii) There is an MSO*-sentence* ϕ_{ex} *so that for any* τ_{ex}*-structure* $\mathfrak{G}$*,* $\mathfrak{G} \models \phi_{\text{ex}}$ *if, and only if,* $\mathfrak{G}$ *is a Tree-Dec Expansion of a graph* $\mathbf{G}$.

We denote the class of structures $\text{TreeExp}(\mathbf{G}, \mathcal{T})$, for $\mathbf{G} \in \mathcal{T}_k$ and $\mathcal{T}$ a tree decomposition of $\mathbf{G}$ of width at most $\mathbf{k}$, by $\text{TreeExp}(\mathcal{T}_k)$. It follows from Lemma 2.3 that $\text{TreeExp}(\mathcal{T}_k)$ is a class of structures of tree width at most $\mathbf{k} + 2$.

COROLLARY 2.4. *For all* $\mathbf{k} \in \mathbb{N}$*, it is decidable whether a given* MSO*-formula is satisfied in some* $\mathfrak{G} \in \text{TreeExp}(\mathcal{T}_k)$.

2.4 Dynamic Programming and MSO-definability MSO is a powerful logic on graphs and many interesting properties of graphs can be expressed naturally in this logic. In fact, for various problems, their mathematical formalisations only use set and individual quantification and therefore *are* already MSO-definitions. However, sometimes algorithms for a certain problem on graph classes of bounded tree width are already known and we don't want to prove an equivalent MSO-definition.

In this section we therefore develop a strong relationship between dynamic programming algorithms as they are typically employed on graph classes of bounded tree width and MSO-definability.

A *typical linear time dynamic programming algorithm* on classes of graphs of bounded tree width, *tlp-algorithm* for short, is an algorithm that, given a graph $\mathbf{G}$ and a tree decomposition $\mathcal{T}$ of $\mathbf{G}$ of width at most $\mathbf{k}$ as input, computes a data record for each node of $\mathcal{T}$, starting at the leaves and proceeding bottom up, so that

- whether or not the algorithm accepts $\mathbf{G}$ only depends on the record computed for the root of $\mathcal{T}$
- there are only constantly many possible data records
- the record computed for a leaf $\mathbf{t}$ of $\mathcal{T}$ only depends on the isomorphism type of $\mathbf{B}_t$
- the record computed for an inner node $\mathbf{t}$ can be computed in constant time from the isomorphism type of $\mathbf{B}_t$ and the records computed for the children.

Many algorithms developed for graphs of small tree width are actually of this form, e.g. standard algorithms for 3-colourability, dominating set, vertex cover and many more. For instance, for the 3-colourability problem, the data record computed for a node $\mathbf{t} \in \mathbf{V}(\mathbf{T})$ of a tree decomposition $(\mathbf{T}, \mathbf{B})$ of a graph $\mathbf{G}$ consists of all possible 3-colourings of $\mathbf{B}_t$ that can be extended to a 3-colouring of the subgraph of $\mathbf{G}$ induced by the vertices contained in bags in the subtree of $\mathbf{T}$ rooted at $\mathbf{t}$. Another example is the algorithm given by Bodlaender and Kloks in [1] which, given a graph $\mathbf{G}$ and a tree decomposition of $\mathbf{G}$ of (possibly non-optimal) width $\mathbf{w}$, and a number $\mathbf{k}$ decides whether $\mathbf{G}$ has tree width at most $\mathbf{k}$ in time exponential in $\mathbf{w}$ and linear in the size of $\mathbf{G}$.

LEMMA 2.4. *Let* $\mathbf{A}$ *be a tlp-algorithm. There is an* MSO *sentence* ϕ_A *such that given a structure* $\text{TreeExp}(\mathbf{G}, \mathcal{T})$ *of a graph* $\mathbf{G}$ *and a tree decomposition* $\mathcal{T}$ *of* $\mathbf{G}$*,* $\mathfrak{G}$ *satisfies* ϕ_A *if, and only if,* $\mathbf{A}$ *accepts the pair* $(\mathbf{G}, \mathcal{T})$.

Proof. Let $\mathbf{C} := \{\mathbf{C}_1, \ldots, \mathbf{C}_l\}$ be the constant set of possible data records used by $\mathbf{A}$. The sentence ϕ_A is defined as $\phi_A := \exists \mathbf{C}_1 \ldots \exists \mathbf{C}_l \psi_A$, where ψ_A is a first-order sentence stating that every node of $\mathcal{T}$ is contained in exactly one set $\mathbf{C}_i$ and that the colours $\mathbf{C}_i$ assigned to the nodes of $\mathcal{T}$ are consistent with the data records computed by $\mathbf{A}$ on $\mathcal{T}$. The latter can be defined in FO, as the data record, and hence the colour of a node $\mathbf{t}$, only depends on the isomorphism type of $\mathbf{B}_t$ and the data records/colours of the successors. Hence, this is finite information that can be coded directly in the first-order formula. □

Recall from above that the Bodlaender, Kloks algorithm from [1] is a tlp-algorithm. Hence, we immediately obtain the following corollary.

COROLLARY 2.5. *For every $\mathbf{k} \in \mathbb{N}$ there is an* MSO-*sentence $\boldsymbol{\phi}_k$ such that a structure* TreeExp$(\mathbf{G}, \mathcal{T}) \in$ TreeExp$(\mathcal{T}_l)$*, for some $\mathbf{l} \geq \mathbf{k}$, is a model of $\boldsymbol{\phi}_k$ if, and only if,* tw$(\mathbf{G}) = \mathbf{k}$.

The following strengthening of Seese's Theorem 2.3 will be useful later on.

THEOREM 2.6. *For every $\mathbf{k}$ it is decidable whether a given* MSO-*formula is satisfied by a graph of tree width exactly $\mathbf{k}$.*

Proof. Let $\boldsymbol{\phi}$ be an MSO-sentence and $\mathbf{k} \in \mathbb{N}$. To decide if $\boldsymbol{\phi}$ has a model of tree width exactly $\mathbf{k}$ we test whether $\boldsymbol{\phi} \wedge \boldsymbol{\phi}_{\text{ex}} \wedge \boldsymbol{\phi}_k$ is satisfiable in TreeExp$(\mathcal{T}_k)$. Here, $\boldsymbol{\phi}_{\text{ex}}$ is the MSO-sentence defined in Lemma 2.3 and $\boldsymbol{\phi}_k$ is the sentence defined in Corollary 2.5. By Corollary 2.4, this is decidable. □

We remark that by the proof of Courcelle's theorem (see [4]), every MSO-definable problem can be solved by a tlp-algorithm on graph classes of bounded tree width.

COROLLARY 2.7. *Let $\mathcal{C}$ be a class of bounded tree width. A problem can be solved by a tlp-algorithm on $\mathcal{C}$ if, and only if, it is* MSO-*definable on the class* TreeExp$(\mathcal{C})$.

3 Computing Excluded Minors

In this section we develop a general machinery for computing minimal sets of excluded minors that will be applied to various minor ideals in Sections 4 and 5.

DEFINITION 3.1. *A class $\mathcal{C}$ of graphs is* layerwise MSO-definable*, if for every $\mathbf{k} \in \mathbb{N}$ we can compute an* MSO-*formula $\boldsymbol{\phi}_k$ defining $\mathcal{C} \cap \mathcal{T}_k$ in* TreeExp$(\mathcal{T}_k)$.

We will show next that if $\mathcal{C}$ is a minor ideal that is layerwise MSO-definable and for which we are given an upper bound on its width, then we can compute the minimal set of excluded minors for $\mathcal{C}$. We will state this result in two equivalent forms, a logical form that we actually prove (cf. Lemma 3.1) and an equivalent algorithmic form.

LEMMA 3.1. *The minimal set of excluded minors is computable for every layerwise* MSO-*definable class of graphs, provided we are given an upper bound on its width.*

Formally, there is an algorithm that, given $\mathbf{w} \in \mathbb{N}$ and (an algorithm computing) a computable function $\mathbf{g} : \mathbb{N} \to$ MSO so that there is a minor ideal $\mathcal{C}$ with

- *for every $\mathbf{k} \in \mathbb{N}$, $\boldsymbol{\phi}_k := \mathbf{g}(\mathbf{k})$ defines $\mathcal{C} \cap \mathcal{T}_k$ in* TreeExp$(\mathcal{T}_k)$ *and*
- *the width of $\mathcal{C}$ is at most $\mathbf{w}$,*

then the algorithm computes $\mathcal{F}(\mathcal{C})$.

Proof. We first show that for any given set $\mathcal{F} := \{\mathbf{F}_1, \ldots, \mathbf{F}_m\}$ we can decide if $\mathcal{F}$ is a set of excluded minors for $\mathcal{C}$. For this, we show that the following two statements are decidable.

(1) Is $(\mathbf{G} \in \mathcal{C} \wedge \exists\, \mathbf{i} : \mathbf{F}_i \preccurlyeq \mathbf{G})$ unsatisfiable?

(2) Is $(\mathbf{G} \notin \mathcal{C} \wedge \forall\, \mathbf{i} : \mathbf{F}_i \not\preccurlyeq \mathbf{G})$ unsatisfiable?

Towards (1), it suffices to check whether $\mathbf{F}_i \in \mathcal{C}$ for some $1 \leq \mathbf{i} \leq \mathbf{m}$. However, as there are no assumptions on the membership problems of $\mathcal{C}$, we give a general method that works for all classes $\mathcal{C}$ satisfying the prerequisites of the lemma. Suppose there is a $\mathbf{G} \in \mathcal{C}$ with a minor $\mathbf{F}_i \preccurlyeq \mathbf{G}$. By Lemma 2.2, there is $\mathbf{G}' \subseteq \mathbf{G}$ of tree width at most $|\mathbf{V}(\mathbf{F}_i)|^2 + 1$ with $\mathbf{F}_i \preccurlyeq \mathbf{G}'$. As $\mathcal{C}$ is closed under taking subgraphs, $\mathbf{G}' \in \mathcal{C}$. Hence, $(\mathbf{G} \in \mathcal{C} \wedge \exists \mathbf{i} : \mathbf{F}_i \preccurlyeq \mathbf{G})$ is satisfiable if, and only if, it is satisfiable by a graph of tree width at most $\mathbf{k}$, where $\mathbf{k} := \max\{|\mathbf{V}(\mathbf{F}_i)|^2 + 1 : \mathbf{F}_i \in \mathcal{F}\}$.

By assumption, the class $\mathcal{C} \cap \mathcal{T}_k$ is MSO-definable in TreeExp$(\mathcal{T}_k)$ by $\boldsymbol{\phi}_{\mathcal{C}} := \mathbf{g}(\mathbf{k})$. Let $\boldsymbol{\phi}_{\mathcal{F}}$ be the MSO-sentence from Example 2.2 saying that every model of $\boldsymbol{\phi}_{\mathcal{F}}$ has at least one graph $\mathbf{F} \in \mathcal{F}$ as a minor. Then there is a $\mathbf{G} \in \mathcal{C}$ with $\mathbf{F}_i \preccurlyeq \mathbf{G}$, for some $1 \leq \mathbf{i} \leq \mathbf{m}$ if, and only if, the sentence $\boldsymbol{\phi}_{\mathcal{C}} \wedge \boldsymbol{\phi}_F$ is satisfiable in TreeExp$(\mathcal{T}_k)$ – which is decidable by Corollary 2.4.

Towards (2), suppose there is a $\mathbf{G} \notin \mathcal{C}$ and $\mathbf{F}_i \not\preccurlyeq \mathbf{G}$ for all $1 \leq \mathbf{i} \leq \mathbf{m}$. As the width of $\mathcal{C}$ is bounded by $\mathbf{w}$, $\mathbf{G}$ has a subgraph $\mathbf{G}'$ of tree width tw$(\mathbf{G}') \leq \mathbf{w}$ so that $\mathbf{G} \notin \mathcal{C}$. Clearly, $\mathbf{F}_i \not\preccurlyeq \mathbf{G}'$ for all $1 \leq \mathbf{i} \leq \mathbf{m}$.

Let $\boldsymbol{\phi}_w := \mathbf{g}(\mathbf{w})$ be the MSO-sentence that defines $\mathcal{C} \cap \mathcal{T}_w$ in TreeExp$(\mathcal{T}_w)$. Then there is a graph $\mathbf{G} \notin \mathcal{C}$ of tree width at most $\mathbf{w}$ with no minor $\mathbf{F}_i \in \mathcal{F}$ if, and only if, $\neg\boldsymbol{\phi}_w \wedge \neg\boldsymbol{\phi}_{\mathcal{F}}$ is satisfiable in TreeExp$(\mathcal{T}_w)$. This is decidable by Corollary 2.4.

Hence, we can decide for each finite set $\mathcal{F}$ of graphs if it is a set of excluded minors for $\mathcal{C}$. Enumerating all finite sets $\mathcal{F}$ of graphs in a way respecting the ordering $\leq_{\mathcal{C}}$ of Definition 2.1, the first set $\mathcal{F}$ of excluded minors for $\mathcal{C}$ will be $\mathcal{F}(\mathcal{C})$. □

For the algorithmic variant of the lemma, we say that a class $\mathcal{C}$ of graphs is *layerwise decidable*, if for every $\mathbf{k} \in \mathbb{N}$ we can compute a tlp-algorithm that decides for any graph $\mathbf{G}$ of tree width at most $\mathbf{k}$ if $\mathbf{G} \in \mathcal{C}$.

COROLLARY 3.1. *The set $\mathcal{F}(\mathcal{C})$ of excluded minors is computable for every layerwise decidable class $\mathcal{C}$ of graphs given an upper bound $\mathbf{w}$ of its width.*

Often we can use a slightly simpler form of the previous lemma and corollary.

COROLLARY 3.2. *There is an algorithm that, given an* MSO-*formula $\boldsymbol{\phi}$ and $\mathbf{w} \in \mathbb{N}$, so that $\boldsymbol{\phi}$ defines a minor ideal $\mathcal{C}$ of width at most $\mathbf{w}$, computes the set $\mathcal{F}(\mathcal{C})$.*

Proof. This follows from Lemma 3.1 by setting $\mathbf{g}(\mathbf{k}) := \boldsymbol{\phi}$ for all $\mathbf{k}$. □

4 Computing Excluded Minors for Specific Classes of Graphs

In this section we present some immediate applications of the machinery developed in Section 3 and show that the set of excluded minors are computable for various natural families of minor ideals.

4.1 Bounded Tree Width and Branch-Width

As a first and simple example we show that, given $\mathbf{k} \in \mathbb{N}$, the set of excluded minors is computable for the class of graphs of tree width at most $\mathbf{k}$ and the class of graphs of branch-width at most $\mathbf{k}$. The following theorem is also an immediate consequence of a result due to Lagergren [16] giving an upper bound on the size of the excluded minors for the class $\mathcal{T}_k$.

THEOREM 4.1. ([16]) *There is an algorithm that, given $\mathbf{k} > 1$, computes $\mathcal{F}(\mathcal{T}_k)$.*

Proof. Let $\mathbf{k} > 1$. Clearly, $\mathcal{T}_k$ has width $\mathbf{k}+1$, as every graph $\mathbf{G}$ of tree width $\geq \mathbf{k}+1$ has a subgraph of tree width $\mathbf{k}+1$. Furthermore, by Corollary 2.5, for every $\mathbf{l} > 1$ we can compute an MSO-sentence $\boldsymbol{\phi}_l$ defining $\mathcal{T}_k \cap \mathcal{T}_l$ in TreeExp($\mathcal{T}_l$). Hence, the theorem follows from Lemma 3.1. □

Another type of graph decompositions related to tree width is the notion of *branch-decompositions* and the associated width measure *branch-width*. Whereas a tree decomposition is a decomposition of the vertex set of a graph, a branch-decompositions decomposes its edge set. We refrain from giving details here and refer the reader to [20]. Let $\mathcal{B}_k$ denote the class of graphs of branch-width at most $\mathbf{k}$.

The branch-width and tree width of a graph are within a constant factor but clearly the sets of excluded minors for $\mathcal{T}_k$ and $\mathcal{B}_k$ differ. In [2], Bodlaender and Thilikos give for any fixed $\mathbf{k} \in \mathbb{N}$ a linear-time algorithm for deciding whether a graph has branch-width exactly $\mathbf{k}$. It follows from their result that there is an tlp-algorithm to decide the branch width of a graph given a tree decomposition of potentially non optimal width. Hence, an analogous reasoning as for Theorem 4.1 establishes the following theorem. The theorem has already been stated in [2], and it also follows easily from an upper bound on the size of the excluded minors for matroids of bounded branch width [13].

THEOREM 4.2. ([2, 13]) *There is an algoriomoat*

given $\mathbf{k} >$

study ways to compute the excluded minors for derived classes of graphs given the excluded minors for the base classes.

5.1 Apices For any given class $\mathcal{C}$ of graphs we can construct the class of graphs $\mathbf{G}$ such that there is a vertex $\mathbf{v} \in \mathbf{V}(\mathbf{G})$ for which $\mathbf{G} \setminus \mathbf{v} \in \mathcal{C}$. We call this vertex an *apex* of $\mathbf{G}$ with respect to $\mathcal{C}$. We denote the class of graphs having an apex with respect to $\mathcal{C}$ by $\mathcal{C}^{\text{apex}}$. The aim of this section is to show the following theorem.

THEOREM 5.1. *If $\mathcal{C}$ is a minor ideal whose set of excluded minors is given, then we can compute the excluded minors of $\mathcal{C}^{\text{apex}}$.*

By iterating this construction, we can compute for any $\mathbf{k}$ the set of excluded minors for the class of graphs $\mathbf{G}$ for which there are $\mathbf{k}$ vertices $\mathbf{v}_1, \dots, \mathbf{v}_k$ such that $\mathbf{G} \setminus \{\mathbf{v}_1, \dots, \mathbf{v}_k\} \in \mathcal{C}$.

We now prove Theorem 5.1. Let $\mathcal{C}$ be a minor ideal. As before, we aim at applying Lemma 3.1. Using Example 2.2, it is easily seen that for each $\mathbf{k} \in \mathbb{N}$, $\mathcal{C}^{\text{apex}} \cap \mathcal{T}_k$ is MSO-definable over TreeExp($\mathcal{T}_k$). Hence, it remains to show that $\mathcal{C}^{\text{apex}}$ has effectively bounded width. We first need some preparations.

5.1.1 Walls, Layouts, and Linkages An *elementary wall of height* $\mathbf{h} \geq 1$ is a graph defined as in Figure 1(a). A *wall of height* $\mathbf{h}$ is a subdivision of an elementary wall of height $\mathbf{h}$. The *perimeter* of a wall is the boundary cycle. A wall *in* a graph $\mathbf{G}$ is a wall $\mathbf{W}$ that is a subgraph of $\mathbf{G}$. Note that, up to homeomorphisms, walls have unique embeddings in the sphere. For walls of height 1, this is obvious, and for walls of height $\mathbf{h} \geq 2$ this follows from a well known theorem due to Tutte stating that 3-connected graphs have unique embeddings, because walls of height ≥ 2 are subdivisons of 3-connected graphs.

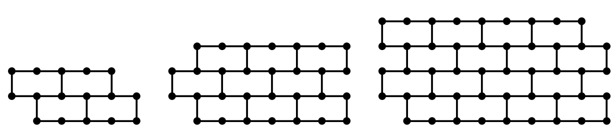

Figure 1: Elementary walls of height 2–4

For a subgraph $\mathbf{D}$ of a graph $\mathbf{G}$, we let $\boldsymbol{\partial}^G\mathbf{D}$ be the set of all vertices of $\mathbf{D}$ that are incident with an edge in $\mathbf{E}(\mathbf{G}) \setminus \mathbf{E}(\mathbf{D})$.

In the following, let $\mathbf{W}$ be a wall of height at least 2 in a graph $\mathbf{G}$, and let $\mathbf{P}$ be the perimeter of $\mathbf{W}$. Let $\mathbf{K}'$ be the unique connected component of $\mathbf{G} \setminus \mathbf{P}$ that contains $\mathbf{W} \setminus \mathbf{P}$. The graph $\mathbf{K} = \mathbf{K}' \cup \mathbf{P}$ is called the *compass* of $\mathbf{W}$ in $\mathbf{G}$. A *layout of* $\mathbf{K}$ *(with respect to the wall* $\mathbf{W}$ *in* $\mathbf{G}$*)* is a family $(\mathbf{C}, \mathbf{D}_1, \dots, \mathbf{D}_m)$ of connected subgraphs of $\mathbf{K}$ such that:

1. $\mathbf{K} = \mathbf{C} \cup \mathbf{D}_1 \cup \dots \cup \mathbf{D}_m$;
2. $\mathbf{W} \subseteq \mathbf{C}$, and there is no separation $(\mathbf{X}, \mathbf{Y})$ of $\mathbf{C}$ of order ≤ 3 with $\mathbf{V}(\mathbf{W}) \subseteq \mathbf{X}$ and $\mathbf{Y} \setminus \mathbf{X} \neq \varnothing$;
3. $\boldsymbol{\partial}^G\mathbf{D}_i \subseteq \mathbf{V}(\mathbf{C})$ for all $\mathbf{i} \in \{1, \dots, \mathbf{m}\}$;
4. $|\boldsymbol{\partial}^G\mathbf{D}_i| \leq 3$ for all $\mathbf{i} \in \{1, \dots, \mathbf{m}\}$;
5. $\boldsymbol{\partial}^G\mathbf{D}_i \neq \boldsymbol{\partial}^G\mathbf{D}_j$ for all $\mathbf{i} \neq \mathbf{j} \in \{1, \dots, \mathbf{m}\}$.

We let $\overline{\mathbf{C}}$ be the graph obtained from $\mathbf{C}$ by adding new vertices $\mathbf{d}_1, \dots, \mathbf{d}_m$ and, for $1 \leq \mathbf{i} \leq \mathbf{m}$, edges between $\mathbf{d}_i$ to the vertices in $\boldsymbol{\partial}^G\mathbf{D}_i$ and edges between all vertices in $\boldsymbol{\partial}^G\mathbf{D}_i$. Hence, for each $\mathbf{i} \in \{1, \dots, \mathbf{m}\}$ the vertex $\mathbf{d}_i$ together with the (at most 3) vertices in $\boldsymbol{\partial}^G\mathbf{D}_i$ form a clique. We call $\overline{\mathbf{C}}$ the *core* of the layout and $\mathbf{D}_1, \dots, \mathbf{D}_m$ its *extensions*. The layout $(\mathbf{C}, \mathbf{D}_1, \dots, \mathbf{D}_m)$ is *flat* if its core $\overline{\mathbf{C}}$ is planar. Note that this implies that the core has an embedding in the plane that extends the "standard planar embedding" of the wall $\mathbf{W}$ (as shown in Figure 1), because the wall $\mathbf{W}$ has a unique embedding into the sphere. We call the wall $\mathbf{W}$ *flat* (in $\mathbf{G}$) if the compass of $\mathbf{W}$ has a flat layout.

The following lemma is (essentially) Lemma (9.8) of [21]. Concerning the uniformity, see the remarks at the end of [21] (on page 109).

TRINITY LEMMA (ROBERTSON AND SEYMOUR [21]). *There are computable functions* $\mathbf{f}, \mathbf{g} : \mathbb{N}^2 \to \mathbb{N}$ *and an algorithm* $\mathbf{A}$ *that, given a graph* $\mathbf{G}$ *and nonnegative integers* $\mathbf{k}, \mathbf{h}$*, computes either*

1. *a tree decomposition of* $\mathbf{G}$ *of width* $\mathbf{f}(\mathbf{k}, \mathbf{h})$*, or*
2. *a* $\mathbf{K}_k$*-minor of* $\mathbf{G}$*, or*
3. *a subset* $\mathbf{X} \subseteq \mathbf{V}(\mathbf{G})$ *with* $|\mathbf{X}| < \binom{k}{2}$*, a wall* $\mathbf{W}$ *of height* $\mathbf{h}$ *in* $\mathbf{G} \setminus \mathbf{X}$*, and a flat layout* $(\mathbf{C}, \mathbf{D}_1, \dots, \mathbf{D}_m)$ *of the compass of* $\mathbf{W}$ *in* $\mathbf{G} \setminus \mathbf{X}$ *such that the tree width of each of the extensions* $\mathbf{D}_1, \dots, \mathbf{D}_m$ *is at most* $\mathbf{f}(\mathbf{k}, \mathbf{h})$.

Furthermore, the running time of the algorithm is bounded by $\mathbf{g}(\mathbf{k}, \mathbf{h}) \cdot |\mathbf{V}(\mathbf{G})|^2$.

A *linkage* $\mathbf{L}$ in a graph $\mathbf{G}$ is a subgraph whose components are paths. For a set $\mathbf{Z} \subseteq \mathbf{V}(\mathbf{G})$, the *effect* of $\mathbf{L}$ on $\mathbf{Z}$ is defined as the partition of $\mathbf{Z} \cap \mathbf{V}(\mathbf{L})$ where two vertices belong to the same class if they belong to the same component of $\mathbf{L}$.

5.1.2 Models in Critical Graphs A graph $\mathbf{G}$ is called *critical* for $\mathcal{C}^{\text{apex}}$, if $\mathbf{G} \notin \mathcal{C}^{\text{apex}}$ but $\mathbf{G}' \in \mathcal{C}^{\text{apex}}$ for every proper subgraph $\mathbf{G}' \subset \mathbf{G}$. Let $\mathcal{F}(\mathcal{C}) := \{\mathbf{H}_1, \dots, \mathbf{H}_m\}$. To show that $\mathcal{C}^{\text{apex}}$ has effectively bounded width we prove that the tree width of its critical graphs is bounded by $\mathbf{f}(\mathbf{H}_1, \dots, \mathbf{H}_m)$ for a computable function $\mathbf{f}$ from finite sets of graphs to natural numbers.

Let $\mathbf{G}$ be a critical graph for $\mathcal{C}^{\text{apex}}$ and let $\mathbf{c} := \max\{|\mathbf{H}_i| : 1 \leq \mathbf{i} \leq \mathbf{m}\}$. We set $\mathbf{k} := \mathbf{c} + 2$ and let $\mathbf{h}$ be "large". We will explain how to choose $\mathbf{h}$ as we go along. For the following arguments, assume that $\mathbf{h}$ is "large enough".

By the Trinity Lemma, there is a computable function $\mathbf{f}(\mathbf{k}, \mathbf{h})$ such that $\mathbf{G}$ either has (1) tree width at most $\mathbf{f}(\mathbf{k}, \mathbf{h})$, or (2) a $\mathbf{K}_k$-minor, or (3) a subset $\mathbf{X} \subseteq \mathbf{V}(\mathbf{G})$ with $|\mathbf{X}| < \binom{k}{2}$, a wall $\mathbf{W}$ of height $\mathbf{h}$ in $\mathbf{G} \setminus \mathbf{X}$, and a flat layout $(\mathbf{C}, \mathbf{D}_1, \dots, \mathbf{D}_m)$ of the compass of $\mathbf{W}$ in $\mathbf{G} \setminus \mathbf{X}$ such that the tree width of each of the extensions $\mathbf{D}_1, \dots, \mathbf{D}_m$ is at most $\mathbf{f}(\mathbf{k}, \mathbf{h})$.

Clearly, Case 2 is impossible, as then $\mathbf{G}$ would have a proper subgraph containing a $\mathbf{K}_{c+1}$ minor. We show next that Case 3 is impossible as well and hence the tree width of $\mathbf{G}$ is bounded by $\mathbf{f}(\mathbf{k}, \mathbf{h})$.

Suppose Case 3 applies. Let $\mathbf{P}$ be the perimeter of the wall $\mathbf{W}$. As $\mathbf{G} \notin \mathcal{C}^{\text{apex}}$, there is a model of some $\mathbf{H}_i$ in $\mathbf{G}$. Let $\mathbf{H}$ be a minimal model of $\mathbf{H}_i$ in $\mathbf{G}$. We use $\mathbf{H}$ both for the minor map and the model of $\mathbf{H}_i$ in $\mathbf{G}$. A vertex $\mathbf{v} \in \mathbf{V}(\mathbf{H}) \cap \mathbf{C}$ is called *important*, if it is in $\mathbf{C} \setminus \mathbf{P}$ and

1. $\mathbf{v}$ has degree at least 3 in $\mathbf{H}$, or
2. $\mathbf{v} \in \mathbf{H}(\mathbf{e})$ for some $\mathbf{e} \in \mathbf{E}(\mathbf{H}_i)$, or
3. $\mathbf{v}$ has a neighbour in $\mathbf{X} \cap \mathbf{V}(\mathbf{H})$, or
4. there is a path in $\mathbf{H}$ from $\mathbf{v}$ to a vertex $\mathbf{u} \in \mathbf{V}(\mathbf{H})$ in an extension $\mathbf{D}$, so that the internal vertices of the path are all contained in $\mathbf{D}$ and the degree of $\mathbf{u}$ in $\mathbf{H}$ is at least 3, or $\mathbf{u} \in \mathbf{H}(\mathbf{e})$ for some $\mathbf{e} \in \mathbf{E}(\mathbf{H}_i)$, or there is an edge in $\mathbf{H}$ from $\mathbf{u}$ to a vertex $\mathbf{x} \in \mathbf{X}$.

From the minimality of the model $\mathbf{H}$ of $\mathbf{H}_i$ in $\mathbf{G}$ and the fact that each $\mathbf{x} \in \mathbf{X} \cap \mathbf{V}(\mathbf{H})$ can have at most $|\mathbf{H}_i|$ neighbours in $\mathbf{V}(\mathbf{H})$, we can infer that the number of important vertices is bounded by $\mathcal{O}(|\mathbf{H}_i|)$. Hence, as $\mathbf{h}$ is large, by the pigeonhole principle there is a subwall $\mathbf{W}'$ of $\mathbf{W}$ whose compass contains no important vertex (and which is still large enough). It follows that the intersection of $\mathbf{H}$ with the compass of $\mathbf{W}'$ is a set $\mathbf{L}_H$ of paths whose endpoints are outside of $\mathbf{W}'$. Now we replace the compass of $\mathbf{W}'$ in $\mathbf{G}$ by its planar core $\overline{\mathbf{C}}'$, i.e. we contract each extension $\mathbf{D}_i$ to a new single vertex $\mathbf{d}_i$ and add edges between any pair of vertices in $\boldsymbol{\partial}\mathbf{D}_i$. Let $\mathbf{G}'$ be the new graph. As $|\boldsymbol{\partial}\mathbf{D}_i| \leq 3$ for all $\mathbf{D}_i$, each extension can be traversed by at most one path in $\mathbf{L}_H$. Hence, $\mathbf{L}_H$ induces a unique linkage $\mathbf{L}$ in $\mathbf{G}'$ with the same effect on the set of important vertices. $\mathbf{L}$ is a linkage between important vertices and hence the number of paths in $\mathbf{L}$ is bounded by $\mathcal{O}(|\mathbf{H}|)$. We show next that we can find another linkage $\mathbf{L}'$ with the same effect on the set of important vertices that only uses the outermost $\mathbf{g}(|\mathbf{L}|)$ *shells* of the wall $\mathbf{W}'$ (cf. Figure 2). The next lemma follows immediately from 3.1 in [19] letting $\Gamma := \mathbf{W}$, $\mathbf{K} := \mathbf{G} \setminus (\mathbf{W} - \mathbf{P})$, and choosing $\mathbf{Z}$ as the set of endpoints of $\mathbf{L}$, as any vertex of $\mathbf{W}$ not in the outer $\mathbf{g}(|\mathbf{L}|)$ circuits is $\mathbf{g}(|\mathbf{L}|)$-insulated from $\mathbf{K}$.

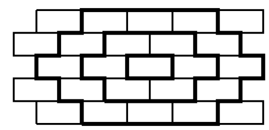

Figure 2: The shells of a wall of height 5

LEMMA 5.1. *There is a computable function* $\mathbf{g} : \mathbb{N} \to \mathbb{N}$ *such that if* $\mathbf{W} \subseteq \mathbf{G}$ *is a wall that is flat in a graph* $\mathbf{G}$ *and* $\mathbf{L}$ *is a linkage in* $\mathbf{G}$ *whose endpoints are disjoint from* $\mathbf{W}$*, then there is a linkage* $\mathbf{L}'$ *in* $\mathbf{G}$ *with the same effect on the endpoints as* $\mathbf{L}$ *such that* $(\mathbf{G} \setminus \mathbf{W}) \cap \mathbf{L}' \subseteq \mathbf{L}$ *and every vertex* $\mathbf{v} \in \mathbf{V}(\mathbf{W}) \cap \mathbf{V}(\mathbf{L}')$ *is contained in the outer* $\mathbf{g}(|\mathbf{L}|)$ *circuits of* $\mathbf{W}$*.*

The linkage $\mathbf{L}'$ also induces a model $\mathbf{H}'$ of $\mathbf{H}_i$ in $\mathbf{G}$ that has an empty intersection with the inner shells of $\mathbf{W}'$. Hence, there is a "large" subwall $\mathbf{W}''$ of $\mathbf{W}'$ that has an empty intersection with $\mathbf{L}'$ (and $\mathbf{H}'$).

For any graph $\mathbf{G}$ we let the *folio* of $\mathbf{G}$ be the class of all minors of $\mathbf{G}$. For $\boldsymbol{\delta} \geq 0$, we say that a graph $\mathbf{H}$ has *detail* at most $\boldsymbol{\delta}$ if $|\mathbf{E}(\mathbf{H})| \leq \boldsymbol{\delta}$ and $|\mathbf{V}(\mathbf{H})| \leq \boldsymbol{\delta}$ and define the $\boldsymbol{\delta}$*-folio* of $\mathbf{G}$ to be the class of all minors of $\mathbf{G}$ of detail at most $\boldsymbol{\delta}$.

In the following, we let $\boldsymbol{\delta} := \max\{|\mathbf{V}(\mathbf{H}_i)|, |\mathbf{E}(\mathbf{H}_i)|\}$. The following lemma is an immediate consequence of (10.3) and the algorithm in (10.4) in [21].

LEMMA 5.2. *There is a computable function* $\mathbf{f} : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ *so that if the height of* $\mathbf{W}''$ *is at least* $\mathbf{f}(|\mathbf{X}|, \boldsymbol{\delta})$*, then there is a subwall* $\mathbf{W}'''$ *of* $\mathbf{W}''$ *of which both middle vertices are irrelevant for the* $\boldsymbol{\delta}$*-folio of* $\mathbf{G}$*.*

The subwall $\mathbf{W}'''$ guaranteed by the previous lemma is called *homogeneous* in [21]. We remark that a homogeneous wall in a graph $\mathbf{G}$ remains homogeneous if we delete a vertex $\mathbf{w}_0$ from $\mathbf{G}$ that is not contained in the compass of the wall. This is used below.

Let $\mathbf{v}_0$ be one of the middle vertices of $\mathbf{W}'''$ and let $\mathbf{G}' := \mathbf{G} \setminus \mathbf{v}_0$. As $\mathbf{G}$ is a critical graph, $\mathbf{G}' \in \mathcal{C}^{\text{apex}}$. Hence, there is a vertex $\mathbf{w}_0$ in $\mathbf{V}(\mathbf{G}')$ so that $\mathbf{G}' \setminus \mathbf{w}_0 \in \mathcal{C}$. Recall that $\mathbf{H}'$ is a model of $\mathbf{H}_i$ in $\mathbf{G}$ that does not contain $\mathbf{v}_0$. Hence, $\mathbf{H}'$ is a model of $\mathbf{H}_i$ in $\mathbf{G}'$ and therefore $\mathbf{w}_0 \in \mathbf{V}(\mathbf{H}')$. Thus, $\mathbf{w}_0$ is not in the compass of $\mathbf{W}'''$.

Let $\mathbf{G}'' := \mathbf{G} \setminus \mathbf{w}_0$. As $\mathbf{G} \notin \mathcal{C}^{\text{apex}}$, $\mathbf{G}'' \notin \mathcal{C}$. Therefore, there is a model of some $\mathbf{H}_i \in \{\mathbf{H}_1, \dots, \mathbf{H}_m\}$ in $\mathbf{G}''$. As $\mathbf{w}_0$ is not contained in the compass of $\mathbf{W}'''$ in $\mathbf{G}$, $\mathbf{W}'''$ is still a flat homogeneous subwall of $\mathbf{G}''$ and hence the two middle vertices of $\mathbf{W}'''$ are still irrelevant for the $\boldsymbol{\delta}$-folio of $\mathbf{G}''$. It follows, that $\mathbf{G}'' \setminus \mathbf{v}_0$ still contains a model of $\mathbf{H}_i$ and therefore $\mathbf{G}'' \notin \mathcal{C}$. But this is a contradiction, as $\mathbf{G}'' \setminus \mathbf{v}_0 = \mathbf{G}' \setminus \mathbf{w}_0$.

This shows that if initially we choose $\mathbf{h}$ large enough so that after applying the various constructions above the wall $\mathbf{W}''$ can be guaranteed to be large enough to apply Lemma 5.2 for our value of $\boldsymbol{\delta}$, then Case 3 of the Trinity Lemma is impossible as well. Hence, the tree width of $\mathbf{G}$ is bounded by $\mathbf{f}(\mathbf{k}, \mathbf{h})$, where $\mathbf{f}$ is the function from the Trinity Lemma. This shows that critical graphs have bounded tree width and concludes the proof of Theorem 5.1.

5.2 Excluded Minors for Unions of Classes The last application of our method we give is to show that we can compute the set of excluded minors for the union of two minor ideals, provided that we are given their excluded minors.

THEOREM 5.2. *There is an algorithm that, given two classes $\mathcal{C}_1, \mathcal{C}_2$ of finite graphs represented by their sets of excluded minors $\mathcal{F}(\mathcal{C}_1) := \{\mathbf{H}_1, \dots, \mathbf{H}_l\}$ and $\mathcal{F}(\mathcal{C}_2) := \{\mathbf{I}_1, \dots, \mathbf{I}_r\}$, computes the set of excluded minors for the union $\mathcal{C} := \mathcal{C}_1 \cup \mathcal{C}_2$.*

Proof. We aim at applying Corollary 3.2. Towards this aim, we first establish an effective bound on the width of $\mathcal{C}$. Let $\mathbf{G} \notin \mathcal{C}$. Hence, there are $\mathbf{H}_i \in \{\mathbf{H}_1, \dots, \mathbf{H}_l\}$ and $\mathbf{I}_j \in \{\mathbf{I}_1, \dots, \mathbf{I}_r\}$ such that $\mathbf{H}_i \preccurlyeq \mathbf{G}$ and $\mathbf{I}_j \preccurlyeq \mathbf{G}$. Let $\mathbf{H}$ and $\mathbf{I}$ be models of $\mathbf{H}_i$ and $\mathbf{I}_j$ in $\mathbf{G}$ such that $|\mathbf{I} \cup \mathbf{H}|$ is as small as possible. We use $\mathbf{H}$ both for the model mapping and the image $\mathbf{H} \subseteq \mathbf{G}$ and likewise for $\mathbf{I}$. It will always be clear what is meant. Let $\mathbf{G}' := \mathbf{H} \cup \mathbf{I} \subseteq \mathbf{G}$. Further, let $\mathbf{k} := \max\{|\mathbf{H}_i|, |\mathbf{I}_j|\}$ and choose $\mathbf{h}$ "large enough", where, as in the previous section, the meaning of large enough will become clear in the course of the proof.

By the Trinity Lemma, $\mathbf{G}'$ either has a) tree width bounded by some $\mathbf{f}(\mathbf{k}, \mathbf{h})$ for some computable function $\mathbf{f}$, or b) a $\mathbf{K}_k$ minor, or c) there is a subset $\mathbf{X} \subseteq \mathbf{V}(\mathbf{G})$ with $|\mathbf{X}| < \binom{k}{2}$, a wall $\mathbf{W}$ of height $\mathbf{h}$ in $\mathbf{G} \setminus \mathbf{X}$, and a flat layout of the compass of $\mathbf{W}$ in $\mathbf{G} \setminus \mathbf{X}$.

Suppose Case c) applies. Note that every vertex in $\mathbf{G}'$ has degree at most $2\mathbf{k}$. If, initially, we choose $\mathbf{h}$ large enough, we can argue as in the previous section and use Lemma 5.2 to find a large homogenous subwall $\mathbf{W}'$ of $\mathbf{W}$ whose compass is disjoint from $\mathbf{X}$. Again by Lemma 5.2, the middle vertex $\mathbf{v}_0$ of $\mathbf{W}'$ is irrelevant for the $\mathbf{k}$-folio of $\mathbf{G}'$ and we can remove it to obtain a graph $\mathbf{G}'' := \mathbf{G} \setminus \mathbf{v}_0$ that still has $\mathbf{I}_j$ and $\mathbf{H}_i$ as a minor, contradicting the minimality of $\mathbf{H}$ and $\mathbf{I}$. Hence, Case 3) is impossible.

For Case b), Lemma 2.2 implies that $\mathbf{G}'$, and hence $\mathbf{G}$, has a subgraph of tree width at most $\mathbf{k}^2+1$ containing a $\mathbf{K}_k$ minor and therefore also an $\mathbf{I}_j$ and an $\mathbf{H}_i$ minor.

Finally, in Case a) the tree width of $\mathbf{G}'$ is bounded by a computable function in $\mathbf{h}$ and $\mathbf{k}$. In either case, we have found a uniform upper bound for the tree width of a subgraph of $\mathbf{G}$ containing an $\mathbf{I}_j$ and an $\mathbf{H}_i$ minor. This shows that $\mathcal{C}$ has effectively bounded with. To conclude the proof we need to show that $\mathcal{C}$ is MSO-definable. But this follows immediately from Example 2.2. ☐

6 Conclusion and Open Problems

We have introduced a general method for computing excluded minors for minor ideals that is based on MSO-definability on graphs of small tree width and the notion of the width of a minor ideal. As straightforward applications of this method we showed that excluded minors for several natural minor ideals are computable, in particular for the classes of graphs of tree width at most $\mathbf{k}$, branch-width at most $\mathbf{k}$, or of bounded genus.

We have also demonstrated the usefullness of our method by giving far more advanced examples. In particular, we showed that excluded minors can be computed for the class of apex graphs over a minor ideal whose excluded minors we know already, or for the union of two minor ideals, provided we are given the excluded minors of the base classes. The latter answers an open question from [5].

We conclude the paper by stating some open problems. In [8], Eppstein introduced minor ideals with the *diameter tree width property*, which is now more commonly known as bounded local tree width. We briefly recall the definition. Let $\mathbf{G}$ be a graph. For every vertex $\mathbf{v} \in \mathbf{V}(\mathbf{G})$ and every $\mathbf{r} \geq 1$ we define the $\mathbf{r}$-*neighbourhood* $\mathbf{N}_r^G(\mathbf{v})$ *of* $\mathbf{v}$ as the set of vertices of distance at most $\mathbf{r}$ from $\mathbf{v}$. The *local tree width* of $\mathbf{G}$ is the function $\mathrm{ltw}(\mathbf{G}, \cdot)$ defined by $\mathrm{ltw}(\mathbf{G}, \mathbf{r}) = \max\left\{ \mathrm{tw}\left(\mathbf{G}[\mathbf{N}_r^G(\mathbf{v})]\right) \ : \ \mathbf{v} \in \mathbf{V}(\mathbf{G})\right\}$. For all nonnegative integers $\boldsymbol{\lambda}, \boldsymbol{\mu}$ we let $\mathcal{L}(\boldsymbol{\lambda}) = \left\{\mathbf{G} \ : \ \forall \mathbf{H} \preccurlyeq \mathbf{G} \ \forall \mathbf{r} \geq 0 : \ \mathrm{ltw}(\mathbf{H}, \mathbf{r}) \leq \boldsymbol{\lambda} \cdot \mathbf{r}\right\}$, and $\mathcal{L}(\boldsymbol{\lambda}, \boldsymbol{\mu}) = \left\{\mathbf{G} \ : \ \exists \mathbf{X} \subseteq \mathbf{V}(\mathbf{G}) \text{ s. th. } |\mathbf{X}| \leq \boldsymbol{\mu} \text{ and } \mathbf{G} \setminus \mathbf{X} \in \mathcal{L}(\boldsymbol{\lambda})\right\}$.

Minor ideals of bounded local tree width have many interesting algorithmic properties. Another reason these classes are interesting is the fact that for every $\mathbf{k} \in \mathbb{N}$ there are $\boldsymbol{\lambda}, \boldsymbol{\mu} \in \mathbb{N}$ such that every graph $\mathbf{G}$ with $\mathbf{K}_k \not\preccurlyeq \mathbf{G}$ has a tree decomposition over $\mathcal{L}(\boldsymbol{\lambda}, \boldsymbol{\mu})$ (see [14]). We do not currently know how to compute upper bounds $\mathbf{w}_\lambda$