Locally Excluding a Minor

Anuj Dawar

University of Cambridge Computer Laboratory

anuj.dawar@cl.cam.ac.uk

Martin Grohe

Institut für Informatik,

Humboldt Universität zu Berlin

grohe@informatik.hu-berlin.de

Stephan Kreutzer

Institut für Informatik,

Humboldt Universität zu Berlin

kreutzer@informatik.hu-berlin.de

February 1, 2007

Abstract

We introduce the concept of locally excluded minors. Graph classes locally excluding a minor generalise the concept of excluded minor classes but also of graph classes with bounded local tree-width and graph classes with bounded expansion.

We show that first-order model-checking is fixed-parameter tractable on any class of graphs locally excluding a minor. This strictly generalises analogous results by Flum and Grohe on excluded minor classes and Frick and Grohe on classes with bounded local tree-width.

As an important consequence of the proof we obtain fixed-parameter algorithms for problems such as dominating or independent set on graph classes excluding a minor, where now the parameter is the size of the dominating set and the excluded minor.

We also study graph classes with excluded minors, where the minor may grow slowly with the size of the graphs and show that again, first-order model-checking is fixed-parameter tractable on any such class of graphs.

1   Introduction

An important task in the theory of algorithms is to find feasible instances of otherwise intractable algorithmic problems. For this purpose, the notion of bounded tree-width has proved to be extremely useful. Many NP-complete problems become tractable on graphs whose tree-width is bounded by a fixed constant. These include k-colourability, Hamiltonicity and the k-dominating and k-vertex cover problems. Courcelle [2] proved a meta-theorem stating that any decision problem definable in monadic second-order logic can be decided in linear time on any class of graphs of bounded tree-width.

Although among the most prominent, tree-width is not the only structural property of graphs that allows for efficient solutions of otherwise intractable problems. Other important restrictions are planarity or bounded degree. Unfortunately, Courcelle's theorem fails for any of these restrictions. Seese [19] was

the first to give a logical meta-theorem for another general class of graphs. He showed that any first-order definable decision problem can be decided in linear time on any class of graphs of bounded degree. A property of planar graphs and graphs of bounded degree that is often used to obtain tractable algorithms is the fact that every neighbourhood of a vertex has low tree-width, i.e. the tree-width of any neighbourhood of a vertex only depends on its diameter. This observation led Eppstein [6] to introduce the notion of bounded local tree-width (which he calls diameter tree-width property). The concept of local tree-width strictly generalises tree-width, planarity and bounded degree.

In [10], Frick and Grohe generalise Seese's result to classes of graphs of bounded local tree-width. They proved that first-order definable decision problems can be decided in linear time on what they call locally tree-decomposable classes of graphs. In the same paper, they show that any first-order definable decision problem can be decided in quadratic time on any class of graphs of bounded local tree-width.

Another important concept of graph structure that has been used to obtain tractable algorithms is the concept of excluded minors. In a series now running to 23 papers, Robertson and Seymour developed their groundbreaking theory on graph minors that culminated in the proof of Wagner's conjecture stating that in every infinite class of finite graphs one graph is a minor of another. In other words, every minor closed class of graphs that is not the class of all graphs can be characterised by a finite set of excluded minors. They also proved that testing whether a fixed graph is a minor of a graph G can be done in cubic time. It follows, that any minor closed class of graphs can be decided in cubic time.

Many parts of the rich and deep theory developed by Robertson and Seymour have found algorithmic applications. In [7], Flum and Grohe proved a meta-theorem similar to the results mentioned above. They showed that any first-order definable decision problem can be decided in polynomial time on any class of graphs excluding a fixed minor. The concept of excluded minors is incomparable to the concept of local tree-width, in fact even to bounded degree. It is therefore a natural question whether there exists a common generalisation of excluded minor classes and bounded local-tree width classes on which we can still efficiently decide first-order definable decision problems. The main contribution of this paper is to introduce such a generalisation. It is based on the following simple concept. Let $\mathcal{C}$ be a class of graphs. Instead of requiring that any graph in $\mathcal{C}$ excludes a fixed minor, we only require that every neighbourhood excludes a minor, depending on its radius. Formally, we require that for every radius r there is a graph H_r so that every r-neighbourhood of a vertex of any member of $\mathcal{C}$ excludes H_r. We call classes of graphs with this property *graph classes locally excluding a minor*. It is easily seen that if $\mathcal{C}$ is a class of graphs of bounded local tree-width or $\mathcal{C}$ is a class of graphs excluding a fixed minor, then $\mathcal{C}$ also locally excludes a minor. But the concept of locally excluded minors is fairly rich. It also generalises the concept of bounded expansion, recently introduced by Nešetřil and de Mendez [20]. The following is the main result on locally excluded minor classes.

Theorem 1.1 *Let $\mathcal{C}$ be a class of graphs locally excluding a minor. Then deciding first-order properties is fixed-parameter tractable on $\mathcal{C}$.*

Here, the exponent of the polynomial is fixed and neither depends on the formula used to define the problem nor on the locally excluded minors. With the

exception of bounded clique-width [3], graph classes locally excluding a minor strictly generalise all classes of graphs known so far on which first-order model checking is fixed-parameter tractable.

We also consider classes $\mathcal{C}$ of graphs such that there is a (slowly growing) function $f : \mathbb{N} \to \mathbb{N}$ with the property that any $G \in \mathcal{C}$ exclude a minor of cardinality at most $f(|G|)$ and show that also on such classes of graphs, first-order model checking is still fixed-parameter tractable.

Theorem 1.2 *There is an unbounded function $f : \mathbb{N} \to \mathbb{N}$ such that deciding first-order properties is fixed-parameter tractable on the class $\mathcal{C}_f$ of graphs G excluding a clique of order at most $f(|G|)$.*

The method we use to show the theorem has further important consequences. It is well known that various intractable problems such as k-dominating set and others are fixed-parameter tractable on classes of graphs excluding a fixed minor, where k is the parameter. However, it was an open problem whether the exponent of the polynomials can be made independent of the excluded minor, i.e. whether these problems can be solved by a fixed-parameter algorithm where the parameter is both k and the excluded minor. The second main contribution of this paper is to give a positive answer to this question.

At the core of many algorithms on excluded minor classes is a deep decomposition theorem by Robertson and Seymour which states that any graph excluding a fixed minor can be decomposed into a tree whose bags are almost embeddable into a surface that almost has bounded genus. In [4], Demaine et al. give a polynomial time algorithm for computing such decompositions. Grohe [12] derived from Robertson and Seymour's work a different decomposition that is often easier to use in the design of algorithms. He showed that any graph excluding a fixed minor can be decomposed into a tree whose bags have bounded local tree-width after removal of a constant number of elements. The main technical contribution of this paper is to show that we can compute such a decomposition in time $f(|H|) \cdot n^{\mathcal{O}(1)}$, where H is the excluded minor. We use this to show the following theorem.

Theorem 1.3 *Let $\mathcal{C}$ be a class of graphs excluding the fixed minor H. Then any first-order definable decision problem can be solved in time $f(|H|, |\varphi|) \cdot n^{\mathcal{O}(1)}$, where f is a computable function, φ is the sentence defining the decision problem, and H is the excluded minor.*

This result immediately implies fixed-parameter algorithms for problems such as the k-dominating set problem, where now the parameter is k and the excluded minor. On the other hand, the new decomposition can also be used directly to obtain faster algorithms for a variety of problems on excluded minor classes.

2 Preliminaries

We denote the set of real numbers by $\mathbb{R}$, the set of integers by $\mathbb{Z}$, and the set of positive integers (natural numbers) by $\mathbb{N}$. For all $n \in \mathbb{N}$, by $[n]$ we denote the set $\{1, \ldots, n\}$, and for all $m, n \in \mathbb{Z}$, by $[m, n]$ we denote the set $\{m, \ldots, n\}$ (the empty set if $m > n$). Occasionally, we use $[0, 1]$ to denote the unit interval

$\{x \in \mathbb{R} \mid 0 \le x \le 1\}$ instead of the set $\{0,1\}$; it will always be clear from the context what is meant.

Graphs are undirected and simple. If G is a graph, we write $V(G)$ for its vertex set and $E(G)$ for its set of edges. For any set $X \subseteq V(G)$ we write $G[X]$ for the subgraph of G induced by X, and we let $G \setminus X = G[V(G) \setminus X]$. For a set $F \subseteq E$, we write $G - F$ for the graph $(V, E \setminus F)$. For every set S, by $K[S]$ we denote the complete graph with vertex set S. Furthermore, for every $k \in \mathbb{N}$ we let $K_k = K[[k]]$.

The *distance* $d^G(v, w)$ between two vertices v, w of a graph G is defined to be the length of the shortest path from v to w. For nonempty sets $X, Y \subseteq V(G)$, we let $d^G(X, Y) = \min\{d^G(x, y) \mid x \in X, y \in Y\}$, and similarly we define the distance $d^G(v, X)$ of a vertex $v \in V(G)$ from a nonempty set $X \subseteq V(G)$. We write $\rho(\mathcal{G}) := \min_{v \in V(G)} \max_{v' \in V(G)} d^G(v, v')$ for the radius of G.

For every nonempty set $X \subseteq V(G)$ and every $r \ge 1$ we define the *r-neighborhood of X* and the *r-sphere around X* by

$$N_r^G(X) = \{v \in V(G) \ : \ d^G(v, X) \le r\},$$
$$S_r^G(X) = \{v \in V(G) \ : \ d^G(v, X) = r\}.$$

For a vertex $v \in V(G)$, we let $N_r^G(v) = N_r^G(\{v\})$ and $S_r^G(v) = S_r^G(\{v\})$. For a nonempty subgraph $H \subseteq G$ we let $N_r^G(H) = N_r^G(V(H))$ and $S_r^G(H) = S_r^G(V(H))$. For 1-neighborhoods and 1-spheres, we write N^G and S^G instead of N_1^G and S_1^G, respectively.

A *separation* of a graph G is a pair (X, Y) of subsets of $V(G)$ such that $G = G[X] \cup G[Y]$. The set $X \cap Y$ is called the *separator* of the separation (X, Y), and its cardinality is the *order* of the separation.

We write $G \preceq H$ to denote that G is a minor of H. A class $\mathcal{C}$ of graphs is *minor-closed* if it is downward closed under $\preceq$. $\mathcal{C}$ is an *excluded minor class* if its minor-closure is not the class of all graphs.

Theorem 2.1 (Robertson and Seymour [17]) *The following problem is fixed-parameter tractable:*

p-Minor	
Input:	Graphs G, H.
Parameter:	$|V(H)|$.
Problem:	Decide if $H \preceq G$.

More precisely, there is a computable function f and an algorithm that solves the problem in time $f(k) \cdot n^3$, where $k = |V(H)|$ and $n = |V(G)|$.

A *curve* in the plane $\mathbb{R}^2$ is the image of a continuous function $f : [0, 1] \to \mathbb{R}^2$ (here $[0, 1]$ denotes the unit interval of real numbers). The *endpoints* of the curve are $f(0)$ and $f(1)$, and its *interior* is the set of all other points. The curve is *simple* if f is one-to-one. A *plane graph* is a graph Γ whose vertices are distinct elements of the plane and whose edges are simple curves such that the endpoints of each edge are the two vertices incident with it, and the interior of each edge is disjoint from the vertex set and from all other edges. Abusing notation, we also write Γ to denote the point set $V(\Gamma) \cup \bigcup E(\Gamma) \subseteq \mathbb{R}^2$. We say that Γ is *embedded* in a subset $X \subseteq \mathbb{R}^2$ if $\Gamma \subseteq X$.

3 Tree-Decompositions

A *tree decomposition* of a graph G is a pair $D = (T, B)$, 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 $B^{-1}(v) = \{t \in V(T) \mid v \in B(t)\}$ is connected in T. The sets B_t, for $t \in V(T)$, are called the *bags* of the decomposition D. For a subset $U \subseteq V(T)$ we let $B(U) = \bigcup_{t \in U} B_t$, and for a subtree or forest $T' \subseteq T$ we let $B(T') = \bigcup_{t \in V(T')} B_t$.

With each bag B_t of a tree decomposition (T, B) of a graph G we associate two graphs:

- The *closure* of B_t is the graph $\overline{B}_t = G[B_t] \cup \bigcup_{u \in S^T(t)} K[B_t \cap B_u]$.

- The *companion* of B_t is the graph $\widehat{B}_t$ obtained from $G[B_t]$ by adding new vertices x_u for all $u \in S^T(t)$ and edges from x_u to all $v \in B_t \cap B_u$.

A tree decomposition is *strongly over* a class $\mathcal{C}$ of graphs if the closures of all its bags belong to $\mathcal{C}$, and it is *weakly over* $\mathcal{C}$ if the companions of all its bags belong to $\mathcal{C}$. The proof of the following lemma is straightforward:

Lemma 3.1 *Let $(T, (B_t)_{t \in V(T)})$ be a decomposition of a graph G, and let $t \in V(T)$. Then for every connected subgraph C of G the graph $\overline{B}_t[V(C) \cap B_t]$ is either empty or connected.*

For a tree T and an edge $e = \{t, u\} \in E(T)$, by T_{tu} and T_{ut} we denote the two connected components of $T - \{e\}$ such that $u \in V(T_{tu})$ and $t \in V(T_{ut})$. (Thus the indices determine which way we are looking from e.) The following lemma is well-known and easy to prove:

Lemma 3.2 *Let (T, B) be a tree decomposition of a graph G and $\{t, u\} \in E(T)$. Then $\big(B(T_{tu}), B(T_{ut})\big)$ is a separation of G with separator $B_t \cap B_u$.*

Let $D = (T, (B_t)_{t \in V(T)})$ be a tree decomposition of a graph G. For every edge $e = \{t, u\} \in E(T)$, we call $\big(B(T_{tu}), B(T_{ut})\big)$ the separation *at* e. We call the collection of all such separations the separations *of* D. Similarly, we call $B_t \cap B_u$ the separator *at* e and the collection of all these separators the separators *of* D. The *order* of the decomposition D is defined to be the maximum of the cardinalities of its separators.

The *width* of the decomposition is the number $\text{width}(D) = \max\{|B_t| \mid t \in V(T)\} - 1$. The *tree width* of a graph G is the number

$$\text{tw}(G) = \min\{\text{width}(D) \mid D \text{ tree decomposition of } G\}.$$

The *local tree width* of G is the function $\text{ltw}(G, \cdot)$ defined by $\text{ltw}(G, r) = \max\left\{\text{tw}\big(G[N_r^G(v)]\big) : v \in V(G)\right\}$. For all nonnegative integers λ, μ we let

$$\mathcal{L}(\lambda) = \big\{G : \forall H \preceq G \ \forall r \geq 0 : \text{ltw}(H, r) \leq \lambda \cdot r\big\},$$

$$\mathcal{L}(\lambda, \mu) = \Big\{G : \exists X \subseteq V(G) \text{ s. th. } \begin{array}{l} |X| \leq \mu \text{ and} \\ G \setminus X \in \mathcal{L}(\lambda) \end{array} \Big\}.$$

The following result is a consequence of Robertson and Seymour's deep structure theorem for graphs with excluded minors proved in [18].

Theorem 3.3 (Grohe [12]) *There are computable functions $\lambda, \mu : \mathbb{N} \rightarrow \mathbb{N}$ such that for every $k \in \mathbb{N}$, every graph G with $K_k \npreceq G$ has a tree decomposition that is strongly over $\mathcal{L}(\lambda(k), \mu(k))$.*

The previous theorem can be used to obtain algorithms for various problems on excluded minor classes. These algorithms usually work along the following lines. Given a graph with a fixed forbidden minor, the first step is to compute the tree decomposition over a class $\mathcal{L}(\lambda, \mu)$. Then, for each block of the decomposition, the μ elements that need to be removed to obtain a graph with small local tree-width are computed. The problem, for instance the vertex cover problem, is then solved on each of these blocks using methods on minor closed classes with bounded local tree-width. The results are then extended to the blocks with the μ removed elements put back in and then to the whole tree. For this approach to work we need to be able to a) compute the tree decomposition over the class $\mathcal{L}(\lambda, \mu)$ and b) compute, for each block, the μ-elements that need to be removed to obtain a graph in $\mathcal{L}(\lambda)$.

Polynomial-time algorithms for both steps are known, but the exponent of the polynomials depend on the excluded minor, for instance on the number μ of elements that need to be removed. In the following two sections, we show that both steps are fixed parameter tractable if the parameter is the excluded minor.

In the next section we give an algorithm to compute the decomposition within the desired running time. In the section thereafter we show that the second step can be computed within suitable time bounds.

4 Computing tree-decompositions over $\mathcal{L}(\lambda, \mu)$

The goal of this section is to prove the following theorem:

Theorem 4.1 *There are computable functions $f, \lambda, \mu : \mathbb{N} \rightarrow \mathbb{N}$ and an algorithm that, given a graph G with $K_k \npreceq G$, computes a tree decomposition of G that is weakly over $\mathcal{L}(\lambda(k), \mu(k))$ in time $f(k) \cdot n^{O(1)}$.*

Let us first give a high level description of the proof. The algorithm to compute the tree decompositions is actually quite simple. Essentially we repeatedly compute small separators in the input graph G to obtain a clique-sum decomposition of G of order k. It turns out that we can take any separator of order k in G, as long as the graph is separated into two parts which are both not too small, where we can take "not too small" as being at least of order $k!$.

We then associate with any such a clique-sum decomposition a certain tree decomposition and show that it can easily be computed given the clique-sum decomposition. We call these k-CMS tree decompositions.

The main and difficult part of the proof consists in showing that these k-CMS tree decompositions actually are tree decompositions that are weakly over $\mathcal{L}(\lambda, \mu)$ for some $\lambda := \lambda(k)$ and $\mu := \mu(k)$.

The proof of Theorem 4.1 is presented in two parts. In Section 4.1 we present the necessary definitions and the algorithm to compute a k-CMS tree decomposition. The proof that this is indeed a tree decomposition weakly over some class $\mathcal{L}(\lambda, \mu)$ is given in Section 4.2.

4.1 Algorithms

A separation (X, Y) of order $k \geq 1$ of some graph is *substantial* if

$$|X|, |Y| > k!.$$

A separation (X, Y) of order 0 is substantial if $|X|, |Y| > 0$. (X, Y) is a *minimum substantial separation* if it is substantial, and if there is no substantial separation of smaller order.

Lemma 4.2 *There is an fpt-algorithm with parameter k that decides whether a given graph G has a substantial separation of order at most k and computes a minimum substantial separation if there is one of order at most k.*

Proof. We shall design a recursive fpt-algorithm for the following more general problem:

Input:	Graph G, sets $P, Q \subseteq V(G)$, and nonnegative integers k, k', ℓ, m with $k' \leq k$ and $\ell, m \leq k!$.
Parameter:	k.
Problem:	Decide whether G has a separation (X, Y) of order k' such that $P \subseteq X$ and $Q \subseteq Y$ and $|X \setminus Y| \geq \ell$ and $|Y \setminus X| \geq m$ and compute such a separation if there is one.

Let G, P, Q, k, k', ℓ, m be an instance of the problem. Without loss of generality we may assume that $k \geq 2$. Let $n = |V(G)|$. If $n < k' + \ell + m$ no substantial separator with the desired properties exists, and the algorithm reports failure. So suppose that $n \geq k' + \ell + m$. If $k' = 0$, it is straightforward to check whether the connected components of G can be arranged into a separation with the desired properties, so we further assume that $k' > 0$.

Using standard techniques, the algorithm computes some separation (X, Y) of G of order k' with $P \subseteq X$ and $Q \subseteq Y$. If no such separation exists, then the algorithm reports failure. If $|X \setminus Y| \geq \ell$ and $|Y \setminus X| \geq m$, the algorithm returns the separation (X, Y). In the following, we assume without loss of generality that $|X \setminus Y| < \ell$. Let $S = X \cap Y$.
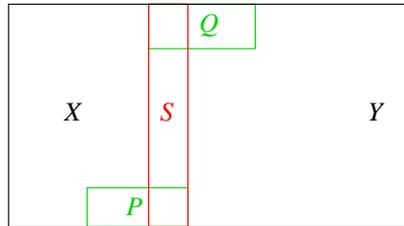
Figure 4.1:

Now for all separations (X_1, X_2) of $G[X]$ of order at most k' with $P \cap X \subseteq X_1$ and $Q \cap X \subseteq X_2$, we recursively call our algorithm on the instance
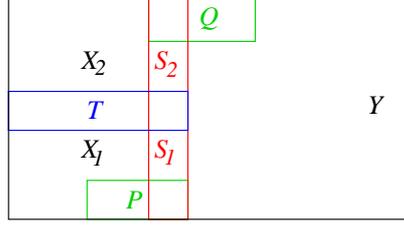
Figure 4.2:
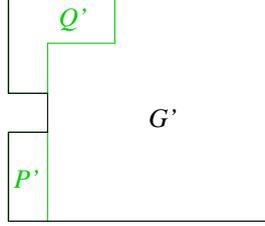
Figure 4.3:

$G', P', Q', k, k'', \ell', m'$ defined as follows: Let $T = X_1 \cap X_2$ and $S_1 = (S \cap X_1) \backslash T$, $S_2 = (S \cap X_2) \setminus T$ (see Figures 4.1–4.3). Let

- $G' = G[Y] \setminus T$,

- $k'' = k' - |T|$,

- $P' = (P \cup S_1) \cap V(G')$,

- $Q' = (Q \cup S_2) \cap V(G')$,

- $\ell' = \max\{0, \ell - |X_1 \setminus (X_2 \cup S)|\}$ and

- $m' = \max\{0, m - |X_2 \setminus (X_1 \cup S)|\}$.

If the recursive call returns a separation (Y_1, Y_2), then our algorithm returns the separation $(X_1 \cup Y_1, X_2 \cup Y_2)$, which can easily be seen to have the desired properties. If the recursive calls fail for all (X_1, X_2), then the algorithm reports failure.

Let $t_k(n, k, \ell, m)$ denote the worst case running time of the algorithm in terms of the parameters n, ℓ, m, k, and let $T(n, k, s) = \max\{t(n, k, \ell, m) \mid 0 \leq \ell, m \leq s, \ell + m = s\}$. (The parameter k' is irrelevant for our running time analysis.) We get the following recurrence in terms of s:

$$
\begin{aligned}
T(n, k, 0) &= n^{O(1)} && \text{for } n \geq 0 \\
T(n, k, s) &= O(1) && \text{for } n < s \\
T(n, k, s) &= O\big(2^{k!} \cdot T_k(n, s - 1)\big) && \text{for } n \geq s \geq 1.
\end{aligned}
$$

To obtain this recurrence, we put a crude upper bound of $2^{k!}$ on the number of separations (X_1, X_2) of $G[X]$ and ignore that not only s, but also n decreases in the recursive calls. Since $s \leq 2k!$, this yields an fpt-bound on the running time. $\qquad \square$

8

Let G be a graph. The graph G is a *clique sum* of graphs A, B (we write $G = A \oplus B$) if there is a separation (X, Y) of G such that $A = G[X] \cup K[X \cap Y]$ and $B = G[Y] \cup K[X \cap Y]$. We write $G = A \oplus_k B$ to indicate that $G = A \oplus B$ and the order of the corresponding separation is at most k, and we similarly write $G = A \oplus_S B$ to indicate that the separator is S. A *clique-sum decomposition* of G is a pair (T, A) where T is a rooted binary tree and H is a mapping that associates a graph A_t with every node $t \in V(T)$ such that:

1. $A_r = G$ for the root r of T;

2. $A_t = A_{t_1} \oplus A_{t_2}$ for every node $t \in V(T)$ with children t_1, t_2.

The separation corresponding to the clique sum $A_t = A_{t_1} \oplus A_{t_2}$ is called the separation *at* node t, and its separator $V(A_{t_1}) \cap V(A_{t_2})$ is called the separator *at* node t. The graphs A_t for the leaves t of T are called the *pieces* of the decomposition. The *order* of a clique-sum decomposition (T, A) is the least k such that $A_t = A_{t_1} \oplus_k A_{t_2}$ for every node $t \in V(T)$ with children t_1, t_2. The *height* of the decomposition is the height of the tree T, that is, the maximum length of a path from the root of T to a leaf.

A clique-sum decomposition (T, A) is an *MS decomposition* if all its separations are minimum substantial. It is a *complete MS decomposition of order k* (for short: *k-CMS decomposition*) if the order of the decomposition is at most k, and if for all leaves t of T the graph A_t has no substantial separation of order at most k.

Lemma 4.3 *There is an fpt-algorithm with parameter k that computes a k-CMS-decomposition for a given graph G.*

Proof. This follows directly from Lemma 4.2. $\qquad\qquad\qquad\qquad\qquad\square$

Definition 4.4 *Let $D = (T, A)$ be a clique-sum decomposition of a graph G.*

A tree decomposition associated with the clique decomposition D is a tree decomposition $D' = (T', A')$ such that

- *the closures of the bags of D' are the pieces of D (and hence the bags of D' are the vertex sets of the pieces of D).*

- *for every node $t \in V(T)$ there is an edge $e \in E(T')$ such that for the separation (X, Y) at t and the separation (X', Y') at e it holds that $X' \cap V(A_t) = X$ and $Y' \cap V(A_t) = Y$.*

- *for every edge $e \in E(T')$ there is a node $t \in V(T)$ such that for the separation (X, Y) at t and the separation (X', Y') at e it holds that $X' \cap V(A_t) = X$ and $Y' \cap V(A_t) = Y$.*

- *D and D' have the same separators and hence the same order.*

Lemma 4.5 *For every clique-sum decomposition $D = (T, A)$ of a graph G there is an associated tree decomposition $D' = (T', A')$.*

Proof. By induction on the height of $t \in V(T)$, we define a tree T'_t whose vertex set is the set of all leaves of T that are in the subtree with root t. Let $t \in V(T)$. If $t \in V(T)$ is a leaf of T, we let T'_t be the tree $(\{t\}, \varnothing)$. So suppose that t has children t_1, t_2. Let (X_1, X_2) be the separation at t such that $V(A_{t_1}) \subseteq X_1$ and

$V(A_{t_2}) \subseteq X_2$. Let $S = X \cap Y$. As S induces a clique in both A_{t_1} and A_{t_2}, there must be leaves u_1, u_2 of T such that for $i = 1, 2$ the leaf u_i is in the subtree of T with root t_i, and $S \subseteq V(A_{u_i})$.

Applying the induction hypothesis, let T'_{t_i} be the trees constructed for T_{t_i}. Thus $u_i \in V(T'_{t_i})$. We let T'_t be the tree obtained from the disjoint union of T'_{t_1} and T'_{t_2} by adding an edge from u_1 to u_2.

Now let $T' = T'_r$ for the root r of T. For every $t \in V(T)$ we let $A'_t = V(A_t)$. It is easy to see that (T', A') is a tree decomposition of G with the desired properties. $\qquad\square$

A *k-CMS tree decomposition* of a graph G is a tree decomposition associated with a k-CMS decomposition.

Corollary 4.6 *There is an fpt-algorithm with parameter k that computes a k-CMS tree decomposition for a given graph G.*

4.2 Correctness

The goal of this section is to show that if G is a graph that has a tree decomposition over $\mathcal{L}(\lambda, \mu)$, then there is a λ' such that any $(\lambda + \mu + 1)$-CMS tree decomposition is a tree decomposition weakly over $\mathcal{L}(\lambda', \mu)$ (see Lemma 4.11). As a consequence, the algorithm presented in the previous section actually computes a tree decomposition weakly over $\mathcal{L}(\lambda', \mu)$.

The first step is to show that pieces of a CMS-decomposition can not be decomposed into two parts that are both not too small.

Lemma 4.7 *Let P be a piece of a k-CMS decomposition of a graph G, and let (X, Y) be a separation of G of order at most k. Then either $|V(P) \cap X| \le k!$ or $|V(P) \cap Y| \le k!$.*

Proof. The proof is by induction on the height of the decomposition. If the height is 0, then $P = G$, and the existence of a separation (X, Y) of order at most k with $|V(P) \cap X| > k!$ and $|V(A) \cap Y| > k!$ would contradict the completeness of the decomposition.

For the induction step, let (T, A) be a k-CMS decomposition of G of height at least 1 of which P is a piece, say, $P = A_\ell$ for some leaf ℓ of T. Let r be the root of T and r_1, r_2 its children such that ℓ is in the subtree rooted in r_1. Let (X_1, X_2) be the separation at r, and let $S = X_1 \cap X_2$ be its separator. For simplicity, we let $A_i = A_{r_i}$. Then $A_i = G[X_i] \cup K[S]$ for $i = 1, 2$. Note that the decomposition (T, A) induces a decomposition of A_1 of smaller height that still has P as a piece. Thus we can apply the induction hypothesis to every separation of A_1 of order at most k. Suppose for contradiction that (Y_1, Y_2) is a separation of G of order at most k such that $|V(P) \cap Y_1| > k!$ and $|(V(P) \cap Y_2| > k!$. Let $T = Y_1 \cap Y_2$ be the separator of this separation.

Let $U = S \cap T$, and for $i = 1, 2$, let $S_i = (S \cap Y_i) \setminus T$ and $T_i = (T \cap X_i) \setminus S$. Furthermore, let $s = |S|$, $t = |T|$, $u = |U|$, $s_i = |S_i|$, $t_i = |T_i|$. Then

$$s_1 + s_2 + u = s \le k, \tag{4.1}$$

$$t_1 + t_2 + u = t \le k. \tag{4.2}$$

Without loss of generality we may further assume that

$$s_1 \le s_2. \tag{4.3}$$

10

Consider the separation $((X_1 \cap Y_1) \cup S_1, (X_1 \cap Y_2) \cup S_1)$ of A_1 with separator $S_1 \cup T_1 \cup U$. The order of this separation is $s_1 + t_1 + u$. Since $V(P) \subseteq X_1$, we have $V(P) \cap ((X_1 \cap Y_2) \cup S_1) = (V(P) \cap Y_2) \cup (V(P) \cap S_1)$ and hence $|V(P) \cap ((X_1 \cap Y_2) \cup S_1)| \geq |V(P) \cap Y_2| > k!$, and similarly $|V(P) \cap ((X_1 \cap Y_1) \cup S_1)| > k!$. We claim that

$$s_1 + t_1 + u > k \geq t. \tag{4.4}$$

To see this, suppose for contradiction that $s_1 + t_1 + u \leq k$. Then $((X_1 \cap Y_1) \cup S_1, (X_1 \cap Y_2) \cup S_1)$ is a separation of A_1 of order at most k, and by the induction hypothesis we either have $|V(P) \cap ((X_1 \cap Y_1) \cup S_1)| \leq k!$ or $|V(P) \cap ((X_1 \cap Y_2) \cup S_1)| \leq k!$. This is a contradiction, which proves (4.4).

(4.4) and (4.2) imply

$$s_1 > t_2. \tag{4.5}$$

Thus by (4.1),

$$s_2 + t_2 + u < s. \tag{4.6}$$

By (4.3), this also implies

$$s_1 + t_2 + u < s. \tag{4.7}$$

Furthermore, $s \geq 2$, because $s > s_1 > t_2 \geq 0$ by (4.7) and (4.5). Remember that $|X_2| > s!$, because the partition (X_1, X_2) is substantial. Thus at least one of the two sets $X_2 \cap Y_1$ and $X_2 \cap Y_2$ is larger than $(s-1)!$. Without loss of generality we assume that

$$|X_2 \cap Y_1| > (s-1)! \tag{4.8}$$

Consider the separation $(X_2 \cap Y_1, X_1 \cup Y_2)$ of G. Its separator is $S_1 \cup T_2 \cup U$ and hence its order is at most $s - 1$ by (4.7). Thus by (4.8), it is substantial. But this contradicts the assumption that S is a minimum substantial separator. $\qquad\square$

The next step is to show that in any MS tree-decomposition, the separators at an edge e are connected in both parts of the separation (see Corollary 4.10). We first prove this for minimum substantial separators (Lemma 4.8) and then extend it to MS-decompositions (Lemma 4.9) and finally to MS tree-decompositions.

Lemma 4.8 *Let (X, Y) be a minimum substantial separation of a graph G and $S = X \cap Y$. Then there is a connected component C of $G[X \setminus Y]$ such that $S^G(C) = S$.*

Proof. Let $k = |S|$ and suppose that $S = \{s_1, \ldots, s_k\}$. The claim is obvious for $k \leq 1$, so assume that $k \geq 2$. Suppose for contradiction that there is no connected component C of $G[X \setminus Y]$ such that $S^G(C) = S$. For each connected component C of $G[X \setminus Y]$, we pick an $i_C \in [k]$ such that $S^G(C) \subseteq S \setminus \{s_i\}$. For all $i \in [k]$, let

$$W_i = \bigcup \{V(C) \mid C \text{ connected component of } G[X \setminus Y] \text{ with } i_C = i\}.$$

Then $(W_i \cup (S \setminus \{s_i\}), V(G) \setminus W_i)$ is a separation of G of order $(k-1)$. By the minimality of S, this separation is not substantial. Hence $W_i \leq (k-1)! - (k-1)$. It follows that

$$k! - k^2 + k = k \cdot ((k-1)! - (k-1)) \geq \sum_{i=1}^k |W_i|$$
$$\geq |\bigcup_{i=1}^k W_i| = |X \setminus S| > k! - k.$$

This is a contradiction. □

Lemma 4.9 *Let (T, A) be an MS-decomposition of a graph G and let $t \in V(T)$. Then for every connected subgraph $C' \subseteq A_t$ there is a connected subgraph $C \subseteq G$ such that $V(C) \cap V(A_t) = V(C')$.*

Proof. Let $t_0, t_1, \ldots, t_p = t$ be the path in T from the root t_0 to t. For $i \in [0, p]$, let $A_i = A_{t_i}$, and for $i < p$, let (X_i, Y_i) be the separation at i with separator $S_i = X_i \cap Y_i$. Without loss of generality, we may assume that $V(A_{i+1}) = X_i$. By induction on $i \in [0, p]$ we prove that there is a connected subgraph C_{p-i} of A_{p-i} such that $V(C_{p-i}) \cap V(A_t) = V(C')$. For $i = 0$, this is trivial. For the induction step, suppose that for some $j \in [0, p-1]$ we have already defined $C_{j+1} \subseteq A_{j+1}$; we shall define C_j. Note that C_{j+1} is not necessarily a subgraph of A_j because some of its edges may be introduced by the clique $K[S_j]$ in A_{j+1}. However, if $V(C_{j+1}) \cap S_j = \varnothing$, then $C_{j+1} \subseteq A_j$, and we let $C_j = C_{j+1}$. So assume that $V(C_{j+1}) \cap S_j \neq \varnothing$.

By Lemma 4.8, there exists a connected component D of $A_j \setminus X_j$ such that $S^{A_j}(D) = S_j$. We let $C_j = A_j[V(C_{j+1}) \cup V(D)]$. Since $V(C_{j+1}) \cap S_j \neq \varnothing$, this graph is connected, and obviously we have $V(C_j) \cap V(A_{j+1}) = V(C_{j+1})$. By the induction hypothesis and since $A_p \subseteq A_{j+1}$, the claim follows. □

Corollary 4.10 *Let $D = (T, A)$ be an MS decomposition of a graph G, and let $D' = (T', A')$ be a tree decomposition of G associated with D.*

Let $\{t, u\} \in E(T')$. Then there is a connected subgraph $C \subseteq G$ such that $S^G(C) = A'_t \cap A'_u$ and

$$V(C) \subseteq A'[T'_{tu}].$$

Proof. Let $X' = A'[T'_{ut}]$, $Y' = A'[T'_{tu}]$, and $S = X' \cap Y' = A'_t \cap A'_u$. Then (X', Y') is the separation at $e = \{t, u\}$, and its separator is S. Since D' is associated with D, there is a node s of T such that for the separation (X, Y) at s it holds that $X' \cap V(A_s) = X$ and $Y' \cap V(A_s) = Y$. Since (X, Y) is a minimum substantial separation of A_s, by Lemma 4.8, there is a connected component C' of $A_s[Y \setminus X]$ such that $S^{A_s}(C') = S$. By Lemma 4.9, there is a connected subgraph C'' of G such that $V(C'') \cap V(A_s) = V(C')$. Then $V(C'') \subseteq Y' \setminus X'$, and $S^G(C'') \supseteq S^{A_s}(C') = S$. Let C be a connected component of $G \setminus X'$ such that $C'' \subseteq C$. Then $S^G(C) = S$. □

We are now ready to prove the main result of this section.

Lemma 4.11 *Let $\lambda, \mu \in \mathbb{N}$ and $\kappa = \lambda + \mu + 1$. Let G be a graph that has a tree decomposition over $\mathcal{L}(\lambda, \mu)$, and let $D = (T, A)$ be a κ-CMS tree decomposition of G. Then the companions of all bags of D are in $\mathcal{L}(2\kappa!, \mu)$.*

Proof. Let $t \in V(T)$. Without loss of generality we assume that $|A_t| > 2\kappa!$; otherwise, it is easy to see that $\widehat{A}_t \in \mathcal{L}(2\kappa!)$. For every $s \in S^T(t)$, let (X_s, X'_s) be the separation at the edge $\{t, s\} \in E(T)$ such that $A_t \subseteq X_s$. (To make this unique, we assume without loss of generality that $A_t \not\subseteq A_s$.) Let $S_s = X_s \cap X'_s$ be the corresponding separator, and let C_s be a connected subgraph of $G[X'_s \setminus X_s]$ such that $S^G(C_s) = S_s$ (see Corollary 4.10).

Let (U, B) be a tree decomposition of G over $\mathcal{L}(\lambda, \mu)$. The order of this decomposition is at most κ, because a graph in $\mathcal{L}(\lambda, \mu)$ cannot contain a $\kappa + 1$-clique. We direct every edge $e = \{u, v\} \in E(U)$ as follows: Let (Y_e, Y'_e) be

the separation at e such that $B_u \subseteq Y_e$. Then (Y_e, Y_e') is a separation of order κ, and hence by Lemma 4.7, either $|Y_e \cap A_t| \leq \kappa!$ or $|Y_e' \cap A_t| \leq \kappa!$, but not both, because $|A_t| > 2\kappa!$. We direct e towards u if $|Y_e \cap A_t| > \kappa!$ and towards v otherwise.

There must exist one node $u \in V(U)$ such that all edges incident with u are directed towards u. Let $\mathcal{U} = S^U(u)$. For every $v \in \mathcal{U}$ we let (Y_v, Y_v') be the separation at the edge $\{u, v\} \in E(U)$ such that $B_u \subseteq Y_v$, and we let $T_v = Y_v \cap Y_{v'}$. Then

$$\overline{B}_u = G[B_u] \cup \bigcup_{v \in \mathcal{U}} K[T_v] \in \mathcal{L}(\lambda, \mu).$$

Note that $V(G) = B_u \cup \bigcup_{v \in \mathcal{U}}(Y_v' \setminus Y_v)$, and actually that this is a disjoint union. By the definition of u, for all $v \in \mathcal{U}$ we have $|A_t \cap Y_v'| \leq \kappa!$.

Let $\mathcal{T} = S^T(t)$. We partition $\mathcal{T}$ into three sets as follows:

$$\begin{aligned}
\mathcal{T}_1 &= \{s \in \mathcal{T} \mid V(C_s) \cap B_u = \varnothing\}, \\
\mathcal{T}_2 &= \{s \in \mathcal{T} \mid V(C_s) \cap B_u \neq \varnothing \text{ and } S_s \subseteq B_u\}, \\
\mathcal{T}_3 &= \{s \in \mathcal{T} \mid V(C_s) \cap B_u \neq \varnothing \text{ and } S_s \not\subseteq B_u\}.
\end{aligned}$$

By Lemma 3.1, for every $s \in \mathcal{T}_2 \cup \mathcal{T}_3$, the graph $\overline{B}_u[B_u \cap V(C_s)]$ is connected. Let H_1 be the graph obtained from $\overline{B}_u$ by deleting all vertices not in $A_t \cup \bigcup_{s \in \mathcal{T}} V(C_s)$ and contracting the connected subgraphs $\overline{B}_u[B_u \cap V(C_s)]$ to single vertices x_s, for all $s \in \mathcal{T}_2 \cup \mathcal{T}_3$. Then H_1 is a minor of $\overline{B}_u$ and hence $H_1 \in \mathcal{L}(\lambda, \mu)$. Note that the vertices and edges of $\overline{B}_u[A_t \cap B_u]$ remain unaffected by the deletions and contractions, hence $\overline{B}_u[A_t \cap B_u] \subseteq H_1$.

For every $v \in \mathcal{U}$, we let $K_v = K[A_t \cap Y_v']$. Then $|V(K_v)| \leq \kappa!$. Let

$$H_2 = H_1 \cup \bigcup_{v \in \mathcal{U}} K_v.$$

It is not hard to prove that $H_2 \in \mathcal{L}(\kappa!, \mu)$, because each of the new cliques K_v intersects H_1 in the clique T_v, and in any tree decomposition of a part of H_1 we can attach a new vertex to the vertex whose bag contains T_v and let the bag of the new vertex contain K_v. We omit the details, because we will prove a stronger statement later.

Claim 1.   1. For every $s \in \mathcal{T}_1$, the set S_s induces a clique in H_1 (and hence in H_2).

     2. For every $s \in \mathcal{T}_2$, the vertex x_s is adjacent to all vertices of S_s in H_2.

     3. For every $s \in \mathcal{T}_3$, the vertex x_s is adjacent to all vertices of $S_s \cap B_u$ in H_2.

Proof: To prove (1), note that if $V(C_s) \cap B_u = \varnothing$, then there exists a $v \in \mathcal{U}$ such that $V(C_s) \subseteq Y_v' \setminus Y_v$. Then $S_s = S^G(C_s) \subseteq T_v$.

To prove (2) and (3), we show that for every $s \in \mathcal{T}_2 \cup \mathcal{T}_3$, the vertex x_s is adjacent to all vertices of $S_s \cap B_u$ in H_2. Let $s \in \mathcal{T}_2 \cup \mathcal{T}_3$. As $G[V(C_s) \cup S^G(C_s)]$ is connected, $B_u \cap (V(C_s) \cup S^G(C_s))$ induces a connected subgraph of $\overline{B}_u$. Recall that $S_s = S^G(C_s)$ and hence $S_s \cap B_u \subseteq S^{\overline{B}_u}(B_u \cap V(C_s))$. In the minor H_1, the set $V(C_s) \cap B_u$ is contracted to the vertex x_s. Hence $S_s \cap B_u \subseteq S^{H_1}(x_s)$. As H_1 is a subgraph of H_2, this implies the claim.

13

Claim 2. Let $s \in \mathcal{T}_3$, and let $v \in \mathcal{U}$ such that $S_s \cap (Y_v' \setminus Y_v) = \varnothing$. Then $x_s \in T_v$.

Proof: Let $D_1, \dots, D_m$ be all connected components of $G \setminus B_u$ that contain a vertex in $S_s \cap Y_v'$. Then for $1 \le i \le m$ we have $V(D_i) \subseteq Y_v'$ and $x_s \in S^G(D_i)$. Thus $x_s \in S^G(Y_v' \setminus Y_v) = T_v$.

Note that Claim 2 implies that for every $s \in \mathcal{T}_3$ the vertex x_s is adjacent to all vertices in $S_s \setminus B_u$ in H_2. Let H_3 be the graph obtained from H_2 by adding, for all $s \in \mathcal{T}_1$, a new vertex x_s and edges from x_s to all vertices in S_s.

Claim 3. $\widehat{A}_t$ is isomorphic to a subgraph of H_3.

Proof: Follows from Claim 1 (2) and (3) and from Claim 2.

Claim 4. $H_3 \in \mathcal{L}(\kappa!, \mu)$.

Proof: Let $X \subseteq V(H_1)$ such that $|X| \le \mu$ and $H_1 \setminus X \in \mathcal{L}(\lambda)$. We shall prove that $H_3 \setminus X \in \mathcal{L}(\kappa!)$. As local tree width is monotone with repect to taking subgraphs, it suffices to prove that for every graph H' that is obtained from $H_3 \setminus X$ by contracting edges and for every $r \ge 1$,

$$\mathrm{ltw}(H', r) \le \kappa! \cdot r. \tag{4.9}$$

To simplify the notation, we first prove (4.9) for $H' = H_3 \setminus X$; it will then be easy to see that essentially the same proof goes through for all H' obtained from $H_3 \setminus X$ by contracting edges. So let $H' = H_3 \setminus X$ and $H_1' = H_1 \setminus X$. Let $r \ge 1$ and $v \in V(H')$. If $r = 1$ and $v = x_s$ for an $s \in \mathcal{T}_1$ with $S_s \cap B_u = \varnothing$, then $N_r^{H'}(v) = \{v\} \cup S_s$ and hence $\mathrm{tw}(H'[N_r^{H'}(v)]) \le |N_r^{H'}(v)| - 1 \le \kappa$. Otherwise, there exists a $v_1 \in V(H_1')$ such that $N_r^{H'}(v) \subseteq N_r^{H'}(v_1)$. Hence without loss of generality, we may assume that $v \in V(H_1')$.

Let $H'' = H'[N_r^{H'}(v)]$ and $H_1'' = H_1'[N_r^{H_1'}(v)]$. Let $(P, (C_p)_{p \in V(P)})$ be a tree decomposition of width at most $\lambda \cdot r$ of H_1''. For every $v \in \mathcal{U}$ there is a vertex $p \in V(P)$ such that $T_v \setminus X \subseteq C_p$, because T_v induces a clique in H_1. If $A_t \cap (Y_v' \setminus (Y_v \cup X)) \ne \varnothing$, we attach a new vertex p_v to the tree P, attach it to p, and let $C_{p_v} = (A_t \cap Y_v') \setminus X$. For every $s \in \mathcal{T}_1$, by Claim 1 the set S_s induces a clique in H_1. Hence there is a node $q \in V(P)$ such that $S_s \setminus X \subseteq C_q$. We attach a new vertex q_s to the tree P, attach it to q, and let $C_{q_s} = (S_s \cup \{x_s\}) \setminus X$. This yields a tree decomposition of H''. As the new bags have size at most $\kappa! \ge \kappa + 1$, the width of this tree decomposition is $\max\{\lambda \cdot r, \kappa! - 1\} \le \kappa! \cdot r$.

We can proceed similarly if H' is obtained from $H_3 \setminus X$ by contracting edges, because cliques remain cliques if we contract edges, and because every minor of $H_1 \setminus X$ is also in $\mathcal{L}(\lambda)$.

As $\widehat{A}_t \subseteq H_3$, this completes the proof of the theorem. $\qquad\square$

Remark 4.12 *Note that we actually proved a stronger bound on the local tree width of the companions than claimed in the statement of the lemma: Let $\widehat{A}_t$ be the companion of a bag. Then for every $r \ge 1$ and every minor A' of $\widehat{A}_t$ we have*

$$\mathrm{ltw}(A', r) \le \max\{\lambda \cdot r, 2\kappa!\}.$$

We are now ready to prove the main theorem of this section.

14

Theorem 4.1. *There are computable functions $f, \lambda, \mu : \mathbb{N} \to \mathbb{N}$ and an algorithm that, given a graph G with $K_k \npreceq G$, computes a tree decomposition of G that is weakly over $\mathcal{L}(\lambda(k), \mu(k))$ in time $f(k) \cdot n^{O(1)}$.*

Proof. Let G be a graph such that $K_k \npreceq G$. By Theorem 3.3 there are $\lambda, \mu \in \mathbb{N}$ such that G has a tree decomposition over $\mathcal{L}(\lambda, \mu)$.

For $\kappa := \lambda + \mu + 1$, Lemma 4.11 implies that any κ-CMS tree decomposition is a tree decomposition weakly over $\mathcal{L}(\lambda', \mu)$, with $\lambda' := 2\kappa!$, and by Theorem 4.3 we can compute such k-CMS tree decompositions in time $f(k) \cdot n^{\mathcal{O}(1)}$, for some computable function f. $\qquad\square$

5 Graphs of almost bounded local tree-width

The goal of this section is to prove that there is an fpt-algorithm that, given a graph in $\mathcal{L}(\lambda, \mu)$, computes a set of at most μ vertices such that the graph obtained by deleting these vertices is in $\mathcal{L}(\lambda')$ for a suitable λ' (Corollary 5.16). The proof is based on a deep structure theorem (Theorem 5.1) by Robertson and Seymour that essentially says that any graph either has small tree-width, or contains a large clique minor, or consists of a large planar wall to which subgraphs of small tree-width are attached in a not too complicated way together with a bounded number of elements which may have arbitrary connections to the rest. As the problem can easily be solved if G has low tree-width and no graph in $\mathcal{L}(\lambda, \mu)$ contains a large clique minor, we only have to deal with the third case. This, however, requires some work. In the next subsection we introduce necessary concepts and state Robertson and Seymour's theorem. In Section 5.2 we prove some facts about linkages in walls that are used in Section 5.3 where we finally present the algorithms.

5.1 Walls and Layouts

An *elementary wall of height $h \geq 1$* is a graph defined as in Figure 5.1. A *wall of height h* is a subdivision of an elementary wall of height h. The *perimeter* of a wall is the boundary cycle (cf. Figure 5.2). A *wall in a graph G* is a wall H that is a subgraph of 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 $h \geq 2$ 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.

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

In the following, let H be a wall of height at least 2 in a graph G, and let P be the perimeter of H. Let K' be the unique connected component of $G \setminus P$ that contains $H \setminus P$. The graph $K = K' \cup P$ is called the *compass* of H in G. A *layout of K (with respect to the wall H in G)* is a family $(C, D_1, \ldots, D_m)$ of connected subgraphs of K such that:

1. $K = C \cup D_1 \cup \ldots \cup D_\ell$;

2. $H \subseteq C$, and there is no separation (X, Y) of C of order ≤ 3 with $V(H) \subseteq X$ and $Y \setminus X \neq \varnothing$;
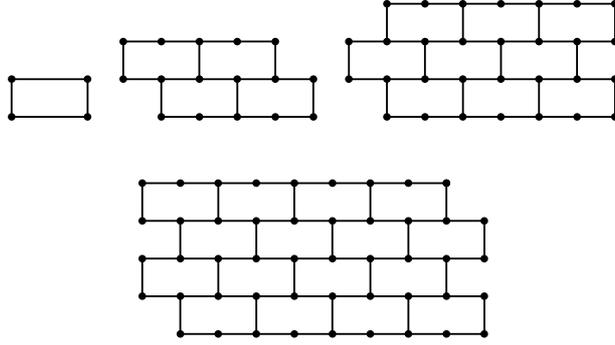
Figure 5.1: Elementary walls of height 1–4
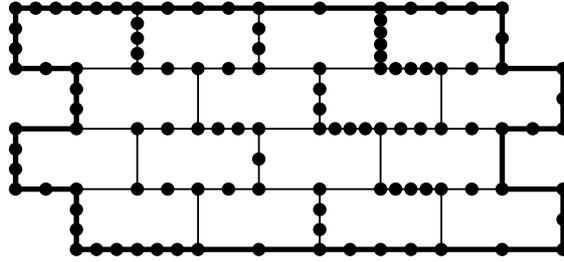
Figure 5.2: A wall of height 4

3. $\partial^G D_i \subseteq V(C)$ for all $i \in [m]$;

4. $|\partial^G D_i| \leq 3$ for all $i \in [m]$;

5. $\partial^G D_i \neq \partial^G D_j$ for all $i \neq j \in [m]$.

We let $\overline{C}$ be the graph obtained from C by adding new vertices $d_1, \ldots, d_m$ and, for $1 \leq i \leq m$, edges between d_i to the vertices in $\partial^G D_i$ and edges between all vertices in $\partial^G D_i$. Hence, for each $i \in [m]$ the vertex d_i together with the (at most 3) vertices in $\partial^G D_i$ form a clique. We call $\overline{C}$ the *core* of the layout and $D_1, \ldots, D_m$ its *extensions*. The layout $(C, D_1, \ldots, D_m)$ is *flat* if its core $\overline{C}$ is planar. Note that this implies that the core has an embedding in the plane that extends the "standard planar embedding" of the wall H (as shown in Figures 5.1 and 5.2), because the wall H has a unique embedding into the sphere. We call the wall H *flat* (in G) if the compass of H has a flat layout.

Theorem 5.1 (Robertson and Seymour [17]) *There are computable functions $f, g : \mathbb{N}^2 \to \mathbb{N}$ and an algorithm A that, given a graph G and nonnegative integers k, h, computes either*

1. *a tree decomposition of G of width $f(k, h)$, or*

2. *a K_k-minor of G, or*

3. *a subset $X \subseteq V(G)$ with $|X| < \binom{k}{2}$, a wall H of height h in $G \setminus X$, and a flat layout $(C, D_1, \ldots, D_m)$ of the compass of H in $G \setminus X$ such that the tree width of each of the extensions $D_1, \ldots, D_m$ is at most $f(k, h)$.*

16

Furthermore, the running time of the algorithm is bounded by $g(k, h) \cdot n^2$, where n is the number of vertices of the input graph G.

Proof. This is (essentially) Lemma (9.8) of [17]. Concerning the uniformity, see the remarks at the end of [17] (on page 109). □

5.2 Linkages in graphs

A *linkage* in graph G is a family of pairwise disjoint paths in G. The *endpoints* of a linkage L are the endpoints of the paths in L, and the *pattern* $\pi(L)$ of L is the matching on the endpoints induced by the paths, that is,

$$\pi(L) = \big\{\{s, t\} \; : \; L \text{ contains a path from } s \text{ to } t\big\}.$$

L is a linkage *between* S and T if $S \cup T$ is the set of endpoints of L, and each path in L has one endpoint in S and one enpoint in T. By $\bigcup L$ we denote the subgraph of G that is the union of the paths in L. The linkage L is *vital* if $V(\bigcup L) = V(G)$ and there is no linkage $L' \neq L$ in G such that $\pi(L') = \pi(L)$.

Lemma 5.2 (Robertson and Seymour [14]) *There is a computable function f such that every graph that has a vital linkage of cardinality k has tree width at most $f(k)$.*

Lemma 5.3 *Let Γ be a plane graph contained in a closed disk D, and let L be a linkage of Γ with the following properties:*

1. *The intersection of Γ with the boundary of D consists of the endpoints of L.*

2. *$V(\bigcup L) = V(\Gamma)$.*

3. *There is no linkage L' of Γ with $\pi(L') = \pi(L)$ and $V(\bigcup L') \neq V(\Gamma)$.*

Then L is vital.

Proof. We prove the lemma by induction on $|L|$. For $|L| = 0$ it is trivial. So suppose that $|L| \geq 1$.

Let C denote the boundary of D. Choose a path $P \in L$ such that one of the two components B_1, B_2 of $C \setminus P$ has an empty intersection with Γ. Say, $B_2 \cap \Gamma = \varnothing$. Let I_2 denote the interior of the circle $B_2 \cup C$. Then $V(\Gamma) \cap I_2 = \varnothing$, because $V(\Gamma) = V(\bigcup L)$ and the paths in L are pairwise disjoint. Furthermore, there is no edge of Γ with interior in I_2, because such an edge would connect two vertices of P and hence allow us to replace P by a strictly shorter path with the same endpoints, which contradicts condition (3). Hence $\Gamma \cap I_2 = \varnothing$. Let D_1 denote the closed disk with boundary $P \cup B_1$. We have just observed that $\Gamma \subseteq D_1$.

Let L' be a linkage of Γ with $\pi(L') = \pi(L)$. We claim that $L = L'$. Let P' be the path in L' that has the same endpoints as P. If $P' \neq P$, then some vertex $v \in P$ is contained in the interior of the disk bounded by $P' \cup B_2$. Then $v \notin \bigcup L'$, which contradicts (3). Hence $P = P'$. Let $\Gamma_1 = \Gamma \setminus P$, $L_1 = L \setminus \{P\}$, and $L_1' = L' \setminus \{P\}$. As $\Gamma_1 \subseteq D_1$, we can apply the induction hypothesis, which implies that $L_1 = L_1'$. Hence $L = L'$. □

17

Let Γ be a plane graph and $C_1, \ldots, C_k$ a sequence of cycles of Γ. We call $C_1, \ldots, C_k$ *concentric* if for all $i \in [k-1]$, the cycle C_i is contained in the interior of C_{i+1} (where the *interior* of a cycle is defined with respect to the drawing of Γ in the plane).

Lemma 5.4 (Robertson and Seymour [15]) *There is a computable function f such that the following holds: Let Γ be a plane graph and L a linkage of Γ. Let $C_1, \ldots, C_{f(k)}$ be concentric cycles of Γ such that all endpoints of the paths in L are contained in the interior of C_1.*

Then there is a linkage L' with the same pattern as L such that all paths in L are contained in the interior of $C_{f(k)}$.

Proof. This is a simplified version of Lemma (3.1) of [15]. For the computability of the function f, see the remarks at the end of [17]. $\qquad\square$

Let Γ be a plane graph, and let $C_1, \ldots, C_k$ be concentric cycles of Γ. For $i \in [k]$, let I_i be the interior of C_i, and let $D_i = I_i \cup C_i$ be the closed disc with boundary C_i. We say that a path P *traverses* $C_1, \ldots, C_k$ if $P \subseteq D_k \setminus I_1$ and the endpoints of P are on C_1 and C_k, respectively. Now suppose P is a path that traverses $C_1, \ldots, C_k$. We think of P as being directed from C_1 to C_k. For each inner point x of P (viewed as a curve in the plane), this gives us an orientation of a small open neighborhood N of x, dividing $N \setminus P$ into a *left* and *right* part. We can extend this orientation to the two endpoints by slightly extending the curve P into the interior of C_1 and the exterior of C_k. Having this orientation, we can distinguish between edges $\Gamma \setminus P$ incident with a vertex of P that *approach P from the left* and edges that *approach P from the right*. Let Γ' be the subgraph of Γ obtained by deleting all edges of $\Gamma \setminus P$ that are incident with a vertex of P and approach P from the left. We call Γ' be the graph obtained from Γ by *cutting $C_1, \ldots, C_k$ along P*.

Lemma 5.5 *Let Γ be a plane graph, and let $C_1, \ldots, C_k$ be concentric cycles of Γ. Let $P_1, \ldots, P_k$ be pairwise disjoint paths in Γ that traverse $C_1, \ldots, C_k$, and let Γ' be the graph obtained from Γ by cutting $C_1, \ldots, C_k$ along P_1. Then $\mathrm{tw}(\Gamma') \geq k - 1$.*

Before we prove the lemma, note that it is easy to see that $\mathrm{tw}(\Gamma) \geq k - 1$ using the robber and cops game: A winning strategy for the robber is to stay on a path P_i and a cycle C_j currently not occupied by the cops.

Proof of Lemma 5.5. We shall prove that for $j \in [k]$ there is a path $Q_j \subseteq C_j \cap \Gamma'$ such that $P_i \cap Q_j \neq \varnothing$ for all $i \in [k]$.

Fix $j \in [k]$. To simplify the notation, we let $P = P_1$ and $C = C_j$. Let $S_1, \ldots, S_m$ be the connected components of $C \setminus P$. For $1 \leq \ell \leq m$, let $\overline{S}_\ell$ be the path extending S_ℓ by the two edges of C that connect S_ℓ with P. Then $C = \bigcup_{\ell=1}^m \overline{S}_\ell$. The crucial observation is that there is precisely one ℓ_0 such that $\overline{S}_{\ell_0}$ leaves P on the right (of P, as defined above) and enters on the left. This can be shown by a fairly straightforward induction on the number of intersections of P and C. Now consider a path P_i for some $i \geq 2$ and suppose for contradiction that P_i does not intersect S_{ℓ_0}. It is easy to see that for all $\ell \in [m]$ with $S_\ell \cap P_i \neq \varnothing$ it holds that $\overline{S}_\ell$ approaches P from the right at both ends. But through such segments, P_i can never reach C_k.

18

We let $Q_j = \overline{S}_{\ell_0} \cap \Gamma'$. Then all P_i (including P_1) have a nonempty intersection with Q_j. Now we can use the robber and cops game to establish that $\mathrm{tw}(\Gamma') \geq k - 1$. $\qquad\square$

The following lemma, which is the main result of this section, says that if we have a linkage in a plane graph that connects vertices on the interior and exterior of a large family of concentric cycles, then there is another linkage with the same pattern that does not cross the "middle cycles" too often. To make this precise, we need more definitions.

Let Γ be a plane graph and $C_1, \ldots, C_\ell$ a sequence of concentric cycles of Γ. For $i \in [\ell]$, let I_i be the interior of C_i and $D_i = I_i \cup C_i$. Let L be a linkage of Γ with endpoints in $I_1 \cup (\mathbb{R}^2 \setminus D_\ell)$. The *traversal linkage* induced by L on $C_1, \ldots, C_\ell$ is the set $T = T(L; C_1, \ldots, C_\ell)$ of all connected components of graphs $P \cap (D_\ell \setminus I_1)$ for $P \in L$. Note that T is a linkage of Γ with endpoints in $V(C_1 \cup C_\ell)$. However, it not necessarily the case that all paths in P traverse $C_1, \ldots, C_\ell$. Also note that the cardinality of T cannot be bounded in terms of $|L|$ and ℓ.

Lemma 5.6 *There is a computable function f such that the following holds for every $k \geq 1$: Let Γ be a plane graph. For some $\ell > 2f(k)$, let $C_1, \ldots, C_\ell$ be a sequence of concentric cycles of Γ. Let L be a linkage of Γ of cardinality k with all endpoints in the interior of C_1 or the exterior of C_ℓ.*

Then there exists a linkage L' of Γ with the same pattern as L such that there are at most $f(k)$ paths in the traversal linkage $T(L'; C_1, \ldots, C_\ell)$ that have a nonempty intersection with one of the cycles $C_{f(k)}, \ldots, C_{\ell-f(k)+1}$.

Proof. Let f_1, f_2 be the functions f of Lemmas 5.2 and 5.4, respectively, and define f by letting $f(k) = \max\{2 \cdot f_1(k) + 1, f_2(k)\}$ for all $k \in \mathbb{N}$.

If L does not contain a path that has one endpoint in the interior of C_1 and one endpoint in the exterior of C_ℓ, then the statement follows by applying Lemma 5.4 twice (once to the interior and once to the exterior). In the following, we assume that there is at least one path in P that has one endpoint in the interior of C_1 and one endpoint in the exterior of C_ℓ.

Let L' be a linkage in Γ with $\pi(L) = \pi(L')$ such that $V(\bigcup L')$ is minimal with respect to set inclusion, that is, there is no linkage L'' of Γ with $\pi(L') = \pi(L'')$ and $V(\bigcup L'') \subset V(\bigcup L')$.

Let Γ' be the topological minor of Γ obtained from the subgraph $\bigcup L' \cup \bigcup_{i=1}^{\ell} C_\ell$ by contracting all segments of the cycles C_i that are not on paths in L' to single edges. We view Γ' as a plane graph that is identical with Γ if considered as a subset of the plane. Let $C_1', \ldots, C_\ell'$ be the cycles of Γ' corresponding to the cycles $C_1, \ldots, C_\ell$.

Then L' is a linkage of Γ' with its endpoints in the interior of C_1' or the exterior of C_ℓ', and there is no linkage L'' of Γ' with $\pi(L') = \pi(L'')$ and $V(\bigcup L'') \subset V(\bigcup L')$. Furthermore, $V(\bigcup L') = V(\Gamma')$. Consider the traversal linkage $T = T(L'; C_1', \ldots, C_\ell')$. Let $P \in T$ be a path with one endpoint on C_1' and one endpoint in C_ℓ'. Such a P exists because L contains a path that has one endpoint in the interior of C_1 and one endpoint in the exterior of C_ℓ. Let Γ'' be the subgraph of Γ' obtained by cutting $C_1', \ldots, C_\ell'$ along P.

Claim 1. L' is a vital linkage of Γ''.

Proof: By the construction, there is no linkage L'' of Γ'' with $\pi(L') = \pi(L'')$ and $V(\bigcup L'') \subset V(\bigcup L')$. It is easy to see that Γ'' can be drawn into a closed disk such that the endpoints of L' are the only points on the boundary of the disk. Hence by Lemma 5.3, L' is vital.

Claim 2. $\operatorname{tw}(\Gamma'') \le f_1(k) < f(k)/2$.

Proof: Follows immediately from Lemma 5.2.

Now let $j \in [f(k), \ell - f(k) + 1]$ and suppose for contradiction that more than $f(k)$ paths in $T(L'; C_1, \ldots, C_\ell)$ have a nonempty intersection with C_j. Then more than $f(k)$ paths in $T(L'; C'_1, \ldots, C'_\ell)$ have a nonempty intersection with C'_j in Γ'. Hence either more than $f(k)/2$ paths in $T(L'; C'_1, \ldots, C'_\ell)$ have a nonempty intersection with all cycles C'_i for all $i \in [j]$ or more than $f(k)/2$ paths in $T(L'; C'_1, \ldots, C'_\ell)$ have a nonempty intersection with all cycles C'_i for all $i \in [j, \ell]$. Then by Lemma 5.5, $\operatorname{tw}(\Gamma'') \ge f(k)/2$. This contradicts Claim 2 $\square$

5.3 Algorithmic Aspects

For every $\ell \ge 1$, the $(\ell \times \ell)$-*pyramid* is the graph obtained from the $(\ell \times \ell)$-grid by adding a new vertex a (the *apex*) and edges from a to all vertices of the grid. For every $\lambda \ge 1$, we let

$$\mathcal{K}(\lambda) = \{G \mid \text{the } (\lambda + 1) \times (\lambda + 1)\text{-pyramid is not a minor of } G\},$$
$$\mathcal{K}(\lambda, \mu) = \{G \mid \exists X \subseteq V(G) : |X| \le \mu \text{ and } G \setminus X \in \mathcal{K}(\lambda)\}.$$

The *bricks* of an elementary wall are the cycles of length 6, and the *bricks* of a wall are the subdivisions of the bricks of the corresponding elementary wall. Two bricks are *adjacent* if they are distinct and have a nonempty intersection. We can assign coordinates to the bricks of a wall as shown in Figure 5.3. In the
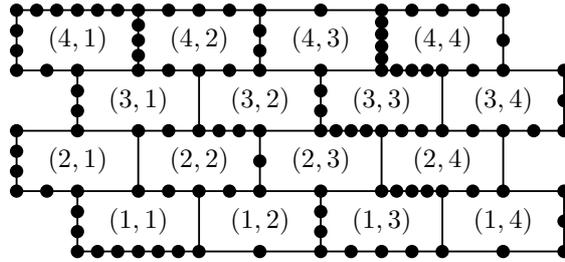
Figure 5.3: The coordinates of the bricks of a wall

folllowing, we always assume that we have fixed the coordinates in our walls. Let H be a wall. For $1 \le i, j \le h$, let B_{ij} denote the brick of H with coordinates (i, j). A subgraph $H' \subseteq H$ is a *subwall* if there exist $h' \le h$ and $i, j \le h - h'$ such that

$$H' = \bigcup_{\substack{i+1 \le i' \le i+h' \\ j+1 \le j' \le j+h'}} B_{i'j'}.$$

Hence a subwall consists of consecutive bricks both in horizontal and vertical direction.

Lemma 5.7 *There is a computable function $f : \mathbb{N}^4 \to \mathbb{N}$ such that the following holds for all $h', \kappa, \lambda, \mu \in \mathbb{N}$: Let $G \in \mathcal{K}(\lambda, \mu)$, $X \subseteq V(G)$ with $|X| \leq \kappa$, and let H be a flat wall of height at least $f(h', \kappa, \lambda, \mu)$ in $G \setminus X$. Then*

1. *either there is a vertex $x \in X$ such that $G \setminus \{x\} \in \mathcal{K}(\lambda, \mu - 1)$,*

2. *or H has a subwall H' of height h' such that the compass of H' in G has an empty intersection with X.*

Furthermore, there is an algorithm with running time $f(h', \kappa, \lambda, \mu) \cdot n^{\mathcal{O}(1)}$ that tests which of the two cases applies and either computes the subwall H' or the element $x \in X$ so that $G \setminus \{x\} \in \mathcal{K}(\lambda, \mu - 1)$.

Proof. Let $g(k, m)$ denote the least n such that if we color $[n]^2$ with k colors then there are subsets $I, J \subseteq [n]$ with $|I| = |J| = m$ such that $I \times J$ is monochromatic. The existence of such a number follows from Ramsey's theorem, but also can easily be proved directly.

We define f by $f(h', \kappa, \lambda, \mu) = (h' + 1) \cdot g(\kappa, (\lambda + 1) \cdot \sqrt{\mu} + 1)$. Let H be a flat wall of height $h \geq f(h', \kappa, \lambda, \mu)$ in $G \setminus X$. For $i, j \in [g(\kappa, (\lambda + 1) \cdot \sqrt{\mu} + 1)]$, let H'_{ij} be the subwall of H of height $h' + 1$ such that the brick in the upper right corner of H'_{ij} has coordinates $(i \cdot (h' + 1), j \cdot (h' + 1))$. Then the subwalls H'_{ij} cover the wall H and have disjoint interiors. Let H_{ij} be the subwall of H'_{ij} of height h' such that the brick in the upper right corner of H_{ij} is the same as the brick in the upper right corner of H'_{ij}. Thus, H_{ij} is the subwall of H'_{ij} with the lowest row and the right-most column removed. By construction, the H_{ij} are disjoint.

Now suppose that (2) does not hold. Then, in particular, for all i, j, the compass of the subwall H_{ij} in G has a nonempty intersection with X. Fix i, j for a moment. As the compass of H_{ij} has a nonempty intersection with X, we can find a shortest path P_{ij} from a vertex y_{ij} of H_{ij} to a vertex $x_{ij} \in X$. As H is flat in $G \setminus X$, we have $V(P_{ij}) \cap V(H) = \{y_{ij}\}$ and $V(P_{ij}) \cap X = \{x_{ij}\}$. We can view the mapping $(i, j) \mapsto x_{ij}$ as a coloring of $[g(\kappa, (\lambda + 1) \cdot \sqrt{\mu} + 1)]^2$ with κ colors. Hence there are subsets $I, J \subseteq [g(\kappa, (\lambda + 1) \cdot \sqrt{\mu} + 1)]$ of size at least $(\lambda + 1) \cdot \sqrt{\mu} + 1$ and an $x \in X$ such that $x = x_{ij}$ for all $i \in I, j \in J$. Pick such I, J, x.

Note that for $i \in I, j \in J$, the paths P_{ij} have pairwise no points other than x in common, for otherwise the wall H would not be flat in $G \setminus X$. Let $G' = H \cup \bigcup_{i \in I, j \in J} P_{ij}$. It is not hard to see that G' has the $((\lambda + 1) \cdot \sqrt{\mu} + 1) \times ((\lambda + 1) \cdot \sqrt{\mu} + 1)$-pyramid with grid vertices y_{ij} and apex x as a minor.

Let $Y \subseteq V(G)$ with $|Y| \leq \mu$ and $G \setminus Y \in \mathcal{K}(\lambda)$. We claim that $x \in Y$. Suppose not. Then the $((\lambda + 1) \cdot \sqrt{\mu} + 1) \times ((\lambda + 1) \cdot \sqrt{\mu} + 1)$-grid has a $(\lambda + 1) \times (\lambda + 1)$-subgrid such that no vertex in Y is contained in any of the connected subgraphs of G' that are contracted to the vertices of the subgrid. Hence $G \setminus Y$ has the $(\lambda + 1) \times (\lambda + 1)$-pyramid with apex x as a minor. This contradicts $G \setminus Y \in \mathcal{K}(\lambda)$. Thus $x \in Y$, which implies (1). $\qquad \square$

A *model* of a graph G in a graph H is a minimal subgraph $M \subseteq H$ such that $G \preceq M$. (Minimality means that there is no $M' \subset M$ with $G \preceq M'$.) With each model M of G in H we associate a mapping that associates a connected subgraph $M(v)$ with each vertex $v \in V(G)$ in such a way that for all distinct $v, w \in V(G)$ the graphs $M(v)$ and $M(w)$ are disjoint, and if $\{v, w\} \in E(G)$ there is an edge from $M(v)$ to $M(w)$ in M. It follows from the minimality condition

in the definition of a model M that $M(v)$ is a tree for every $v \in V(G)$. Note the slight abuse of notation: We use M to denote both a subgraph of H and a mapping from $V(G)$ to the set of connected subgraphs of H.

Lemma 5.8 *Let $\lambda \geq 1$. Let G be a graph, H a wall of height h in G, and $(C, D_1, \ldots, D_m)$ a flat layout of the compass K of H such that every extension D_i has tree width at most $\lambda - 4$. Furthermore, let G' be the graph obtained from G by replacing the compass K of H by $\overline{C}$. That is, for every $i \in [m]$ we delete $D_i \setminus \partial^G D_i$, add a new vertex d_i, and add edges between all pairs of vertices of $\partial^G D_i$ and between d_i and all vertices of $\partial^G D_i$.*

If the $(\lambda + 1) \times (\lambda + 1)$-pyramid is a minor of G then it is also a minor of G'.

Proof. Suppose that the $(\lambda + 1) \times (\lambda + 1)$-pyramid is a minor of G, and let M be a model of the pyramid in G.

Claim 1. For every $i \in [m]$, there is at most one vertex v of the $(\lambda+1) \times (\lambda+1)$-pyramid such that $M(v) \subseteq (D_i \setminus \partial^G D_i)$.

Proof: Let $D = D_i$ and $\overline{D} = D \cup K_{\partial^G D}$. Then $\mathrm{tw}(\overline{D}) \leq \lambda - 2$, and therefore $\overline{D}$ does not have the $(\lambda - 1) \times (\lambda - 1)$-grid as a minor. Hence there are at least two vertices x_1, x_2 of the pyramid such that $M(x_i) \cap V(D) = \varnothing$. Suppose for contradiction that there are two vertices y_1, y_2 of the pyramid with $M(y_i) \subseteq (D \setminus \partial^G D)$. There are at least 4 internally disjoint paths from $\{x_1, x_2\}$ to $\{y_1, y_2\}$ in the pyramid (one through the apex and at least 3 in the grid. This contradicts $|\partial^G D| \leq 3$. $\dashv$

Let M' be defined by letting

$$M'(x) = \begin{cases} M(x) \cap V(C) & \text{if } M(x) \cap V(C) \neq \varnothing, \\ (\{d_i\}, \varnothing) & \text{if } M(x) \subseteq D_i \setminus \partial^G D_i \\ & \quad \text{for some } i \in [m]. \end{cases}$$

Here $(\{d_i\}, \varnothing)$ denotes the graph with vertex d_i and no edges. It is easy to see that M' induces a model of the $(\lambda + 1) \times (\lambda + 1)$-pyramid in G'. $\square$
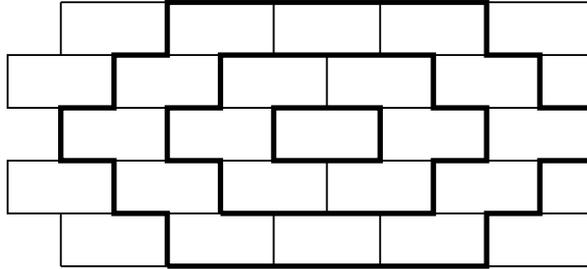
Figure 5.4: The shells of a wall of height 5

Let H be a wall of height h. We view H as a plane graph (drawn into the plane in the natural way). The *central brick* of H is the brick with coordinates $(\lceil h/2 \rceil, \lceil h/2 \rceil)$. The *shells* of H are the cycles $S_0, \ldots, S_{\lceil h/2 \rceil - 1}$, where S_0 is

boundary cycle of the central brick and for $r \geq 1$, the cycle S_r is the shortest cycle disjoint from S_{r-1} with S_{r-1} in its interior. The index r is called the *radius* of the shell. The closed disk bounded by S_r is denoted by D_r and its interior by I_r. Furthermore, we let $\mathcal{S}_0$ be the set consisting of the central brick, and for $r \geq 1$, we let $\mathcal{S}_r$ be the set of all bricks between S_{r-1} and S_r. We shall use the notation S_r, D_r, I_r and $\mathcal{S}_r$ for the rest of this section.

Let $r \geq 1$. The bricks in $\mathcal{S}_r$ are arranged in a cycle. For every brick $B \in \mathcal{S}_r$ there is a unique brick adjacent to B on the *left* of B in clockwise order and a unique brick adjacent to B on the *right* of B in clockwise order. Furthermore, every brick in $\mathcal{S}_r$ is adjacent to one or two bricks in $\mathcal{S}_{r-1}$. Those bricks in $\mathcal{S}_r$ that are only adjacent to one vertex in $\mathcal{S}_{r-1}$ are called *cornerbricks*. Note that $\mathcal{S}_r$ contains exactly six cornerbricks. Finally, a straightforward induction shows that $|\mathcal{S}_r| = 6r$ and that there are exactly $r - 1$ non-cornerbricks between any two adjacent cornerbricks.

We call a tuple $(x_1, \ldots, x_m) \in V(S_r)^m$ *accessible from the interior* if $x_1, \ldots, x_m$ appear on the cycle S_r in clockwise order, and if every brick of $\mathcal{S}_r$ contains at most one vertex x_i. If $r \leq \lceil h/2 \rceil - 2$, we call $(x_1, \ldots, x_m) \in V(S_r)^m$ *accessible from the exterior* if $x_1, \ldots, x_m$ appear on the cycle S_r in clockwise order, and if every brick of $\mathcal{S}_{r+1}$ contains at most one vertex x_i. We call a tuple $(x_1, \ldots, x_m)$ of vertices of the outermost shell $S_{\lceil h/2 \rceil - 1}$ accessible from the exterior if it is accessible from the exterior in all extensions of H by another layer of bricks.

Lemma 5.9 *Let $h, k, \ell, m \in \mathbb{N}$ such that $m \leq k \leq \ell < \lceil h/2 \rceil$ and $\ell - k - 1 \geq (3/2)m$. Let H be a wall of height h, and let $(y_1, \ldots, y_m) \in V(S_k)^m$ be accessible from the exterior and $(y_1', \ldots, y_m') \in V(S_\ell)^m$ be accessible from the interior.*

Then there is a linkage L in $H \cap (D_\ell \setminus I_k)$ with pattern $\big\{ \{y_i, y_i'\} \ : \ i \in [m] \big\}$.

Proof. Let $\tilde{H} = H \cap (D_k \setminus I_i)$. Let $Y = \{y_1, \ldots, y_m\}$ and $Y' = \{y_1', \ldots, y_m'\}$. Observe that for every linkage L between Y and Y' there is an integer p with $-m/2 < p \leq m/2$ such that for all $i \in [m]$ the path in L with endpoint $y_i \in Y$ has endpoint $y'_{((i+p) \bmod m)+1} \in Y'$. This follows easily from the planarity of H. We call p the *offset* of L.

Claim 1. Let $(x_1, \ldots, x_m)$ be an m-tuple on some shell S_r.

1. If $r < \lceil h/2 \rceil - 1$ and $(x_1, \ldots, x_m)$ is accessible from the exterior, then there are m pairwise disjoint paths from $\{x_1, \ldots, x_m\}$ to S_{r-1} in $\bigcup \mathcal{S}_{r+1}$ whose endpoints form a tuple on S_{r-1} that is accessible from the exterior.

2. If $r > m$ and $(x_1, \ldots, x_m)$ is accessible from the interior, then there are m pairwise disjoint paths from $\{x_1, \ldots, x_m\}$ to S_{r-1} in $\bigcup \mathcal{S}_r$ whose endpoints form a tuple on S_{r-1} that is accessible from the inside.

Proof: The first claim is obvious. For the second, observe that between any two cornerbricks of $\mathcal{S}_r$ there are at least m non-corner bricks of $\mathcal{S}_r$, and hence whenever a cornerbrick is occupied, each adjacent side of the hexagon $\mathcal{S}_r$ has at least one brick that contains none of the vertices x_i. This gives us enough room for the desired paths.

Claim 2. There is a linkage between Y and Y'.

Proof: By Menger's theorem, it suffices to prove that there is no separator

of order less than m that separates Y from Y'. This follows easily from Claim 1.

Let $\{P_1, \ldots, P_m\}$ be a linkage between Y and Y' such that for each $i \in [m]$ the vertex y_i is an endpoint of P_i. Let Σ be the cylinder $\{z \in \mathbb{R}^2 \mid 1 \leq ||z|| \leq 2\}$, and let $S = \{z \in \mathbb{R}^2 \mid ||z|| = 1\}$, $S' = \{z \in \mathbb{R}^2 \mid ||z|| = 2\}$ be the boundary cycles. In the following, we shall use polar coordinates to describe the points in Σ. We may assume that $\tilde{H}$ is embedded into Σ such that:

- $\tilde{H} \cap S = Y$, and for all $i \in [m]$, the polar coordinates of y_i are $(1, 2\pi \cdot i/m)$.

- For all $i \in [m]$, the path P_i is the straight line $\{(r, 2\pi \cdot i/m) \mid 1 \leq r \leq 2\}$.

- $\tilde{H} \cap S' = Y'$.

We shall complete the proof by applying a theorem about the existence of linkages on cylinders due to Robertson and Seymour [16]. A continuous function $f : [0, 1] \to \Sigma$ is $\tilde{H}$-*normal* if $f([0,1]) \cap \tilde{H} \subseteq V(\tilde{H})$. Of course here $[0, 1]$ denotes the unit interval $\{x \in \mathbb{R} \mid 0 \leq x \leq 1\}$ and not the set $\{0, 1\}$. We let $\ell(f)$ be the number of times f intersects $\tilde{H}$, that is, $\ell(f) = |\{x \mid f(x) \in V(\tilde{H})\}|$. Furthermore, we let $\vartheta(f)$ be $1/2\pi$ times the total angle (measured counterclockwise) that is covered by a ray from the origin through $f(x)$ as x ranges from 0 to 1. For $1 \leq i \leq m$, let $x_i = (1, 2\pi \cdot i/m + \pi/m)$ and $x_i' = (2, 2\pi \cdot i/m + \pi/m)$. A *helix* is an $\tilde{H}$-normal curve h such that $\ell(h)$ is finite, $h(0) = x_i$ for some $i \in [m]$, and $h(1) = x_j$ for some $j \in [m]$.

Let p be the offset of the linkage $\{P_1, \ldots, P_m\}$. Theorem (5.10) of [16] implies that for every $q \in [-p, m-p]$ there is a linkage with offset q if and only if every helix h satisfies

$$m \cdot \vartheta(h) - \ell(h) \leq q \leq m \cdot \vartheta(h) + \ell(h)$$

The existence of a linkage with offset 0 and hence with the desired pattern follows from the next claim:

Claim 3. Let $h : [0, 1] \to \Sigma$ be a helix. Then

$$m \cdot \vartheta(h) - \ell(h) \leq -p \leq m \cdot \vartheta(h) + \ell(h) \tag{5.1}$$

Proof: Recall that $|p| \leq m/2$. It is easy to see that there is a family of $2m$ pairwise disjoint paths in $\tilde{H}$ from S_k to S_ℓ. If the helix h winds around the cylinder i times, it must intersect all these paths i times and hence $\ell(h) \geq 2m \cdot i$. Thus

$$\ell(h) \geq 2m \cdot \lfloor |\vartheta(h)| \rfloor,$$

and both inequalities in (5.1) hold if $|\vartheta(h)| \geq 1$. If $|\vartheta(h)| < 1$, we use the observation that h has a nonempty intersection with the shells $S_{k+1}, \ldots, S_{\ell-1}$. Thus $\ell(h) \geq (3/2)m$ and therefore

$$\ell(h) - |m \cdot \vartheta(h)| \geq \frac{m}{2} \geq |p|,$$

which implies (5.1).

$\square$

Lemma 5.10 *Let $h, k, \ell, m \in \mathbb{N}$ such that $m \le k \le \ell < \lceil h/2 \rceil$ and $\ell - k - 1 \ge (5/2)m$. Let H be a flat wall of height h in a graph G, and let $y_1, \ldots, y_m \in V(S_k)$, $y'_1, \ldots, y'_m \in V(S_\ell)$ such that:*

1. *The tuple $(y_1, \ldots, y_m)$ is accessible from the exterior.*

2. *The vertices $y'_1, \ldots, y'_m$ appear on S_ℓ in that clockwise order.*

3. *There is a linkage L' such that every path in L' has precisely one end-point in $\{y'_1, \ldots, y'_m\}$ and a nonempty intersection with each of the cycles $S_\ell, \ldots, S_{\ell - m + 1}$.*

Then there is a linkage L in $G \cap (D_\ell \setminus I_k)$ with pattern $\big\{ \{y_i, y'_i\} \,:\, i \in [m] \big\}$.

Proof. Let $Y' = \{y'_1, \ldots, y'_m\}$, and let $X \subseteq V(S_{\ell - m + 1})$ be a set of vertices that contains exactly one vertex from every brick in $\mathcal{S}_{\ell - m + 1}$.

Claim 1. There is a family of m pairwise disjoint paths in $G \cap (D_\ell \setminus I_{\ell - m + 1})$ from Y' to X.

Proof: Suppose for contradiction that there is a separator Z of cardinality at most $m - 1$ that separates Y' from X. Then there is a path P in the linkage L' and a shell $S = S_j$ for some $j \in [\ell, \ell - m + 1]$ such that $Z \cap V(P) = Z \cap V(S) = \varnothing$. Furthermore, it is easy to see that there are m pairwise disjoint paths from S to X. In particular, there is one path Q from S to X with $V(Q) \cap Z = \varnothing$.

Then we obtain a path from Y' to X that does not intersect Z by following P from Y to S, then going to Q on S, and then following Q from S to X. This contradicts Z being a separator of Y' and X.

Now the lemma follows from Lemma 5.9. $\qquad\square$

Let Γ be a plane graph embedded in a closed disk D, and let $x_1, \ldots, x_k$ be the vertices of Γ that appear on the boundary of D in that clockwise order. Let M be a model of Γ in a wall H of height h. Then M is *boundary preserving* if:

- M is contained in the disk bounded by some shell S of H.

- For $i \in [k]$, the intersection of $M(x_i)$ and S has exactly one vertex y_i, and the tuple $(y_1, \ldots, y_k)$ is accessible from the exterior.

- For all $v \in V(\Gamma) \setminus \{x_1, \ldots, x_k\}$, the intersection of $M(v)$ and S is empty.

We say that M *avoids* a set X of vertices of H if $V(M) \cap X = \varnothing$.

Lemma 5.11 *There is a computable function $f : \mathbb{N}^2 \to \mathbb{N}$ such that the following holds for all $m, n \in \mathbb{N}$: Let H be a wall of height $h \ge f(m, n)$ and $X \subseteq V(H)$ with $|X| \le m$. Let Γ be a plane graph with $|V(\Gamma)| \le n$ that is drawn in a closed disk D. Then there is a boundary preserving model M of Γ in H that avoids X.*

Proof. Without loss of generality we may assume that D is the square $\{(x_1, x_2) \in \mathbb{R}^2 \mid 0 \le x_1, x_2 \le 1\}$.

Recall that the ℓ_1-distance between $x = (x_1, x_2), y = (y_1, y_2) \in \mathbb{R}^2$ is $|x_1 - y_1| + |x_2 - y_2|$. Choose $\varepsilon > 0$ such that:

- For all $v, w \in V(\Gamma)$ we have $\ell_1(v, w) > 2\varepsilon \cdot n$.

- For all edges $e, f \in E(\Gamma)$ and all points $x \in e, y \in f$, either $\ell_1(x, y) \geq 2\varepsilon$, or there exists a vertex $v \in V(\Gamma)$ such that $\ell_1(x, v) < 2\varepsilon \cdot n$ and $\ell_1(y, v) < 2\varepsilon \cdot n$.

The existence of such an epsilon follows from a simple compactness argument. Then it is easy to see that Γ has a model in the $\lceil 1/\varepsilon \rceil \times \lceil 1/\varepsilon \rceil$-grid: Think of the grid as being embedded into the unit square in the natural way. Each vertex occupies an $n \times n$ subgrid around its position in the square, and the edges are routed as closely as possible to their curve in the square.

If we further subdivide the grid into an $\lceil (m+1)/\varepsilon \rceil \times \lceil (m+1)/\varepsilon \rceil$ grid, we can easily modify the model in such a way that it avoids a given set X of at most m vertices. Observing that a shell of radius r in a wall has a subwall of height r in its interior, and that each wall of height r has an $r \times r$ grid minor that essentially has the same topology as the wall, it follows that there is a model of Γ in every wall H of height at least $2\lceil (m+1)/\varepsilon \rceil$ that avoids a given set X of at most m vertices. The boundary vertices will appear on the perimeter of H in the right order. Using similar arguments as in Lemma 5.9, it is not hard to turn this model into a bondary preserving model.

It remains to prove that $\lceil 1/\varepsilon \rceil$ has a computable upper bound in terms of m and n. One easy way to obtain such a bound is to enumerate all n-vertex graphs and all their planar drawings, choose a suitable ε for each of them, and then take the maximum $\lceil 1/\varepsilon \rceil$ for all the εs. Here we may view planar drawings as combinatorial objects such as rotation systems, which are then turned into actual drawings in some systematic way. Of course this may give us values of ε which are smaller than necessary, but this does not matter. $\square$

Lemma 5.12 *There is a computable function $f : \mathbb{N}^2 \to \mathbb{N}$ such that for all $\lambda, \mu \geq 1$ the following holds: Let G be a graph, H a wall of height $f(\lambda, \mu)$ in G, and $(C, D_1, \ldots, D_m)$ be a flat layout of the compass K of H such that every extension D_i has tree width at most $\lambda - 4$. Furthermore, let z be a vertex of the central brick of H. Then*

$$G \in \mathcal{K}(\lambda, \mu) \leftrightarrow G \setminus \{z\} \in \mathcal{K}(\lambda, \mu)$$

and if $X \subseteq V(G) \setminus \{z\}$ is such that $|X| \leq \mu$ and $G \setminus (X \cup \{z\}) \in K(\lambda)$ then $G \setminus X \in K(\lambda)$.

Proof. Let f_1, f_2 be the functions f of Lemmas 5.6 and 5.11, respectively. Let

$$\begin{aligned}
\nu &= 10(\lambda + 1)^2, \\
\xi &= 2f_1(\nu) \cdot (f_1(\nu) + 1), \\
\rho &= f_2\big(2f_1(\nu) + 3\nu, \mu + 1\big), \\
h &= \xi + \rho + (5/2) \cdot (2f_1(\nu) + 3\nu),
\end{aligned}$$

and let $f(\lambda, \mu) = h$. Let H be a flat wall of height h in G. Furthermore, let z be a vertex of the central brick of H.

If $G \in \mathcal{K}(\lambda, \mu)$ then trivially $G \setminus \{z\} \in \mathcal{K}(\lambda, \mu)$, because $G \setminus \{z\}$ is a minor of G.

For the backward direction, assume that $G \setminus \{z\} \in \mathcal{K}(\lambda, \mu)$. Let $X \subseteq V(G)$ with $|X| \leq \mu$ and $G \setminus (X \cup \{z\}) \in \mathcal{K}(\lambda)$. Suppose for contradiction that

$G \setminus X \notin \mathcal{K}(\lambda)$. Let $M \subseteq G$ be a model of the $(\lambda + 1) \times (\lambda + 1)$-grid in G. We decompose the graph M into a family of stars and a linkage as follows: Let B be the set of vertices of M of degree at least 3. For every $b \in B$, let N_b denote the subgraph of M with vertex set $N^M(b)$ an edges between b and its neighbors. Let $N_B = \bigcup_{b \in B} N_b$. Let L be the set of connected components of $M \setminus N_B$. Then vertices of graphs in L have degree at most 2. Hence L is a linkage. It is easy to see that $|V(N_B)| \leq \nu$, $|E(N_B)| \leq \nu$, and $|L| \leq \nu$.

By Lemma 5.8 we may assume without loss of generality that each of the extensions D_i of the flat layout of the compass K of H consist of just one vertex. Then K is a planar graph. We fix some embedding into the plane that extends the natural embedding of H and view K as a plane graph. Let $S_1, \ldots, S_{h/2}$ be the shells of H. By the pigeonhole principle, there must be an

$$ i \in \left[\lceil h/2 \rceil - 1 - \nu \cdot \xi, \lceil h/2 \rceil - 1 - \xi \right] $$

such that no $v \in N_B$ is contained in $D_{i+\xi} \setminus I_{i+1}$, where, as usual, D_j denote the closed disk bounded by the shell S_j of H and I_j denotes its interior (for every $j \in [[\lceil h/2 \rceil - 1]]$). Choose such an i and consider the sequence $S_{i+1}, \ldots, S_{i+\xi}$ of concentric cycles. All endpoints of the linkage L are in N_B and hence in the interior of S_{i+1} or the exterior of $S_{i+\xi}$. By Lemma 5.6, there is a linkage L' of G with the same pattern as L such that at most $f_1(\nu)$ paths in the traversal linkage $T(L'; S_{i+1}, \ldots, S_{i+\xi})$ have a nonempty intersection with one of the $2f_1(\nu) \cdot (f_1(\nu) + 1)$ cycles $S_{i+f_1(\nu)}, \ldots, S_{i+\xi-f_1(\nu)}$. We choose such an L'. Applying the pigeonhole principle again, we find a $i' \in [i + f_1(\nu), i + \xi - 2f_1(\nu)]$ and subset $T' \subseteq T(L'; S_{i+1}, \ldots, S_{i+\xi})$ of cardinality $|T'| \leq f(k)$ such that:

- For every $P \in T'$ and every $j \in [i'+1, \ldots, i'+2f_1(\nu)]$ we have $P \cap S_j \neq \varnothing$.

- For every $P \in T(L'; S_{i+1}, \ldots, S_{i+\xi})$, if there exists a $j \in [i' + 1, \ldots, i' + 2f_1(\nu)]$ such that $P \cap S_j \neq \varnothing$, then $P \in T'$.

Let $\sigma = i' + 2f_1(\nu)$. Let $L_1 = \{P \cap D_\sigma \mid P \in L'\}$ and $L_2 = \{P \cap (\mathbb{R}^2 \setminus I_\sigma) \mid P \in L'\}$. Then $|L_1| \leq |T'| + \nu \leq f_1(\nu) + \nu$, because each path in L_1 is either a path of L' that is fully contained in $D_{i'+f_1(\nu)}$ or it contains a path in T'. Let $N_1 = N_B \cap D_\sigma$ and $N_2 = N_B \setminus N_1 = N_B \cap (\mathbb{R}^2 \setminus I_\sigma)$. Let $M_1 = N_1 \cup \bigcup L_1$ and $M_2 = N_2 \cup \bigcup L_2$. Then $M_1 \cup M_2 = \bigcup L' \cup N_B$ is a model of the $(\lambda+1) \times (\lambda+1)$-grid in G.

M_1 is a plane graph drawn in the disk D_σ. By contracting all the paths in L_1 to single edges, we obtain a plane graph M_1', which is identical with M_1 if viewed as a point set in the plane. Since $|L_1| \leq f_1(\nu) + \nu$ and $|V(N_B)| \leq \nu$, we have $|V(M_1)| \leq 2f_1(\nu) + 3\nu$. By Lemma 5.11, there is a boundary preserving model M_1' of M_1 in the disk D_ρ that avoids $X \cup \{z\}$. Suppose that the boundary shell of M_1' is S_ρ. Let $y_1, \ldots, y_m$ be the vertices of M_1 on the boundary cycle S_σ, and let $y_1', \ldots, y_m'$ be the corresponding vertices of M_1' on S_ρ. Then $m \leq 2f_1(\nu) + 3\nu$, and hence $\sigma - \rho - 1 \geq (5/2)m$. By Lemma 5.10, there is a linkage L_3 of G with pattern $\{\{y_i, y_i'\} : i \in [m]\}$ such that $\bigcup L_3 \subseteq D_\sigma \setminus I_\rho$. Then $M_1' \cup \bigcup L_3 \cup M_2$ is a model of the $(\lambda + 1) \times (\lambda + 1)$-grid in $G \setminus (X \cup \{z\})$. This contradicts our assumption that $G \setminus (X \cup \{z\}) \in \mathcal{K}(\lambda)$. □

We are now ready to prove the main result of this section.

Theorem 5.13 *There are computable functions $f, g : \mathbb{N}^2 \to \mathbb{N}$ and an algorithm that, given a graph $G \in \mathcal{K}(\lambda, \mu)$, computes a set $X \subseteq V(G)$ with $|X| \leq \mu$ such that $G \setminus X \in \mathcal{K}(g(\lambda, \mu))$ in time $f(\lambda, \mu) \cdot n^{O(1)}$.*

Proof. Let $k := (\lambda + 1)^2 + \mu + 2$ and let h be "big enough", so that the recursive calls to the algorithms work. (It is easy but tedious to compute the correct value for h.) By Theorem 5.1, G either

1. has tree-width at most $f(k, h)$, or

2. contains a K_k-minor, or

3. there is a subset $X \subseteq V(G)$ with $|X| < \binom{k}{2}$, a wall H of height h in $G \setminus X$, and a flat layout $(C, D_1, \ldots, D_m)$ of the compass of H in $G \setminus X$ such that the tree width of each of the extensions $D_1, \ldots, D_m$ is at most $f(k, h)$.

As no graph in $\mathcal{K}(\lambda, \mu)$ can contain a $K_{(\lambda+1)^2+\mu+2}$ as a minor and the problem can easily be solved for graphs of bounded tree-width, we only have to deal with Case 3). In this case, the algorithm in Theorem 5.1 actually returns the wall and its layout. We can now apply the algorithm from Lemma 5.7. It either returns an element $x \in X$ so that $G \setminus \{x\} \in \mathcal{K}(\lambda, \mu - 1)$ or a subwall H' whose compass has an empty intersection with X. In the first case we have found one of the μ elements and repeat the process on the graph $G \setminus \{x\}$. In the latter case, we call the algorithm recursively on the smaller graph $G \setminus \{z\}$ for some vertex z of the central brick of H'. By Lemma 5.12, $G \in \mathcal{K}(\lambda, \mu)$ if, and only if, $G \setminus \{z\} \in \mathcal{K}(\lambda, \mu)$ and, in addition, if the recursive call to the algorithm on $G \setminus \{z\}$ returns a set X then, by Lemma 5.12 again, $G \setminus X \in \mathcal{K}(\lambda)$. $\square$

The previous theorem gives an fpt algorithm to compute, given a graph $G \in \mathcal{K}(\lambda, \mu)$, a set X of vertices with $|X| \leq \mu$ such that $G \setminus X \in \mathcal{K}(\lambda)$. However, we aim at computing for a given graph $G \in \mathcal{L}(\lambda, \mu)$ a set X so that $G \setminus X \in \mathcal{L}(\lambda)$. The following lemma by Eppstein [6] and Demaine and Hajiaghayi [5] solves this problem.

Lemma 5.14 *There is a computable function $f : \mathbb{N} \to \mathbb{N}$ such that for all $\lambda \in \mathbb{N}$:*
$$\mathcal{L}(\lambda) \subseteq \mathcal{K}(\lambda) \subseteq \mathcal{L}(f(\lambda)).$$

The lemma easily extends to $K(\lambda, \mu)$ and $\mathcal{L}(\lambda, \mu)$.

Corollary 5.15 *There is a computable function $f : \mathbb{N} \to \mathbb{N}$ such that for all $\lambda, \mu \in \mathbb{N}$:*
$$\mathcal{L}(\lambda, \mu) \subseteq \mathcal{K}(\lambda, \mu) \subseteq \mathcal{L}(f(\lambda), \mu).$$

From this, the main result of this section follows immediately.

Corollary 5.16 *There are computable functions $f, g : \mathbb{N}^2 \to \mathbb{N}$ and an algorithm that, given a graph $G \in \mathcal{L}(\lambda, \mu)$, computes a set $X \subseteq V(G)$ with $|X| \leq \mu$ such that $G \setminus X \in \mathcal{L}(g(\lambda, \mu))$ in time $f(\lambda, \mu) \cdot n^{O(1)}$.*

6 FO-model checking on graphs with excluded minors

In [7], Flum and Grohe show that the model-checking problem for first-order logic is fixed-parameter tractable with parameter φ on any class of graphs with an excluded minor. In their proof, Flum and Grohe use the tree-decomposition of graphs excluding a fixed minor guaranteed by Theorem 3.3. The proof can easily be modified to work with tree-decompositions that are weakly over $\mathcal{L}(\lambda, \mu)$. As a consequence, we immediately get the following theorem.

Theorem 6.1 *The following problem*

FO-MODEL-CHECKING	
Input:	Graphs G, H such that $H \not\preceq G$ and $\varphi \in$ FO.
Parameter:	$|\varphi| + |H|$.
Problem:	Decide wether $G \models \varphi$.

is fixed-parameter tractable.

The theorem implies that problems such as the dominating or independent set problem become fixed-parameter tractable when the parameter is the excluded minor and the size of the independent set etc. This improves over previously know results, where the minor was not part of the parameter and determined the exponent of the polynomials.

Another consequence of the methods developed in the previous sections is the following result. For any function $f : \mathbb{N} \to \mathbb{N}$, let $\mathcal{C}_f$ be the class of graphs G such that the excluded clique number of G is at most $f(|G|)$.

Theorem 6.2 *There is an unbounded function $f : \mathbb{N} \to \mathbb{N}$ such that first-order model checking is fixed-parameter tractable on $\mathcal{C}_f$.*

The algorithms presented in the previous sections depend in various ways on the excluded minor H. For instance H determines the numbers λ and μ used throughout the sections. We therefore refrain from giving explicit bounds on the function f whose existence is proved in Theorem 6.2.

7 Locally Excluding a Minor

In [7] Flum and Grohe prove that first-order model checking is fixed-parameter tractable on any class of graphs with an excluded minor. In the same year, Frick and Grohe [10] established the analogous result for graph classes with bounded local tree-width. As the two structural properties are incomparable, i.e. there are classes of graphs excluding a minor but with unbounded local tree-width and vice-versa, it is a natural question, whether there is a common generalisation of excluded minors and bounded local tree-width on which first-order model checking is still fixed-parameter tractable. In this section we present such a generalisation and show that it also generalises graph classes of bounded expansion, a notion recently introduced by Nešetřil and de Mendez [20].

7.1 Definition

Definition 7.1 *A class $\mathcal{C}$ of graphs* locally excludes a minor *if for every $r \in \mathbb{N}$ there is a graph H_r so that if $G \in \mathcal{C}$ and $v \in V(G)$ then $H \npreceq N_r^G(v)$, i.e. H is not a minor of the r-neighbourhood of v in G.*

It is easily seen that any class of graphs with bounded local tree-width locally excludes a minor as does any class of graphs excluding a fixed minor. Another example are graph classes with bounded expansion, introduced by Nešetřil and de Mendez [20]. We recall the definition.

Definition 7.2 *Let G be a graph. A* ball *of G is a subset of vertices inducing a connected subgraph. We denote the set of all families of balls of a graph G by $\mathcal{B}(G)$. Let $\mathcal{P} := \{B_1, \ldots, B_k\}$ be a family of balls of G. The* radius *$\rho(\mathcal{P})$ of $\mathcal{P}$ is defined as $\rho(\mathcal{P}) := \max_{X \in \mathcal{P}}(\rho(G[X]))$. The* quotient *$G/\mathcal{P}$ is defined as the graph with vertex set $1, \ldots, k$ and an edge between vertex i and j if, and only if, there is an edge $\{v, v'\} \in E(G)$ with $v \in B_i$ and $v' \in B_j$.*

Clearly, the quotients of a graph are precisely its minors (up to isomorphism). However, the notion of quotient allows us to define the radius of a quotient, and hence of a minor, as the radius of the family of balls the quotient is taken over.

Definition 7.3 *For every graph G and every $r \in \mathbb{N}$ we define the* greatest reduced average density (grad) *$\nabla_r(G)$ of G with radius r as*

$$\nabla_r(G) := \max\left\{ \frac{|E(G/\mathcal{P})|}{|V(G/\mathcal{P})|} : \mathcal{P} \in \mathcal{B}(G), \rho(\mathcal{P}) \leq r \right\}.$$

We also define $\nabla(G) := \max_r \nabla_r(G) = \max\{\frac{|E(H)|}{|V(H)|} : H \preceq G\}$.

Definition 7.4 *A class $\mathcal{C}$ of graphs has* bounded expansion *if there exists a function $b : \mathbb{N} \to \mathbb{N}$ such that $\nabla_r(G) \leq b(r)$ for all $G \in \mathcal{C}$ and $r \in \mathbb{N}$.*

Note that any class of graphs excluding a fixed minor has bounded expansion (indeed, bounded by a constant) and every class of graphs with bounded degree has bounded expansion (by an exponential function). However, the class of graph classes with bounded expansion does not include all graph classes with bounded local tree-width. The next lemma shows that any class of graphs with bounded expansion also locally excludes minors.

Lemma 7.5 *Let $\mathcal{C}$ be a class of graphs with bounded expansion and let $b : \mathbb{N} \to \mathbb{N}$ be a function witnessing this. Then for every $r \in \mathbb{N}$ there exists a graph H_r so that if $G \in \mathcal{C}$ and $v \in V(G)$ then $H_r \npreceq N_r^G(v)$.*

Proof. Let $G \in \mathcal{C}$ be a graph and $r \in \mathbb{N}$. As the expansion of $\mathcal{C}$ is bounded by b, we have $\nabla_r(G) \leq b(r)$. Let $N := N_r(v)$ be the r-neighbourhood of a vertex $v \in V(G)$. Hence, $\frac{|E(N)|}{|V(N)|} \leq b(r)$ and the same holds for all minors of N. It follows that $K_{b(r)+2}$ is not a minor of N. $\square$

Corollary 7.6 *If $\mathcal{C}$ is a class of graphs with*

- *bounded expansion or*

- *bounded local tree-width or*

- *excluding a fixed minor,*

then $\mathcal{C}$ locally excludes a minor. The converse is not true, i.e. there are classes of graphs locally excluding a minor with unbounded expansion or unbounded local tree-width or whose minor closure is the class of all graphs.

Note, however, that the concept of excluding slowly growing minors as we considered in Section 6 is incomparable to locally excluding a minor.

7.2 First-Order Model-Checking

We show next that first-order model checking is fixed-parameter tractable on any class of graphs locally excluding a minor.

Gaifman [11] showed that any first-order sentence is equivalent to a Boolean combination of basic-local sentences. We recall the definition. For every $r \geq 0$ we will use formulas $d(x, y) \leq r$ and $d(x, y) > r$ to say that the distance between x and y is at most r and greater than r, respectively. Clearly, these are easily first-order definable. If $\varphi(x)$ is a first-order formula, then $\varphi^{N_r(x)}(x)$ is the formula obtained from φ by relativising the quantifiers in φ to the r-neighbourhood of x, i.e. replacing $\forall y \psi$ by $\forall y(d(x, y) \leq r \rightarrow \psi)$ and $\exists y \psi$ by $\exists y(d(x, y) \leq r \wedge \psi)$. A formula $\psi(x)$ of the form $\varphi^{N_r(x)}(x)$ is called r-*local*. The essential property of an r-local formula is that its truth value at a vertex x in G only depends on the r-neighbourhood of x in G.

Theorem 7.7 (Gaifman [11]) *Every first-order sentence is equivalent to a Boolean combination of basic-local sentences, i.e. a Boolean combination of sentences of the form*

$$\exists x_1 \ldots \exists x_k \big(\bigwedge_{1 \leq i < j \leq k} d(x_i, x_j) > 2r \wedge \bigwedge_{1 \leq i \leq k} \vartheta(x_i) \big),$$

for suitable $r, k > 0$ and an r-local formula $\vartheta(x)$.

We are now ready to prove the main result of this section.

Theorem 7.8 *Let $\mathcal{C}$ be a class of graphs locally excluding a minor. Then the following problem*

FO-MODELCHECKING ON $\mathcal{C}$	
Input:	$G \in \mathcal{C}$, $\varphi \in$ FO.
Parameter:	$\lvert \varphi \rvert$.
Problem:	Decide $G \models \varphi$.

is fixed-parameter tractable.

Proof. Let $\varphi \in$ FO be a sentence. By Gaifman's theorem 7.7, φ is equivalent to a Boolean combination of basic local sentences. Hence, to prove the theorem, it suffices to only consider the case where φ is a sentence of the form

$$\exists x_1 \ldots \exists x_k \bigwedge_{1 \leq i < j \leq k} d(x_i, x_j) > 2 \cdot r \wedge \bigwedge_{1 \leq i \leq k} \vartheta(x_i)$$

for $r, k \geq 1$ and an r-local formula ϑ.

As $\mathcal{C}$ locally excludes a minor, there is for every $s \in \mathbb{N}$ a graph H_s such that H_s is excluded in every s-neighbourhood of vertices in any member of $\mathcal{C}$. Let $G \in \mathcal{C}$.

The first step of the evaluation algorithm is to compute the set $P \subseteq V(G)$ of vertices v such that $G[N_r(v)] \models \vartheta(v)$. As $H_r \npreceq G[N_r(v)]$, Theorem 6.1 implies that there is a computable function $f : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ such that checking whether $G[N_r(v)] \models \vartheta(v)$ can be done in time $f(|H_r|, |\vartheta|) \cdot |N_r(v)|^{\mathcal{O}(1)}$.

It remains to find a set of k elements of P whose distance is pairwise $> 2r$. For this, we proceed as follows. Set $Q := P$ and set $l := 0$. While $Q \neq \varnothing$ and $l < k$, choose an arbitrary element $a_l \in Q$, increase l to $l+1$ and remove $N_r^G(a_l)$ from Q. If this process stops with $l = k$, we can accept, as then $\{a_1, \ldots, a_k\}$ is the required set. If $l = 0$, then $P := \varnothing$ and therefore ϑ is false in every r-neighbourhood of a vertex in G and hence $G \not\models \varphi$. Finally, if $0 < l < k$, we know that every $v \in P$ is contained in the $2r$-neighbourhood of some a_i, $1 \leq i \leq l$. Let $N := G\big[N_{2r}^G(\{a_1, \ldots, a_l\})\big]$. By construction, the radius of N is at most $2r \cdot l$ and hence $H := H_{2rl} \npreceq N$. By Theorem 6.1, there is a function $g : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ such that we can test whether $(N, P) \models \psi$ in time $g(|H|, |\psi|) \cdot |N|^{\mathcal{O}(1)}$, where $\psi := \exists x_1 \ldots \exists x_k \big(\bigwedge_{i=1}^k Px_i \wedge \bigwedge_{1 \leq i < j \leq k} d(x_i, x_j) > r\big)$. If $(N, P) \models \psi$, then there is a set of k vertices in P pairwise far apart and we can accept. Otherwise we reject.

Note that k and hence $|\psi|, |H_r|$, and $|H|$ only depend on φ and hence are bounded by the parameter $|\varphi|$. The algorithm correctly determines whether $G \models \varphi$ and has a total running time of $h(|\varphi|) \cdot |G|^{\mathcal{O}(1)}$, where $h : \mathbb{N} \to \mathbb{N}$ is a function that dominates the functions f and g above. $\qquad\square$

As an immediate consequence of this we obtain that on classes $\mathcal{C}$ locally excluding a minor the following problems are fixed parameter tractable with parameter H: For every fixed graph H decide whether for a graph $G \in \mathcal{C}$: H has a homomorphism to G; H is a subgraph of G; H is an induced subgraph of G.

Furthermore, problems such as independent or dominating set and many others are fixed-parameter tractable on any class of graphs locally excluding a minor.

8 Conclusions

We introduce the notion of graph classes locally excluding a minor and prove that deciding first-order properties of such classes is fixed-parameter tractable. The result is particularly interesting because it unifies incomparable previous results for classes of bounded local tree width [10] and for classes with excluded minors [7] in a natural way. But the result is considerably stronger than just a combination of those two. In particular, it also covers all classes of bounded expansion.

To prove the result, we need to strengthen the fixed-parameter tractability result for classes with excluded minors [7] in such a way that the size of the excluded minor can now be taken as a parameter in the running time analysis. This implies fixed-parameter tractability results for problems such as dominating set and independent set, now parameterized by the size of the desired solution and the size of the excluded minor. Even though both problems have been intensely

studied on restricted graph classes including classes with excluded minors (see e.g. [13, 1, 9, 8] and the references there), the existence of such algorithms was not known before.

Let us finally remark that algorithmic meta theorems like ours are not meant to be practical, as usually the dependence of the running time on the formula size is nonelementary and the hidden constants are enormous. One reason for the interest in such results is that they often provide an easy way to quickly check if a concrete problem is fixed-parameter tractable (see our remarks on dominating set and independent set above). The more significant reason for our interest in such meta theorems is that they yield a better understanding of the limits of general algorithmic techniques and, in some sense, the limits of tractability. In particular, they clarify the interactions between logic and combinatorial structure, which we believe to be fundamental for computational complexity.

References

[1] J. Alber, H. Fernau, and R. Niedermeier. Parameterized complexity: exponential speed-up for planar graph problems. *Journal of Algorithms*, 52(1):26–56, 2004.

[2] B. Courcelle. Graph rewriting: An algebraic and logic approach. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, pages 194 – 242. Elsevier, 1990.

[3] B. Courcelle, J. Makowsky, and U. Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory Comput. Systems*, 33:125–150, 2000.

[4] E. Demaine, M. Hajiaghayi, and K. Kawarabayashi. Algorithmic graph minor theory: Decomposition, approximation, and coloring. In *Symposium on Foundations of Computer Science (FOCS)*, pages 637–646, 2005.

[5] E. Demaine and M.T. Hajiaghayi. Equivalence of local treewidth and linear local treewidth and its algorithmic applications. In *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 840–849, 2004.

[6] D. Eppstein. Diameter and treewidth in minor-closed graph families. *Algorithmica*, 27:275–291, 2000.

[7] J. Flum and M. Grohe. Fixed-parameter tractability, definability, and model checking. *SIAM Journal on Computing*, 31:113 – 145, 2001.

[8] J. Flum and M. Grohe. *Parameterized Complexity Theory*. 2006.

[9] F. Fomin and D.M. Thilikos. A simple and fast approach for solving problems on planar graphs. In *Symp. on Theoretical Aspects of Computer Science (STACS)*, pages 56–67, 2004.

[10] M. Frick and M. Grohe. Deciding first-order properties of locally tree-decomposable structures. *Journal of the ACM*, 48:1184 – 1206, 2001.

[11] H. Gaifman. On local and non-local properties. In J. Stern, editor, *Herbrand Symposium, Logic Colloquium '81*, pages 105 – 135. North Holland, 1982.

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

[13] M. Hajiaghayi. *The bidimensionality Theory and Its Algorithmic Applications*. PhD thesis, Department of Mathematics, Massachusetts Institute of Technology, 2005.

[14] N. Robertson and P.D. Seymour. Graph minors XXI. Graphs with unique linkages. To appear.

[15] N. Robertson and P.D. Seymour. Graph minors XXII. Irrelevant vertices in linkage problems. To appear.

[16] N. Robertson and P.D. Seymour. Graph minors VI. Disjoint paths across a disc. *Journal of Combinatorial Theory, Series B*, 41:115–138, 1986.

[17] N. Robertson and P.D. Seymour. Graph minors XIII. The disjoint paths problem. *Journal of Combinatorial Theory, Series B*, 63:65–110, 1995.

[18] N. Robertson and P.D. Seymour. Graph minors XVI. Excluding a non-planar graph. *Journal of Combinatorial Theory, Series B*, 77:1–27, 1999.

[19] D. Seese. Linear time computable problems and first-order descriptions. *Mathematical Structures in Computer Science*, 2:505 – 526, 1996.

[20] J. Nešetřil and P. Ossona de Mendez. The grad of a graph and classes with bounded expansion. In *International Colloquium on Graph Theory*, pages 101 – 106, 2005.