ULB FACULTÉ DES SCIENCES
UNIVERSITÉ LIBRE DE BRUXELLES

Minor-closed classes of graphs : isometric embeddings, cut dominants and ball packings

Thesis presented by Carole MULLER
in fulfilment of the requirements of the PhD Degree in Sciences ("Docteur en sciences")
Année académique 2020-2021

Supervisor : Professor Samuel FIORINI
Co-supervisor : Professor Gwenaël JORET

Fonds National de la Recherche Luxembourg

FACULTÉ DES SCIENCES

UNIVERSITÉ LIBRE DE BRUXELLES

Minor-closed classes of graphs : isometric embeddings, cut dominants and ball packings

Thesis presented by Carole MULLER

in fulfilment of the requirements of the PhD Degree in Sciences (”Docteur en sciences”)

Année académique 2020-2021

Supervisor : Professor Samuel FIORINI

Co-supervisor : Professor Gwenaël JORET

Thesis jury :

Samuel FIORINI (Université libre de Bruxelles, Supervisor)
Gwenaël JORET (Université libre de Bruxelles, Secretary, Co-supervisor)
Michele D'ADDERIO (Université libre de Bruxelles, Chair)
Antonios VARVITSIOTIS (Singapore University of Technology and Design)
Stefan WELTGE (Technical University Munich)

Fonds National de la Recherche Luxembourg

Summary

A class of graphs is closed under taking minors if for each graph in the class and each minor of this graph, the minor is also in the class. By a famous result of Robertson and Seymour, we know that characterizing such a class can be done by identifying a finite set of minimal excluded minors, that is, graphs which do not belong to the class and are minor-minimal for this property.

In this thesis, we study three problems in minor-closed classes of graphs. The first two are related to the characterization of some graph classes, while the third one studies a packing-covering relation for graphs excluding a minor.

In the first problem, we study isometric embeddings of edge-weighted graphs into metric spaces. In particular, we consider ℓ_2- and ℓ_∞-spaces. Given a weighted graph, an isometric embedding maps the vertices of this graph to vectors such that for each edge of the graph the weight of the edge equals the distance between the vectors representing its ends. We say that a weight function on the edges of the graph is a realizable distance function if such an embedding exists. The minor-monotone parameter $f_p(G)$ determines the minimum dimension k of an ℓ_p-space such that any realizable distance function of G is realizable in ℓ_p^k. We characterize graphs with large $f_p(G)$ value in terms of unavoidable minors for $p = 2$ and $p = \infty$. Roughly speaking, a family of graphs gives unavoidable minors for a minor-monotone parameter if these graphs "explain" why the parameter is high.

The second problem studies the minimal excluded minors of the class of graphs such that $\varphi(G)$ is bounded by some constant k, where $\varphi(G)$ is a parameter related to the cut dominant of a graph G. This unbounded polyhedron contains all points that are componentwise larger than or equal to a convex combination of incidence vectors of cuts in G. The parameter

$\varphi(G)$ is equal to the maximum right-hand side of a facet-defining inequality of the cut dominant of G in minimum integer form. We study minimal excluded graphs for the property $\varphi(G) \leqslant 4$ and provide also a new bound of $\varphi(G)$ in terms of the vertex cover number.

The last problem has a different flavor as it studies a packing-covering relation in classes of graphs excluding a minor. Given a graph G, a ball of center v and radius r is the set of all vertices in G that are at distance at most r from v. Given a graph and a collection of balls, we can define a hypergraph $\mathcal{H}$ such that its vertices are the vertices of G and its edges correspond to the balls in the collection. It is well-known that, in the hypergraph $\mathcal{H}$, the transversal number $\tau(\mathcal{H})$ is at least the packing number $\nu(\mathcal{H})$. We show that we can bound $\tau(\mathcal{H})$ from above by a linear function of $\nu(\mathcal{H})$ for every graphs G and ball collections $\mathcal{H}$ if the graph G excludes a minor, solving an open problem by Chepoi, Estellon et Vaxès.

Résumé

Une classe de graphes est close par mineurs si, pour tout graphe dans la classe et tout mineur de ce graphe, le mineur est également dans la classe. Par un fameux théorème de Robertson et Seymour, nous savons que caractériser une telle classe peut être fait à l'aide d'un nombre fini de mineurs exclus minimaux. Ceux-ci sont des graphes qui n'appartiennent pas à la classe et qui sont minimaux dans le sens des mineurs pour cette propriété.

Dans cette thèse, nous étudions trois problèmes à propos de classes de graphes closes par mineurs. Les deux premiers sont reliés à la caractérisation de certaines classes de graphes, alors que le troisième étudie une relation de "packing-covering" dans des graphes excluant un mineur.

Pour le premier problème, nous étudions des plongements isométriques de graphes dont les arêtes sont pondérées dans des espaces métriques. Principalement, nous nous intéressons aux espaces ℓ_2 et ℓ_∞. Étant donné un graphe pondéré, un plongement isométrique associe à chaque sommet du graphe un vecteur dans l'autre espace de sorte que pour chaque arête du graphe le poids de celle-ci est égal à la distance entre les vecteurs correspondant à ses sommets. Nous disons qu'une fonction de poids sur les arêtes est une fonction de distances réalisable s'il existe un tel plongement. Le paramètre $f_p(G)$ détermine la dimension k minimale d'un espace ℓ_p telle que toute fonction de distances réalisable de G peut être plongée dans ℓ_p^k. Ce paramètre est monotone dans le sens des mineurs. Nous caractérisons les graphes tels que $f_p(G)$ a une grande valeur en termes de mineurs inévitables pour $p = 2$ et $p = \infty$. Une famille de graphes donne des mineurs inévitables pour un invariant monotone pour les mineurs, si ces graphes "expliquent" pourquoi l'invariant est grand.

Le deuxième problème étudie les mineurs exclus minimaux pour la classe de graphes avec $\varphi(G)$ borné par une constante k, où $\varphi(G)$ est un paramètre lié

au dominant des coupes d'un graphe G. Ce polyèdre contient tous les points qui, composante par composante, sont plus grands ou égaux à une combination convexe des vecteurs d'incidence de coupes dans G. Le paramètre $\varphi(G)$ est égal au membre de droite maximum d'une description linéaire du dominant des coupes de G en forme entière minimale. Nous étudions les mineurs exclus minimaux pour la propriété $\varphi(G) \leqslant 4$ et montrons une nouvelle borne sur $\varphi(G)$ en termes du "vertex cover number".

Le dernier problème est d'un autre type. Nous étudions une relation de "packing-covering" dans les classes de graphes excluant un mineur. Étant donné un graphe G, une boule de centre v et de rayon r est l'ensemble de tous les sommets de G qui sont à distance au plus r de v. Pour un graphe G et une collection de boules donnés nous pouvons définir un hypergraphe $\mathcal{H}$ dont les sommets sont ceux de G et les arêtes correspondent aux boules de la collection. Il est bien connu que dans l'hypergraphe $\mathcal{H}$, le "transversal number" $\tau(\mathcal{H})$ vaut au moins le "packing number" $\nu(\mathcal{H})$. Nous montrons une borne supérieure sur $\nu(\mathcal{H})$ qui est linéaire en $\tau(\mathcal{H})$, résolvant ainsi un problème ouvert de Chepoi, Estellon et Vaxès.

Acknowledgements

A PhD is a long and exciting journey. I am grateful for the last four years in which I had many opportunities to meet and work with great mathematicians and computer scientists.

First of all, I have to thank my supervisors Samuel Fiorini and Gwenaël Joret for their support throughout the years. I was lucky to have them already supervise my Master Thesis and without them I may not have started my PhD. They accompanied me trough this adventure and encouraged me to go on whenever I felt like I reached my limits. They allowed me to broaden my mathematical horizon by traveling to various conferences, workshops, and summer schools, which I deeply appreciated.

Samuel, you were one of my first professors at university and I still remember how you started the very first class. During the years you have taught me much about polytopes and related topics. I have learned a lot from you during many discussions, and without you I wouldn't present a chapter about cut dominants. I deeply appreciated working with you and your ERC team, the reading seminars you organized, and the workshops abroad. Furthermore, I have to thank you for your financial support for my travels during the first two years.

It is thanks to Gwen that I turned my research towards graph theory. Your class about graph theory really made me curios and I wanted to acquire further knowledge in that domain. Moreover, you let me follow some Algo lunch seminars back then, which made me realize that I did not want to stop my studies right after my Master degree. During my Master's thesis and PhD your office door has always been open to me if I needed to discuss anything. Furthermore, you gave me the possibility to do reserach with some of your visitors as well, which led to Chapter 5 of this thesis.

Thank you both for everything!

Second, I thank Michele D'Adderio, Antonios Varvitsiotis and Stefan Weltge for accepting to be in my jury. Furthermore, I thank Michele and Jean Cardinal for being part of my comité d'accompagnement.

My PhD was funded by the Fonds National de la Recherche, Luxembourg by an AFR research grant (grant number 11628910). I thank the FNR for giving me the possibility to spend four years exploring the world as a researcher.

Doing research should not be a lonely activity. Even though for the last one and a half year we were more or less banned from our offices, I want to thank all my collegues from the eighth floor and beyond for sharing lunches and various discussions. Organising the Brussels Summer School of Mathematics and the SPP seminars was very exciting and I thank all co-organisers and speakers for their work.

Many thanks should go to friends who ensured that research visits and conferences did not end with the last talk of the day. I thank Marthe for welcoming me in Bordeaux in January 2020. Together with the other visitors and locals I spent a great time doing maths and tasting all kinds of delicious food. A good card game, a piece of chocolate, a cup of tea or coffee, or a glass of beer should never be missing wherever reasearch is done.

I thank my co-authors (by paper) Gwenaël, Samuel, and Tony [41], François, Gwenaël, Louis, Nicolas, William, and Wouter [11], Alexandra, Carla, Jakub, Jonathan, and Marthe [8], as well as Samuel for the results presented in Chapter 4.

Life is not all about work. I thank the musicians from the ULB orchestra and from Crazy Cube for letting me play along. Furthermore, I thank all my long-time friends for their support during the years. I thank Yannick for the many shared car rides from and to Brussels. I should also thank more recent friends who shared laughs and games with me in the last 1,5 year. Credits should also go to my family, who have always supported me, and my furry motivators, namely Struppi, Kaffi, and Schocki. Lastly, I thank Gilles for his support, patience, and love.

Contents

Chapter 1

Introduction and main results of this thesis

Although graphs have been studied for the last 300 years, it is a fast-growing topic nowadays. Indeed, with the appearance of computers and an interconnected world, graphs have raised to be one of the standard tools to model and study connections in many areas of everyday life. For instance, finding a best itinerary is done by computing a shortest path in a graph modeling the streets in the relevant part of the world. Another common example of a graph is the family tree, which is used in genealogy to represent parent-child relations. These graphs are studied because of their practical purposes. However, mathematicians often study graphs for their own sake or to prove other theoretical results, without a real-life application in mind. Structural graph theory, combinatorial optimization, combinatorics, spectral graph theory, game theory, complexity are only some areas in which graphs are studied from very different theoretical points of view. This is why the interest in graphs has gained much popularity in the last century.

In this thesis we build on two important results of the last decades. First, we study two applications of the Graph Minor Theorem of Robertson and Seymour [67]. These problems are related to isometric embeddings in metric spaces, and to cut dominants, respectively. Second, we establish an Erdős-Pósa property for balls in graphs excluding a minor. This property is named after Erdős and Pósa who established a similar relation for packing and covering cycles in graphs in 1965 [39]. Both these papers have had profound impact on research in graph theory during the last decades (see Sections 2.3

and 12.6 in [31]).

First, we look further into the Graph Minor Theorem, which is a key result used in two chapters. After that, we also introduce the Erdős-Pósa property, which we establish for balls in graphs excluding a minor.

1.1 The Graph Minor Theorem

Robertson and Seymour published a series of twenty-three papers from 1983 to 2010 establishing several milestones in graph structure theory. Among them is Wagner's conjecture, which was proved in 2004 [67]. Wagner conjectured in 1970 [76] that for every infinite set of finite graphs, one of its members if isomorphic to a minor of another. Recall that a minor of a graph G is a graph H that can be obtained from G by edge deletions and contractions and vertex deletions in any order. Theorem 1.1 below, the *Graph Minor Theorem*, is equivalent to Wagner's conjecture.

Given a class of graphs $\mathcal{G}$, we say that $\mathcal{G}$ is *closed under taking minors* or *minor-closed* if given a graph $G \in \mathcal{G}$, every proper minor H of G is also in the class, $H \in \mathcal{G}$. We say that a graph G is an *excluded minor of* $\mathcal{G}$ if $G \notin \mathcal{G}$. G is a *minimal excluded minor* if G is an excluded minor of $\mathcal{G}$ and G is minor-minimal, that is for every proper minor H of G we have $H \in \mathcal{G}$. Observe that given a minor-closed class of graphs $\mathcal{G}$, any graph that contains a minimal excluded minor of $\mathcal{G}$ as a minor is not in the class by transitivity of the minor relation. Such a graph is called an *excluded* or *forbidden* minor. The set of minimal excluded minors for $\mathcal{G}$ is also sometimes referred to as the *obstruction set of* $\mathcal{G}$ in the literature. It is non-trivial that the obstruction set for every minor-closed class of graphs is always finite. Robertson and Seymour [67] proved exactly this.

Theorem 1.1 (Graph Minor Theorem)**.** *Let $\mathcal{G}$ be a minor-closed class of graphs. Then the set of minimal excluded minors for $\mathcal{G}$ is finite.*

We will also talk of *minor-closed properties* in the following. We say that a property is minor-closed or closed under taking minors if the class of graphs satisfying this property is closed under taking minors. An easy example of a minor-closed class of graphs is the class of forests. By definition, a graph is a forest if it does not contain a cycle. Hence, K_3 is the (only) minimal excluded minor.

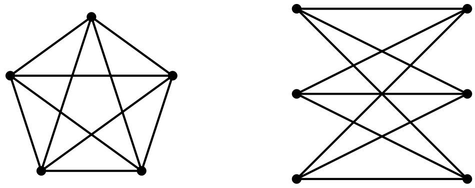

Figure 1.1. The graphs K_5 and $K_{3,3}$ are the minimal excluded minors for planarity.

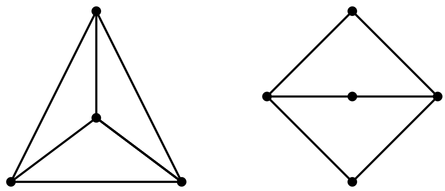

Figure 1.2. The graphs K_4 and $K_{2,3}$ are the minimal excluded minors for outerplanarity.

Planar graphs form a minor-closed class that has two minimal excluded minors. It is easy to verify that, given a planar graph G, all minors of G are also planar. Furthermore, the graphs K_5 and $K_{3,3}$ are not planar. Wagner [75] showed in 1937 that a graph is planar if and only if it does not contain a K_5 or $K_{3,3}$ graph as a minor. The graphs K_5 and $K_{3,3}$ are shown in Figure 1.1.

As for outerplanar graphs, which are planar graphs such that there exists a drawing with all vertices on the outer face, it is known that the graphs K_4 and $K_{2,3}$ shown in Figure 1.2 are the minimal excluded minors, see [31, Exercise 23 in Chapter 4]. More generally, graphs of bounded genus also form a minor-closed class implying that there exists a finite set of minimal excluded minors for the set of graphs with genus at most k for every fixed $k \in \mathbb{N}$.

Our examples may suggest that usually the list of minimal excluded minors is small for minor-closed properties. However, there exist also classes of graphs for which thousands of minimal excluded minors are known and, despite this, completeness of the set has not yet been proven. Such examples include the class of $Y\Delta Y$-reducible graphs or the class of apex graphs, which are graphs such that there exists a vertex whose deletion results in a planar graph. Examples of apex graphs include K_5 and $K_{3,n}$ for all $n \in \mathbb{N}$. A graph is $Y\Delta Y$-reducible if it can be reduced to isolated vertices

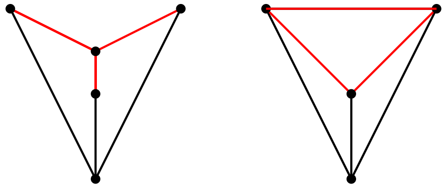

Figure 1.3. A Y–Δ operation consists of deleting a degree-3 vertex and its incident edges, and adding edges between all pairs of its neighbors (from left to right). A Δ–Y operation consists of deleting the edges of a triangle and adding a vertex that is adjacent to the three vertices of the triangle (from right to left).

by suppressing degree-2 vertices, Y–to–Δ or Δ–to–Y operations, deleting degree-1 vertices, loops, and parallel edges (that may be created by previous operations) in any order. Y–to–Δ and Δ–to–Y are shown in Figure 1.3. For $Y\Delta Y$-reducible graphs Yu [81] showed that there are more than 68 billion minimal excluded minors, whereas Pierce [62] showed that there exist at least 157 minimal excluded minors for apex graphs.

The Graph Minor Theorem is of interest not only for its contributions in structural graph theory but also has algorithmic consequences. For instance, Robertson and Seymour [66] showed that verifying whether a graph contains a fixed graph H as a minor can be done in cubic time. Hence, it follows from the Graph Minor Theorem that there exists a polynomial algorithm checking membership in a given minor-closed class.

1.2 Applications of the Graph Minor Theorem

In Chapter 3 we answer a question that is closely related to finding minimal excluded minors. We study a minor-monotone graph invariant, denoted by $f_p(G)$ for a graph G with $p \in [1, \infty]$, that can take unbounded integer values. This invariant is related to isometric embeddings of edge-weighted graphs into metric spaces. We focus on the cases $p = 2$ and $p = \infty$. Instead of identifying the minimal excluded minors for some fixed k, we aim to find a function g and some minors H_k such that each graph with $f_p(G) \geqslant g(k)$ contains an H_k minor, where $f_p(H_k) > k$.

In Chapter 4 we will study the minimal excluded minors for another minor-monotone graph invariant $\varphi(G)$ that is related to the dominant of the cut polytope of a graph G.

1.2.1 Isometric Embeddings

The first application of the Graph Minor Theorem that we consider is related to isometric embeddings of graphs in metric spaces.

Metric spaces are a well-studied topic of mathematics and their properties are studied from multiple points of view. One of them is the study of how a given metric space embeds into another one. Recall that a *metric space* (X, d) consists of a *set of points* X and a *metric* $d : X \times X \to \mathbb{R}_{\geqslant 0}$. That is, for all $x, y, z \in X$, (i) $d(x, y) = d(y, x)$, (ii) $d(x, y) = 0$ if and only if $x = y$, and (iii) $d(x, y) \leqslant d(x, z) + d(z, y)$. We focus on the metric spaces $\ell_p = (\mathbb{R}^k, d_p)$ with $p = 2$ and $p = \infty$. Recall that $\|x\|_p = (\sum_{i=1}^{k} |x|^p)^{1/p}$ if $p \in [1, \infty)$ and $\|x\|_\infty = \max_{i \in [k]} |x_i|$, where we let $[k] = \{1, \dots, k\}$ for $k \in \mathbb{N}$, and $\mathbb{N}$ denotes the set of non-negative integers. We set $d_p(x, y) = \|x - y\|_p$ for all $p \in [1, \infty]$.

An *embedding* of a metric space (X, d) in another metric space (X', d') is a map $\phi : X \to X'$. We say that an embedding is *isometric* if $d(x, y) = d'(\phi(x), \phi(y))$ for all $x, y \in X$. Observe that isometric embeddings are very restrictive, which is why relaxations of isometric embeddings have been studied.

We consider isometric embeddings of semi-metric spaces for which we do not require all distances to be preserved. A semi-metric space satisfies the same conditions as a metric space except we accept zero distances between two points. Observe that we can encode the distances that we want to preserve using a weighted graph. It is an easy exercise to show that every (semi-)metric space corresponds to a weighted complete graph. However, not every weighted graph can be completed to a weighted complete graph corresponding to a (semi-)metric space. In order to be able to do so, we need that the weight function satisfies some conditions.

We say that a weight function is a *distance function on G* if $d : E(G) \to \mathbb{R}_+$ is such that for each edge uv and every path $P = v_0 v_1 \dots v_r$ with $v_0 = u$ and $v_r = v$, $d(uv) \leqslant d(P) = \sum_{i=1}^{r} d(v_{i-1} v_i)$. If $d : E \to \mathbb{R}$ is a distance function, we say that (G, d) is a *metric graph*.

An *isometric embedding* of a metric graph (G, d) in ℓ_p^k is a map $\phi : V(G) \to \mathbb{R}^k$ such that $d_p(\phi(v), \phi(w)) = d(vw)$ for all edges $vw \in E(G)$. For each

$p \in [1, \infty]$ and graph G, a distance function $d : E(G) \to \mathbb{R}_{\geqslant 0}$ is ℓ_p-*realizable* if it has an isometric embedding in ℓ_p^k for some k. If d is ℓ_p-realizable, we define the invariant $f_p(G, d)$ to be the least integer k such that (G, d) can be isometrically embedded in ℓ_p^k. The ℓ_p-*dimension* of G is defined to be $f_p(G) = \sup_d f_p(G, d)$, where the supremum is over all ℓ_p-realizable distance functions d on G.

It can be shown that the class of graphs G satisfying $f_p(G) \leqslant k$ for some fixed k is closed under taking minors. Hence, we know by the Graph Minor theorem that there exists a finite list of minimal excluded minors for each of these classes. A question of interest is therefore to determine these set of minimal excluded minors for small values of k. Fiorini, Huynh, Joret, and Varvitsiotis [42] showed that there exist two minimal excluded minors for $f_p(G) \leqslant k$ for $k = 2$, and for $p = 1$ or $p = \infty$. These graphs are shown in Figure 3.3 on page 26. The study of case $k = 3$ and $p = \infty$ was part of my Master's Thesis [60] and I provided a partial list of excluded minors. For most of these graphs it is not known whether they are minimal.

Our main result is inspired by the Grid Minor Theorem for treewidth. Robertson and Seymour [65] established the following result in their long series of papers about graph minors. The treewidth of a graph $\mathrm{tw}(G)$ can take integer values and describes in some sense how tree-like a graph is. It is well known that a square $k \times k$-grid has treewidth k. Furthermore, by [65], there exists a function $f : \mathbb{N} \to \mathbb{R}$ such that every graph with $\mathrm{tw}(G) \geqslant f(k)$ has a $k \times k$-grid minor. The original function in [65] is super exponential in k. Chekuri and Chuzhoy [20] improved the function to a polynomial of k and the current best function is due to Chuzhoy and Tan [23].

In this thesis we show a similar result for the invariant $f_p(G)$ with $p = 2$ and $p = \infty$. We identify graphs H_k such that for every $k \in \mathbb{N}$, $f_p(H_k) > k$, and show the existence of a function $g : \mathbb{N} \to \mathbb{R}$ such that and every graph G with $f_p(G) \geqslant g(k)$ has an H_k minor. In this sense, large H_k minors are unavoidable in graphs with large $f_p(G)$ value. For $p = 2$ there is one graph H_k for every $k \in \mathbb{N}$ as for treewidth, but for $p = \infty$ there are four graphs H_k for each k such that a graph with $f_p(G) \geqslant g(k)$ contains at least one of these four minors.

1.2.2 Cut Dominants

Recall that a *cut* in a graph G is a subset of edges whose removal disconnects the graph. Cuts are an important topic in graph theory.

Before we explain the invariant that we study in Chapter 4 we show how cuts and embeddings into ℓ_1-spaces are related. Proposition 4.2.2 in the book of Deza and Laurent [29] asserts that a weighted graph is realizable in ℓ_1 if and only if the weight function is a non-negative combination of cuts. Given a graph G and a cut $\delta(S)$ with $S \subseteq V(G)$ we can embed the cut isometrically into $\mathbb{R}$. The vector $\phi \in \mathbb{R}$ such that $\phi(v) = 1$ if $v \in S$ and $\phi(v) = 0$ if $p \notin S$ is such that $d(\phi(u), \phi(v)) = |\phi(v) - \phi(u)| = 1$ if $uv \in \delta(S)$ and $d(\phi(u), \phi(v)) = 0$ otherwise. Observe that we can also derive the cut $\delta(S)$ from the vector $\phi \in \mathbb{R}$. Similarly, we can embed a non-negative combination of cuts $\sum_{i \in [k]} \lambda_i \delta(S_i)$ into ℓ_1^k by setting $\phi(v)_i = \lambda_i$ if $v \in S_i$ and $\phi(v)_i = 0$ otherwise. Hence, by considering a fixed coordinate of the embedding in ℓ_1^k, we can derive a coefficient λ_i as well as the cut $\delta(S_i)$.

The *cut polytope* of a graph G is the convex hull of the incidence vectors of all cuts in G and is defined in $\mathbb{R}^{E(G)}$. The *cut dominant* is obtained by adding the non-negative orthant $\mathbb{R}^E_+(G)$ to the cut polytope. We let $\varphi(G)$ be the greatest right-hand side coefficient in a minimum integer linear description of the dominant of the cut polytope of G. A minimum integer linear description is such that each row has integer coefficients and the greatest common divisor of each row is 1. It is known that $\varphi(G) \in \{1\} \cup 2\mathbb{N}$ for all graphs by a result of Conforti, Rinaldi, and Wolsey [27]. It has been shown that the class of graphs satisfying $\varphi(G) \leqslant k$ is closed under taking minors for all fixed $k \in \mathbb{N}$. Hence, it is possible to characterize these graphs with minimal excluded minors. It is an easy exercise to show that the set of minimal excluded minors for $\varphi(G) \leqslant 1$ is $\{K_3\}$. The minimal excluded minors for $\varphi(G) \leqslant 2$ were determined by Conforti, Fiorini, and Pashkovich [26]. They are the pyramid and the prism graph, shown in Figure 1.4. We will focus on the case $k = 4$ as it is the smallest value of k for which the set of minimal excluded minors is unknown.

A motivation for studying the cut dominant is its relation with the traveling salesman problem. Indeed, the vertices of the subtour elimination polyhedron correspond exactly to the facets of the cut dominant. Cornuéjols, Fonlupt and Naddef [28] showed that the graphs G without a prism, pyramid, or Θ minor are exactly the graphs for which the graphical salesman polytope coincides with the subtour elimination relaxation. The graph Θ is

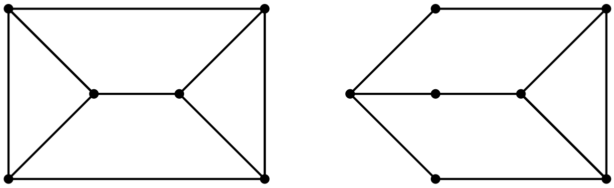

Figure 1.4. The prism on the left and the pyramid on the right are the minimal excluded minors for $\varphi(G) \leqslant 2$.

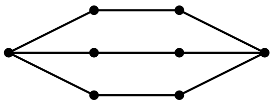

Figure 1.5. The Θ graph.

shown in Figure 1.5. The graphical traveling salesman polytope is a relaxation of the traveling salesman polytope for which we consider any tour in a graph G, instead of only Hamiltonian cycles. Furthermore, Conforti, Fiorini, and Pashkovich [26] showed that the subtour elimination relaxation is integer if and only if the graph does not contain a prism or pyramid minor.

1.3 Ball packings

Besides the Graph Minor Theorem, the other influential result that motivates our findings in Chapter 5 is due to Erdős and Pósa. In their paper from 1965 [39], they showed that for every graph G with at most k vertex-disjoint cycles there exists a set of at most $O(k \log k)$ vertices whose removal yields an acyclic graph. Furthermore, they showed that the bound is asymptotically best possible.

During the years, many mathematicians have generalized the result to other subgraphs. In order to state what we mean by the Erdős-Pósa property we need to have a look at a subject from combinatorial optimization, namely packings and coverings.

Given a ground set V and a collection of subsets $\mathcal{S} = \{S \mid S \subseteq V\}$ we can ask the two following questions.

1. What is the maximum size of a subset $\mathcal{P}$ of $\mathcal{S}$ such that all sets $S \in \mathcal{P}$ are disjoint?

2. What is the minimum size of a subset X of V such that each $S \in \mathcal{S}$ contains at least one element from X?

A subset $\mathcal{P}$ such that all sets $S \in \mathcal{P}$ are disjoint is called a *packing*. The *packing number* $\nu(\mathcal{S})$ equals the maximum size of a packing whose members are in $\mathcal{S}$.

A subset X such that each $S \in \mathcal{S}$ contains at least one element from X is called a *transversal* (aka a *covering*). We say that X covers or hits $\mathcal{S}$. The *transversal number* $\tau(\mathcal{S})$ is the minimum size of a transversal.

Since every transversal X contains at least one element per set $S \in \mathcal{P}$ in a packing $\mathcal{P}$, we see that the transversal number is always at least the packing number.

$$\tau(\mathcal{S}) \geqslant \nu(\mathcal{S})$$

However, in general we cannot bound $\tau(\mathcal{S})$ from above by a function of $\nu(\mathcal{S})$ for all possible set systems $\mathcal{S}$. We say that a family of set systems satisfies the *Erdős-Pósa property* if there exists a function $f : \mathbb{N} \to \mathbb{N}$ such that $\nu(\mathcal{S}) \leqslant f(\tau(\mathcal{S}))$ for all possible set systems.

In the original Erdős-Pósa paper, it is shown that the set system formed by the (vertex sets) of the cycles of any graph G satisfies this property with a function $f \in O(k \log k)$. Robertson and Seymour [65] showed that the set of subgraphs that contain a fixed planar graph H as a minor do satisfy the Erdős-Pósa property with an exponential function f_H depending on the minor H. In the original paper [39], Erdős and Pósa' considered the case $H = K_3$. Cames van Batenburg, Huynh, Joret, and Raymond [15] recently improved the function to $f_H \in O(k \log k)$, which is best possible.

We consider a different setting in Chapter 5. Instead of establishing an Erdős-Pósa property for some fixed minors in any graph, we want to find a relation for any balls in graphs excluding a fixed minor. Hence, we pack a different object in graphs. Furthermore, the function we obtain does not depend on the balls we pack but on the host graph which we consider.

A ball centered at a vertex v and with radius r is the set of all vertices that are at distance at most r from v, where the distance $d_G(u, v)$ is the length of a shortest path from u to v in G.

$$B_r(v) := \{u \in V(G) \mid d_G(u, v) \leqslant r\}$$

Figure 1.6 shows a packing and a transversal in a graph in which we consider all balls of radius 2.

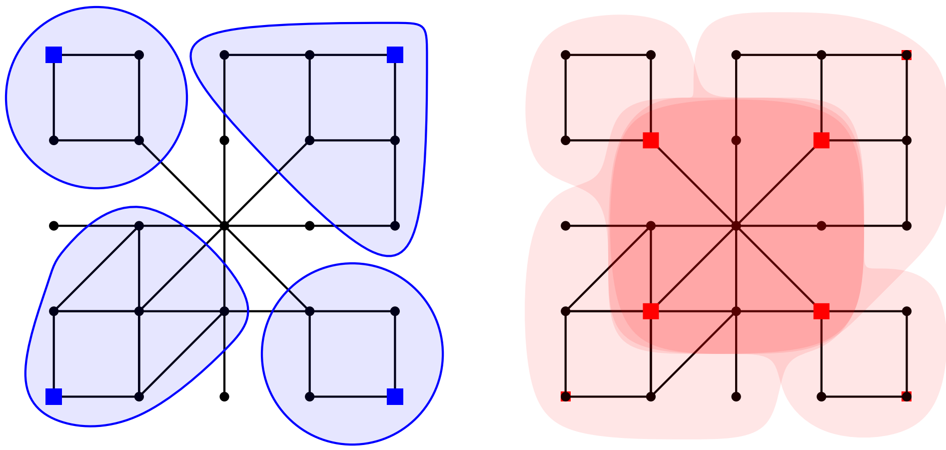

Figure 1.6. At the left there is a packing with 2-balls in blue, whose centers are blue squares. At the right is a transversal of size four for all 2-balls. Each red square hits the 2-balls whose center lies in the red shade.

We show that balls, even with different radii, satisfy the Erdős-Pósa property in K_t-minor-free graphs for fixed t.

1.4 Contributions of the thesis

1.4.1 Isometric embeddings

In Chapter 3 we take the approach of seeking *unavoidable minors* for the invariants $f_2(G)$ and $f_\infty(G)$. That is, for each $k \in \mathbb{N}$, we look for a finite collection $\mathcal{U}_p^k$ of graphs H_k and an integer $c_p(k)$, such that every graph $H_k \in \mathcal{U}_p^k$ of the collection satisfies $f_p(H_k) > k$, and every graph G with $f_p(G) > c_p(k)$ has a minor in $\mathcal{U}_p^k$ for $p \in \{2, \infty\}$.

If $p = 2$, we show that triangular grids are unavoidable minors for $f_2(G)$. The triangular grid $\triangle_7$ is shown in Figure 1.7.

Theorem 1.2. *There exists a function $g_{1.2}(k) = O(k^9 \operatorname{polylog}(k))$ such that every graph G with $f_2(G) > g_{1.2}(k)$ contains a $\triangle_{k+2}$ minor. Moreover, every graph G that contains a $\triangle_{k+2}$ minor has $f_2(G) > k$.*

It turns out that the case $p = \infty$ is much more challenging. Indeed, in

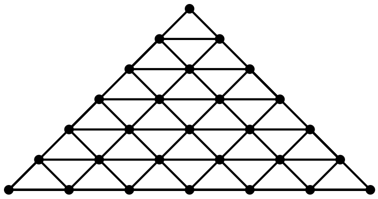

Figure 1.7. The grid $\triangle_7$ satisfies $f_2(\triangle_7) > 5$.

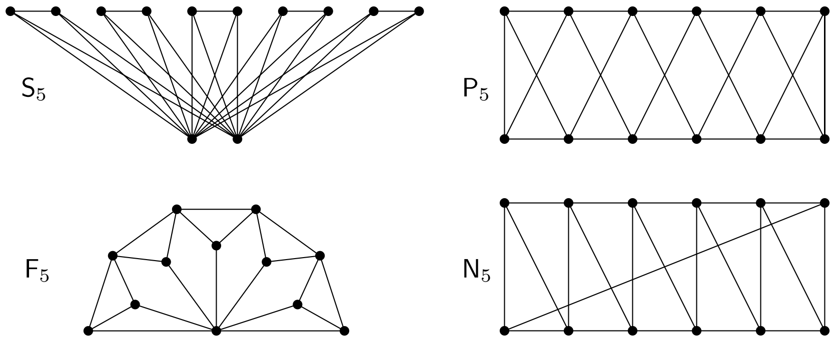

Figure 1.8. The graphs S_5, P_5, F_5 and N_5.

that case, the set $\mathcal{U}_\infty^k$ consists of four graphs whose construction will be detailed in Chapter 3. Figure 1.8 shows the graphs in $\mathcal{U}_\infty^5$. The graphs S_k, F_k, P_k can be obtained by gluing k copies of K_4 along a same edge, edges having exactly one vertex in common, and edges with no vertex in common, respectively. The graphs N_k are obtained from a ladder by adding a diagonal edge and contracting some edges. The main contribution to Chapter 3 is the following theorem for $p = \infty$.

Theorem 1.3. *There exists a computable function $g_{1.3} : \mathbb{N} \to \mathbb{R}$ such that for every $k \in \mathbb{N}$, every graph G with $f_\infty(G) > g_{1.3}(k)$ contains a $\mathcal{U}_\infty^k$ minor. Moreover, every graph G that contains a $\mathcal{U}_\infty^k$ minor has $f_\infty(G) > k$.*

Furthermore, we include a partial list of excluded minors for $f_\infty(G) \leqslant 3$ in Appendix A. However, we are not able to prove minimality for most of these graphs, nor can we prove that the list of excluded minors is complete.

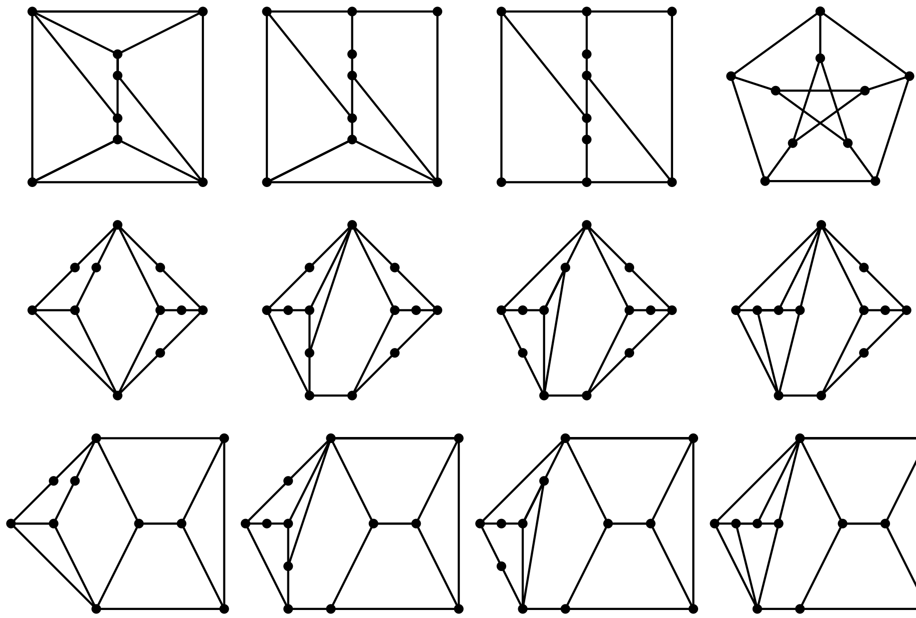

Figure 1.9. The 12 known minimal excluded minors for $\varphi(G) \leqslant 4$.

1.4.2 Cut dominants

Our contributions in Chapter 4 are twofold. First, we show a new bound on $\varphi(G)$ as a function of the vertex cover number $\tau(G)$, which is the minimum size of a set of vertices $X \subseteq V(G)$ such that every edge of G is incident with some vertex in X. We remark that all logarithms in this thesis are natural logarithms.

Theorem 1.4. *There exists a constant c such that, letting $g : \mathbb{N} \to \mathbb{R}$ denote the function $g(x) = 2^{cx \log x}$ we have $\varphi(G) \leqslant g(\tau(G))$ for all graphs G.*

Second, we establish several results regarding minimal excluded minors for $\varphi(G) \leqslant 4$. We present 12 graphs that satisfy $\varphi(G) > 4$ and are minor-minimal with this property, see Figure 1.9. Three of them were already known by Conforti [25], whereas the other nine are new. Furthermore, we present some insights suggesting that we know already all minor-minimal graphs with $\varphi(G) > 4$ that are not internally 3-connected. However, we do not have a complete proof for this. A graph is *internally 3-connected* if every 2-cutset separates exactly one vertex from the rest of the graph.

Conjecture 1.5. *The graphs in Figure 1.9 form the complete list of minimal excluded minors for* $\varphi(G) \leqslant 4$.

1.4.3 Ball packings

Recall that a graph is K_t-minor-free if it does not contain a K_t minor. In Chapter 5 we consider packings and transversals of balls in K_t-minor-free graphs.

Our main theorem of Chapter 5 states that the graphs excluding a K_t-minor satisfy the Erdős-Pósa property for balls with a linear bounding function, for every fixed $t \geqslant 1$.

Theorem 1.6. *For every integer* $t \geqslant 1$*, there is a constant* c_t *such that for every* K_t*-minor-free graph* G *and any collection of balls in* G *with packing number at most* k *the transversal number is at most* $c_t \cdot k$.

Chapter 2

Basic notions about graphs and polyhedra

In this chapter we briefly recall the basic definitions about graphs and polyhedra that we use in the thesis. Alternatively, a reader who has not been introduced to basic notions from graph theory of polyhedral theory can have a look at the references [29, 31, 83]. Proficient readers may move on to Chapter 3 immediately. We note that logarithms in this thesis are natural, and the base of the natural logarithm is denoted by e. Furthermore, we set $[n] = \{1, 2, \ldots, n\}$ for $n \in \mathbb{N}$.

2.1 Graphs

A *simple graph* $G = (V, E)$ consists of a pair of sets $V = V(G)$ and $E = E(G)$. The elements of $V(G)$ are called *vertices* and the elements of $E(G)$ are called *edges* and are unordered pairs of elements of $V(G)$. Besides simple graphs, we consider also directed graphs and multigraphs, which we define next. A graph is *directed* if the edges are ordered pairs of vertices. We say that two edges are *parallel* if they have the same ends, and a *loop* is an edge whose ends coincide. These can occur in both the undirected and directed case. A graph is a *multigraph* if it contains loops or multiple edges. Notice that multigraphs can also be directed. In this thesis we will consider simple, finite graphs unless stated otherwise. A simple graph on n vertices with all possible edges is called a *complete graph* and denoted by K_n.

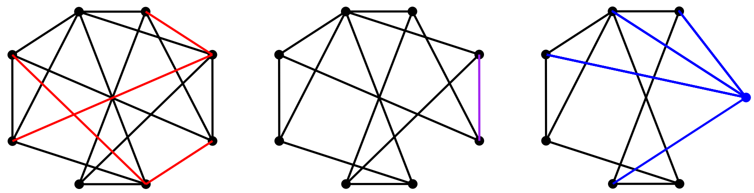

Figure 2.1. The middle graph is a subgraph of the left graph, where the red edges have been deleted. The right graph has been obtained from the middle graph by contracting the purple edge (on the right). The blue edges in the right graph are those obtained from the contraction. The right graph is a minor of both the middle and the left graph.

Hypergraphs are a generalization of graphs where edges are allowed to contain any number of vertices. Formally, a *hypergraph* $\mathcal{H} = (V, \mathcal{E})$ is a pair of vertices $V = V(\mathcal{H})$ and *hyperedges* $\mathcal{E} = \mathcal{E}(\mathcal{H})$, where hyperedges are any subsets of vertices. The *rank* of a hypergraph is the maximum number of vertices in an hyperedge.

We mostly use lower-case letters such as u, v, or $v_1, v_2, \dots, v_r$ to denote vertices and denote an edge e in a simple graph with vertices u and v by uv instead of $\{u, v\}$. If an edge e contains a vertex v, we say that the vertex v is *incident* to the edge e. Two vertices u and v that form an edge uv are *adjacent*. Similarly, two edges sharing a vertex are *adjacent*. The *degree* $d(v)$ of a vertex is the number of edges that are incident to v, or equivalently the number of its neighbors. The *average degree* of a graph is $\mathrm{ad}(G) = \frac{2|E(G)|}{|V(G)|}$.

Given a graph G there are several operations we can define on G. Besides adding vertices or edges, we can also remove them. *Deleting an edge* consists of removing the edge e from $E(G)$. *Deleting a vertex* v consists of removing all edges incident to v and then the vertex v from $V(G)$. *Contracting an edge* uv consists of creating a vertex w, adding all the edges zw such that $uz \in E(G)$ or $vz \in E(G)$ and then deleting the vertices u and v. A graph H that can be obtained from a graph G by vertex deletions, edge deletions, and edge contractions in any order is called a *minor* of G. If the graph H can be obtained from G by vertex and edge deletions only, we say that the graph H is a *subgraph* of G. Figure 2.1 shows an example for a subgraph and edge contraction.

Given two graphs G and H, we want tools to compare these two graphs. We say that G and H are *isomorphic* if there exists a bijection $\phi : V(G) \to V(H)$

such that uv is an edge of G if and only if $\phi(u)\phi(v)$ is an edge of H.

A *path* $P = v_0v_1 \cdots v_r$ of length r in a graph G is a sequence of $r+1$ distinct vertices such that $v_{i-1}v_i$ is an edge of G for all $i \in [r]$. The vertices v_0 and v_r are the *ends* of the path P. If the first and last vertices are adjacent we say that $v_1 \cdots v_r$ is a *cycle*. Note that we can also consider a cycle or path as being a set of edges.

Using paths, we can introduce connectivity in a graph. A graph G is *connected* if for any two vertices u and v there exists a path from u to v. We say that such a path is an *u–v path*. If $u \in A$ and $v \in B$ for some sets of vertices A and B and if no internal vertex of the path is in $A \cup B$, we also talk of an *A-B path*, u–B path, or A–v path. The *distance between two vertices u and v* is defined to be the length of a shortest u–v path if there exists some path, and infinity otherwise.

A *cutset* $X \subseteq V$ is a set of vertices such that $G - X$ is disconnected, where $G - X$ is the graph obtained by deleting all vertices in X from G. A *cut* Y is a set of edges such that $G \setminus Y$ is disconnected, where $G \setminus Y$ is the graph obtained by deleting all edges in Y from G. A graph is *k-connected* if $|V(G)| \geqslant k+1$ and there exists no cutset of size strictly less than k.

The maximal connected subgraphs of a graph G form the *connected components* $G_1, \ldots, G_r$ of G. Similarly, we can define blocks in a graph. A *block* is a maximal connected subgraph of G without a cutvertex, which is a cutset which is a single vertex.

The next notions are related to cutsets in a graph. A *k-separation* of a graph G is an ordered pair (G_1, G_2) of edge-disjoint subgraphs of G with $G = G_1 \cup G_2$, $|V(G_1) \cap V(G_2)| = k$, and $E(G_1)$, $E(G_2)$, $V(G_2) \setminus V(G_1)$, $V(G_1) \setminus V(G_2)$ all non-empty. observe that the vertices in $V(G_1) \cap V(G_2)$ form a k-cutset, which is a cutset of size k. A *k-sum* is a graph G obtained by gluing two graphs G_1 and G_2 along a common clique K of size k and then possibly deleting some edges of K.

A notion that we refer to in Chapter 3 for the Eulidean case is *treewidth*. This graph invariant describes how treelike a graph is. The invariant treewidth is defined such that trees and forests have treewidth 1. Robertson and Seymour [64] defined tree decompositions and treewidth as follows.

Let G be a graph, T a tree, and let $\mathcal{V} = (X_t)_{t \in T}$ be a family of vertex sets $X_t \subseteq V(G)$ indexed by the nodes t in T. The pair $(T, \mathcal{V})$ is called a *tree-decomposition of G* if it satisfies the following three conditions:

1. $V(G) = \cup_{t \in T} X_t$;
2. for every edge $e \in G$ there exists a $t \in T$ such that both ends of e lie in X_t;
3. for $t, t', t'' \in V(T)$, if t' is on the path of T between t and t'' then $X_t \cap X_{t''} \subseteq X_{t'}$.

The *width of the tree-decomposition* $(T, \mathcal{V})$ is

$$\max_{t \in T} |X_t| - 1 \ ,$$

and the *treewidth* $\mathrm{tw}(G)$ of G is the minimum width of any tree-decomposition of G.

It is well-known that a rectangular $r \times r$-grid $\square_r$ satisfies $\mathrm{tw}(\square_r) = r$. The grid $\square_r$ is the graph defined such that $V(\square_r) = \{v_{i,j} \mid i, j \in [r]\}$ and two vertices $v_{i,j}$ and $v_{k,\ell}$ are linked by an edge if and only if $|i - k| + |j - \ell| = 1$. Similarly, the *triangular grid* $\triangle_r$ has vertex set $V(\triangle_r) = \{v_{i,j} \mid i, j \in [r],\ i \leqslant j\}$ and edge set $E(\triangle_r) = \{v_{i,j}v_{k,\ell} \mid v_{i,j}, v_{k,\ell} \in V(\triangle_r),\ (i-k, j-\ell) \in \{\pm(1,0), \pm(0,1), \pm(1,1)\}\}$. In order to illustrate the concept of treewidth, we show the following upper bound on the treewidth of $\triangle_r$. Lemma 2.1 will be used later in Chapter 3 for obtaining unavoidable minors for large Euclidean dimension.

Lemma 2.1. $\mathrm{tw}(\triangle_r) \leqslant r - 1$ *for all* $r \geqslant 3$.

Proof. In order to show the bound $\mathrm{tw}(\triangle_r) \leqslant r - 1$ it is sufficient to find a tree decomposition $(T, \mathcal{V})$ of width at most $r - 1$. That is, all bags $X \in \mathcal{V}$ contain at most r vertices.

First, notice that the graph $\triangle_r \setminus \{v_{1,1}, v_{r,r}\}$ has a tree-decompostion where T is a path and each bag contains at most $r - 1$ vertices. Indeed, let us say the first bag of the path contains vertices $\{v_{1,i}, v_{2,2}, v_{2,3} \mid 1 < i < r\}$. Then for $i = 2, \ldots r - 3$, we replace $v_{1,i}$ by $v_{2,i+1}$ in the following bag. Notice that the last bag now contains $v_{2,1}, \ldots, v_{2,r-2}$. Now, for each row $j = 2, \ldots, r-1$, we may add a bag containing exactly $\{v_{j,1}, \ldots, v_{j,r-j}, v_{j+1,1}\}$. Now, for each each $i = 1, \ldots, j - 1$ we replace $v_{j,i}$ with $v_{j+1,i+1}$ in the following bag start all over for the next j.

In order to obtain a tree decomposition for $\triangle_r$ it is sufficient to add the bags $\{v_{1,1}, v_{1,2}, v_{2,2}\}$ and $\{v_{r,r}, v_{r-1,r-1}, v_{r-1,r}\}$ and make them adjacent to some bag containing $\{v_{1,2}, v_{2,2}\}$ and $\{v_{r-1,r-1}, v_{r-1,r}\}$, respectively. $\square$

2.2 Polyhedra

Finally, we introduce some notions from polyhedral theory and linear programming. We will need these notions in Chapters 4 and 5.

Polytopes and polyhedra are finitely generated convex sets that can be defined in two different ways. We can define polytopes (respectively polyhedra) either as convex hulls of points (respectively, convex hulls plus a convex cone), or as intersections of closed half-spaces. It is well-known that both definitions are equivalent, see [83, Theorem 1.1]. The difference between a polytope and a polyhedron is that we ask that a polytope is bounded. Notice that every polytope is also a polyhedron. For the rest of the section we will work in the vector space $\mathbb{R}^n$ unless stated otherwise.

In the *vertex description* a *polytope* P is defined as the convex hull of a finite set of points. We write

$$P = \operatorname{conv}(\{x_1, \ldots, x_k\}) = \left\{ \sum_{i=1}^{k} \lambda_i x_i \mid \sum_i \lambda_i = 1, \lambda_i \geqslant 0 \quad \forall i \in [k] \right\}.$$

A *polyhedron* Q is the Minkowski sum of a convex hull and a conical hull, both finitely generated,

$$\begin{aligned} Q &= \operatorname{conv}(\{x_1, \ldots, x_k\}) + \operatorname{cone}(\{y_1, \ldots, y_r\}) \\ &= \left\{ \sum_{i=1}^{k} \lambda_i x_i + \sum_{j=1}^{r} \mu_j y_j \mid \sum_i \lambda_i = 1, \lambda_i \geqslant 0 \quad \forall i \in [k], \mu_j \geqslant 0 \quad \forall j \in [r] \right\} \end{aligned}$$

If the description is non-redundant (that is, no point x_i is a convex combination of the points x_h with $h \neq i$ plus a conical combination of y_j), then we say that $x_1, \ldots, x_k$ are the *vertices* of the polytope P, respectively the polyhedron Q. The vectors y_i that are non-zero are called *rays* of the polyhedron Q. Notice that Q is a polytope if and only if it has no rays.

A polyhedron P (or polytope) defined by a *linear description* is the intersection of m closed half-spaces, where $m \in \mathbb{N}$. As every closed half-space can be described by an inequality of the form $a_i^T x \leqslant b_i$ (with $a_i \in \mathbb{R}^n$, $b_i \in \mathbb{R}$), we can collect the m inequalities as a system $Ax \leqslant b$ (with $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$) such that a point is contained in the polyhedron if and only if it satisfies all given inequalities. An inequality $a^T x \leqslant b$ is *valid for* P if all points $x \in P$ satisfy the inequality. We say that the linear description of P

is *non-redundant* if deleting any inequality from the system gives rise to a strictly larger polyhedron.

In Chapter 4, we will focus on 0/1 *polyhedra*, which are polyhedra whose vertices and rays have 0/1 coefficients, and on linear descriptions of these polyhedra in *minimum integer form*. That is, every inequality has integer coefficients and the coefficients of a given inequality have greatest common divisor equal to 1, which is possible because 0/1 polyhedra are in particular rational.

The *dimension* $\dim(P)$ of P is the dimension of its affine hull, which contains all linear combinations of points of P. Let $a^T b \leqslant b$ be a valid inequality for P. The set $F = \{x \in P \mid a^T x = b\}$ is a *face* of P. Notice that every face of P is a polyhedron, contained in P. If $F = \{v\}$ is a 0-dimensional face, we say that v is a *vertex* of P. If the face F has dimension $\dim(F) = \dim(P) - 1$, it is called a *facet*. If the face F has dimension $\dim(F) = \dim(P) - 2$, it is called a *ridge*.

In Chapter 5, we will use the fractional packing number and the fractional transversal number, which are defined via dual linear programs. A *linear program* is the task of maximizing of minimizing a linear function under linear equality or inequality constraints. Given a *primal* linear program

$$\begin{aligned} \max \quad & c^T x \\ \text{s.t.} \quad & Ax \leqslant b \\ & x \geqslant 0, \end{aligned}$$

where $A \in \mathbb{R}^{m \times n}$ is a matrix, $b \in \mathbb{R}^m$, $c \in \mathbb{R}^n$ are column vectors, and $x \in \mathbb{R}^n$ is the variable, its *dual* linear program is defined to be

$$\begin{aligned} \min \quad & b^T y \\ \text{s.t.} \quad & A^T y \geqslant c \\ & y \geqslant 0, \end{aligned}$$

where $y \in \mathbb{R}^m$ is the variable. By the weak duality theorem, we have $\max\{c^T x \mid Ax \leqslant b, x \geqslant 0\} \leqslant \min\{b^T y \mid A^T y \geqslant c, y \geqslant 0\}$ for all primal-dual pairs. Furthermore, if both problems are realizable, and the maximum and the minimum are finite, then we have $\max\{c^T x \mid Ax \leqslant b, x \geqslant 0\} = \min\{b^T y \mid A^T y \geqslant c, y \geqslant 0\}$ by the strong duality theorem, see [24].

Chapter 3

Isometric embeddings

This chapter is based on joint work with Samuel Fiorini, Tony Huynh, and Gwenaël Joret, see the paper *Unavoidable minors for graphs with large* ℓ_p*-dimension* which has been published in Discrete and Computational Geometry [41].

Spaces are omnipresent in mathematics. Mathematics students are confronted to vector spaces, topological spaces, differential geometry, combinatorics and many more areas of mathematics where they study the different behaviors of their favorite space. There are many ways to study their behavior. Topology focuses on the shape and local behavior without a notion of length or units. In a topologist's mind a donut behaves the same way as a coffee cup as both can be obtained from the sphere by gluing a handle and then deforming the object to obtain the desired shape. In contrast to topology stands metric geometry, where one introduces a unit notion and uses it to measure some properties of the space differently. When a space X is equipped with a metric d we talk of a *metric space*. Using metric spaces we can compare two mathematical objects with the newly acquired tools. For instance, two triangles that have the same side lengths are said to be isometric, whereas if the lengths match only up to some scaling factor they are similar.

An important area in metric theory is the study of how metric spaces compare to each other. Whenever two metric spaces are defined on a different ground set or with distinct metrics, we can wonder whether these lead to the same metric space or whether they behave similarly. One way to do so is by *isometric embeddings*. Given two metric spaces (X, d) and (X', d'),

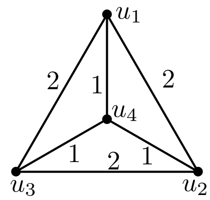

Figure 3.1. The graph K_4 with a distance function that cannot be embedded in ℓ_2. Indeed, in such an embedding, the vertices u_1, u_2, u_3 would have to form an equilateral triangle of side length 2. Furthermore, $d(u_i, u_4) = 1$ for every $i \in [3]$ and $d(u_1, u_2) = d(u_2, u_3) = d(u_3, u_1) = 2$ imply that u_4 is the midpoint of the three segments $[u_1, u_2]$, $[u_2, u_3]$, and $[u_3, u_1]$, a contradiction.

an *isometric embedding* is a map $\phi : X \to X'$ such that distances between any two points are preserved. That is $d(x, y) = d'(\phi(x), \phi(y))$ for every $x, y \in X$. Notice that only few pairs of metric spaces admit an isometric embedding from one into the other as a necessary condition is that there exists a bijection from X to X'.

Recall that the d_p metric on $\mathbb{R}^n$ is defined as $d_p(x, y) = (\sum_{i\in[n]} |x_i - y_i|^p)^{1/p}$ for $p \in [1, \infty[$ and $d_\infty(x, y) = \max_{i\in[n]} |x_i - y_i|$. An ℓ_p-space is a space $(\mathbb{R}^n, d_p)$ for some n. If we want to emphasize that $n = k$ we write ℓ_p^k-space. It is a classic result that two spaces ℓ_p^k and $\ell_{p'}^{k'}$ are isometric, that is there exists an isometric embedding from one to the other, if and only if $p = p'$ and $k = k'$.

We are mostly interested in ℓ_p-spaces with $p \in \{2, \infty\}$ in this chapter. We say that a metric space (X, d) is ℓ_p*-realizable* if there exists some dimension k such that (X, d) can be isometrically embedded in ℓ_p^k. Observe that some metric spaces are not ℓ_p-realizable. For instance, Figure 3.1 illustrates a 4-point metric space that cannot be embedded in Euclidean space. However, every metric space can be isometrically embedded into an ℓ_∞-space. We let $f_p(n)$ be the smallest integer k such that every ℓ_p-realizable n-point metric space can be embedded in ℓ_p^k.

Ball [3] studied isometric embeddings of n-point metric spaces into ℓ_p^k-spaces. His main result is that every ℓ_p-realizable n-point metric space (X, d) can be embedded in an ℓ_p-space of dimension at most $\binom{n}{2}$ for all $p \in [1, \infty]$, that is, $f_p(n) \leqslant \binom{n}{2}$ for all $p \in [1, \infty]$. Moreover, he showed $f_\infty(n) \geqslant n - cn^{3/4}$ where $c \in \mathbb{R}$ is a constant. This improved a result by Witsenhausen [79], who showed that $f_\infty(n) \geqslant \frac{2}{3}n$. Rödl and Ruciński [68] later showed that there exists a constant c such that $f_\infty(n) \geqslant n - c\log_2 n$ for every $n \in \mathbb{N}$. On the other hand, Holsztyński [51] showed that every n-point metric space can be isometrically embedded in ℓ_∞^{n-2}, that is $f_\infty(n) \leqslant n-2$. We will show this bound in Lemma 3.15. Notice that this bound is better than the bound for general ℓ_p-spaces by Ball, $f_p(n) \leqslant \binom{n}{2}$ for all $p \in [1,\infty]$. Barvinok [5] was interested in Euclidean dimension and showed $f_2(n) \leqslant \frac{\sqrt{8\binom{n}{2}+1}-1}{2}$.

One way of generalizing isometric embeddings is by allowing the distances to vary a little bit. One such approach was taken by Bourgain [10] using *distortion*. The distortion of a map $f : X \to X'$ is the smallest value $\alpha \geqslant 1$ for which there exists an $r > 0$ such that for all $x, y \in X$, $r \cdot d(x,y) \leqslant d'(f(x), f(y)) \leqslant \alpha r \cdot d(x,y)$. He showed that $O(\log n)$-distortion can be achieved when embedding n-point metric spaces in an ℓ_p-space for fixed p, and that it is best possible. More precisely, every n-point metric space can be embedded into an $\ell_p^{O(\log^2 n)}$-space with $O(\log n)$-distortion and there is some n-point metric space that cannot be embedded with smaller distortion. Since Bourgain's breakthrough in 1985, there have been several improvements of this result. In particular, the dimension of the space in which we can embed with distortion $O(\log n)$ was improved from $O(\log^2 n)$ to $O(\log n)$ by Abraham, Bartal, and Neiman [1], which is best possible.

Another way to generalize isometric embeddings is to require only a subset of distances to be preserved, which is the perspective we take. Given an n-point metric space, we ask only for a subset of the distances to be isometrically embedded, whereas we do not give any condition on the other distances. This implies that the map we consider may not be injective because we can map two points without a prescribed distance to the same point in the target space.

We also want to include the possibility that the distance between two points may be zero. Hence, we will work with semi-metric spaces. Recall that in a metric space (X, d), the metric $d : X \to \mathbb{R}_+$ satisfies the following conditions for all $x, y, z \in X$.

(i) $d(x,y) = d(y,x)$

(ii) $d(x,y) = 0$ if and only if $x = y$

(iii) $d(x,y) \leqslant d(x,z) + d(z,y)$.

We say that a space (X, d) is a semi-metric space if d satisfies the above conditions except maybe (ii). Observe that every metric space is also a semi-metric space and that a given semi-metric space can be transformed into a metric space by identifying sets of vertices that are mutually at distance zero from one another.

We can encode an embedding, where we fix only a subset of distances that are isometrically embedded, by an edge-weighted graph on n vertices, where an edge is included if and only if the corresponding distance has to be preserved, and the weight of this edge is the prescribed distance. Notice that not every weighted graph can be obtained in that way. The following definition characterizes the weighted graphs that can be derived from semi-metric spaces.

Let $G = (V, E)$ be a graph. We say that a weight function on the edges of G, $d : E(G) \to \mathbb{R}_+$, is a *distance function on* G if $d(e) \geqslant 0$ for each edge $uv \in E(G)$, and if for every path $v_0 v_1 \dots v_r$ with $v_0 = u$ and $v_r = v$, we have $d(uv) \leqslant \sum_{i=1}^{r} d(v_{i-1}v_i)$. If d is a distance function, we say that (G, d) is a *metric graph*.

Observe that, given a metric graph (G, d) we can find a semi-metric space that leads to this metric graph. Indeed, for all edges uv we assume that the distance in the semi-metric space is $d(uv)$. If two vertices u and v do not form an edge, we can set the distance of the corresponding elements in the metric space to be the length of a shortest path from u to v.

Given a metric graph (G, d) and $p \in [1, \infty]$, we say that d is *ℓ_p-realizable* if (G, d) has an isometric embedding in ℓ_p^k for some k. Notice that this embedding may not be injective. If d is ℓ_p-realizable, we define the invariant $f_p(G, d)$ to be the least integer k such that (G, d) can be isometrically embedded into ℓ_p^k. The *ℓ_p-dimension* of G is defined to be $f_p(G) = \sup_d f_p(G, d)$, where the supremum is over all ℓ_p-realizable distance functions d on G. Notice that we have $f_p(K_n) = f_p(n)$ because every n-point semi-metric space can be encoded as an n-vertex metric graph. We remark that in the special case $p = \infty$, the supremum is taken over all distance functions on G, as every n-point metric space can be isometrically embedded into ℓ_∞^{n-2}, which we will prove in Lemma 3.15.

An important observation in our study is that the class of graphs G satisfying $f_p(G) \leqslant k$ is closed under taking minors.

Lemma 3.1. *Let G be a graph and let H be a minor of G. Then $f_p(H) \leqslant f_p(G)$ for all $p \in [1, \infty]$.*

Proof. Fix a distance function d_H on H. Assume first that H has been obtained from G by contracting one edge uv. We define $d_G : E(G) \to \mathbb{R}$ such that $d_G(uv) = 0$ and $d_G(e) = d_H(e)$ for every other edge $e \in E(H)$. Notice that d_G is a distance function on G and that in any embedding u and v have the same image. Thus, an embedding of (G, d_G) is also an embedding of (H, d_H). Assume now that H has been obtained from G by deleting the edge uv. In this case, we set $d_G(uv)$ to be the length of a shortest u–v path. This way we ensure that d_G is a distance function on G. Again, any embedding of (G, d_G) is also an embedding of (H, d_H). Therefore, it follows that $f_p(H) \leqslant f_p(G)$ for every $p \in [1, \infty]$. □

Combining Lemma 3.1 with the Graph Minor Theorem, Theorem 1.1 we get the following theorem.

Theorem 3.2. *For each p and k, the property $f_p(G) \leqslant k$ has a* finite *set of minimal excluded minors.*

Notice that $\ell_p^1 = \ell_q^1$ for all $p, q \in [1, \infty]$ because all metrics d_p are the same in one dimension. Furthermore, all forests are realizable in one dimension. Indeed, we can embed each tree independently of one another. Assume that ϕ is an isometric embedding of a fixed tree in ℓ_p^1. We can fix a root and assume that it is embedded at 0 and then, following a BFS order, embed every child c at the prescribed distance from its parent p, that is $\phi(c) - \phi(p) = d(c, p)$.

Furthermore, we can show that K_3 cannot be realized in one dimension. Assuming that the vertices of K_3 are u, v, w, we get a contradiction when the edge weights are $d(uv) = d(uw) = d(vw) = 1$. Indeed, on the line $d(uv) = d(uw) = 1$ implies that either v and w coincide or are at distance 2, contradicting $d(vw) = 1$.

Every graph that is not a forest contains a cycle, and thus a K_3-minor. This implies that for all $p \in [1, \infty]$, K_3 is the only minimal excluded minor for $f_p(G) \leqslant 1$. It is a natural question to look for the minimal excluded minors of $f_p(G) \leqslant k$ when k is small.

We are going to focus on the cases $p = 2$ and $p = \infty$. The complete sets of minimal excluded minors are known in the Euclidean case $p = 2$ for dimensions $k = 1, 2, 3$. Belk and Connelly [6, 7] showed that in these cases $\{K_3\}$, $\{K_4\}$, $\{K_5, K_{2,2,2}\}$ are the respective sets of minimal excluded minors. In the special case of Euclidean spaces, we have some convenient

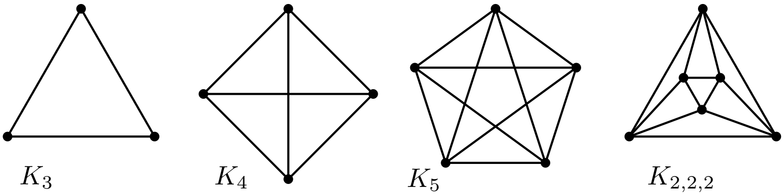

Figure 3.2. The graphs K_3, K_4, K_5 and $K_{2,2,2}$.

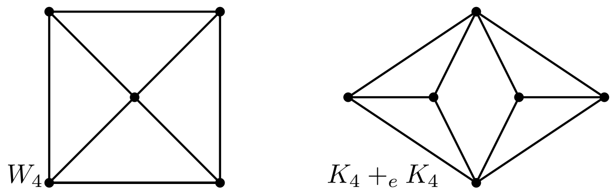

Figure 3.3. The minimal excluded minors for $f_\infty(G) \leqslant 2$.

properties such as the fact that rotating points does not affect the pairwise distances of these points. For other values of p this is no longer true, which makes these values of p more challenging.

Fiorini, Huynh, Joret, and Varvitsiotis [42] determined that W_4, the wheel on 5 vertices, and the graph $K_4 +_e K_4$ (see Figure 3.3) are the only minimal excluded minors for $f_\infty(G) \leqslant 2$. They showed also that the case $f_1(G) \leqslant 2$ has the same minimal excluded minors as $f_\infty(G) \leqslant 2$. As far as we know, the complete set of minimal excluded minors for $f_p(G) \leqslant k$ is unknown for all other values of p and k.

It is plausible that determining any further set of minimal excluded minors will require significant effort, especially in dimension 3 or higher (see [60]). In Section 3.8 we briefly discuss graphs that are not realizable in ℓ_∞^3.

Instead of obtaining the minimal excluded minors of the property $f_p(G) \leqslant k$ for some constants p and k, we take a different approach and seek collections of *unavoidable minors*. That is, we want to identify a family of graphs for which there exists a function $g : \mathbb{N} \to \mathbb{R}$ such that any graph G with $f_p(G) \geqslant g(k)$ contains at least a graph H from this family with $f_p(H) > k$ as a minor.

An example of a theorem involving unavoidable minors is the famous Grid Minor Theorem of Robertson and Seymour [65]. This theorem states that square grids are unavoidable for large treewidth. Precisely, the treewidth of the $k \times k$-grid $\square_k$ is $\mathrm{tw}(\square_k) = k$ and there exists a function $g : \mathbb{N} \to \mathbb{R}$ such that every graph G with treewidth at least $\mathrm{tw}(G) \geqslant g(k)$ has a $\square_{k+1}$ minor H that satisfies $\mathrm{tw}(\square_{k+1}) > k$. Chekuri and Chuzhoy [20], and later Chuzhoy and Tan [23] have improved the function g to a polynomial, $g = O(k^9 \operatorname{polylog}(k))$.

We show that triangular grids $\triangle_{k+2}$ are unavoidable minors when embedding metric graphs in ℓ_2-spaces. Precisely, we show $f_2(\triangle_{k+2}) > k$ and there exists a function $g : \mathbb{N} \to \mathbb{R}$ such that every graph G with $f_2(G) \geqslant g(k)$ has a $\triangle_{k+2}$ minor, see also Theorem 1.2 in Section 3.1.

Theorem 1.2. *There exists a function $g_{1.2}(k) = O(k^9 \operatorname{polylog}(k))$ such that every graph G with $f_2(G) > g_{1.2}(k)$ contains a $\triangle_{k+2}$ minor. Moreover, every graph G that contains a $\triangle_{k+2}$ minor has $f_2(G) > k$.*

Embedding metric graphs isometrically in ℓ_∞ is much more challenging as we need four graphs H with $f_\infty(H) > k$ in our family that, together, are unavoidable. That is, there exists a function g such that if $f_\infty(G) \geqslant g(k)$ then the graph G contains a minor H that is one of the four graphs in the family with $f_\infty(H) > k$. Most of the chapter is devoted to the case $p = \infty$ and our main result is Theorem 1.3 that identifies unavoidable minors for $p = \infty$.

In order to state the main theorem of this chapter we present the construction of the unavoidable graphs. The graph S_k is obtained by gluing the k copies of K_4 along one common edge. The graph P_k is obtained by picking a perfect matching $\{e_i, f_i\}$ in each copy of K_4, and identifying f_i and e_{i+1} for all $i \in [k-1]$. The graph F_k is constructed similarly, except that we take e_i and f_i to be incident edges. Edges are identified in such a way that the common end of e_i and f_i is identified to the common end of e_{i+1} and f_{i+1} for all $i \in [k-1]$. The notation for these first three families reflect the fact that the corresponding copies of K_4 are arranged as a star, path, and fan, respectively. Notice that $\mathsf{S}_2 = \mathsf{P}_2 = \mathsf{F}_2 = K_4 +_e K_4$, which is one of the minimal excluded minors for $f_\infty(G) \leqslant 2$. Next, we define our final family of graphs. The graph N_k is the graph with $V(\mathsf{N}_k) = \{v_0, \dots, v_k\} \cup \{w_0, \dots, w_k\}$ and

$$E(\mathsf{N}_k) = \{v_{i-1}v_i, v_iw_i, v_{i-1}w_i, w_{i-1}w_i \mid i \in [k]\} \cup \{v_0w_0, w_0v_k\}.$$

For each $k \in \mathbb{N}$, we let $\mathcal{U}_\infty^k = \{\mathsf{S}_k, \mathsf{P}_k, \mathsf{F}_k, \mathsf{N}_k\}$. The graphs of $\mathcal{U}_\infty^5$ are shown in Figure 1.8 on page 11. We say that a graph G *contains a $\mathcal{U}_\infty^k$ minor* if it contains $\mathsf{S}_k, \mathsf{F}_k, \mathsf{P}_k$ or N_k as a minor. Our main theorem shows that if $f_\infty(G)$ is large, then G necessarily contains a $\mathcal{U}_\infty^k$ minor with large k.

Theorem 1.3. There exists a computable function $g_{1.3} : \mathbb{N} \to \mathbb{R}$ such that for every $k \in \mathbb{N}$, every graph G with $f_\infty(G) > g_{1.3}(k)$ contains a $\mathcal{U}_\infty^k$ minor. Moreover, every graph G that contains a $\mathcal{U}_\infty^k$ minor has $f_\infty(G) > k$.

Let $\mathcal{S} = \bigcup_k \{\mathsf{S}_k\}, \mathcal{F} = \bigcup_k \{\mathsf{F}_k\}, \mathcal{P} = \bigcup_k \{\mathsf{P}_k\}$, and $\mathcal{N} = \bigcup_k \{\mathsf{N}_k\}$. For a class of graphs $\mathcal{C}$ and $p \in [1, \infty]$, we let $f_p(\mathcal{C}) = \max\{f_p(G) \mid G \in \mathcal{C}\}$, if this number is finite, and $f_p(\mathcal{C}) = \infty$, otherwise. As an immediate corollary, our main theorem gives an exact characterization of all minor-closed classes $\mathcal{C}$ with $f_\infty(\mathcal{C}) = \infty$.

Corollary 3.3. *For all minor-closed classes of graphs $\mathcal{C}$, $f_\infty(\mathcal{C}) = \infty$ if and only if $\mathcal{S} \subseteq \mathcal{C}$ or $\mathcal{F} \subseteq \mathcal{C}$ or $\mathcal{P} \subseteq \mathcal{C}$ or $\mathcal{N} \subseteq \mathcal{C}$.*

To prove the corollary it is sufficient to observe that if no class $\mathcal{S}, \mathcal{F}, \mathcal{P}, \mathcal{N}$ is included in $\mathcal{C}$, then we can determine the maximum value k for which some graph in $\mathcal{C}$ contains a S_k, F_k, P_k or N_k graph. Now we can apply Theorem 1.3 to bound $f_\infty(G) \leqslant g_{1.3}(k+1)$ for all $G \in \mathcal{C}$.

The chapter is organized as follows. In Section 3.1, we establish that grids are unavoidable minors for large ℓ_2-dimension. In Section 3.2, we give a more combinatorial definition of ℓ_∞-dimension. In Section 3.3, we establish some lemmas on ℓ_∞-dimension to be used later.

We establish the second part of our main result, Theorem 1.3, in Section 3.4, by constructing on each graph $G \in \mathcal{U}_\infty^k$ a distance function d that allows us to show $f_\infty(G, d) > k$ in a simple, combinatorial way.

In order to prove the first part of Theorem 1.3, we consider a graph G without a $\mathcal{U}_\infty^k$ minor and set out to prove that we can upper bound $f_\infty(G)$ by some integer $g_{1.3}(k)$.

It is straightforward to show that the ℓ_∞-dimension of a graph is the maximum ℓ_∞-dimension of one of its blocks (see Lemma 3.12). Therefore, we may assume that G is 2-connected. In Section 3.5, we prove that we can essentially assume that G is 3-connected. This part relies on SPQR trees.

The 3-connected case is the part of the proof requiring most of the work. The proof techniques here are mostly graph-theoretic, and may be of independent interest. This is done in Section 3.6 and Section 3.7.

Finally, in Section 3.8, we conclude with some remarks about the minimal excluded minors for $f_\infty(G) \leqslant 3$.

3.1 Euclidean dimension

The goal of this section is to establish that grids are a collection of unavoidable minors for large Euclidean dimension, which is the analogue of Theorem 1.3 for ℓ_2-dimension.

Let $r \in \mathbb{N}$. Recall that the *square grid* graph $\square_r$ is the graph with vertex set $[r] \times [r]$, where (i, j) is adjacent to (i', j') if and only if $|i - i'| + |j - j'| = 1$. The *triangular grid* graph $\triangle_r$ has vertex set $V(\triangle_r) = \{v_{i,j} \mid i, j \in [r],\ i \leqslant j\}$ and edge set $E(\triangle_r) = \{v_{i,j}v_{k,\ell} \mid v_{i,j}, v_{k,\ell} \in V(\triangle_r),\ (i - k, j - \ell) \in \{\pm(1, 0), \pm(0, 1), \pm(1, 1)\}\}$.

Let G and H be graphs such that H is a minor of G. Then G contains an H*-model*, that is, a collection $\{X_v \mid v \in V(H)\}$ of disjoint subsets $X_v \subseteq V(G)$ each inducing a connected subgraph of G such that for every edge $vw \in E(H)$ there is an edge of G with one end in X_v and the other in X_w. The sets X_v are called the *vertex images*. The following is the main result of this section.

Theorem 1.2. There exists a function $g_{1.2}(k) = O(k^9 \operatorname{polylog}(k))$ such that every graph G with $f_2(G) > g_{1.2}(k)$ contains a $\triangle_{k+2}$ minor. Moreover, every graph G that contains a $\triangle_{k+2}$ minor has $f_2(G) > k$.

In order to prove the first part of Theorem 1.2, we use the by now standard notion of *treewidth* (see [31] for the definition). We let $\operatorname{tw}(G)$ denote the treewidth of a graph G. As observed by Belk and Connelly [7], $f_2(G) \leqslant \operatorname{tw}(G)$ holds for all graphs G. Thus, if $f_2(G) > c$, then $\operatorname{tw}(G) > c$.

By the grid theorem [65], there is a function $\gamma(k)$ such that every graph G with $\operatorname{tw}(G) \geqslant \gamma(k)$ contains $\square_k$ as a minor. In fact, one can take $\gamma(k) = O(k^9 \operatorname{polylog}(k))$ by very recent results [23] (see [20] for the original polynomial grid theorem). Furthermore, it is easy to check that $\square_{2k+2}$ has a $\triangle_{k+2}$ minor, for all $k \in \mathbb{N}$. Figure 3.4 illustrates this for $k = 4$. Therefore, in Theorem 1.2, we may take $g_{1.2}(k) = \gamma(2k + 2)$. This proves the first part of the theorem. Notice that for all $r \in \mathbb{N}$, $\triangle_r$ has $\square_m$ as a subgraph, where $m = \lfloor \frac{r-1}{2} \rfloor$. Thus, excluding triangular grids is equivalent to excluding rectangular grids within a factor of 2.

We now prove the second part of Theorem 1.2, see Lemma 3.4 below. We remark that Eisenberg-Nagy, Laurent and Varvitsiotis [37] prove a similar result for a related invariant called *extreme Gram dimension*. This is a variant of the *Gram dimension* of a graph, that is studied and compared

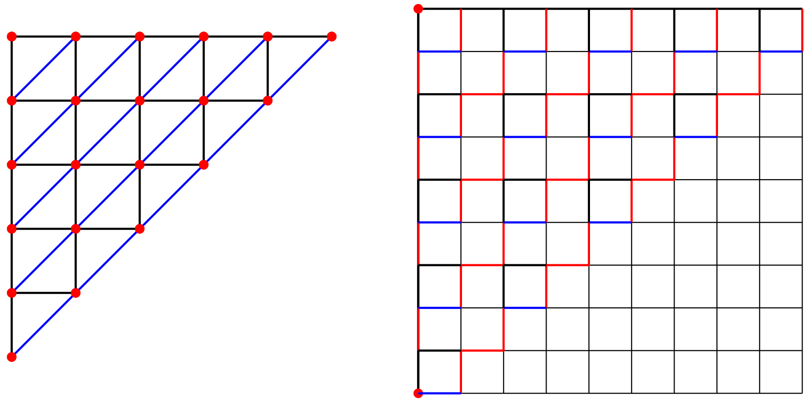

Figure 3.4. On the left is $\triangle_6$. On the right is a $\triangle_6$-model in $\square_{10}$. Vertex images are displayed in red, and edges between the vertex images in black or blue.

to the Euclidean dimension in Laurent and Varvitsiotis [56]. The idea of considering a triangular grid instead of a rectangular one comes from [37], and our induction-based proof is inspired by their proof. However, to our knowledge, the results of [56] and [37] do not imply our next lemma.

Lemma 3.4. *For all $r \in \mathbb{N}$, $f_2(\triangle_r) \geqslant r-1$.*

Proof. Let $e_1, \ldots, e_r$ be the r standard basis vectors in $\mathbb{R}^r$. We recursively define an embedding $\phi : V(\triangle_r) \to \mathbb{R}^r$ by $\phi(v_{1,j}) = e_j$ for all $j \in [r]$ and $\phi(v_{i,j}) = \frac{1}{2}\phi(v_{i-1,j-1}) + \frac{1}{2}\phi(v_{i-1,j})$ for all $2 \leqslant i \leqslant j$. We define an ℓ_2-realizable distance function $d : E(\triangle_r) \to \mathbb{R}_+$ from the embedding ϕ, by letting $d(vv') = ||\phi(v) - \phi(v')||_2$ for each $vv' \in E(\triangle_r)$.

Now consider an arbitrary isometric embedding ψ of $(\triangle_r, d)$ in some Euclidean space $\mathbb{E}$. By our choice of the distance function, $\psi(v_{i,j})$ is the midpoint of $\psi(v_{i-1,j-1})$ and $\psi(v_{i-1,j})$ for every $i \geqslant 2$. Hence, the whole embedding ψ is entirely determined by the r points $q_j = \psi(v_{1,j})$, and lies in the affine hull of $q_1, \ldots, q_r$. By applying an appropriate isometry, we may assume that $\mathbb{E} = \{x \in \mathbb{R}^r \mid \sum_i x_i = 1\}$. We claim that $||q_i - q_j||_2 = \sqrt{2}$ for all distinct $i, j \in [r]$. Hence, these r points are the vertices of a regular simplex, which implies $f_2(G, d) \geqslant r-1$.

The proof is by induction on r. Since the statement is clear for $r = 2$, we may assume that $r \geqslant 3$. Observe that the induced subgraphs $\triangle_r - \{v_{i,r} \mid i \in [r]\}$ and $\triangle_r - \{v_{i,i} \mid i \in [r]\}$ are both isomorphic to $\triangle_{r-1}$. By the inductive hypothesis, this implies that $q_1, \ldots, q_{r-1}$ are equidistant, and $q_2, \ldots, q_r$ are equidistant. Thus, it remains to show $||q_1 - q_r||_2 = \sqrt{2}$.

Since $||q_i - q_j||_2 = \sqrt{2}$ for all distinct $i, j \in [r-1]$, by applying an appropriate isometry we may assume that $q_k = e_k$ for all $k \in [r-1]$.

Let $x_1, \ldots, x_r \in \mathbb{R}$ denote the coordinates of q_r in $\mathbb{R}^r$. The following constraints hold:

$$\sum_i x_i = 1\,, \tag{3.1}$$

$$\sum_i x_i^2 = 1 + 2x_k \quad \forall 2 \leqslant k \leqslant r-1\,. \tag{3.2}$$

The first constraint is due to the fact that $q_r \in \mathbb{E}$, and the second is equivalent to $||\psi(v_{1,r}) - \psi(v_{1,k})||_2^2 = ||\phi(v_{1,r}) - \phi(v_{1,k})||_2^2$ (for $2 \leqslant k \leqslant r-1$), which holds by induction. Notice that $x_2 = x_3 = \cdots = x_{r-1}$ follows from (3.2). Since $v_{r-1,r-1}v_{r-1,r}$ is an edge of $\triangle_r$,

$$||\psi(v_{r-1,r-1}) - \psi(v_{r-1,r})||_2^2 = ||\phi(v_{r-1,r-1}) - \phi(v_{r-1,r})||_2^2\,. \tag{3.3}$$

Since $\psi(v_{1,j}) = \phi(v_{1,j})$ for all $j \in [r-1]$, $\psi(v_{i,j}) = \phi(v_{i,j})$ for all $i \leqslant j \leqslant r-1$. Hence, we can rewrite the left-hand side of (3.3) as

$$\begin{aligned} ||\psi(v_{r-1,r-1}) - \psi(v_{r-1,r})||_2^2 &= ||\phi(v_{r-1,r-1}) - \psi(v_{r-1,r})||_2^2 \\ &= ||(\phi(v_{r-1,r-1}) - \phi(v_{r-1,r})) - (\psi(v_{r-1,r}) - \phi(v_{r-1,r}))||_2^2 \end{aligned}$$

Thus, (3.3) holds if and only if

$$||\psi(v_{r-1,r}) - \phi(v_{r-1,r})||_2^2 = 2\left\langle \phi(v_{r-1,r}) - \phi(v_{r-1,r}), \psi(v_{r-1,r}) - \phi(v_{r-1,r})\right\rangle\,. \tag{3.4}$$

By induction, we see that, for all $i \in [r-1]$,

$$\psi(v_{i,r}) - \phi(v_{i,r}) = \frac{1}{2^{i-1}}(\psi(v_{1,r}) - \phi(v_{1,r})) = \frac{1}{2^{i-1}}(q_r - e_r)\,.$$

Using this, we can rewrite the left-hand side of (3.4):

$$\begin{aligned}||\psi(v_{r-1,r}) - \phi(v_{r-1,r})||_2^2 &= \left(\frac{1}{2^{r-2}}\right)^2 ||q_r - e_r||_2^2 \\ &= \frac{1}{2^{2r-4}}(||q_r||_2^2 + ||e_r||_2^2 - 2\langle q_r, e_r\rangle) \\ &= \frac{1}{2^{2r-4}}(1 - 2x_2 + 1 - 2x_r)\,.\end{aligned}$$

Notice that, since $x_2 = x_3 = \ldots = x_{r-1}$,

$$q_r - e_r = x_2\mathbf{1} + (x_1 - x_2)e_1 + (x_r - x_2 - 1)e_r\,,$$

where $\mathbf{1}$ is the all-ones vector. Also, an easy induction on i shows that

$$\langle \phi(v_{i,i}), e_1\rangle = \frac{1}{2^{i-1}} = \langle \phi(v_{i,r}), e_r\rangle\ ,$$

and thus

$$\begin{aligned}\langle \phi(v_{i,i}) - \phi(v_{i,r}), e_1\rangle &= \frac{1}{2^{i-1}}, \text{ and} \\ \langle \phi(v_{i,i}) - \phi(v_{i,r}), e_r\rangle &= -\frac{1}{2^{i-1}}\,.\end{aligned}$$

Now, we can rewrite the right-hand side of (3.4) as

$$\begin{aligned}&\frac{1}{2^{r-3}}\langle \phi(v_{r-1,r}) - \phi(v_{r-1,r}), q_r - e_r\rangle \\ &\quad = \frac{1}{2^{r-3}}\langle \phi(v_{r-1,r}) - \phi(v_{r-1,r}), x_2\mathbf{1} + (x_1 - x_2)e_1 + (x_r - x_2 - 1)e_r\rangle \\ &\quad = \frac{1}{2^{r-3}}\left(0 + \frac{1}{2^{r-2}}(x_1 - x_2) - \frac{1}{2^{r-2}}(x_r - x_2 - 1)\right)\,.\end{aligned}$$

Hence, (3.4) can be rewritten

$$\begin{aligned}\frac{1}{2^{2r-4}}(1 - 2x_2 + 1 - 2x_r) = \frac{1}{2^{r-3}}\left(\frac{1}{2^{r-2}}(x_1 - x_2) - \frac{1}{2^{r-2}}(x_r - x_2 - 1)\right) \\ \iff x_2 = -x_1\,.\end{aligned}$$

Now,

$$\begin{aligned}||q_r - q_1||_2^2 = ||q_r - e_1||_2^2 = \sum_i x_i^2 + 1 - 2x_1 &= (1 - 2x_2) + 1 - 2x_1 \\ &= (1 + 2x_1) + 1 - 2x_1 = 2\,.\end{aligned}$$

□

It is easy to check that $\mathrm{tw}(\triangle_r) \leqslant r - 1$ for all $r \geqslant 3$, see Lemma 2.1 on page 18. Thus, Lemma 3.4 implies that $f_2(\triangle_r) = r - 1$ for all $r \geqslant 3$. Moreover, since every planar graph is a minor of a sufficiently large triangular grid, Theorem 1.2 immediately yields the following corollary.

Corollary 3.5. *For all minor-closed classes of graphs $\mathcal{C}$, $f_2(\mathcal{C}) = \infty$ if and only if $\mathcal{C}$ contains all planar graphs.*

3.2 Alternative view of ℓ_∞-dimension

In this section, we provide a more combinatorial definition of ℓ_∞-dimension. The equivalence follows by considering potentials on a weighted auxilliary digraph.

Let D be a digraph with edge weights $l : A(D) \to \mathbb{R}$. A *potential on* (D, l) is a function $p : V(D) \to \mathbb{R}$ such that $p(w) - p(v) \leqslant l(v, w)$ for all arcs $(v, w) \in A(D)$.

Now consider a metric graph (G, d). Let (D, l) be the (edge)-weighted digraph obtained from (G, d) by bidirecting all edges and setting $l(v, w) = l(w, v) = d(vw)$ for all edges $vw \in E(G)$. Note that $p : V(D) \to \mathbb{R}$ is a potential on (D, l) if and only if $|p(w) - p(v)| \leqslant d(vw)$ for all edges $vw \in E(G)$.

For convenience, we let $D(G)$ and $l(d)$ denote the digraph and edge weights defined above, respectively. Thus, the weighted digraph (D, l) we are considering can also be denoted $(D(G), l(d))$ when more precision is required.

Recall that distances in ℓ_∞^k are given by $d_\infty(x, y) = \max_{i \in [k]} |x_i - y_i|$. Hence $d_\infty(x, y) = \delta$ if and only if $|x_i - y_i| \leqslant \delta$ for all $i \in [k]$ and there exists some index $j \in [k]$ for which $|x_j - y_j| = \delta$. Therefore, (G, d) has an isometric embedding ϕ in ℓ_∞^k if and only if there exist k potentials $p_i : V(G) \to \mathbb{R}$ on (D, l) such that for each edge vw there is at least one index $j \in [k]$ with $|p_j(w) - p_j(v)| = d(vw)$. This can be seen by taking $p_i(v)$ to be the i-th coordinate of $\phi(v)$, for all $i \in [k]$ and $v \in V(G)$.

We say that a set of arcs $F \subseteq A(D)$ is a *flat set* of (G, d) if there exists a potential $p : V \to \mathbb{R}$ on (D, l) such that $p(w) - p(v) = -d(vw)$ if and only if $p(v) - p(w) = d(vw)$ for all arcs $(v, w) \in F$. Given a set $F \subseteq A(D)$, consider the modified edge weights $l_F : A(D) \to \mathbb{R}$ such that

$$l_F(v, w) = \begin{cases} d(vw) & \text{if } (v, w) \notin F \\ -d(vw) & \text{if } (v, w) \in F\,. \end{cases}$$

When necessary, we denote these edge weights by $l_F(d)$. Then $F \subseteq A(D)$ is a flat set of (G, d) if and only if $(D, l_F) = (D(G), l_F(d))$ admits a potential. By the well-known characterization of the existence of potentials, this is equivalent to the non-existence of a negative weight directed cycle in (D, l_F). That is, $F \subseteq A(D)$ is a flat set if and only if (D, l_F) does not contain a negative directed cycle. In proofs, we will often use the notation $\langle G, d; F\rangle$ to denote $(D(G), l_F(d))$. Notice that F is a flat set if and only if $F' = \{(w, v) \mid (v, w) \in F\}$ is a flat set.

We say that a flat set $F \subseteq A(D)$ *covers* an edge $vw \in E(G)$ if F contains (v, w) or (w, v). A *flat covering* of (G, d) is a collection $\mathcal{F} = \{F_1, \ldots, F_k\}$ of flat sets such that every edge $vw \in E(G)$ is covered by at least one F_i. Then, (G, d) has an isometric embedding into ℓ_∞^k if and only if (G, d) has a flat covering of size at most k. To construct an embedding given a flat covering, we pick a potential p_i on $\langle G, d; F_i\rangle$ for each flat set F_i, and use these potentials to define the embedding coordinatewise. That is, each potential p_i associated to F_i gives us the i-th coordinate of the vertices in the embedding. Notice that the potentials respect the maximum differences given by the distance function d. Furthermore, because each edge is covered by some potential, the vertices of this edge are at exact distance in the corresponding coordinate. Hence we get an embedding of (G, d). For the other direction, it is sufficient to realize that each coordinate of an embedding defines a potential. Furthermore, for each edge at least one of the potentials defined by the coordinates is such that the distance between the vertices is attained with equality, that is the edge is covered by this potential. Thus, the coordinates define a flat covering of size k.

In our terminology, the ℓ_∞-dimension $f_\infty(G)$ is the least integer k such that for each distance function d, the metric graph (G, d) has a flat covering of size at most k.

3.3 Metric tools

In this section, we present several general results related to distance functions and flat coverings.

Given a vertex v of a graph G, we let $N(v) = \{w \in V(G) \mid vw \in E(G)\}$ denote the neighborhood of v in G.

Lemma 3.6. *Let (G, d) be a metric graph and let $v \in V(G)$. The set*

$F = \{(v, w) \mid w \in N(v)\}$ is a flat set of (G, d).

Proof. Let C be an arbitrary directed cycle in $\langle G, d; F\rangle$. The cycle C uses at most one arc of F. Thus, at most one arc of C has negative weight in $\langle G, d; F\rangle$, and all other arcs of C have non-negative weight. Since d is a distance function, it follows that C has non-negative weight in $\langle G, d; F\rangle$. Thus, F is a flat set of (G, d), as required. □

A *vertex cover* of a graph G is a set of vertices $X \subseteq V(G)$ such that every edge of G is incident with some vertex in X. The *vertex cover number* of G, denoted $\tau(G)$, is the size of a smallest vertex cover of G. By Lemma 3.6, $f_\infty(G)$ is at most the vertex cover number of G.

Lemma 3.7 ([42], Lemma 9)**.** *For every graph G, $f_\infty(G) \leqslant \tau(G)$.*

Clearly, if d is a distance function on G, and H is a subgraph of G, then the restriction of d to $E(H)$ is a distance function on H. We denote it by $d|_H$. Conversely, sometimes we can define a distance function on a graph from distance functions on certain subgraphs, see Lemma 3.8 below.

A *k-sum* is a graph G obtained by gluing two graphs G_1 and G_2 along a common clique K of size k and then possibly deleting some edges of K. We use the following notation for 1-sums and 2-sums. We write $G = G_1 +_v G_2$ if $G = G_1 \cup G_2$ with $V(G_1) \cap V(G_2) = \{v\}$. Now let $e = vw$ be an edge. We write $G = G_1 \oplus_e G_2$ if $G = G_1 \cup G_2$ with $V(G_1) \cap V(G_2) = \{v, w\}$ and $e \in E(G_1) \cap E(G_2)$. Also, we denote by $G_1 +_e G_2$ the graph $G_1 \oplus_e G_2$ minus the edge e.

Lemma 3.8. *Let $G = G_1 \oplus_f G_2$. For $i \in [2]$, let d_i be a distance function on G_i. If $d_1(f) = d_2(f)$, then the function $d : E(G) \to \mathbb{R}_{\geqslant 0}$ defined by $d(e) = d_i(e)$ if $e \in E(G_i)$ is a distance function on G.*

Proof. Let vw be any edge of G. Without loss of generality, we may suppose $vw \in E(G_1)$. Let P be a v–w path in G. If P is contained in G_1 then $d(P) = d_1(P) \geqslant d_1(vw) = d(vw)$. Otherwise, P uses both ends of f and we may decompose P into a path P_1 from v to an end of f with $E(P_1) \subseteq E(G_1)$, a path P_2 between the two ends of f with $E(P_2) \subseteq E(G_2)$ and a path P_1' from the other end of f to w with $E(P_1') \subseteq E(G_1)$. Then we get $d(P) = d(P_1) + d(P_2) + d(P_1') \geqslant d(P_1) + d(f) + d(P_1') \geqslant d(vw)$, where the first inequality uses that d_2 is a distance function, and the second inequality uses that d_1 is a distance function. □

Similarly, every subset of a flat set is flat, and if F is a flat set of (G, d), then F is also a flat set of $(H, d|_H)$, for all subgraphs H of G with $F \subseteq A(D(H))$. The following lemma gives conditions under which a flat set of a subgraph is a flat set of the entire graph.

Lemma 3.9. *Let G be a graph obtained by gluing two graphs G_1 and G_2 along a common clique K. Let d be a distance function on G and $d_i = d|_{G_i}$ its restriction to G_i, where $i \in [2]$. If F is a flat set of (G_j, d_j) for some $j \in [2]$, then F is also a flat set of (G, d). Conversely, if F is a flat set of (G, d) then $F_i = F \cap A(D(G_i))$ is a flat set of (G_i, d_i) for all $i \in [2]$.*

Proof. For the first part, it suffices to show that $\langle G, d; F\rangle$ does not contain a negative weight directed cycle. Let C be a minimum weight directed cycle in $\langle G, d; F\rangle$ such that $V(C)$ is inclusion-wise minimal. We may assume that C contains some arc of F, since otherwise C is disjoint from F and has non-negative weight. Thus C intersects $A(D(G_j))$.

We claim that C must be fully contained in $D(G_j)$. Otherwise, C contains a directed path P from v to w, where $v, w \in K$, that is internally disjoint from $D(G_j)$. By replacing P with the arc (v, w) we obtain a new directed cycle C' in $\langle G, d; F\rangle$ whose weight is at most that of C and such that $V(C') \subsetneq V(C)$, a contradiction.

Since C is contained in $D(G_j)$ and F is a flat set of (G_j, d_j), C has non-negative weight in $\langle G_j, d_j; F\rangle$ and thus in $\langle G, d; F\rangle$.

For the second part, notice that F_i is a flat set of (G, d) because $F_i \subseteq F$ and F is a flat set of (G, d). Since G_i is a subgraph of G, F_i is also clearly a flat set of (G_i, d_i). □

Lemma 3.10. *Let F be a flat set of a metric graph (G, d) and u and v be vertices of G. Let P_1 be a directed path from u to v and let P_2 be a directed path from v to u. Then at least one of P_1 and P_2 has non-negative weight in $\langle G, d; F\rangle$.*

Proof. Consider the directed closed walk obtained by concatenating P_1 and P_2. This directed closed walk decomposes into directed cycles. If P_1 and P_2 both have negative weight in $\langle G, d; F\rangle$, then at least one of these directed cycles has negative weight in $\langle G, d; F\rangle$. But this contradicts the fact that F is a flat set. □

In [42], the following result is proved.

Lemma 3.11 ([42])**.** *For every graph G with $f_\infty(G) \geqslant 2$ and every edge $e \in E(G)$,*

$$f_\infty(G) = f_\infty(G +_e K_3) = f_\infty(G \oplus_e K_3).$$

Hence, deleting a degree-2 vertex v and adding a new edge between the neighbors of v (if there was none) does not change $f_\infty(G)$, provided the resulting graph is not a forest. We will refer to this operation as *suppressing a degree-2 vertex*. It follows that for all $k \geqslant 2$, the minimal excluded minors for $f_\infty(G) \leqslant k$ have minimum degree at least 3.

We will use the following bounds on $f_\infty(G)$ when G is a k-sum.

Lemma 3.12. *For all graphs G_1 and G_2 (for which the k-sums below exist),*

$$f_\infty(G_1 +_v G_2) = \max\{f_\infty(G_1), f_\infty(G_2)\} \tag{3.5}$$

and

$$f_\infty(G_1 +_{vw} G_2) \leqslant f_\infty(G_1 \oplus_{vw} G_2) \leqslant f_\infty(G_1) + f_\infty(G_2) - 1\,. \tag{3.6}$$

Moreover,

$$f_\infty(G) \leqslant f_\infty(G_1) + f_\infty(G_2) \tag{3.7}$$

whenever G is a k-sum of G_1 and G_2.

Proof. Observe that (3.7) follows from Lemma 3.9. Next, we prove (3.5). Let $k = \max\{f_\infty(G_1), f_\infty(G_2)\}$. Since f_∞ is minor-monotone, it is clear that $f_\infty(G_1 +_v G_2)$ is at least k. The next paragraph proves that it is at most k.

Let d be a distance function on $G_1 +_v G_2$. For $i \in [2]$, let $d_i = d|_{G_i}$. Then d_i is a distance function on G_i. For $i \in [2]$, let ϕ_i be any isometric embedding of (G_i, d_i) into ℓ_∞^k. After translating one of the embeddings if necessary, we may assume that $\phi_1(v) = \phi_2(v)$. It is easy to see that the function $\phi : V(G_1 +_v G_2) \to \mathbb{R}^k$ obtained by setting $\phi(w) = \phi_i(w)$ if $w \in V(G_i)$ for $i \in [2]$ is an isometric embedding of $(G_1 +_v G_2, d)$ into ℓ_∞^k.

Finally, we prove (3.6). The first inequality in (3.6) is trivial since $G_1 +_{vw} G_2$ is a minor of $G_1 \oplus_{vw} G_2$. To prove the second inequality, consider a distance function d on G. For $i \in [2]$, let $d_i = d|_{G_i}$ be the corresponding distance function of G_i.

Let $\mathcal{F}_i$ be a minimum size flat covering of (G_i, d_i). By Lemma 3.9, each set in $\mathcal{F}_1 \cup \mathcal{F}_2$ is flat in (G, d). For $i \in [2]$, let F_i be a flat set in $\mathcal{F}_i$ covering

vw. By reversing arcs if necessary, we may assume both F_1 and F_2 contain (v, w). We may also assume that neither F_1 nor F_2 contains (w, v), since otherwise we get $d(vw) = 0$. In this case, we can contract the edge vw and use (3.5).

We claim that $F_1 \cup F_2$ is a flat set of (G, d). Let C be an arbitrary directed cycle in $\langle G, d; F_1 \cup F_2 \rangle$. For $i \in [2]$, let C_i be the directed cycle obtained by restricting C to $D(G_i)$ and possibly adding (v, w) or (w, v) (possibly $C_i = \emptyset$). Let $l = l_{F_1 \cup F_2}(d)$ be the edge weights on $\langle G, d; F_1 \cup F_2 \rangle$ and $l_i = l_{F_i}(d_i)$ be the edge weights on $\langle G_i, d_i; F_i \rangle$. Notice that $l(v, w) = -d(vw)$ and $l(w, v) = d(vw)$. Then $l(C) = l(C_1) + l(C_2) = l_1(C_1) + l_2(C_2) \geqslant 0 + 0 = 0$ since l_i is the restriction of l to $A(D(G_i))$ and F_i is flat in (G_i, d_i). Thus, C has non-negative weight and $F_1 \cup F_2$ is a flat set of (G, d), as claimed.

Now $\mathcal{F} = \{F_1 \cup F_2\} \cup (\mathcal{F}_1 \cup \mathcal{F}_2) \setminus \{F_1, F_2\}$ is a flat covering of (G, d) of size at most $|\mathcal{F}_1| + |\mathcal{F}_2| - 1 \leqslant f_\infty(G_1) + f_\infty(G_2) - 1$. □

Let (G, d) be a metric graph. We say that two edges e and f of G are *incompatible*, if there is no flat set of (G, d) that covers both of them. Note that two such edges are necessarily independent, by Lemma 3.6. A simple but crucial observation is that if (G, d) contains k pairwise incompatible edges, then $f_\infty(G) \geqslant k$. The following lemma provides sufficient conditions under which two edges are incompatible.

Lemma 3.13. *Let (G, d) be a metric graph and let v_1v_2, w_1w_2 be two independent edges of G. If for all $i, j \in [2]$, there exist paths $P_{i,j}$ between v_i and w_j such that $d(P_{1,1}) + d(P_{2,2}) < d(v_1v_2) + d(w_1w_2)$ and $d(P_{1,2}) + d(P_{2,1}) < d(v_1v_2) + d(w_1w_2)$, then v_1v_2 and w_1w_2 are incompatible.*

Proof. Suppose F is a flat set covering v_1v_2 and w_1w_2. Suppose first $(v_1, v_2), (w_1, w_2) \in F$. Consider the closed directed walk W that starts at v_1, takes (v_1, v_2), follows $P_{2,1}$ to w_1, takes (w_1, w_2) and then follows $P_{1,2}$ back to v_1. The weight of W in $\langle G, d; F \rangle$ is at most $d(P_{1,2}) + d(P_{2,1}) - d(v_1v_2) - d(w_1w_2) < 0$. Thus, W contains a negative weight directed cycle, which contradicts that F is flat.

By symmetry the remaining case is $(v_1, v_2), (w_2, w_1) \in F$. Again it is easy to find a negative weight directed walk W in $\langle G, d; F \rangle$ using the fact that $d(P_{1,1}) + d(P_{2,2}) < d(v_1v_2) + d(w_1w_2)$. Hence, F cannot simultaneously cover the edges v_1v_2 and w_1w_2, as claimed. □

Finally, we also need the fact that $f_\infty(K_4) = 2$. We also show here that $f_\infty(K_n) \leqslant n-2$, which has been claimed in the introduction.

Lemma 3.14 ([79], 4.2)**.** *$f_\infty(K_4) = 2$.*

Proof. Let v_1, v_2, v_3, v_4 be the vertices of K_4 such that $d(v_1v2) + d(v_3v4) \geqslant d(v1v3) + d(v_2v4) \geqslant d(v_1v_4) + d(v_2v_3)$. It is easy to check that $\{(v_1, v_2), (v_1, v_3), (v_4, v_2)\}$ and $\{(v_3, v_4), (v_1, v_4), (v_3, v_2)\}$ are flat sets. Hence we get $f_\infty(K_4) = 2$. □

Lemma 3.15. *$f_\infty(K_n) \leqslant n-2$.*

Proof. Let $v_1, \dots, v_n$ be the vertices of K_n. For $i \in [n-4]$, the set $\{v_iv_j \mid i < j\}$ is a flat set by Lemma 3.6. The remaining edges form a K_4-subgraph and can be covered with two flatsets by Lemma 3.14. Hence, a total of $n-2$ flatsets is sufficient to cover all edges of K_n. □

In order to illustrate the concepts introduced in the last two sections, we briefly describe a polynomial reduction from computing the chromatic number of a graph H to computing $f_\infty(G, d)$ given a metric graph (G, d). This proves that the latter problem is NP-hard. We remark that there is a different reduction using the Partition problem which shows that the problem of deciding if $f_\infty(G, d) \leqslant 1$ given a metric graph (G, d) is NP-complete (see [69]).

Lemma 3.16. *Deciding $f_\infty(G, d) = k$ is NP-hard.*

Proof. Let H be a graph. We construct a metric graph (G, d) by replacing each vertex $v \in V(H)$ by two adjacent vertices $v_1, v_2 \in V(G)$, and each edge $vw \in E(H)$ by a $K_{2,2}$ in G with edge set $\{v_iw_j \mid i \in [2],\ j \in [2]\}$. The distance function d is defined by $d(v_1v_2) = 2$ for all $v \in V(H)$ and $d(v_iw_j) = 1$ for all $vw \in E(H)$, $i \in [2]$ and $j \in [2]$. We claim that $f_\infty(G, d) = \chi(H)$.

To see that $f_\infty(G, d) \geqslant \chi(H)$, notice that edges v_1v_2 and w_1w_2 are incompatible whenever $vw \in E(H)$. Thus every size-k flat covering of (G, d) gives a k-coloring of H.

Finally, $f_\infty(G, d) \leqslant \chi(H)$, since for every stable set S in G, $\{(v_1, v_2) \mid v \in S\} \cup \{(u_i, v_1) \mid i \in [2],\ uv \in E(H),\ v \in S\} \cup \{(v_2, w_j) \mid j \in [2],\ vw \in E(H),\ v \in S\}$ is a flat set of (G, d). Hence, every k-coloring of H gives a size-k flat covering of (G, d). □

3.4 Certificates of large ℓ_∞-dimension

In this section, we show that if $H \in \mathcal{U}_\infty^k = \{\mathsf{S}_k, \mathsf{P}_k, \mathsf{F}_k, \mathsf{N}_k\}$, then $f_\infty(H) > k$. It follows that if a graph G contains a $\mathcal{U}_\infty^k$ minor, then $f_\infty(G) > k$. Therefore, the existence of one of these four minors is a certificate that $f_\infty(G) > k$. Conversely, our main theorem shows that if $f_\infty(G) \geqslant g_{1.3}(k)$, then G necessarily contains one of these four minors. We also prove that $\mathsf{S}_k, \mathsf{P}_k$, and F_k are minimal excluded minors for the property $f_\infty(G) \leqslant k$, that is, all their proper minors have ℓ_∞-dimension at most k.

We begin by proving that for each $H \in \{\mathsf{S}_k, \mathsf{P}_k, \mathsf{F}_k\}$, $f_\infty(H) = k+1$. We first prove the upper bound.

Lemma 3.17. *For all $k \in \mathbb{N}$ and all $H \in \{\mathsf{S}_k, \mathsf{P}_k, \mathsf{F}_k\}$, $f_\infty(H) \leqslant k+1$.*

Proof. We proceed by induction on k. The base case follows by Lemma 3.14, since $\mathsf{S}_1 = \mathsf{P}_1 = \mathsf{F}_1 = K_4$. Next note that $\mathsf{S}_k = \mathsf{S}_{k-1} +_e K_4, \mathsf{P}_k = \mathsf{P}_{k-1} +_e K_4$, and $\mathsf{F}_k = \mathsf{F}_{k-1} +_e K_4$. Therefore, we are done by induction and Lemmas 3.12 and 3.14. ☐

Theorem 3.18. *For all $k \in \mathbb{N}$, $f_\infty(\mathsf{S}_k) = k+1$.*

Proof. By Lemma 3.17, it suffices to show $f_\infty(\mathsf{S}_k) \geqslant k+1$. Since $S_1 = K_4$, by Lemma 3.14, we may assume $k \geqslant 2$. We now give a distance function d on S_k, which is illustrated in Figure 3.5, such that there are $k+1$ incompatible edges in (S_k, d).

Let $V(\mathsf{S}_k) = \{v, w\} \cup \{v_1, w_1, \dots, v_k, w_k\}$ where v, w, v_i, w_i are the vertices of the ith copy of K_4. We define d as follows:

$$\begin{aligned} d(vv_1) = d(ww_1) &= 4k\,, && \\ d(vv_i) = d(ww_i) &= 2(k+i-1) && \text{for all } i \in [k],\ i \neq 1\,, \\ d(wv_i) = d(vw_i) &= k+i-1 && \text{for all } i \in [k]\,, \\ d(v_iw_i) &= 3(k+i-1) && \text{for all } i \in [k]\,. \end{aligned}$$

First, we show that d is a distance function. For this, let (G, d') be obtained from (S_k, d) by adding the edge vw of length $d'(vw) = 3k$. Observe that

$$G = K_4 \oplus_{vw} K_4 \oplus_{vw} \cdots \oplus_{vw} K_4,$$

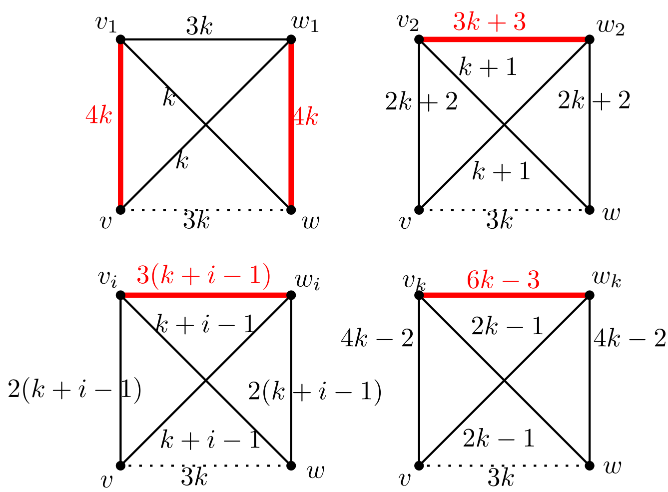

Figure 3.5. (S_k, d) as in the proof of Theorem 3.18. The red edges are pairwise incompatible. Vertices with the same label are identified.

where K_4 appears k times in the righthand side. It is easy to see that the restriction of d' to each K_4 subgraph of G is a distance function. Therefore, by Lemma 3.8, d' is a distance function on G. Since d is a restriction of d' to S_k it follows that d is a distance function on S_k.

We now show that the $k+1$ edges $vv_1, ww_1, v_2w_2, v_3w_3, \ldots, v_kw_k$ are pairwise incompatible. For this, we make repeated use of Lemma 3.13.

First, consider vv_1 and ww_1. Observe that $d(vv_1)+d(ww_1) = 8k$. However, $d(vw_1) + d(wv_1) = 2k < 8k$ and $d(v_1w_1) + d(vv_2w) = 6k + 3 < 8k$, since $k \geqslant 2$. By Lemma 3.13, vv_1 and ww_1 are incompatible.

Next, consider vv_1 and v_iw_i with $i \in \{2, \ldots, k\}$. Observe that $d(vv_1) + d(v_iw_i) = 7k+3i-3$. However, $d(vv_i)+d(w_iwv_1) = 5k+2i-2 < 7k+3i-3$ and $d(vw_i) + d(v_iwv_1) = 3k + 2i - 2 < 7k + 3i - 3$. Hence, by Lemma 3.13, vv_1 and v_iw_i are incompatible.

By symmetry, ww_1 and v_iw_i are also incompatible for each $i \in \{2, \ldots, k\}$.

Finally, consider v_iw_i and v_jw_j for $2 \leqslant i < j \leqslant k$. Observe that $d(v_iw_i) + d(v_jw_j) = 6k+3i+3j-6$. However, $d(v_iwv_j)+d(w_ivw_j) = 4k+2i+2j-4 < 6k+3i+3j-6$, and $d(v_ivw_j)+d(w_iwv_j) = 6k+4i+2j-6 < 6k+3i+3j-6$ since $i < j$. Hence, by Lemma 3.13, v_iw_i and v_jw_j are incompatible, which completes the proof. □

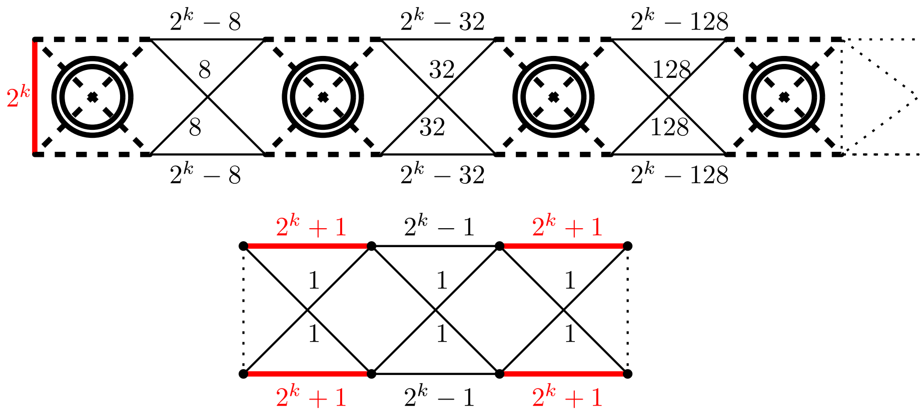

Figure 3.6. The top half of the figure depicts the distance function on P_k used in the proof of Theorem 3.19. The dashed crosses with a double circle are each to be replaced with the metric graph shown in the bottom half of the figure.

Theorem 3.19. *For all $k \in \mathbb{N}$, $f_\infty(\mathsf{P}_k) = k + 1$.*

Proof. Again, $f_\infty(\mathsf{P}_k) \leqslant k + 1$ follows from Lemma 3.17. We label the vertices of the topmost path of P_k as $v_0, v_1, \ldots, v_k$ and the vertices of the bottommost path of P_k as $w_0, w_1, \ldots, w_k$. Thus $V(\mathsf{P}_k) = \{v_0, v_1, \ldots, v_k\} \cup \{w_0, w_1, \ldots, w_k\}$ and $E(\mathsf{P}_k) = \{v_0w_0, v_kw_k\} \cup \{v_{i-1}v_i, v_{i-1}w_i, w_{i-1}v_i, w_{i-1}w_i \mid i \in [k]\}$. For the lower bound, consider the following distance function d, which is illustrated in Figure 3.6 (we take $i \in [k]$):

$$
\begin{aligned}
d(v_0w_0) = d(v_kw_k) &= 2^k\,, \\
d(v_{i-1}v_i) = d(w_{i-1}w_i) &= 2^k + 1 && \text{if } i \equiv 1 \pmod 2\,, \\
d(v_{i-1}v_i) = d(w_{i-1}w_i) &= 2^k - 1 && \text{if } i \equiv 2 \pmod 4\,, \\
d(v_{i-1}v_i) = d(w_{i-1}w_i) &= 2^k - 2^{1+i/2} && \text{if } i \equiv 0 \pmod 4\,, \\
d(v_{i-1}w_i) = d(w_{i-1}v_i) &= 2^{1+i/2} && \text{if } i \equiv 0 \pmod 4\,, \\
d(v_{i-1}w_i) = d(w_{i-1}v_i) &= 1 && \text{if } i \not\equiv 0 \pmod 4\,.
\end{aligned}
$$

Let (G, d') be obtained from (P_k, d) by adding edges v_iw_i with $d'(v_iw_i) = 2^k$ for all $i \in [k-1]$. Notice that for all i, the length of a shortest path between v_i and w_i in (P_k, d) is 2^k. Therefore, (P_k, d) is a metric graph if and only if (G, d') is a metric graph. Observe that the restriction of d' to every K_4

subgraph of G is a distance function. Therefore, (G, d') and hence also (P_k, d) is a metric graph by Lemma 3.8.

Consider the matching $M = \{v_{i-1}v_i, w_{i-1}w_i \mid i \equiv 1 \pmod 2\}$. If k is even, then we also add the edge $v_k w_k$ to M. Thus $|M| = k+1$ always. We claim that the edges of M are pairwise incompatible. To see this, let $e = xx'$ and $f = yy'$ be distinct edges of M. Let P be a shortest x–y path, and P' be a shortest x'–y' path. We claim that $d(P) + d(P') \leqslant 2 \cdot 2^k$ (see next paragraph for a proof). However, $d(e) + d(f) > 2 \cdot 2^k$ because $e, f \in M$. Therefore, by Lemma 3.13, e and f are incompatible. Since $|M| = k+1$, $f_\infty(\mathsf{P}_k) \geqslant k+1$, as required.

To prove the claim, we split the discussion into two cases. A *segment* in P_k is any subgraph induced by $\{v_i, w_i \mid i = 4q + r,\ r \in \{0,1,2,3\},\ i \leqslant k\}$ for some q. If e and f belong to the same segment, then it is easy to see that $d(P) + d(P') \leqslant 2 \cdot 2^k$. (Notice that sometimes $d(P) = 2^k + 1$ and $d(P') = 2^k - 1$.) Now if a and b are any two vertices in distinct segments (indexed by q and s, with $q < s$), then there is a a–b path Q such that

$$
\begin{aligned}
d(Q) &\leqslant 1 + 1 + 1 + 2^{2q+3} + 1 + 1 + 1 + \cdots \\
&\qquad + 2^{2s-1} + 1 + 1 + 1 + (2^k - 2^{2s+1}) + 1 + 1 + 1 \\
&\leqslant \underbrace{(3s+3)}_{\leqslant 1+2+4+2^{2s}} + 2^3 + 2^5 + \cdots + 2^{2s-1} - 2^{2s+1} + 2^k \\
&\leqslant \sum_{i=0}^{2s} 2^i - 2^{2s+1} + 2^k \leqslant 2^k \,.
\end{aligned}
$$

It follows that $d(P) + d(P') \leqslant 2 \cdot 2^k$ in this case too. □

Theorem 3.20. *For all $k \in \mathbb{N}$, $f_\infty(\mathsf{F}_k) = k+1$.*

Proof. For all $i \in [k]$, we label the vertices of the ith copy of K_4 in F_k as $v_0, v_{2i-1}, v_{2i}, v_{2i+1}$. Remember that in order to obtain F_k we form the 2-sum of these k copies of K_4 and delete every edge that is in two consecutive copies. Thus $V(\mathsf{F}_k) = \{v_j \mid j \in \{0, \ldots, 2k+1\}\}$ and $E(\mathsf{F}_k) = \{v_0v_1, v_0v_{2k+1}\} \cup \{v_0v_{2i}, v_{2i-1}v_{2i}, v_{2i-1}v_{2i+1}, v_{2i}v_{2i+1}\}$.

By Lemma 3.17, it suffices to show $f_\infty(\mathsf{F}_k) \geqslant k+1$. Consider the following distance function d on F_k:

$$
\begin{aligned}
d(v_0v_1) &= 1\,, && \\
d(v_0v_{2i}) &= 1 && \text{for } i \in [k]\,,\\
d(v_{2i-1}v_{2i+1}) &= 1 && \text{for } i \in [k]\,,\\
d(v_{2i}v_{2i+1}) &= i && \text{for } i \in [k]\,,\\
d(v_{2i}v_{2i-1}) &= i+1 && \text{for } i \in [k]\,,\\
d(v_0v_{2k+1}) &= k+1\,. &&
\end{aligned}
$$

As before, by Lemma 3.8, we can prove that d is a distance function. Notice that v_0 is at distance $i+1$ from v_{2i+1} for each $i \in [k-1]$.

Consider the matching $M = \{v_0v_{2k+1}\} \cup \{v_{2i}v_{2i-1} \mid i \in [k]\}$ in (F_k, d). See Figure 3.7 for an illustration of the distance function d and the matching M in F_5.

Finally, we need to show that all $k+1$ edges of M are pairwise incompatible. If $i < j$, then

$$
\begin{aligned}
d(v_{2i}, v_{2j-1}) &\leqslant j-1,\\
d(v_{2i-1}, v_{2j}) &\leqslant i+2,\\
d(v_{2i}, v_{2j}) &\leqslant 2,\\
d(v_{2i-1}, v_{2j-1}) &\leqslant j-i
\end{aligned}
$$

and

$$
\begin{aligned}
d(v_0, v_{2i}) &\leqslant 1,\\
d(v_{2k+1}, v_{2i-1}) &\leqslant k-i+1\\
d(v_0, v_{2i-1}) &\leqslant i+1,\\
d(v_{2k+1}, v_{2i}) &\leqslant k.
\end{aligned}
$$

□

Theorem 3.21. *For all $k \geqslant 2$, $\mathsf{S}_k, \mathsf{P}_k, \mathsf{F}_k$ are minimal excluded minors for the property $f_\infty(G) \leqslant k$.*

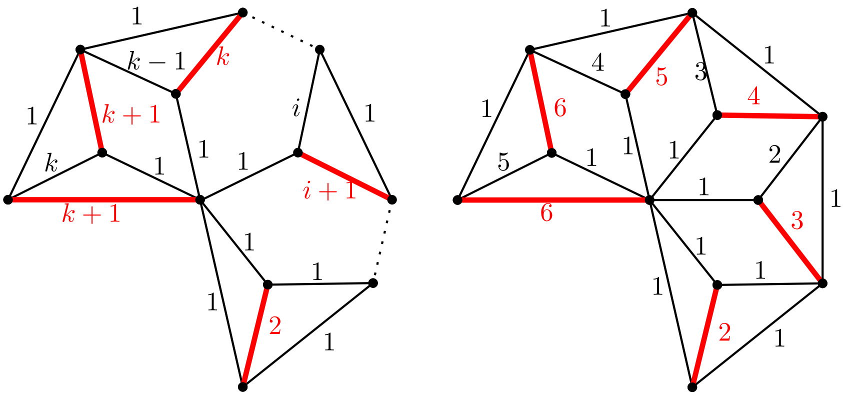

Figure 3.7. (F_k, d) as in the proof of Theorem 3.20 and (F_5, d). The red edges are pairwise incompatible.

Proof. Let H be one of $\mathsf{S}_k, \mathsf{P}_k, \mathsf{F}_k$. By Theorems 3.18, 3.19, and 3.20, we know $f_\infty(H) > k$.

When deleting or contracting an edge in H, we get a minor H' which can be expressed as a 2-sum of two graphs H_1, H_2 with the following properties. First, $H_1 \in \{\mathsf{S}_\ell, \mathsf{P}_\ell, \mathsf{F}_\ell\}$ for some $\ell < k$ (and H_1 is of the same type as H). Second, H_2 has a degree-2 vertex and recursively suppressing the degree-2 vertices from H_2 results in a graph H_2' such that $H_2' \in \{\mathsf{S}_m, \mathsf{P}_m, \mathsf{F}_m\}$ for some $m \leqslant k - l - 1$ (again H_2' is of the same type as H), or H_2' is a single edge (this corresponds to the case $m = 0$).

By Lemma 3.12 and Lemma 3.17,

$$\begin{aligned} f_\infty(H') &\leqslant f_\infty(H_1) + f_\infty(H_2) - 1 = f_\infty(H_1) + f_\infty(H_2') - 1 \\ &\leqslant (l+1) + (m+1) - 1 \leqslant k. \end{aligned}$$

Thus, H is a minimal excluded minor for $f_\infty(G) \leqslant k$. □

Theorem 3.22. *For all $k \in \mathbb{N}$, $f_\infty(\mathsf{N}_k) \geqslant k + 1$.*

Proof. Let $V(\mathsf{N}_k) = \{v_0, \ldots, v_k\} \cup \{w_0, \ldots, w_k\}$ and

$$E(\mathsf{N}_k) = \{v_{i-1}v_i, v_iw_i, v_{i-1}w_i, w_{i-1}w_i \mid i \in [k]\} \cup \{v_0w_0, w_0v_k\}.$$

Consider the distance function d such that $d(w_0v_k) = d(v_{i-1}v_i) = d(w_{i-1}w_i) = 1$, $d(v_{i-1}w_i) = k$ for all $i \in [k]$ and $d(v_iw_i) = k + 1$ for all

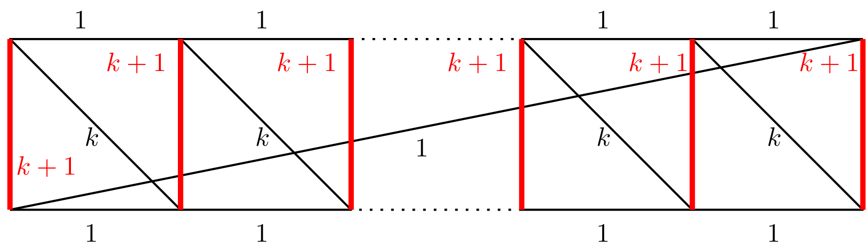

Figure 3.8. (N_k, d) as in the proof of Theorem 3.22.

$i = 0, \ldots, k$. It is easy to check that d is indeed a distance function. Let $M = \{v_i w_i \mid i = 0, \ldots, k\}$. See Figure 3.8 for an illustration of (N_k, d) and M, where $v_0 \cdots v_k$ and $w_0 \cdots w_k$ are the topmost and bottommost paths, respectively.

We claim that the edges in M are pairwise incompatible. To see this, first observe that the shortest v_i–v_j and w_i–w_j paths both have weight $|j - i| \leqslant k$ since all edges in these paths have weight 1, hence the cumulative weight of these paths is at most $2k$. If $i > j$, then

$$\begin{aligned} &d(v_i v_{i+1} \cdots v_k w_0 w_1 \cdots w_j) + d(v_j v_{j+1} \cdots v_{i-1} w_i) \\ &= (k - i + j + 1) + (i - j - 1 + k) = 2k. \end{aligned}$$

This shows that there exist a v_i–w_j path and a v_j–w_i path of cumulative weight $2k$. Since $d(v_i w_i) + d(v_j w_j) = 2k + 2$, the conditions of Lemma 3.13 are satisfied and we get that $v_i w_i$ and $v_j w_j$ are incompatible for all $i \neq j$. Hence, $f_\infty(\mathsf{N}_k) \geqslant k + 1$. □

Since N_k is 3-connected, it is difficult to adapt the proof of Theorem 3.21 to show that N_k is also a minimal excluded minor for the property $f_\infty(G) \leqslant k$. However, we conjecture that this is true.

3.5 2-connected graphs

In this section, we show that it is enough to prove our main theorem, Theorem 1.3, for 3-connected graphs. To do so, we introduce a variant of SPQR trees in Section 3.5.1. In section 3.5.2, we show that in a graph $G_1 +_e G_2$ obtained as a 2-sum of two graphs G_1 and G_2, we can merge flat sets from G_1 and G_2 under some conditions. In Section 3.5.3, we present several lemmas that show how to bound $f_\infty(H)$, where H is obtained by gluing several 2-connected graphs on a given graph. At the end of this section, we also show how to complete the proof of Theorem 1.3 under some additional assumptions.

3.5.1 Contracted SPQR trees

In this context we need to consider *multigraphs* that are minors of a simple 2-connected graph, that is, parallel edges resulting from edge contractions are kept. (Loops on the other hand are not important for our purposes and thus can safely be discarded.) SPQR trees were introduced in [30] as a way to decompose a 2-connected graph across its 2-separations. They are defined as follows.

Let G be a (simple) 2-connected graph. The *SPQR tree* T_G *of* G is a tree each of whose node $a \in V(T_G)$ is associated with a multigraph H_a which is a minor of G. Each vertex $x \in V(H_a)$ is a vertex of G, that is, $V(H_a) \subseteq V(G)$. Each edge $e \in E(H_a)$ is classified either as a *real* or *virtual* edge. By the construction of an SPQR tree each edge $e \in E(G)$ appears in exactly one minor H_a as a real edge, and each edge $e \in H_a$ which is classified real is an edge of G. The SPQR tree T_G is defined recursively as follows.

1. If G is 3-connected, then T_G consists of a single *R-node* a for which we have $H_a = G$. All edges of H_a are real in this case.

2. If G is a cycle, then T_G consists of a single S-*node* for which $H_a = G$. Again, all edges of H_a are real in this case.

3. Otherwise G has a cutset $\{x, y\}$ such that the vertices x and y have degree at least 3. In this case we construct T_G inductively. First we add a P-*node* a to T_G, for which H_a is the graph consisting of the single edge xy. The edge xy of H_a is real if xy is an edge of G, and virtual otherwise. Next we consider the connected components $C_1, \ldots, C_r$ ($r \geqslant 2$) of $G - \{x, y\}$. Let G_i be the graph $G[V(C_i) \cup \{x, y\}]$ with the additional edge xy if it is not already there. Since we include the edge xy, each G_i is 2-connected and we can construct the corresponding SPQR tree T_{G_i} by induction. Let a_i be the (unique) node in T_{G_i} for which xy is a real edge in H_{a_i}. In order to construct T_G, we make xy a virtual edge in the node a_i, and connect a_i to a in T_G. Finally,

we add parallel virtual edges xy to H_a so that it has exactly r virtual edges xy.

Notice that minors corresponding to S-nodes and R-nodes are simple graphs, whereas those corresponding to P-nodes are multigraphs consisting of two vertices linked by at least two virtual edges and possibly a real one. To each edge ab of the SPQR tree T_G corresponds a unique virtual edge $e \in E(H_a) \cap E(H_b)$ with ends $x, y \in V(G)$. Thus we can define a corresponding multigraph $H_{a,b}$ which is the minor of G obtained by taking the 2-sum of H_a and H_b in which the edge e is deleted. (To be precise, one virtual edge xy from each of H_a and H_b is deleted in the operation, other copies of xy, if any, are kept in the resulting graph.) Similarly, we can define a unique minor of G for each *subtree* of T_G by performing one 2-sum operation as described above for each edge of the subtree.

Let G be a 2-connected graph, and let T_G be the SPQR tree of G. We define the *contracted SPQR tree* T'_G as the tree obtained from T_G by contracting every maximal connected subtree of T_G each of whose nodes is either a S-node or a P-node, see Figure 3.9 for an example. We call the new nodes resulting from the contraction *O-nodes*. Each node a of T'_G has a unique corresponding minor H_a of G. If a is an R-node, then we keep the same minor as in T_G. Otherwise, a is an O-node and H_a is the minor of G corresponding to the subtree of T_G that was contracted to node a of T'_G.

We quickly give some standard terminology before stating our first result of the section. The *length* of a path in G is its number of edges. The *diameter* of a graph G is the maximum length of a shortest path between any two vertices.

Lemma 3.23. *Let G be a 2-connected graph with minimum degree at least 3.*

1. *Every O-node in T'_G corresponds to a 2-connected treewidth-2 graph.*
2. *All leaves of T'_G are R-nodes.*
3. *If the diameter of T'_G is at least $6k$, then G contains P_k or F_k as a minor.*

Proof. (1) Let o be an O-node of T'_G. Its corresponding minor H_o is obtained by 2-sums from cycles corresponding to S-nodes, and parallel edges corresponding to P-nodes. Hence H_o is 2-connected and has treewidth 2.

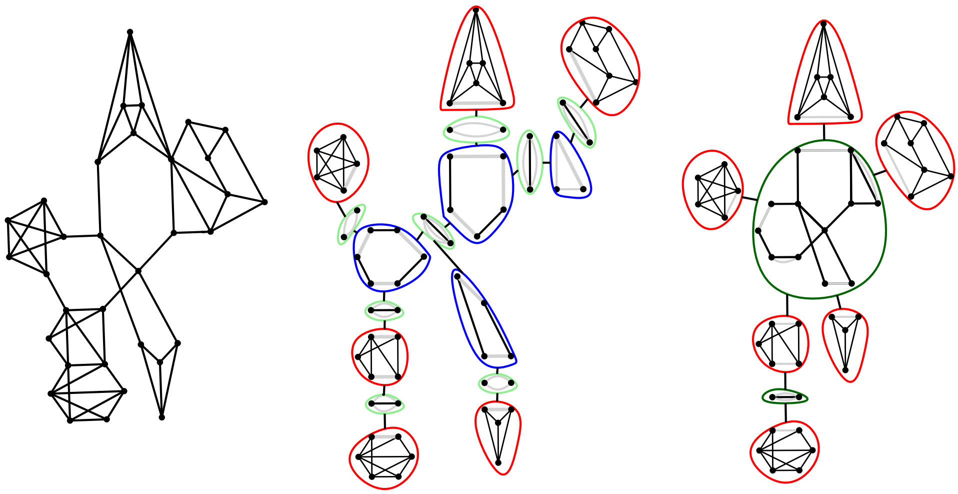

Figure 3.9. An example of a 2-connected graph G, its SPQR tree T_G, and the contracted SPQR tree T'_G.

(2) Suppose for a contradiction that some leaf o of T'_G is an O-node. Since a P-node cannot be a leaf in T_G, the subtree corresponding to o in T_G has at least one leaf s which is an S-node. Because s is a leaf, H_s contains exactly one virtual edge. Since H_s is a cycle of length at least 3, there is at least one degree-2 vertex in G, a contradiction.

(3) Let $P = a_0 \cdots a_m$ be a maximum length path in T'_G. By maximality, P is a leaf-to-leaf path in T'_G, a_i is an R-node for even i and an O-node for odd i, and m is even.

For $i \in [m-1]$, we let x_i and y_i be the ends of the virtual edge in $E(H_{a_i}) \cap E(H_{a_{i+1}})$. Since H_{a_i} is 2-connected, exchanging x_i and y_i if necessary we may assume that for each $i \in [m-1]$, H_{a_i} contains an x_{i-1}–x_i path P_i and a y_{i-1}–y_i path Q_i such that P_i and Q_i are vertex-disjoint.

Let $i \in [m-1]$ with i even. Let us emphasize that the vertices $x_{i-1}, x_i, y_{i-1}, y_i$ are not necessarily all distinct. We call a K_4-model in H_{a_i} *good* if the intersections of the four vertex images with these vertices fall in one of the following cases:

- $\{x_{i-1}\}, \{x_i\}, \{y_{i-1}\}, \{y_i\}$, or
- $\{x_{i-1}, x_i\}, \{y_{i-1}\}, \{y_i\}, \emptyset$ with $x_{i-1} \neq x_i$, or

- $\{x_i\}, \{y_{i-1}\}, \{y_i\}, \emptyset$ with $x_{i-1} = x_i$, or
- $\{x_{i-1}\}, \{x_i\}, \{y_{i-1}, y_i\}, \emptyset$ with $y_{i-1} \neq y_i$, or
- $\{x_{i-1}\}, \{x_i\}, \{y_i\}, \emptyset$ with $y_{i-1} = y_i$.

We claim that H_{a_i} has a good K_4-model for each even $i \in [m-1]$. To see this, let $C_i = P_i + Q_i + x_{i-1}y_{i-1} + x_iy_i$. First suppose $V(C_i) = V(H_{a_i})$. Since H_{a_i} is 3-connected, there is an edge $e \in E(H_{a_i})$ distinct from $x_{i-1}y_{i-1}$ and x_iy_i between $V(P_i)$ and $V(Q_i)$, and another edge f such that $C_i \cup \{e, f\}$ is a subdivision of K_4. Then $C_i + e + f$ contains a good K_4-model. Assume now that $V(C_i) \subsetneq V(H_{a_i})$. It follows that there is a component of $H_{a_i} - V(C_i)$ that sends edges to three vertices of C_i which are neither all in $V(P_i)$ nor all in $V(Q_i)$; otherwise $H_{a_i} - \{x_{i-1}, x_i\}$ or $H_{a_i} - \{y_{i-1}, y_i\}$ would be disconnected. Thus, H_{a_i} has a good K_4-model whose vertex images are a single component of $H_{a_i} - V(C_i)$ and three disjoint connected subgraphs of C_i.

We say that a good K_4-model in H_{a_i} is *type*-0 if x_{i-1}, x_i, y_{i-1}, and y_i are in distinct vertex images, *type*-1 if x_{i-1} and x_i are in the same vertex image, and *type*-2 if y_{i-1} and y_i are in the same vertex image. We pick a good K_4-model in each even $i \in [m-1]$. Since $m \geqslant 6k$, at least k of these good K_4-models are of the same type, say type-t for some $t \in \{0, 1, 2\}$.

We obtain the required minor of G as follows. First, for each even $i \in [m-1]$ such that H_{a_i} contains a type-t good K_4-model, we contract the vertex images of the K_4-model and delete the vertices not belonging to any vertex image. Second, for each index $i \in [m-1]$ not yet considered, we contract the edges in $E(P_i) \cup E(Q_i)$ and delete all other vertices of H_{a_i}. Note that this second step has the effect of 2-summing the type-t good K_4-models. Therefore, we obtain a P_k minor in G, if $t = 0$, and a F_k minor in G in the other two cases. □

3.5.2 Extending flat sets in 2-connected graphs

We now develop some more tools to handle 2-separations in graphs. Assume that $G = G_1 \oplus_e G_2$ with $e = vw$. The goal is to improve the bounds for $f_\infty(G)$ given in Lemma 3.12. Recall that the proof of Lemma 3.12 relies on the fact that it is possible to merge a flat set F_1 of (G_1, d_1) and a flat set F_2 of (G_2, d_2) into one flat set $F_1 \cup F_2$ of (G, d) whenever $(v, w) \in F_1 \cap F_2$.

Here is another proof of this fact. Let (D, l), (D_1, l_1) and (D_2, l_2) denote the weighted digraphs obtained by bidirecting (G, d), (G_1, d_1) and (G_2, d_2) respectively. For $i \in [2]$, consider a potential p_i on (D_i, l_i) such that $p_i(x) - p_i(y) = d(xy)$ for all $(x, y) \in F_i$. Since $(v, w) \in F_1 \cap F_2$, we have $p_1(v) - p_1(w) = p_2(v) - p_2(w) = d(vw)$. Hence, it is possible to shift one of the potentials in order to satisfy $p_1(v) = p_2(v)$ and $p_1(w) = p_2(w)$. The potential $p_1 \cup p_2 : V(G) \to \mathbb{R}$ on (D, l) such that $(p_1 \cup p_2)(u) = p_i(u)$ if $u \in V(G_i)$ for $i \in [2]$ witnesses that $F_1 \cup F_2$ is a flat set.

Suppose now that the flat sets F_1, F_2 of (G_1, d_1) and (G_2, d_2) are such that $(v, w) \in F_1$ but $(v, w), (w, v) \notin F_2$. The previous idea does not work anymore since we could have $|p_2(v) - p_2(w)| < d(vw)$. Hence, we can no longer combine the potentials p_1 and p_2. However, there possibly exists a potential p'_1 for $F_1 \setminus \{(v, w)\}$ such that $p'_1(v) - p'_1(w) = p_2(v) - p_2(w)$. In that case, $p'_1 \cup p_2$ is a potential for $(F_1 \cup F_2) \setminus \{(v, w)\}$ on (D, l). It follows that in this case $(F_1 \cup F_2) \setminus \{(v, w)\}$ is a flat set.

We now introduce the notion of *compressible edges*, which are edges for which we can apply the idea of the previous paragraph. In this context, it is helpful to switch from directed notions to undirected notions. We call a set F of edges of G *flattenable* (in (G, d)) if some orientation of F is a flat set in (G, d), that is, if there exists a potential p on (D, l) such that $|p(v) - p(w)| = d(vw)$ for all $vw \in F$. Let $F \subseteq E(G)$ be flattenable in (G, d). An edge subset $\Gamma \subseteq F$ is said to be *compressible* in F if for all $\lambda \in [0, 1]^{\Gamma}$ there exists a potential p on (D, l) such that $|p(v) - p(w)| = \lambda(vw) \cdot d(vw)$ for all $vw \in \Gamma$ and $|p(v) - p(w)| = d(vw)$ for all $vw \in F \setminus \Gamma$. We define a *frame* in (G, d) as a pair (Γ, F) where $\Gamma \subseteq F \subseteq E(G)$, F is flattenable in (G, d) and Γ is compressible in F.

Notice that subsets of flattenable sets are flattenable, and that $f_\infty(G)$ is the least integer k such that for every distance function d the edges of the metric graph (G, d) can be partitioned into k flattenable sets.

The next lemma follows directly from the formal definition of compressible edges.

Lemma 3.24. *Let $G = G_1 \oplus_{vw} G_2$, and let d be a distance function on G. For $i \in [2]$, let d_i be the restriction of d to G_i and let (Γ_i, F_i) be a frame in (G_i, d_i).*

(i) If $vw \in (F_1 \setminus \Gamma_1) \cap (F_2 \setminus \Gamma_2)$ then $(\Gamma_1 \cup \Gamma_2, F_1 \cup F_2)$ is a frame in (G, d).

(ii) If $vw \in \Gamma_1 \cup \Gamma_2$ then $((\Gamma_1 \cup \Gamma_2) \setminus \{vw\}, (F_1 \cup F_2) \setminus \{vw\})$ is a frame in

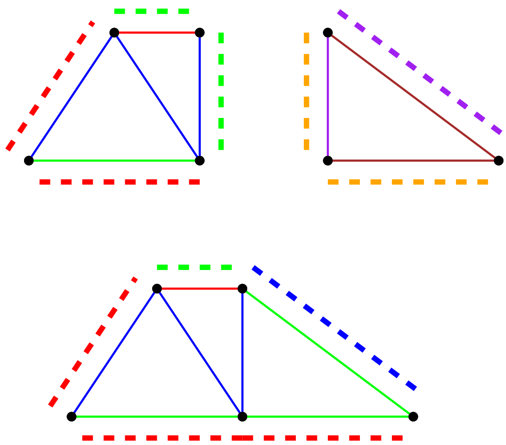

Figure 3.10. Illustration of the proof of Lemma 3.25: G is a 2-sum of $G' = K_4 - e$ and K_3. Each color defines a frame (Γ, F) in the corresponding graph. Edges of $F \setminus \Gamma$ are straight and edges of Γ are fat dashed. The distance function is defined by taking the corresponding Euclidean distance in the figure.

(G, d).

We will now use this lemma to improve some bounds given by Lemma 3.12. For simplicity, we call *gluing* the 2-sum operation where the edge involved in the 2-sum is kept. Let H be a graph obtained by gluing graphs $G_1, \ldots, G_m$ on distinct edges of a graph G. That is, there are distinct edges $e_1, \ldots, e_m$ such that $H = G \oplus_{e_1} G_1 \cdots \oplus_{e_m} G_m$. The bound obtained by applying Lemma 3.12 is $f_\infty(H) \leqslant f_\infty(G) + \sum_{i\in[m]} (f_\infty(G_i) - 1)$. We provide better bounds in the following cases. First, when G is a 2-connected outerplanar graph and all G_i are glued on edges of its outer cycle. Second, when G is a 2-connected treewidth-2 graph and H has no S_k minor.

Lemma 3.25. *Let G be a 2-connected outerplanar graph drawn in the plane with outer cycle C. Let H be obtained from G by gluing graphs $G_1, \ldots, G_m$ on distinct edges of C. Let $M = \max_{i\in[m]} f_\infty(G_i)$. Then $f_\infty(H) \leqslant 3M$.*

Proof. We will show that G satisfies the following property:

($\star$) *For every distance function d on G, there exist three frames (Γ_j, F_j), $j \in [3]$, in (G, d) such that each edge of G is in at least one flattenable set F_j, and each edge of its outer cycle C is in* *exactly two flattenable sets F_j and in exactly one compressible set Γ_j.*

For $i \in [m]$, let $\{v_i, w_i\} = V(G_i) \cap V(G)$. Thus, $v_i w_i$ is an edge of C. Without loss of generality, we may assume that $v_i w_i$ is an edge of H.

Now let d be some distance function on H. We will slightly abuse notation and let d also denote the restriction of this distance function to G. For $i \in [m]$, let d_i denote the restriction of d to G_i.

Assuming ($\star$), we can find three frames (Γ_j, F_j), $j \in [3]$, in (G, d) as above. For each $i \in [m]$, let $F_1^i, \ldots, F_M^i$ be a partition of the edges of (G_i, d_i) into flattenable set. By Lemma 3.24, for every $j \in [3]$ and $k \in [M]$,

$$\left(F_j \cup \bigcup_{i \in I_j} F_k^i\right) \setminus \{v_i w_i \mid i \in I_j\}$$

is a flattenable set in (H, d), where $I_j = \{i \in [m] \mid v_i w_i \in \Gamma_j\}$. These $3M$ flattenable sets cover the edges of (H, d), which implies $f_\infty(G) \leqslant 3M$.

To prove the lemma, it remains to show that the claimed frames (F_j, Γ_j), $j \in [3]$ exist in (G, d). We can assume that all inner faces of the drawing of G are triangular faces (if not, add extra edges). We show the result by induction on the number of vertices.

The base case is given by $G = K_3$. Let $V(K_3) = \{v_1, v_2, v_3\}$. Without loss of generality, we can assume $d(v_1 v_2) \leqslant d(v_1 v_3) \leqslant d(v_2 v_3)$. It is easy to show that $(\Gamma_1, F_1) = (\{v_1 v_2, v_1 v_3\}, \{v_1 v_2, v_1 v_3\})$, $(\Gamma_2, F_2) = (\{v_2 v_3\}, \{v_2 v_1, v_2 v_3\})$, and $(\Gamma_3, F_3) = (\emptyset, \{v_3 v_1, v_3 v_2\})$ are frames in (G, d). For instance, one can use Lemma 3.6 to see that each F_j is flattenable, and a direct verification to see that each Γ_j is compressible in F_j. Thus K_3 satisfies ($\star$).

Now for the inductive case, suppose that G has at least four vertices. Let v be a degree-2 vertex of G (which exists since G is outerplanar and 2-connected), and consider the graph $G' = G - v$. Let v_1, v_2 be the two neighbors of v in G, with $d(vv_1) \geqslant d(vv_2)$. Let C' be the cycle obtained from the outer cycle C in G by shortcutting the path $v_1 v v_2$ to $v_1 v_2$.

By induction, ($\star$) holds for G'. Let (Γ'_j, F'_j), $j \in [3]$ denote the corresponding frames. Consider three frames (Γ''_j, F''_j), $j \in [3]$ for the triangle $v v_1 v_2 v$, as described in the base case of the induction.

By permuting the indices if necessary, we may assume that $v_1 v_2$ is in $(F'_1 \setminus \Gamma'_1) \cap (F''_1 \setminus \Gamma''_1)$, Γ'_2 and Γ''_3. By Lemma 3.24, $(\Gamma_1, F_1) = (\Gamma'_1 \cup \Gamma''_1, F'_1 \cup F''_1)$

and, for $j \in \{2,3\}$, $(\Gamma_j, F_j) = ((\Gamma'_j \cup \Gamma''_j) \setminus \{v_1v_2\}, (F'_j \cup F''_j) \setminus \{v_1v_2\})$ are all frames in (G, d). See Figure 3.10 for an illustration. It is straightforward to check that these frames satisfy the required condition for G. □

3.5.3 Handling several 2-cutsets simultaneously

Before proceeding, we require the following easy lemma. Let $K_4 - e$ be the graph obtained from K_4 by deleting an edge.

Lemma 3.26 ([42]). *Let G be a 2-connected graph with distinct vertices u and v such that $\deg_G(w) \geqslant 3$ for all $w \in V(G) \setminus \{u, v\}$. Then G has a $K_4 - e$ minor where u and v are contracted to the ends of e.*

Let G be a graph together with a subset of $E(G)$ called *glued edges*. We say that G has a *k-glumpkin* minor if G contains k glued edges in parallel as a minor, that is, if there is a way of choosing a connected subgraph H of G containing at least k glued edges, and of contracting all but k edges of H in such a way that the resulting minor consists of k parallel glued edges. A k-glumpkin minor is *rooted* at a glued edge r if it contains r. If H is obtained by gluing graphs $G_1, \dots, G_m$ on distinct edges of G, an edge $e \in E(G)$ is a *glued edge* if $e \in E(G) \cap E(G_i)$ for some $i \in [m]$. The parameter we are really interested in is the largest S_k minor in H. However, the next lemma relates S_k minors in H to k-glumpkin minors in G.

Lemma 3.27. *Let H be obtained by gluing 2-connected graphs $G_1, \dots, G_m$ on distinct edges of a graph G such that H has minimum degree at least 3. If G has a k-glumpkin minor, then H has an S_k-minor.*

Proof. Let u_iv_i be the glued edge of G_i. Since H has minimum degree at least 3, $\deg_{G_i}(w) \geqslant 3$ for all $w \in V(G_i) \setminus \{u_i, v_i\}$. By Lemma 3.26, G_i has a K_4 minor containing the glued edge u_iv_i, for all $i \in [m]$. Therefore, since G has a k-glumpkin minor, H has an S_k-minor. □

Lemma 3.28. *For all $k, M \in \mathbb{N}$, let $g_{3.28}(k, M) = 3^k M$. Let H be a graph obtained from a 2-connected outerplanar graph G by gluing 2-connected graphs $G_1, \dots, G_m$ on distinct edges of G. Let C be the outercycle of G and let $M = \max_{i \in [m]} f_\infty(G_i)$. If there exists a glued edge $r \in E(C)$ such that G does not contain a k-glumpkin minor rooted at r, then $f_\infty(H) \leqslant g_{3.28}(k, M)$.*

Proof. We proceed by induction on k. The case $k = 1$ is vacuous. If $k = 2$, then by 2-connectivity, r is the only glued edge of G. Since G is outerplanar, $f_\infty(G) \leqslant 2$ and so by Lemma 3.12, $f_\infty(H) \leqslant M + 1 \leqslant g_{3.28}(2, M)$. Therefore, we may assume $k \geqslant 3$. A subpath of $C - r$ is *good* if its ends are connected by a glued edge. Let $P_1, \ldots P_p$ be the maximal (under inclusion) good subpaths of $C - r$. Since G is outerplanar, P_i and P_j are internally-disjoint for $i \neq j$. By maximality, every glued edge has both of its ends on some P_i.

Let G'_i be the subgraph of G induced by $V(P_i)$. Let e_i be the glued edge connecting the ends of P_i. Since G does not contain a k-glumpkin minor rooted at r, G'_i does not contain a $(k-1)$-glumpkin minor rooted at e_i. Let H_i be the subgraph of H induced by G'_i and all the graphs G_j that are glued to some edge of G'_i. By induction, $f_\infty(H_i) \leqslant 3^{k-1}M$ for all $i \in [p]$. Let C' be the cycle obtained from C by replacing P_i with e_i for each $i \in [p]$. Let G' be the subgraph of G induced by the vertices of C'. Notice that G' is a 2-connected outerplanar graph with outer cycle C', and H can be obtained from G' by gluing the graphs H_i on edges of C'. By Lemma 3.25,

$$f_\infty(H) \leqslant 3 \cdot \max_{i \in [p]} f_\infty(H_i) \leqslant 3 \cdot 3^{k-1}M = g_{3.28}(k, M). \qquad \square$$

We now generalize Lemma 3.28 to 2-connected treewidth-2 graphs.

Lemma 3.29. *For all $k, M \in \mathbb{N}$, let $g_{3.29}(k, M) = 3^{k^2}M$. Let G be a 2-connected treewidth-2 graph and let H be obtained by gluing 2-connected graphs $G_1, \ldots, G_m$ on distinct edges of G. Let $M = \max_{i \in [m]} f_\infty(G_i)$. If for some glued edge r, G does not contain a k-glumpkin minor rooted at r, then $f_\infty(H) \leqslant g_{3.29}(k, M)$.*

Proof. We proceed by lexicographic induction on $(k, |V(H)|)$. Let r be a glued edge such that G does not contain a k-glumpkin minor rooted at r.

The case $k = 1$ is vacuous. Suppose $k = 2$. Since G is 2-connected and does not have a 2-glumpkin minor rooted at r, edge r must be the only glued edge of G. Since G is 2-connected and has treewidth 2, $f_\infty(G) \leqslant 2$. By Lemma 3.12, $f_\infty(H) \leqslant M + 1 \leqslant g_{3.29}(2, M)$. Therefore, we may assume $k \geqslant 3$. If $\deg_H(w) = 2$ for some vertex $w \in V(H)$, then we can suppress w by Lemma 3.11 and apply induction. Therefore, we may assume H has minimum degree at least 3.

Since G is 2-connected, there is a cycle in G containing r. Let C be a longest cycle in G such that $r \in E(C)$. Let $\mathcal{E}$ be an ear decomposition of G beginning with C. (See for instance [31] for background about ear decompositions.) The *ear-decomposition tree* $T(\mathcal{E})$ of $\mathcal{E}$ is the rooted tree, whose vertices are the ears in $\mathcal{E}$, defined recursively as follows. The root of $T(\mathcal{E})$ is C. The parent of an ear P is the closest ear Q to C (in $T(\mathcal{E})$) such that both ends of P are on Q. (Such an ear Q is guaranteed to exist since G has treewidth 2 and is 2-connected.)

Let $P_1, \dots, P_\ell$ be the set of C-ears of $\mathcal{E}$. Let $T_1, \dots, T_\ell$ be the subtrees of $T(\mathcal{E})$ rooted at $P_1, \dots, P_\ell$, respectively. For each $i \in [\ell]$, let x_i and y_i be the ends of P_i on C. Let R_i be the x_i–y_i path in C containing r and let S_i be the other x_i–y_i path in C. Notice that $|E(S_i)| \geqslant |E(P_i)|$, by maximality of C. If P_i is an edge, then since G is simple, $|E(S_i)| \geqslant 2$. Otherwise, $|E(S_i)| \geqslant |E(P_i)| \geqslant 2$. Therefore, for all $i \in [\ell]$, $|E(S_i)| \geqslant 2$.

We claim that for all $i \in [\ell]$, $V(S_i)$ contains the ends of a glued edge. Suppose not. Among all S_i such that $V(S_i)$ does not contain the ends of a glued edge, choose S_j so that S_j is inclusion-wise minimal. Since G has treewidth 2 and is 2-connected, for all $i \neq j$, $S_i \subseteq S_j$, $S_j \subseteq S_i$, or S_i and S_j are internally-disjoint. By the minimality of S_j, each internal vertex of S_j has degree 2 in H. However, this contradicts that H has minimum degree at least 3.

For each $i \in [\ell]$, let G'_i be the union of all ears in T_i together with the edge $e_i = x_i y_i$, which we declare to be glued. Since $V(S_i)$ contains the ends of a glued edge and R_i contains r, the graph G'_i does not contain a $(k-1)$-glumpkin minor rooted at e_i; otherwise, G contains a k-glumpkin minor rooted at r. Note that each G'_i contains at least one glued edge other than e_i since H has minimum degree at least 3. Let H_i be the graph obtained from G'_i by gluing all G_j such that the glued edge of G_j belongs to G'_i. By induction, $f_\infty(H_i) \leqslant g_{3.29}(k-1, M)$, for all $i \in [\ell]$. Let $e_{i+1}, \dots, e_L$ be the glued edges in $E(C)$.

Observe that H is obtained by gluing graphs $H_1, \dots H_L$ onto edges of an outerplanar graph G' with outercycle C, where $M' = \max_{i \in [L]} f_\infty(H_i) = \max\{M, g_{3.29}(k-1, M)\} = g_{3.29}(k-1, M)$. Since G does not contain a k-glumpkin minor rooted at r, neither does G'. Applying Lemma 3.28 to G' gives

$$f_\infty(H) \leqslant g_{3.28}(k, g_{3.29}(k-1, M)) = 3^k(3^{(k-1)^2} M) \leqslant g_{3.29}(k, M). \qquad \square$$

Lemma 3.29 yields the following corollary.

Lemma 3.30. *For all $k, M \in \mathbb{N}$, let $g_{3.30}(k, M) = 3^{k^2} M$. Let G be a 2-connected treewidth-2 graph and let H be obtained by gluing 2-connected graphs $G_1, \ldots, G_m$ on distinct edges of G. If H does not contain an S_k minor and $M = \max_{i\in[m]} f_\infty(G_i)$, then $f_\infty(H) \leqslant g_{3.30}(k, M)$.*

Proof. We proceed by induction on $|V(H)|$. If $\deg_H(w) = 2$ for some $w \in V(H)$, then by Lemma 3.11, we can suppress w and apply induction. Since H does not contain an S_k minor, G does not contain a k-glumpkin minor, by Lemma 3.27. In particular, for each glued edge r, G does not contain a k-glumpkin minor rooted at r. By Lemma 3.29, $f_\infty(H) \leqslant g_{3.29}(k, M) = g_{3.30}(k, M)$. □

The following is the main result of this section.

Lemma 3.31. *Suppose there exist computable functions $g_{3.47} : \mathbb{N} \to \mathbb{R}$ and $g_{3.48} : \mathbb{N} \times \mathbb{N} \to \mathbb{R}$ satisfying the two following conditions.*

1. *$f_\infty(G) \leqslant g_{3.47}(k)$ for every 3-connected graph G not containing a $\mathcal{U}^k_\infty$ minor.*
2. *$f_\infty(H) \leqslant g_{3.48}(k, M)$ for every graph H containing no $\mathcal{U}^k_\infty$ minor, obtained by gluing 2-connected graphs $G_1, \ldots, G_m$ on distinct edges of a 3-connected graph G_0, where $M = \max_{i\in[m]} f_\infty(G_i)$.*

Then there exists a computable function $g_{1.3} : \mathbb{N} \to \mathbb{R}$ such that $f_\infty(G) \leqslant g_{1.3}(k)$ for all graphs G without a $\mathcal{U}^k_\infty$ minor.

Proof. We define $g_{1.3}(k)$ as follows. For all $k, M \in \mathbb{N}$, let $\alpha(k, M)$ be the maximum of $g_{3.30}(k, M)$ and $g_{3.48}(k, M)$. Define $\gamma_0(k) = g_{3.47}(k)$. For all $i, k \in \mathbb{N}$ recursively define $\gamma_i(k) = \alpha(k, \gamma_{i-1}(k))$. Finally, let $g_{1.3}(k) = \gamma_{6k}(k)$.

Let G be a graph without a $\mathcal{U}^k_\infty$ minor. By Lemma 3.12, we may assume that G is 2-connected. By Lemma 3.11, we can assume that G has no degree-2 vertices. Let T_G be the SPQR tree of G and let $T = T'_G$ be the contracted SPQR tree, see Lemma 3.23.

Pick an arbitrary root node r in T. For each node b of T, we denote by T_b the subtree of T rooted at b and by H_b the minor of G corresponding to that subtree. Note that $G = H_r$. By Lemma 3.23, every leaf of T is an R-node. Hence, each leaf u of T corresponds to a 3-connected minor H_u of G. By our first assumption, $f_\infty(H_u) \leqslant g_{3.47}(k) = \gamma_0(k)$. Let a be some inner node of T and let $a_1, \dots, a_\ell$ denote its children. Let $M_a = \max_{j \in [\ell]} f_\infty(H_{a_j})$. If a is an O-node, then by Lemma 3.30, $f_\infty(H_a) \leqslant g_{3.30}(k, M_a)$. If a is a R-node, then $f_\infty(H_a) \leqslant g_{3.48}(k, M_a)$ by our second assumption. In either case, $f_\infty(H_a) \leqslant \alpha(k, M_a)$. It follows that if i is the maximum length of an a to leaf path of T, then $f_\infty(H_a) \leqslant \gamma_i(k)$. By Lemma 3.23, the height of T is at most $6k$. Therefore, $f_\infty(G) = f_\infty(H_r) \leqslant \gamma_{6k}(k) = g_{1.3}(k)$. □

We will establish the existence of $g_{3.47}$ and $g_{3.48}$ in Lemmas 3.47 and 3.48, respectively. Lemmas 3.31, 3.47, and 3.48 and the results from Section 3.4 together establish Theorem 1.3, which we now restate:

Theorem 1.3. There exists a computable function $g_{1.3} : \mathbb{N} \to \mathbb{R}$ such that for every $k \in \mathbb{N}$, every graph G with $f_\infty(G) > g_{1.3}(k)$ contains a $\mathcal{U}^k_\infty$ minor. Moreover, every graph G that contains a $\mathcal{U}^k_\infty$ minor has $f_\infty(G) > k$.

Proof. For the first part of the theorem, by Lemmas 3.31, 3.47, and 3.48, there exists a computable function $g_{1.3} : \mathbb{N} \to \mathbb{R}$ such that $f_\infty(G) \leqslant g_{1.3}(k)$ for all graphs G without a $\mathcal{U}^k_\infty$ minor. Thus, every graph G satisfying $f_\infty(G) > g_{1.3}(k)$ contains a $\mathcal{U}^k_\infty$ minor.

For the second part of the theorem, it is shown in Section 3.4 that each of the four graphs G in $\mathcal{U}^k_\infty$ satisfies $f_\infty(G) > k$. Since $f_\infty(G)$ is monotone w.r.t. minors, it follows that $f_\infty(G) > k$ for every graph G containing a $\mathcal{U}^k_\infty$ minor. □

3.6 3-connected graphs

The results in this section are purely graph theoretical and may be of independent interest. In particular, we prove several lemmas which give sufficient conditions under which a graph contains some specific graphs as minors. We also introduce a reduction operation, called *fan-reduction*. The main result of the section is that if G is a 3-connected, fan-reduced graph having no $\mathcal{U}^k_\infty$ minor, then the vertex cover number of G, $\tau(G)$, is bounded by a function of k.

Before proceeding, we quickly review some graph theoretical terminology. Let A, B be subsets of vertices of a graph G. An *A–B path* is a path P in G such that the ends of P are in A and B respectively, and no internal vertex

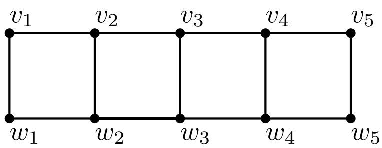

Figure 3.11. The ladder L_5.

of P is in $A \cup B$. If H is a subgraph of G then an *H-path* is a path P in G such that the ends of P are in H but no other vertex nor edge of P is in H.

The *n-ladder* L_n is the graph on $2n$ vertices with vertex set $V = \{v_i \mid i \in [n]\} \cup \{w_i \mid i \in [n]\}$ and edge set $E = \{v_i w_i \mid i \in [n]\} \cup \{v_i v_{i+1}, w_i w_{i+1} \mid i \in [n-1]\}$ (see Figure 3.11). By repeatedly suppressing degree-2 vertices, we can reduce L_n to the graph K_3. This implies that $f_\infty(\mathsf{L}_n) = 2$ for all $n \geqslant 2$ by Lemma 3.11.

Lemma 3.32. *For all $k \in \mathbb{N}$, let $g_{3.32}(k) = 12k^2 + 7k$. If G is a 3-connected graph containing a $g_{3.32}(k)$-ladder as a minor, then G contains N_k, P_k, or F_k as a minor.*

Proof. Since L_n has maximum degree 3, every graph with an L_n minor also contains an L_n subdivision. Let S be a subgraph of G isomorphic to a subdivision of L_n with $n = g_{3.32}(k)$. We say that the vertices of S that do not correspond to internal vertices of a subdivided edge are *branch vertices*. We name these branch vertices $\{v_i \mid i \in [n]\} \cup \{w_i \mid i \in [n]\}$ as in the definition of L_n given above. A *rung* is a path in S corresponding to an edge of L_n of the form $v_i w_i$, for some $i \in [n]$. We say that an S-path P *crosses a rung* R, if the ends of P are in different components of $S - V(R)$. A rung is *crossed* if it is crossed by some S-path, and is *uncrossed* otherwise.

If there exists an S-path in G that crosses at least $2k+1$ rungs, then G contains an N_k minor, and we are done. Hence, we may assume that each S-path crosses at most $2k$ rungs of S.

We say that the path in S from v_1 to v_n avoiding all w_i for $i \in [n]$ is the *upper path* of S. Similarly the *lower path* is the path in S from w_1 to w_n avoiding all vertices v_i for $i \in [n]$. For each $i \in \{2, \dots, n-1\}$, let S_ℓ^i and S_r^i be the components of $S - \{v_i, w_i\}$ that contain v_1 and v_n, respectively.

Suppose there are $8k+1$ uncrossed rungs $R_1, \dots, R_{8k+1}$. For each $i \in [8k+1]$, let $v_{i'}$ and $w_{i'}$ be the ends of R_i. We may assume that $i' < j'$ for all $i < j$. Since G is 3-connected, $G - \{v_{i'}, w_{i'}\}$ is connected. Therefore, there is a path

P in $G-\{v_{i'}, w_{i'}\}$ from $V(S_\ell^{i'})$ to $V(S_r^{i'})$. Since R_i is uncrossed, P must use an internal vertex of R_i. Thus, there exists a vertex $y_i \in V(R_i) \setminus \{v_{i'}, w_{i'}\}$ that is connected by an S-path P_i to some vertex $z_i \notin V(R_i)$.

By symmetry and pigeonhole, there is a subset I of size k of $\{2, 4, \ldots, 8k\}$ such that $z_i \in V(S_r^{i'})$ and z_i is not on the lower path of S, for all $i \in I$. Since R_i is uncrossed for all $i \in [8k+1]$ it follows that $z_i \in V(S_\ell^{(i+1)'}) \cup V(R_{i+1})$. For the same reason, P_i and P_j are vertex-disjoint for all distinct $i, j \in I$. Therefore, $S \cup \bigcup_{i \in I} P_i$ contains an F_k minor.

We may hence assume that S contains at most $8k$ uncrossed rungs. Thus, S contains at least $n - 8k = 12k^2 - k$ crossed rungs. Since $12k^2 - k = 1 + (4k+1)(3k-1)$, there is a subset J of $[n]$ of size $3k$ such that for all distinct $i, j \in J$, $|i - j| \geqslant 4k + 1$ and R_i is crossed. For each $i \in J$, let P_i be an S-path crossing R_i. Let ℓ_i and r_i be the ends of P_i in S_ℓ^i and S_r^i, respectively.

We say that P_i is of *type* v if ℓ_i and r_i are both on the upper path, *type* w if ℓ_i and r_i are both on the lower path, and *type* p otherwise. Since $|J| = 3k$, there is a subset J' of J of size k such that P_i is of the same type T for all $i \in J'$. Recall that each S-path crosses at most $2k$ rungs and $|i-j| \geqslant 4k+1$ for all distinct $i, j \in J'$. Therefore, if $i, j \in J'$ and $i < j$, then r_i is to the left of ℓ_j. Moreover, for the same reason, P_i and P_j are vertex-disjoint for all distinct $i, j \in J'$. Therefore, $S \cup \bigcup_{i \in J'} P_i$ contains an F_k minor if $\mathsf{T} \in \{v, w\}$ and $S \cup \bigcup_{i \in J'} P_i$ contains a P_k minor if $\mathsf{T} = p$. □

For each $k \in \mathbb{N}$, the *k-fan* is the graph consisting of a k-vertex path called its *outer path*, plus a universal vertex called its *center*. The edges connecting the center to the ends of the k-vertex path are called the *boundary edges* of the k-fan. A *fan* is a graph isomorphic to a k-fan for some k.

Let H be a fan, and assume that G has an H-model. We say that the H-model is *rooted at x, y* if x and y are contained in the vertex images of vertices a and b of H, respectively, and ab is a boundary edge of the fan.

Lemma 3.33. *For all $k, q \in \mathbb{N}$, let $g_{3.33}(k, q) = 3(8k^3)^q$. Let G be a graph and let $P = p_1 \cdots p_r$ be a path in G of length at least $g_{3.33}(k, q)$ such that $V(G) \setminus V(P)$ is a stable set. Then at least one of the following holds:*

1. *G has a k-fan minor;*

2. *there is a model of the q-fan in G rooted at p_2, p_{r-1} and avoiding p_1, p_r;*

3. *there are non-consecutive indices s, t with $1 < s < t < r$ such that $\{p_s, p_t\}$ separates in G the p_s–p_t subpath of P from the other vertices of P.*

Proof. The proof is by induction on q. For the base case $q = 1$, observe $g_{3.33}(k, 1) \geqslant 24$, for all $k \in \mathbb{N}$. Thus, it suffices to take p_2 and the p_3–p_{r-1} subpath of P as the two vertex images to obtain a model of the 1-fan rooted at p_2, p_{r-1} and avoiding p_1, p_r.

For the inductive step, assume $q > 1$. Let $S = V(G) \setminus V(P)$. We may assume that every vertex in S has degree at most $k - 1$ in G, since otherwise there is a k-fan minor in G. Note that $g_{3.33}(k, q) = 8k^3 \cdot g_{3.33}(k, q - 1)$. A *jump* is a pair (a, b) of indices $a, b \in [r]$ with $b \geqslant a + 2$ such that either $p_a p_b \in E(G)$ (*type 1*) or p_a and p_b have a common neighbor in S (*type 2*). For definiteness, if both conditions are satisfied then (a, b) is considered to be of type 1. To each jump (a, b) of type 2 we associate a corresponding *middle vertex* $w \in S$ adjacent to both a and b, that is chosen arbitrarily. A jump (a, b) is called an *outer jump* if $a = 1$ or $b = r$; otherwise, (a, b) is an *inner jump*. In what follows we will be mostly interested in inner jumps.

Case 1: There exists an inner jump (a, b) with $b - a \geqslant k \cdot g_{3.33}(k, q - 1)$. Let (a, b) be such a jump. If (a, b) is of type 2, we first modify it as follows. Let w be the middle vertex of (a, b). Since w has degree at most $k - 1$, it follows that there exists a jump (a', b') with $b' - a' \geqslant k \cdot g_{3.33}(k, q - 1)/(k - 2) \geqslant g_{3.33}(k, q - 1)$ such that w is adjacent to $p_{a'}$ and $p_{b'}$ but to no vertex lying strictly in between them on P. We rename (a', b') to (a, b).

Let G' be the minor of G obtained by contracting the p_1–p_a subpath of P into p_a and the p_b–p_r subpath of P into p_b. Let P' be the path obtained from P by performing these contractions. We regard p_a and p_b as the ends of P'. Note that $V(G') \setminus V(P')$ is a stable set in G'. Since P' has length $b - a \geqslant g_{3.33}(k, q - 1)$, by induction at least one of the following holds:

1. G' has a k-fan minor;
2. there is a model $\mathcal{M}'$ of the $(q - 1)$-fan in G' rooted at p_{a+1}, p_{b-1} and avoiding p_a, p_b;
3. there are non-consecutive indices s, t with $a < s < t < b$ such that $\{p_s, p_t\}$ separates in G' the p_s–p_t subpath of P' from the other vertices of P'.

In the first case, we are done since G' is a minor of G. In the second case, $\mathcal{M}'$ is also such a model in G since the two subpaths that were contracted in the definition of G' resulted in vertices p_a, p_b. By symmetry, we may assume that the vertex image V_0 corresponding to the center of the fan contains p_{a+1}.

Recall that $2 \leqslant a < b \leqslant r-1$, since (a, b) is an inner jump. Let L and R be the p_2–p_a and p_b–p_{r-1} subpaths of P, respectively. Let w be the middle vertex of (a, b) if (a, b) is type 2. Let $R' = R$ if R is type 1, and $R' = R \cup \{w\}$ if (a, b) is type 2. In either case, observe that L and R' are connected by an edge. By construction, $V(L) \cup V(R)$ is disjoint from all vertex images of $\mathcal{M}'$. Since w is not adjacent to any internal vertex of P', $\{w\}$ is also disjoint from all vertex images of $\mathcal{M}'$. Finally, the edges $p_a p_{a+1}$ and $p_{b-1} p_b$ connect $V(L)$ and $V(R)$ to the vertex images of $\mathcal{M}'$ containing p_{a+1} and p_{b-1}, respectively. Therefore, $(\mathcal{M}' \setminus \{V_0\}) \cup \{V_0 \cup L, R'\}$ is a model of the q-fan in G rooted at p_2, p_{r-1} and avoiding p_1, p_r, as desired.

It remains to consider the third case. Suppose s, t are non-consecutive indices with $a < s < t < b$ such that $\{p_s, p_t\}$ separates in G' the p_s–p_t subpath of P' from the other vertices of P'. Given how G' was obtained from G, this is also true in G. That is, $\{p_s, p_t\}$ separates in G the p_s–p_t subpath of P from the other vertices of P, as desired.

Case 2: $b - a < k \cdot g_{3.33}(k, q-1)$ for all inner jumps (a, b). Let us introduce one more definition. A *jump sequence* is a sequence $(a_1, b_1), \ldots, (a_\ell, b_\ell)$ of inner jumps with $\ell \geqslant 1$ satisfying $a_i < a_{i+1} < b_i < b_{i+1}$ for each $i \in [\ell-1]$, and $b_i \leqslant a_{i+2}$ for each $i \in [\ell-2]$. Its *length* is ℓ and its *spread* is $b_\ell - a_1$.

Case 2.1: There exists a jump sequence of spread at least $2k^2 \cdot g_{3.33}(k, q-1)$. Let $(a_1, b_1), \ldots, (a_\ell, b_\ell)$ be a jump sequence of spread at least $2k^2 \cdot g_{3.33}(k, q-1)$ and with ℓ *minimum*. For each $i \in [\ell]$, if (a_i, b_i) is of type 2 let $w_i \in S$ be the middle vertex of (a_i, b_i).

We claim that all middle vertices w_i defined above are distinct. Indeed, assume $w_i = w_j$ for some $i, j \in [\ell]$ with $i < j$. Then (a_i, b_j) is also an inner jump, and $(a_1, b_1), \ldots, (a_{i-1}, b_{i-1}), (a_i, b_j), (a_{j+1}, b_{j+1}), \ldots, (a_\ell, b_\ell)$ is a jump sequence, as the reader can easily check. But the latter jump sequence has length at most $\ell - 1$ and yet its spread is also $b_\ell - a_1$, contradicting our choice of the original jump sequence.

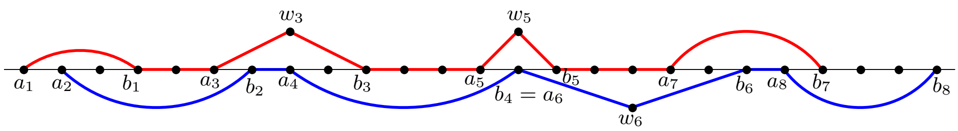

Figure 3.12. Illustration of a k-fan-model obtained from a jump sequence $(a_1, b_1), \ldots, (a_{2k}, b_{2k})$ for $k = 4$. The blue path is the vertex image for the center of the fan, and the red path corresponds to the outer path. Edges incident to the center of the fan map to the first edge of the subpath of P from a_{2i} to b_{2i-1}.

Since $b_i - a_i \leqslant k \cdot g_{3.33}(k, q-1)$ for each $i \in [\ell]$, we have

$$2k^2 \cdot g_{3.33}(k, q-1) \leqslant b_\ell - a_1 \leqslant \sum_{i \in [\ell]} (b_i - a_i) \leqslant \ell k \cdot g_{3.33}(k, q-1),$$

implying $\ell \geqslant 2k$. Now, one can obtain a k-fan-model using the jump sequence $(a_1, b_1), \ldots, (a_{2k}, b_{2k})$ as illustrated in Figure 3.12.

Case 2.2: All jump sequences have spread less than $2k^2 \cdot g_{3.33}(k, q-1)$. Let

$$\begin{aligned} M = &\{2, r-1\} \cup \{i \in [r] \mid (1, i) \text{ is an outer jump}\} \\ &\cup \{i \in [r] \mid (i, r) \text{ is an outer jump}\}. \end{aligned}$$

If there are k outer jumps of the form $(1, i)$ then G has a k-fan minor, and the same is true for those of the form (i, r). Thus we may assume that $|M| \leqslant 2k$. By the pigeonhole principle, there are two indices $i, j \in M$ with $i < j$ and $M \cap [i+1, j-1] = \emptyset$ such that

$$j - i \geqslant \frac{r-1}{|M|-1} \geqslant \frac{g_{3.33}(k, q)}{2k} = 4k^2 \cdot g_{3.33}(k, q-1).$$

If there exists an inner jump (a, b) with $a < i < b$, let $(a_1, b_1), \ldots, (a_\ell, b_\ell)$ be a jump sequence such that $a_1 < i < b_1$ and maximizing its spread, and let $s = b_\ell$. If no such jump exists, simply let $s = i$.

We claim that there is no inner jump (a, b) with $a < s < b$. This is obviously true if $s = i$, so assume $s \neq i$, and consider the corresponding jump sequence $(a_1, b_1), \ldots, (a_\ell, b_\ell)$ defined above. Arguing by contradiction, suppose that there is an inner jump (a, b) with $a < s < b$. If $a \leqslant a_1$ then (a, b) is a jump sequence with $a < i < b$ and spread $b - a > b_\ell - a_1$, contradicting our choice of the jump sequence. If $a_1 < a$ then letting $\ell' \in [\ell]$ be the smallest index such that $a < b_{\ell'}$ (which is well defined since $a < b_\ell$), we deduce that $(a_1, b_1), \ldots, (a_{\ell'}, b_{\ell'}), (a, b)$ is a jump sequence with $a_1 < i < b_1$ and of spread $b - a_1 > b_\ell - a_1$, again a contradiction. Hence, no inner jump (a, b) with $a < s < b$ exists, as claimed.

Next, if there exists an inner jump (a, b) with $a < j < b$, let $(a'_1, b'_1), \ldots, (a'_{\ell'}, b'_{\ell'})$ be a jump sequence such that $a'_{\ell'} < j < b'_{\ell'}$ and maximizing its spread, and let $t = a'_1$. If no such jump exists, simply let $t = j$. By a symmetric argument, there is no inner jump (a, b) with $a < t < b$.

Recall that every jump sequence has spread strictly less than $2k^2 \cdot g_{3.33}(k, q-1)$. Thus, $s - i \leqslant 2k^2 \cdot g_{3.33}(k, q-1) - 1$ and $j - t \leqslant 2k^2 \cdot g_{3.33}(k, q-1) - 1$. It follows that

$$t - s \geqslant j - i - 4k^2 \cdot g_{3.33}(k, q-1) + 2 \geqslant 2.$$

In other words, $[s+1, t-1]$ is not empty. Since $[s+1, t-1] \subseteq [i+1, j-1]$ and $M \cap [i+1, j-1] = \emptyset$, there is no outer jump $(1, b)$ with $b \in [s+1, t-1]$ and there is no outer jump (a, r) with $a \in [s+1, t-1]$. Since we already established that there is no inner jump (a, b) with $a < s < b$ or $a < t < b$, we deduce that the two indices s, t satisfy the third outcome of the claim. That is, s and t are non-consecutive indices with $1 < s < t < r$ such that $\{p_s, p_t\}$ separates in G the p_s–p_t subpath of P from the other vertices of P. □

As an easy corollary of Lemma 3.33, we obtain the following strengthening of Lemma 4.7 in [52].[1]

Lemma 3.34. *For all $k \in \mathbb{N}$, let $g_{3.34}(k) = 3(8k^3)^k$. Let G be a graph with no k-fan minor. Let P be a path in G of length at least $g_{3.34}(k)$ such that $V(G) \setminus V(P)$ is a stable set. Then there exist two non-consecutive internal vertices u, v of P such that $\{u, v\}$ separates in G the u–v subpath of P from the other vertices of P.*

Proof. Note that $g_{3.34}(k) = g_{3.33}(k, k)$. The lemma follows by applying Lemma 3.33 to G and P, and noting that the first two outcomes of Lemma 3.33 are impossible since G has no k-fan minor. □

[1] The latter lemma works under the assumption that G does not have the graph consisting of two vertices linked by k parallel edges as a minor, which is more restrictive than just forbidding a k-fan minor. Nevertheless, the two proofs are based on a similar strategy.

Next, we introduce two lemmas about 3-connected graphs containing subdivisions of large fans as subgraphs. Given a graph G, we say that F is a *fan subdivision in* G if F is a subgraph of G isomorphic to a subdivision of a fan. Moreover, we say that F is a *maximal* fan subdivision in G if F is maximal with respect to subgraph inclusion. That is, for every fan subdivision F' in G such that $F \subseteq F' \subseteq G$, we have $F = F'$.

Lemma 3.35. *For all $k \in \mathbb{N}$, let $g_{3.35}(k) = 8k^4 + 4k^3 + 10k$. If G is a 3-connected graph and F is a maximal fan subdivision in G such that at least $g_{3.35}(k)$ of the edges of the fan are subdivided, then G has an L_k, S_k or F_k minor.*

Proof. Let F^* denote the m-fan such that F is a subdivision of F^*, where v_0 is the center of F^* and $v_1 \cdots v_m$ is the outer path of F^*.

In the following we consider the graph H obtained from G by performing the following two operations. First, we contract each component of $G - V(F)$ into a vertex. Second, for each edge e of F^* that is subdivided at least once in F, we contract the corresponding path P of F into a 2-edge path, that is, we leave just one subdivision vertex. We call this subdivision vertex v_i^1 if $e = v_0v_i$ for some $i \in [m]$, and v_i^2 if $e = v_iv_{i+1}$ for some $i \in [m-1]$.

Hence, each vertex of H is of the form v_i, v_i^1, v_i^2, or results from the contraction of a component of $G - V(F)$. We denote by F' the fan subdivision in H that is the image of F, that is, which is obtained from F by the above contractions. Observe that F' is a *maximal* fan subdivision in H. Indeed, if some fan subdivision in H strictly contained F' then that fan subdivision could be mapped to a fan subdivision in G strictly containing F, contradicting the maximality of F.

We will establish the following key property of H:

> ($\star$) *If u_i is a vertex of H of the form v_i^1 or v_i^2, then there is an F'-path P_i in H of length at most 2 connecting u_i to another vertex u_i' of F' distinct from its two neighbors in F' and from v_0.*

Suppose ($\star$) does not hold for some v_i^1. Then $\{v_0, v_i\}$ is a size-2 cutset of H separating v_i^1 from every vertex v_j with $j \notin \{0, i\}$ (here we implicitly use that $m \geqslant 2$, since F^* has at least $g_{3.35}(k) \geqslant 2$ edges). By the construction of H, the set $\{v_0, v_i\}$ is also a cutset of G separating v_i^1 from every vertex v_j with $j \notin \{0, i\}$. However, this contradicts the fact that G is 3-connected.

The remaining case is if ($\star$) does not hold for some v_i^2. Here we first observe that v_i^2 is not adjacent to v_0 in H, because otherwise this would contradict the maximality of F' in H. For the same reason, there is no length-2 path from v_i^2 to v_0 in H going through a vertex in $V(H) \setminus V(F')$. Using these two observations, we can proceed similarly as in the proof for v_i^1. This concludes the proof of ($\star$).

Now, we color each edge of F' blue, and each remaining edge of H red. Consider the graph H^* obtained from H as follows. Every edge of the form $v_i^1 v_i$ is contracted to the vertex v_i, every edge of the form $v_i^2 v_i$ is contracted to the vertex v_i, and finally, for every vertex $w \in V(H) \setminus V(F')$, we select a neighbor of w distinct from v_0 in the current graph (which exists) and contract the corresponding edge. Finally, we delete all red edges incident to v_0. Loops and parallel edges resulting from edge contractions are deleted as always, but if a red edge parallel to a blue edge is created, we keep the blue edge and delete the red edge. Thus, the blue subgraph of H^* is exactly the fan F^*. Let R^* denote the red subgraph of H^*. We regard R^* as a spanning subgraph of H^*, and thus R^* may have isolated vertices.

If R^* has a vertex of degree at least $2k+1$, then that vertex is not v_0 (since v_0 is not incident to any red edge), and it is then easily seen that H^* has an S_k minor. Thus we may assume that the maximum degree of R^* is at most $2k$.

If R^* has a matching of size k^3, then by Pigeonhole and Erdős-Szekeres [38], R^* has a matching $M = \{v_{a_i} v_{b_i} : i \in [k]\}$ of size k that satisfies one of the following three conditions:

1. $a_1 < a_2 < \cdots < a_k < b_1 < b_2 < \cdots < b_k$, or
2. $a_1 < a_2 < \cdots < a_k < b_k < b_{k-1} < \cdots < b_1$, or
3. $a_1 < b_1 < a_2 < b_2 < \cdots < a_k < b_k$.

In the first two cases, we see that H^* has an L_k minor (obtained by combining M with the v_{a_1}–v_{a_k} and v_{b_1}– v_{b_k} subpaths of the outer path of H^*). In the third case, we see that H^* has an F_k minor. Hence we may assume that R^* has no matching of size k^3.

It follows that R^* has a vertex cover of size at most $2k^3$. However, since R^* has maximum degree at most $2k$, it follows in turn that at most $2k^3(2k+1)$ vertices of R^* have non-zero degrees in R^*.

Recall that v_i^1 and v_i^2 (if they exist) are the only 2 vertices of F' that are contracted to v_i in F^*. Since F^* has at least $g_{3.35}(k)$ edges that are subdivided in F' and $g_{3.35}(k)/2 - 2k^3(2k+1) = 5k$, there exists $I \subseteq [m]$ with $|I| = k$ such that the following holds:

- there is a vertex u_i of the form v_i^1 or v_i^2 in H, for each $i \in I$;
- v_i has degree 0 in R^* for all $i \in I$, and
- $|i - j| \geqslant 5$ for all $i, j \in I$ with $i \neq j$.

Now, consider an index $i \in I$ and its associated subdivision vertex u_i in H. By ($\star$), there is an F'-path P_i in H of length at most 2 connecting u_i to another vertex u_i' of F' distinct from its two neighbors in F' and from v_0. The (one or two) edges of P_i are red and are not incident to v_0, and they disappeared in the edge contraction operations leading to the graph H^*. It follows that u_i' is very close to u_i in $F' - v_0$, namely u_i' must be one of v_{i-1}, v_{i+1}, or one of the subdivision vertices $v_{i-1}^1, v_{i+1}^1, v_{i-1}^2, v_i^2, v_{i+1}^2$ (if they exist).

Since the paths P_i and P_j are vertex disjoint for all $i, j \in I$ with $i \neq j$ (which follows from the fact that v_i and v_j have degree 0 in R^*), and since $|i - j| \geqslant 5$, combining F' with these k paths we can see that H contains an F_k minor. □

Let F be an m-fan with center v_0 and outer path $v_1 \cdots v_m$. Suppose that F is a subgraph of a graph G. We say that F is *reducible in* G if $m \geqslant 5$ and all vertices $v_2, \ldots, v_{m-1}$ have degree exactly 3 in G. The *F-reduction* of G is the minor of G obtained by contracting the edges of the path $v_3 \cdots v_{m-1}$. Thus, the resulting graph has $m - 4$ fewer vertices than G.

A reducible fan subgraph in G is said to be *maximal in* G if it is not a proper subgraph of any other reducible fan subgraph of G. Observe that if F_1 and F_2 are two distinct maximal reducible fan subgraphs of G then F_1 and F_2 are almost vertex disjoint in the following sense: F_2 contains none of the internal vertices of the outer path of F_1, and vice versa. We define the *fan-reduction* of G as the minor of G obtained by simultaneously performing all F-reductions for all maximal reducible fan subgraphs F of G. By the previous observation, this minor is well-defined. We say that G is *fan-reduced* if G does not contain a reducible fan subgraph. Observe that the fan-reduction of G is fan-reduced.

Lemma 3.36. *For all $k \in \mathbb{N}$, let $g_{3.36}(k) = 20k^5 + 14k^4 + 2k^3 + 5k$. If G is a 3-connected fan-reduced graph containing a $g_{3.36}(k)$-fan as a subgraph, then G contains an $\mathsf{S}_k, \mathsf{F}_k$ or L_k minor.*

Proof. Consider an m-fan subgraph F in G with center v_0, outer path $v_1 \cdots v_m$, and $m = g_{3.36}(k)$. Let H be obtained from G by contracting each component of $G - V(F)$ into a vertex. We color the edges of F blue and the remaining edges of H red as in the proof of Lemma 3.35, and define H^* in exactly the same way. The only difference here is that no edge of F needs to be contracted since F is already a fan. In the notation used in the proof of Lemma 3.35, here we have $F = F' = F^*$. Let R^* denote the red spanning subgraph of H^*.

If R^* has a vertex of degree at least $2k + 1$ or a matching of size k^3, then we find one of our target minors, exactly as in the proof of Lemma 3.35. Thus we may assume that this does not happen, implying that at most $2k^3(2k + 1)$ vertices of R^* have non-zero degrees in R^*.

Since $(m - 2k^3(2k + 1))/(2k^3(2k + 1) + 1) \geqslant 5k$ there is an index $i \in [m - 5k]$ such that none of $v_{i+1}, \ldots, v_{i+5k}$ is incident to a red edge in H^*. For each $\ell \in [k]$, there must be an index $j \in \{i + 5(\ell - 1) + 2, i + 5(\ell - 1) + 3, i + 5(\ell - 1) + 4\}$ such that v_j is incident to a red edge of H. Otherwise, $v_{i+5(\ell-1)+1}, \ldots, v_{i+5(\ell-1)+5}$ together with v_0 form a reducible fan in G. Since all red edges incident to v_j in H disappeared when constructing H^*, it follows that v_j is adjacent in H to a vertex $w_\ell \in V(H) \setminus V(F)$ such that the neighbors of w_ℓ in H are a subset of $\{v_0, v_{j-1}, v_j, v_{j+1}\}$. Furthermore, w_ℓ must be adjacent to at least three of these four vertices, since otherwise G would not be 3-connected. Now, combining F with the k vertices $w_1, \ldots, w_k$ we see that H contains an F_k minor. □

Combining the two previous lemmas, we obtain the following lemma.

Lemma 3.37. *For all $k \in \mathbb{N}$, let $g_{3.37}(k) = g_{3.36}(k)(g_{3.35}(k)+1)+g_{3.35}(k)$. If G is a 3-connected, fan-reduced graph containing a subdivision of a $g_{3.37}(k)$-fan as a subgraph, then G has an $\mathsf{S}_k, \mathsf{F}_k$ or L_k minor.*

Proof. Since G contains a $g_{3.37}(k)$-fan subdivision, G contains a maximal m-fan subdivision F with $m \geqslant g_{3.37}(k)$. If at least $g_{3.35}(k)$ edges of the m-fan are subdivided in F, then, by Lemma 3.35, G contains an $\mathsf{L}_k, \mathsf{S}_k$ or F_k minor. Otherwise, F contains an m'-fan as a subgraph with $m' \geqslant$

$(g_{3.37}(k) - g_{3.35}(k))/(g_{3.35}(k)+1) = g_{3.36}(k)$, and by Lemma 3.36, G contains an $\mathsf{L}_k, \mathsf{S}_k$ or F_k minor. □

The next lemma is standard, we include the proof nevertheless for completeness.

Lemma 3.38. *For all $k \in \mathbb{N}$, let $g_{3.38}(k) = k^{k^2+2}$. If G is a graph with a $g_{3.38}(k)$-fan minor, then G contains a subdivision of a k-fan as a subgraph, or G contains an L_k minor.*

Proof. Let G be a graph containing an m-fan F as minor with $m = g_{3.38}(k)$. Let v_0 be the center of F and $v_1 \cdots v_m$ be the outer path. Let $\{X_i \mid i \in \{0, 1, \ldots, m\}\}$ denote an F-model in G, with X_i denoting the vertex image of v_i.

For every edge $v_i v_j$ of F we choose vertices x_i^j, x_j^i of X_i, X_j, respectively, such that $x_i^j x_j^i \in E(G)$. Let T be a subtree of $G[X_0 \cup \{x_i^0 \mid i \in [m]\}]$ such that the leaves of T are exactly the vertices x_i^0 for $i \in [m]$. If T contains a vertex of degree at least k, then G contains a subdivision of a k-fan. Thus we may assume that T has maximum degree less than k.

Now, suppress all degree-2 vertices in T, giving a tree T'. Thus every non-leaf vertex of T' has degree between 3 and $k-1$ in T'. In particular, $k \geqslant 4$. Choose an arbitrary non-leaf vertex r of T'. Since T' has $m \geqslant (k-1)^{k^2+2}$ leaves and maximum degree at most $k-1$, it follows that there is a leaf of T' at distance at least $\log_{k-1}|T'| - 1 \geqslant \log_{k-1}(k-1)^{k^2+2} - 1 = k^2+1$ from r in T'.

Consider the path P' of T' from r to that leaf, minus the leaf, and let P denote the corresponding path of T. By construction, there are k^2 vertex-disjoint $V(P)$–$\{x_i^0 \mid i \in [m]\}$ paths in the graph $G[X_0 \cup \{x_i^0 \mid i \in [m]\}]$. Applying Erdős-Szekeres we then find an L_k minor in G. □

Lemma 3.39. *For all $k \in \mathbb{N}$, let $g_{3.39}(k) = g_{3.34}(g_{3.38}(g_{3.37}(g_{3.32}(k))))$. If G is a 3-connected, fan-reduced graph with no $\mathcal{U}_\infty^k$ minor, then the maximum length of a path in G is at most $g_{3.39}(k)$.*

Proof. By Lemmas 3.38, 3.37 and 3.32, we deduce that G has no m-fan minor, where $m = g_{3.38}(g_{3.37}(g_{3.32}(k)))$. Arguing by contradiction, suppose G has a path P of length more than $g_{3.39}(k) = g_{3.34}(m)$.

Let $C_1, \dots, C_p$ denote the components of $G - V(P)$. Let H be the graph obtained from G by contracting each component C_i into a vertex c_i. Note that H has no m-fan minor, since H is a minor of G. By Lemma 3.34, applied to the graph H and path P, there exist two non-consecutive internal vertices u, v of P such that $\{u, v\}$ separates in H the uv-subpath of P from the other vertices of P. However, the same remains true in G, by construction of H. Therefore, $\{u, v\}$ is a cutset of G, contradicting the fact that G is 3-connected. □

In the following we will use another reduction operation for 3-connected graphs. Let G be a 3-connected graph and let $h \geqslant 3$ be a fixed integer. Let $T_1, \dots, T_\ell$ be an enumeration of all stable sets of G satisfying the following conditions for each $i \in [\ell]$,

- $|T_i| \geqslant h + 1$,
- there exists $S_i \subseteq V(G)$ with $|S_i| \leqslant h$ such that for all $v \in T_i$, the set of neighbors of v in G is exactly S_i,
- T_i is inclusion-wise maximal with respect to the above two properties.

Observe that by maximality, the sets $T_1, \dots, T_\ell$ are pairwise disjoint. Let G' be the graph obtained from G by removing all vertices in T_i except $h+1$ of them, for each $i \in [\ell]$. Clearly, G' does not depend on which $h+1$ vertices remain in each T_i. We call G' the *h-reduction* of G. Note that, since G is 3-connected, G' is also 3-connected. If G' is the graph G itself, that is, no vertex was removed in the process, then we say that G is *h-reduced*.

Lemma 3.40. *Let G be a 3-connected graph, let $h \geqslant 3$, and let G' be the h-reduction of G. Then $\tau(G') = \tau(G)$.*

Proof. Since G' is a subgraph of G, $\tau(G') \leqslant \tau(G)$. It remains to show that $\tau(G') \geqslant \tau(G)$.

Let $T_1, \dots, T_\ell$ and $S_1, \dots, S_\ell$ be as in the definition of h-reduction. Let W be a minimum-size vertex cover of G'. We claim $\bigcup_{i \in [\ell]} S_i \subseteq W$. By contradiction, suppose that there exists a vertex $w \in S_i \setminus W$ for some $i \in [\ell]$. Then all edges incident to w have to be covered with all $h+1$ vertices of T_i remaining in G'. However, S_i has at most h vertices. Hence, replacing these $h+1$ vertices of T_i with the at most h vertices of S_i in W gives a smaller vertex cover, a contradiction.

Now, we note that W is also a vertex cover of G, implying that $\tau(G') \geqslant \tau(G)$. To see this, observe that all edges of G that are not in G' are of the form vw with $v \in T_i$ and $w \in S_i$, and every such edge vw is covered by $w \in S_i \subseteq W$. □

Let G be a connected graph and let T be a depth-first search (DFS) tree of G from some vertex r of G. We see T as being rooted at r, and define the usual notions of ancestors and descendants: w is an *ancestor* of v if w is on the r–v path in T, in which case we say that v is a *descendant* of w. Note that these relations are not strict: v is both an ancestor and a descendant of itself. By definition of DFS trees, all edges vw of G are such that either v is a strict ancestor of w in T or v is a strict descendant of w in T.

Lemma 3.41. *For all $k, p \in \mathbb{N}$, let $g_{3.41}(k,p) = ((p+1)2^p + kp^3)^{p+1}$. Let G be a 3-connected graph such that the longest path in G has length at most p, G is p-reduced, and G has no S_k minor. Then $|V(G)| \leqslant g_{3.41}(k,p)$.*

Proof. Let T be a DFS tree of G rooted at some vertex r of G. First we claim that for every vertex v of G, at most $(p+1)2^p$ children of v in T are leaves of T. Indeed, for each such leaf w, the neighborhood of w in G is a subset of the set X of ancestors of v in T. Since G is p-reduced, at most $p+1$ of these leaves have the same neighborhood in G. Moreover, $|X| \leqslant p$, since T has no path of length more than p, implying that there are at most 2^p choices for the neighborhood of w. This implies the claim.

Let
$$d = (p+1)2^p + k(p-1)\binom{p-1}{2} + 1.$$
If T has maximum degree at most d, then since T has at most $p+1$ levels,
$$|V(G)| = |V(T)| \leqslant \sum_{i=0}^{p} d^i = \frac{d^{p+1}-1}{d-1} \leqslant d^{p+1} \leqslant g_{3.41}(k,p),$$
as desired. Hence, it is enough to show that T has maximum degree at most d. For each $x \in V(T)$, we let T_x be the subtree of T rooted at x. Note that if x has at least two children, then the set of ancestors A of x is a cutset of G. Since G is 3-connected, $|A| \geqslant 3$. Partitioning the vertices of T into levels according to their distances from the root, it follows that there is only one vertex on each of the first 3 levels. We argue by contradiction and suppose that there is a vertex v of T having at least d children in T. Since $d \geqslant 2$,

the set X of ancestors of v is a cutset of G with $|X| \geqslant 3$. This implies that v is at distance at least 2 from the root r of T.

Let w be the ancestor of v closest to r in T that is adjacent in G to at least one vertex in T_v. Let P be the w–v path in T. If w has a neighbor in G which is a strict descendant of v, we let v_0 denote a child of v whose subtree T_{v_0} contains a neighbor of w, and let w_0 denote such a neighbor. Otherwise, we just let $v_0 = w_0 = v$. Let C denote the cycle of G obtained by adding the edge ww_0 to the w–w_0 path of T.

Recall that at most $(p+1)2^p$ children of v are leaves of T. Enumerate the non-leaf children of v that are distinct from v_0 as $v_1, \dots, v_q$; thus, $q \geqslant d - (p+1)2^p - 1 = k(p-1)\binom{p-1}{2}$.

Fix some index $i \in [q]$, and let x_i denote a child of v_i in T. We will construct a special K_4-model in G using the cycle C and some vertices of the subtree T_{v_i}. The four vertex images of this K_4-model are denoted V_i, X_i', P_i^1, P_i^2. We proceed with their definitions in the next few paragraphs.

First, observe that every edge out of $V(T_{x_i})$ in $G - v_i$ has its other end in P, by our choice of w. Choose a vertex x_i' in $V(T_{x_i})$ having a neighbor p_i^2 in $V(P)$, with p_i^2 as close to v on P as possible (thus possibly $p_i^2 = v$).

Since G is 3-connected, there is an $\{x_i'\}$–$V(P)$ path Q_i in the graph $G - \{v_i, p_i^2\}$. Let p_i^1 denote the end of Q_i in $V(P)$. Note that all vertices of $Q_i - p_i^1$ are in $V(T_{x_i})$. Also, p_i^1 is a strict ancestor of p_i^2 by our choice of p_i^2.

For a walk W and vertices a, b of W, we write aWb to denote the a–b subwalk of W. If W_1 and W_2 are walks such that W_1 ends at the same vertex that W_2 starts, we let W_1W_2 denote the concatenation of W_1 and W_2.

Next, let R_i be a $\{v_i\}$–$(V(P) \cup V(Q_i))$ path in the graph $G - \{v, x_i'\}$, and let y_i denote its end distinct from v_i. We choose R_i so that y_i is as close as possible to $V(P)$ in the graph $P \cup Q_i$. Let S_i denote the v_i–x_i' path in T. If s_i is the last vertex of R_i contained in S_i, we replace R_i by $S_i s_i R_i$. The definitions of the four vertex images V_i, X_i', P_i^1, P_i^2 depend on whether $y_i \in V(P)$ or not.

First suppose that $y_i \in V(P)$. We define $V_i = V(R_i) \setminus \{y_i\}$ and $X_i' = (V(S_i) \setminus V(R_i)) \cup (V(Q_i) \setminus \{p_i^1\})$. Notice that there is an edge e_i of S_i with one end in V_i and the other in X_i'. The two sets P_i^1, P_i^2 will be a partition of the vertices of the cycle C, chosen as follows. If y_i is a strict ancestor of p_i^2, let P_i^1 be the vertices of the p_i^1–y_i path of T, and let $P_i^2 = V(C) \setminus P_i^1$. If, on

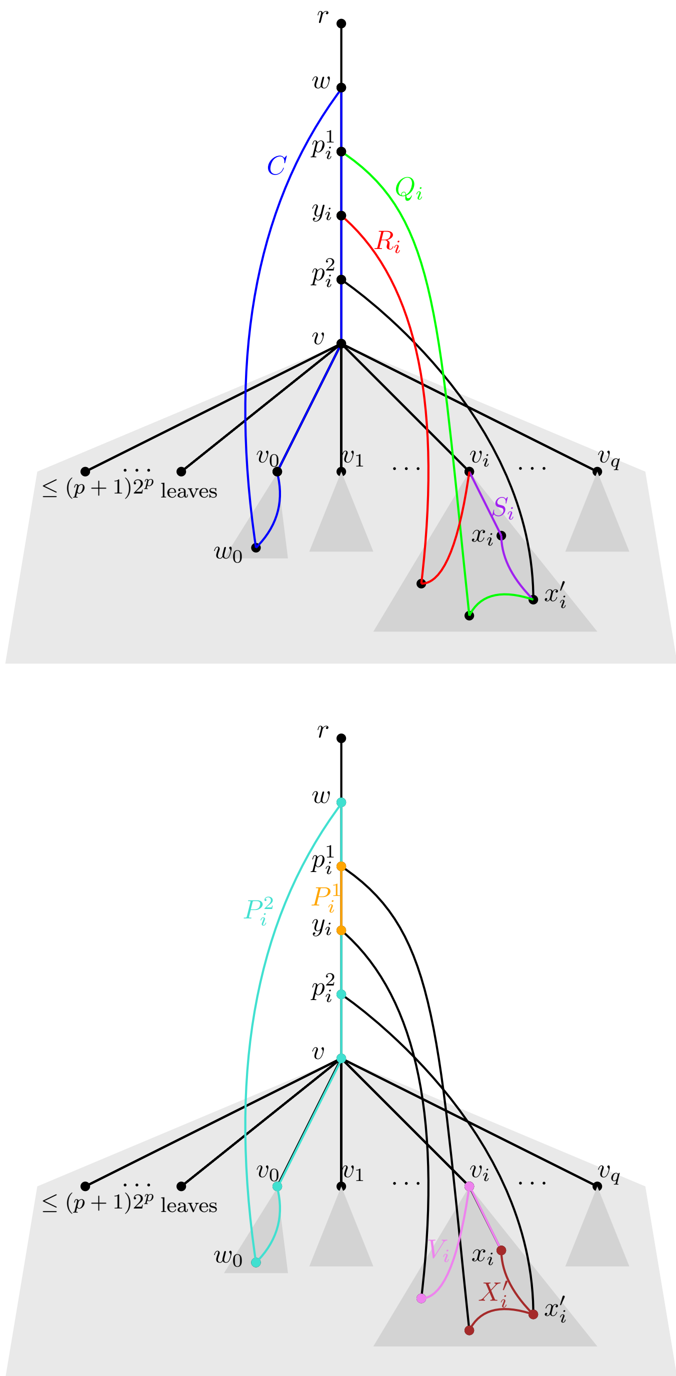

Figure 3.13. The case $y_i \in V(P)$ of the proof of Lemma 3.41.

the other hand, y_i is a descendant of p_i^2, let P_i^2 be the vertices of the p_i^2–y_i path of T, and let $P_i^1 = V(C) \setminus P_i^2$. This case is illustrated in Figure 3.13.

We now argue that the sets V_i, X_i', P_i^1, P_i^2 do form a K_4-model in this case. These sets are connected, there is an edge between P_i^1 and P_i^2 (because of the cycle C), there is an edge between X_i' and P_i^j for $j \in [2]$ (because $p_i^j \in P_i^j$), there is an edge between V_i and X_i' (namely, e_i), and finally there is an edge between V_i and P_i^j for $j \in [2]$ (because one of v, y_i is in P_i^1 and the other is in P_i^2). This concludes the case where $y_i \in V(P)$.

Next, suppose that $y_i \notin V(P)$. In this case, y_i is a vertex of $Q_i - p_i^1$. Consider an $\{v_i\}$–$V(Q_i)$ path R_i' in $G - \{v, y_i\}$. Note that, by our choice of R_i, the path R_i' avoids $V(P)$, and thus all its vertices are in $V(T_{v_i})$. Furthermore, the end y_i' of R_i' distinct from v_i must be in the subpath $x_i' Q_i y_i - \{y_i\}$, again by our choice of R_i.

Define

$$\begin{aligned} V_i &= (V(R_i) \setminus \{y_i\}) \cup (V(R_i') \setminus \{y_i'\}) \\ X_i' &= V(x_i' Q_i y_i) \setminus \{y_i\} \\ P_i^1 &= V(y_i Q_i p_i^1) \\ P_i^2 &= V(C) \setminus \{p_i^1\} \end{aligned}$$

Using the previous observations, one can check that V_i, X_i', P_i^1, P_i^2 form a K_4-model in this case as well. This case is illustrated in Figure 3.14.

This ends the definitions of the vertex images V_i, X_i', P_i^1, P_i^2. Observe that, in all cases, the only vertices of these sets *not* in the subtree T_{v_i} are the vertices of the cycle C.

Now, there are at most $\binom{p-1}{2}$ choices for p_i^1 and p_i^2. Furthermore, when $y_i \in V(P)$, there are at most $p-2$ choices for vertex y_i. Seeing the possibility that $y_i \notin V(P)$ as another 'choice', and using that $q \geqslant k(p-1)\binom{p-1}{2}$, we conclude that there is a set I of k distinct indices $i \in [q]$ that have the same pair (p_i^1, p_i^2), that agree on whether $y_i \in V(P)$, and furthermore that have the same vertex y_i in case $y_i \in V(P)$. Letting $P^j = \bigcup_{i \in I} P_i^j$ for $j \in [2]$, we then see that P^1, P^2 together with the sets V_i, X_i' for $i \in I$ define an S_k-model in G, a contradiction. □

Lemma 3.42. *For all $k \in \mathbb{N}$, let $g_{3.42}(k) = g_{3.41}(k, g_{3.39}(k))$. If G is a 3-connected, fan-reduced graph having no $\mathcal{U}_\infty^k$ minor, then $\tau(G) \leqslant g_{3.42}(k)$.*

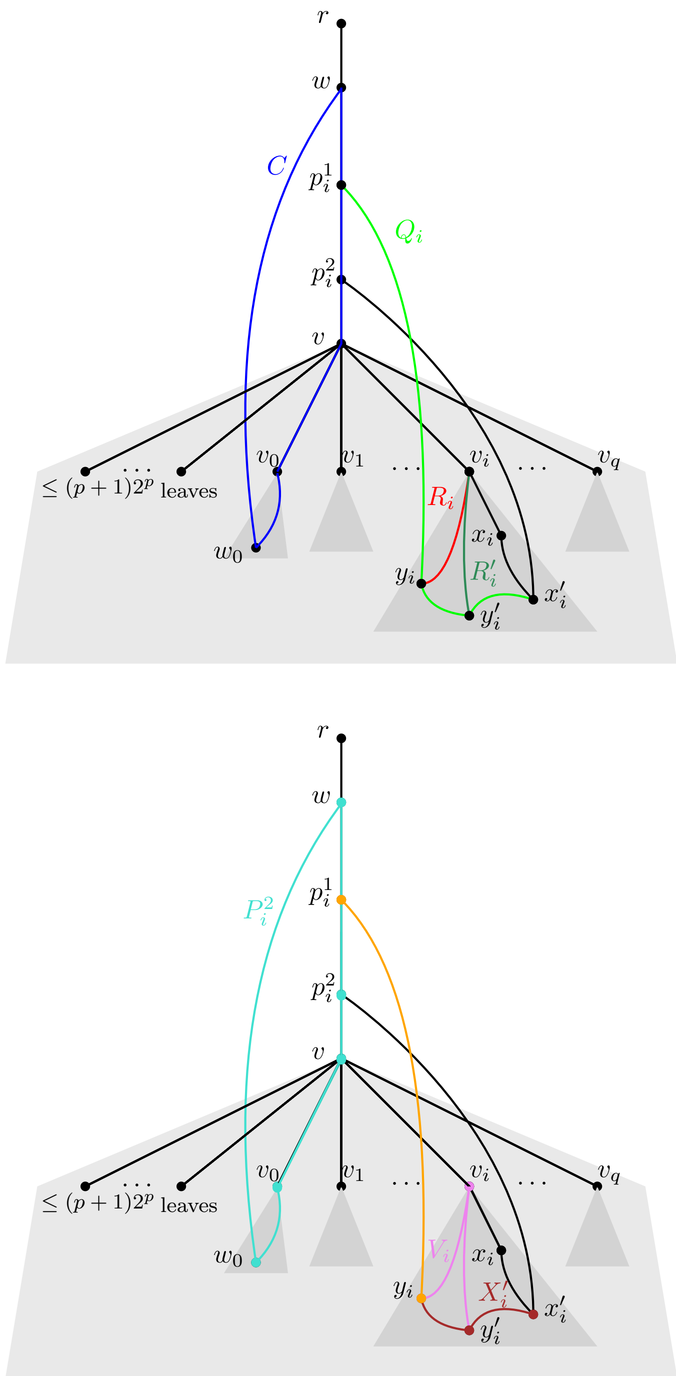

Figure 3.14. The case $y_i \in V(Q_i)$ of the proof of Lemma 3.41.

Proof. By Lemma 3.39, the maximum length of a path in G is at most $p = g_{3.39}(k)$ since G is 3-connected, and does not have a $\mathcal{U}^k_\infty$ minor. Let G' be the p-reduction of G. Notice that G' is 3-connected, has no S_k minor and the length of a longest path in G' is bounded by p. Hence, by Lemma 3.41, $\tau(G') \leqslant |V(G')| \leqslant g_{3.41}(k,p)$. Now, by Lemma 3.40,

$$\tau(G) = \tau(G') \leqslant g_{3.41}(k,p) = g_{3.41}(k, g_{3.39}(k)) = g_{3.42}(k). \qquad \square$$

3.7 Finishing the proof

Recall that to prove our main result, Theorem 1.3, it suffices to establish the existence of the functions $g_{3.47}$ and $g_{3.48}$ from Lemma 3.31. We do this in Lemmas 3.47 and 3.48 at the end of this section. Before doing so, we require a few more lemmas. The *wheel* W_n is the graph obtained by adding a universal vertex to a cycle of length n.

Lemma 3.43. *$f_\infty(W_n) \leqslant 4$, for all $n \geqslant 3$.*

Proof. Let v_0 be the universal vertex of W_n and $W_n - v_0 = C = v_1 \cdots v_n v_1$. Let d be an arbitrary distance function on W_n. Define $\mathcal{S}$ to be the set of inclusion-wise minimal subsets S of $E(C)$ such that S is not flattenable in (W_n, d). Let d' be d restricted to $E(C)$. Let $\mathcal{S}_1$ be the sets in $\mathcal{S}$ that are not flattenable in (C, d'), and let $\mathcal{S}_2 = \mathcal{S} \setminus \mathcal{S}_1$.

Fix $S \in \mathcal{S}_2$ and let $\vec{S}$ be an orientation of S such that $\vec{S}$ is flat in (C, d'). Let the length function of $\langle W_n, d; \vec{S}\rangle$ be l, and Z be a negative directed cycle in $\langle W_n, d; \vec{S}\rangle$. Since S is flattenable in (C, d'), Z must use the vertex v_0. By renaming vertices, we may assume that Z is of the form $v_0 v_1 \cdots v_k v_0$. Let $P = v_1 \cdots v_k$ and $Q = v_k \cdots v_n v_1$. We abuse notation and regard P, Q, and C as subsets of edges or arcs whenever convenient.

Since $\vec{S}$ is flat in (C, d'), $l(C) \geqslant 0$. Combining this with $l(Z) < 0$ gives

$$d(v_0v_1) + d(v_0v_k) < l(Q) \leqslant d(Q) \text{ and } d(v_0v_1) + d(v_0v_k) < l(P) \leqslant d(P). \tag{3.8}$$

Let H_1 and H_2 be the subgraphs of W_n induced by $\{v_0, v_1, \ldots, v_k\}$ and $\{v_0, v_k, \ldots, v_n, v_1\}$, respectively. Let d_i be the restriction of d to H_i. Clearly, each (H_i, d_i) can be covered by two flat sets F_i^1, F_i^2. By (3.8), every negative directed cycle W in $\langle W_n, d; F_i^j\rangle$ can be shortened to a negative directed cycle

W' in $\langle H_i, d_i; F_i^j \rangle$ for all $i, j \in [2]$. Therefore, F_i^j is also flat in (W_n, d) for all $i, j \in [2]$. Thus, (W_n, d) has a flat cover of size 4.

We may therefore assume that $\mathcal{S}_2 = \emptyset$. That is, every set in $\mathcal{S}$ is not flattenable in (C, d'). Let U be the set of edges of W_n incident to v_0. Note that U is flattenable in (W_n, d) by Lemma 3.6. If $\mathcal{S}_1 = \emptyset$, then $E(C)$ is flattenable in (W_n, d), and so $E(W_n)$ is the union of two flattenable sets, $E(C)$ and U. Therefore, we may assume $\mathcal{S}_1 \neq \emptyset$ and choose $T \in \mathcal{S}_1$. Let $X \subseteq E(C)$. Observe that if $\sum_{e \in X} d(e) \leqslant \frac{1}{2} d(C)$, then X is flattenable in (C, d'). It follows that for every $X \subseteq E(C)$, at least one of X or $E(C) \setminus X$ is flattenable in (C, d'). Since T is not flattenable in (C, d'), $E(C) \setminus T$ is flattenable in (C, d'). Since $\mathcal{S}_2 = \emptyset$, $E(C) \setminus T$ is flattenable in (W_n, d). By minimality, T is the union of two flattenable sets T_1 and T_2 of (W_n, d). Thus, $E(W_n) = (E(C) \setminus T) \cup T_1 \cup T_2 \cup U$, as required. ☐

We now generalize Lemma 3.43. This generalization is analogous to Lemma 3.30 for 2-connected treewidth-2 graphs.

Lemma 3.44. *Let H be a graph obtained by gluing 2-connected graphs $G_1, \ldots, G_m$ on distinct edges of the wheel W_n, such that H has no S_k minor. Let $M = \max_{i \in [m]} f_\infty(G_i)$. Then $f_\infty(H) \leqslant (k+7)M$.*

Proof. Let $W_n - v_0 = C = v_1 \cdots v_n$. We proceed by induction on $|V(H)|$. By Lemma 3.11, we may assume that H has minimum degree at least 3. Let E_0 be the set of glued edges incident to v_0. If $|E_0| \geqslant k$, then W_n has a k-glumpkin minor. By Lemma 3.27, H contains an S_k minor, which is a contradiction. Thus, $|E_0| \leqslant k - 1$.

Let d be an arbitrary distance function on H, and d_W be the restriction of d to W_n. By Lemma 3.43, (W_n, d_W) has a flat cover of size 4, say F_1, F_2, F_3, F_4. Let F_0 be the set of arcs of $D(W_n)$ incident to v_0. For each $i \in [4]$, let Γ_i^+, Γ_i^- be such that $\Gamma_i^+ \cup \Gamma_i^- = F_i \setminus F_0$ and $(v_{j+1}, v_j) \notin \Gamma_i^+$, $(v_j, v_{j+1}) \notin \Gamma_i^-$ for all $j \in \mathbb{Z}/n\mathbb{Z}$. Since every two arcs of $\Gamma_i^\pm$ are both forward or both backward arcs of every directed cycle of $D(W_n)$, $(\Gamma_i^\pm, F_i)$ is a frame of (W_n, d_W) for all $i \in [4]$. Let H' be the graph obtained from W_n by only gluing along glued edges belonging to $E(C)$. By Lemma 3.6 and Lemma 3.24, $f_\infty(H') \leqslant 1 + 8M$. Since $|E_0| \leqslant k - 1$, Lemma 3.12 implies that

$$f_\infty(H) \leqslant f_\infty(H') + (k-1)(M-1) \leqslant (k+7)M.$$ ☐

We now apply our results about wheels to fan-reduced graphs. Recall that every graph can be obtained from its fan-reduction by replacing fan gadgets by fans.

Lemma 3.45. *Let F be a reducible fan of a graph G, and let G' be the F-reduction of G. Then $f_\infty(G) \leqslant f_\infty(G') + 4$.*

Proof. Let v_0 be the center of F, and $v_1 \cdots v_k$ be its outer path. When performing the F-reduction, we rename vertices such that v_0 is still the center and $v_1 v_2 v_{k-1} v_k$ is the outer path of the reduced fan. Let W_{k-2} be the wheel graph on $k-1$ vertices, where v_0 is the universal vertex, and $v_2 v_3 \cdots v_{k-1} v_2$ is the outer cycle. Let H be the graph obtained by performing the 3-sum of G' with W_{k-2} along the clique $v_0 v_2 v_{k-1}$. Note that G is obtained from H by deleting the edge $v_2 v_{k-1}$. Hence, $f_\infty(G) \leqslant f_\infty(H)$. By Lemma 3.43, $f_\infty(W_{k-2}) \leqslant 4$. Therefore, applying Lemma 3.12,

$$f_\infty(G) \leqslant f_\infty(H) \leqslant f_\infty(G') + f_\infty(W_{k-2}) \leqslant f_\infty(G') + 4. \qquad \square$$

Lemma 3.46. *Let G be a graph, G' be the fan-reduction of G, and t be the number of reduced fans in G'. Then, $t \leqslant \tau(G')$ and $f_\infty(G) \leqslant 5\tau(G')$.*

Proof. Suppose F' is a reduced fan in G', where v_0 is the center and $v_1 \cdots v_4$ is the outer path. Note that every vertex cover of G' must use at least one of v_2 or v_3. Since $\{v_2, v_3\}$ is disjoint from all other reduced fans, we conclude that $t \leqslant \tau(G')$. For the second part, first observe that $f_\infty(G') \leqslant \tau(G')$, by Lemma 3.7. By repeatedly applying Lemma 3.45 to each maximal reducible fan of G,

$$f_\infty(G) \leqslant f_\infty(G') + 4t \leqslant 5\tau(G'). \qquad \square$$

Lemma 3.47. *For all $k \in \mathbb{N}$, let $g_{3.47}(k) = 5g_{3.42}(k)$. If G is a 3-connected graph with no $\mathcal{U}_\infty^k$ minor, then $f_\infty(G) \leqslant g_{3.47}(k)$.*

Proof. Let G' be the fan-reduction of G. By Lemmas 3.46 and 3.42,

$$f_\infty(G) \leqslant 5\tau(G') \leqslant 5g_{3.42}(k) = g_{3.47}(k). \qquad \square$$

Lemma 3.48. *For all $k, M \in \mathbb{N}$, let $g_{3.48}(k, M) = (2k+11)Mg_{3.42}(k)$. Let G be a 3-connected graph and let H be a graph obtained by gluing 2-connected graphs $G_1, \ldots, G_m$ on distinct edges of G such that H has no $\mathcal{U}_\infty^k$ minor. Let $M = \max_{i \in [m]} f_\infty(G_i)$. Then $f_\infty(H) \leqslant g_{3.48}(k, M)$.*

Proof. We proceed by induction on $|E(H)|$. By Lemma 3.11, we may assume that H has minimum degree at least 3. Let $\mathcal{F}$ be the set of maximal reducible fans in G. Let G' be the fan-reduction of G and let $\mathcal{F}'$ be the set of reduced fans in G'. If F is a fan with center v_0 and outerpath $v_1 \cdots v_m$, we define $I(F) = V(F) \setminus \{v_0, v_1, v_m\}$. Let X' be a vertex cover of G' and set $X = X' \setminus \bigcup_{F' \in \mathcal{F}'} I(F')$. We regard X as a subset of vertices of G. Let Γ be the set of glued edges of G and Γ_X be the set of edges of Γ incident to a vertex in X.

If $|\Gamma_X| > (k-1)\tau(G')$, then there is a vertex $x \in X$ incident to at least k glued edges $xy_1, \ldots, xy_k$. Since G is 3-connected, there is a tree in $G - x$ containing $\{y_1, \ldots, y_k\}$. Therefore, G contains a k-glumpkin minor that is obtained by contracting the tree to a single vertex. By Lemma 3.27, H contains an S_k minor, which is a contradiction. Hence, $|\Gamma_X| \leqslant (k-1)\tau(G')$.

Let $F \in \mathcal{F}$ with center v_0 and outerpath $v_1 \cdots v_m$. Let F^+ be the graph obtained from F by adding the edge $v_1 v_m$ (if it is not already present) and gluing all G_i whose glued edge is contained in $E(F)$.

Let G^X be obtained from G by gluing all G_i whose glued edge belongs to Γ_X and replacing each $F \in \mathcal{F}$ by a triangle, Δ_F. Let H^+ be obtained from G^X by simultaneously taking the clique-sum of F^+ and G^X along Δ_F for all $F \in \mathcal{F}$. Notice that H is a subgraph of H^+.

By Lemma 3.46, $f_\infty(G) \leqslant 5\tau(G')$. Since $|\Gamma_X| \leqslant (k-1)\tau(G')$, by Lemma 3.12

$$f_\infty(G^X) \leqslant f_\infty(G) + (k-1)(M-1)\tau(G') \leqslant (k+4)M\tau(G').$$

Since G' is a 3-connected fan-reduced graph not containing a $\mathcal{U}^k_\infty$ minor, by Lemma 3.42, $\tau(G') \leqslant g_{3.42}(k)$. By Lemma 3.44, $f_\infty(F^+) \leqslant (k+7)M$, for all $F \in \mathcal{F}$. Finally, $|\mathcal{F}| \leqslant \tau(G')$, by Lemma 3.46. Putting this altogether,

$$\begin{aligned}
f_\infty(H) &\leqslant f_\infty(H^+) \\
&\leqslant f_\infty(G^X) + (k+7)M\tau(G') \\
&\leqslant (k+4)M\tau(G') + (k+7)M\tau(G') \\
&= (2k+11)M\tau(G') \\
&\leqslant (2k+11)M g_{3.42}(k) \\
&= g_{3.48}(k, M).
\end{aligned}$$

□

3.8 Minimal excluded Minors for ℓ_∞–dimension 3

Another approach of research is to establish the complete lists of minimal excluded minors for the property $f_p(G) \leqslant k$ for small $k \in \mathbb{N}$ and some value for p. We will now focus on the case $p = \infty$ and $k = 3$. By the Robertson-Seymour theorem, Theorem 1.1, we know that there exists a finite set of minimal excluded minors for the property $f_\infty(G) \leqslant 3$.

Several results of this chapter show respectively how to obtain upper and lower bounds for a given graph. For instance, we can derive from Lemma 3.15 that $f_\infty(G) \leqslant |V(G)| - 2$. The lower bounds are obtained when knowing that a graph contains some minor with big f_∞ value. We can use these results to find restrictions for minimal excluded minors. As $f_\infty(G) \leqslant \tau(G)$, we know that any minimal excluded minor for $f_\infty(G) \leqslant 3$ has vertex cover number at least four. Also, we know by Lemma 3.11 that a minor minimal graph has no two adjacent degree-2 vertices.

In my Master thesis [60], I considered the case $f_\infty(G) \leqslant 3$ and provided a list of graphs that are not realizable in ℓ_∞^3. However, it is not known for all these graphs whether they are minimal, and whether the list is complete. During the first few months of my PhD research, we investigated this problem further and noticed that some graphs had some common minors, which are also not realizable in ℓ_∞^3. Furthermore, we identified some more graphs that are not realizable in ℓ_∞^3. As there are no efficient tools yet to prove whether a given graph is a minimal excluded minor it is not known whether we can improve the current list (in the sense that we identify a minor of one of the graphs as being not realizable in ℓ_∞^3).

Another hard problem is proving the completenessness of the list of minors that we know of. No attempt in that direction has been made because of a lack of efficient tools.

In Appendix A we give a list of metric graphs that are not realizable in ℓ_∞^3. They are listed in the form `v1 v2 w` where `v1` and `v2` form an edge that has weight `w`.

Chapter 4

Cut Dominants

This chapter is based on unpublished joint work with Samuel Fiorini.

Cuts in graphs are a well studied subject in graph theory and combinatorial optimization. A *cut* in a graph is a set of edges whose removal disconnects the graph. Formally, $X \subseteq E(G)$ is a cut if $G \setminus X$ is not connected. Usually, we want to find a minimum or maximum cut in a graph with non-negative edge-weights, that is a cut which minimizes or maximizes the sum of the weights of its edges. We will focus mainly on the min-cut problem although the max-cut problem is relevant, too. For instance, in statistical physics, the max-cut problem gives the minimizers of the Hamiltonian of the Ising model [4], which was introduced in the 1920s.

We discuss s–t cuts first. An *s–t cut* is a cut that separates two fixed vertices s and t. Formally, given two vertices s and t in graph G, an s–t cut is a set of edges X such that s and t are in different connected components of $G \setminus X$.

Computing a minimum s–t cut can be done in polynomial time by using the Edmonds-Karp algorithm [36] combined with the max-flow min-cut theorem. The algorithm uses an augmenting path method for maximum flow introduced by Ford and Fulkerson [45]. The algorithm of Edmonds and Karp runs in $O(nm^2)$ time, where n is the number of vertices and m is the number of edges of the input graph G. Note that this algorithm can be extended to find a global minimum cut in a graph by computing a minimum s–t cut for all pairs s, t of distinct vertices in the graph and taking the best possible cut.

A more efficient approach to compute a minimum s–t cut for all possible pairs s, t of vertices is to use a Gomory-Hu tree [49]. Such a tree can be computed by performing $n-1$ maximum flow computations, where n is the number of vertices of the graph. Let G be an edge-weighted graph. An edge-weighted tree T with the same vertex set as G is a *Gomory-Hu tree* for G if for every two vertices s, t the minimum s–t cut in T has same weight as the minimum s–t cut in G. Observe that by taking the cheapest edge of T, we find a global minimum cut in G.

When it comes to computing a (global) minimum cut, the fastest deterministic algorithm is due to Ibaraki and Nagamochi [61]. Their algorithm operates in two steps which are repeated $n-1$ times. The first step consists of finding an appropriate order of the vertices. In a second step, they compute the weight of a cut separating the last vertex of the ordering from all the other vertices and put it in a list along with the vertex, indexed by the current step. Then, they shrink the two last two vertices of the ordering to one vertex and continue with the first step in the resulting smaller graph until only two vertices remain. Finally, the global minimum cut is given by the minimum weight of a cut in the list. Their algorithm runs in $O(nm)$ time, where m is the number of edges of G. In the randomized case, the fastest algorithm is due to Karger and Stein [53], which runs in $O(n^2 \log^3 n)$ time.

Schrijver [70] notices that for many combinatorial optimization problems the three following properties are related to one another. First, the existence of a polynomial time algorithm. Second, the existence of a min-max relation for the problem. Third, a "nice" polyhedral description, in the sense that the linear description is well understood. Most problems in Schrijver's book [70] satisfy all of these properties.

We have already seen that there exist several polynomial time algorithms for the min-cut problem. Hence, we should look for a polyhedral description (or a min-max relation) in order to gain further understanding of the problem. Recall that the cut dominant is defined as the Minkowski sum of the cut polytope and the non-negative orthant. Soving the min-cut problem in a graph G can be done by minimizing a linear function on the cut dominant of G. This is one of the reasons why we are interested in understanding the facets of the cut dominant.

Despite the fact that a complete characterization of the facets of the cut dominant is not known, we know that the cut dominant has polynomial extension complexity, see [18]. Hence, the linear description of the cut dominant is easier to understand than the linear description of the cut polytope, which is known to have exponential extension complexity [44]. Roughly speaking, this shows that taking the Minkowski sum of the cut polytope with the non-negative orthant suppresses the part of the cut polytope that is hardest to understand geometrically.

Another reason to understand the geometry of the cut dominant is that it is the blocking polar of the subtour elimination relaxation of the TSP polytope. That is, the vertices of the subtour elimination relaxation correspond to the facets of the cut dominant. Gaining a good understanding of these vertices is important in many algorithms solving the TSP. See the book of Applegate, Bixby, Chvátal, and Cook [2] for exact algorithms, or the book of Williamson and Shmoys [78] for approximation algorithms. For more recent work in approximation algorithms, see for instance [48], [71] and [73].

However, it is an open problem to determine the exact linear description of the cut dominant, or the vertices of the subtour elimination relaxation, in general. These are known for some graph classes of bounded treewidth, such as trees or series-parallel graphs, see [28]. However, these graphs are very restrictive. So, another approach is to study graphs for which the right-hand side of the inequalities defining the cut dominant is bounded. This is why we consider the parameter $\varphi(G)$, which is defined to be the maximum right-hand side of a non-trivial facet-defining inequality of the cut dominant of G in minimum integer form.

In order to state our results formally we recall the definitions of the cut dominant and the parameter we are studying. Given a graph $G = (V, E)$ we define $\delta(S) := \{uv \mid u \in S, v \notin S\}$ if $S \subseteq V(G)$. Its incidence vector $\chi^{\delta(S)}$ is such that $\chi_e^{\delta(S)} = 1$ if exactly one end of e is in S, and $\chi_e^{\delta(S)} = 0$ otherwise. The *cut polytope* $\text{cut}(G)$ is defined as the convex hull of all incidence vectors of the cuts in G. That is, $\text{cut}(G) = \text{conv}\{\chi^{\delta(S)} \mid \emptyset \neq S \subsetneq V(G)\}$. Observe that the cut polytope is a 0/1-polytope because all its vertices are 0/1-vectors. Hence, $\text{cut}(G) \subseteq \mathbb{R}_+^{E(G)}$.

The *cut dominant* $\text{cutdom}(G)$ *of* G is defined as the dominant of the cut polytope, that is

$$\text{cutdom}(G) = \text{cut}(G) + \mathbb{R}_+^{E(G)}.$$

Recall from Chapter 2 that the cut dominant, as all 0/1 polyhedra, has a linear description whose constraints are in minimum integer form. That is, the coefficients and right-hand side of every inequality of the system are integers without common factor.

Let $\{\sum_{e\in E(G)} c_i(e)x_e \geqslant \lambda_i\}_{i\in I}$ be a non-redundant linear description of $\mathrm{cutdom}(G)$ in minimum integer form. In general, little is known about this linear description. However, Conforti, Rinaldi, and Wolsey [27] showed that $\lambda_i \in 2\mathbb{N} \cup \{1\}$ for all graphs and all $i \in I$. We are interested in the largest right-hand side λ_i. If G is a connected graph, we let

$$\varphi(G) := \max_{i\in I} \lambda_i$$

and if G is not connected we let $\varphi(G)$ be the maximum of $\varphi(H)$ over all connected components H of G. We point out that this definition is slightly different from the one used in previous papers.[1]

Observe that $\varphi(G) \in \{1\} \cup 2\mathbb{N}$ for all graphs G because these are the only values that right-hand side coefficients can take [27]. Moreover, it is known that all graphs G satisfying $\varphi(G) \leqslant k$ form a minor-closed class for every $k \in \mathbb{N}$, see Lemma 4.1. Hence, by the Graph Minor Theorem, Theorem 1.1, there exists a finite set of minimal excluded minors for the property $\varphi(G) \leqslant k$ for every $k \in \mathbb{N}$.

Observe that $\varphi(G) = 0$ if G has no edge and $\varphi(G) \geqslant 1$ if G has at least one edge. It is easy to check that $\varphi(G) \leqslant 1$ if G is a forest. Furthermore, we can show $\varphi(K_3) = 2$, which implies that K_3 is the only minimal excluded minor for $\varphi(G) \leqslant 1$. Indeed, to see $\varphi(K_3) = 2$ it is sufficient to verify that the largest right-hand side λ_i in the minimum integer form linear description of $\mathrm{cutdom}(G)$ is exactly 2, see Figure 4.1.

Conforti, Fiorini and Pashkovich [26] showed that $\varphi(G) \leqslant 2$ if and only if G has no pyramid or prism minor. These minimal excluded minors are shown in Figure 1.4 on page 8.

In this chapter we are interested in an excluded minor characterization of the graphs satisfying $\varphi(G) \leqslant 4$. We introduce 12 graphs that are minimal excluded minors for this property. Furthermore, we prove some properties of minor-minimal graphs G with $\varphi(G) > 4$ that are not internally 3-connected. The main theorem of Section 4.4, Theorem 4.22, shows that they need to

[1] [27] and [26] defined $\varphi(G) := \max_{i\in I} \lambda_i$ for all graphs G. However, with this definition the classes of graphs satisfying $\varphi(G) \leqslant k$ are not closed under taking minors (even for $k = 0$). Indeed, any disconnected graph satisfies $\varphi(G) = 0$ even though its connected components H, which are also minors of G, satisfy $\varphi(H) > 0$.

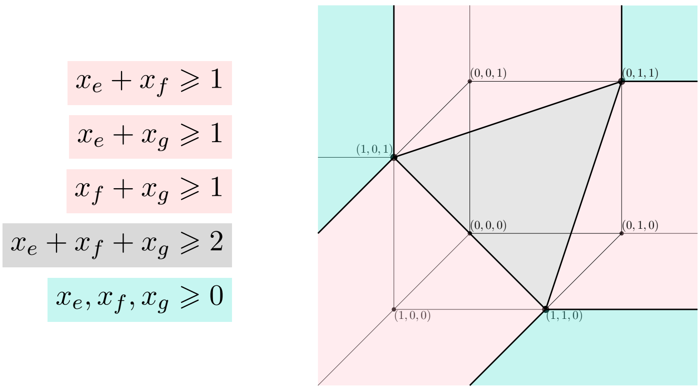

Figure 4.1. Linear description in minimum integer form of cutdom(K_3). The first three equations define the red facets, the next the gray facet, and the last three equations the blue facets.

satisfy $\varphi(G) = 8$. Moreover, we bound $\varphi(G)$ as a function of $\tau(G)$, the vertex cover number of G.

In Section 4.1 we give an overview of the properties of facets of the cut dominant for general graphs. Using these tools, we bound $\varphi(G)$ as a function of $\tau(G)$ in Section 4.2. In Section 4.3 we introduce some minimal excluded minors for $\varphi(G) \leqslant 4$. After that, in Section 4.4, we show some properties that are satisfied by minor-minimal graphs G with $\varphi(G) > 4$. Section 4.5 introduces amplifiers, which can be used to double the $\varphi(G)$-value of graphs. Finally, in Section 4.6 we state some open questions and conjectures.

4.1 General results about facets of cut dominants and $\varphi(G)$

We start with a lemma adapted from [26, Lemma 3] which allows us to actually apply the Graph Minor Theorem to find minimal excluded minors for $\varphi(G) \leqslant k$, $k \in \mathbb{N}$.

Lemma 4.1. *Let G be a graph and let H be a minor of G. Then $\varphi(H) \leqslant \varphi(G)$.*

Remark that we consider only weighted graphs (G, c) with non-negative weights. The reason is that every facet-defining inequality is of the form

$x_e \geqslant 0$ for some edge $e \in E(G)$, which we call a *trivial* inequality, or $\sum_{e \in E(G)} c(e)x_e \geqslant \lambda(G,c)$ for some $c : E(G) \to \mathbb{Q}_+$, where $\lambda(G,c)$ denotes the minimum weight of a cut in (G,c), and moreover $\lambda(G,c) > 0$.

We say that a *family $\mathcal{F}$ of non-empty proper subsets of $V(G)$ defines linearly independent minimum cuts* in (G,c) if the following conditions are satisfied.

1. For every $S \in \mathcal{F}$ the cut $\delta(S)$ is a minimum cut in (G,c).

2. The incidence vectors $\chi^{\delta(S)}$ of all cuts $\delta(S)$ with $S \in \mathcal{F}$ are linearly independent.

A family of subsets $\mathcal{F}$ is *laminar* if for every sets $S, S' \in \mathcal{F}$ either S and S' are disjoint or one set is completely contained in the other one.

The following result by Cornuéjols, Fonlupt, and Naddef [28] characterizes the facet-defining inequalities of cutdom(G).

Theorem 4.2 (Characterization of facet-defining inequalities of cutdom(G))**.** *Let $\sum_{e \in E(G)} c(e)x_e \geqslant k$ be a valid inequality for* cutdom(G) *with $k > 0$, and r edges with $c(e) \neq 0$. Then the inequality $\sum_{e \in E(G)} c(e)x_e \geqslant k$ is facet-defining if and only if $\lambda(G,c) = k$, and there exists a family $\mathcal{F}$ of r subsets of $V(G)$ defining linearly independent minimum cuts in (G,c). Furthermore, the family $\mathcal{F}$ can be chosen to be laminar.*

It is also known that the dimension of the face determined by a valid inequality $\sum_{e \in E(G)} c(e)x_e \geqslant k$ with $k > 0$ equals $a + b - 1$, where a is the number of edges e with $c(e) = 0$ and b counts the number of linearly independent minimum cuts in the graph G^c, where G^c denotes the graph whose vertices are those of G and whose edges are the edges $e \in E(G)$ such that $c(e) \neq 0$. In the following of the chapter we will assume that all inequalities we consider are valid for the cut dominant of the graph we consider.

The next two lemmas combine several results from earlier work, see Cornuéjols, Fonlupt, and Naddef [28], Conforti, Rinaldi, and Wolsey [27], Conforti, Fiorini, and Pashkovich [26].

Given a facet-defining inequality of cutdom(G) and a corresponding laminar family $\mathcal{F}$ defining linearly independent minimum cuts, we can define the *level* of each set $S \in \mathcal{F}$ recursively. Sets that do not contain any other set have level zero, $level(S) = 0$. For all other sets S we set $level(S) = 1 + \max_{S_i} level(S_i)$, where we take the maximum over all sets $S_i \subsetneq S$. We call any set $S \in \mathcal{F}$ such that *level*$(S) = i$ a *level-i set*, for all $i \in \mathbb{N}$. The following lemma lists several properties that are satisfied by any facet-defining inequality of $\text{cutdom}(G)$.

Lemma 4.3. *Let $G = (V, E)$ be a graph, $\sum_{e \in E} c(e)x_e \geqslant k$ be a facet-defining inequality for* $\text{cutdom}(G)$ *with $k > 0$, and let $\mathcal{F}$ be a family defining $|E|$ linearly independent minimum cuts in (G, c). Then the following hold.*

1. *If the facet-defining inequality is in minimum integer form, then $k \in \{1\} \cup 2\mathbb{N}$.*
2. *For every $S \in \mathcal{F}$, the induced subgraphs $G[S]$ and $G[\overline{S}]$ are both connected.*
3. *The graph $G^c = (V, E^c)$, where $E^c = \{e \in E \mid c(e) \neq 0\}$, is simple.*
4. *For every $e \in E$ there exists at least one $S \in \mathcal{F}$ such that $e \in \delta(S)$.*
5. *If the family $\mathcal{F}$ is laminar, then every level-0 set is a singleton.*

It follows from Lemma 4.3 that $\varphi(G) \in \{1\} \cup 2\mathbb{N}$ for every graph G. Furthermore, if G has no edge, then $\varphi(G) = 0$. While the previous lemma is valid for any facet-defining inequalities of $\text{cutdom}(G)$ for any graph G, the next lemma focuses on so-called witnesses in minor-minimal graphs with $\varphi(G) > k$. A *witness* for a minor-minimal graph G with $\varphi(G) > k$ is a non-trivial facet-defining inequality of $\text{cutdom}(G)$ such that, when put in minimum integer form, has right-hand side strictly greater than k.

Lemma 4.4. *Let G be a minor-minimal graph with $\varphi(G) > k$. Let $\sum_{e \in E(G)} c(e)x_e \geqslant k$ be a witness for $\varphi(G) > k$. Let $\mathcal{F}$ be a family defining $|E(G)|$ linearly independent minimum cuts in (G, c). Then the following assertions hold.*

1. *$c(e) > 0$ for every $e \in E(G)$.*
2. *$c(e) \leqslant k/2$ for every $e \in E(G)$.*
3. *Every level-1 set in $\mathcal{F}$ is of the form $S = \{u, v\}$ with $uv \in E(G)$, such that $c(uv) = c(\delta(u) \setminus \{uv\}) = c(\delta(v) \setminus \{uv\}) = k/2$.*
4. *For every $e \in E(G)$ there exists at least two minimum cuts $\delta(S)$, $S \in \mathcal{F}$, such that $e \in \delta(S)$.*

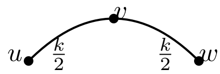

Figure 4.2. The incident edges of a degree-2 vertex have weight $k/2$ if $\sum c(e)x_e \geqslant k$ is a witness for $\varphi(G) > k$.

Notice that Lemma 4.4 is invariant under scaling. This also holds for Lemma 4.6 below.

The next lemma shows that the support of any facet has only linearly many edges. It follows from Theorem 4.2 and the well-known fact that a laminar family of subsets of a set of size n contains at most $2n-3$ subsets.

Lemma 4.5. *Let G be a graph with n vertices and $\sum c(e)x_e \geqslant k$ be a facet-defining inequality of* $\text{cutdom}(G)$*. Then at most $2n-3$ edges satisfy $c(e) > 0$. Consequently, if G is minor-minimal with fixed $\varphi(G)$, then G has at most $2n-3$ edges.*

Our next lemma is a consequence of the second assertion of Lemma 4.4.

Lemma 4.6. *Let G be a minor-minimal graph with $\varphi(G) > k$ and let $\sum c(e)x_e \geqslant k$ be a witness for $\varphi(G) > k$. Then every edge uv incident to a degree-2 vertex v satisfies $c(uv) = k/2$.*

Proof. Let v be a vertex of degree 2 in G with neighbors u and w. We know that the weight of a minimum cut in (G, c) is k. Hence we get $c(\delta(v)) = c(uv) + c(vw) \geqslant k$. Furthermore, $c(uv) \leqslant k/2$ and $c(vw) \leqslant k/2$ by Lemma 4.4. Hence, $c(uv) = c(vw) = k/2$. The situation is shown in Figure 4.2. □

We finish this section with several results about 1-separations. Recall that a *k-separation* of a graph G is an ordered pair (G_1, G_2) of edge-disjoint subgraphs of G with $G = G_1 \cup G_2$, $|V(G_1) \cap V(G_2)| = k$, and $E(G_1)$, $E(G_2)$, $V(G_2) \setminus V(G_1)$, $V(G_1) \setminus V(G_2)$ all non-empty.

If (G_1, G_2) is a 1-separation of G and $V(G_1) \cap V(G_2) = \{v\}$ we write $G = G_1 +_v G_2$. The following lemma is taken from [26]. We immediately derive two quick corollaries from it that we use later.

Lemma 4.7 (Remark 9, [26])**.** *Let G be a graph such that $G = G_1 +_v G_2$.*

Let k be a fixed positive integer. Let

$$\sum_{e\in E(G_1)} c_i^1(e)x_e \geqslant k \qquad \textit{for } i \in I$$
$$x_e \geqslant 0 \qquad \textit{for } e \in E(G_1)$$

and

$$\sum_{e\in E(G_2)} c_j^2(e)x_e \geqslant k \qquad \textit{for } j \in J$$
$$x_e \geqslant 0 \qquad \textit{for } e \in E(G_2)$$

be irredundant systems of inequalities describing cutdom(G_1) *and* cutdom(G_2) *respectively. Then the following system of inequalities provides an irredundant description of* cutdom(G)*.*

$$\sum_{e\in E(G_1)} c_i^1(e)x_e + \sum_{e\in E(G_2)} c_j^2(e)x_e \geqslant k \qquad \textit{for } i \in I, j \in J$$
$$x_e \geqslant 0 \qquad \textit{for } e \in E(G)$$

Observe that the crucial point of this lemma is that two non-trivial facet-defining inequalities for cutdom(G_1) and cutdom(G_2), respectively, with the same right-hand side k can be combined to form a facet-defining inequality of cutdom(G) with right-hand side k, and that all non-trivial facets of cutdom(G) can be obtained in that way.

Corollary 4.8. *Let $G = G_1 +_v G_2$ be a graph. If $\varphi(G_1) \leqslant 1$ and $\varphi(G_2) \geqslant 1$, or $\varphi(G_1) \leqslant 2$ and $\varphi(G_2) \geqslant 2$, then $\varphi(G) = \varphi(G_2)$.*

Proof. We treat the case where $\varphi(G_1) \leqslant 2$ and $\varphi(G_2) \geqslant 2$. The other case is easier and left to the reader.

First, notice that $\varphi(G) \geqslant \varphi(G_2)$ since G_2 is a minor of $G = G_1 +_v G_2$.

Consider a facet-defining inequality $\sum_{e\in E(G_2)} c_j^2(e)x_e \geqslant k$ for cutdom(G_2) in minimum integer form with $k \geqslant 1$. Hence, we have $k \leqslant \varphi(G_2)$. Let $\sum_{e\in E(G_1)} c_i^1(e)x_e \geqslant k$ be a non-trivial facet-defining inequality for cutdom(G_1). Observe that $c_i^1(e) \in \{0, k/2, k\}$ for all $e \in E(G_1)$ because

$\varphi(G_1) \leqslant 2$. Furthermore, if k is even all coefficients of this inequality are integer. By Lemma 4.7, the inequality $\sum_{e\in E(G_1)} c_i^1(e)x_e + \sum_{e\in E(G_2)} c_j^2(e)x_e \geqslant k$ is facet-defining for cutdom(G). Observe that if $k \geqslant 2$, then k is even by Lemma 4.3 and the coefficients $c_i^1(e)$ and $c_j^2(e)$ are integer for all $e \in E(G)$. Otherwise, the coefficients $c_i^1(e)$ are half-integral and the right-hand side of the equation written in minimum integer form is at most 2.

As every facet-defining inequality for cutdom(G) can be obtained in that way, it follows that the right-hand side k of any inequality defining a facet for cutdom(G) written in minimum integer form is $k \leqslant \varphi(G_2)$ or $k \leqslant 2$. As $\varphi(G_2) \geqslant 2$, we get $\varphi(G) \leqslant \varphi(G_2)$. □

Lemma 4.9. *Let G be a minor-minimal graph with $\varphi(G) > k$, where $k \geqslant 1$. Then the minimum degree of G is at least 2.*

Proof. Assume by contradiction that G has a vertex u with degree 1. Observe that the neighbor v of u is a cutvertex of G. Hence, we can write $G = G_1 +_v G_2$, where $G_1 = (\{u,v\},\{uv\})$ and $G_2 = G - u$. Observe that $\varphi(G_1) = 1$, which implies $\varphi(G) = \varphi(G_2)$ by Corollary 4.8. □

4.2 Bounding $\varphi(G)$ as a function of $\tau(G)$

Bounding a new parameter as a function of a known parameter is a popular approach when trying to understand the structure of graphs. We show that we can bound $\varphi(G)$ as a function of the number of vertices or the vertex cover number $\tau(G)$. It is known that $\varphi(G)$ cannot be bounded as a function of treewidth (see [26]). Indeed, Conforti et al. give a construction of such a family of graphs that has unbounded $\varphi(G)$-value and constant treewidth.

Lemma 4.10. *There exists a constant c_1 such that such that $\varphi(G) \leqslant 2^{c_1 n \log n}$ for all n-vertex graphs G.*

Proof. Note that it is sufficient to show the bound if G is a minor-minimal graph with fixed $\varphi(G)$. Let $\sum_{e\in E(G)} c(e)x_e \geqslant k$ be a facet-defining inequality of cutdom(G) in minimum integer form such that $k = \varphi(G)$. We know by [26, Lemma 14] that $0 < c(e) \leqslant \varphi(G)/2$ for all $e \in E(G)$. Hence, $\sum c(e)x_e \geqslant k$ defines also a facet of the convex hull of non-empty cuts in G, which is a 0/1-polytope of dimension $m = |E(G)|$. By [82, Corollary 26], we can bound the largest integer coefficient in $\sum c(e)x_e \geqslant k$ by $\frac{m^{m/2}}{2^{m-1}}$. Since

$m \leqslant 2n-3$ by Lemma 4.5, and since the largest coefficient in the inequality is $\varphi(G)$, we get $\varphi(G) \leqslant 2^{c_1 n \log n}$. □

Theorem 1.4. There exists a constant c_2 such that, letting $g : \mathbb{N} \to \mathbb{R}$ denote the function $g(x) = 2^{c_2 x \log x}$ we have $\varphi(G) \leqslant g(\tau(G))$ for all graphs G.

Proof. We may assume that G is a minor-minimal graph with fixed $\varphi(G)$. Indeed, $\tau(G)$ is a minor-monotone parameter and the function we consider is non-decreasing. Hence, if there exists a minor H of G with $\varphi(H) = \varphi(G)$, we obtain $\varphi(G) = \varphi(H) \leqslant g(\tau(H)) \leqslant g(\tau(G))$.

Observe that if $\varphi(G) \leqslant 1$, then we have $\varphi(G) \leqslant \tau(G)$. Hence we may assume $\varphi(G) \geqslant 2$.

Let $k = \varphi(G) - 2$ and let $\sum_{e \in E(G)} c(e)x_e \geqslant k$ be a witness for $\varphi(G) > k$. Let $\mathcal{F}$ be a family of vertex subsets defining $|E(G)|$ linearly independent minimum cuts $\{\delta(S) \mid S \in \mathcal{F}\}$ of (G, c). Let X be a vertex cover of G such that $|X| = \tau(G) = x$. Let $Y = V \setminus X$ and $y = |Y|$. Let $n = |V(G)|$. Clearly, $n = x + y$.

By Lemma 4.9, G has minimum degree at least 2. Let $Y_2 = \{v \in Y \mid \deg(v) = 2\}$, $y_2 = |Y_2|$, and $Y_{\geqslant 3} = Y \setminus Y_2 = \{v \in Y \mid \deg(v) \geqslant 3\}$, $y_{\geqslant 3} = |Y_{\geqslant 3}|$. The total number of vertices in the graph G is $n = x + y_2 + y_{\geqslant 3}$. We want to bound n as a function of x and apply Lemma 4.10 in a second step. For this, we bound $y_{\geqslant 3}$ and y_2 separately in this order.

By Lemma 4.5, $|\mathcal{F}| = |E(G)| \leqslant 2n - 3 = 2(x + y_2 + y_{\geqslant 3}) - 3$. Furthermore we know $2y_2 + 3y_{\geqslant 3} \leqslant |E(G)|$. This implies $y_{\geqslant 3} \leqslant 2x - 3$.

To bound y_2 in terms of x consider the graph H with vertex set X, that has one edge with endpoints u and v for each vertex $w \in Y_2$ whose neighbors in G are u and v. Notice that $u, v \in X$ since X is a vertex cover. Notice also that $c(uw) = c(uv) = k/2$ by Lemma 4.6, since w has degree 2 in G.

We claim that H is a forest of cacti, or equivalently, that every block of H is an edge or a cycle. Toward a contradiction, assume that H has two vertices u and v that are linked by three internally disjoint paths Q_1, Q_2, Q_3 in H.

For $i \in [3]$, let P_i denote the u–v path of G that corresponds to Q_i. Thus P_i has twice as many edges as P_i and every other vertex of P_i belongs to Y_2. Notice that every minimum cut $\delta(S)$ containing one edge of P_i contains exactly two edges of the same path P_i and no further edge. Indeed, if a cut contains exactly one edge of P_i, then it necessarily also contains an edge from the other two paths. As all these edges have weight $k/2$, such a cut is not minimal.

Let o be an arbitrary internal vertex of P_3. By Theorem 4.2, we may assume that the family $\mathcal{F}$ is laminar. Furthermore, we may also assume that $o \notin S$ for each $S \in \mathcal{F}$ (replacing each set by its complement, when necessary). Let $\mathcal{F}_1$ denote the subfamily of $\mathcal{F}$ consisting of all sets S such that $\delta(S)$ contains some edge of P_1. Observe that each $S \in \mathcal{F}_1$ is contained in the set of internal vertices of P_1.

Let $p_1 = 2q_1$ denote the number of edges of P_1. Observe that there are at least p_1 sets in $\mathcal{F}_1$, since otherwise the cuts $\delta(S)$, $S \in \mathcal{F}$ do not form a basis of minimum cuts, that is, all minimum cuts are linearly independent and the set $\mathcal{F}$ is maximum. This implies that $\mathcal{F}_1$ contains a set S with $level(S) > 0$, and hence a level-1 set $\{u', v'\}$. By Lemma 4.4, we conclude that P_1 has at least three consecutive degree-2 vertices, which contradicts the minimality of G. Indeed, it is easily seen that contracting an edge whose ends are degree-2 vertices can be contracted while keeping a witness.

It is an easy exercise to show that the number of edges in a forest of cacti is at most twice the number of its vertices. Hence, $|E(H)| \leqslant 2|V(H)|$ and $y_2 \leqslant 2x$. This leads to $n = x + y_2 + y_{\geqslant 3} \leqslant x + 2x + (2x - 3) \leqslant 5x$. Finally, by Lemma 4.10 we get that there exist constants c_1 and c_2 such that $\varphi(G) \leqslant 2^{c_1 5x \log(5x)} \leqslant 2^{c_2 x \log x}$. □

4.3 Some minimal excluded minors for $\varphi(G) \leqslant 4$

In this section we present some minimal excluded minors for $\varphi(G) \leqslant 4$. It is possible to verify by hand for these graphs that the given weight function and minimum cuts satisfy the conditions of a witness. This shows $\varphi(G) \geqslant k$ for some k. In order to verify $\varphi(G) \leqslant k$, we computed the minimum linear description of cutdom(G) with the program `Panda` [58] and verified that the biggest right-hand side is $\varphi(G)$. Also, the minor-minimality has been checked by computing the minimum linear descriptions of the graphs obtained by deleting or contracting an edge.

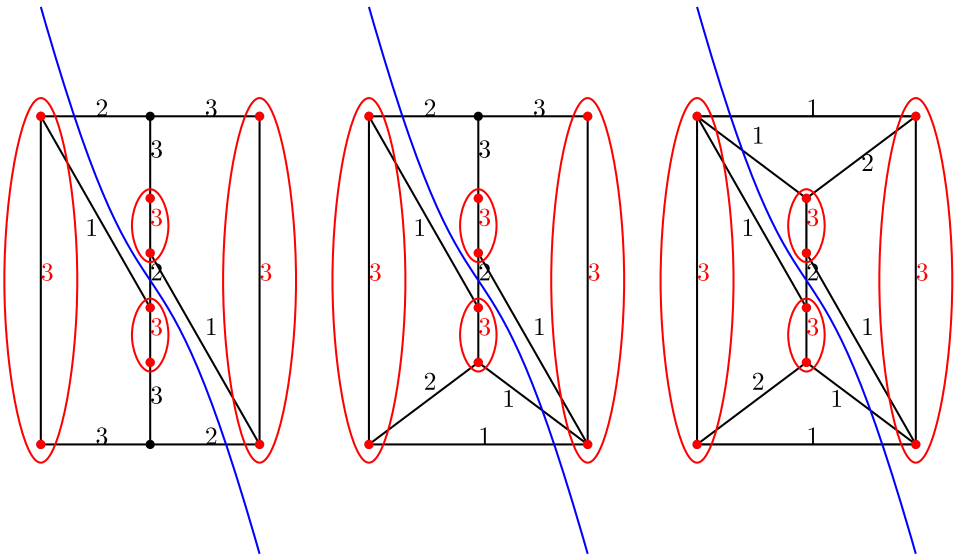

Figure 4.3. Three graphs that are known to satisfy $\varphi(G) = 6$. Level-0 sets are shown as red vertices, level-1 sets are red and cuts from level-2 sets are blue.

4.3.1 Internally 3-connected graphs

Recall that a graph is *internally* 3*-connected* if every 2-cutset is such that it separates exactly one vertex from all other vertices of the graph.

The first three graphs that we consider were already known to be minor-minimal graphs with $\varphi(G) > 4$ [25]. Cecchetto also mentionned these graphs in her Master thesis [19]. The three graphs G are shown in Figure 4.3 together with a witness for $\varphi(G) > 4$ and a family defining $|E(G)|$ minimum cuts that show $\varphi(G) \geqslant 6$.

Observe that these graphs can be obtained from one another by Δ–to–Y operations and all have the same number of edges and the same structure of minimum cuts. Cecchetto proved the following result concerning Δ–to–Y operations.

Lemma 4.11. *[Proposition 3.9.1 in [19]] Let $G = (V, E)$ be a minor-minimal graph with $\varphi(G) > k$ and $\mathcal{F}$ be a laminar family such that $\{\delta(S) | S \in \mathcal{F}\}$ is a basis of minimum cuts in (G, c) for some witness $\sum_{e \in E(G)} c(e)x_e \geqslant k$. Suppose there is a 3-cycle $C = \{v_1, v_2, v_3\}$ such that $v_1 \in S_1$, $v_2 \in S_2$, $v_3 \in S_3$, with $S_1, S_2, S_3 \in \mathcal{F}$ and the three sets S_1, S_2, S_3 form a partition*

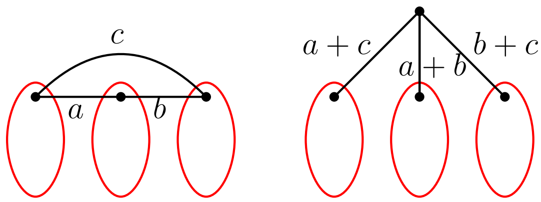

Figure 4.4. The Δ–to–Y operation preserves linearly independent minimum cuts.

of V. Let G' be the graph obtained from G by replacing the 3-cycle C with a claw as in Figure 4.4. Then the linearly independent minimum cuts in G correspond to linearly independent minimum cuts in G'.

This lemma suggests that a minor-minimal graph G can sometimes be transformed into another minor-minimal graph G' by Δ–to–Y operations. However, as observed by Cecchetto, Δ–to–Y operations could potentially transform a witness to a non-witness because of divisibility issues. We remark that the behavior of Y–to–Δ operations in graphs without further conditions can be more complicated, since no graph in the Petersen family besides the Petersen graph itself has $\varphi(G) > 4$.

The Petersen graph has some other interesting properties related to the cut dominant. It is the only known minor-minimal graph with $\varphi(G) > 4$ which is non-planar and has more than one witness. (Notice however that the different witnesses are images of a unique witness by the automorphism group of the Petersen graph.) This contradicts the first part of Conjecture 3.10.1 in [19] stating that there exists only one witness for every minor-minimal graph G with $\varphi(G) > k$, where $k \geqslant 4$ is arbitrary. The second part of the conjecture stating that $\varphi(G) = k + 2$, for every such graph, is disproved by the graphs in Section 4.3.2.

The Petersen graph is shown in Figure 4.5 with the weights given by a witness as well as the corresponding laminar family defining 15 linearly independent minimum cuts.

4.3.2 Not internally 3-connected graphs

Besides the graphs in Figures 4.3 and 4.5, we know eight more graphs. They are shown in Figure 4.6. Their structure is very different from that of the previous graphs. Indeed, these graphs can be obtained as 2-sums of the

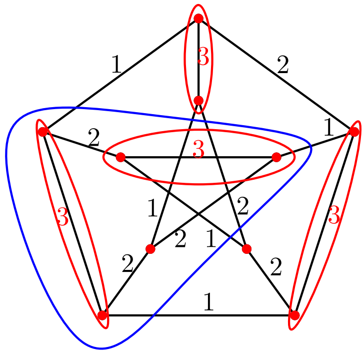

Figure 4.5. The Petersen graph G satisfies $\varphi(G) = 6$. All vertices are level-0 cuts, the level-1 cuts are in red, and the level-2 cut is in blue.

prism or pyramid graph and some other graph that we call an amplifier. The minimal excluded minors that we obtain all satisfy $\varphi(G) = 8$.

We will prove in Section 4.5 that the four amplifiers are such that they increase $\varphi(G)$ at least by a factor 2 when glued along an *odd edge* of a minor-minimal graph H with $\varphi(H) > k$. An odd edge has odd weight in an inequality, which in minimum integer form has right-hand side $\varphi(H)$, see Theorem 4.24.

4.4 Properties of minimal excluded minors for $\varphi(G) \leqslant 4$

Recall that we say that a facet of cutdom(G) is *trivial* if it is defined by a non-negativity inequality $x_e \geqslant 0$ for some $e \in E(G)$.

Lemma 4.12. *A graph G satisfies $\varphi(G) \leqslant 4$ if and only if every non-trivial facet of* cutdom(G) *can be defined by a (unique) inequality of the form $\sum c(e)x_e \geqslant 4$, where $c \in \mathbb{N}^{E(G)}$.*

Proof. The "if" part is obvious. We prove the "only if" part. Suppose that G is a graph with $\varphi(G) \leqslant 4$ and let F be a non-trivial facet of cutdom(G). Let $\sum_{e \in E(G)} c(e)x_e \geqslant k$ denote the inequality in minimum integer form that defines F. By Conforti, Rinaldi and Wolsey [27], $k \in \{1, 2, 4\}$. Hence, $4/k$ is integer and we can multiply the inequality by $4/k$ in order to give it the

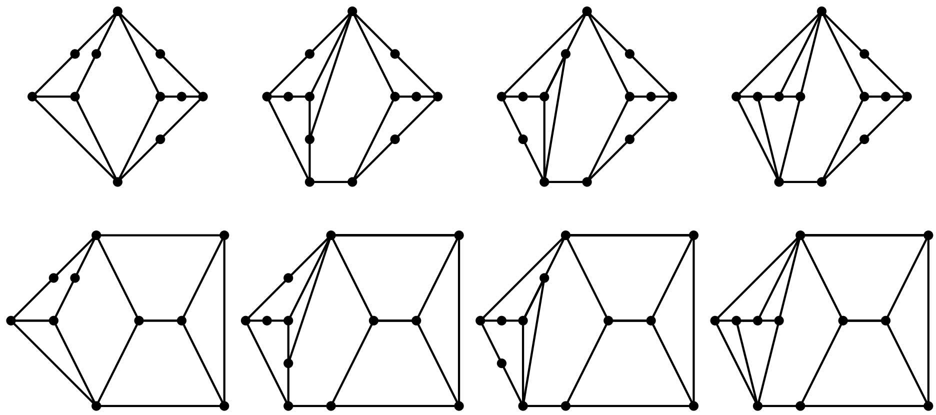

Figure 4.6. The eight known minimal excluded minors satisfying $\varphi(G) = 8$.

desired form. The resulting inequality still defines F. □

Combining Lemma 4.12 with results of Conforti, Fiorini and Pashkovich, see Lemma 4.3, we obtain the following corollary.

Lemma 4.13. *Every minor-minimal graph with $\varphi(G) > 4$ is simple and 2-connected.*

Proof. Let G be a minor-minimal graph with $\varphi(G) > 4$. [26, Remark 7] says that a minor-minimal graph contains no loops and [26, Remark 8] excludes parallel edges. Hence G is simple.

Now, suppose that G is a minor-minimal graph with a cutvertex v such that $G = G_1 +_v G_2$. By minor-minimality of G we have $\varphi(G_1) \leqslant 4$ and $\varphi(G_2) \leqslant 4$. It follows directly from Lemmas 4.7 and 4.12 that $\varphi(G) \leqslant 4$, a contradiction.

□

In the following of this section, we will focus on 2-connected graphs that have 2-separations. Recall that a 2-*separation* of a graph G is an ordered pair (G_1, G_2) of edge-disjoint subgraphs of G with $G = G_1 \cup G_2$, $|V(G_1) \cap V(G_2)| = 2$, and $E(G_1)$, $E(G_2)$, $V(G_2) \setminus V(G_1)$, $V(G_1) \setminus V(G_2)$ all non-empty.

The goal of this section is to show that minor-minimal graphs G with $\varphi(G) > 4$ that are not internally 3-connected satisfy $\varphi(G) = 8$. This contradicts the Conjecture 3.10.1 in [19] that all excluded minors for $\varphi(G) \leqslant 4$ satisfy $\varphi(G) = 6$.

Observe that non-internally 3-connected graphs admit a 2-separation (G_1, G_2) such that both G_1 and G_2 have at least four vertices. Indeed, if all 2-separations (G_1, G_2) in a graph G are such that one of G_1 or G_2 is a path of length 2, then the graph G is internally 3-connected.

Let G be a graph that has a 2-*separation* (G_1, G_2). For $i \in [2]$, we let $G'_i := G_i + e_i$ be the graph G_i with the edge $e'_i = uv$ added. If u and v are adjacent in $V(G_i)$, then we add a parallel edge $e'_i = uv$.

Let $\sum_{e\in E(G)} c(e)x_e \geqslant k$ be a witness for a minor-minimal graph G with $\varphi(G) > k$. For $i \in [2]$, we define c_i to be the restriction of c to $E(G_i)$. Let λ_i be the minimum weight cut of an u–v cut in (G_i, c_i). We define the weight function c'_i on $E(G'_i)$ such that $c'_i(e) = c_i(e)$ for all $e \in E(G_i)$ and $c'_i(e'_i) = \lambda_{3-i}$. We say that (G_1, G_2) is a 2-*separation of type* (λ_1, λ_2) of the graph (G, c). We may also talk of (λ_1, λ_2)-*separation* if the graphs G_1 and G_2 are clear from the context.

Our first result can be directly adapted from the proof of the "2-cutset lemma" of Conforti et al. [26, Lemma 20] by replacing the right-hand side 2 in their paper by 4. We do not include the adapted proof as it does not have any new tools or ideas.

Lemma 4.14. *Let G be a minor-minimal graph with $\varphi(G) > 4$ with a 2-separation (G_1, G_2) such that $V(G_1) \cap V(G_2) = \{u, v\}$, and let $\sum_{e\in E(G)} c(e)x_e \geqslant 4$ be a witness for $\varphi(G) > 4$. For $i \in [2]$, we define $c_i \in \mathbb{Q}_+^{E(G_i)}$, $G'_i := G_i + e'_i$ and $c'_i \in \mathbb{Q}_+^{E(G'_i)}$ as above. The following properties hold:*

(1) There exists a minimum cut of (G, c) that separates u and v.

(2) Up to exchanging G_1 and G_2, we may assume that c_1 is non-integer and c_2 is integer. Then, $\sum_{e\in E(G'_1)} c'_1(e)x_e \geqslant 4$ defines a ridge of cutdom(G'_1) *and $\sum_{e\in E(G'_1)} c'_2(e)x_e \geqslant 4$ defines a facet of* cutdom(G'_2).

(3) The vertices u and v are not adjacent in G.

Notice that our statement of Lemma 4.14 is not exactly the same as Lemma 20 in [26]. In particular, they do not list the properties that we want to show in the statement. However, in the proof of [26, Lemma 20], Conforti et al. show the properties of Lemma 4.14 for minor-minimal graphs with $\varphi(G) > 2$ before showing their statements. Assertion (i) of [26, Lemma 20] follows from Lemma 4.19, which we prove later.

When we talk about a 2-separation (G_1, G_2) of type (λ_1, λ_2) such that the corresponding witness has right-hand side k, we will always assume that some weights c_1 on G_1 are fractional while the weights c_2 on G_2 are all integer. It follows from the previous lemmas that λ_1 and λ_2 are non-zero positive integers with $\lambda_1 + \lambda_2 = 4$. Observe that if $\lambda_1 = 0$ or $\lambda_2 = 0$ then u or v is a cutvertex of G, contradicting Lemma 4.13. Thus, there are three possible values (λ_1, λ_2) for $k = 4$, namely $(1, 3)$,$(2, 2)$, and $(3, 1)$.

The 2-separations of type $(2, 2)$ turn out to be much easier to handle that the other types of separations. Lemma 4.15 shows that assertion (ii) of [26, Lemma 20] holds for $(2, 2)$-separations. Again, we follow the proof of [26, Lemma 20] very closely.

Lemma 4.15. *Let G be a minor-minimal graph with $\varphi(G) > 4$, and let $\sum_{e \in E(G)} c(e)x_e \geqslant 4$ be a witness for $\varphi(G) > 4$. If G has a 2-separation (G_1, G_2) with $V(G_1) \cap V(G_2) = \{u, v\}$ of type $(2, 2)$, then the integer side (G_2, c_2) is a path uwv of length 2. Moreover, $c(uw) = c(wv) = 2$.*

Proof. Recall that c_i is the restriction of c to $E(G_i)$ for $i \in [2]$, λ_i is the minimum weight cut of an u–v cut in (G_i, c_i), and $c'_i \in \mathbb{Q}_+^{E'_i}$ is such that $c'_i(e) = c_i(e)$ for all $e \in E(G_i)$ and $c'_i(e'_i) = \lambda_{3-i}$. Let $\mathcal{F}$ be a family defining $E(G)$ linearly independent minimum cuts in (G, c).

Let $\delta(S^*)$ be a fixed u–v cut such that each u–v cut satisfies (4.1).

$$\delta(S) \cap E_1 = \delta(S^*) \cap E_1 \text{ or } \delta(S) \cap E_2 = \delta(S^*) \cap E_2. \tag{4.1}$$

We let M be the non-singular matrix whose rows are the characteristic vectors of the cuts $\delta(S)$ for each $S \in \mathcal{F}$. For $i \in [2]$, let M_i be the submatrix of M induced by the rows whose intersection with $E(G_{3-i})$ is either empty or equal to $\delta(S^*) \cap E(G_{3-i})$. Observe that M_1 and M_2 have full row-rank since they are row-induced submatrices of M. Notice that they only have one row in common, namely the one of the cut $\delta(S^*)$. Thus, $\text{rk}(M_1) + \text{rk}(M_2) = \text{rk}(M) + 1 = |E(G)| + 1 = |E(G_1)| + |E(G_2)| + 1$.

For $i \in [2]$, we define the column vector $\xi_{3-i} \in \{0, 1\}$ as follows. ξ_{3-i} has one entry for each row of M_i, a 1 for entries corresponding tu uv-cuts, and a

0 otherwise. For each $e \in E(G_{3-i})$, the column in M_i indexed by e is equal to ξ_i if $e \in \delta(S^*)$ and the zero vector otherwise. Removing all columns corresponding to edges of $E(G_{3-i})$ from M_i and adding a single copy of the column ξ_{3-i} indexed by the edge e'_i results in a matrix M'_i. Notice that each row of M'_i corresponds to a cut of G'_i, and that each of these cuts is minimum with respect to c'_i, see (4.1).

Suppose that G_2 has more than three vertices. Consider the graph H with vertices $V(G_1) \cup \{w\}$ and edges $E(G_1) \cup \{uw, wv\}$. Notice that H is a proper minor of G because it can be obtained by contracting $E(G_2 \setminus \{u, v\})$ to a single vertex. We define c_H such that $c_H(e) = c(e)$ if $e \in E(G_1)$ and $c(uw) = c(wv) = 2$. The inequality $\sum_{e \in E(H)} c_H(e) x_e \geqslant 4$ is valid for cutdom(H) by [26, Lemma 11]. Furthermore, because $c(\delta_G(S^*) \cap E(G_1)) = 2$, the cut $\delta_H(S^* \cap V(G_1))$ is a minimum cut in H with respect to c_H. Let M_H be the matrix obtained from M'_1 by reindexing the column of e'_2 by uw, adding a new column indexed by wv, and adding two lines. The first line corresponds to the cut $\delta_H((S^* \cap V(G_1)) \cup \{w\})$ and the second line corresponds to the cut $\delta_H(w) = \{uw, wv\}$. Note that both cuts are minimum with respect to c_H. We leave it to the reader to check that $\text{rk}(M_H) = \text{rk}(M_1) + 2 = |E_1| + 2$. Hence, the rows define $|E_1|+2$ linearly independent minimum cuts of $(H, c_H$. Thus $\sum_{e \in E(H)} c_H(e) x_e \geqslant k$ defines a facet of cutdom(H). The vector c_H is integral because of the minor-minimality of G. Since c_2 is integral as well, it follows that c is an integral vector, a contradiction.

Thus, $|G_2| = 3$, that is V_2 consists of three vertices u, v, w and edges uw, wv. Since w is a degree-2 vertex, we have $c(uw) = c(wv) = 2$ by Lemma 4.6. □

It is possible to generalize the next lemma to minor-minimal graphs with $\varphi(G) > k$. However, we only include a proof if $k = 4$.

Lemma 4.16. *Let G be a minor-minimal graph with $\varphi(G) > 4$. Then no two degree-2 vertices are adjacent in G.*

Proof. Let v_1 and v_2 be two adjacent degree-2 vertices. If v_1 and v_2 have a common neighbor v_0, then the graph G is of the form $G_0 +_{v_0} K_3$, where G_0 is the graph obtained from G by deleting v_1 and v_2. By Corollary 4.8 we get $\varphi(G) = \varphi(G_0)$, which contradicts the minor-minimality of G.

Thus, we may assume that v_0 and v_3 are distinct neighbors of v_1 and v_2, respectively. Note that $\{v_0, v_3\}$ is a 2-cutset of G. We have $G = G_1 \cup G_2$, where $G_1 = G - \{v_1, v_2\}$ and G_2 is the path $v_0 v_1 v_2 v_3$. Note that G_1 has

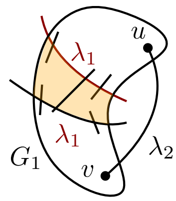

Figure 4.7. The filled orange area shows $S_1 \setminus S_2$. Observe that $c(\delta(S_1 \setminus S_2)) \leqslant 2\lambda_1$.

at least four vertices because $\varphi(G) \geqslant 4$ implies that G has at least six vertices. Let $\sum c(e)x_e \geqslant 4$ be a witness of $\varphi(G) > 4$. By Lemma 4.15, we have $c(v_0v_1) = c(v_1v_2) = c(v_2v_3) = 2$. This implies that $\{v_0, v_3\}$ is a $(2,2)$-separator for (G_1, G_2). By Lemma 4.15 either G_1 or G_2 has exactly 3 vertices, a contradiction. □

The following lemma can be generalized for larger values of k if the minor-minimal graph G with $\varphi(G) > k$ is such that, given a witness for $\varphi(G) > k$, there exists a (λ_1, λ_2)-separation with $\lambda_1 \neq \lambda_2$ and $\lambda_1 + \lambda_2 = k$.

Lemma 4.17. *Let G be a minor-minimal graph with $\varphi(G) > 4$ such that (G_1, G_2) is a (λ_1, λ_2)-separation in (G, c), where $\sum_{e \in E(G)} c(e)x_e \geqslant 4$ is a witness for $\varphi(G) > 4$. If $\lambda_1 < \lambda_2$, then there exists a unique minimum u–v cut in (G_1, c_1). If $\lambda_1 > \lambda_2$, then there exists a unique minimum u–v cut in (G_2, c_2).*

Proof. It is sufficient to prove the statement if $\lambda_1 < \lambda_2$ as the argument is symmetric and does not depend on integrality of the edges. Recall that $\lambda_1 + \lambda_2 = 4$. By contradiction, let $\delta(S_1)$ and $\delta(S_2)$ be two u–v cuts in (G_1, c_1). We may assume $\emptyset \neq S_1 \setminus S_2$ because we can exchange S_1 and S_2. Notice that $\delta(S_1 \setminus S_2)$ is a cut in (G_1, c_1) and (G, c) because $S_1 \setminus S_2$ is non-empty. Furthermore, $c(\delta(S_1 \setminus S_2)) \leqslant c(\delta(S_1)) + c(\delta(S_2)) = \lambda_1 + \lambda_1 < \lambda_1 + \lambda_2 = k$. This contradicts that a minimum cut in (G, c) has weight 4. The situation is illustrated in Figure 4.7. □

The next lemma implies in particular that gluing three graphs on a K_3 graph and deleting all edges of K_3 does not result in a graph G that is minor-minimal with $\varphi(G) > 4$. After that, we show that any 2-cutset $\{u, v\}$ in a minor-minimal graph with $\varphi(G) > 4$ is such that $G - \{u, v\}$ has exactly two connected components.

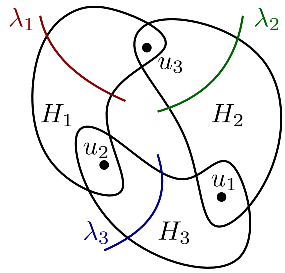

Figure 4.8. The graph G is has three 2-cutsets.

Lemma 4.18. *Let G be a minor-minimal graph with $\varphi(G) > 4$. There exists no triple of vertices $\{u_1, u_2, u_3\}$ such that any two of them form a 2-cutset of G and $G - \{u_1, u_2, u_3\}$ has exactly three connected components.*

Proof. Suppose that we can write $G = H_1 \cup H_2 \cup H_3$ with $V(H_i) \cap V(H_j) = \{u_k\}$ for all distinct indices $i, j, k \in [3]$, $E(G) = E(H_1) \cup E(H_2) \cup E(H_3)$, and $E(H_i) \cap E(H_j) = \emptyset$ for $i, j \in [3]$. The graphs H_1, H_2, H_3 are shown in Figure 4.8.

Let $\sum_{e \in E(G)} c(e)x_e \geqslant 4$ be a witness for $\varphi(G) > 4$. Let c_i be the restriction of c to $E(H_i)$ for $i \in [3]$. By applying Lemma 4.14 to the 2-separations $(H_1, H_2 \cup H_3)$, $(H_2, H_1 \cup H_3)$, and $(H_3, H_1 \cup H_2)$ we get that exactly one of c_1, c_2, c_3 is fractional. We may assume that c_1 is non-integer and that c_2 and c_3 are integer.

For distinct indices $i, j, k \in [3]$, let λ_i be the weight of a minimum u_j–u_k cut in (H_i, c_i). Note that $\lambda_i \in \{1, 2, 3\}$ because if $\lambda_i = 0$, H_i is disconnected, and if $\lambda_i = 4$ then there is no minimum cut separating u_j and u_k, contradicting Lemma 4.14. Notice also that $\lambda_i + \lambda_j \geqslant 4$ for every $i \neq j \in [3]$. Otherwise $G - \{u_1, u_2, u_3\}$ consists of more than three components because the weight of a minimum cut is at least 4. We can also assume $\lambda_2 \geqslant \lambda_3$.

Case 1: $\lambda_1 = 1$. Observe that we have $\lambda_2 = \lambda_3 = 3$ because otherwise there exists a cut of weight strictly less than 4 in (G, c), a contradiction.

For $i = 2, 3$, let $H'_i = H_i \cup \{u_1 u_{5-i}\}$, $c'_i : E(H'_i) \to \mathbb{Q}_+$ be such that $c'_i(e) = c_i(e)$ for all $e \in E(H_i)$ and $c'_i(u_1 u_{5-i}) = 1$. Let $H_4 = H_2 \cup H_3$. Let $H'_4 = H_4 \cup \{u_2 u_3\}$. Let c_4 be the restriction of c to $E(H_4)$ and let $c'_4 : E(H'_4) \to \mathbb{R}$ be such that $c'_4(u_2 u_3) = 1$ and $c'_4(e) = c_4(e)$ for all $e \in E(H_4)$. Notice that $\sum c'_i(e)x_e \geqslant 4$ defines a facet of $\mathrm{cutdom}(H'_i)$ for

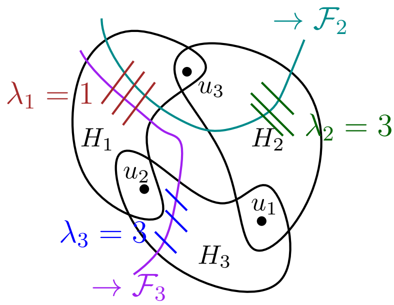

Figure 4.9. Two u_2–u_3 cuts that are in $\mathcal{F}_2$ and $\mathcal{F}_3$, respectively.

$i \in \{2, 3, 4\}$ by Lemma 4.14.

Let $\mathcal{F}$ be a family defining $|E(G)|$ linearly independent minimum cuts in (G, c). We partition $\mathcal{F}$ into sets $\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3$ in such a way that

- if $\delta(S) \subseteq E(H_i)$, then $S \in \mathcal{F}_i$, for $i \in [3]$;
- if $\delta(S)$ is a minimum u_1–u_3 cut, then $S \in \mathcal{F}_2$;
- if $\delta(S)$ is a minimum u_1–u_2 cut, then $S \in \mathcal{F}_3$.

Figure 4.9 illustrates an u_1–u_3 and am u_1–u_2 cut. Note that any u_2–u_3 cut is either an u_1–u_2 cut or an u_1–u_3 cut. Thus $\{\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3\}$ is such that every $S \in \mathcal{F}$ is contained in exactly one $\mathcal{F}_i$, $i \in [3]$.

Consider the 2-cutset $\{u_2, u_3\}$. Note that (H_1, H_4) is a $(1, 3)$-separation with cutset $\{u_2, u_3\}$. By Lemma 4.14 we may assume $|\mathcal{F}_1| = |E(H_1)|$ and $|\mathcal{F}_2 \cup \mathcal{F}_3| = |E(H_4)| + 1$. Indeed, this is because c_1' defines a ridge of $\operatorname{cutdom}(H_1')$ and c_4' defines a facet of $\operatorname{cutdom}(H_4')$. Furthermore, we have $|\mathcal{F}_i| \leqslant |E(H_i)| + 1$ for $i = 2, 3$ because $|E(H_4)| = |E(H_2)| + |E(H_3)|$ and by applying Lemma 4.14 to the cutsets $\{u_1, u_3\}$ and $\{u_1, u_2\}$, respectively. Thus we may assume $|\mathcal{F}_2| = |E(H_2)| + 1$ and $|\mathcal{F}_3| = |E(H_3)|$. We claim that the graph H obtained from G by contracting $E(H_3)$ to a single vertex $u = u_1 = u_3$ contradicts the minimality of G.

Let $c_H : E(H) \to \mathbb{R}$ be such that $c_H(e) = c(e)$ for every $e \in E(H)$. Let M_G be the matrix whose columns correspond to edges of G and the rows correspond to the cuts defined by $\mathcal{F}$. Let M_H be the matrix obtained from M_G by the following operations. We delete all rows corresponding to cuts defined by $\mathcal{F}_3$, and all columns corresponding to $E(H_3)$.

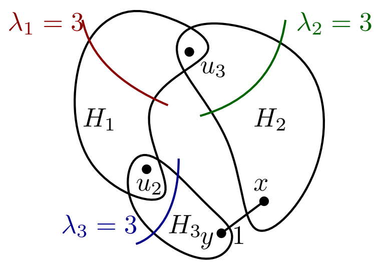

Figure 4.10. The situation in Case 3.1.

Notice that the matrix M_H has full rank. Furthermore, each row corresponds to a minimum cut in (H, c_H). Thus $\sum c_H(e)x_e \geqslant 4$ defines a facet of cutdom(H). Finally, observe that c_H is not integral because $E(H_1) \subseteq E(H)$ and for some edge $e^* \in E(H_1)$ we have that $c_H(e^*) = c(e^*)$ is fractional, which implies $\varphi(H) > 4$. This contradicts the minimality of G.

Case 2: $\lambda_1 = 2$. Then $(H_1, H_2 \cup H_3)$ is a $(2,2)$-separation, implying $G = H_1 +_{u_2u_3} K_3$ by Lemma 4.15. However, $H_2 \cup H_3$ contains at least 5 vertices because $\{u_1, u_2\}$ and $\{u_1, u_3\}$ are cutsets, a contradiction.

Case 3: $\lambda_1 = 3$. Either $\lambda_2 = 1$ or $\lambda_3 = 1$. We may assume $\lambda_3 = 1$. As c_3 is integer this implies that H_3 contains a bridge xy.

Case 3.1: $x, y \neq u_2$. We redefine H_2 and H_3 as in Figure 4.10. We want to show that the graph H obtained by contracting the edges of H_3 to a single vertex satisfies $\varphi(G') > 4$, contradicting the minimality of G.

Note that now $G = H_1 \cup H_2 \cup H_3 \cup \{xy\}$ and $\lambda_i = 3$ for each $i \in [3]$. Let $\mathcal{F}$ be a family defining $|E(G)|$ linearly independent minimum cuts in (G, c). We can partition $\mathcal{F}$ into $\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3$ such that $S \in \mathcal{F}_i$ if and only if $\delta(S) \subseteq E(H_i) \cup \{xy\}$.

As in Case 1, we can assume that $|\mathcal{F}_3| = |E(H_3)|$ by considering the 2-separations given by $\{u_1, u_3\}$, $\{u_1, u_2\}$ and $\{u_2, u_3\}$. By the same arguments as above, contracting all edges in H_3 results in a graph H with $\varphi(H) > 4$ that contradicts the minimality of G.

Case 3.2: $y = u_2$. Note that $\{u_1, x, u_3\}$ is a triple such that any two vertices define a 2-cutset of G and such that $\{x, u_3\}$ is a $(1,3)$-separator. Hence, by redefining H_1 and H_3 as in Figure 4.11 we may assume that we are in the situation $\lambda_1 = 1$. This concludes the proof.

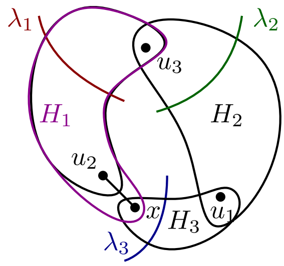

Figure 4.11. The situation in Case 3.2.

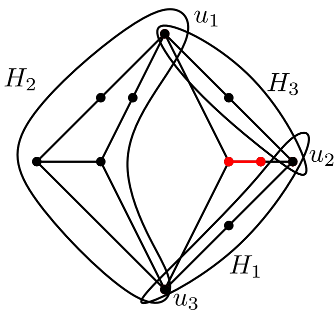

Figure 4.12. The red vertices are in none of H_1, H_2, H_3.

□

Observe that the condition that $G-\{u_1,u_2,u_3\}$ has exactly three connected components is necessary. Indeed, the graph minimal excluded graph G with $\varphi(G) > 4$ in Figure 4.12 has three vertices u_1, u_2, u_3 such that any two of them form 2-cutsets, but $G - \{u_1, u_2, u_3\}$ has four connected components.

Lemma 4.19. *Let G be a minor-minimal graph with $\varphi(G) > 4$. Let $\{u, v\}$ be a 2-cutset in G. Then $G - \{u, v\}$ has exactly two connected components.*

Proof. Assume that $G-\{u, v\}$ has r connected components, $r \neq 2$. We may assume $r \geqslant 3$. Let $C_1, \ldots, C_r$ be the connected components of $G - \{u, v\}$. Let G_i be the graph induced by $\{u, v\} \cup V(C_i)$ for every $i \in [r]$. Let G'_i be the graph G_i with the edge uv added.

Let $\sum c(e)x_e \geqslant 4$ be a witness for $\varphi(G) > 4$. Let c_i be the restriction of c to $E(G_i)$. Let λ_i be the value of a minimum u–v cut in (G_i, c_i). Let $c'_i : E(G'_i) \to \mathbb{Q}_+$ be such that $c'_i(e) = c_i(e)$ if $e \in E_i$ and $c'_i(uv) = \sum_{j \neq i} \lambda_j$.

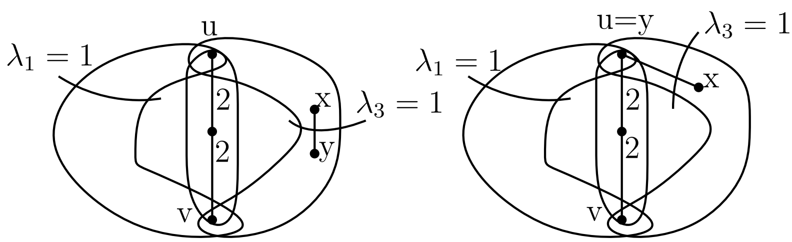

Figure 4.13. The bridge xy is such that $u = x$.

By Lemma 4.14, we may assume that

1. $r \leqslant 4$ because there exists a minimum u–v cut, $\sum_{i \in [r]} \lambda_i = 4$, $\lambda_i > 0$ and λ_i is integer for $i \in [r]$.

2. $c(e)$ is integer for all $e \in E(G) \setminus E(G_1)$

3. Some edge $e^* \in E(G_1)$ is such that $c(e^*)$ is fractional.

There are three cases we need to discuss.

Case 1: $r = 4$**.** Then $\lambda_i = 1$ for all $i \in [4]$. Let $H_1 = G_1 \cup G_2$ and $H_2 = G_3 \cup G_3$. Notice that both graphs H_1 and H_2 consist of at least four vertices because each G_i has at least one vertex distinct from $\{u, v\}$. Notice that (H_1, H_2) is a $(2, 2)$-separation. Thus, by Lemma 4.14 H_2 consists of 3 vertices, a contradiction.

Case 2: $r = 3$ **and** $\lambda_1 = 2$**.** Let $H_2 = G_2 \cup G_3$. Notice that (G_1, H_2) is a $(2, 2)$-separation of G. As before, Lemma 4.14 implies that H_2 has 3 vertices, contradicting that H_2 is the union of G_2 and G_3.

Case 3: $r = 3$ **and** $\lambda_1 = 1$**.** Suppose $\lambda_2 = 2$ and $\lambda_3 = 1$. By Lemma 4.15, G_2 consists of three vertices u, v, w and two edges uw and wv. Since $\lambda_3 = 1$ and $c_3(e)$ is integer for every $e \in E_3$, there exists a bridge xy in G_3 with $c_3(xy) = 1$, see Figure 4.13. Notice that by Lemma 4.18, one of x, y is u or v. We may assume that $y = u$.

Furthermore, there exists a unique set of edges $X \subseteq E(G_1)$ such that X separates u and v in G_1 and $c(X) = \sum_{e \in X} c(e) = 1$ by Lemma 4.17.

Let $\mathcal{F}$ be a family of sets defining $|E(G)|$ linearly independent minimum cuts $\{\delta(S) \mid S \in \mathcal{F}\}$ in (G, c). We may subdivide these cuts into the following three sets.

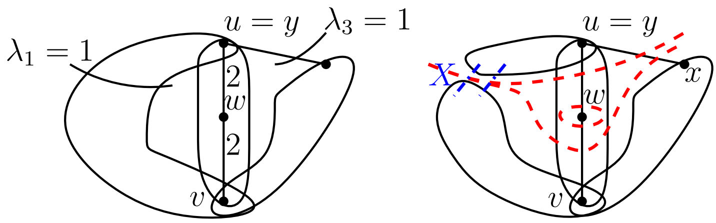

Figure 4.14. The cuts $\delta(S)$ with $S \in \mathcal{F}_2$ are shown in dashed red.

- $\mathcal{F}_1$ contains $S \in \mathcal{F}$ if $\delta(S) \subseteq E(G_1)$.
- $\mathcal{F}_2$ contains $S \in \mathcal{F}$ if $uw \in \delta(S)$ or $wv \in \delta(S)$, or both.
- $\mathcal{F}_3$ contains $S \in \mathcal{F}$ if $\delta(S) \subseteq E(G_3)$.

We may assume $\mathcal{F}_2 = \{S_1, S_2, S_3\}$ with $\delta(S_1) = \{X \cup \{uw, ux\}\}$, $\delta(S_1) = \{X \cup \{vw, ux\}\}$, and $\delta(S_1) = \{uw, vw\}$. The cuts $\delta(S)$ with $S \in \mathcal{F}_2$ are shown in Figure 4.14.

Let H be the graph obtained from G by contracting the edge ux to the vertex u and deleting the vertex w. Note that H is a minor of G and we may see the edges of H as a subset of the edges of G. Let $c_H : E(H) \to \mathbb{Q}_+$ be such that $c_H(e) = c(e)$ for every edge $e \in E(H)$. Observe that there exists some $e^* \in E(H)$ such that $c_H(e^*)$ is non-integer. That is because the edges of G_1 are all contained in H. We claim that c_H defines a facet of cutdom(H). If so, c_H witnesses that H is a graph with $\varphi(H) > 4$. Since H is a proper minor of G, H contradicts the minor-minimality of G. Thus Case 3 cannot occur either.

We need to prove that c_H defines a facet of cutdom(H). For this we show that there exists a family $\mathcal{H}$ of $|E(H)|$ linearly independent minimum cuts in (H, c_H). First, we describe $\mathcal{H}$, then we prove the linear independence by contradiction.

Let $\mathcal{H}_1 = \mathcal{F}_1$. Note that $\mathcal{H}_1$ is well-defined because G_1 is a induced subgraph of H. Now, for all $S \in \mathcal{F}_3$, we let $T_S \subseteq V(H)$ be such that in $\delta_H(T_S)$ the edge $ux \in \delta(S)$ is replaced by X if $ux \in \delta(S)$. That is, we add some vertices of H_1 to S to obtain T_S in this case. Note that all cuts in $\mathcal{H}_1$ and $\mathcal{H}_3$ are distinct, $|\mathcal{F}_1| = |\mathcal{H}_1|$, and $|\mathcal{F}_3| = |\mathcal{H}_3|$. We set $\mathcal{H} = \mathcal{H}_1 \cup \mathcal{H}_3$. As $|\mathcal{F}_2| = 3$, $\mathcal{H}$ contains $|E(G)| - 3 = |E(H)|$ cuts.

It remains to prove that the cuts in $\mathcal{H}$ are linearly independent. For this, we assume that they are linearly depend and show that this implies that the cuts in G are also linearly dependent, a contradiction. Suppose that the cuts $\{\delta(T) \mid T \in \mathcal{H}\}$ are linearly dependent. Recall that all cuts $\delta(T) \in \mathcal{H}_1$ contain no edge of G_3 and all cuts $\delta(T) \in \mathcal{H}_3$ contain either all edges of X (and no other edge $e \in E(G_1)$) or no edge of G_1 at all. Hence, the only edges that can be part of cuts from both $\mathcal{H}_1$ and $\mathcal{H}_3$ are in X.

Let $\alpha \in \mathbb{R}^{\mathcal{H}}$, $\alpha \neq 0$ such that $\sum_{T \in \mathcal{H}} \alpha_T \chi^{\delta(T)} = 0$. Notice that

$$\sum_{T \in \mathcal{H}_1} \alpha_T \chi^{\delta(T)} = - \sum_{T \in \mathcal{H}_3} \alpha_T \chi^{\delta(T)}.$$

As all cuts in $\mathcal{H}_3$ contain either all or none edges of X, we may assume

$$\sum_{T \in \mathcal{H}_1} \alpha_T \chi^{\delta(T)} = \chi^X$$

and

$$\sum_{T \in \mathcal{H}_3} \alpha_T \chi^{\delta(T)} = -\chi^X$$

by rescaling α. Indeed, if the rescaling is not possible we have

$$\sum_{T \in \mathcal{H}_1} \alpha_T \chi^{\delta(T)} = 0 \cdot \chi^X.$$

Then the incidence vectors of the cuts $\delta(T)$ defined by the sets $T \in \mathcal{H}_1$ are linearly dependent, which contracts linear independence of the cuts defined by $\mathcal{F}$ because $\mathcal{H}_1 = \mathcal{F}_1 \subseteq \mathcal{F}$.

We can use α to get coefficients $\beta \in \mathbb{R}^{\mathcal{F}}$ such that $\sum_{S \in \mathcal{F}} \beta_S \chi^{\delta(S)} = 0$, showing that c is not a facet-defining inequality of cutdom(G). Let $\beta \in \mathbb{R}^{\mathcal{F}}$ be such that

- $\beta_S = \alpha_{T_S}$ if $S \in \mathcal{F}_1$,
- $\beta_S = -\alpha_{T_S}$ if $S \in \mathcal{F}_3$,
- $\beta_{S_1} = -1/2$,
- $\beta_{S_2} = -1/2$,
- $\beta_{S_3} = 1/2$.

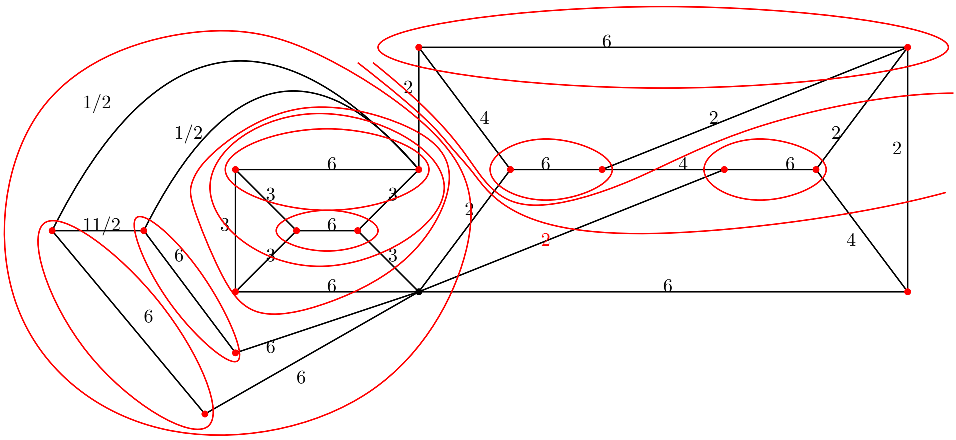

Figure 4.15. A minor-minimal graph with $\varphi(G) > 12$ that is obtained by gluing three graphs along an edge. The weights define a facet of cutdom(G) such that the red vertices and red sets form a family defining $|E(G)|$ linearly dependent minimum cuts.

Note that $\sum_{S\in\mathcal{F}_1}\beta_S\chi^{\delta(S)} = \chi^X$ and $\sum_{S\in\mathcal{F}_3}\beta_S\chi^{\delta(S)} = \chi^{\{ux\}}$ by definition of β. Hence, $\sum_{S\in\mathcal{F}}\beta_S\chi^{\delta(S)} = 0$. Since $\beta_{S_i} \neq 0$ for $i \in [3]$, this contradicts the linear independence of the cuts $\{\delta(S) \mid S \in \mathcal{F}\}$.

□

Lemma 4.19 cannot be generalized for general k because we can construct minor-minimal graphs with a 2-cutset giving three connected components, see Figure 4.15. We assume that it is even possible to construct minor-minimal graphs such that a 2-cutset can have any number of connected components.

A consequence of Lemmas 4.18 and 4.19 is that, given any 2-cutset $\{u, v\}$ in a minor-minimal graph G with $\varphi(G) > 4$, there exists exactly one 2-separation (G_1, G_2) such that $V(G_1) \cap V(G_2) = \{u, v\}$.

Before we move on to the next lemma, observe that by Lemmas 4.17 and 4.18, it follows that to every $(3, 1)$-separation in a minor-minimal graph with $\varphi(G) = 4$ can be associated one $(1, 3)$-separation as illustrated in Figure 4.11. Indeed, if the integer side has a cut of weight 1, then it must be a bridge that is incident to a vertex of the 2-cutset and we find a "nearby" 2-cutset defining a $(1, 3)$-separation.

We will now show some properties that the related graphs G'_1 and G'_2 of a 2-separation (G_1, G_2) of type $(1,3)$ need to satisfy. We start with an observation about G'_2.

Lemma 4.20. *Let G be a minor-minimal graph with $\varphi(G) > 4$. Let (G_1, G_2) be a separation of G of type $(1,3)$ or $(3,1)$. Then $\varphi(G'_2) = 4$.*

Proof. Observe that G'_2 is a minor of G. This implies $\varphi(G'_2) \leqslant 4$ as G is a minor-minimal graph with $\varphi(G) > 4$. By contradiction, assume $\varphi(G'_2) \leqslant 2$. Then $\sum_{e \in E(G'_2)} c'_2(e) x_e \geqslant 4$ defines a facet of $\text{cutdom}(G'_2)$ by Lemma 4.14. Hence, $c'_2(e) \in \{0, 2, 4\}$ for every edge $e \in E(G'_2)$ and thus λ_2 is even. This contradicts that (G_1, G_2) is a separation of type $(1,3)$ or $(3,1)$. □

Recall that a ridge R of a polyhedron is a face with dimension $d-2$, where d is the dimension of the polyhedron.

Lemma 4.21. *Let G be a graph. Let $uv \in E(G)$. Assume that there exists a ridge of* $\text{cutdom}(G)$ *defined with a unique inequality*

$$\sum_{e \in E(G)} c(e) x_e \geqslant k$$

such that $k \geqslant 2$, $c(uv) = k/2$ and $c(e) > 0$ for all edges $e \in E(G)$. Assume that there exists a minimum cut separating u and v in (G, c).

Let $H = G +_{uv} K_3$ with $V(K_3) = \{u, v, w\}$. Let $c_H : E(H) \to \mathbb{Q}_+$ be such that $c_H(e) = c(e)$ if $e \in E(G) \setminus \{uv\}$ and $c_H(uw) = c_H(wv) = k/2$. Then

$$\sum_{e \in E(H)} c_H(e) x_e \geqslant k$$

defines a facet of $\text{cutdom}(H)$*. Furthermore, if $k = 4$ and $c_H(e)$ is fractional for some $e \in E(H)$, then $\varphi(H) > 4$.*

Proof. First, observe that the weight of any cut in (H, c_H) is at least k because we replace an edge of (G, c) of weight $k/2$ by a path of two edges of weight $k/2$ in (H, c_H). Moreover, the minimum weight of a cut in (H, c_H) is k.

Furthermore, observe that we can write $c = \alpha f_1 + (1-\alpha) f_2$ for some $0 < \alpha < 1$, where f_1 and f_2 define facets of $\text{cutdom}(G)$ with $f_1(uv) < k/2$ and $f_2(uv) > k/2$. This is because we assume that there exists a unique inequality $\sum_{e \in E(G)} c(e) x_e \geqslant k$ defining the ridge such that $c(uv) = k/2$.

Let $\mathcal{F}_G$ be a family defining $|E(G)|-1$ linearly independent minimum cuts in (G,c). $\mathcal{F}_G$ exists by the assumption that $\sum_{e\in E(G)} c(e)x_e \geqslant k$ is a ridge. By assumption, there exists a minimum cut separating u and v in (G,c). Hence, we may assume that there exists $S^* \in \mathcal{F}_G$ such that $uv \in \delta(S^*)$.

Let $\mathcal{F}_H = \mathcal{F}_G \cup \{\{w\}, S^* \cup \{w\}\}$. Observe that $\mathcal{F}_H$ contains $|E(H)| = |E(G)|+1$ sets and that each corresponding cut is a minimum cut in (H, c_H). In order to show that $\sum_{e\in E(H)} c_H(e)x_e \geqslant k$ defines a facet of $\text{cutdom}(H)$, it is sufficient to show that $\{\delta(S) \mid S \in \mathcal{F}_H\}$ is a set of linearly independent minimum cuts.

Let M_G be the matrix whose columns correspond to $E(G)$ and the rows to $\chi^{\delta(S)}$ with $S \in \mathcal{F}_G$. Figure 4.16 shows the matrix M_G. Observe that M_G has rank $|E(G)|-1$ because $\sum_{e\in E(G)} c(e)x_e \geqslant k$ defines a ridge of $\text{cutdom}(G)$.

We claim that we can express the column of uv as a unique linear combination of the other columns. Indeed, the facets f_1 and f_2 such that $c = \alpha f_1 + (1-\alpha) f_2$ satisfy $M_G \cdot f_1 = k \cdot \mathbf{1}$ and $M_G \cdot f_2 = k \cdot \mathbf{1}$, where $\mathbf{1}$ is the all-one vector. Hence, $M_G \cdot (f_1 - f_2) = \mathbf{0}$ and $f_1(e) - f_2(e) \neq 0$ because $f_1(uv) < k/2 < f_2(uv)$. This implies that the matrix $\tilde{M}_G$ obtained from M_G by dropping the column for uv still has rank $|E(G)-1|$.

Let M_H be the matrix whose columns correspond to $E(H)$ and the rows correspond to $\chi^{\delta(S)}$ with $S \in \mathcal{F}_H$. Figure 4.17 shows the matrix M_H. Notice that the unit vectors e_{uw} and e_{wv} are in the span of M_H since

$$e_{uv} = \frac{1}{2}(\chi^{\delta(S^*)} + \chi^{\delta(\{w\})} - \chi^{\delta(S^*\cup\{w\})})$$

and

$$e_{vw} = \frac{1}{2}(\chi^{\delta(S^*\cup\{w\})} + \chi^{\delta(\{w\})} - \chi^{\delta(S^*)}).$$

This implies that the rank of M_H is equal to 2 plus the rank of the matrix obtained from M_H by removing the columns for uv and vw. This last matrix is $\tilde{M}_G$, and hence

$$\text{rk}(M_H) = |E(G)| - 1 + 2 = |E(G)| + 1 = |E(H)|.$$

Hence, M_H has full rank, which implies that $\sum_{e\in E(H)} c_H(e)x_e \geqslant k$ defines a facet of $\text{cutdom}(H)$.

Finally, observe that if $c_H(e)$ is fractional for some $e \in E(H)$, then $\sum_{e\in E(H)} c_H(e)x_e \geqslant k$ is a witness for $\varphi(H) > k$. □

$E(G) \setminus \delta(S^*)$	$\delta(S^*) \setminus \{uv\}$	uv	
$* \cdots *$	$* \cdots *$	$*$	
$\vdots$	$\vdots$	$\vdots$	$\lvert\mathcal{F}_G\rvert - 1$ cuts
$* \cdots *$	$* \cdots *$	$*$	
$0 \cdots 0$	$1 \cdots 1$	1	$\delta(S^*)$

Figure 4.16. The matrix M_G illustrating the cuts in the graph (G, c).

$E(G) \setminus \delta(S^*)$	$\delta(S^*) \setminus \{uv\}$	uw	wv	
$* \cdots *$	$* \cdots *$	$*$	0	
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\lvert\mathcal{F}_G\rvert - 1$ cuts
$* \cdots *$	$* \cdots *$	$*$	0	
$0 \cdots 0$	$1 \cdots 1$	1	0	$\delta(S^*)$
$0 \cdots 0$	$1 \cdots 1$	0	1	$\delta(S^* \cup \{w\})$
$0 \cdots 0$	$0 \cdots 0$	1	1	$\delta(\{w\})$

Figure 4.17. The matrix M_H illustrating the cuts in the graph (H, c_H).

We can now show that the minor-minimal graphs which are non-internally 3-connected disprove Conjecture 3.10.1 from [19].

Theorem 4.22. *A minor-minimal graph G with $\varphi(G) > 4$ with a (λ_1, λ_2)-separation (G_1, G_2) with $\lambda_1 \neq \lambda_2$ satistisfies $\varphi(G) = 8$.*

Proof. First recall that we can assume $(\lambda_1, \lambda_2) = (1, 3)$. Let $\sum_{e \in E(G)} c(e) x_e \geqslant 4$ be a witness for $\varphi(G) > 4$ such that the inequality, when put in minimum integer form, has right-hand side $\varphi(G)$.

Assume as before that $V(G_1) \cap V(G_2) = \{u, v\}$. For $i \in [2]$, let G_i' be the graphs with the edge uv added, and let c_i be the restrictions of c to $E(G_i)$, respectively, and let c_i' be the weight function on $E(G_i')$ such that $c_i'(uv) = \lambda_{3-i}$ and $c_i'(e) = c_i(e)$ for all edges $e \in E(G_i)$. By Lemma 4.14, we know that $\sum_{e \in E(G_1')} c_1'(e) x_e \geqslant 4$ defines a ridge of $\text{cutdom}(G_1')$. Hence, we can write $c_1' = \alpha f_1 + (1 - \alpha) f_2$, where

$$\sum_{e \in E(G_1')} f_1(e) x_e \geqslant 4$$

and

$$\sum_{e \in E(G_1')} f_2(e) x_e \geqslant 4$$

define facets of $\text{cutdom}(G_1')$, and $0 < \alpha < 1$. Observe that $c_1'(uv) = 3$ and $f_i(e) \in \{0, 1, 2, 3, 4\}$ for every $i \in [2]$ and $e \in E(G_1')$. Furthermore, we may assume $f_1(uv) \leqslant c_1'(uv) = 3 \leqslant f_2(uv)$. We can determine all possible triples $(f_1(uv), f_2(uv), \alpha)$ satisfying the following conditions.

$$\begin{cases} 3 = \alpha f_1(uv) + (1 - \alpha) f_2(uv) \\ f_1(uv) \in \{0, 1, 2, 3\} \\ f_2(uv) \in \{3, 4\} \\ 0 < \alpha < 1 \end{cases} \tag{4.2}$$

Table 4.1 shows all possible triples.

We can rule out the Case 1 since if $f_1(uv) = f_2(uv) = 3$, we could change the cost $c(e)$ of all edges $e \in E(G_1)$ to $f_1(e)$ while keeping the same family of minimum cuts, contradicting the fact that $\sum_{e \in E(G)} c(e) x_e \geqslant 4$ is facet-defining.

	$f_1(uv)$	$f_2(uv)$	α
Case 1	3	3	$]0,1[$
Case 2	2	4	$\frac{1}{2}$
Case 3	1	4	$\frac{1}{3}$
Case 4	0	4	$\frac{1}{4}$

Table 4.1: The possible values satisfying (4.2).

Case 2 leads to a half-integral weight function c'_1. Hence, c is half-integral as well, which implies $\varphi(G) = 8$.

Observe that in Case 3, the weight function $\tilde{c}'_1 = \frac{2}{3}f_1 + \frac{1}{3}f_2$ defines the same ridge as c'_1 and has $\tilde{c}'_1(uv) = 2$. Hence, we can apply Lemma 4.21 to G'_1 and $\tilde{c}'_1$ in order to get a graph $H = G'_1 +_{uv} K_3$ that is a minor of G. Moreover, observe that $\tilde{c}'_1(e)$ is fractional if and only if $c'_1(e)$ is fractional for every edge $e \in E(G_1)$, because $x + 2y \not\equiv 0 \pmod 3$ if and only if $2x + y \not\equiv 0 \pmod 3$. Since $c'_1(e)$ is fractional for some edge $e \in E(G_1)$, the weight function c_H that we obtain from Lemma 4.21 is also fractional. This contradicts the minor-minimality of G.

Case 4 is similar to Case 3. This time, we define $\tilde{c}'_1 = \frac{1}{4}f_1 + \frac{3}{4}f_2$. Again, we apply Lemma 4.21 to G'_1 and $\tilde{c}'_1$ to obtain a proper minor H of G and a facet-defining weight function c_H. By minimality of G, c_H is integer and thus $f_1(e) + 3f_2(e) \equiv 0 \pmod 4$ for each $e \in E(G_1)$. Hence, $3f_1(e) + f_2(e) \equiv 0 \pmod 2$ and c'_1 is half-integral. We again conclude $\varphi(G) = 8$. □

Before we state the last lemma of this section, we remark that there exist two possible ways of gluing two graphs G_1 and G_2 on a common edge. We can identify (u_1, v_1) with (u_2, v_2), or (u_1, v_1) with (v_2, u_2).

Lemma 4.23. *Let G be a minor-minimal graph with $\varphi(G) > 4$. Let $\sum_{e\in E(G)} c(e)x_e > 4$ be a witness for $\varphi(G) > 4$. Assume that G has a 2-separation (G_1, G_2) with $\lambda_1 = 1$ in (G, c) and $\varphi(G'_1) = 2$. Then G'_1 is the graph on the left in Figure 4.18.*

Proof. Let $c'_1 : E(G'_1) \to \mathbb{Q}_+$ be defined as before. We claim that c'_1 is of the form $c'_1 = 1/4f_1 + 3/4f_2$, where f_1 and f_2 are weight functions defining facets of $\operatorname{cutdom}(G'_1)$ as in Case 4 from the proof of Theorem 4.22. Indeed, $f_i(e) \in \{0, 2, 4\}$ for every $e \in E(G_1)$ and $i \in [2]$ by the assumption

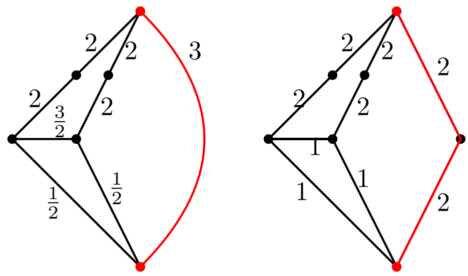

Figure 4.18. On the left the weighted graph (G'_1, c'_1) with $\varphi(G'_1) = 2$. On the right the graph (H, c_H) obtained by applying Lemma 4.21 to the graph G'_1 and the red edge uv.

$\varphi(G'_1) = 2$. Hence, the Cases 1 and 3 from the proof of Theorem 4.22 cannot happen. If $c'_1 = 1/2f_1 + 1/2f_2$, then $c'_1(e)$ is integer for every $e \in E(G_1)$. This contradicts the fact that there is some edge $e^* \in E(G_1)$ with $c'_1(e^*)$ fractional.

Hence, $c'_1 = 1/4f_1 + 3/4f_2$. Now, observe that $\tilde{c}'_1 = 1/2f_1 + 1/2f_2$ is such that $\tilde{c}'_1(uv) = 2$. We can apply Lemma 4.21 to the graph G'_1 and obtain the graph H, which is obtained from G'_1 by subdividing the edge uv once and creating the vertex w, and the weight function c_H. Recall that $c_H(uw) = c_H(wv) = 2$ and $c_H(e) = c_1(e)$ for every $e \in E(G_1)$. By Lemma 4.21, we know that $\sum_{e \in E(H)} c_H(e)x_e \geqslant 4$ is a facet-defining inequality for $\text{cutdom}(H)$.

Observe that $\varphi(H) = 4$. By contradiction, assume $\varphi(H) = 2$. Hence, $c_H(e) = 1/2f_1(e) + 1/2f_2(e) \in \{0, 2, 4\}$ for every edge $e \in E(H)$, which implies that $c'_1(e) = 1/4f_1(e) + 3/4f_2(e)$ is integer for every $e \in E(G_1)$, a contradiction.

It remains to identify the graphs H such that $\varphi(H) = 4$ and contracting an edge incident to a degree-2 vertex results in a graph H' with $\varphi(H') = 2$. Observe that H contains a prism or pyramid minor by [26, Theorem 5]. These minors are shown in Figure 1.4 on page 8. If H has a proper pyramid minor or a prism minor, then contracting an edge incident to a degree-2 vertex results in a graph H' with a pyramid or prism minor which satisfies $\varphi(H') = 4$, a contradiction.

Hence, H is the pyramid graph, see the right graph in Figure 4.18 and contracting an edge incident to a degree-2 vertex results in the graph G'_1 shown on the left in the figure. □

4.5 Amplifiers

In this section, we introduce so-called amplifier graphs, which are graphs A with a marked edge, and show that gluing this graph along the marked edge to some graph G may increase $\varphi(G)$ if done in a certain way.

We say that a pair (A, uv) is an *amplifier*, where A is a graph called *amplifier graph* and $uv \in E(\mathsf{A})$, if the following conditions are satisfied.

1. $\varphi(\mathsf{A}) \geqslant 2$.
2. There exist facet-defining inequalities $\sum_{e \in E(\mathsf{A})} f_i(e) x_e \geqslant 1$ of $\text{cutdom}(\mathsf{A})$ for $i \in [2]$ with $f_1(uv) \leqslant 1/2 < f_2(uv) \leqslant 1$ such that $c_\alpha = \alpha f_1 + (1-\alpha) f_2$ is a weight function defining a ridge of $\text{cutdom}(\mathsf{A})$ for every $0 < \alpha < 1$. Note that this implies $c_\alpha(e) \in [0, 1]$ for every $e \in E(\mathsf{A})$.
3. For all fixed values of k and a, where k is an even integer and $k/2 < a < k$, there exists a unique value of α such that $kc_\alpha(uv) = a$. Furthermore, if a is an odd integer, then there exists an edge $e^* \in E(\mathsf{A})$ such that $kc_\alpha(e^*)$ is fractional.
4. There exists a minimum cut $\delta(S^*)$ with $S^* \subsetneq V(\mathsf{A})$ such that $\delta(S^*)$ separates u and v in (A, c_α) for every value of $\alpha \in]0, 1[$.

We remark that in Property 4, the cut $\delta(S^*)$ separating u and v in (A, c_α) is unique if $c_\alpha(uv) > k/2$. Otherwise, we could combine the two minimum cuts to obtain a cut with strictly smaller weight, similarly as in Lemma 4.17.

The next result is the main theorem of this section.

Theorem 4.24. *Let H be a graph such that $\sum_{e \in E(H)} c_H(e) x_e \geqslant k$ is a facet-defining inequality of* $\text{cutdom}(H)$ *in minimum integer form with $k = \varphi(H) \geqslant 2$. Let uv be an edge of H such that $c_H(uv) < k/2$ and $c_H(uv)$ is odd. Let (A, uv) be an amplifier.*

Then the graph G obtained by gluing H and A along the edge uv, $G = \mathsf{A} +_{uv} H$ satisfies $\varphi(G) \geqslant 2\varphi(H)$.

Proof. Let c_A be the ridge of $\text{cutdom}(\mathsf{A})$ such that c_A is of the form $c_A = kc_\alpha$, where α is such that $kc_\alpha(uv) = k - a$ and $a = c_H(uv)$, which are both odd integers. Let $c : E(G) \to \mathbb{R}$ be the weight function on G such that

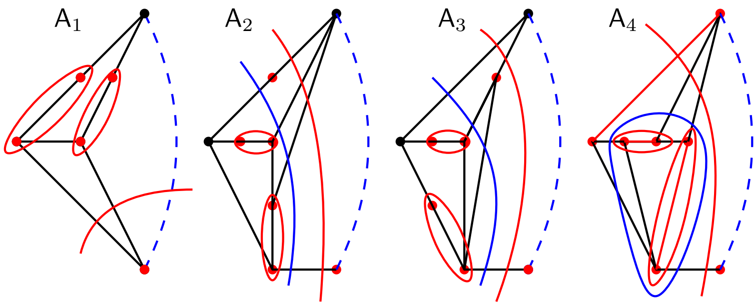

Figure 4.19. The four amplifiers (A_i, uv) for $i \in [4]$ with the edge uv shown in dashed blue. The red vertices are level-0 sets and the red sets are level-1 and level-2 sets.

X	$E(\mathsf{A}) \setminus (X \cup \{uv\})$	$E(H) \setminus \{uv\}$	
$* \cdots *$	$* \cdots *$	$0 \cdots 0$	
$\vdots$	$\vdots$	$\vdots$	cuts from $\mathcal{A} \setminus \{S^*\}$
$* \cdots *$	$* \cdots *$	$0 \cdots 0$	
$1 \cdots 1$	$0 \cdots 0$	$* \cdots *$	
$\vdots$	$\vdots$	$\vdots$	u–v cuts from $\mathcal{F}_H$
$1 \cdots 1$	$0 \cdots 0$	$* \cdots *$	
$0 \cdots 0$	$0 \cdots 0$	$* \cdots *$	
$\vdots$	$\vdots$	$\vdots$	other cuts from $\mathcal{F}_H$
$0 \cdots 0$	$0 \cdots 0$	$* \cdots *$	

Figure 4.20. The matrix M_G illustrating the cuts in the graph (G, c).

$c(e) = c_H(e)$ if $e \in E(H) \setminus \{uv\}$ and $c(e) = c_A(e)$ if $e \in E(\mathsf{A}) \setminus \{uv\}$. We claim that $\sum_{e \in E(G)} c(e)x_e \geqslant k$ defines a facet of $\operatorname{cutdom}(G)$.

Notice that by Property 3, some edge $e^* \in E(\mathsf{A}) \setminus \{uv\}$ is such that $c_A(e^*) = kc_\alpha(e^*)$ is fractional. Hence, we also have that $c(e^*)$ is fractional and that the inequality $\sum_{e \in E(G)} c(e)x_e \geqslant k$ has some fractional coefficients.

Now, put this inequality in minimum integer form. Notice that we need to multiply the inequality with some number of the form $q/c(e^*)$, where q is an integer, since otherwise the coefficient of e^* is not integer. We may write $q/c(e^*) = s/t$, where the fraction s/t is irreducible.

Assume that $t > 1$. Observe that $\frac{s}{t}c_H(e)$ integer implies that $\frac{c_H(e)}{t}$ is integer as well for every edge $e \in E(H) \setminus \{uv\}$. Similarly, k/t is integer. It follows that $c_H(uv)/t = a/t$ is integer because there exists a minimum u–v cut in $(H, \frac{c_H}{t})$ of integer weight k/t and all edges $e \in E(H) \setminus \{uv\}$ have integer weight $\frac{c_H(e)}{t}$. Hence, the inequality $\sum_{e \in E(H)} c_H(e)x_e \geqslant k$ is not in minimum integer form, a contradiction.

This implies $t = 1$ (and $s > 1$) and the minimum integer form of the facet-defining inequality of $\operatorname{cutdom}(G)$ has right-hand side sk. Hence, $\varphi(G) \geqslant sk \geqslant 2\varphi(H)$.

We want to show that $\sum_{e \in E(G)} c(e)x_e \geqslant k$ defines a facet of $\operatorname{cutdom}(G)$. For this, we need to show that there exists a family $\mathcal{F}$ defining $|E(G)|$ linearly independent minimum cuts of weight k in (G, c).

The weight of a minimum cut in (G, c) is k by our choice of the weight function. Indeed, any cut in (G, c) not separating u and v has weight at least k because all cuts in (A, c_A) and (H, c_H) have minimum weight k. If a cut separates u and v, then its restriction to (A, c_A) has weight at least a and its restriction to (H, c_H) has weight at least $k - a$. Thus any cut in (G, c) has weight at least k.

Next, we show that there exist $|E(G)|$ minimum cuts in (G, c). Let $\mathcal{A}$ be a family defining $|E(\mathsf{A})| - 1$ linearly independent minimum cuts in (A, c_A). Let $\mathcal{F}_H$ be any family defining $|E(H)|$ linearly independent minimum cuts in (H, c_H). Note that we may assume that there exists a minimum cut separating u and v in (H, c_H) by Lemma 4.3 because c_H defines a non-trivial facet of $\operatorname{cutdom}(H)$. Moreover, by Property 4 in the definition of amplifier, we may assume that there is a set $S^* \in \mathcal{A}$ such that $\delta(S^*)$ is a minimum u–v cut, and S^* is unique. We let $X = \delta(S^*) \setminus \{uv\}$ in the graph G. We construct a family $\mathcal{F} = \{S \subsetneq V(G) \mid c(\delta(S)) = k\}$ defining minimum

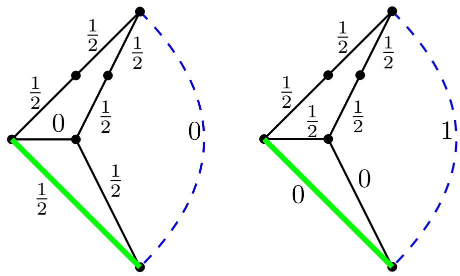

Figure 4.21. The two facets of (A_1, uv). The edge uv is in dashed blue, the edge e^* is in fat green.

cuts as follows. For each set $S_A \in \mathcal{A} \setminus \{S^*\}$, we include S_A in $\mathcal{F}$. For each set $S_H \in \mathcal{F}_H$ we include S_H in $\mathcal{F}$ if $uv \notin \delta(S_H)$. If $uv \in \delta(S_H)$, then we add the set S'_H to $\mathcal{F}$ such that $\delta(S'_H) = \delta(S_H) \setminus \{uv\} \cup X$ in G. Observe that $\delta(S'_H)$ is a minimum cut in (G, c) because $c(X) = c_A(X) = a = c_H(uv)$. Hence, $\mathcal{F}$ contains $(|E(\mathsf{A})| - 1) - 1 + |E(H)| = |E(\mathsf{A})| + |E(H)| - 2 = |E(G)|$ minimum cuts.

To complete the proof it remains to show that the cuts defined by $\mathcal{F}$ are linearly independent. This is done as in the proof of Lemma 4.21 by showing that the matrix M_G shown in Figure 4.20 has full rank. We leave the details to the reader. □

We will now introduce four amplifiers (A_i, uv), $i \in [4]$. When applying Theorem 4.24 to any of these amplifiers and the prism or pyramid graph, the resulting graph is a minor-minimal graph with $\varphi(G) > 4$. These are shown in Figure 4.6 on page 96. Figure 4.19 shows the four amplifiers graphs $\mathsf{A}_1, \mathsf{A}_2, \mathsf{A}_3, \mathsf{A}_4$ together with a linearly independent family of minimum cuts for the ridge (without the weights).

In Figures 4.21, 4.22, 4.23, and 4.24 on page 118, we show for each amplifier graph A_i, $i \in [4]$, the weights of the two facets from which we can obtain the ridge of A_i with the minimum cuts from Figure 4.19 such that (A_i, uv) satisfies the four conditions on page 115, where uv is a well-chosen edge. The edge uv is in dashed blue in each figure, and the fat green edge corresponds to the edge e^* such that $kc_\alpha(e^*)$ is fractional if k is even and $kc_\alpha(uv)$ is an odd integer.

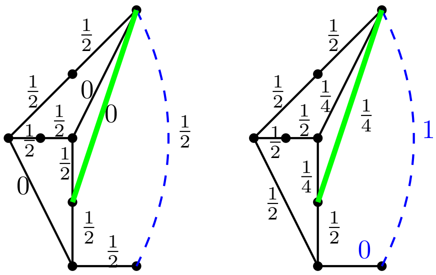

Figure 4.22. The two facets of (A_2, uv). The edge uv is in dashed blue, the edge e^* is in fat green.

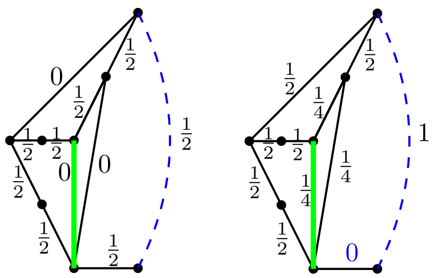

Figure 4.23. The two facets of (A_3, uv). The edge uv is in dashed blue, the edge e^* is in fat green.

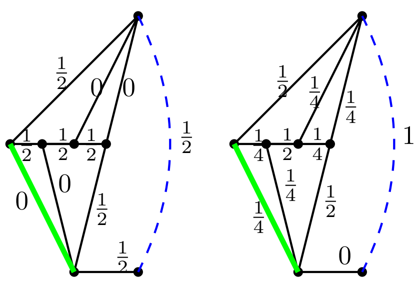

Figure 4.24. The two facets of (A_4, uv). The edge uv is in dashed blue, the edge e^* is in fat green.

4.6 Further research directions

Unfortunately, we are not yet able to prove that the graphs in Figure 1.9 on page 12 form the complete set of minimal excluded minors for the class of graphs satisfying $\varphi(G) \leqslant 4$. We recall the results that we have established in this chapter and state some conjectures. Furthermore, we will give some open questions and ideas for further research.

First, observe that the laminar family $\mathcal{F}$ defining linearly independent minimum cuts in all our minor-minimal graphs with $\varphi(G) > 4$ contain at least one level-2 set. This is a property that can be shown for all minor-minimal graphs with $\varphi(G) > 4$. Notice also that, for instance, the minor-minimal graph obtained by gluing A_2 to the prism has two level-2 sets, while the Petersen graph has one level-2 set. Hence, minor-minimal graphs with $\varphi(G) > 4$ do not share the same structure of the laminar family $\mathcal{F}$. An open question is to determine all possible structures of the laminar family $\mathcal{F}$.

Question 4.25. *What are the possible structures for a laminar family $\mathcal{F}$ defining linearly independent minimum cuts in a graph (G, c), where G is a minor-minimal graph with $\varphi(G) > k$, where k is an even integer with $k \geqslant 4$, and $\sum_{e \in E(G)} c(e)x_e \geqslant k$ is a witness for $\varphi(G) > k$?*

Furthermore, we can ask how level-2 sets behave in minor-minimal graphs with $\varphi(G) > k$. For level-0 sets it is known that these correspond to singletons, while level-1 sets contain two adjacent vertices. We can show that, in minor-minimal graphs, the following assertions are true. A level-2 set S contains at least four vertices which form a cycle. If $|S| \geqslant 5$, then the subgraph induced by S contains a cycle of length at least five. Furthermore, if $\delta(S)$ is a matching of size three, then the cycle passes through the vertices incident to the edges of $\delta(S)$.

Now, consider the known internally 3-connected excluded minors for $\varphi(G) \leqslant 4$ from Section 4.3.1. Observe that each of them consists of two induced cycles and at most five edges between these cycles. It can be shown that any graph consisting of two cycles and a matching of size five between them contains a minor from Section 4.3.1. This limits the candidates for internally 3-connected graphs G that are minor-minimal with $\varphi(G) > k$. As we do not have evidence for the existence of other minor-minimal internally 3-connected graphs G with $\varphi(G) > 4$, we conjecture that our graphs are the only possible ones.

Conjecture 4.26. *The graphs in Figure 4.3 on page 93 and in Figure 4.5 on page 95 are the only internally* 3*-connected graphs G such that G is a minor-minimal graph with $\varphi(G) > 4$.*

In Section 4.4 we have shown several properties that minor-minimal graphs G with $\varphi(G) > 4$ satisfy when $G = (G_1, G_2)$ is a 2-separation of G of type (λ_1, λ_2) with $V(G_1) \cap V(G_2) = \{u, v\}$. We give a list of the main results in Table 4.2. For this, we use the same notations and definitions as in Section 4.4.

Lemma 4.13	The graph G is simple and 2-connected.
Lemma 4.19	The graph $G - \{u, v\}$ has exactly two connected components.
Lemma 4.15	If $(\lambda_1, \lambda_2) = (2, 2)$, then G_2 is a path of length 2.
Lemma 4.20	If $(\lambda_1, \lambda_2) = (1, 3)$, then graph G_2' satisfies $\varphi(G_2') = 4$.
Lemma 4.23	If $(\lambda_1, \lambda_2) = (1, 3)$ and $\varphi(G_1') = 2$, then G_1' is the graph in Figure 4.18 on page 114.
Theorem 4.22	If $(\lambda_1, \lambda_2) = (1, 3)$, then graph G satisfies $\varphi(G) = 8$.

Table 4.2: Overview of the results in Section 4.4. Here G is a minor-minimal graph with $\varphi(G) > 4$ that has a 2-cutset $\{u, v\}$.

The following conjecture is related to Section 4.4 and amplifiers introduced in Section 4.5. Observe that Conjecture 1.5 on page 12 implies the following conjecture, and that Conjectures 4.26 and 4.27 imply Conjecture 1.5.

Conjecture 4.27. *Let G be a minor-minimal graph with $\varphi(G) > 4$ such that (G_1, G_2) is a* 2*-separation of type* $(1, 3)$*. Then G_1' is one of $\mathsf{A}_1, \mathsf{A}_2, \mathsf{A}_3, \mathsf{A}_4$ and G_2' is the prism or pyramid graph. Furthermore, let uv be the marked edge of (A_i, uv) for $i \in [4]$ and let uv be an edge of a triangle in the prism of pyramid graph. Then the graph G is obtained by gluing A_i and the prism or pyramid graph along the edge uv and deleting that edge.*

The next question is related to the previous question. We cannot prove that the four amplifiers we know are the only ones, even with $\varphi(\mathsf{A}) \leqslant 4$. Indeed, if some other amplifier exists with $\varphi(\mathsf{A}) \leqslant 4$, then the previous conjecture is false. As we can construct arbitrary large graphs with amplifiers, it could be of interest to identify other amplifiers.

Question 4.28. *Are there amplifiers* (A, uv) *with* $\varphi(\mathsf{A}) = k$ *for every even* $k \geqslant 2$ *?*

Of course, identifying the minimal excluded minors of the classes of graphs G such that $\varphi(G) \leqslant k$ is an open problem for $k \geqslant 4$. Our last question asks whether it is possible to obtain a result similar to our main result in Chapter 3. That is, do there exist sets of graphs $\mathcal{U}_k$ for every even integer k and a function $f : \mathbb{N} \to \mathbb{R}$ such that each set $\mathcal{U}_k$ has bounded size independently of k, every graph $H \in \mathcal{U}$ is such that $\varphi(H) > k$ and every graph G with $\varphi(G) > f(k)$ has a minor H with $H \in \mathcal{U}_k$?

Question 4.29. *Is it possible to characterize graphs with large* $\varphi(G)$ *value in terms of unavoidable minors ?*

Chapter 5

Ball packings

This chapter is based on joint work with Nicolas Bousquet, Wouter Cames van Batenburg, Louis Esperet, Gwenaël Joret, William Lochet, and François Pirot, see the paper *Packing and covering balls in graphs excluding a minor*, which has been published in Combinatorica [11].

In this chapter, we study a problem about packing and transversals. These problems appear in different flavors and applications. They can be studied from a purely combinatorial point of view, but also have many applications in daily life.

In a combinatorial setting we consider a ground set or universe E and a collection $\mathcal{S}$ of subsets of E. A *packing* $\mathcal{P}$ is a subcollection of $\mathcal{S}$ such that for all distinct $S_i, S_j \in \mathcal{P}$ their intersection is empty, that is $S_i \cap S_j = \emptyset$. A packing is *maximal* if it is inclusion-wise maximal, that is, every set $S \in \mathcal{S} \setminus \mathcal{P}$ has non-empty intersection with some set of $\mathcal{P}$. A packing is *maximum* if it has maximum cardinality among all packings. We let $\nu(\mathcal{S})$ denote the size of a maximum packing in $\mathcal{S}$.

Another problem we will study is the *hitting set problem.* In this problem we ask to find a minimum set $X \subseteq E$ such that X meets every member of $\mathcal{S}$. That is, we ask that each set in $\mathcal{S}$ contains at least one element from X. We say that X is a *transversal* and that the elements of X *hit* $\mathcal{S}$. The minimum cardinality of X is denoted by $\tau(\mathcal{S})$.

In combinatorial optimization, the problem of finding a maximum packing or a minimum transversal are closely related. Indeed, it is well-known that a transversal has at least the size of a maximum packing because we want the transversal to hit each set of the packing. However, it is unclear whether in general there exists a similar relation that bounds $\nu(\mathcal{S})$ as a function of $\tau(\mathcal{S})$ for a given E and $\mathcal{S}$. If we can find such a function, we say that the problem satisfies the *Erdős-Pósa property*. The name goes back to Erdős and Pósa, who showed in 1965 [39] that in any graph with at most k vertex-disjoint cycles, only $O(k \log k)$ vertices are needed to hit all the cycles and that this is best possible.

We will focus on packing and transversal problems in graphs. For this, we consider a hypergraph $\mathcal{H}$ whose vertices correspond to the vertices of the original graph and the hyperedges of the hypergraph correspond to the objects (usually subgraphs) that we want to pack and hit. Our ground set is $V(\mathcal{H})$ and the collection of subsets we consider is $E(\mathcal{H})$.

During the years, many papers appeared that studied generalizations of the result of Erdős and Pósa to other graph minors. Moreover, alternative proofs of the original Erdős-Pósa theorem have been given, see for instance [16, 65] or [31, Chapters 2.3 and 12.6].

Robertson and Seymour [65] showed that the class of graphs containing a fixed planar graph H as a minor satisfies the Erdős-Pósa property. That is, for each graph H, there exists a function $f_H : \mathbb{N} \to \mathbb{R}$ such that, for every graph G and every positive integer k, the graph G has k vertex-disjoint subgraphs each containing H as a minor, or there exists a subset X of vertices of G with $|X| \leqslant f_H(k)$ such that $G - X$ has no H-minor. The function given by the Robertson-Seymour result is exponential in k. It has recently been improved to a $O(k \log k)$-bound by Cames van Batenburg, Huynh, Joret, and Raymond [15], which is tight by the original bound from Erdős and Pósa [39] if H has a cycle. The case of forest minors was studied by Fiorini, Joret, and Wood [43] who showed that the bound is $O(k)$.

In this chapter, we will consider packings and transversals of balls. There are some differences with the above cited papers. Instead of packing and hitting minors of a given graph we consider balls in graphs that exclude a minor. Furthermore, our Erdős-Pósa property function depends on t, which is such that G has no K_t minor, and does not depend on the balls we pack. Observe that for the above results the function depends only on the minor H we want to pack and is independent of the graph G.

Given a graph $G = (V, E)$, an integer $r \geqslant 0$, and a vertex $v \in V$, we denote by $B_r(v)$ the *ball of radius* r in G centered in v, that is

$$B_r(v) := \{u \in V(G) \mid d_G(u, v) \leqslant r\},$$

where $d_G(u, v)$ denotes the distance between u and v in G (we will omit the subscript G when the graph is clear from the context). If the ball has radius r we say that it is an *r-ball.* We say that a hypergraph $\mathcal{H}$ is a *ball hypergraph* of G if $\mathcal{H}$ has vertex set $V = V(G)$ and each edge of $\mathcal{H}$ is a ball $B_r(v)$ in G for some integer r and some vertex $v \in V$. If all the balls forming the edges of $\mathcal{H}$ have the same radius r, we say that $\mathcal{H}$ is an *r-ball hypergraph* of G. Remark that a r-ball hypergraph (or a ball hypergraph) does not require to include all possible balls. Figure 1.6 on page 10 shows an example of a planar graph where we pack and hit balls of radius 2.

The problem we study goes back to 2001. Gavoille, Peleg, Raspaud, and Sopena [47] conjectured that there exists a constant c such that in every planar graph of diameter at most $2r$, all r-balls can be hit with c vertices, and showed the lower bound $c \geqslant 4$. Their conjecture was proved in 2007 by Chepoi, Estellon, and Vaxès [22], and later extended to graphs embeddable on a fixed surface with a bounded number of apices by Borradaile and Chambers [9].

Note that G has diameter at most $2r$ if and only if there are no two disjoint balls of radius r in G. Thus, these results state equivalently the existence of a universal constant c such that for every $r \geqslant 0$ and every planar (or more generally bounded genus) graph G, if the r-ball hypergraph $\mathcal{H}$ consisting of all balls of radius r satisfies $\nu(\mathcal{H}) = 1$, then $\tau(\mathcal{H}) \leqslant c$. With this interpretation in mind, Chepoi, Estellon, and Vaxès [13] conjectured the following generalization in 2007 (see also [40]).

Conjecture 5.1 (Chepoi, Estellon, and Vaxès [13])**.** *There exists a constant c such that for every integer $r \geqslant 0$, every planar graph G, and every r-ball hypergraph $\mathcal{H}$ of G, we have $\tau(\mathcal{H}) \leqslant c \cdot \nu(\mathcal{H})$.*

If one considers all metric spaces obtained as standard graph-metrics of planar graphs, then Conjecture 5.1 states that these metric spaces satisfy the so-called *bounded covering-packing property* [21]. Recently, Chepoi, Estellon, and Naves [21] showed that other metric spaces do have this property, including the important case of Busemann surfaces. (Quoting [21], the latter are roughly the geodesic metric spaces homeomorphic to $\mathbb{R}^2$ in which the distance function is convex; they generalize Euclidean spaces, hyperbolic spaces, Riemannian manifolds of global nonpositive sectional curvatures, and CAT(0) spaces.)

Going back to Conjecture 5.1, let us emphasize that a key aspect of this conjecture is that the constant c is independent of the radius r. If c is allowed to depend on r, then the conjecture is known to be true. In fact, it holds more generally for all graph classes with bounded expansion, as shown by Dvořák [34].

Some evidence for Conjecture 5.1 was given by Bousquet and Thomassé [12], who proved that it holds with a polynomial bound instead of a linear one. More generally, they proved that for every integer $t \geqslant 1$, there exists a constant c_t such that for every integer $r \geqslant 0$, every K_t-minor free graph G, and every r-ball hypergraph $\mathcal{H}$ of G, we have $\tau(\mathcal{H}) \leqslant c_t \cdot \nu(\mathcal{H})^{2t+1}$.

The main result of this chapter is that Conjecture 5.1 is true, and furthermore it is not necessary to assume that all the balls have the same radius. The following theorem is equivalent to Theorem 1.6.

Theorem 5.2 (Main result). *For every integer $t \geqslant 1$, there is a constant c_t such that $\tau(\mathcal{H}) \leqslant c_t \cdot \nu(\mathcal{H})$ for every K_t-minor-free graph G and every ball hypergraph $\mathcal{H}$ of G.*

A set S of vertices of a graph G is r-*dominating* if each vertex of G is at distance at most r from S, and r-*independent* if any two vertices of S are at distance at least $2r+1$ apart in G. Note that if we take $\mathcal{H}$ to be the r-ball hypergraph consisting of all balls of radius r in G, Theorem 5.2 has the following interesting graph-theoretic interpretation: if G is K_t-minor-free, then the minimum size of an r-dominating set is at most c_t times the maximum size of an r-independent set in G.

5.1 Proof idea and content of the chapter

Among the tools that we use to prove our main theorem, Theorem 5.2, are some that are related to the fractional packing and transversal numbers which are upper and lower bounds of the packing and transversal numbers, respectively. We can express the packing number of a hypergraph $\mathcal{H}$ as an integer program.

$$\nu(\mathcal{H}) = \max \sum_{e \in \mathcal{E}(\mathcal{H})} w_e$$

$$\text{given that} \begin{cases} \sum_{e \ni v} w_e \leqslant 1 & \text{for every vertex } v \text{ of } \mathcal{H}, \\ w_e \in \{0,1\} & \text{for every edge } e \text{ of } \mathcal{H}. \end{cases}$$

The *fractional packing number* $\tau^*(\mathcal{H})$ is obtained by considering the linear relaxation of the above program.

$$\nu^*(\mathcal{H}) = \max \sum_{e \in \mathcal{E}(\mathcal{H})} w_e$$

$$\text{given that } \begin{cases} \sum_{e \ni v} w_e \leqslant 1 & \text{for every vertex } v \text{ of } \mathcal{H}, \\ w_e \geqslant 0 & \text{for every edge } e \text{ of } \mathcal{H}. \end{cases}$$

Similarly, the transversal number $\nu(\mathcal{H})$ is given by the following integer program and $\nu^*(\mathcal{H})$ by its linear relaxation.

$$\tau(\mathcal{H}) = \min \sum_{v \in V(\mathcal{H})} w_v$$

$$\text{given that } \begin{cases} \sum_{v \in e} w_v \geqslant 1 & \text{for every edge } e \text{ of } \mathcal{H}, \\ w_v \in \{0,1\} & \text{for every vertex } v \text{ of } \mathcal{H}. \end{cases}$$

$$\tau^*(\mathcal{H}) = \min \sum_{v \in V(\mathcal{H})} w_v$$

$$\text{given that } \begin{cases} \sum_{v \in e} w_v \geqslant 1 & \text{for every edge } e \text{ of } \mathcal{H}, \\ w_v \geqslant 0 & \text{for every vertex } v \text{ of } \mathcal{H}. \end{cases}$$

Observe that the linear relaxations of the packing and transversal problems are dual linear programs one of another. As both programs have a finite optimum, the strong duality theorem tells us that $\nu(\mathcal{H}) \leqslant \nu^*(\mathcal{H}) = \tau^*(\mathcal{H}) \leqslant \tau(\mathcal{H})$ for every hypergraph $\mathcal{H}$. Our main result states that there exists a constant c such that $\tau(\mathcal{H}) \leqslant c\nu(\mathcal{H})$ for every ball hypergraph $\mathcal{H}$. This result is obtained in several steps, following a bootstrapping approach.

Our proof of Theorem 5.2 relies on the existence of some function f_t such that $\tau(\mathcal{H}) \leqslant f_t(\nu(\mathcal{H}))$. Hence, we use the Erdős-Pósa property of the ball hypergraphs of K_t-minor-free graphs in the proof when $\nu(\mathcal{H})$ is not 'too big'. However, showing this property was an open problem. This was known for r-ball hypergraphs by the result of Bousquet and Thomassé [12] but their proof method does not extend to the case of balls of arbitrary radii. For this reason, as a first step toward proving Theorem 5.2, we prove Theorem 5.3 below establishing said Erdős-Pósa property. We also note that, while the bounding function in Theorem 5.3 is not optimal, it is a near linear bound of the form $\tau(\mathcal{H}) \leqslant c_t \cdot \nu(\mathcal{H}) \log \nu(\mathcal{H})$ where c_t is a small explicit constant polynomial in t. This is in contrast with the constant c_t in our proof of Theorem 5.2 which is large, exponential in t. Thus, the bound in Theorem 5.3 is better for small values of $\nu(\mathcal{H})$.

Theorem 5.3 (Near linear bound)**.** *Let G be a graph with no K_t-minor and such that every minor of G has average degree at most d. Then for every ball hypergraph $\mathcal{H}$ of G,*

$$\tau(\mathcal{H}) \leqslant 2\mathrm{e}(t-1)\, d \cdot \nu(\mathcal{H}) \cdot \log(11\mathrm{e}\, d \cdot \nu(\mathcal{H})).$$

In particular, $\tau(\mathcal{H}) \leqslant ct^2\sqrt{\log t} \cdot \nu(\mathcal{H}) \cdot \log(t \cdot \nu(\mathcal{H}))$ for some absolute constant $c > 0$, and if G is planar then $\tau(\mathcal{H}) \leqslant 48\,\mathrm{e} \cdot \nu(\mathcal{H}) \cdot \log(66\,\mathrm{e} \cdot \nu(\mathcal{H}))$.

In order to obtain Theorem 5.3, we want to bound $\tau^*(\mathcal{H})$ as a function of $\tau(\mathcal{H})$, and $\nu(\mathcal{H})$ as a function of $\nu^*(\mathcal{H})$. The second bound is new and given by the following theorem.

Theorem 5.4 (Fractional version)**.** *Let G be a graph and let d be the maximum average degree of a minor of G. Then for every ball hypergraph $\mathcal{H}$ of G, we have $\nu^*(\mathcal{H}) \leqslant \mathrm{e}\, d \cdot \nu(\mathcal{H})$.*
In particular, if G is planar then $\nu^*(\mathcal{H}) \leqslant 6\mathrm{e} \cdot \nu(\mathcal{H})$ and if G has no K_t-minor then $\nu^*(\mathcal{H}) \leqslant c \cdot t\sqrt{\log t} \cdot \nu(\mathcal{H})$, for some absolute constant $c > 0$.

The bound of $\tau^*(\mathcal{H})$ as a function of $\tau(\mathcal{H})$ is given by a classical result using bounded VC-dimension of hypergraphs [12, 32]. We can show that the VC-dimension of ball hypergraphs of G is bounded when G excludes a minor. As $\tau^*(\mathcal{H}) = \nu^*(\mathcal{H})$, we can combine these results to obtain Theorem 5.3.

We note that results on the VC-dimension of ball hypergraphs in graphs excluding a minor have also been used recently to obtain improved algorithms for the computation of the diameter in sparse graphs [33, 57].

The proofs of Theorems 5.2, 5.3, and 5.4 are constructive, and can be transformed into efficient algorithms producing transversals (in the case of Theorems 5.2 and 5.3) or matchings (in the case of Theorem 5.4) of the desired size.

The chapter is organized as follows. Sections 5.2 and 5.3 are devoted to technical lemmas that will be used in our proofs. Theorems 5.3 and 5.4

are proved in Section 5.4. Theorem 5.2 is proved in Section 5.5. Finally, we conclude the chapter in Section 5.6 with a construction suggesting that Theorem 5.2 does not extend way beyond proper minor-closed classes.

5.2 Hypergraphs, balls, and minors

We will need two technical lemmas, whose proofs are very similar to the proof of [12, Theorem 4] and [22, Proposition 1]. We start with Lemma 5.5, which will be used in the proofs of Theorem 5.4 and Theorem 5.2. We first need the following definitions.

We say that two balls B_1 and B_2 are *incomparable* if $B_1 \subsetneq B_2$ and $B_2 \subsetneq B_1$. Consider two intersecting and incomparable balls $B_1 = B_{r_1}(v_1)$ and $B_2 = B_{r_2}(v_2)$ in a graph G, and let $d := d_G(v_1, v_2)$. A *median vertex* of B_1 and B_2 is any vertex u lying on a shortest path between v_1 and v_2, at distance $\lfloor \frac{r_1-r_2+d}{2} \rfloor$ from v_1 and at distance $\lceil \frac{r_2-r_1+d}{2} \rceil$ from v_2, or symmetrically at distance $\lceil \frac{r_1-r_2+d}{2} \rceil$ from v_1 and at distance $\lfloor \frac{r_2-r_1+d}{2} \rfloor$ from v_2. Since B_1 and B_2 intersect, we have $r_1 + r_2 \geqslant d$ and since B_1 and B_2 are incomparable, we have $r_2 \leqslant r_1 + d$ and $r_1 \leqslant r_2 + d$, and in particular $\lfloor \frac{r_1-r_2+d}{2} \rfloor \geqslant 0$ and $\lceil \frac{r_2-r_1+d}{2} \rceil \geqslant 0$ (so the distances above are well defined). Moreover, $\lfloor \frac{r_1-r_2+d}{2} \rfloor = \lfloor \frac{2r_1-r_1-r_2+d}{2} \rfloor \leqslant r_1$ and $\lceil \frac{r_2-r_1+d}{2} \rceil = \lceil \frac{2r_2-r_1-r_2+d}{2} \rceil \leqslant r_2$, so any median vertex of B_1 and B_2 lies in $B_1 \cap B_2$. Finally, note that by the definition of a median vertex u of B_1 and B_2,

- for every $\{i, j\} = \{1, 2\}$ we have $r_j - d(v_j, u) \leqslant r_i - d(v_i, u) + 1$, and
- if $v_1 = v_2$ (which implies $B_1 = B_2$ since the balls are incomparable), then $u = v_1 = v_2$.

Lemma 5.5. *Let G be a graph, let $S = \{B_i = B_{r_i}(s_i)\}_{i \in [n]}$ be a set of n pairwise incomparable balls in G, with pairwise distinct centers, and let $E_S \subseteq \binom{S}{2}$ be a subset of pairs of intersecting balls $\{B_i, B_j\} \subseteq S$, each of which is associated with a median vertex $x_{\{i,j\}}$ of B_i and B_j, and such that the only balls of S containing $x_{\{i,j\}}$ are B_i and B_j. Then the graph $H = (S, E_S)$ is a minor of G.*

Proof. Let us fix a total ordering $\prec$ on the vertices of G. In the proof, all distances are in the graph G, so we write $d(u, v)$ instead of $d_G(u, v)$ for the sake of readability. For every pair of balls $\{B_i, B_j\} \in E_S$, we write x_{ij} or x_{ji} instead of $x_{\{i,j\}}$, for the sake of readability (x_{ij}, x_{ji}, and $x_{\{i,j\}}$ all correspond to the same median vertex of B_i and B_j). We also let $P(s_i, x_{ij})$ be a shortest path from s_i to x_{ij}, and we assume that the sequence of vertices from s_i to x_{ij} on the path is minimum with respect to the lexicographic order induced by $\prec$ (among all shortest paths from s_i to x_{ij}). By the assumptions, we know that $P_{ij} := P(s_i, x_{ij}) \cup P(s_j, x_{ij})$ is a shortest path from s_i to s_j.

For every $i \in [n]$, we define

$$\mathcal{T}_i := \bigcup_{j\,:\,\{B_i,B_j\}\in E_S} P(s_i, x_{ij}).$$

Claim 1. For every $i \in [n]$, $\mathcal{T}_i$ is a tree.

Assume for the sake of contradiction that there is a cycle C in $\mathcal{T}_i$. Observe that, by construction, if uv is an edge of $\mathcal{T}_i$ then $|d(s_i, u) - d(s_i, v)| = 1$. Let y be a vertex of C maximizing $d(s_i, y)$, and let z_1, z_2 denote its two neighbors in C. Then $d(s_i, z_1) = d(s_i, z_2) = d(s_i, y) - 1$, and there exist j_1, j_2 such that $z_1 y$ is an edge of $P(s_i, x_{ij_1})$ and $z_2 y$ is an edge of $P(s_i, x_{ij_2})$. Let P_1 and P_2 be the subpaths from s_i to y of $P(s_i, x_{ij_1})$ and $P(s_i, x_{ij_2})$, respectively. Then P_1 and P_2 are two different paths from s_i to y, and one of them is not minimum either in terms of length, or with respect to the lexicographic order induced by $\prec$. This contradicts the definition of $P(s_i, x_{ij_1})$ and $P(s_i, x_{ij_2})$.

Claim 2. For every two pairs of balls $\{B_i, B_k\}, \{B_j, B_\ell\} \in E_S$ with $i \neq j$, if $P(s_i, x_{ik})$ and $P(s_j, x_{j\ell})$ intersect in some vertex y such that $d(y, x_{ik}) \leqslant d(y, x_{j\ell})$, then $j = k$ and $y = x_{ij}$.

Note that $d(s_j, x_{ik}) \leqslant d(s_j, y) + d(y, x_{ik}) \leqslant d(s_j, y) + d(y, x_{j\ell}) = d(s_j, x_{j\ell})$. Since $x_{j\ell}$ is a median vertex of B_j and B_ℓ, we have $d(s_j, x_{j\ell}) \leqslant r_j$, which implies that $d(s_j, x_{ik}) \leqslant r_j$ and thus $x_{ik} \in B_j$. By definition, x_{ik} is only contained in the balls B_i and B_k of S and thus $j = k$. If we also have $i = \ell$, then necessarily $y = x_{ij}$.

From now on, we assume that $i \neq \ell$. Since $P_{ij} = P(s_i, x_{ij}) \cup P(s_j, x_{ij})$ is a shortest path containing the vertex y, the s_j–y section of that path (which contains x_{ij}) has the same length as the s_j–y section of $P(s_j, x_{j\ell})$.

Replacing the latter section by the former, we obtain a shortest path from s_j to $x_{j\ell}$ containing x_{ij}, which we denote $Q(s_j, x_{j\ell})$. As a consequence,

$$d(x_{j\ell}, x_{ij}) = d(x_{j\ell}, s_j) - d(s_j, x_{ij}) \leqslant r_j - d(s_j, x_{ij}) \leqslant r_i - d(s_i, x_{ij}) + 1,$$

where the last inequality follows from the definition of x_{ij}. We now use the fact that y appears on the path $P(s_i, x_{ij})$ and on the x_{ij}–$x_{j\ell}$ section of $Q(s_j, x_{j\ell})$, and obtain

$$\begin{aligned} d(s_i, x_{j\ell}) &\leqslant d(s_i, y) + d(y, x_{j\ell}) = d(s_i, x_{ij}) + d(x_{ij}, x_{j\ell}) - 2d(y, x_{ij}) \\ &\leqslant r_i + 1 - 2d(y, x_{ij}). \end{aligned}$$

Since $x_{j\ell} \notin B_i$ by definition (and so $d(s_i, x_{j\ell}) > r_i$), this implies that $y = x_{ij}$, as desired.

This claim immediately implies that for every $i, j \in [n]$ with $i \neq j$, we have $V(\mathcal{T}_i) \cap V(\mathcal{T}_j) = \{x_{ij}\}$ if $\{B_i, B_j\} \in E_S$, and $V(\mathcal{T}_i) \cap V(\mathcal{T}_j) = \emptyset$ otherwise. Another consequence is that for every $\{B_i, B_j\} \in E_S$, the vertex x_{ij} is a leaf in at least one of the two trees $\mathcal{T}_i$ and $\mathcal{T}_j$ (since otherwise there exist $k \neq j$ and $\ell \neq i$ such that $x_{ij} \in P(s_i, x_{ik})$ and $x_{ij} \in P(s_j, x_{j\ell})$, which readily contradicts Claim 2 above).

In the subgraph $\bigcup_{i\in[n]} \mathcal{T}_i$ of G, for each $i \in [n]$ we contract each edge of $\mathcal{T}_i$ except the ones incident to a leaf of $\mathcal{T}_i$. It follows from the paragraph above that the resulting graph is precisely a graph obtained from $H = (S, E_S)$ by subdividing each edge at most once, and thus H is a minor of G. □

The next result has a very similar proof[1], but the setting is slightly different. It will be used in the proof of Theorem 5.2.

Lemma 5.6. *Let G be a graph and $S = \{B_i = B_{r_i}(s_i)\}_{i\in[n]}$ be a set of n pairwise vertex-disjoint balls in G, and let $E_S \subseteq \binom{S}{2}$ be a subset of pairs of balls $\{B_i, B_j\} \subseteq S$, each of which is associated with a ball $B_{\{i,j\}} \notin S$ of G which intersects only B_i and B_j in S. Then the graph $H = (S, E_S)$ is a minor of G.*

Proof. Let us fix a total ordering $\prec$ on the vertices of G. As before, all distances are in the graph G, and we write $d(u, v)$ instead of $d_G(u, v)$. For

[1] Despite our best effort, we have not been able to prove the two results at once in a satisfactory way, i.e. with a proof that would be both readable and shorter than the concatenation of the two existing proofs.

every $\{B_i, B_j\} \in E_S$ we write B_{ij} or B_{ji} interchangeably for $B_{\{i,j\}}$, and we denote by x_{ij} the center of the ball B_{ij}, and by r_{ij} its radius ($x_{ij} = x_{ji}$ and $r_{ij} = r_{ji}$). We can assume that the centers x_{ij} are chosen so that the radii r_{ij} are minimal (among all balls of G not in S that intersect only B_i and B_j in S).

We let $P(s_i, x_{ij})$ be the shortest path from s_i to x_{ij} which minimizes the sequence of vertices from s_i to x_{ij} with respect to the lexicographic ordering induced by $\prec$ (among all shortest paths from s_i to x_{ij}). Observe that $P(s_i, x_{ij})$ and $P(s_j, x_{ij})$ only intersect in x_{ij} (if not, we could replace x_{ij} by a vertex that is closer to s_i and s_j and reduce the radius r_{ij} accordingly – the new ball B_{ij} would still intersect B_i and B_j, and no other ball of S, and this would contradict the minimality of r_{ij}). We may also assume that $r_i + r_{ij} - 1 \leqslant d(s_i, x_{ij}) \leqslant r_i + r_{ij}$ (otherwise we can again move x_{ij} and decrease r_{ij} accordingly).

For every $i \in [n]$, we define

$$\mathcal{T}_i \coloneqq \bigcup_{j\,:\,\{B_i,B_j\}\in E_S} P(s_i, x_{ij}).$$

Claim 1. For every $i \in [n]$, $\mathcal{T}_i$ is a tree.

The proof is exactly the same as that of Claim 1 in the proof of Lemma 5.5 (we do not repeat it here).

On the path $P(s_i, x_{ij})$, we let $z_{i,ij}$ be the vertex at distance r_i from s_i (and since $x_{ij} = x_{ji}$ we use $z_{i,ij}$ and $z_{i,ji}$ interchangeably). Note that $r_{ij} - 1 \leqslant d(x_{ij}, z_{i,ij}) \leqslant r_{ij}$, since otherwise we could move x_{ij} and decrease r_{ij} accordingly. In particular, $d(x_{ij}, z_{j,ij}) - 1 \leqslant d(x_{ij}, z_{i,ij}) \leqslant d(x_{ij}, z_{j,ij}) + 1$.

Claim 2. For two pairs of balls $\{B_i, B_k\}, \{B_j, B_\ell\} \in E_S$, with $i \neq j$, if $P(s_i, x_{ik})$ and $P(s_j, x_{j\ell})$ intersect in some vertex y such that $d(y, z_{i,ik}) \leqslant d(y, z_{j,j\ell})$, then $i = \ell$ and $y = x_{ij}$.

We first argue that y appears after $z_{j,j\ell}$ when traversing $P(s_j, x_{j\ell})$ from s_j to $x_{j\ell}$. Indeed, otherwise we would have

$$d(s_j, z_{i,ik}) \leqslant d(s_j, y) + d(y, z_{i,ik}) \leqslant d(s_j, y) + d(y, z_{j,j\ell}) = d(s_j, z_{j,j\ell}) = r_j,$$

which means that B_i and B_j intersect, contradicting the assumptions that $i \neq j$ and all balls in S are vertex-disjoint. So y lies on the $z_{j,j\ell}$–$x_{j\ell}$ section of $P(s_j, x_{j\ell})$, and we infer that

$$d(x_{j\ell}, z_{i,ik}) \leqslant d(x_{j\ell}, y)+d(y, z_{i,ik}) \leqslant d(x_{j\ell}, y)+d(y, z_{j,j\ell}) = d(x_{j\ell}, z_{j,j\ell}) \leqslant r_{j\ell}.$$

It follows that the ball $B_{j\ell}$ intersects the ball B_i. By the assumption, this means that $i = \ell$, and thus $s_\ell = s_i$ and $z_{j,j\ell} = z_{j,ij}$. We now argue that y lies in the $z_{i,ik}$–x_{ik} section of $P(s_i, x_{ik})$. Suppose for a contradiction that y appears strictly before $z_{i,ik}$ when traversing $P(s_i, x_{ik})$ from s_i to x_{ik}. By definition of $z_{i,ik}$, it then follows that $d(s_i, y) \leqslant r_i - 1$. On the other hand

$$d(s_j, y) = d(s_j, x_{ij}) - d(y, x_{ij}) = d(s_j, z_{j,ij}) + d(z_{j,ij}, x_{ij}) - d(y, x_{ij}).$$

Since $d(z_{j,ij}, x_{ij}) \leqslant d(z_{i,ij}, x_{ij}) + 1 \leqslant d(y, x_{ij}) + 1$, it follows that $d(s_j, y) \leqslant d(s_j, z_{j,ij}) + 1 = r_j + 1$. Hence $d(s_i, s_j) \leqslant d(s_i, y) + d(y, s_j) \leqslant r_i + r_j$, so B_i and B_j intersect, a contradiction. We conclude that y lies in the $z_{i,ik}$–x_{ik} section of $P(s_i, x_{ik})$, and thus $d(x_{ik}, y) + d(y, z_{i,ik}) = d(x_{ik}, z_{i,ik})$.

Recall that by the initial assumption of the claim, combined with $i = \ell$, we have $d(y, z_{i,ik}) \leqslant d(y, z_{j,ij})$. Assume first that $d(y, z_{i,ik}) = d(y, z_{j,ij})$. Then

$$d(x_{ik}, z_{j,ij}) \leqslant d(x_{ik}, y) + d(y, z_{j,ij}) = d(x_{ik}, y) + d(y, z_{i,ik}) \leqslant r_{ik},$$

which implies that B_j intersects B_{ik}. Thus $j = k$, $P(s_i, x_{ik}) = P(s_i, x_{ij})$, and $P(s_j, x_{j\ell}) = P(s_j, x_{ij})$. Since these two paths have only x_{ij} in common, in this case we conclude that $y = x_{ij}$. We can now assume that $d(y, z_{i,ik}) \leqslant d(y, z_{j,ij}) - 1$. Recall that by definition of x_{ij}, we have $d(x_{ij}, z_{i,ij}) \geqslant d(x_{ij}, z_{j,ij}) - 1$, which implies that

$$d(y, z_{i,ik})+d(y, x_{ij}) \leqslant d(y, z_{j,ij})-1+d(y, x_{ij}) = d(z_{j,ij}, x_{ij})-1 \leqslant d(z_{i,ij}, x_{ij}).$$

Since $z_{i,ik}$ and $z_{i,ij}$ are both at distance r_i from s_i and $P(s_i, x_{ij})$ is a shortest path from s_i to x_{ij}, it follows that the concatenation of the s_i–y section of $P(s_i, x_{ik})$ and the y–x_{ij} section of $P(s_j, x_{ij})$ is a shortest path from s_i to x_{ij} (containing y). As y is also on a shortest path from s_j to x_{ij}, if we had $d(y, x_{ij}) > 0$, then we could replace x_{ij} by y and reduce r_{ij} to $r_{ij} - d(y, x_{ij})$ (B_{ij} would still intersect B_i and B_j and only these balls of S), which would contradict the minimality of r_{ij}. It follows that $y = x_{ij}$, as desired.

As in the proof of Lemma 5.5, the claim implies that for $i \neq j \in [n]$, $\mathcal{T}_i \cap \mathcal{T}_j = \{x_{ij}\}$ if $\{B_i, B_j\} \in E_S$, and otherwise the trees $\mathcal{T}_i$ and $\mathcal{T}_j$ are vertex-disjoint. Another direct consequence is that for every $\{B_i, B_j\} \in E_S$, the vertex x_{ij} is a leaf in at least one of the two trees $\mathcal{T}_i$ and $\mathcal{T}_j$. As before, we can contract the edges of each tree $\mathcal{T}_i$ not incident to a leaf of $\mathcal{T}_i$, and the resulting graph is precisely a graph obtained from $H = (S, E_S)$ by subdividing each edge at most once, and thus H is a minor of G. □

5.3 Hypergraphs and density

A *partial hypergraph* of $\mathcal{H}$ is a hypergraph obtained from $\mathcal{H}$ by removing a (possibly empty) subset of the edges. In addition to hypergraphs, it will also be convenient to consider *multi-hypergraphs*, i.e. hypergraphs $\mathcal{H} = (V, \mathcal{E})$ where $\mathcal{E}$ is a *multiset* of edges. The *rank* of a hypergraph or multi-hypergraph $\mathcal{H}$ is the maximum cardinality of an edge of $\mathcal{H}$.

We start with a useful tool, inspired by [46] (see also [14]), itself inspired by the Crossing lemma. Given a graph $G = (V, E)$, we denote by $\mathrm{ad}(G)$ the average degree of G, that is $\mathrm{ad}(G) = 2|E|/|V|$.

Lemma 5.7. *Let $\mathcal{H} = (V, \mathcal{E})$ be a multi-hypergraph of rank at most $k \geqslant 2$ on n vertices, and let $E \subseteq \binom{V}{2}$ be a set of pairs of vertices $\{u, v\}$ of V such that there exists an edge e_{uv} of $\mathcal{H}$ containing u and v. (Note that we allow that $e_{uv} = e_{xy}$ for two different pairs $\{u, v\}$ and $\{x, y\}$.) Then the graph (V, E) contains a subgraph H such that $\mathrm{ad}(H) \geqslant \frac{2|E|}{nek}$ and for every edge uv of H, the corresponding edge e_{uv} of $\mathcal{H}$ contains no vertex from $V(H) - \{u, v\}$.*

Proof. Let $\mathbf{H}$ be the (random) graph obtained by selecting each vertex of $\mathcal{H}$ independently with probability $1/k$, and keeping a single edge (of cardinality 2) between u and v whenever the only selected vertices of e_{uv} are u and v. Then we have

$$\mathbb{E}\left(|V(\mathbf{H})|\right) = \frac{n}{k}, \quad \text{and}$$
$$\mathbb{E}\left(|E(\mathbf{H})|\right) \geqslant |E| \cdot \frac{1}{k^2}\left(1 - \frac{1}{k}\right)^{k-2} \geqslant \frac{|E|}{\mathrm{e}k^2},$$

since $k \geqslant 2$. It follows that $\mathbb{E}\left(2|E(\mathbf{H})| - \frac{2|E|}{nek}|V(\mathbf{H})|\right) \geqslant 0$. In particular, there exists a subgraph H of (V, E) such that $\mathrm{ad}(H) \geqslant \frac{2|E|}{nek}$ and for every edge uv of H, the edge e_{uv} of $\mathcal{H}$ contains no vertex from $V(H) - \{u, v\}$, as desired. □

The proof actually gives a randomized algorithm producing the graph H. This algorithm can easily be derandomized using the method of conditional expectations, giving a deterministic algorithm running in time $O(|E| + n)$.

Given a hypergraph $\mathcal{H}$ and a matching $\mathcal{B}$ in $\mathcal{H}$, we define the *packing-hypergraph* $\mathcal{P}(\mathcal{H}, \mathcal{B})$ as the hypergraph with vertex set $\mathcal{B}$, in which a subset $\mathcal{B}' \subseteq \mathcal{B}$ is an edge if some edge of $\mathcal{H}$ intersects all the edges in $\mathcal{B}'$ and no other edge of $\mathcal{B}$.

Lemma 5.8. *Let G be a graph such that each minor of G has average degree at most d, let $\mathcal{H}$ be a ball hypergraph of G, and let $\mathcal{B}$ be a matching of size n in $\mathcal{H}$. For every integer $k \geqslant 2$, the number of edges of cardinality at most k in the packing-hypergraph $\mathcal{P}(\mathcal{H}, \mathcal{B})$ is at most*

$$(1 + dek)^{k-1} \cdot n.$$

Proof. Let $\mathcal{P}'$ be the partial hypergraph of $\mathcal{P}(\mathcal{H}, \mathcal{B})$ induced by the edges of cardinality at most k. Let H be the graph with vertex set $\mathcal{B}$ in which two distinct vertices are adjacent if they are contained in an edge of $\mathcal{P}'$ (i.e. an edge of $\mathcal{P}(\mathcal{H}, \mathcal{B})$ of cardinality at most k). Let m be the number of edges of H. Applying Lemma 5.7 to $\mathcal{P}'$, we obtain a subgraph H' of H of average degree at least $\frac{2m}{nek}$, and such that for any pair x, y of adjacent vertices in H', there is an edge of $\mathcal{P}'$ that contains x and y and no other vertex of H'. The vertices of H' correspond to a subset S of pairwise disjoint balls of G (since $\mathcal{B}$ is a matching), and each edge of H' corresponds to a ball of G that intersects some pair of balls of S (and does not intersect any other ball of S).

By Lemma 5.6, H' is a minor of G, so in particular $\frac{2m}{nek} \leqslant \mathrm{ad}(H') \leqslant d$, and hence $m \leqslant \frac{1}{2}\, dekn$. It follows that H contains a vertex of degree at most dek, and the same is true for every induced subgraph of H (since we can replace $\mathcal{B}$ in the proof by any subset of $\mathcal{B}$). As a consequence, H is $\lfloor dek \rfloor$-degenerate. It is a folklore result that ℓ-degenerate graphs on n vertices have at most $\binom{\ell}{t-1} n$ cliques of size t (see for instance [80, Lemma 18], where the proof gives a linear time algorithm to enumerate all the cliques of size t when t and ℓ are fixed), and hence there are at most

$$n \cdot \sum_{i=1}^{k} \binom{\lfloor dek \rfloor}{i-1} \leqslant n \cdot (1 + dek)^{k-1}$$

cliques of size at most k in H, which is an upper bound on the number of edges of cardinality at most k in $\mathcal{P}(\mathcal{H}, \mathcal{B})$. □

Note that the proof gives an $O(n)$ time algorithm enumerating all edges of cardinality at most k in the packing-hypergraph $\mathcal{P}(\mathcal{H}, \mathcal{B})$, when k and d are fixed (note that since H is $\lfloor dek \rfloor$-degenerate, it contains a linear number of edges, and thus the application of Lemma 5.7 takes $O(n)$ time).

5.4 Fractional packings of balls

We now prove Theorem 5.4. The proof is inspired by ideas from [63].

Proof of Theorem 5.4. Let $\mathcal{H}$ be a ball hypergraph of G. Since $\nu^*(\mathcal{H})$ is attained and is a rational number (recall that $\nu^*(\mathcal{H})$ is the solution of a linear program with integer coefficients), there exists a multiset $\mathcal{B}$ of p balls of G, such that every vertex $v \in V(G)$ is contained in at most q balls of $\mathcal{B}$, and $\nu^*(\mathcal{H}) = p/q$ (see for instance [63], where the same argument is applied to fractional cycle packings). We may assume that q is arbitrarily large (by taking k copies of each ball of $\mathcal{B}$, with multiplicities, for some arbitrarily large constant k), so in particular we may assume that $q \geqslant 2$. We may also assume that G contains at least one edge (i.e. $d \geqslant 1$), otherwise the result clearly holds. Enumerate all the balls in $\mathcal{B}$ as $B_1, B_2, \ldots, B_p$ (and recall that since $\mathcal{B}$ is a multiset, some balls B_i and B_j might coincide). We may assume that there is no pair of balls B_i, B_j such that $B_i \subsetneq B_j$ (otherwise we can replace B_j by B_i in $\mathcal{B}$, and we still have a fractional matching). It follows that the balls of $\mathcal{B}$ are pairwise incomparable (as defined at the beginning of Section 5.2). For any two intersecting balls B_i and B_j we define x_{ij} as a median vertex of B_i and B_j (also defined at the beginning of Section 5.2). Recall that it implies in particular that whenever B_i and B_j intersect, $x_{ij} \in B_i \cap B_j$, and if B_i and B_j coincide then x_{ij} is the center of B_i and B_j.

We let $\mathcal{G}$ be the intersection graph of the balls in $\mathcal{B}$, that is $V(\mathcal{G}) = \mathcal{B}$ and two vertices $B_i, B_j \in \mathcal{B} = V(\mathcal{G})$ with $i \neq j$ are adjacent in $\mathcal{G}$ if and only if $B_i \cap B_j \neq \varnothing$. (In particular, there is an edge linking B_i and B_j when B_i and B_j are two copies of the same ball.) Let m be the number of edges of $\mathcal{G}$. Let $\mathcal{B}^*$ denote the multi-hypergraph with vertex set $\mathcal{B}$, where for every vertex of G of the form x_{ij} there is a corresponding edge consisting of the balls in $\mathcal{B}$ that contain x_{ij}. Note that two distinct such vertices could possibly define the same edge, which is why edges in $\mathcal{B}^*$ could have multiplicities greater than 1. The multi-hypergraph $\mathcal{B}^*$ has rank at most q and contains p vertices. Note moreover that the number of pairs of vertices B_i, B_j of $\mathcal{B}^*$

with $i \neq j$ such that there exists an edge of $\mathcal{B}^*$ containing B_i and B_j is precisely m.

By Lemma 5.7 applied to the multi-hypergraph $\mathcal{B}^*$, we obtain a graph $H = (S, E_S)$ satisfying the following properties:

- $S \subseteq \mathcal{B}$;
- for each edge $\{B_i, B_j\} \in E_S$, x_{ij} is contained in B_i and B_j but in no other ball from S, and
- $\mathrm{ad}(H) \geqslant \frac{2m}{peq}$.

We would like to apply Lemma 5.5 to H but this is not immediately possible, since some balls of S might coincide (recall that $\mathcal{B}$ is a multiset), and therefore the centers of the balls of S might not be pairwise distinct. However, observe that if two balls of S coincide, then by definition the two corresponding vertices of H have degree either 0 or 1 in H (and in the latter case the two vertices are adjacent in H). Indeed, if two balls B_i, B_j of S coincide and B_i is adjacent to B_k in H with $k \neq j$, then the only balls of S containing x_{ik} are B_i and B_k, contradicting the fact that x_{ik} is also in B_j.

Let $S_1 \subseteq S$ be the subset of balls of S having multiplicity 1 in S. Since no ball of $\mathcal{B}$ is a strict subset of another ball of $\mathcal{B}$, the centers of the balls of S_1 are pairwise distinct. As a consequence of the previous paragraph, if we consider the subgraph H_1 of H induced by S_1, then $\mathrm{ad}(H) \leqslant \max(1, \mathrm{ad}(H_1))$.

By Lemma 5.5 applied to the set of balls S_1 in G, we obtain that H_1 is a minor of G and thus $\mathrm{ad}(H_1) \leqslant d$. It follows that $\frac{2m}{peq} \leqslant \mathrm{ad}(H) \leqslant \max(1, d) \leqslant d$ (since $d \geqslant 1$). This implies that the average degree $2m/p$ of $\mathcal{G}$ is at most $\mathrm{e}\, dq$. By the Caro-Wei inequality [17, 77] (or Turán's theorem [74]), it follows that $\mathcal{G}$ contains an independent set of size at least

$$\frac{|V(\mathcal{G})|}{\mathrm{ad}(\mathcal{G}) + 1} \geqslant \frac{p}{\mathrm{e}\, dq + 1} = \frac{\nu^*(\mathcal{H})}{\mathrm{e}\, d + 1/q}.$$

An independent set in $\mathcal{G}$ is precisely a matching in $\mathcal{H}$, and thus $\nu(\mathcal{H}) \geqslant \frac{1}{\mathrm{e}\, d + 1/q} \cdot \nu^*(\mathcal{H})$ and $\nu^*(\mathcal{H}) \leqslant (\mathrm{e}\, d + 1/q) \cdot \nu(\mathcal{H})$. Since we can assume that q is arbitrarily large, it follows that $\nu^*(\mathcal{H}) \leqslant \mathrm{e}\, d \cdot \nu(\mathcal{H})$, as desired.

The rest of the result follows from well known results on the average degree of graphs. On the one hand, an easy consequence of Euler's formula is that planar graphs have average degree at most 6. On the other hand, it was proved by Kostochka [54] and Thomason [72] that every K_t-minor-free graph has average degree $O(t\sqrt{\log t})$. □

The linear program for ν^* has coefficients in $\{0,1\}$, and can thus be solved in time $O(n^3)$, since we can assume that the balls have pairwise distinct centers (and so the number of variables and inequalities is linear in the number of vertices). The associated rational coefficients w_e can thus be found in $O(n^3)$. It is then convenient to define w'_e as the largest $\frac{\ell}{n} \leqslant w_e$ with $\ell \in \mathbb{N}$. Note that the coefficients (w'_e) still satisfy the inequalities of the linear program for ν^*, and their sum is at least $\nu^* - 1$ since we can assume that there are at most n balls (since there centers are pairwise distinct). There is a small loss on the multiplicative constant (compared to the statement of Theorem 5.4), but we can now assume that in the proof we have $q \leqslant n$ and thus $p \leqslant n^2$ and $m = O(n^3)$. It follows that the application of Lemma 5.7 can be done in $O(m) = O(n^3)$ time, and the construction of a stable set of suitable size in $\mathcal{G}$ can also be done in $O(m) = O(n^3)$ time. Therefore, the proof of Theorem 5.4 gives an $O(n^3)$ time algorithm constructing a matching of size $\Omega(\nu^*(\mathcal{H}))$ in $\mathcal{H}$.

The *VC-dimension* of a hypergraph $\mathcal{H}$ is the cardinality of a largest subset X of vertices such that for every $X' \subseteq X$, there is an edge e in $\mathcal{H}$ such that $e \cap X = X'$. Bousquet and Thomassé [12] proved the following result.

Theorem 5.9. *If G has no K_t-minor, then every ball hypergraph $\mathcal{H}$ of G has VC-dimension at most $t-1$.*

A classical result is that for hypergraphs of bounded VC-dimension, $\tau = O(\tau^* \log \tau^*)$. We will use the following precise bound of Ding, Seymour, and Winkler [32].

Theorem 5.10. *If a hypergraph $\mathcal{H}$ has VC-dimension at most δ, then*

$$\tau(\mathcal{H}) \leqslant 2\delta\tau^*(\mathcal{H}) \log(11\tau^*(\mathcal{H})).$$

Combining Theorems 5.4, 5.9, and 5.10, and using that $\nu^*(\mathcal{H}) = \tau^*(\mathcal{H})$, we obtain Theorem 5.3 as a direct consequence.

As before, the linear program for τ^* has coefficients in $\{0,1\}$, and can thus be solved in time $O(n^3)$, since we can assume that the balls have pairwise distinct centers (and so the number of variables and inequalities is linear in the number of vertices). The associated rational coefficients w_v can thus be found in time $O(n^3)$. Using algorithmic versions of Theorem 5.10 (see [50, 59]) and the coefficients (w_v), a transversal of $\mathcal{H}$ of size $O(\tau^* \log \tau^*) = O(\nu \log \nu)$ can be found by a randomized algorithm sampling $O(\tau^* \log \tau^*)$ vertices according to the distribution given by (w_v), or a deterministic algorithm running in time $O(n({\tau^*}^2 \log \tau^*)^t)$. So the overall complexity of obtaining a transversal of the desired size is $O(n^3)$ (randomized) and $O(n^3 + n(\tau^* \log \tau^*)^t)$ (deterministic). In the remainder of the paper, the result will be used when τ^* is a fixed constant, in which case the complexity of the deterministic algorithm is also $O(n^3)$.

5.5 Linear bound

In this section we prove Theorem 5.2. Recall that by Theorem 5.3, there is a (monotone) function f_t such that $\tau(\mathcal{H}) \leqslant f_t(\nu(\mathcal{H}))$ for every ball hypergraph $\mathcal{H}$ of a K_t-minor-free graph. In the proof, we write d_t for the supremum of the average degree of G taken over all graphs G excluding K_t as a minor. Recall that $d_t = O(t\sqrt{\log t})$ [54, 72].

Let $t \geqslant 1$ be an integer and let $c_t := 2 \cdot (1 + \frac{3}{2} d_t^2 \mathrm{e})^{3d_t/2} \cdot f_t(\frac{3}{2} d_t)$. We will prove that every ball hypergraph $\mathcal{H}$ of a K_t-minor-free graph satisfies $\tau(\mathcal{H}) \leqslant c_t \cdot \nu(\mathcal{H})$.

Proof of Theorem 5.2. We prove the result by induction on $k := \nu(\mathcal{H})$. The result clearly holds if $k = 0$ so we may assume that $k \geqslant 1$. If $k \leqslant \frac{3}{2} d_t$ then by the definition of f_t we have $\tau(\mathcal{H}) \leqslant f_t(\frac{3}{2} d_t) \leqslant c_t \leqslant c_t \cdot k$, as desired.

Assume now that $k \geqslant \frac{3}{2} d_t$ and for every ball hypergraph $\mathcal{H}'$ of a K_t-minor-free graph with $\nu(\mathcal{H}') < k$, we have $\tau(\mathcal{H}') \leqslant c_t \cdot \nu(\mathcal{H}')$. Let G be a K_t-minor-free graph and $\mathcal{H}$ be a ball hypergraph of G with $\nu(\mathcal{H}) = k$. Our goal is to show that $\tau(\mathcal{H}) \leqslant c_t \cdot k$. Note that we can assume that $\mathcal{H}$ is *minimal*, in the sense that no edge of $\mathcal{H}$ is contained in another edge of $\mathcal{H}$ (otherwise we can remove the larger of the two from $\mathcal{H}$, this does not change the matching number nor the transversal number).

Consider a maximum matching $\mathcal{B}$ (of cardinality k) in $\mathcal{H}$. Let $\mathcal{E}_1$ be the set consisting of all the edges of $\mathcal{H}$ that intersect at most $\frac{3}{2} d_t$ edges of $\mathcal{B}$. By Lemma 5.8, the packing-hypergraph $\mathcal{P}(\mathcal{H}, \mathcal{B})$ contains at most $(1 + \frac{3}{2} d_t^2 \mathrm{e})^{3d_t/2} \cdot k$ edges of cardinality at most $\frac{3}{2} d_t$. For each such edge e of

$\mathcal{P}(\mathcal{H}, \mathcal{B})$, consider the corresponding subset $\mathcal{B}_e$ of at most $\frac{3}{2}d_t$ edges of $\mathcal{B}$, and the subset $\mathcal{E}_e$ of edges of $\mathcal{H}$ that intersect each ball of $\mathcal{B}_e$, and no other ball of $\mathcal{B}$. Denoting by $\mathcal{H}_e$ the partial hypergraph of $\mathcal{H}$ with edge set $\mathcal{E}_e$, observe that by the maximality of the matching $\mathcal{B}$ we have $\nu(\mathcal{H}_e) \leqslant \frac{3}{2}d_t$ (since in $\mathcal{B}$, replacing the edges of $\mathcal{B}_e$ by a matching of $\mathcal{E}_e$ again gives a matching of $\mathcal{H}$). It follows that $\tau(\mathcal{H}_e) \leqslant f(\frac{3}{2}d_t)$. And thus, if we denote by $\mathcal{H}_1$ the partial hypergraph of $\mathcal{H}$ with edge set $\mathcal{E}_1$, we have

$$\tau(\mathcal{H}_1) \leqslant (1 + \tfrac{3}{2}d_t^2\mathrm{e})^{3d_t/2} \cdot f(\tfrac{3}{2}d_t) \cdot k = \tfrac{1}{2}c_t \cdot k.$$

Consider now the subset $\mathcal{E}_2$ consisting of all the edges of $\mathcal{H}$ that intersect more than $\frac{3}{2}d_t$ edges of $\mathcal{B}$, and let $\mathcal{H}_2$ be the partial hypergraph of $\mathcal{H}$ with edge set $\mathcal{E}_2$. Note that $\mathcal{E}_1$ and $\mathcal{E}_2$ partition the edge set of $\mathcal{H}$ and thus $\tau(\mathcal{H}) \leqslant \tau(\mathcal{H}_1) + \tau(\mathcal{H}_2)$. Let $\mathcal{B}_2$ be a maximum matching in $\mathcal{H}_2$, and let $\ell = \nu(\mathcal{H}_2) = |\mathcal{B}_2|$. Let H be the (bipartite) intersection graph of the edges of $\mathcal{B} \cup \mathcal{B}_2$, i.e. each vertex of H corresponds to an edge of $\mathcal{B} \cup \mathcal{B}_2$, and two vertices are adjacent if the corresponding edges intersect. (The graph is bipartite because $\mathcal{B}$ and $\mathcal{B}_2$ are matchings.)

Note that since H is bipartite, for every two distinct edges $\{B, B'\}$ and $\{C, C'\}$ of H, the sets $B \cap B'$ and $C \cap C'$ are disjoint. Moreover, no ball of $\mathcal{B} \cup \mathcal{B}_2$ is a subset of another ball of $\mathcal{B} \cup \mathcal{B}_2$, and thus the balls of $\mathcal{B} \cup \mathcal{B}_2$ are pairwise incomparable (as defined at the beginning of Section 5.2). So, enumerating the balls in $\mathcal{B} \cup \mathcal{B}_2$ as $B_1, B_2, \ldots, B_n$, we can choose, for each edge $\{B_i, B_j\}$ of H, a median vertex x_{ij} of B_i and B_j (also defined at the beginning of Section 5.2). Recall that $x_{ij} \in B_i \cap B_j$, and thus it follows from the property above that the only balls of $\mathcal{B} \cup \mathcal{B}_2$ containing x_{ij} are B_i and B_j. By Lemma 5.5, H is a minor of G and thus has average degree at most d_t. On the other hand, the vertices of H corresponding to the edges of $\mathcal{B}_2$ have degree at least $\frac{3}{2}d_t$ in H, and thus

$$\tfrac{3}{2}d_t \cdot \ell \leqslant \tfrac{1}{2}\,\mathrm{ad}(H)(k + \ell) \leqslant \tfrac{1}{2}d_t \cdot (k + \ell),$$

where the central term counts the number of edges of H. It follows that $\nu(\mathcal{H}_2) = \ell \leqslant \frac{k}{2}$, and thus by the induction hypothesis we have $\tau(\mathcal{H}_2) \leqslant c_t \cdot \nu(\mathcal{H}_2) \leqslant c_t \cdot \frac{k}{2}$. As a consequence,

$$\tau(\mathcal{H}) \leqslant \tau(\mathcal{H}_1) + \tau(\mathcal{H}_2) \leqslant \tfrac{1}{2}c_t \cdot k + c_t \cdot \tfrac{k}{2} = c_t \cdot k,$$

which concludes the proof of Theorem 5.2. $\square$

The first part of the proof of Theorem 5.2 uses Theorem 5.3 when ν (and thus τ^*, by Theorem 5.4) is bounded by a function of the constant t, and in this case, by the discussion after the proof of Theorem 5.3, a transversal of the desired size can be found deterministically in time $O(n^3)$.

The second part of the proof of Theorem 5.2 can be made constructive by performing the following small modification. We observe that we have not quite used the fact that $\mathcal{B}$ is a *maximum* matching of $\mathcal{H}$, simply that it has the property that, for any edge e in the packing-hypergraph $\mathcal{P}(\mathcal{H}, \mathcal{B})$ of cardinality at most $\frac{3}{2}d_t$, the matching number of $\mathcal{H}_e$ is bounded. As we have explained after Lemma 5.8, such edges can be enumerated in linear time when t is fixed. We can then compute each $\tau^*(\mathcal{H}_e) = \nu^*(\mathcal{H}_e)$ in time $O(n^3)$ and if this value is more than $\mathrm{e}\, d_t \cdot |e|$, then we can find a matching of size more than $|e| = |\mathcal{B}_e|$ in $\mathcal{H}_e$ in time $O(n^3)$ by Theorem 5.4, and replace $\mathcal{B}_e$ by this larger matching in $\mathcal{B}$, thus increasing the size of $\mathcal{B}$ (this can be done at most $\nu(\mathcal{H})$ times). On the other hand, if for all the (linearly many) edges e as above, we have $\tau^*(\mathcal{H}_e) \leqslant \mathrm{e}\, d_t \cdot |e| = O({d_t}^2)$, then by Theorem 5.3, we can find a transversal of size $O({d_t}^2 \log d_t)$ in each hypergraph $\mathcal{H}_e$ in time $O(n^3)$. So overall we find a matching $\mathcal{B}$ that has the desired property, and a transversal of the partial hypergraph of $\mathcal{H}$ with edge set $\mathcal{E}_1$ of the desired size in time $O(\nu(\mathcal{H}) \cdot n^4)$. Taking the induction step into account (which divides ν by at least 2), we obtain a deterministic algorithm constructing a transversal of size $O(\nu(\mathcal{H}))$ in $\mathcal{H}$, in time $O(\sum_{i \geqslant 0} \frac{1}{2^i} \cdot \nu(\mathcal{H}) \cdot n^4) = O(\nu(\mathcal{H}) \cdot n^4)$, when t is a fixed constant.

5.6 Conclusion

The proof of Theorem 5.2 gives a bound of the order of $\exp(t \log^{3/2} t)$ for the constant c_t. It would be interesting to improve this bound to a polynomial in t.

It is also natural to wonder whether Theorem 5.2 remains true in a setting broader than proper minor-closed classes. Natural candidates are graphs of bounded maximum degree, graphs excluding a topological minor, k-planar graphs, classes with polynomial growth (meaning that the size of each ball is bounded by a polynomial function of its radius, see e.g. [55]), and classes with strongly sublinear separators (or equivalently, classes with polynomial expansion [35]). We now observe that in all these cases, the associated ball hypergraphs do not satisfy the Erdős-Pósa property, even if all the balls have the same radius. That is, we can find r-ball hypergraphs in these classes with bounded ν and unbounded τ. Our construction shows that this is true even in the seemingly simple case of subgraphs of a grid with all diagonals (i.e. strong products of two paths).

Fix two integers k, ℓ with $k \geqslant 3$, and ℓ sufficiently large compared to k and divisible by $2(\binom{k}{2} - 1)$. Given k vertices $v_0, v_1, \dots v_{k-1}$, an *ℓ-broom with root v_0 and leaves $v_1, \dots, v_{k-1}$* is a tree T of maximum degree 3 with root v_0 and leaves $v_1, \dots, v_{k-1}$ such that

1. each leaf is at distance ℓ from the root v_0,
2. the ball of radius $\ell/2$ centered in v_0 in T is a path (called the *handle* of the broom), and
3. the distance between every two vertices of degree 3 in T is sufficiently large compared to k.

We now construct a graph $G_{k,\ell}$ as follows. We start with a set X of k vertices $x_1, \dots, x_k$, and a path of $\binom{k}{2}$ vertices with vertex set $Y = \{y_{\{i,j\}} \mid 1 \leqslant i < j \leqslant k\}$, disjoint from X. We then subdivide each edge of the latter path $\frac{\ell}{2}\frac{1}{\binom{k}{2}-1} - 1$ times, so that the subdivided path has length $\ell/2$. Finally, for each $1 \leqslant i \leqslant k$, we add an ℓ-broom T_i with root x_i and leaves $Y_i = \{y_{\{i,j\}} \mid j \neq i\}$.

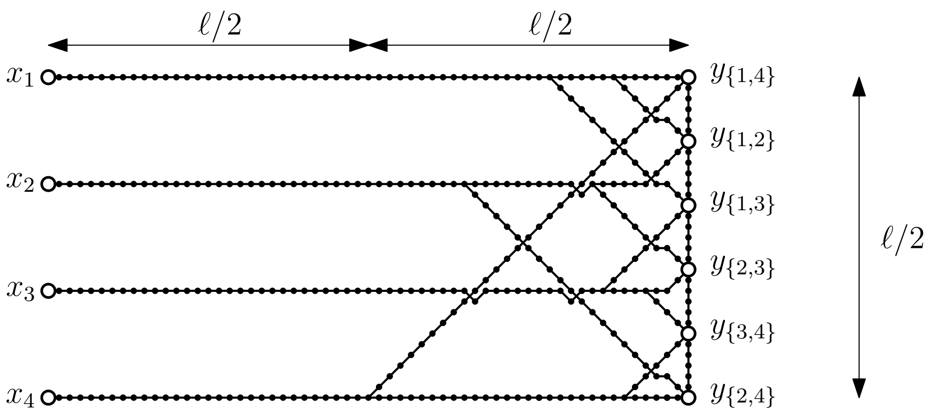

Figure 5.1. An embedding of the graph $G_{4,\ell}$ in the 2-dimensional grid with all diagonals (the grid itself is not depicted for the sake of clarity).

We first claim that $G_{k,\ell}$ is a subgraph of the 2-dimensional grid with all diagonals (i.e. the strong products of two paths). To see this, place X on a single column on the left, and Y on another column on the right (in the sequence given by the path), at distance ℓ from the column of X, then draw each of the brooms in the plane (with crossings allowed). Since the distance between two vertices of degree 3 in a broom is sufficiently large compared to k, we can safely embed each topological crossing in the strong product of two edges (see Figure 5.1 for an example).

Let $\mathcal{H}_{k,\ell}$ be the ℓ-ball hypergraph of $G_{k,\ell}$ obtained by considering all the balls of radius ℓ in $G_{k,\ell}$. We first observe that $\nu(\mathcal{H}_{k,\ell}) = 1$: this follows from the fact that each ball of radius ℓ centered in a vertex that does not belong to the handle of a broom contains all the vertices of Y, while every two vertices on the handles of two brooms T_i and T_j are at distance at most ℓ from $y_{\{i,j\}}$. Finally, for every two vertices x_i and x_j of X, note that $y_{\{i,j\}}$ is the unique vertex of $G_{k,\ell}$ lying at distance at most ℓ from x_i and x_j, and thus $\tau(\mathcal{H}_{k,\ell}) \geqslant \frac{k}{2}$. It follows that there is no function f such that $\tau(\mathcal{H}) \leqslant f(\nu(\mathcal{H}))$ for every ball hypergraph of a subgraph of the strong product of two paths (even when all the balls in the ball hypergraph have the same radius).

Bibliography

[1] I. Abraham, Y. Bartal, and O. Neiman. Advances in metric embedding theory. *Advances in Mathematics*, 228(6):3026–3126, 2011.

[2] D. L. Applegate, R. E. Bixby, V. Chvátal, and W. J. Cook. *The traveling salesman problem.* Princeton university press, 2011.

[3] K. Ball. Isometric embedding in lp-spaces. *European Journal of Combinatorics*, 11(4):305 – 311, 1990.

[4] F. Barahona, M. Grötschel, M. Jünger, and G. Reinelt. An application of combinatorial optimization to statistical physics and circuit layout design. *Operations Research*, 36(3):493–513, 1988.

[5] A. I. Barvinok. Problems of distance geometry and convex properties of quadratic maps. *Discrete & Computational Geometry*, 13(2):189–202, 1995.

[6] M. Belk. Realizability of graphs in three dimensions. *Discrete & Computational Geometry*, 37(2):139–162, 2007.

[7] M. Belk and R. Connelly. Realizability of graphs. *Discrete & Computational Geometry*, 37(2):125–137, 2007.

[8] M. Bonamy, C. Groenland, C. Muller, J. Narboni, J. Pekárek, and A. Wesolek. A note on connected greedy edge colouring. *Discrete Applied Mathematics*, 304:129–136, 2021.

[9] G. Borradaile and E. W. Chambers. Covering nearly surface-embedded graphs with a fixed number of balls. *Discrete & Computational Geometry*, 51(4):979–996, 2014.

[10] J. Bourgain. On Lipschitz embedding of finite metric spaces in Hilbert space. *Israel Journal of Mathematics*, 52(1):46–52, Mar 1985.

[11] N. Bousquet, W. Cames van Batenburg, L. Esperet, G. Joret, W. Lochet, C. Muller, and F. Pirot. Packing and covering balls in graphs excluding a minor. *Combinatorica*, 41(3):299–318, 2021.

[12] N. Bousquet and S. Thomassé. VC-dimension and Erdős–Pósa property. *Discrete Mathematics*, 338(12):2302–2317, 2015.

[13] P. J. Cameron. Problems from CGCS Luminy, May 2007. *European Journal of Combinatorics*, 31(2):644–648, 2010.

[14] W. Cames van Batenburg, L. Esperet, and T. Müller. Coloring Jordan regions and curves. *SIAM Journal on Discrete Mathematics*, 31(3):1670–1684, 2017.

[15] W. Cames van Batenburg, T. Huynh, G. Joret, and J.-F. Raymond. A tight Erdős-Pósa function for planar minors. *Advances in Combinatorics*, page 10807, 2019.

[16] W. Cames van Batenburg, G. Joret, and A. Ulmer. Erdős-Pósa from ball packing. *SIAM Journal on Discrete Mathematics*, 34(3):1609–1619, 2020.

[17] Y. Caro. New results on the independence number. Technical report, Technical Report, Tel-Aviv University, 1979.

[18] R. D. Carr, G. Konjevod, G. Little, V. Natarajan, and O. Parekh. Compacting cuts: a new linear formulation for minimum cut. *ACM Transactions on Algorithms (TALG)*, 5(3):1–16, 2009.

[19] F. Cecchetto. On the structure of minimum cuts of a graph, 2019.

[20] C. Chekuri and J. Chuzhoy. Polynomial bounds for the Grid-Minor Theorem. *Journal of the ACM (JACM)*, 63(5):40:1–40:65, 2016.

[21] V. Chepoi, B. Estellon, and G. Naves. Packing and covering with balls on Busemann surfaces. *Discrete & Computational Geometry*, 57(4):985–1011, 2017.

[22] V. Chepoi, B. Estellon, and Y. Vaxès. Covering planar graphs with a fixed number of balls. *Discrete & Computational Geometry*, 37(2):237–244, 2007.

[23] J. Chuzhoy and Z. Tan. Towards tight(er) bounds for the excluded grid theorem. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1445–1464. ACM.

[24] V. Chvatal and V. Chvatal. *Linear programming.* Macmillan, 1983.

[25] M. Conforti. Private communication, 2019.

[26] M. Conforti, S. Fiorini, and K. Pashkovich. Cut dominants and forbidden minors. *SIAM journal on discrete mathematics*, 30(3):1571–1589, 2016.

[27] M. Conforti, G. Rinaldi, and L. Wolsey. On the cut polyhedron. *Discrete Mathematics*, 277(1-3):279–285, 2004.

[28] G. Cornuéjols, J. Fonlupt, and D. Naddef. The traveling salesman problem on a graph and some related integer polyhedra. *Mathematical programming*, 33(1):1–27, 1985.

[29] M. Deza and M. Laurent. *Geometry of Cuts and Metrics.* Springer, 1997.

[30] G. Di Battista and R. Tamassia. On-line maintenance of triconnected components with SPQR-trees. *Algorithmica*, 15(4):302–318, Apr 1996.

[31] R. Diestel. *Graph Theory, 4th Edition*, volume 173 of *Graduate texts in mathematics.* Springer, 2012.

[32] G.-L. Ding, P. Seymour, and P. Winkler. Bounding the vertex cover number of a hypergraph. *Combinatorica*, 14(1):23–34, 1994.

[33] G. Ducoffe, M. Habib, and L. Viennot. Diameter computation on H-minor free graphs and graphs of bounded (distance) VC-dimension. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1905–1922. SIAM, 2020.

[34] Z. Dvořák. Constant-factor approximation of the domination number in sparse graphs. *European Journal of Combinatorics*, 34(5):833–840, 2013.

[35] Z. Dvorák and S. Norin. Strongly sublinear separators and polynomial expansion. *SIAM Journal on Discrete Mathematics*, 30(2):1095–1101, 2016.

[36] J. Edmonds and R. M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *Journal of the ACM (JACM)*, 19(2):248–264, 1972.

[37] M. Eisenberg-Nagy, M. Laurent, and A. Varvitsiotis. Forbidden minor characterizations for low-rank optimal solutions to semidefinite programs over the elliptope. *J. Comb. Theory, Ser. B*, 108:40–80, 2014.

[38] P. Erdös and G. Szekeres. A combinatorial problem in geometry. *Compositio Math.*, 2:463–470, 1935.

[39] P. Erdös and L. Pósa. On independent circuits contained in a graph. *Canadian Journal of Mathematics*, 17:347–352, 1965.

[40] B. Estellon. *Algorithmes de couverture et d'augmentation de graphes sous contraintes de distance.* PhD thesis, Aix-Marseille 2, 2007.

[41] S. Fiorini, T. Huynh, G. Joret, and C. Muller. Unavoidable minors for graphs with large ℓ_p-dimension. *Discrete & Computational Geometry*, 66(1):301–343, 2021.

[42] S. Fiorini, T. Huynh, G. Joret, and A. Varvitsiotis. The excluded minors for isometric realizability in the plane. *SIAM J. Discrete Math.*, 31(1):438–453, 2017.

[43] S. Fiorini, G. Joret, and D. R. Wood. Excluded forest minors and the Erdős–Pósa property. *Combinatorics, Probability and Computing*, 22(5):700–721, 2013.

[44] S. Fiorini, S. Massar, S. Pokutta, H. R. Tiwary, and R. D. Wolf. Exponential lower bounds for polytopes in combinatorial optimization. *Journal of the ACM (JACM)*, 62(2):1–23, 2015.

[45] L. R. Ford and D. R. Fulkerson. A simple algorithm for finding maximal network flows and an application to the Hitchcock problem. *Canadian journal of Mathematics*, 9:210–218, 1957.

[46] J. Fox and J. Pach. Touching strings. *manuscript, 2012.*

[47] C. Gavoille, D. Peleg, A. Raspaud, and E. Sopena. Small k-dominating sets in planar graphs with applications. In *International Workshop on Graph-Theoretic Concepts in Computer Science*, pages 201–216. Springer, 2001.

[48] S. O. Gharan, A. Saberi, and M. Singh. A randomized rounding approach to the traveling salesman problem. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 550–559. IEEE, 2011.

[49] R. E. Gomory and T. C. Hu. Multi-terminal network flows. *Journal of the Society for Industrial and Applied Mathematics*, 9(4):551–570, 1961.

[50] D. Haussler and E. Welzl. ϵ-nets and simplex range queries. *Discrete & Computational Geometry*, 2(2):127–151, 1987.

[51] W. Holsztynski. Rn as a universal metric space. In *Abstract 78T-G56, Notices AMS*, volume 25, 1978.

[52] G. Joret, C. Paul, I. Sau, S. Saurabh, and S. Thomassé. Hitting and harvesting pumpkins. *SIAM Journal on Discrete Mathematics*, 28(3):1363–1390, 2014.

[53] D. R. Karger and C. Stein. A new approach to the minimum cut problem. *Journal of the ACM (JACM)*, 43(4):601–640, 1996.

[54] A. V. Kostochka. Lower bound of the Hadwiger number of graphs by their average degree. *Combinatorica*, 4(4):307–316, 1984.

[55] R. Krauthgamer and J. R. Lee. The intrinsic dimensionality of graphs. *Combinatorica*, 27(5):551–585, 2007.

[56] M. Laurent and A. Varvitsiotis. A new graph parameter related to bounded rank positive semidefinite matrix completions. *Math. Program.*, 145(1-2):291–325, 2014.

[57] J. Li and M. Parter. Planar diameter via metric compression. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 152–163, 2019.

[58] S. Lörwald and G. Reinelt. PANDA: a software for polyhedral transformations. *EURO Journal on Computational Optimization*, pages 1–12, 2015.

[59] J. Matousek. Approximations and optimal geometric divide-and-conquer. *Journal of Computer and System Sciences*, 50(2):203–208, 1995.

[60] C. Muller. Excluded Minors for Isometric Embeddings of Graphs in ℓ_∞^k-spaces. Master's thesis, Université Libre de Bruxelles (ULB), 2017.

[61] H. Nagamochi and T. Ibaraki. Computing edge-connectivity in multigraphs and capacitated graphs. *SIAM Journal on Discrete Mathematics*, 5(1):54–66, 1992.

[62] M. Pierce. *Searching for and classifying the finite set of minor-minimal non-apex graphs*. PhD thesis, Citeseer, 2014.

[63] B. A. Reed and F. B. Shepherd. The Gallai-Younger conjecture for planar graphs. *Combinatorica*, 16(4):555–566, 1996.

[64] N. Robertson and P. D. Seymour. Graph minors. III. Planar tree-width. *Journal of Combinatorial Theory, Series B*, 36(1):49–64, 1984.

[65] N. Robertson and P. D. Seymour. Graph minors. V. Excluding a planar graph. *J. Combin. Theory Ser. B*, 41(1):92–114, 1986.

[66] N. Robertson and P. D. Seymour. Graph minors. XIII. The disjoint paths problem. *Journal of combinatorial theory, Series B*, 63(1):65–110, 1995.

[67] N. Robertson and P. D. Seymour. Graph minors. XX. Wagner's conjecture. *Journal of Combinatorial Theory, Series B*, 92(2):325–357, 2004.

[68] V. Rödl and A. Ruciński. Bipartite coverings of graphs. *Combinatorics, Probability and Computing*, 6(3):349–352, 1997.

[69] J. B. Saxe. Embeddability of weighted graphs in k-space is strongly NP-hard. In *Proc. of 17th Allerton Conference in Communications, Control and Computing, Monticello, IL*, pages 480–489, 1979.

[70] A. Schrijver. *Combinatorial optimization: polyhedra and efficiency*, volume 24. Springer Science & Business Media, 2003.

[71] A. Sebő and J. Vygen. Shorter tours by nicer ears: 7/5-approximation for the graph-TSP, 3/2 for the path version, and 4/3 for two-edge-connected subgraphs. *Combinatorica*, 34(5):597–629, 2014.

[72] A. Thomason. An extremal function for contractions of graphs. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 95, pages 261–265. Cambridge University Press, 1984.

[73] V. Traub, J. Vygen, and R. Zenklusen. Reducing path TSP to TSP. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 14–27, 2020.

[74] P. Turán. On an external problem in graph theory. *Mat. Fiz. Lapok*, 48:436–452, 1941.

[75] K. Wagner. Über eine Eigenschaft der ebenen Komplexe. *Mathematische Annalen*, 114(1):570–590, 1937.

[76] K. Wagner. Graphentheorie: BI-Hochschultaschenbücher. 1970.

[77] V. Wei. A lower bound on the stability number of a simple graph. Technical report, Bell Laboratories Technical Memorandum 81-11217-9, Murray Hill, NJ, 1981.

[78] D. P. Williamson and D. B. Shmoys. *The design of approximation algorithms*. Cambridge university press, 2011.

[79] H. Witsenhausen. Minimum dimension embedding of finite metric spaces. *Journal of Combinatorial Theory, Series A*, 42(2):184 – 199, 1986.

[80] D. R. Wood. Cliques in graphs excluding a complete graph minor. *The Electronic Journal of Combinatorics*, 23(3):P3–18, 2016.

[81] Y. Yu. More forbidden minors for wye-delta-wye reducibility. *The electronic journal of combinatorics*, pages R7–R7, 2006.

[82] G. M. Ziegler. Lectures on 0/1-polytopes. In *Polytopes—combinatorics and computation*, pages 1–41. Springer, 2000.

[83] G. M. Ziegler. *Lectures on polytopes*, volume 152. Springer Science & Business Media, 2012.

Appendix A

Excluded minors for ℓ_∞^3

We present 38 excluded minors for $f_\infty(G) \leqslant 3$. The graphs are shown in Figure A.1. The tables below are such that the first two columns represent the vertices of the graph which form an edge and the last column contains the length of that edge. For instance, the first graph G has vertices $V(G) = \{0, 1, 2, 3, 4, 5, 6, 7\}$ and edges

$$E(G) = \{01, 02, 03, 12, 13, 24, 25, 34, 35, 45, 26, 27, 36, 37, 67\}.$$

The edge 01 has length 74. These distance functions are such they cannot be isometrically embedded in ℓ_∞^3.

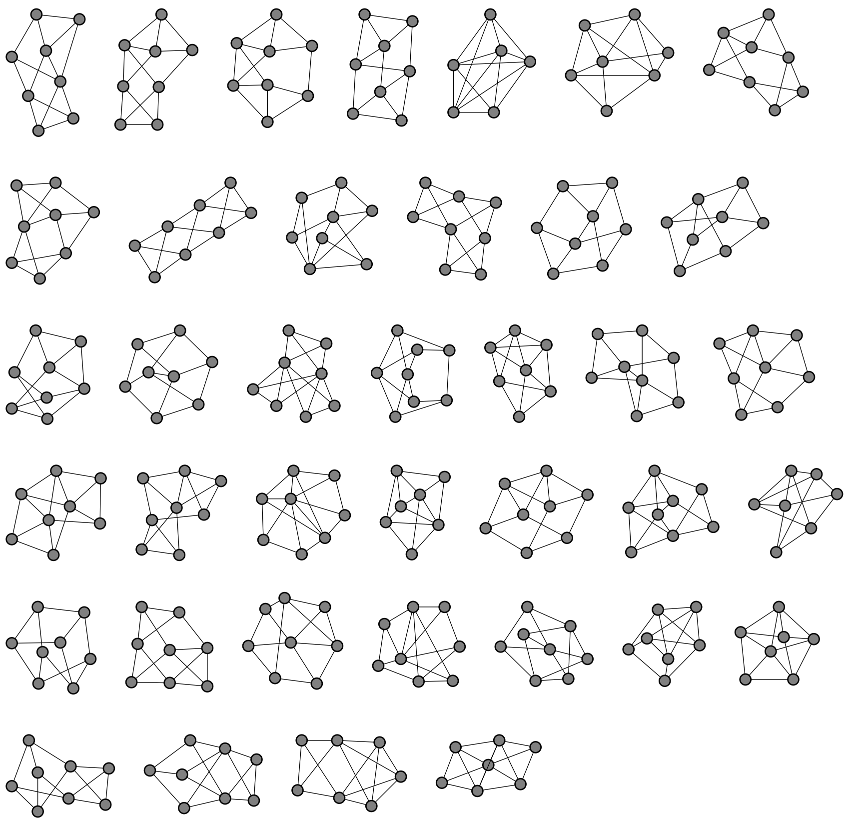

Figure A.1. The 38 known excluded graphs for $f_\infty(G) \leqslant 3$.

```
0 1 112
0 2 73
0 3 108
1 2 69
1 3 9
2 4 55
2 5 27
3 4 59
3 5 53
4 5 82
2 6 17
2 7 83
3 6 84
3 7 5
6 7 79
```

```
0 1 74
0 2 88
0 3 33
1 2 37
1 3 77
2 4 67
2 5 42
3 4 43
3 5 13
4 6 17
4 7 69
5 6 73
5 7 13
6 7 62
```

```
0 1 3
2 3 3
4 5 2
6 7 2
0 2 1
0 4 1
0 6 1
0 7 1
1 2 2
3 4 1
5 6 1
1 3 1
3 5 1
5 7 1
```

```
0 1 4
2 3 4
4 5 4
6 7 4
0 2 1
2 4 1
4 6 1
1 3 1
3 5 1
5 7 1
0 3 3
2 5 3
4 7 3
6 1 1
```

```
0 1 4
2 3 4
4 5 4
6 7 4
0 2 1
2 4 1
4 6 1
1 3 1
3 5 1
5 7 1
0 3 3
2 7 2
1 4 2
5 6 3
```

```
0 1 4
2 3 4
4 5 4
6 7 5
0 2 1
2 4 1
4 6 1
1 3 1
3 5 1
5 7 1
0 3 3
2 5 3
4 7 4
1 4 2
3 6 3
```

```
0 1 2
2 3 4
4 5 4
6 7 4
0 2 1
2 4 1
4 6 1
1 7 1
3 5 1
5 7 1
0 7 1
1 2 1
3 4 3
5 6 3
```

0 1 4	0 1 4	0 1 7	0 1 7
2 3 4	2 3 4	2 3 8	2 3 8
4 5 4	4 5 4	4 5 8	4 5 8
6 7 4	6 7 4	6 7 7	6 7 8
0 2 1	0 2 2	0 2 2	0 2 2
2 4 1	0 4 1	2 4 2	2 4 3
4 6 1	0 6 1	2 6 6	2 6 5
1 7 3	1 3 2	1 3 2	1 3 2
3 5 1	1 5 1	1 5 3	1 5 3
5 7 1	1 7 1	1 7 3	1 7 3
0 7 1	2 5 2	0 3 6	0 3 6
1 4 2	4 7 2	3 6 2	2 7 3
3 4 3	6 1 3	0 5 4	3 6 4
5 6 3	0 3 2	4 7 3	3 4 5
			2 5 5

0 1 68	0 1 68	0 1 19	0 1 44
1 2 85	0 2 62	0 2 10	0 2 20
2 0 55	0 6 96	1 2 9	1 2 25
0 3 72	0 4 58	0 3 13	0 3 27
0 4 94	1 2 12	1 3 13	1 3 17
0 5 13	2 3 53	2 4 9	2 4 49
0 6 65	3 4 28	3 4 18	3 4 18
1 3 128	4 1 23	2 5 33	2 5 20
1 4 109	5 1 56	3 5 12	3 5 51
2 5 68	5 2 48	4 5 24	4 5 40
2 6 11	5 3 5	6 0 24	6 5 35
2 4 46	6 2 72	6 1 20	6 1 35
4 5 105	6 3 62	6 4 13	6 2 34
5 6 64	6 4 89	6 3 11	6 3 27
3 4 22			

```
1 7 100
5 6 134
2 3 132
0 4 100
0 1 50
0 2 61
0 3 71
0 5 72
0 6 62
0 7 50
1 2 45
3 4 35
4 5 35
6 7 45
1 4 52
7 4 54
```

```
1 7 106
5 6 77
2 3 81
0 2 22
0 3 59
0 5 15
0 6 84
0 7 72
1 2 37
3 4 68
4 5 88
6 7 81
1 4 14
6 4 11
```

```
0 1 2
2 3 2
6 7 4
0 2 2
1 2 4
0 4 4
1 4 2
2 5 2
3 5 4
3 4 2
4 6 2
4 7 2
5 7 1
2 6 1
```

```
0 1 137
2 3 131
4 5 131
6 7 96
0 2 34
1 2 114
1 4 60
2 5 18
3 5 113
3 4 44
0 6 110
0 7 49
4 7 30
3 6 57
```

```
0 1 10
2 3 2
4 5 2
6 7 10
0 2 5
1 2 5
0 4 5
1 4 5
2 5 10
3 5 8
3 4 10
0 6 5
0 7 5
5 7 1
3 6 1
```

```
2 3 7
4 5 3
6 7 10
0 2 5
1 2 10
0 4 10
1 4 5
2 5 3
3 5 10
3 4 7
0 6 8
0 7 2
4 6 2
1 7 3
```

```
0 1 25
1 2 13
2 3 9
3 4 34
4 5 11
5 0 4
0 3 19
1 4 13
2 5 29
6 0 16
6 1 29
6 3 24
6 4 24
```

```
0 1 3
1 2 10
2 3 12
3 4 8
4 5 4
5 0 18
0 3 25
1 4 25
2 5 24
6 0 26
6 4 4
6 3 12
6 2 17
```

```
2 3 43
3 4 11
4 5 68
5 0 10
0 3 50
1 4 5
2 5 36
6 0 60
6 1 13
6 2 14
7 2 104
7 1 125
7 0 65
```

```
0 1 3
0 2 4
0 3 1
1 2 1
1 3 4
2 5 1
2 7 1
3 4 1
3 6 1
4 5 3
4 7 1
5 6 1
6 7 3
```

```
1 2 90
2 3 74
3 1 87
6 7 144
1 4 49
1 5 58
2 4 60
2 5 32
3 4 64
3 5 42
6 4 87
6 5 24
7 4 100
7 5 158
```

```
0 1 63
0 2 82
0 3 60
0 4 41
0 5 41
1 2 75
1 3 58
1 4 35
1 5 49
2 3 22
2 4 48
2 5 60
3 4 49
3 5 82
```

```
0 1 30
2 3 30
4 5 30
0 2 10
2 4 10
4 1 10
1 3 10
3 5 10
5 0 10
6 7 20
6 3 15
6 4 5
7 3 5
7 4 15
```

```
0 1 30
2 3 30
4 5 30
0 2 10
2 4 10
4 1 10
1 3 10
3 5 10
5 0 10
6 7 30
6 0 10
6 4 15
7 4 15
7 3 10
```

```
0 1 6
2 3 9
0 2 2
2 4 3
4 1 9
1 3 3
3 5 2
5 0 9
6 7 6
6 5 2
6 4 2
7 4 6
7 3 2
```

```
0 1 1
2 3 2
4 5 8
0 2 8
4 1 5
1 3 7
3 5 4
5 0 2
6 7 9
6 2 1
6 4 1
7 4 8
7 5 3
```

0	1	3
0	3	3
0	5	2
1	2	9
2	3	3
3	4	11
4	5	6
5	6	11
6	1	7
7	0	11
7	6	2
7	2	6
7	4	3

0	1	53
0	2	33
0	5	30
1	2	83
2	3	7
3	4	85
4	5	15
5	6	93
6	1	14
7	0	82
7	6	15
7	3	47
7	4	73

0	1	91
0	2	25
0	3	27
1	2	88
1	3	71
2	4	53
2	7	24
3	5	3
3	6	78
3	7	56
4	7	29
4	5	24
5	6	81
6	7	28

0	1	130
0	2	95
0	3	28
1	2	45
1	3	112
2	4	36
2	5	26
3	4	37
3	6	2
4	6	39
4	7	59
5	6	39
5	7	70
6	7	31

0	1	4
2	3	4
4	5	4
6	7	4
0	3	1
1	2	1
2	4	1
3	5	1
7	5	1
6	4	1
7	1	1
6	0	1

6	7	2
0	1	5
0	2	2
0	4	1
1	2	3
1	3	2
2	3	5
2	5	2
3	5	3
3	6	1
3	7	1
4	5	5
4	6	1
4	7	1

0	1	116
0	2	103
0	3	114
0	4	105
1	2	24
2	3	113
3	4	42
4	1	47
5	1	136
5	2	122
5	3	47
6	2	120
6	3	130
6	4	135