Embedding simply connected 2-complexes in 3-space, and further results on infinite graphs and matroids

Habilitationsschrift

vorgelegt im
Fachbereich Mathematik
Universität Hamburg

von
Johannes Carmesin

Cambridge
2017, version from 2019

To Sarah

Contents

Introduction

It is quite well understood which graphs can be embedded in the plane. For example, Kuratowski's theorem from 1930 says that a graph can be embedded in the plane if and only if it does not contain a graph from Figure 1 as a minor[1]. In the first part of this thesis we prove a 3-dimensional analogue of Kuratowski's

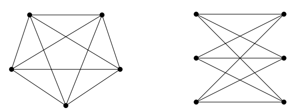

Figure 1: The graphs K_5 (on the left) and $K_{3,3}$ (on the right).

theorem. This answers questions of Lovász, Pardon and Wagner. The author regards this as the main result of this thesis and refers the reader to Chapter 1 for a detailed introduction.

In addition to that first part, this thesis has two more parts. These consist of seven chapters that each are self contained.

End boundaries of infinite graphs have proven to be an important tool in Infinite Graph Theory. Here we prove a conjecture of Andreae from 1981 that implies that end-degrees of infinite directed graphs exist.

For undirected 1-ended graphs, we construct tree-decompositions that display the end and respect the symmetries of the graph. This can be applied to prove a conjecture of Halin from 2000 and solves a recent problem of Boutin and Imrich.

Furthermore, we characterise the classes of infinite graphs with bounded colouring number in terms of forbidden obstructions.

In a nutshell, matroids are common generalisations of graphs and vector spaces. The connection between matchings in bipartite graphs, Menger's theorem about vertex-disjoint paths and base packing and covering is transparent from Edmonds' theorem about the intersection of matroids. In 1990 Nash-Williams proposed a possible extension of Edmonds' theorem to infinite ma-

[1] In the context of planar graphs, the minor relation is just the subgraph relation combined with planar duality.

troids. We show that like Edmonds' theorem, this conjecture is equivalent to other natural problems such as base packing or matroid union. Furthermore it implies the Erdős-Menger-Conjecture, which had been open for 50 years until it was proved by Aharoni and Berger in 2009. Our new perspectives allow us to prove Nash-Williams' conjecture in various special cases.

A lot of matroid theory focuses on matroids representable over vector spaces. We develop a compactness method that allows us to lift many of the foundational theorems about such representable matroids to the infinite setting. A related construction answers a question of Bruhn, Diestel, Kriesell, Pendavingh and Wollan.

This thesis is based on the twelve papers [25], [26], [27], [28], [29], [19], [73],[60], [16],[6],[14], [15]. Five of these are already published or are accepted in journals: one in Combinatorica, one in Discrete Mathematics and three in Journal of Combinatorial Theory.

The papers of the first part are single authored. The other parts are based on joint work with various subsets of Elad Aigner-Horev, Nathan Bowler, Jan-Oliver Fröhlich, Julian Pott, Péter Komjáth, Florian Lehner, Rögnvaldur Möller and Christian Reiher. I am grateful to all these coauthors for fruitful cooperation, in particular to Nathan Bowler.

Part I

Embedding simply connected 2-complexes in 3-space

Abstract

We characterise the embeddability of simply connected locally 3-connected 2-dimensional simplicial complexes in 3-space in a way analogous to Kuratowski's characterisation of graph planarity, by excluded minors. This answers questions of Lovász, Pardon and Wagner.

Introduction

In 1930, Kuratowski proved that a graph can be embedded in the plane if and only if it has none of the two non-planar graphs K_5 or $K_{3,3}$ as a minor[2]. The main result of this chapter may be regarded as a 3-dimensional analogue of this theorem.

Kuratowski's theorem gives a way how embeddings in the plane could be understood through the minor relation. A far reaching extension of Kuratowski's theorem is the Robertson-Seymour theorem [83]. Any minor closed class of graphs is characterised by the list of minor-minimal graphs not in the class. This theorem says that this list always must be finite. The methods developed to prove this theorem are nowadays used in many results in the area of structural graph theory [36] – and beyond; recently Geelen, Gerards and Whittle extended the Robertson-Seymour theorem to representable matroids by proving Rota's conjecture [43]. Very roughly, the Robertson-Seymour structure theorem establishes a correspondence between minor closed classes of graphs and classes of graphs almost embeddable in 2-dimensional surfaces.

In his survey on the Graph Minor project of Robertson and Seymour [67], in 2006 Lovász asked whether there is a meaningful analogue of the minor relation in three dimensions. Clearly, every graph can be embedded in 3-space[3].

One approach towards this question is to restrict the embeddings in question, and just consider so called linkless embeddings of graphs, see [82] for a survey. Instead of restricting embeddings, one could also put some additional structure on the graphs in question. Indeed, Wagner asked how an analogue of the minor relation could be defined on general simplicial complexes [96].

Unlike in higher dimensions, a 2-dimensional simplicial complex has a topological embedding in 3-space if and only if it has a piece-wise linear embedding if and only if it has a differential embedding [11, 55, 72, 77]. In [68], Matoušek, Sedgewick, Tancer and Wagner proved that the embedding problem for 2-dimensional simplicial complexes in 3-space is decidable. In August 2017, de Mesmay, Rieck, Sedgwick and Tancer complemented this result by showing that this problem is NP-hard [33].

This might suggest that if we would like to get a structural characterisation of embeddability, we should work inside a subclass of 2-dimensional simplicial complexes. And in fact such questions have been asked: in 2011 at the internet

[2] A *minor* of a graph is obtained by deleting or contracting edges.

[3] Indeed, embed the vertices in general position and embed the edges as straight lines.

forum ‘MathsOverflow’ Pardon[4] asked whether there are necessary and sufficient conditions for when contractible 2-dimensional simplicial complexes embed in 3-space. The *link graph* at a vertex v of a simplicial complex is the incidence graph between edges and faces incident with v. He notes that if embeddable the link graph at any vertex must be planar. This leads to obstructions for embeddability such as the cone over the complete graph K_5, see Figure 2. – But there are different obstructions of a more global character, see Figure 3. All their link graphs are planar – yet they are not embeddable.

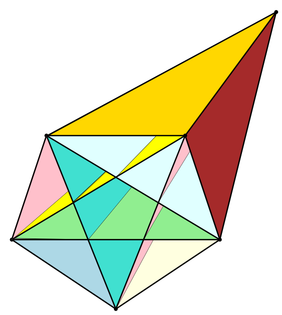

Figure 2: The cone over K_5. Similarly as the graph K_5 does not embed in 2-space, the cone over K_5 does not embed in 3-space.

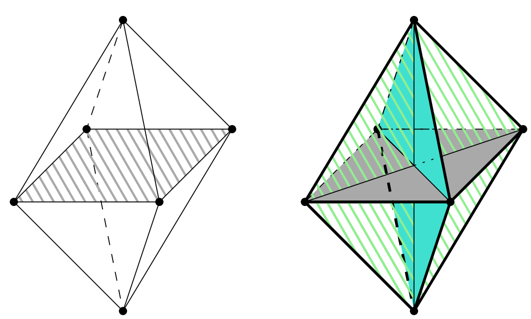

Figure 3: The octahedron obstruction, depicted on the right, is obtained from the octahedron with its eight triangular faces by adding 3 more faces of size 4 orthogonal to the three axis. If we add just one of these 4-faces to the octahedron, the resulting 2-complex is embeddable as illustrated on the left. A second 4-face could be added on the outside of that depicted embedding. However, it can be shown that the octahedron with all three 4-faces is not embeddable.

Addressing these questions, we introduce an analogue of the minor relation

[4]John Pardon confirmed in private communication that he asked that question as the user ‘John Pardon’.

	delete	contract
edge		
face		

Figure 4: For each of the four corners of the above diagram we have one space minor operation.

for 2-complexes and we use it to prove a 3-dimensional analogue of Kuratowski's theorem characterising when simply connected 2-dimensional simplicial complexes (topologically) embed in 3-space.

More precisely, a *space minor* of a 2-complex is obtained by successively deleting or contracting edges or faces, and splitting vertices. See Figure 4 and Figure 5. The precise details of these definitions are given in Section 1.5; for example contraction of edges is only allowed for edges that are not loops[5] and we only contract faces of size at most two.

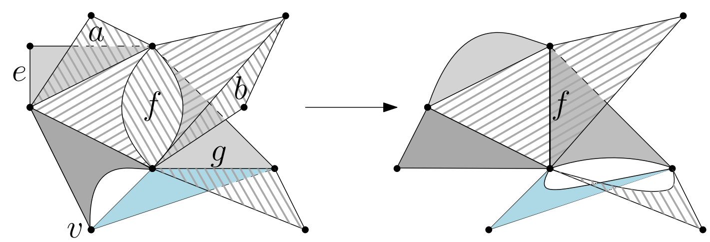

Figure 5: The complex on the right is a space minor of the complex on the left. Indeed, for that just delete the faces labelled a and b, contract the edge e and contract the face f, and delete the edge g and split the vertex v.

It will be quite easy to see that space minors preserve embeddability in 3-space and that this relation is well-founded. The operations of face deletion and face contraction correspond to the minor operations in the dual matroids of simplicial complexes in the sense of Chapter 4.

The main result of this chapter is the following.

Theorem 0.0.1. *Let C be a simply connected locally 3-connected 2-dimensional simplicial complex. The following are equivalent.*

- *C embeds in 3-space;*
- *C has no space minor from the finite list $\mathcal{Z}$.*

The finite list $\mathcal{Z}$ is defined explicitly in Subsection 1.5.3 below. The members of $\mathcal{Z}$ are grouped in six natural classes. Here a (2-dimensional) simplicial

[5] *Loops* are edges that have only a single endvertex. While contraction of edges that are not loops clearly preserves embeddability in 3-space, for loops this is not always the case.

complex is *locally 3-connected* if all its link graphs are connected and do not contain separators of size one or two. In Chapter 5, we extend Theorem 0.0.1 to simplicial complexes that need not be locally 3-connected. For general simplicial complexes, not necessarily simply connected ones, the proof implies that a locally 3-connected simplicial complex has an embedding into some 3-manifold if and only if it does not have a minor from $\mathcal{L}$.

We are able to extend Theorem 0.0.1 from simply connected simplicial complexes to those whose first homology group is trivial.

Theorem 0.0.2. *Let C be a locally 3-connected 2-dimensional simplicial complex such that the first homology group $H_1(C, \mathbb{F}_p)$ is trivial for some prime p. The following are equivalent.*

- *C embeds in 3-space;*
- *C is simply connected and has no space minor from the finite list $\mathcal{Z}$.*

In general there are infinitely many obstructions to embeddability in 3-space. Indeed, the following infinite family of obstructions appears in Theorem 0.0.2.

Example 0.0.3. Given a natural number $q \geq 2$, the q-folded cross cap consists of a single vertex, a single edge that is a loop and a single face traversing the edge q-times in the same direction. It can be shown that q-folded cross caps cannot be embedded in 3-space.

A more sophisticated infinite family is constructed in Chapter 4.

This chapter is subdivided into five chapters, which are self-contained except in those few cases, where we point it out explicitly. In what follows we summarise roughly the content of the other four chapters. The results of Chapter 2 give combinatorial characterisations when simplicial complexes embed in 3-space, which are used in the proofs of Theorem 0.0.1 and Theorem 0.0.2.

As mentioned above, the main result of Chapter 5 is an extension of Theorem 0.0.1 to simply connected simplicial complexes. This relies on Chapter 1 and Chapter 2.

Chapter 3 is purely graph-theoretic and its results are used as a tool in Chapter 4.

In Chapter 4, we prove an extension of the main theorem of Chapter 5 that goes beyond the simply connected case. And we additionally prove the following. Like Kuratowski's theorem, Whitney's theorem is a characterisation of planarity of graphs. In Chapter 4 we prove a 3-dimensional analogue of that theorem.

In Chapter 4, we prove an extension of the main theorem of Chapter 5 that goes beyond the simply connected case. And we additionally prove the following. Like Kuratowski's theorem, Whitney's theorem is a characterisation of planarity of graphs. In Chapter 4 we prove a 3-dimensional analogue of that theorem.

Contents

Chapter 1

A Kuratowski-type characterisation

This chapter is organised as follows. Most of this chapter is concerned with the proof of Theorem 0.0.2, which implies Theorem 0.0.1. In Section 1.1, we introduce 'planar rotation systems' and state a theorem of Chapter 2 that relates embeddability of simply connected simplicial complexes to existence of planar rotation systems. In Section 1.2 we define the operation of 'vertex sums' and use it to study rotation systems. In Section 1.3 we relate the existence of planar rotation systems to a property called 'local planarity'. In Section 1.4 we characterise local planarity in terms of finitely many obstructions. In Section 1.5 we introduce space minors and prove Theorem 0.0.1 and Theorem 0.0.2.

For graphs[1] we follow the notation of [36]. Beyond that a *2-complex* is a graph (V, E) together with a set F of closed trails[2], called its *faces*. In this chapter we follow the convention that each vertex or edge of a simplicial complex or a 2-complex is incident with a face. The definition of *link graphs* naturally extends from simplicial complexes to 2-complexes with the following addition: we add two vertices in the link graph $L(v)$ for each loop incident with v. We add one edge to $L(v)$ for each traversal of a face at v.

1.1 Rotation systems

Rotation systems of 2-complexes play a central role in our proof of Theorem 0.0.1. In this section we introduce them and prove some basic properties of them.

[1] In this chapter graphs are allowed to have loops and parallel edges.

[2] A *trail* is sequence $(e_i | i \leq n)$ of distinct edges such that the endvertex of e_i is the starting vertex of e_{i+1} for all $i < n$. A trail is *closed* if the starting vertex of e_1 is equal to the endvertex of e_n.

A rotation system of a graph G is a family $(\sigma_v|v \in V(G))$ of cyclic orientations[3] σ_v of the edges incident with the vertices v [71]. The orientations σ_v are called *rotators*. Any rotation system of a graph G induces an embedding of G in an oriented (2-dimensional) surface S. To be precise, we obtain S from G by gluing faces onto (the geometric realisation of) G along closed walks of G as follows. Each directed edge of G is in one of these walks. Here the direction $\vec{a}$ is directly before the direction $\vec{b}$ in a face f if the endvertex v of $\vec{a}$ is equal to the starting vertex of $\vec{b}$ and b is just after a in the rotator at v. The rotation system is *planar* if that surface S is a disjoint union of 2-spheres. Note that if the graph G is connected, then for any rotation system of G, also the surface S is connected.

A *rotation system of a (directed[4]) 2-complex C* is a family $(\sigma_e|e \in E(C))$ of cyclic orientations σ_e of the faces incident with the edge e. A rotation system of a 2-complex C *induces* a rotation system at each of its link graphs $L(v)$ by restricting to the edges that are vertices of the link graph $L(v)$; here we take $\sigma(e)$ if e is directed towards v and the reverse of $\sigma(e)$ otherwise.

A rotation system of a 2-complex is *planar* if all induced rotation systems of link graphs are planar. In Chapter 2 we prove the following, which we use in the proof of Theorem 0.0.1.

Theorem 1.1.1. *[Theorem 2.2.1] A simply connected simplicial complex has an embedding in $\mathbb{S}^3$ if and only if it has a planar rotation system.*

Given a 2-complex C, its link graph $L(v)$ is *loop-planar* if it has a planar rotation system such that for every loop ℓ incident with v the rotators at the two vertices e_1 and e_2 associated to ℓ are reverse – when we apply the following bijection between the edges incident with e_1 and e_2. If f is an edge incident with the vertex e_1 whose face of C consists only of the loop ℓ, then f is an edge between e_1 and e_2 and the bijection is identical at that edge. If the face f is incident with more edges than ℓ, it can by assumption traverse ℓ only once. So there are precisely two edges for that traversal, one incident with e_1, the other with e_2. These two edges are in bijection.

A 2-complex C is *locally planar* if all its link graphs are loop-planar. Clearly, a 2-complex that has a planar rotation system is locally planar. However, the converse is not true.

Let $C = (V, E, F)$ be a 2-complex and let x be a non-loop edge of C, the 2-complex obtained from C by *contracting* x (denoted by C/x) is obtained from C by identifying the two endvertices of x, deleting x from all faces and then deleting x, formally: $C/x = ((V, E)/x, \{f - x|f \in F\})$.

Let C be a 2-complex and x be a non-loop edge of C, and $\Sigma = (\sigma_e|e \in E(C))$ be a rotation system of C. The *induced* rotation system of C/x is $\Sigma_x = (\sigma_e|e \in E(C) - x)$. This is well-defined as the incidence relation between edges of C/x

[3] A *cyclic orientation* is a bijection to an oriented cycle.

[4] A *directed* 2-complex is a 2-complex together with a choice of direction at each of its edges and a choice of orientation at each of its faces. All 2-complexes considered in this chapter are directed. In order to simplify notation we will not always say that explicitly.

and faces is the same as in C. Planarity of rotation systems is preserved under contractions:

Lemma 1.1.2. *If Σ is planar, then Σ_x is planar.*

Conversely, for any planar rotation system Σ' of C/x, if the non-loop edge x is not a cutvertex of any of the two link graphs at its endvertices, there is a planar rotation system of C inducing Σ'.[5]

Hence the class of 2-complexes that have planar rotation systems is closed under contractions. As noted above it contains the class of locally planar 2-complexes, which is clearly not closed under contractions. However, if we close the later class under contractions, then they do agree – in the locally 3-connected case as follows.

Lemma 1.1.3. *A locally 3-connected 2-complex has a planar rotation system if and only if all contractions are locally planar.* [6]

We remark that by Lemma 1.2.4 below the class of locally 3-connected 2-complexes is closed under contractions.

1.2 Vertex sums

In this short section we prove some elementary facts about an operation we call 'vertex sum' which is used in the proof of Theorem 0.0.1.

Let H_1 and H_2 be two graphs with a common vertex v and a bijection ι between the edges incident with v in H_1 and H_2. The *vertex sum* of H_1 and H_2 over v given ι is the graph obtained from the disjoint union of H_1 and H_2 by deleting v in both H_i and adding an edge between any pair (v_1, v_2) of vertices $v_1 \in V(H_1)$ and $v_2 \in V(H_2)$ such that v_1v and v_2v are mapped to one another by ι, see Figure 1.1.

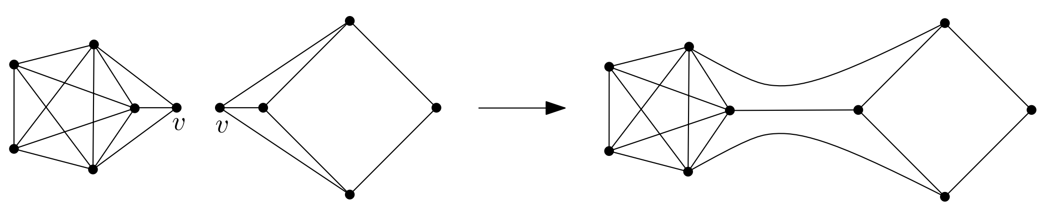

Figure 1.1: The vertex sum of the two graphs on the left is the graph on the right.

Let C be a 2-complex with a non-loop edge e with endvertices v and w.

Observation 1.2.1. *The link graph of C/e at e is the vertex sum of the link graphs $L(v)$ and $L(w)$ over the common vertex e.* $\square$

[5] This lemma is proved in Section 1.2.

[6] Lemma 1.1.3 will follow from Lemma 1.3.1 below.

Lemma 1.2.2. *Let G be a graph that is a vertex sum of two graphs H_1 and H_2 over the common vertex v. Let $(\sigma_x^i | x \in V(H_i))$ be a planar rotation system of H_i for $i = 1, 2$ such that σ_v^1 is the inverse of σ_v^2. Then $(\sigma_x^i | x \in V(H_i) - v, i = 1, 2)$ is a planar rotation system of G.*

Proof sketch. This is a consequence of the topological fact that the connected sum of two spheres is the sphere. □

Lemma 1.2.3. *Let G be a graph that is a vertex sum of two graphs H_1 and H_2 over the common vertex v. Assume that the vertex v not a cutvertex of H_1 or H_2. Assume that G has a planar rotation system Σ. Then there are planar rotation systems of H_1 and H_2 that agree with Σ at the vertices in $V(G) \cap V(H_i)$ and that are reverse at v.*

Proof. Since the vertex v is not a cutvertex of the graph H_2, the graph H_1 can be obtained from the graph G by contracting the connected vertex set $V(H_2) - v$ onto a single vertex. Now let a plane embedding ι of G be given that is induced by the rotation system Σ. Since contractions can be performed within the plane embedding ι, there is a planar rotation system Σ_1 of the graph H_1 that agrees with Σ at all vertices in $V(H_1) - v$.

Since the vertex v is not a cutvertex of H_1 or H_2, the cut X of G consisting of the edges between $V(H_1) - v$ and $V(H_2) - v$ is actually a bond of the graph G. The bond X is a circuit o of the dual graph of G with respect to the embedding ι. And the rotator at v of the embedding Σ_1 is equal (up to reversing) to the cyclic orientation of the edges on the circuit o. Similarly, we construct a planar rotation system Σ_2 of H_2 that agrees with Σ at all vertices in $V(H_2) - v$, and the rotator at the vertex v is the other orientation of the circuit o. This completes the proof. □

Proof of Lemma 1.1.2. This is a consequence of Lemma 1.2.2 and Lemma 1.2.3. □

Lemma 1.2.4. *Let G be a graph that is a vertex sum of two graphs H_1 and H_2 over the common vertex v. Let $k \geq 2$. If H_1 and H_2 are k-connected*[7]*, then so is G.*

Proof. Suppose for a contradiction that there is a set of less than k vertices of G such that $G \setminus X$ is disconnected. Let Y be the set of edges incident with v (suppressing the bijection between the edges incident with v in H_1 and H_2 in our notation). As H_1 is k-connected, the set Y contains at least k edges. If $k > 2$, then since no H_i has parallel edges, no two edges in Y share a vertex. Thus in this case the set Y contains k edges that are vertex disjoint. If $k = 2$, then either one H_i consists of a single class of parallel edges and the lemma is immediate; or else, there are two disjoint edges of Y – here this is true as Y

[7]Given $k \geq 2$, a graph with at least $k + 1$ vertices is *k-connected* if the removal of less than k vertices does not make it disconnected. Moreover it is not allowed to have loops and if $k > 2$, then it is not allowed to have parallel edges.

considered as a subgraph of G is a bipartite graph with at least two vertices on either side each having degree at least one.

Hence by the pigeonhole principle, there is an edge e in Y such that no endvertex of e is in X. Let C be the component of $G \setminus X$ that contains e. Let C' be a different component of $G \setminus X$. Let i be such that H_i contains a vertex w of C'.

In H_i this vertex w and an endvertex of e are separated by $X + v$. As H_i is k-connected, we deduce that all vertices of X are in H_i. Then the connected graph H_{i+1} is a subset of C. Hence the vertex w and an endvertex of e are separated by X in H_i. This is a contradiction to the assumption that H_i is k-connected. $\square$

In our proof we use the following simple fact.

Lemma 1.2.5. *Let G be a graph with a minor H. Let v and w be vertices of G contracted to the same vertex of H. Then there is a minor G' of G such that v and w are contracted to different vertices of G' and their branch vertices are joined by an edge e and $H = G'/e$.* $\square$

1.3 Constructing planar rotation systems

The aim of this section is to prove the following lemma, which is used in the proof of Theorem 0.0.1. This lemma roughly says that a 2-complex has a planar rotation system if and only if certain contractions are locally planar. A *chord* of a cycle o is an edge not in o joining two distinct vertices in o but not parallel to an edge of o. A cycle that has no chord is *chordless*.

Lemma 1.3.1. *Let C be a locally 3-connected 2-complex. Assume that the following 2-complexes are locally planar: C, for every non-loop edge e the contraction C/e, and for every non-loop chordless cycle o of C and some $e \in o$ the contraction $C/(o - e)$.*

Then C has a planar rotation system.

First we show the following.

Lemma 1.3.2. *Let C be a 2-complex with an edge e with endvertices v and w. Assume that the link graphs $L(v)$ and $L(w)$ at v and w are 3-connected and that the link graph $L(e)$ of C/e at e is planar. Then for any two planar rotation systems of $L(v)$ and $L(w)$ the rotators at e are reverse of one another or agree.*

Proof. Let $\Sigma = (\sigma_x | x \in (L(v) \cup L(w)) - e)$ be a planar rotation system of $L(e)$. By Lemma 1.2.3 there is a rotator τ_e at e such that $(\sigma_x | x \in L(v) - e)$ together with τ_e is a planar rotation system of $L(v)$ and $(\sigma_x | x \in L(w) - e)$ together with the inverse of τ_e is a planar rotation system of $L(w)$.

Since $L(v)$ and $L(w)$ are 3-connected, their planar rotation system are unique up to reversing and hence the lemma follows. $\square$

Let C be a locally 3-connected 2-complex such that C and for every non-loop e all contractions C/e are locally planar. We pick a planar rotation system $(\sigma_e^v | e \in V(L(v)))$ at each link graph $L(v)$ of C. By Lemma 1.3.2, for every edge e of C with endvertices v and w the rotators σ_e^v and σ_e^w are reverse or agree. We colour the edge e green if they are reverse and we colour it red otherwise.

A *pre-rotation system* is such a choice of rotation systems such that all edges are coloured green. The following is an immediate consequence of the definitions.

Lemma 1.3.3. *C has a pre-rotation system if and only if C has a planar rotation system.* □

Lemma 1.3.4. *Let o be a cycle of C and e an edge on o. Assume that the link graph $L[o, e]$ of $C/(o - e)$ at e is loop-planar. Then the number of red edges of o is even.*

Proof. Since $L[o, e]$ is loop-planar, by Lemma 1.2.3 there are planar rotation systems of all link graphs of vertices of C on o such that for every edge $x \in o$ with endvertices v and w the rotators σ_x^v and σ_x^w are reverse. Hence there are assignments of planar rotation systems to the link graphs at vertices of o such the number of red edges on o is zero.

Since all link graphs are 3-connected, the planar rotation systems are unique up to reversing. Reversing a rotation system flips the colours of all incident edges. Hence for any assignment of planar rotation systems the number of red edges of o must be even. □

Proof of Lemma 1.3.1. By Lemma 1.3.3, it suffices to construct a pre-rotation system, that is, to construct suitable rotation systems at each link graph of C.

We may assume that C is connected. We pick a spanning tree T of C with root r. At the link graph at r we pick an arbitrary planar rotation system. Now we define a rotation system $(\sigma_e^v | e \in V(L(v)))$ at some vertex v assuming that for the unique neighbour w of v nearer to the root in T we have already defined a rotation system $(\sigma_e^w | e \in V(L(w)))$. Let e be the edge between v and w that is in T. By Lemma 1.3.2, there is a planar rotation system $(\sigma_e^v | e \in V(L(v)))$ of the link graph $L(v)$ such that the rotators σ_e^v and σ_e^w are reverse. As C is connected, this defines a planar rotation system at every vertex of C. It remains to show that every edge e of C is green with respect to that assignment. This is true by construction if e is in T.

Lemma 1.3.5. *Every edge e of C that is not in T and is not a loop is green.*

Proof. Let o_e be the fundamental cycle of e with respect to T. We prove by induction on the number of edges of o_e that e is green. The base case is that o_e is chordless. Then by assumption the link graph $L[o, e]$ of $C/(o - e)$ at e is loop-planar. So the number of red edges on o_e is even by Lemma 1.3.4. As shown above all edges of o_e except for possibly e are green. So e must be green.

Thus we may assume that o_e has chords. By shortcutting along chords we obtain a chordless cycle o_e' containing e such that each edge x of o_e' not in o_e is a chord of o_e. Thus each such edge x is not in T and not a loop. Since no chord x

can be parallel to e, the corresponding fundamental cycles o_x have each strictly less edges than o_e. Hence by induction all the edges x are green. Thus all edges of o'_e except for possibly e are green. Similarly as in the base case we can now apply Lemma 1.3.4 to the chordless cycle o'_e to deduce that e is green. □

Sublemma 1.3.6. *Every loop ℓ of C is green.*

Proof. Let v be the vertex incident with ℓ. As the link graph $L(v)$ is 3-connected and loop-planar each of its (two) planar rotation systems must witness that $L(v)$ is loop-planar. Hence the rotation system we picked at $L(v)$ witnesses that $L(v)$ is loop planar. Thus ℓ is green. □

As all edges of C are green with respect to Σ, the family Σ is a pre-rotation system of C. Hence C has a planar rotation system by Lemma 1.3.3. □

1.4 Marked graphs

In this section we prove Lemma 1.4.9 and Lemma 1.4.19 which are used in the proof of Theorem 0.0.1. More precisely, these lemmas characterise when a 2-complex is locally planar in terms of finitely many obstructions.

A *marked graph* is a graph G together with two of its vertices v and w and three pairs $((a_i, b_i)|i = 1, 2, 3)$ of its edges, where the a_i are incident with v and the b_i are incident with w. We stress that we allow $a_i = b_i$.

Given a 2-complex C, a link graph $L(x)$ of C, a loop ℓ of C incident with x and three distinct faces f_1, f_2, f_3 of C traversing ℓ, the marked graph *associated* with (x, ℓ, f_1, f_2, f_3) is the graph $L(x)$ together with the two vertices v and w of $L(x)$ corresponding to ℓ. The traversal of each face f_i of ℓ corresponds to edges a_i and b_i incident with v and w, respectively. As f_i is a closed trail in C, each vertex of $L(x)$ is incident with at most one edge corresponding to f_i. Hence a_i and b_i are defined unambiguously. Note that if f_i consists only of ℓ, then $a_i = b_i$. This completes the definition of the associated marked graph $(G, v, w, ((a_i, b_i)|i = 1, 2, 3))$.

A marked graph $(G, v, w, ((a_i, b_i)|i = 1, 2, 3))$ is *planar* if there is a planar rotation system $(\sigma_x|x \in V(G))$ of G such that σ_v restricted to (a_1, a_2, a_3) is the inverse permutation of σ_w restricted to (b_1, b_2, b_3) – when concatenated with the bijective map $b_i \mapsto a_i$. The next lemma characterises loop-planarity.

Lemma 1.4.1. *A 3-connected link graph $L(x)$ is loop-planar if and only if it is a planar graph and all its associated marked graphs are planar marked graphs.*

Proof. Clearly, if $L(x)$ is loop-planar, then all its link graphs and all their associated marked graphs are planar. Conversely assume that a link graph $L(x)$ and all its associated marked graphs are planar. Then $L(v)$ has a planar rotation system Σ. As $L(x)$ is 3-connected, this rotation system is unique up to reversing. Hence any planar rotation system witnessing that some associated marked graph is planar is equal to Σ or its inverse. By reversing that rotation system if necessary, we may assume that it is equal to Σ. Hence Σ is a planar rotation system that witnesses that $L(x)$ is loop-planar. $\square$

Corollary 1.4.2. *A locally 3-connected 2-complex C is locally planar if and only if all its link graphs and all their associated marked graphs are planar.*

Proof. By definition, a 2-complex is locally planar if all its link graphs are loop-planar. $\square$

A marked graph $(G, v, w, ((a_i, b_i)|i = 1, 2, 3))$ is *3-connected* if G is 3-connected. We abbreviate $A = \{a_1, a_2, a_3\}$ and $B = \{b_1, b_2, b_3\}$.

A *marked minor* of a marked graph $(G, v, w, ((a_i, b_i)|i = 1, 2, 3))$ is obtained by doing a series of the following operations:

- contracting or deleting an edge not in $A \cup B$;
- replacing an edge $a_i \in A \setminus B$ and an edge $b_j \in B \setminus A$ that are in parallel by a single new edge which is in that parallel class. In the reduced graph, this new edge is a_i and b_j.
- the above with 'serial' in place of 'parallel'.
- apply the bijective map $(v, A) \mapsto (w, B)$.

Lemma 1.4.3. *Let $\hat{G} = (G, v, w, ((a_i, b_i)|i = 1, 2, 3))$ be a marked graph such that G is planar. Let $\hat{H}$ be a 3-connected marked minor of $\hat{G}$. Then $\hat{G}$ is planar if and only if $\hat{H}$ is planar.*

Before we can prove this, we need to recall some facts about rotation systems of graphs. Given a graph G with a rotation system $\Sigma = (\sigma_v|v \in V(G))$ and an edge e. The rotation system *induced* by Σ on $G - e$ is $(\sigma_v - e|v \in V(G))$. Here $\sigma_v - e$ is obtained from the cyclic ordering σ_v by deleting the edge e. The rotation system *induced* by Σ on G/e is $(\sigma_v|v \in V(G/e) - e)$ together with σ_e defined as follows. Let v and w be the two endvertices of e. Then σ_e is obtained from the cyclic ordering σ_v by replacing the interval e by the interval $\sigma_w - e$ (in such a way that the predecessor of e in σ_v is followed by the successor of e in σ_w). Summing up, Σ induces a rotation system at every minor of G. Since the class of plane graphs[8] is closed under taking minors, rotation systems induced by planar rotation systems are planar.

Proof of Lemma 1.4.3. Let Σ be a planar rotation system of G. Let Σ' be the rotation system of the graph H of $\hat{H}$ induced by Σ. As mentioned above, Σ' is planar.

Moreover, Σ witnesses that $\hat{G}$ is a planar marked graph if and only if Σ' witnesses that $\hat{H}$ is a planar marked graph. Hence if $\hat{G}$ is planar, so is $\hat{H}$. Now assume that $\hat{H}$ is planar. Since H is 3-connected, it must be that Σ' witnesses that the marked graph $\hat{H}$ is planar. Hence the marked graph $\hat{G}$ is planar. $\square$

[8] A *plane graph* is a graph together with an embedding in the plane.

Our aim is to characterise when 3-connected marked graphs are planar. By Lemma 1.4.3 it suffices to study that question for marked-minor minimal 3-connected marked graphs; we call such marked graphs *3-minimal*.

It is reasonable to expect – and indeed true, see below – that there are only finitely many 3-minimal marked graphs. In the following we shall compute them explicitly.

Let $\hat{G} = (G, v, w, ((a_i, b_i)|i = 1, 2, 3))$ be a marked graph. We denote by V_A the set of endvertices of edges in A different from v. We denote by V_B the set of endvertices of edges in B different from w.

Lemma 1.4.4. *Let $\hat{G} = (G, v, w, ((a_i, b_i)|i = 1, 2, 3))$ be 3-minimal. Unless G is K_4, every edge in $E(G) \setminus (A \cup B)$ has its endvertices either both in V_A or both in V_B.*

Proof. By assumption G is a 3-connected graph with at least five vertices such that any proper marked minor of $\hat{G}$ is not 3-connected. Let e be an edge of G that is not in $A \cup B$. By Bixby's Lemma [76, Lemma 8.7.3] either $G - e$ is 3-connected[9] after suppressing serial edges or G/e is 3-connected after suppressing parallel edges.

Sublemma 1.4.5. *There is no 3-connected graph H obtained from $G - e$ by suppressing serial edges.*

Proof. Suppose for a contradiction that there is such a graph H. As G is 3-connected, every class of serial edges of $G-e$ has size at most two. By minimality of G, there is no marked minor of $\hat{G}$ with graph H. Hence one of these series classes has to contain two edges in A or two edges in B. By symmetry, we may assume that e has an endvertex x that is incident with two edges e_1 and e_2 in A. As G is 3-connected these two adjacent edges of A can only share the vertex v. Thus $x = v$. This is a contradiction to the assumption that e_1 and e_2 are in series as v is incident with the three edges of A. □

By Sublemma 1.4.5 and Bixby's Lemma, we may assume that the graph H obtained from G/e by suppressing parallel edges is 3-connected. By minimality of G, there is no marked minor of $\hat{G}$ with graph H. Hence G/e has a nontrivial parallel class. And it must contain two edges e_1 and e_2 that are both in A or both in B. By symmetry we may assume that e_1 and e_2 are in A. Since G is 3-connected, the edges e, e_1 and e_2 form a triangle in G. The common vertex of e_1 and e_2 is v. Thus both endvertices of e are in V_A. □

A consequence of Lemma 1.4.4 is that every 3-minimal marked graph has at most most 12 edges. However, we can say more:

[9]The notion of '3-connectedness' used in [76, Lemma 8.7.3] is slightly more general than the notion used here. Indeed, the additional 3-connected graphs there are subgraphs of K_3 or subgraphs of $U_{1,3}$ – the graph with two vertices and three edges in parallel. It is straightforward to check that these graphs do not come up here as they cannot be obtained from a 3-connected graph with at least 5 vertices by a single operation of deletion or contraction (and simplification as above).

Corollary 1.4.6. *Let $\hat{G} = (G, v, w, ((a_i, b_i)|i = 1, 2, 3))$ be 3-minimal. Then G has at most five vertices.*

Proof. Let G_A be the induced subgraph with vertex set $V_A + v$. Let G_B be the induced subgraph with vertex set $V_B + w$. Note that $G = G_A \cup G_B$. If G_A and G_B have at least three vertices in common, then G has at most five vertices as G_A and G_B both have at most four vertices. Hence we may assume that G_A and G_B have at most two vertices in common. As G is 3-connected, the set of common vertices cannot be a separator of G. Hence $G_A \subseteq G_B$ or $G_B \subseteq G_A$. Hence G has at most four vertices in this case. □

An *unlabelled marked graph* is a graph G together with vertices v and w and edge sets A and B of size three such that all edges of A are incident with v and all edges in B are incident with w. The *underlying* unlabelled marked graph of a marked graph $(G, v, w, ((a_i, b_i)|i = 1, 2, 3))$ is G together with v, w and the sets $A = \{a_1, a_2, a_3\}$ and $B = \{b_1, b_2, b_3\}$. Informally, an unlabelled marked graph is a marked graph without the bijection between the sets A and B. For a planar 3-connected unlabelled marked graph, there are three bijections between A and B for which the associated marked graph is planar as a marked graph. For the other three bijections it is not planar.

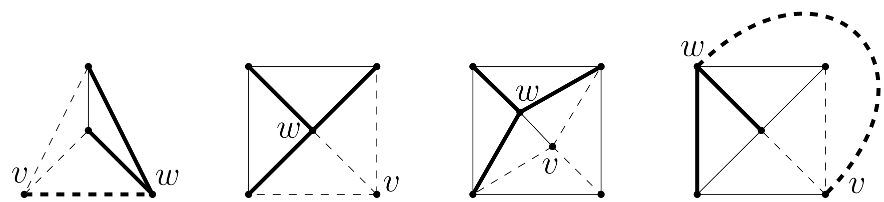

Figure 1.2: The four unlabelled marked graphs in $\mathcal{X}$. The edges in A are depicted dotted, the ones in B are bold.

Marked graphs $\hat{G} = (G, v, w, ((a_i, b_i)|i = 1, 2, 3))$ associated to link graphs always have the property that the vertices v and w are distinct. 3-minimal marked graphs need not have this property. Of particular interest to us is the class $\mathcal{X}$ depicted in Figure 1.2; indeed, they describe the 3-connected marked graphs with the property that $v \neq w$ that are marked minor minimal with G planar, as shown in the following. We shall refer to the four members of $\mathcal{X}$ in the linear ordering given by accessing Figure 1.2 from left to right (and say things like 'the first member of $\mathcal{X}$').

Lemma 1.4.7. *Let $\hat{G} = (G, v, w, ((a_i, b_i)|i = 1, 2, 3))$ be a 3-connected marked graph with $v \neq w$ and G planar. Then $\hat{G}$ has a marked minor that has an underlying unlabelled marked graph in $\mathcal{X}$.*

Proof. By Corollary 1.4.6, $\hat{G}$ has a marked minor minimal 3-connected marked minor $\hat{H} = (H, v, w, ((a_i, b_i)|i = 1, 2, 3))$, where H has at most five vertices.

Sublemma 1.4.8. *The only 3-connected planar graphs with at most five vertices are K_4, the 4-wheel and K_5^-.*

Proof. Since K_4 is the only 3-connected graph with less than five vertices, it suffices to consider the case where the graph K in question has five vertices. As five is an odd number and K has minimum degree 3, K has a vertex v of degree 4. Hence $K - v$ is 2-connected. Hence it has to contain a 4-cycle. Thus K has the 4-wheel as a subgraph. Thus K is the 4-wheel, K_5^- or K_5. As K is planar, it cannot be K_5. $\square$

By Sublemma 1.4.8, H is K_4, the 4-wheel or K_5^-. In the following we treat these cases separately. As above we let $A = \{a_1, a_2, a_3\}$ and $B = \{b_1, b_2, b_3\}$.

Case 1: $H = K_4$. If the vertices v and w of H are distinct, then the underlying unlabelled marked graph of $\hat{H}$ is the first member of $\mathcal{X}$ and the lemma is true in this case. Suppose for a contradiction that $v = w$. Then each edge incident with v is in A and B. Let H' be the marked graph obtained from $\hat{H}$ by replacing each edge incident with v by two edges in parallel, one in A, one in B. It is clear that H' is a marked minor of $\hat{G}$. By applying Lemma 1.2.5 to the graph of H', we deduce that G has K_5 as a minor. This is a contradiction to the assumption that G is planar.

Case 2A: H is the 4-wheel and $v \neq w$.

Subcase 2A1: v or w is the center of the 4-wheel. By applying the bijective map $(v, A) \mapsto (w, B)$ if necessary, we may assume that w is the center. Our aim is to show that the underlying unlabelled marked graph of $\hat{H}$ is the second member of $\mathcal{X}$. As v has degree three, A is as desired. By Lemma 1.4.4, the two edges on the rim not in A must have both their endvertices in V_B. Hence B is as desired. Thus the underlying unlabelled marked graph of $\hat{H}$ is the second member of $\mathcal{X}$.

Subcase 2A2: v and w are adjacent vertices on the rim. We shall show that this case is not possible. Suppose for a contradiction that it is possible.

We denote by e the edge on the rim not incident with v or w. One endvertex has distance two from v, the other has distance two from w. Hence the endvertices of e cannot both be in V_A or both be in V_B. This is a contradiction to Lemma 1.4.4.

Subcase 2A3: v and w are opposite vertices on the rim. We shall show that this case is not possible. Suppose for a contradiction that it is possible.

There is an edge incident with the center not incident with v or w. Deleting that edge and suppressing the vertex of degree two gives a marked graph whose graph is K_4. Hence $\hat{H}$ is not minimal in that case, a contradiction. This completes Case 2A.

Case 2B: H is the 4-wheel and $v = w$. By Lemma 1.4.4, every edge not in $A \cup B$ must have both endvertices in V_A or V_B. Hence v can only be the center of the 4-wheel. By the minimality of $\hat{H}$ and by Lemma 1.4.4, each edge of the rim has both its endvertices in V_A or in V_B. At most two edges of the rim can have all their endvertices in V_A and in that case these edges are adjacent on the rim. The same is true for V_B.

We denote the vertices of the rim by $(v_i|i \in \mathbb{Z}_4)$, where v_iv_{i+1} is an edge. By symmetry, we may assume that v_1 is the unique vertex of the rim not in V_A. Then v_3 must be the unique vertex of the rim not in V_B. It follows that the edges vv_2 and vv_4 are in A and B. Let H' be the marked graph obtained from $\hat{H}$ by replacing each of vv_2 and vv_4 by two edges in parallel, one in A, one in B. It is clear that H' is a marked minor of $\hat{G}$. Let H'' be the marked graph obtained from H' by applying Lemma 1.2.5. The underlying unlabelled marked graph of H'' is the third member of $\mathcal{X}$.

Case 3: H is K_5^-.

We shall show that the underlying unlabelled marked graph of $\hat{H}$ is the forth graph of $\mathcal{X}$. H has three vertices of degree four, which lie one a common 3-cycle. Removing any edge of that 3-cycle gives a graph isomorphic to the 4-wheel. Hence by minimality of $\hat{H}$, it must be that this 3-cycle is a subset of $A \cup B$. In particular, v and w are distinct vertices on that 3-cycle. Up to symmetry, there is only one choice for v and w. By applying the map $(v, A) \mapsto (w, B)$ if necessary, we may assume that A contains at least two edges of that 3-cycle.

We denote the two vertices of H of degree three by u_1 and u_2. We denote the vertex of degree four different from v and w by x. By exchanging the roles of u_1 and u_2 if necessary, we may assume that $A = \{vw, vx, vu_1\}$.

Recall that $wx \in B$. The endvertex u_2 of the edge vu_2 is not in V_A and this edge cannot be in B. Hence by Lemma 1.4.4, both its endvertices must be in V_B. Hence $vw \in B$ and $wu_2 \in B$. Summing up $B = \{wx, vw, wu_2\}$. Thus in this case the underlying unlabelled graph of $\hat{H}$ is the forth graph of $\mathcal{X}$. $\square$

By $\mathcal{Y}$ we denote the class of marked graphs that are not planar as marked graphs and whose underlying unlabelled marked graphs are isomorphic to a member of $\mathcal{X}$ – perhaps after applying the bijective map $(v, A) \mapsto (w, B)$. We consider two marked graphs the same if they have the have the same graph and the same bijection between the sets A and B (although the elements in A might have different labels). Hence for each $X \in \mathcal{X}$, there are precisely three marked graphs in $\mathcal{Y}$ with underlying unlabelled marked graph X, one for each of the three bijections between A and B that are not compatible with any rotation system of the graph of X (which is 3-connected). Thus $\mathcal{Y}$ has twelve elements.

Summing up we have proved the following.

Lemma 1.4.9. *A locally 3-connected 2-complex is locally planar if and only if all its link graphs are planar and all their associated marked graphs do not have a marked minor from $\mathcal{Y}$.*

Proof. Since no marked graph in $\mathcal{Y}$ is planar, it is immediate that if a 2-complex is locally planar, then all its link graphs are planar and all their associated marked graphs do not have a marked minor from $\mathcal{Y}$.

For the other implication it suffices to show that any 3-connected link graph $L(x)$ that is planar but not loop-planar has an associated marked graph that has a marked minor in $\mathcal{Y}$. By Lemma 1.4.1, $L(x)$ has an associated marked graph $\hat{G}$ that is not planar. By Lemma 1.4.7, $\hat{G}$ has a marked minor $\hat{H}$ whose underlying unlabelled marked graph is in $\mathcal{X}$. By Lemma 1.4.3, $\hat{H}$ is not planar. Hence $\hat{H}$ is in $\mathcal{Y}$. $\square$

Lemma 1.4.9 has already the following consequence, which characterises embeddability in 3-space by finitely many obstructions.[10]

Corollary 1.4.10. *Let C be a simply connected locally 3-connected 2-complex. Let C' be a contraction of C to a single vertex v. Then C has an embedding into $\mathbb{S}^3$ if and only if no marked graph associated to the link graph at v has a marked minor in the finite set $\mathcal{Y}$.*

Proof. By Theorem 1.1.1, C is embeddable if and only if it has a planar rotation system. By Lemma 1.1.3 C has a planar rotation system if and only if C' is locally planar. Hence Corollary 1.4.10 follows from Lemma 1.4.9. $\square$

In the following we will deduce from Lemma 1.4.9 a more technical analogue. A *strict marked graph* is a marked graph $(G, v, w, ((a_i, b_i)|i = 1, 2, 3))$ together with a bijective map between the edges incident with v and the edges incident with w that maps a_i to b_i. A *strict marked minor* is obtained by deleting edges not incident with v or w or deleting an edge not in $A \cup B$ incident with v and the edge it is bijected to, and contracting edges if they have an endvertex x of degree two such that x is neither equal to v or w nor x is adjacent to v or w. We also allow to apply the bijective map $(v, A) \mapsto (w, B)$.

Remark 1.4.11. We call this relation the 'strict marked minor relation' as it is more restrictive than the 'marked minor relation'.

The proof of the next lemma is technical. We invite the reader to skip it when first reading the paper.

Lemma 1.4.12. *There is a finite set $\mathcal{Y}'$ of strict marked graphs such that a strict marked graph has a strict marked minor in $\mathcal{Y}'$ if and only if its marked graph has a marked minor in $\mathcal{Y}$.*

Proof. The *underlyer* of a strict marked graph $\hat{Y}$ is the the underlying unlabelled marked graph of the strict marked graph $\hat{Y}$. We define $\mathcal{Y}'$ and reveal the precise definition in steps during the proof. Now we reveal that by $\mathcal{Y}'$ we denote the class of strict marked graphs with underlyer in $\mathcal{X}_5$ – perhaps after applying the bijective map $(v, A) \mapsto (w, B)$. The set $\mathcal{X}_5$, however, is revealed later. We abbreviate 'strict marked minor' by *5-minor*. We define *0-minors* like 'marked minors' but on the larger class of strict marked graphs where we additionally allow that edges incident with v or w have no image under ι. (This is necessary for this class to be closed under 0-minors). Let $\mathcal{X}_0 = \mathcal{X}$.

Let $\hat{Y}$ be a strict marked graph. In this language, it suffices to show that $\hat{Y}$ has a 0-minor with underlyer in $\mathcal{X}_0$ if and only if $\hat{Y}$ has a 5-minor with underlyer

[10] As turns out, Corollary 1.4.10 is too weak to be used directly in our proof of Theorem 0.0.1. Indeed, in our proof it will not always be possible to contract C onto a single vertex but we need to choose the edges we contract carefully (using the additional information provided in Lemma 1.3.1).

in $\mathcal{X}_5$. We will show this in five steps. In the n-th step we define *n-minors* and a set $\mathcal{X}_n$ of unlabelled marked graphs and prove that $\hat{Y}$ has an $(n-1)$-minor with underlyer in $\mathcal{X}_{n-1}$ if and only if $\hat{Y}$ has an n-minor with underlyer in $\mathcal{X}_n$.

Starting with the first step, we define *1-minors* like '0-minors' where we do not allow to contract edges incident with v or w. We define $\mathcal{X}_1$ and reveal it during the proof of the following fact.

Sublemma 1.4.13. *$\hat{Y}$ has a 0-minor with underlyer in $\mathcal{X}_0$ if and only if $\hat{Y}$ has a 1-minor with underlyer in $\mathcal{X}_1$.*

Proof. Assume that $\hat{Y}$ has a 0-minor $\hat{Y}_0$ with underlyer in $\mathcal{X}_0$. So there is a 1-minor $\hat{Y}_1$ of $\hat{Y}$ so that we obtain $\hat{Y}_0$ from $\hat{Y}_1$ by contracting edges incident with v or w. We reveal that $\mathcal{X}_1$ is a superset of $\mathcal{X}_0$. Hence we may assume that there is an edge of $\hat{Y}_1$ that is not in $\hat{Y}_0$. By symmetry, we may assume that it is incident with v. We denote that edge by e_v, see Figure 1.3.

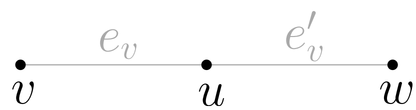

Figure 1.3: The situation of the proof of Sublemma 1.4.13.

We may assume that $\hat{Y}_1$ is minimal, that is, it has no proper 1-minor that has a 0-minor isomorphic to $\hat{Y}_0$. Applying this to $\hat{Y}_1 - e_v$, yields that there must be an edge e'_v incident with v in $\hat{Y}_0$ that in $\hat{Y}_1$ is not incident with v but the other endvertex of e_v. In particular, the edge e'_v is not in A. Let u be the common vertex of e_v and e'_v.

Next we show that u is only incident with e_v and e'_v in Y_1. By going through the four unlabelled marked graphs in $\mathcal{X}_0 = \mathcal{X}$, we check that there is at most one edge incident with v but not in A. Hence u can only be incident with edges not in $\hat{Y}_0 - e'_v$. Moreover the connected component of $Y_1 \setminus Y_0$ containing u can only contain v and vertices not incident with any edge of Y_0. Thus by the minimality of $\hat{Y}_1$, this connected component only contains the edge e_v. So u is only incident with e_v and e'_v.

Since u has degree 2, $\hat{Y}_1/e'_v$ has a 0-minor isomorphic to $\hat{Y}_0$. By the minimality of $\hat{Y}_1$, it must be that $\hat{Y}_1/e'_v$ is not 1-minor of it. Hence e'_v has to be incident with w.

Suppose for a contradiction that there is an edge e_v and an edge e_w defined as e_v with 'w' in place of 'v'. Then as each member of $\mathcal{X}$ has at most one edge between v and w, it must be that $e'_v = e'_w$. This is a contradiction as e'_v is incident with w but not with v in $\hat{Y}_1$ and for e'_w it is the other way round.

Summing up, we have shown that $\hat{Y}_1$ is either equal to $\hat{Y}_0$ or otherwise $\hat{Y}_0$ has an edge e between v and w and $\hat{Y}_1$ is obtained by subdividing that edge. This edge e cannot be in $A \cap B$.

Now we reveal that we define $\mathcal{X}_1$ from $\mathcal{X}$ by adding two more unlabelled marked graphs as follows, see Figure 1.4. The first we get from the second

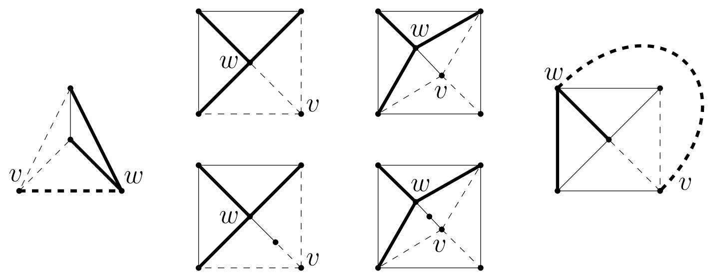

Figure 1.4: The six unlabelled marked graphs in $\mathcal{X}_1$. The edges in A are depicted dotted, the ones in B are bold.

member by subdividing the edge between v and w and let the subdivision edge incident with v remain in A. The second we get from the third member by subdividing the edge between v and w.

From this construction it follows that if $\hat{Y}$ has a 0-minor $\hat{Y}_0$ with underlyer in $\mathcal{X}_0$, then the 1-minor $\hat{Y}_1$ of $\hat{Y}$ defined above has an underlyer in $\mathcal{X}_1$. Hence $\hat{Y}$ has a 1-minor with underlyer in $\mathcal{X}_1$ if and only if it has a 0-minor with underlyer in $\mathcal{X}_0$. $\square$

Starting with the second step, we define *2-minors* like '1-minors' where we only allow to delete edges incident with v and w in the pairs given by the bijection ι – and if they are not in $A \cup B$. We obtain $\mathcal{X}_2$ from $\mathcal{X}_1$ by adding the following unlabelled marked graphs. For each member of $\mathcal{X}_1$ such that all edges incident with v or w are in $A \cup B$ we add no new member. There is one member in $X \in \mathcal{X}_1$ that has an edge incident with w not in $A \cup B$ but every edge incident with v is in A. We add new members obtained from X by adding one more edge incident to v and one other vertex; this may be a vertex of $X - v$ or a new vertex. All other members of $X' \in \mathcal{X}_1$ have the property that they have exactly one edge incident with v not in $A \cup B$ and exactly one edge incident with w not in $A \cup B$. We add new members to $\mathcal{X}_2$ obtained from such an X' by adding two more non-loop edges, one incident with v, the other incident with w.[11] This completes the definition of $\mathcal{X}_2$.

Sublemma 1.4.14. *$\hat{Y}$ has a 1-minor with underlyer in $\mathcal{X}_1$ if and only if $\hat{Y}$ has a 2-minor with underlyer in $\mathcal{X}_2$.*

Proof. By construction, if $\hat{Y}$ has a 2-minor with underlyer in $\mathcal{X}_2$, then it has a 1-minor with underlyer in $\mathcal{X}_1$. Now conversely assume that $\hat{Y}$ has a 1-minor $\hat{Y}_1$ with underlyer in $\mathcal{X}_1$. We define $\hat{Y}_2$ like '$\hat{Y}_1$' except that we only delete edges incident with v or w if also their image under ι is deleted. It remains to show

[11]There are some technical conditions we could further force these newly added edges to satisfy. For example, there are ways in which we could add two edges to the forth member of $\mathcal{X}$ such that the resulting unlabelled marked graph has another member of $\mathcal{X}$ as a strict marked minor. This would give rise to a slightly stronger version of Lemma 1.4.12 and thus of Theorem 0.0.1. To simplify the presentation we do not do it here.

that the underlyer of $\hat{Y}_2$ is in $\mathcal{X}_2$, that is, the graph Y_2 has no loops. This is true as the graph Y_1 has no loops and the additional edges of Y_2 are incident with v or w. So they cannot be loops as no edge of $\hat{Y}$ incident with v or w is contracted by the definition of 1-minor. $\square$

Starting with the third step, we define *3-minors* like '2-minors' where we do not allow to replace parallel or serial pairs of edges in $A \cup B$ as in the second and third operation of marked minor. Each member of $\mathcal{X}_2$ has at most one edge in $A \cap B$. We obtain $\mathcal{X}_3$ from $\mathcal{X}_2$ by adding two new member for each $X \in \mathcal{X}_2$ that has an edge e in $A \cap B$. The first one we obtain by replacing the edge e by two edges in parallel, one in $A \setminus B$ and the other in $B \setminus A$. The second member we construct the same with 'parallel' replaced by 'serial'. The following is immediate.

Sublemma 1.4.15. *$\hat{Y}$ has a 2-minor with underlyer in $\mathcal{X}_2$ if and only if $\hat{Y}$ has a 3-minor with underlyer in $\mathcal{X}_3$.* $\square$

We define *4-minors* like '3-minors' where we only allow to contract edges e if they have an endvertex x of degree two (and as before e is not incident with v or w).

Definition 1.4.16. We say that a graph H is obtained from a graph G by *coadding* the edge e of H at the vertex z of G if $H/e = G$, and the edge e is contracted onto the vertex z of G, and e is not a loop in H.

We obtain $\mathcal{X}_3$ from $\mathcal{X}_4$ by adding all marked graphs obtained from marked graphs in $\mathcal{X}_3$ by coadding edges e at vertices different from v and w such that both endvertices of e have degree at least three. We remark that $\mathcal{X}_4$ is finite as any coadding of such an edge strictly reduces the degree-sequence of the graph in the lexicographical order.

Sublemma 1.4.17. *$\hat{Y}$ has a 3-minor with underlyer in $\mathcal{X}_3$ if and only if $\hat{Y}$ has a 4-minor with underlyer in $\mathcal{X}_4$.*

Proof. Clearly every marked graph in $\mathcal{X}_4$ has a 3-minor in $\mathcal{X}_3$. Now assume that $\hat{Y}$ has a 3-minor $\hat{H}$ with underlyer in $\mathcal{X}_3$. We do the minors as before but only contract edges contracted before if they have an endvertex of degree two; and if they have an endvertex of degree one or are loops, we delete them instead. The resulting strict marked graph $\hat{G}$ has $\hat{H}$ as a 3-minor; namely we just need to contract the edges in $E(G) \setminus E(H)$. However, both endvertices of these edges have degree at least three and they are not loops; that is, G can be obtained from H by coadding edges. Thus $\hat{G}$ is in $\mathcal{X}_4$. So $\hat{Y}$ has a 4-minor in $\mathcal{X}_4$. $\square$

We define *5-minors* like '4-minors', where we additionally require that the endvertex x of degree two is not adjacent to v or w. We let $\mathcal{X}_5$ to consists of those marked graphs obtained from a marked graph of $\mathcal{X}_4$ by subdividing each edge incident with v or w at most once.

Sublemma 1.4.18. *$\hat{Y}$ has a 4-minor with underlyer in $\mathcal{X}_4$ if and only if $\hat{Y}$ has a 5-minor with underlyer in $\mathcal{X}_5$.*

Proof. Clearly every marked graph in $\mathcal{X}_5$ has a 4-minor in $\mathcal{X}_4$. 5-minors are slightly more restricted than 4-minors in that there are a few edges we are not allowed to contract. These edges have an endvertex of degree two that is adjacent to v or w. Hence if $\hat{Y}$ has a 4-minor with underlyer in $\mathcal{X}_4$, and we do the minors as before but do not contract the edges forbidden for 5-minors, we get a strict marked graph with underlyer in $\mathcal{X}_5$, which then is a 5-minor of $\hat{Y}$. □

It is clear from the definitions that 5-minors are just strict marked minors. By Sublemma 1.4.13, Sublemma 1.4.14, Sublemma 1.4.15, Sublemma 1.4.17 and Sublemma 1.4.18, any strict marked graph has a strict marked minor with underlyer in $\mathcal{X}_5$ if and only if its marked graph has a marked minor with underlyer in $\mathcal{X}_0$. This completes the proof. □

The set $\mathcal{Y}'$ is defined explicitly in the proof of Lemma 1.4.12. We fix the set $\mathcal{Y}'$ as defined in that proof. The following is analogue to Lemma 1.4.9 for strict marked minors.

Lemma 1.4.19. *A locally 3-connected 2-complex is locally planar if and only if all its link graphs are planar and all their associated strict marked graphs do not have a strict marked minor from $\mathcal{Y}'$.*

Proof. This is a direct consequence of Lemma 1.4.9 and Lemma 1.4.12. □

1.5 Space minors

In this sections we introduce 'space minors' and prove Theorem 0.0.1 and Theorem 0.0.2.

1.5.1 Motivation

Our approach towards Lovász question mentioned in the Introduction is based on the following two lines of thought.

The first line is as follows. Suppose that a 2-complex C can be embedded in $\mathbb{S}^3$ then we can define a dual graph G of the embedding as follows. Its vertices are the components of $\mathbb{S}^3 \setminus C$ and its edges are the faces of C; each edge is incident with the two components of $\mathbb{S}^3 \setminus C$ touched by its face. It would be nice if the minor operations on the dual graph would correspond to minor operations on C.

The operation of contraction of edges of G corresponds to deletion of faces. But which operation corresponds to deletion of edges of G? If the face of C corresponding to the edge of G is incident with at most two edges of C, then this is the operation of contraction of faces (that is, identify the two incident edges along the face). For faces of size three, however, it is less clear how such an operation could be defined.

The second line of thought is that we would like to define the minor operation such that we can prove an analogue of Kuratowski's theorem – at least in the simply connected case.

Corollary 1.4.10 above is already a characterisation of embeddability in 3-space by finitely many obstructions. However, the reduction operations are not directly operations on 2-complexes (some are just defined on their link complexes). But does Corollary 1.4.10 imply such a Kuratowski theorem? Thus our aim is to define three operations on 2-complexes that correspond to

1. contraction of edges that are not loops[12];
2. deletion of edges in link graphs;
3. contraction of edges in link graphs.

So we make our first operation to be just the first one: contraction of edges that are not loops. A natural choice for the second operation is deletion of faces. This very often corresponds to deletion of edges in the link graph. In some cases however it may happen that a face corresponds to more than one edge in a link graph. This is a technicality we will consider later. Also note that contraction of edges and deletion of faces are "dual"; that is, given a 2-complex C embedded in 3-space and the dual complex D (this is the dual graph G defined above with a face attached for every edge e of C to the edges of G incident with e), contracting an edge in C results in deleting a face in D, and vice versa. This is analogous to the fact that deleting an edge in a plane graph corresponds to contracting that edge in the plane dual.

For the third operation we have some freedom. One operation that corresponds to 3 is the inverse operation of contracting an edge. However this would not be compatible with the first line of thought and we are indeed able to make such a compatible choice as follows.

If an edge of the link graph corresponds to a face of C that is incident with only two edges of C, then contracting that face corresponds to contracting the corresponding edge in the link graph. It is not clear, however, how that definition could be extended to faces of size three (in particular if all edges incident with that face are loops; which we have to deal with as we allow contractions of edges of C).

Our solution is the following. Essentially, we are able to show that in order to construct a bounded obstruction in any non-embeddable 2-dimensional simplicial complex (which is the crucial step in a proof of a Kuratowski type theorem) that is nice enough, we only need to contract faces incident with two edges but not those of size three! Here 'nice enough' means simply connected and locally 3-connected. Both these conditions can be interpreted as face maximality conditions on the complex, see Theorem 2.8.1. 'Essentially' here means that additionally we have to allow for the following two (rather simple) operations.

If the link graph at a vertex v of a 2-complex C is disconnected, the 2-complex obtained from C by *splitting* the vertex v is obtained by replacing v

[12]Contractions of loops do not preserve embeddability in general (as $\mathbb{S}^3 / \mathbb{S}^1 \not\cong \mathbb{S}^3$).

by one new vertex for each connected component K of the link graph that is incident with the edges and faces in K.

Given an edge e in a 2-complex C, the 2-complex obtained from C by *deleting the edge* e is obtained from C by replacing e by parallel edges such that each new edge is incident with precisely one face (for an example, see the deletion of the edge g in Figure 5).

Remark 1.5.1. (On a variant of space minors and Theorem 0.0.1). In our proof we only ever split vertices directly after deleting edges or faces, and after such a deletion we can without changing the proof always split the incident vertices. Hence we could modify these two operations so that we always afterwards additionally split all vertices incident with the deleted edge or face. This way we would only have four space minor operations, one for each corner of Figure 4. And Theorem 0.0.1 would be true in this form.

Formally, let f be a face of size two in a 2-complex C, the 2-complex C/f obtained from C by *contracting* the face f is obtained from C by replacing the face f and its two incident edges by a single edge (also denoted by f). This new edge is incident with all faces that are incident with one of the two edges of f – and it is incident with the same vertices as f.

1.5.2 Basic properties

A *space minor* of a 2-complex is obtained by successively performing one of the five operations.

1. contracting an edge that is not a loop;
2. deleting a face (and all edges or vertices only incident with that face);
3. contracting a face of size one[13] or two if its two edges are not loops;
4. splitting a vertex;
5. deleting an edge (we also refer to that operation as 'forgetting the incidences of an edge').

Remark 1.5.2. A little care is needed with contractions of faces. This can create faces traversing edges multiple times. In this chapter, however, we do not contract faces consisting of two loops and we only perform these operations on 2-complexes whose faces have size at most three. Hence it could only happen that after contraction some face traverses an edge twice but in opposite direction. Since faces have size at most three, these traversals are adjacent. In this case we omit the two opposite traversals of the edge from the face. We delete faces incident with no edge. This ensures that the class of 2-complexes with faces of size at most three is closed under face contractions.

[13] Although we do not need it in our proofs, it seems natural to allow contractions of faces of size one.

A 2-complex is *3-bounded* if all its faces are incident with at most three edges. The closure of the class of simplicial complexes by space minors is the class of 3-bounded 2-complexes.

It is easy to see that the space minor operations preserve embeddability in $\mathbb{S}^3$ (or in any other 3-dimensional manifold) and the first three commute when defined.[14]

Lemma 1.5.3. *The space minor relation is well-founded.*

Proof. The *face degree* of an edge e is the number of faces incident with e. We consider the sum S of all face degrees ranging over all edges. None of the five above operations increases S. And 1, 2 and 3 always strictly decrease S. Hence we can apply 1, 2 or 3 only a bounded number of times.

Since no operation increases the sizes of the faces, the total number of vertices and edges incident with faces is bounded. Operation 4 increases the number of vertices and preserves the number of edges. For operation 5 it is the other way round. Hence we can also only apply[15] 4 and 5 a bounded number of times. □

Lemma 1.5.4. *If a 2-complex C has a planar rotation system, then all its space minors do.*

Proof. By Lemma 1.1.2 existence of planar rotation systems is preserved by contracting edges that are not loops. Clearly the operations 2, 4 and 5 preserve planar rotation systems as well. Since contracting a face of size two corresponds to locally in the link graph contracting the corresponding edges, contracting faces of size two preserves planar rotation systems as noted after Lemma 1.4.3. The operation that corresponds to contracting a face of size one in the link graph is explained in Figure 1.5. It clearly preserves embeddings in the plane. Thus contracting a face of size one also preserves planar rotation systems. □

1.5.3 Generalised Cones

In this subsection we define the list $\mathcal{Z}$ of obstructions appearing in Theorem 0.0.1 and prove basic properties of the related constructions.

Given a graph G without loops and a partition P of its vertex set into connected sets, the *generalised cone* over G with respect to P is the following (3-bounded) 2-complex C. Let H be the graph obtained from G by contracting each class of P to a single vertex and then removing some of the loops (and keeping parallel edges). The vertices of C are the vertices of H together with one extra vertex, which we call the *top (of the cone).* The edges of C are the edges of H together with one edge for each vertex e of G joining the top with the vertex of H that corresponds to the partition class containing e. We have one face for every edge f of G. If that edge is not an edge of H, its two endvertices

[14] In order for the contraction of a face to be defined we need the face to have at most two edges. This may force contractions of edges to happen before the contraction of the face.

[15] We exclude applications of 4 to a vertex whose link graph is connected and applications of 5 to edges incident with a single face.

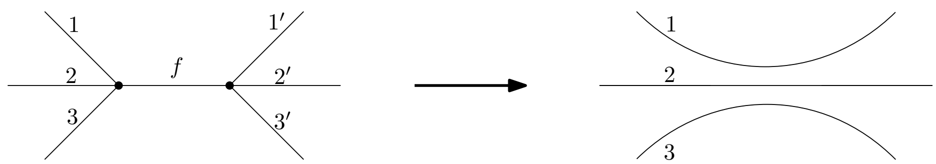

Figure 1.5: The operation that in the link graph corresponds to contracting a face f only incident with a single edge ℓ. The edge ℓ must be a loop. Hence in the link graph we have two vertices for ℓ which are joined by the edge f. On the left we depicted that configuration. Contracting f in the complex yields the configuration on the right. Formally, we delete f and both its endvertices and add for each face x of size at least two traversing ℓ an edge as follows. Before the contraction, the link graph contains two edges corresponding to the traversal of x of ℓ. These edges have precisely two distinct endvertices that are not vertices corresponding to ℓ. We add an edge between these two vertices.

in G are in the same partition class; this face is only incident with the two edges of C corresponding to these vertices. Otherwise the face is additionally incident with the edge of H that corresponds to f.

Example 1.5.5. The generalised cone construction has as a special case the cone construction; indeed we can just pick P to consist only of singletons. However, this construction has more flexibility, for example if G is connected and simple and P just consists of a single vertex, the construction gives a 2-complex with only two vertices such that G is the link graph at both vertices.

Lemma 1.5.6. *Let C be a 3-bounded 2-complex with a vertex v. If C has no loop incident with v, then C has a space minor that is a generalised cone whose link graph at the top is $L(v)$.*

Proof. We obtain C_1 from C by deleting all faces not incident with v. We obtain C_2 from C_1 by forgetting all incidences at the edges not incident with v. We obtain C_3 from C_2 by splitting all vertices different from v. It remains to prove the following.

Sublemma 1.5.7. *C_3 is a generalised cone over $L(v)$ with top v.*

Proof. Let w be a vertex of C_3 different from v. Since every face of C_3 has size two or three and is incident with v, there is an edge e with endvertices v and w. Let $P[w]$ be the set of those vertices e' of $L(v)$ such that there is a path from e to e' all of whose edges are faces of size two in C_3 or else faces of size three in C_3 that contain a loop. By construction, every edge in $P[w]$ is incident with w. Any edge in the link graph $L(w)$ of C_3 with only one endvertex in $P[w]$ must be a face of C_3 of size three. Let x be the endvertex of such an edge not in $P[w]$. Then x is an edge of C_3 that is not incident with v; and thus is only incident with a single face; that is, x has degree one in $L(w)$. Hence the connected component of e is contained in $P[w]$ and these attached leaves. By our construction $L(w)$ is connected, so $P[w]$ is a connected subset of $L(w)$ and is equal to the set of edges of C_3 between v and w.

It follows that C_3 is (isomorphic to) a generalised cone over the link graph $L(v)$ at the top v with respect to the partition $(P[w]|w \in V(C_3) - v)$. □

□

Lemma 1.5.8. *Let C be a generalised cone over a graph L, and let f be an edge of L. Then there is a space minor of C that is a generalised cone over $L - f$.*

Proof. We denote the top of C by v. We obtain C_1 from C by deleting the face f. We obtain C_2 from C_1 by splitting[16] all vertices of C_1 that are incident with f in C except for v. It is straightforward to check that C_2 is a generalised cone over $L - f$. □

The following is immediate from the definition of generalised cones.

Observation 1.5.9. *Let C be a generalised cone and e be an edge of C not incident with the top. If e is not a loop, then the space minor C/e is a generalised cone (over the same graph).*

Proof. We denote the two endvertices of the edge e by x and y. The link graph $L(e)$ at e in C/e is obtained from the link graphs $L(x)$ and $L(y)$ of C by gluing them together at the degree-one-vertex e, and then suppressing the vertex e. Thus the link graph $L(e)$ is connected. So C/e is a generalised cone over the same graph as the generalised cone C. □

Lemma 1.5.10. *Let C be a generalised cone over a graph L. Assume L contains an edge f that has an endvertex e of degree two. Then there is a space minor of C that is (isomorphic to) a generalised cone over L/f.*

Proof. We denote the edge incident with the vertex e aside from f by f'. The edges f and f' are in series in the graph L. Now we consider f and f' as faces of C. If one of these faces has size two, we contract it and denote the resulting complex by C'. It is then straightforward to check that C' is a generalised cone over L/f or L/f', respectively. As these two graphs are isomorphic, C' is a generalised cone over L/f.

Hence we may assume that the two faces f and f' have size three. We denote the edge of f not incident with the top v of C by x, and the edge of f' not incident with v by x'.

Recall that the edge e is incident with the top v of C. We denote the endvertex of e aside from v by w. In the link graph $L(w)$ at w, the vertex e is only incident with the edges f and f', and the other endvertices of these edges have degree one. Hence the connected component of $L(w)$ containing e is a path of length two. As $L(w)$ is connected by the definition of generalised cones, the link graph $L(w)$ must be equal to a path of length two with the vertices e, x and x'. So the edge x is not a loop.

[16] If the face f has size three, this additional splitting is trivial and hence not necessary.

Thus by Observation 1.5.9 $C'' = C/x$ is a generalised cone over L. In this generalised cone the face f has degree two. Hence C''/f is a generalised cone over the graph L/f. $\square$

Lemma 1.5.11. *Let C be a generalised cone and H be a subdivision of the link graph at the top that has no loops. Then C has a space minor that is a generalised cone over H.*

Proof. By Lemma 1.5.8, there is a space minor C_1 of C that is a generalised cone over a graph L' so that H can be obtained from L' by suppressing vertices of degree two. By Lemma 1.5.10, there is a space minor C_2 of C_1 that is a generalised cone over the graph H. $\square$

In the following we introduce 'looped generalised cones' and prove for them analogues of Lemma 1.5.6 and Lemma 1.5.11.

A *looped generalised cone* is obtained from a generalised cone by attaching a loop at the top of the cone, adding some faces of size one only containing that loop and adding the incidence with the loop to some existing faces of size two. This is well-defined as all faces of a generalised cone are incident with the top. The following is proved analogously to Lemma 1.5.6[17].

Lemma 1.5.12. *Let C be a 3-bounded 2-complex and let v be a vertex. If C has precisely one loop e incident with v, then C has a space minor that is a looped generalised cone whose link graph at the top is $L(v)$.* $\square$

We prove the following analogue of Lemma 1.5.11 for looped generalised cones. Given a graph L together with two specified vertices v and w, a *strict subdivision* of (L, v, w) is obtained by successively deleting edges from L or contracting edges that have an endvertex x of degree two such that neither x is equal to v or w nor x is adjacent to v or w. Given a looped generalised cone, we refer to the link graph at the top together with the two vertices of that link graph corresponding to the loop as the *specific link graph at the top*.

Lemma 1.5.13. *Let C be a looped generalised cone and let (H, v, w) be a strict subdivision of the specific link graph at the top. Assume that H has no loops. Then C has a space minor that is a looped generalised cone such that (H, v, w) is the specific link graph at the top.*

Proof. The proof of Lemma 1.5.13 is analogous to that of Lemma 1.5.11. Indeed, the analogous proof of that of Lemma 1.5.8 shows the following.

Sublemma 1.5.14. *Let C be a looped generalised cone over a graph L, and let f be an edge of L. Then there is a space minor of C that is a looped generalised cone over $L - f$.*

[17]The statement analogue to Sublemma 1.5.7 is that 'C_3 is a looped generalised cone over $L(v)$ with top v'. By the proof of that sublemma it follows that the 2-complex C_3/e, obtained from C_3 by contracting the loop e, is a generalised cone. Using the definition of looped generalised cone, it follows that C_3 is a looped generalised cone with the desired property.

Similarly like Lemma 1.5.10 we prove the following.

Sublemma 1.5.15. *Let C be a looped generalised cone with (L, v, w) as the specific graph at the top. Assume L contains an edge f that has an endvertex e of degree two such that e is neither equal to v or w nor adjacent to v or w. Then there is a space minor of C that is (isomorphic to) a looped generalised cone over L/f.*

Proof. Take the proof of Lemma 1.5.10 and replace 'generalised cone' by 'looped generalised cone'. The analogous statement of Observation 1.5.9 for looped generalised cones is also true. □

Thus we can apply the proof of Lemma 1.5.11 to prove Lemma 1.5.13, where we refer to Sublemma 1.5.14 in place of Lemma 1.5.8 and to Sublemma 1.5.15 in place of Lemma 1.5.10. □

Let $\mathcal{Z}_1$ be the set of generalised cones over the graphs K_5 or $K_{3,3}$. Let $\mathcal{Z}_2$ be the set of looped generalised cones such that some member of $\mathcal{Y}'$ is a strict marked graph associated to the link graph at the top. Let $\mathcal{Z}$ be the union of $\mathcal{Z}_1$ and $\mathcal{Z}_2$.

1.5.4 A Kuratowski theorem

In this subsection we prove Theorem 0.0.1. First we prove the following.

Theorem 1.5.16. *Let C be a simply connected locally 3-connected 2-dimensional simplicial complex. Then C has a planar rotation system if and only if C has no space minor from the finite list $\mathcal{Z}$.*

Proof. If C has a planar rotation system, it cannot have a space minor in $\mathcal{Z}$. Indeed, every complex Z in $\mathcal{Z}$ has a link graph that is not loop planar. Hence no Z in $\mathcal{Z}$ has a planar rotation system by Lemma 1.1.3. Since by Lemma 1.5.4 the class of 2-complexes with planar rotation systems is closed under space minors, C cannot have a space minor in $\mathcal{Z}$.

Now conversely assume that the simplicial complex C has no space minor in $\mathcal{Z}$. Suppose for a contradiction that C has no planar rotation system. Then by Lemma 1.3.1, there is a 3-bounded space minor C' that is not locally planar, where C' is either C, or for some (non-loop) edge e the contraction C/e or there is a (non-loop) chordless cycle o of C and some $e \in o$ such that $C' = C/(o - e)$. We distinguish two cases.

Case 1: C or C/e are not locally planar. Since C has no parallel edges or loops by assumption, here C' has no loop. Hence C' has a vertex v such that the link graph $L(v)$ at v is not planar. By Lemma 1.5.6 C' has a space minor that is a generalised cone such that the link graph at the top is $L(v)$. By Kuratowski's theorem, $L(v)$ has a subdivision isomorphic to K_5 or $K_{3,3}$. So by Lemma 1.5.11 C' has a space minor that is a generalised cone over K_5 or $K_{3,3}$. So C has a space minor in $\mathcal{Z}_1$, which is the desired contradiction.

Case 2: Not Case 1. So $C' = C/(o - e)$. Let v be the vertex of C' corresponding to $o - e$. Since we are not in Case 1, all link graphs at vertices of C are loop planar. In particular, it must be the link graph at v that is not loop planar.

Sublemma 1.5.17. *If the link graph $L(v)$ of C' is not planar, there is an edge $e' \in o - e$ such that the link graph at e' in C/e' is not planar.*

Proof. We prove the contrapositive. So assume that for every edge $e' \in o - e$ the link graph at e' of C/e' is planar. Since C is locally 3-connected, the planar rotation systems of the link graphs $L(w)$ at the vertices w of o are unique up to reversing. By Lemma 1.3.2 these rotation systems are reverse or agree at any rotator of a vertex in $o - e$.

Note that $L(v)$ is the vertex sum of the link graphs $L(w)$ along the set $o - e$ of gluing vertices. Thus by reversing some of these rotation systems if necessary, we can apply Lemma 1.2.2 to build a planar rotation system of $L(v)$. In particular, $L(v)$ is planar. □

By Sublemma 1.5.17 and since we are not in Case 1, the link graph $L(v)$ is planar but not loop planar.

Since C has no loops and parallel edges and o is chordless, in this case C' can only have the loop e. Thus by Lemma 1.5.12 C' has a space minor C'' that is a looped generalised cone such that the link graph at the top is $L(v)$.

Since C is locally 3-connected by assumption and by Lemma 1.2.4 the link graph $L(v)$ is 3-connected. So by Lemma 1.4.1 there is a marked graph $\hat{G}$ associated to $L(v)$ that is not planar. Let G' be a strict marked graph associated to $L(v)$ with marked graph $\hat{G}$. By Lemma 1.4.19 G' has a strict marked minor $\hat{Y} = (Y, x, z, (a_i, b_i)|i = 1, 2, 3; \iota)$ in $\mathcal{Y}'$, where x and z are the vertices in $L(v)$ corresponding to the loop at v. By the definition of strict subdivision, we have that (Y, x, z) is a strict subdivision of the specific link graph $(L(v), x, z)$ at the top of C''. So by Lemma 1.5.13 C' has a space minor that is a looped generalised cone such that $\hat{Y}$ is a strict marked graph associated to the top. So C has a space minor in $\mathcal{Z}_2$, which is the desired contradiction. □

Proof of Theorem 0.0.1. By Theorem 1.1.1 a simply connected simplicial complex is embeddable in $\mathbb{S}^3$ if and only if it has a planar rotation system. So Theorem 0.0.1 is implied by Theorem 1.5.16. □

Proof of Theorem 0.0.2. By Theorem 2.2.2 a simplicial complex with $H_1(C, \mathbb{F}_p) = 0$ is embeddable if and only if it is simply connected and it has a planar rotation system. So Theorem 0.0.2 is implied by Theorem 1.5.16. □

1.6 Concluding remarks

The proof of Theorem 0.0.1 yields that quite a few properties are equivalent. This is summarised in the following.

Theorem 1.6.1. *Let C be a simply connected locally 3-connected 2-dimensional simplicial complex. The following are equivalent.*

1. *C has an embedding in the 3-sphere;*
2. *C has an embedding in some 3-manifold;*
3. *C has a planar rotation system;*
4. *all contractions of C are locally planar;*
5. *no contraction has a link graph that has K_5 or $K_{3,3}$ as a minor or a marked minor of the 12 marked graphs in the list $\mathcal{Y}$ defined in Section 1.4;*
6. *C has no space minor from the finite list $\mathcal{Z}$ defined in Subsection 1.5.3.*

Proof. The equivalence between 1, 2 and 3 is proved in Chapter 2. The equivalence between 3 and 4 is proved in Lemma 1.1.3. The equivalence between 1 and 5 is Corollary 1.4.10. Finally, the equivalence between 1 and 6 is Theorem 0.0.1. □

Theorem 0.0.2 is a structural characterisation of which locally 3-connected 2-dimensional simplicial complex C whose first homology group is trivial embed in 3-space. Does this have algorithmic consequences? The methods of this chapter give an algorithm that checks in linear[18] time whether a locally 3-connected 2-dimensional simplicial complex has a planar rotation system. For general 2-dimensional simplicial complex we obtain a quadratic algorithm, see Chapter 5 for details. But how easy is it to check whether C is simply connected? For simplicial complexes in general this is not decidable; indeed for every finite presentation of a group one can build a 2-dimensional simplicial complex that has that fundamental group. However, for simplicial complexes that embed in some 3-manifold; that is, that have a planar rotation system, this problem is known as the sphere recognition problem. Recently it was shown that sphere recognition lies in NP [59, 84] and co-NP assuming the generalised Riemann hypothesis [57, 103]. It is an open question whether there is a polynomial time algorithm.

[18]Linear in the number of faces of C.

Chapter 2

Rotation systems

2.1 Abstract

We prove that 2-dimensional simplicial complexes whose first homology group is trivial have topological embeddings in 3-space if and only if there are embeddings of their link graphs in the plane that are compatible at the edges and they are simply connected.

2.2 Introduction

Here we give combinatorial characterisations for when certain simplicial complexes embed in 3-space. This completes the proof of a 3-dimensional analogue of Kuratowski's characterisation of planarity for graphs, started in Chapter 1.

A (2-dimensional) simplicial complex has a topological embedding in 3-space if and only if it has a piece-wise linear embedding if and only if it has a differential embedding [11, 55, 77].[1] Perelman proved that every compact simply connected 3-dimensional manifold is isomorphic to the 3-sphere $\mathbb{S}^3$ [79, 80, 81]. In this chapter we use Perelman's theorem to prove a combinatorial characterisation of which simply connected simplicial complexes can be topologically embedded into $\mathbb{S}^3$ as follows.

The *link graph* at a vertex v of a simplicial complex is the graph whose vertices are the edges incident with v and whose edges are the faces incident with v and the incidence relation is as in C, see Figure 2.1. Roughly, a *planar rotation system* of a simplicial complex C consists of cyclic orientations $\sigma(e)$ of the faces incident with each edge e of C such that there are embeddings in the plane of the link graphs such that at vertices e the cyclic orientations of the incident edges agree with the cyclic orientations $\sigma(e)$. It is easy to see that if a simplicial complex C has a topological embedding into some 3-dimensional manifold, then

[1] However this is not equivalent to having a linear embedding, see [20], and [69] for further references.

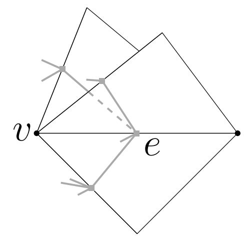

Figure 2.1: The link graph at the vertex v is indicated in grey. The edge e projects down to a vertex in the link graph. The faces incident with e project down to edges.

it has a planar rotation system. Conversely, for simply connected simplicial complexes the existence of planar rotation systems is enough to characterise embeddability into $\mathbb{S}^3$:

Theorem 2.2.1. *Let C be a simply connected simplicial complex. Then C has a topological embedding into $\mathbb{S}^3$ if and only if C has a planar rotation system.*

The main result of this chapter is the following extension of Theorem 2.2.1.

Theorem 2.2.2. *Let C be a simplicial complex such that the first homology group $H_1(C, \mathbb{F}_p)$ is trivial for some prime p. Then C has a topological embedding into $\mathbb{S}^3$ if and only if C is simply connected and it has a planar rotation system.*

This implies characterisations of topological embeddability into $\mathbb{S}^3$ for the classes of simplicial complexes with abelian fundamental group and simplicial complexes in general, see Section 2.8 for details.

The chapter is organised as follows. After reviewing some elementary definitions in Section 2.3, in Section 2.4, we introduce rotation systems, related concepts and prove basic properties of them. In Sections 2.5 and 2.6 we prove Theorem 2.2.1. The proof of Theorem 2.2.2 in Section 2.7 makes use of Theorem 2.2.1. Further extensions are derived in Section 2.8.

2.3 Basic definitions

In this short section we recall some elementary definitions that are important for this chapter.

A *closed trail* in a graph is a cyclically ordered sequence $(e_n | n \in \mathbb{Z}_k)$ of distinct edges e_n such that the starting vertex of e_n is equal to the endvertex of e_{n-1}. An (abstract) (2-dimensional) *complex* is a graph[2] G together with a family of closed trails in G, called the *faces* of the complex. We denote complexes C by triples $C = (V, E, F)$, where V is the set of *vertices*, E the set of *edges* and F the set of faces. We assume furthermore that every vertex of a

[2]In this chapter graphs are allowed to have parallel edges and loops.

complex is incident with an edge and every edge is incident with a face. The *1-skeleton* of a complex $C = (V, E, F)$ is the graph (V, E). A *directed* complex is a complex together with a choice of direction at each of its edges and a choice of orientation at each of its faces. For an edge e, we denote the direction chosen at e by $\vec{e}$. For a face f, we denote the orientation chosen at f by $\vec{f}$.

Examples of complexes are (abstract) (2-dimensional) simplicial complexes. In this chapter all simplicial complexes are directed – although we will not always say it explicitly. A *(topological) embedding* of a simplicial complex C into a topological space X is an injective continuous map from (the geometric realisation of) C into X. In our notation we suppress the embedding map and for example write '$\mathbb{S}^3 \setminus C$' for the topological space obtained from $\mathbb{S}^3$ by removing all points in the image of the embedding of C.

In this chapter, a *surface* is a compact 2-dimensional manifold (without boundary)[3]. Given an embedding of a graph in an oriented surface, the *rotation system* at a vertex v is the cyclic orientation[4] of the edges incident with v given by 'walking around' v in the surface in a small circle in the direction of the orientation. Conversely, a choice of rotation system at each vertex of a graph G defines an embedding of G in an oriented surface as explained in Chapter 1.

A *cell complex* is a graph G together with a set of directed walks such that each direction of an edge of G is in precisely one of these directed walk es. These directed walks are called the *cells*. The geometric realisation of a cell complex is obtained from (the geometric realisation of) its graph by gluing discs so that the cells are the boundaries of these discs. The geometric realisation is always an oriented surface. Note that cell complexes need not be complexes as cells are allowed to contain both directions of an edge. The *rotation system of* a cell complex C is the rotation system of the graph of C in the embedding in the oriented surface given by C.

2.4 Rotation systems

In this section we introduce rotation systems of complexes and some related concepts.

The *link graph* of a simplicial complex C at a vertex v is the graph whose vertices are the edges incident with v. The edges are the faces incident[5] with v. The two endvertices of a face f are those vertices corresponding to the two edges of C incident with f and v. We denote the link graph at v by $L(v)$.

A *rotation system* of a directed complex C consists of for each edge e of C a cyclic orientation[6] $\sigma(e)$ of the faces incident with e.

Important examples of rotation systems are those *induced* by topological embeddings of complexes C into $\mathbb{S}^3$ (or more generally in some 3-manifold); here for an edge e of C, the cyclic orientation $\sigma(e)$ of the faces incident with e

[3]We allow surfaces to be disconnected.

[4]A *cyclic orientation* is a choice of one of the two orientations of a cyclic ordering.

[5]A face is incident with a vertex if there is an edge incident with both of them.

[6]If the edge e is only incident with a single face, then $\sigma(e)$ is empty.

is the ordering in which we see the faces when walking around some midpoint of e in a circle of small radius[7] – in the direction of the orientation of $\mathbb{S}^3$. It can be shown that $\sigma(e)$ is independent of the chosen circle if small enough and of the chosen midpoint.

Such rotation systems have an additional property: let $\Sigma = (\sigma(e)|e \in E(C))$ be a rotation system of a simplicial complex C induced by a topological embedding of C in the 3-sphere. Consider a ball of small radius around a vertex v. We may assume that each edge of C intersects the boundary of that ball in at most one point and that each face intersects it in an interval or not at all. The intersection of the boundary of the ball and C is a graph: the link graph at v. Hence link graphs of complexes with induced rotation systems must always be planar. And even more: the cyclic orientations $\sigma(e)$ at the edges of C form – when projected down to a link graph to rotators at the vertices of the link graph – a rotation system at the link graph, see Figure 2.2.

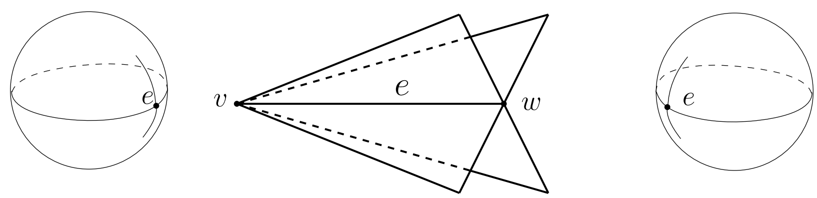

Link at v The edge e with its incident faces Link at w

Figure 2.2: The cyclic orientation $\sigma(e)$ of the faces incident with the edge e of C projects down to rotators at e in the link graphs at either endvertex of e. In these link graphs, the projected rotators at e are reverse.

Next we shall define 'planar rotation systems' which roughly are rotation systems satisfying such an additional property. The cyclic orientation $\sigma(e)$ at the edge e of a rotation system defines a rotation system $r(e, v, \Sigma)$ at each vertex e of a link graph $L(v)$: if the directed edge $\vec{e}$ is directed towards v we take $r(e, v, \Sigma)$ to be $\sigma(e)$. Otherwise we take the inverse of $\sigma(e)$. As explained in Section 2.3, this defines an embedding of the link graph into an oriented surface. The *link complex* for (C, Σ) at the vertex v is the cell complex obtained from the link graph $L(v)$ by adding the faces of the above embedding of $L(v)$ into the oriented surface. By definition, the geometric realisation of the link complex is always a surface. To shortcut notation, we will not distinguish between the link complex and its geometric realisation and just say things like: 'the link complex is a sphere'. A *planar rotation system* of a directed simplicial C is a rotation system such that for each vertex v all link graphs are a disjoint union of spheres. The paragraph before shows the following.

[7]Formally this means that the circle intersects each face in a single point and that it can be contracted onto the chosen midpoint of e in such a way that the image of one such contraction map intersects each face in an interval.

Observation 2.4.1. *Rotation systems induced by topological embeddings of locally connected*[8] *simplicial complexes in 3-manifolds are planar.* □

Next we will define the *local surfaces of a topological embedding* of a simplicial complex C into $\mathbb{S}^3$. The local surface at a connected component of $\mathbb{S}^3 \setminus C$ is the following. Pour concrete into this connected component. The surface of the concrete is a 2-dimensional manifold. The local surface is the simplicial complex drawn at the surface by the vertices, edges and faces of C. Note that if an edge e of G is incident with more than two faces that are on the surface, then the surface will contain at least two clones of the edge e, see Figure 2.3.

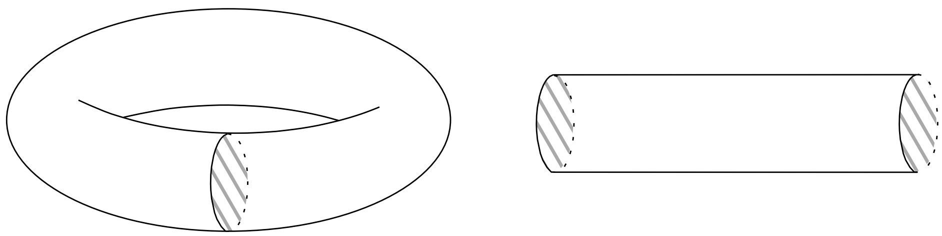

Figure 2.3: On the left we depicted the torus with an additional face attached on a genus reducing curve in the inside. On the right we depicted the local surface of its inside component. It is a sphere and contains two copies of the newly added face (and its incident edges).

Now we will define *local surfaces for a pair* (C, Σ) consisting of a complex C and one of its rotation systems Σ. Lemma 2.4.4 below says that under fairly general circumstances the local surfaces of a topological embedding are the local surfaces of the rotation system induced by that topological embedding. The set of faces of a local surface will be an equivalence class of the set of orientations of faces of C. The *local-surface-equivalence relation* is the symmetric transitive closure of the following relation. An orientation $\vec{f}$ of a face f is *locally related* via an edge e of C to an orientation $\vec{g}$ of a face g if f is just before g in $\sigma(e)$ and e is traversed positively by $\vec{f}$ and negatively by $\vec{g}$ and in $\sigma(e)$ the faces f and g are adjacent. Here we follow the convention that if the edge e is only incident with a single face, then the two orientations of that face are related via e. Given an equivalence class of the local-surface-equivalence relation, the *local surface* at that equivalence class is the following complex whose set of faces is (in bijection with) the set of orientations in that equivalence class. We obtain the complex from the disjoint union of the faces of these orientations by gluing together two of these faces f_1 and f_2 along two of their edges if these edges are copies of the same edge e of C and f_1 and f_2 are related via e. Of course, we glue these two edges in a way that endvertices are identified only with copies of the same

[8] Observation 2.4.1 is also true without the assumption of 'local connectedness'. In that case however the link complex is disconnected. Hence it is no longer directly given by the drawing of the link graph on a ball of small radius as above.

vertex of C. Hence each edge of a local surface is incident with precisely two faces. Hence its geometric realisation is always a is a surface. Similarly as for link complexes, we shall just say things like 'the local surface is a sphere'.

Observation 2.4.2. *Local surfaces of planar rotation systems are always connected.* □

A *(2-dimensional) orientation* of a complex C such that each edge is in precisely two faces is a choice of orientation of each face of C such that each edge is traversed in opposite directions by the chosen orientation of the two incident faces. Note that a complex whose geometric realisation is a surface has an orientation if and only its geometric realisation is orientable.

Observation 2.4.3. *The set of orientations in a local-surface-equivalence class defines an orientation of its local surface.*

In particular, local surfaces are cell complexes. □

We will not use the following lemma in our proof of Theorem 2.2.1 However, we think that it gives a better intuitive understanding of local surfaces. We say that a simplicial complex C is *locally connected* if all link graphs are connected.

Lemma 2.4.4. *Let C be a connected and locally connected complex embedded into $\mathbb{S}^3$ and let Σ be the induced planar rotation system. Then the local surfaces of the topological embedding are equal to the local surfaces for (C, Σ).* □

There is the following relation between vertices of local surfaces and faces of link complexes.

Lemma 2.4.5. *Let Σ be a rotation system of a simplicial complex C. There is a bijection ι between the set of vertices of local surfaces for (C, Σ) and the set of faces of link complexes for (C, Σ), which maps each vertex v' of a local surface cloned from the vertex v of C to a face f of the link complex at v such that the rotation system at v' is an orientation of f.*

Proof. The set of faces of the link complex at v is in bijection with the set of v-equivalence classes; here the *v-equivalence relation* on the set of orientations of faces of C incident with v is the symmetric transitive closure of the relation 'locally related'. Since we work in a subset of the orientations, every v-equivalence class is contained in a local-surface-equivalence class. On the other hand the set of all clones of a vertex v of C contained in a local surface S is in bijection with the set of v-equivalence classes contained in the local-surface-equivalence class of S. This defines a bijection ι between the set of vertices of local surfaces for (C, Σ) and the set of faces of link complexes for (C, Σ).

It is straightforward to check that ι has all the properties claimed in the lemma. □

Corollary 2.4.6. *Given a local surface of a simplicial complex C and one of its vertices v' cloned from a vertex v of C, there is a homeomorphism from a neighbourhood around v' in the local surface to the cone with top v' over the face boundary of $\iota(v')$ that fixes v' and the edges and faces incident with v' in a neighbourhood around v'.* □

The definitions of link graphs and link complexes can be generalised from simplicial complexes to complexes as follows. The *link graph* of a complex C at a vertex v is the graph whose vertices are the edges incident with v. For any traversal of a face of the vertex v, we add an edge between the two vertices that when considered as edges of C are in the face just before and just after that traversal of v. We stress that we allow parallel edges and loops. Given a complex C, any rotation system Σ of C defines rotation systems at each link graph of C. Hence the definition of link complex extends.

2.5 Constructing piece-wise linear embeddings

In this section we prove Theorem 2.5.4 below, which is used in the proof of Theorem 2.2.1.

Throughout this section we fix a connected and locally connected simplicial complex C with a rotation system Σ. An associated topological space $T(C,\Sigma)$ is defined as follows. For each local surface S of (C,Σ) we take an embedding into $\mathbb{S}^3$ as follows. Let g be the genus of S. We start with the unit ball in $\mathbb{S}^3$ and then identify g disjoint pairs of discs through the outside.[9] The constructed surface is isomorphic to S, so this defines an embedding of S. Each local surface is oriented and we denote by $\hat{S}$ the topological space obtained from $\mathbb{S}^3$ by deleting all points on the outside of S. We obtain $T(C,\Sigma)$ from the simplicial complex C by gluing onto each local surface S the topological space $\hat{S}$ along S.

We remark that associated topological spaces may depend on the chosen embeddings of the local surfaces S into $\mathbb{S}^3$. However, if all local surfaces are spheres, then any two associated topological spaces are isomorphic and in this case we shall talk about 'the' associated topological space.

Clearly, associated topological spaces $T(C,\Sigma)$ are compact and connected as C is connected.

Lemma 2.5.1. *The rotation system Σ is planar if and only if the associated topological space $T(C,\Sigma)$ is a 3-dimensional manifold.*

Proof. Observation 2.4.1 implies that if $T(C,\Sigma)$ is a 3-dimensional manifold, then Σ is planar. Conversely, now assume that Σ is a planar rotation system. We have to show that there is a neighbourhood around any point x of $T(C,\Sigma)$ that is isomorphic to the closed 3-dimensional ball B_3.

If x is a point not in C, this is clear. If x is an interior point of a face f, we obtain a neighbourhood of x by gluing together neighbourhoods of copies of x in the local surfaces that contain an orientation of f. Each orientation of f is contained in local surfaces exactly once. Hence we glue together the two orientations of f and clearly x has a neighbourhood isomorphic to B_3.

Next we assume that x is an interior point of an edge e. Some open neighbourhood of x is isomorphic to the topological space obtained from gluing together for each copy of e in a local surface, a neighbourhood around a copy x'

[9] We have some flexibility along which paths on the outside we do these identifications but we do not need to make particular choices for our construction to work.

of x on those edges. A neighbourhood around x' has the shape of a piece of a cake, see Figure 2.4

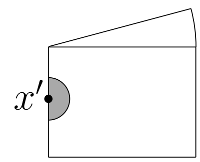

Figure 2.4: A piece of a cake. This space is obtained by taking the product of a triangle with the unit interval. The edge e is mapped to the set of points corresponding to some vertex of the triangle.

First we consider the case that x has several copies. As $\sigma(e)$ is a cyclic orientation, these pieces of a cake are glued together in a cyclic way along faces. Since each cyclic orientation of a face appears exactly once in local surfaces, we identify in each gluing step the two cyclic orientations of a face. Informally, the overall gluing will be a 'cake' with x as an interior point. Hence a neighbourhood of x is isomorphic to B_3. If there is only one copy of x', then the copy of e containing x' is incident with the two orientations of a single face. Then we obtain a neighbourhood of x by identifying these two orientations. Hence there is a neighbourhood of x isomorphic to B_3.

It remains to consider the case where x is a vertex of C. We obtain a neighbourhood of x by gluing together neighbourhoods of copies of x in local surfaces. We shall show that we have one such copy for every face of the link complex for (C, Σ) and a neighbourhood of x in such a copy is given by the cone over that face with x being the top of the cone, see Figure 2.5. We shall show

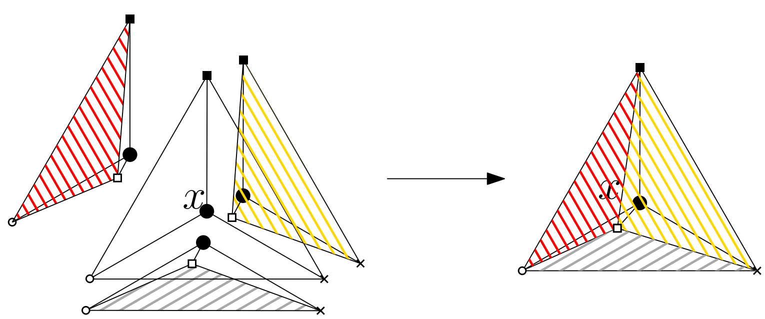

Figure 2.5: In this example the link complex of x is a tetrahedron. The three faces visible in our drawing are highlighted in red, gold and grey. On the left we see how the four cones over the faces of the link complex are pasted together to form the cone over the link complex depicted on the right.

that the glued together neighbourhood is the cone over the link complex with x at the top. Since Σ is planar and C is locally connected, the link complex is isomorphic to the 2-sphere. Since the cone over the 2-sphere is a 3-ball, the neighbourhood of x has the desired type.

Now we examine this plan in detail. By Lemma 2.4.5 and Corollary 2.4.6, the copies are mapped by the bijection ι to the faces of the link complex at x and a neighbourhood around such a copy x' is isomorphic to the cone with top x' over the face $\iota(x')$. We glue these cones over the faces $\iota(x')$ on their faces that are obtained from edges of $\iota(x')$ by adding the top x'.

The glued together complex is isomorphic to the cone over the complex S obtained by gluing together the faces $\iota(x')$ along edges, where we always glue the edge the way round so that copies of the same vertex of the local incidence graph are identified. Hence the vertex-edge-incidence relation and the edge-face-incidence relation of S are the same as for the link complex at x. The same is true for the cyclic orderings of edges on faces. So S is equal to the link complex at x.

Hence a neighbourhood of x is isomorphic to a cone with top x over the link complex at x. Since Σ is a planar rotation system, the link complex is a disjoint union of spheres. As C is locally connected, it is a sphere. Thus its cone is isomorphic to B_3. □

Lemma 2.5.2. *If C is simply connected, then so is any associated topological space $T(C, \Sigma)$.*

Proof. This is a consequence of Van Kampen's Theorem [56, Theorem 1.20]. Indeed, we obtain T from $T(C, \Sigma)$ by deleting all interior points of the sets $\hat{S}$ for local surfaces S that are not in a small open neighbourhood of C. This can be done in such a way that T has a deformation retract to C, and thus is simply connected. Now we recursively glue the spaces $\hat{S}$ back onto T. In each step we glue a single space $\hat{S}$. Call the space obtained after n gluings T_n.

The fundamental group of $\hat{S}$ is a quotient of the fundamental group of the intersection of T_n and $\hat{S}$ (this follows from the construction of the embedding of the surface S into the space $\hat{S}$). And the fundamental group of T_n is trivial by induction. So we can apply Van Kampen's Theorem to deduce that the gluing space T_{n+1} has trivial fundamental group. Hence the final gluing space $T(C, \Sigma)$ has trivial fundamental group. So it is simply connected. □

The converse of Lemma 2.5.2 is true if all local surfaces for (C, Σ) are spheres.

Lemma 2.5.3. *If all local surfaces for (C, Σ) are spheres and the associated topological space $T(C, \Sigma)$ is simply connected, then so is C.*

Proof. Let φ an image of $\mathbb{S}^1$ in C. Since $T(C, \Sigma)$ is simply connected, there is a homotopy from φ to a point of C in $T(C, \Sigma)$. We can change the homotopy so that it avoids an interior point of each local surface of the embedding. Since each local surface is a sphere, for each local surface without the chosen point there is a continuous projection to its boundary. Since these projections are continuous, the concatenation of them with the homotopy is continuous. Since this concatenation is constant on C this defines a homotopy of φ inside C. Hence C is simply connected. □

We conclude this section with the following special case of Theorem 2.2.1.

Theorem 2.5.4. *A locally connected simplicial complex C has a planar rotation system Σ if and only if $T(C,\Sigma)$ is a 3-manifold. And if C is simply connected, then $T(C,\Sigma)$ must be the 3-sphere.*

Proof. By treating different connected components separately, we may assume that C is connected. The first part follows from Lemma 2.5.1. The second part follows from Lemma 2.5.2 and Perelman's theorem [79, 81, 80] that any compact simply connected 3-manifold is isomorphic to the 3-sphere. □

Remark 2.5.5. We used Perelman's theorem in the proof of Theorem 2.5.4. On the other hand it together with Moise's theorem [72] that every compact 3-dimensional manifold has a triangulation implies Perelman's theorem: let M be a simply connected 3-dimensional compact manifold. Let T be a triangulation of M. And let C be the simplicial complex obtained from T by deleting the 3-dimensional cells. Let Σ be the rotation system given by the embedding of C into T. It is clear from that construction that T is equal to the triangulation given by the embedding of C into $T(C,\Sigma)$. Hence we can apply Lemma 2.5.3 to deduce that C is simply connected. Hence by Theorem 2.5.4 the topological space $T(C,\Sigma)$, into which C embeds, is isomorphic to the 3-sphere. Since $T(C,\Sigma)$ is isomorphic to M, we deduce that M is isomorphic to the 3-sphere.

2.6 Cut vertices

In this section we deduce Theorem 2.2.1 from Theorem 2.5.4 proved in the last section. Given a prime p, a simplicial complex C is *p-nullhomologous* if every directed cycle of C is generated over $\mathbb{F}_p$ by the boundaries of faces of C. Note that a simplicial complex C is p-nullhomologous if and only if the first homology group $H_1(C,\mathbb{F}_p)$ is trivial. Clearly, every simply connected simplicial complex is p-nullhomologous.

A vertex v in a connected complex C is a *cut vertex* if the 1-skeleton of C without v is a disconnected graph[10]. A vertex v in an arbitrary, not necessarily connected, complex C is *a cut vertex* if it is a cut vertex in a connected component of C.

Lemma 2.6.1. *Every p-nullhomologous simplicial complex without a cut vertex is locally connected.*

Proof. We construct for any vertex v of an arbitrary simplicial complex C such that the link graph $L(v)$ at v is not connected and v is not a cut vertex a cycle containing v that is not generated by the face boundaries of C.

[10]We define this in terms of the 1-skeleton instead of directly in terms of C for a technical reason: The object obtained from a simplicial complex by deleting a vertex may have edges not incident with faces. So it would not be a 2-dimensional simplicial complex in the terminology of this chapter.

Let e and g be two vertices in different components of $L(v)$. These are edges of C and let w and u be their endvertices different from v. Since v is not a cut vertex, there is a path in C between u and w that avoids v. This path together with the edges e and g is a cycle o in C that contains v.

Our aim is to show that o is not generated by the boundaries of faces of C. Suppose for a contradiction that o is generated. Let F be a family of faces whose boundaries sum up to o. Let F_v be the subfamily of faces of F that are incident with v. Each face in F_v is an edge of $L(v)$ and each vertex of $L(v)$ is incident with an even number (counted with multiplicities) of these edges except for e and g that are incident with an odd number of these faces. Let X be the connected component of the graph $L(v)$ restricted to the edge set F_v that contains the vertex e. We obtain X' from X by adding $k-1$ parallel edges to each edge that appears k times in F_v. Since X' has an even number of vertices of odd degree also g must be in X. This is a contradiction to the assumption that e and g are in different components of $L(v)$. Hence o is not generated by the boundaries of faces of C. This completes the proof. □

Given a connected complex C with a cut vertex v and a connected component K of the 1-skeleton of C with v deleted, the *complex attached* at v centered at K has vertex set $K+v$ and its edges and faces are those of C all of whose incident vertices are in $K+v$.

Lemma 2.6.2. *A connected simplicial complex C with a cut vertex v has a piece-wise linear embedding into $\mathbb{S}^3$ if and only if all complexes attached at v have a piece-wise linear embedding into $\mathbb{S}^3$.*

Proof. If C has an embedding into $\mathbb{S}^3$, then clearly all complexes attached at v have an embedding. Conversely suppose that all complexes attached at v have an embedding into $\mathbb{S}^3$. Pick one of these complexes arbitrarily, call it X and fix an embedding of it into $\mathbb{S}^3$. In that embedding pick for each component of C remove v except that for X a closed ball contained in $\mathbb{S}^3$ that intersects X precisely in v such that all these closed balls intersect pairwise only at v. Each complex attached at v, has a piece-wise linear embedding into the 3-dimensional unit ball as they have embeddings into $\mathbb{S}^3$ such that some open set is disjoint from the complex. Now we attach these embeddings into the balls of the embedding of X inside the reserved balls by identifying the copies of v. This defines an embedding of C. □

Recall that in order to prove Theorem 2.2.1 it suffices to show that any simply connected simplicial complex C has a piece-wise linear embedding into $\mathbb{S}^3$ if and only if C has a planar rotation system.

Proof of Theorem 2.2.1. Clearly if a simplicial complex is embeddable into $\mathbb{S}^3$, then it has a planar rotation system. For the other implication, let C be a simply connected simplicial complex and Σ be a planar rotation system. We prove the theorem by induction on the number of cut vertices of C. If C has no cut vertex, it is locally connected by Lemma 2.6.1. Thus it has a piece-wise linear embedding into $\mathbb{S}^3$ by Theorem 2.5.4.

Hence we may assume that C has a cut vertex v. As C is simply connected, every complex attached at v is simply connected. Hence by the induction hypothesis each of these complexes has a piece-wise linear embedding into $\mathbb{S}^3$. Thus C has a piece-wise linear embedding into $\mathbb{S}^3$ by Lemma 2.6.2. $\square$

2.7 Local surfaces of planar rotation systems

The aim of this section is to prove Theorem 2.2.2. A shorter proof is sketched in Remark 2.7.10 using algebraic topology. As a first step in that direction, we first prove the following.

Theorem 2.7.1. *Let C be a locally connected p-nullhomologous simplicial complex that has a planar rotation system. Then all local surfaces of the planar rotation system are spheres.*

Before we can prove Theorem 2.7.1 we need some preparation. The complex *dual* to a simplicial C with a rotation system Σ has as its set of vertices the set of local surfaces of Σ. Its set of edges is the set of faces of C, and an edge is incident with a vertex if the corresponding face is in the corresponding local surface. The faces of the dual are the edges of C. Their cyclic ordering is as given by Σ. In particular, the edge-face-incidence-relation of the dual is the same as that of C but with the roles of edges and faces interchanged.

Moreover, an orientation $\vec{f}$ of a face f of C corresponds to the direction of f when considered as an edge of the dual complex D that points towards the vertex of D whose local-surface-equivalence class contains $\vec{f}$. Hence the direction of the dual complex C induces a direction of the complex D. By $\Sigma_C = (\sigma_C(f)|f \in E(D))$ we denote the following rotation system for D: for $\sigma_C(f)$ we take the orientation $\vec{f}$ of f in the directed complex C.

In this chapter we follow the convention that for edges of C we use the letter e (with possibly some subscripts) while for faces of C we use the letter f. In return, we use the letter f for the edges of a dual complex of C and e for its faces.

Lemma 2.7.2. *Let C be a connected and locally connected simplicial complex. Then for any rotation system, the dual complex D is connected.*

Proof. Two edges of C are *C-related* if there is a face of C incident with both of them. And they are *C-equivalent* if they are in the transitive closure of the symmetric relation 'C-related'. Clearly, any two C-equivalent edges of C are in the same connected component. If C however is locally connected, also the converse is true: any two edges in the same connected component are C-equivalent. Indeed, take a path containing these two edges. Any two edges incident with a common vertex are C-equivalent as C is locally connected. Hence any two edges on the path are C-equivalent.

We define *D-equivalent* like 'C-equivalent' with 'D' in place of 'C'. Now let f and f' be two edges of D. Let e and e' be edges of C incident with f and f', respectively. Since C is connected and locally connected the edges e and e'

are C-equivalent. As C and D have the same edge/face incidence relation, the edges f and f' of D are D-equivalent. So any two edges of D are D-equivalent. Hence D is connected. $\square$

First, we prove the following, which is reminiscent of euler's formula.

Lemma 2.7.3. *Let C be a locally connected p-nullhomologous simplicial complex with a planar rotation system and D the dual complex. Then*

$$|V(C)| - |E| + |F| - |V(D)| \geq 0$$

Moreover, we have equality if and only if D is p-nullhomologous.

Proof. Let Z_C be the dimension over $\mathbb{F}_p$ of the cycle space of C. Similarly we define Z_D. Let r be the rank of the edge-face-incidence matrix over $\mathbb{F}_p$. Note that $r \leq Z_D$ and that $r = Z_C$ as $H_1(C, \mathbb{F}_p) = 0$. So $Z_D - Z_C \geq 0$. Hence it suffices to prove the following.

Sublemma 2.7.4.

$$|V(C)| - |E| + |F| - |V(D)| = Z_D - Z_C$$

Proof. Let k_C be the number of connected components of C and k_D be the number of connected components of D. Recall that the space orthogonal to the cycle space (over $\mathbb{F}_p$) in a graph G has dimension $|V(G)|$ minus the number of connected components of G. Hence $Z_C = |E| - |V(C)| + k_C$ and $Z_D = |F| - |V(D)| + k_D$. Subtracting the first equation from the second yields:

$$|V(C)| - |E| + |F| - |V(D)| + (k_D - k_C) = Z_D - Z_C$$

Since the dual complex of the disjoint union of two simplicial complexes (with planar rotation systems) is the disjoint union of their dual complexes, $k_C \leq k_D$. By Lemma 2.7.2 $k_C = k_D$. Plugging this into the equation before, proves the sublemma. $\square$

This completes the proof of the inequality. We have equality if and only if $r = Z_D$. So the 'Moreover'-part follows. $\square$

Our next goal is to prove the following, which is also reminiscent of euler's formula but here the inequality goes the other way round.

Lemma 2.7.5. *Let C be a locally connected simplicial complex with a planar rotation system Σ and D the dual complex. Then:*

$$|V(C)| - |E| + |F| - |V(D)| \leq 0$$

with equality if and only if all link complexes for (D, Σ_C) are spheres.

Before we can prove this, we need some preparation. By a we denote the sum of the faces of link complexes for (C, Σ). By a' we denote the sum over the faces of link complexes for (D, Σ_C). Before proving that a is equal to a' we prove that it is useful by showing the following.

Claim 2.7.6. *Lemma 2.7.5 is true if $a = a'$ and all link complexes for (D, Σ_C) are connected.*

Proof. Given a face f of C, we denote the number of edges incident with f by $deg(f)$. Our first aim is to prove that

$$2|V(C)| = 2|E| - \sum_{f \in F} \deg(f) + a \tag{2.1}$$

To prove this equation, we apply Euler's formula [36] in the link complexes for (C, Σ). Then we take the sum of all these equations over all $v \in V(C)$. Since Σ is a planar rotation system, all link complexes are a disjoint union of spheres. Since C is locally connected, all link complexes are connected and hence are spheres. So they have euler characteristic two. Thus we get the term $2|V(C)|$ on the left hand side. By definition, a is the sum of the faces of link complexes for (C, Σ).

The term $2|E|$ is the sum over all vertices of link complexes for (C, Σ). Indeed, each edge of C between the two vertices v and w of C is a vertex of precisely the two link complexes for v and w.

The term $\sum_{f \in F} \deg(f)$ is the sum over all edges of link complexes for (C, Σ). Indeed, each face f of C is in precisely those link complexes for vertices on the boundary of f. This completes the proof of (2.1).

Secondly, we prove the following inequality using a similar argument. Given an edge e of C, we denote the number of faces incident with e by $\deg(e)$.

$$2|V(D)| \geq 2|F| - \sum_{e \in E} \deg(e) + a' \tag{2.2}$$

To prove this, we apply Euler's formula in link complexes for (D, Σ_C), and take the sum over all $v \in V(D)$. Here we have '$\geq$' instead of '$=$' as we just know by assumption that the link complexes are connected but they may not be a sphere. So we have $2|V(D)|$ on the left and a' is the sum over the faces of link complexes for (D, Σ_C).

The term $2|F|$ is the sum over all vertices of link complexes for (D, Σ_C). Indeed, each edge of D between the two different vertices v and w of D is a vertex of precisely the two link complexes for v and w. A loop gives rise to two vertices in the link graph at the vertex it is attached to.

The term $\sum_{e \in E} \deg(e)$ is the sum over all edges of link complexes for (D, Σ_C). Indeed, each face e of D is in the link complex at v with multiplicity equal to the number of times it traverses v. This completes the proof of (2.2).

By assumption, $a = a'$. The sums $\sum_{f \in F} \deg(f)$ and $\sum_{e \in E} \deg(e)$ both count the number of nonzero entries of A, so they are equal. Subtracting (2.2) from (2.1), rearranging and dividing by 2 yields:

$$|V(C)| - |E| + |F| - |V(D)| \leq 0$$

with equality if and only if all link complexes for (D, Σ_C) are spheres. □

Hence our next aim is to prove that a is equal to a'. First we need some preparation.

Two cell complexes C and D are *(abstract) surface duals* if the set of vertices of C is (in bijection with) the set of faces of D, the set of edges of C is the set of edges of D and the set of faces of C is the set of vertices of D. And these three bijections preserve incidences.

Lemma 2.7.7. *Let C be a simplicial complex and Σ be a rotation system and let D be the dual. The surface dual of a local surface S for (C, Σ) is equal to the link complex for (D, Σ_C) at the vertex ℓ of D that corresponds to S.*

Proof. It is immediate from the definitions that the vertices of the link complex $\bar{L}$ at ℓ are the faces of S. The edges of S are triples $(e, \vec{f}, \vec{g})$, where e is an edge of C and $\vec{f}$ and $\vec{g}$ are orientations of faces of C that are related via e and are in the local-surface-equivalence class for S. Hence in D, these are triples $(e, \vec{f}, \vec{g})$ such that $\vec{f}$ and $\vec{g}$ are directions of edges that point towards ℓ and f and g are adjacent in the cyclic ordering of the face e. This are precisely the edges of the link graph $L(\ell)$. Hence the link graph $L(\ell)$ is the dual graph[11] of the cell complex S.

Now we will use the Edmonds-Hefter-Ringel rotation principle, see [71, Theorem 3.2.4], to deduce that the link complex $\bar{L}$ at ℓ is the surface dual of S. We denote the unique cell complex that is a surface dual of S by S^*. Above we have shown that $\bar{L}$ and S^* have the same 1-skeleton. Moreover, the rotation systems at the vertices of the link complex $\bar{L}$ are given by the cyclic orientations in the local-surface-equivalence class for S. By Observation 2.4.3 these local-surface-equivalence classes define an orientation of S. So $\bar{L}$ and S^* have the same rotation systems. Hence by the Edmonds-Hefter-Ringel rotation principle $\bar{L}$ and S^* have to be isomorphic. So $\bar{L}$ is a surface dual of S. $\square$

Proof of Lemma 2.7.5. Let C be a locally connected simplicial complex and Σ be a rotation system and let D be the dual. Let Σ_C be as defined above. By Observation 2.4.2 and Lemma 2.7.7 every link complex for (D, Σ_C) is connected. By Claim 2.7.6, it suffices to show that the sum over all faces of link complexes of C with respect to Σ is equal to the sum over all faces of link complexes for D with respect to Σ_C. By Lemma 2.7.7, the second sum is equal to the sum over all vertices of local surfaces for (C, Σ). This completes the proof by Lemma 2.4.5. $\square$

Proof of Theorem 2.7.1. Let C be a p-nullhomologous locally connected simplicial complex that has a planar rotation system Σ. Let D be the dual complex. Then by Lemma 2.7.5 and Lemma 2.7.3, C and D satisfy Euler's formula, that is:

$$|V(C)| - |E| + |F| - |V(D)| = 0$$

[11] The *dual graph* of a cell complex C is the graph G whose set of vertices is (in bijection with) the set of faces of C and whose set of edges is the set of edges of C. And the incidence relation between the vertices and edges of G is the same as the incidence relation between the faces and edges of C.

Hence by Lemma 2.7.5 all link complexes for (D, Σ_C) are spheres. By Lemma 2.7.7 these are dual to the local surfaces for (C, Σ). Hence all local surfaces for (C, Σ) are spheres. $\square$

The following theorem gives three equivalent characterisations of the class of locally connected simply connected simplicial complexes embeddable in $\mathbb{S}^3$.

Theorem 2.7.8. *Let C be a locally connected simplicial complex embedded into $\mathbb{S}^3$. The following are equivalent.*

1. *C is simply connected;*
2. *C is p-nullhomologous for some prime p;*
3. *all local surfaces of the planar rotation system induced by the topological embedding are spheres.*

Proof. Clearly, 1 implies 2. To see that 2 implies 3, we assume that C is p-nullhomologous. Let Σ be the planar rotation system induced by the topological embedding of C into $\mathbb{S}^3$. By Theorem 2.7.1 all local surfaces for (C, Σ) are spheres.

It remains to prove that 3 implies 1. So assume that C has an embedding into $\mathbb{S}^3$ such that all local surfaces of the planar rotation system induced by the topological embedding are spheres. By treating different connected components separately, we may assume that C is connected. By Lemma 2.4.4 all local surfaces of the topological embedding are spheres. Thus 3 implies 1 by Lemma 2.5.3.

$\square$

Remark 2.7.9. Our proof actually proves the strengthening of Theorem 2.7.8 with 'embedded into $\mathbb{S}^3$' replaced by 'embedded into a simply connected 3-dimensional compact manifold.' However this strengthening is equivalent to Theorem 2.7.8 by Perelman's theorem.

Recall that in order to prove Theorem 2.2.2, it suffices to show that every p-nullhomologous simplicial complex C has a piece-wise linear embedding into $\mathbb{S}^3$ if and only if it is simply connected and C has a planar rotation system.

Proof of Theorem 2.2.2. Using an induction argument on the number of cut vertices as in the proof of Theorem 2.2.1, we may assume that C is locally connected. If C has a piece-wise linear embedding into $\mathbb{S}^3$, then it has a planar rotation system and it is simply connected by Theorem 2.7.8. The other direction follows from Theorem 2.2.1. $\square$

Remark 2.7.10. One step in proving Theorem 2.2.2 was showing that if a simplicial complex whose first homology group is trivial embeds in $\mathbb{S}^3$, then it must be simply connected. In this section we have given a proof that only uses elementary topology. We use these methods again in Chapter 4.

However there is a shorter proof of this fact, which we shall sketch in the following. Let C be a simplicial complex embedded in $\mathbb{S}^3$ such that one local surface of the embedding is not a sphere. Our aim is to show that the first homology group of C cannot be trivial.

We will rely on the fact that the first homology group of $X = \mathbb{S}^3 \setminus \mathbb{S}^1$ is not trivial. It suffices to show that the homology group of X is a quotient of the homology group of C. Since here by Hurewicz's theorem, the homology group is the abelisation of the fundamental group, it suffices to show that the fundamental group $\pi_1(X)$ of X is a quotient of the fundamental group $\pi_1(C)$.

We let C_1 be a small open neighbourhood of C in the embedding of C in $\mathbb{S}^3$. Since C_1 has a deformation retract onto C, it has the same fundamental group. We obtain C_2 from C_1 by attaching the interiors of all local surfaces of the embedding except for one – which is not a sphere. This can be done by attaching finitely many 3-balls. Similar as in the proof of Lemma 2.5.2, one can use Van Kampen's theorem to show that the fundamental group of C_2 is a quotient of the fundamental group of C_1. By adding finitely many spheres if necessary and arguing as above one may assume that remaining local surface is a torus. Hence C_2 has the same fundamental group as X. This completes the sketch.

2.8 Embedding general simplicial complexes

There are three classes of simplicial complexes that naturally include the simply connected simplicial complexes: the p-nullhomologous ones that are included in those with abelian fundamental group that in turn are included in general simplicial complexes. Theorem 2.2.2 characterises embeddability of p-nullhomologous complexes. In this section we prove embedding results for the later two classes. The bigger the class gets, the stronger assumptions we will require in order to guarantee topological embeddings into $\mathbb{S}^3$.

A *curve system* of a surface S of genus g is a choice of at most g genus reducing curves in S that are disjoint. An *extension* of a rotation system Σ is a choice of curve system at every local surface of Σ. An extension of a rotation system of a complex C is *simply connected* if the topological space obtained from C by gluing[12] a disc at each curve of the extension is simply connected. The definition of a *p-nullhomologous extension* is the same with 'p-nullhomologous' in place of 'simply connected'.

Theorem 2.8.1. *Let C be a connected and locally connected simplicial complex with a rotation system Σ. The following are equivalent.*

1. *Σ is induced by a topological embedding of C into $\mathbb{S}^3$.*
2. *Σ is a planar rotation system that has a simply connected extension.*

[12] We stress that the curves need not go through edges of C. 'Gluing' here is on the level of topological spaces not of complexes.

3. *We can subdivide edges of C, do baricentric subdivision of faces and add new faces such that the resulting simplicial complex is simply connected and has a topological embedding into $\mathbb{S}^3$ whose induced planar rotation system Σ' 'induces' Σ.*

Here we define that 'Σ' induces Σ' in the obvious way as follows. Let C be a simplicial complex obtained from a simplicial complex C' by deleting faces. A rotation system $\Sigma = (\sigma(e)|e \in E(C))$ of C is *induced* by a rotation system $\Sigma' = (\sigma'(e)|e \in E(C))$ of C' if $\sigma(e)$ is the restriction of $\sigma'(e)$ to the faces incident with e. If C is obtained from contracting edges of C' instead, a rotation system Σ of C is *induced* by a rotation system Σ' of C' if Σ is the restriction of Σ' to those edges that are in C. If C' is obtained from C by a baricentric subdivision of a face f we take the same definition of 'induced', where we make the identification between the face f of C and all faces of C' obtained by subdividing f. Now in the situation of Theorem 2.8.1, we say that Σ' *induces* Σ if there is a chain of planar rotation systems each inducing the next one starting with Σ' and ending with Σ.

Before we can prove Theorem 2.8.1, we need some preparation. The following is a consequence of the Loop Theorem [78, 55].

Lemma 2.8.2. *Let X be an orientable surface of genus $g \geq 1$ embedded topologically into $\mathbb{R}^3$, then there is a genus reducing circle*[13] *γ of X and a disc D with boundary γ and all interior points of D are contained in the interior of X.*

Corollary 2.8.3. *Let X be an orientable surface of genus $g \geq 1$ embedded topologically into $\mathbb{R}^3$, then there are genus reducing circles γ_1,..., γ_g of X and closed discs D_i with boundary γ_i such that the D_i are disjoint and the interior points of the discs D_i are contained in the interior of X.*

Proof. We prove this by induction on g. In the induction step we cut of the current surface along D. Then we the apply Lemma 2.8.2 to that new surface. □

Proof of Theorem 2.8.1. 1 is immediately implied by 3.

Next assume that Σ is induced by a topological embedding of C into $\mathbb{S}^3$. Then Σ is clearly a planar rotation system. It has a simply connected extension by Corollary 2.8.3. Hence 1 implies 2.

Next assume that Σ is a planar rotation system that has a simply connected extension. We can clearly subdivide edges and do baricentric subdivision and change the curves of the curve system of the simply connected extension such that in the resulting simplicial complex C' all the curves of the simply connected extension closed are walks in the 1-skeleton of C'. We define a planar rotation system Σ' of C' that induces Σ as follows. If we subdivide an edge, we assign to both copies the cyclic orientation of the original edge. If we do a baricentric subdivision, we assign to all new edges the unique cyclic orientation of size two. Iterating this during the construction of C' defines $\Sigma' = (\sigma'(e)|e \in E(C))$, which

[13]A *circle* is a topological space homeomorphic to $\mathbb{S}^1$.

clearly is a planar rotation system that induces Σ. By construction Σ' has a simply connected extension such that all its curves are walks in the 1-skeleton of C' .

Informally, we obtain C'' from C' by attaching a disc at the boundary of each curve of the simply connected extension. Formally, we obtain C'' from C' by first adding a face for each curve γ in the simply connected extension whose boundary is the closed walk γ. Then we do a baricentric subdivision to all these newly added faces. This ensures that C'' is a simplicial complex. Since C is locally connected, also C'' is locally connected. Since the geometric realisation of C'' is equal to the geometric realisation of C, which is simply connected, the simplicial complex C'' is simply connected.

Each newly added face f corresponds to a traversal of a curve γ of some edge e of C'. This traversal is a unique edge of the local surface S to whose curve system γ belongs. For later reference we denote that copy of e in S by e_f.

We define a rotation system $\Sigma'' = (\sigma''(e)|e \in E(C))$ of C'' as follows. All edges of C'' that are not edges of C' are incident with precisely two faces. We take the unique cyclic ordering of size two there.

Next we define $\sigma''(e)$ at edges e of C' that are incident with newly added faces. If e is only incident with a single face of C', then e is only in a single local surface and it only has one copy in that local surface. Since the curves at that local surface are disjoint. We could have only added a single face incident with e. We take for $\sigma''(e)$ the unique cyclic orientation of size two at e.

So from now assume that e is incident with at least two faces of C'. In order to define $\sigma''(e)$, we start with $\sigma'(e)$ and define in the following for each newly added face in between which two cyclic orientations of faces adjacent in $\sigma'(e)$ we put it. We shall ensure that between any two orientations we put at most one new face. Recall that two cyclic orientations $\vec{f_1}$ and $\vec{f_2}$ of faces f_1 and f_2, respectively, are adjacent in $\sigma'(e)$ if and only if there is a clone e' of e in a local surface S for (C', Σ') containing $\vec{f_1}$ and $\vec{f_2}$ such that e' is incident with $\vec{f_1}$ and $\vec{f_2}$ in S. Let f be a face newly added to C'' at e. Let γ_f be the curve from which f is build and let S_f be the local surface that has γ_f in its curve system. Let e_f be the copy of e in S_f that corresponds to f as defined above. when we consider f has a face obtained from the disc glued at γ_f. We add f to $\sigma'(e)$ in between the two cyclic orientations that are incident with e_f in S_f. This completes the definition of Σ''. Since the copies e_f are distinct for different faces f, the rotation system Σ'' is well-defined. By construction Σ'' induces Σ. We prove the following.

Sublemma 2.8.4. *Σ'' is a planar rotation system of C''.*

Proof. Let v be a vertex of C''. If v is not a vertex of C', then the link graph at v is a cycle. Hence the link complex at v is clearly a sphere. Hence we may assume that v is a vertex of C'.

Our strategy to show that the link complex S'' at v for (C'', Σ'') is a sphere will be to show that it is obtained from the link complex S' for (C', Σ') by adding edges in such a way that each newly added edge traverses a face of S' and two newly added edges traverse different face of S'.

So let f be a newly added face incident with v of C'. Let x and y be the two edges of f incident with v. We make use of the notations γ_f, S_f, x_f and y_f defined above. Let v_f be the unique vertex of S_f traversed by γ_f in between x_f and y_f. By Lemma 2.4.5 there is a unique face z_f of S' mapped by the map ι of that lemma to v_f. And x and y are vertices in the boundary of z_f. The edges on the boundary of z_f incident with x and y are the cyclic orientations of the faces that are incident with x_f and y_f in S_f. Hence in S'' the edge f traverses the face z_f.

It remains to show that the faces z_f of S' are distinct for different newly added faces f of C''. For that it suffices by Lemma 2.4.5 to show that the vertices v_f are distinct. This is true as curves for S_f traverse a vertex of S_f at most once and different curves for S_f are disjoint.

□

Since Σ'' is a planar rotation system of the locally connected simplicial complex C'' and C'' is simply connected, Σ'' is induced by a topological embedding of C'' into $\mathbb{S}^3$ by Theorem 2.5.4. Hence 2 implies 3. □

A natural weakening of the property that C is simply connected is that the fundamental group of C is abelian. Note that this is equivalent to the condition that every chain that is p-nullhomologous is simply connected.

Theorem 2.8.5. *Let C be a connected and locally connected simplicial complex with abelian fundamental group. Then C has a topological embedding into $\mathbb{S}^3$ if and only if it has a planar rotation system Σ that has a p-nullhomologous extension.*

In order to prove Theorem 2.8.5, we prove the following.

Lemma 2.8.6. *A p-nullhomologous extension of a planar rotation system of a simplicial complex C with abelian fundamental group is a simply connected extension.*

Proof. Let C' be the topological space obtained from C by gluing discs along the curves of the p-nullhomologous extension. The fundamental group π' of C' is a quotient of the fundamental group π of C, see for example [56, Proposition 1.26]. Since π is abelian by assumption, also π' is abelian. That is, it is equal to its abelisation, which is trivial by assumption. Hence C' is simply connected. □

Proof of Theorem 2.8.5. If C has a topological embedding into $\mathbb{S}^3$, then by Theorem 2.8.1 it has a planar rotation system that has a p-nullhomologous extension. If C has a planar rotation system that has a p-nullhomologous extension, then that extension is simply connected by Lemma 2.8.6. Hence C has a topological embedding into $\mathbb{S}^3$ by the other implication of Theorem 2.8.1. □

Chapter 3

Constraint minors

3.1 Abstract

We characterise the following property by six obstructions: given a graphic matroid M and a set X of its elements, when is M the cycle matroid of a graph G such that X is a connected edge set in G?

3.2 Introduction

For a purely graph-theoretic introduction read Section 3.3.

Tutte [93] proved that a matroid can be represented by a graph if and only if it has no minor isomorphic to $U_{2,4}$, the fano-plane, the dual fano-plane or the dual matroids of the two nonplanar graphs K_5 or $K_{3,3}$. The topic of this chapter is the following related reconstruction question: given a graphic matroid M and a set X of its elements, when is M the cycle matroid of a graph G such that X is a connected edge set in G? Our motivation for studying that question is that in Chapter 4 it arises when characterising embeddability in 3-space of certain 2-complexes by excluded minors.

A *constraint matroid* is a pair (M, X), where M is a matroid and X is a set of elements of M. A constraint matroid (M, X) is *realisable* if M is the cycle matroid of a graph G such that X is a connected edge set in G. The class of constraint matroids (M, X) that are realisable is closed under contracting arbitrary elements and deleting elements not in X. A constraint matroid obtained by these operations from (M, X) is a *constraint minor* of (M, X). In this chapter we characterise the class of the realisable (graphic) constraint matroids by excluded constraint minors.

Theorem 3.2.1. *A graphic constraint matroid is realisable if and only if it does not have one of the six constraint minors depicted in Figure 3.1, Figure 3.2 or Figure 3.3.*

All these six obstructions are 3-connected and graphic. So we just depict their unique graphs. Theorem 3.2.1 can be restated in purely graph theoretic terms, see Theorem 3.3.1 below.

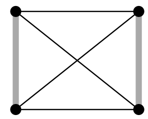

Figure 3.1: The constraint K_4. The edges in X are depicted grey.

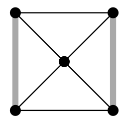

Figure 3.2: The constraint wheel. The edges in X are depicted grey.

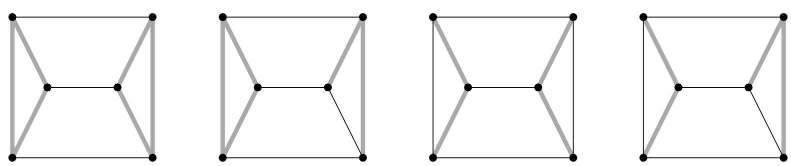

Figure 3.3: The four constraint prisms. The edges in X are depicted grey.

3.3 A graph theoretic perspective

Although Theorem 3.2.1 is about matroids, most of this chapter is about the following equivalent graph theoretic version.

A *constraint graph* is a pair (G, X), where G is a graph and X is an edge set of G. A constraint graph is *constraint connected* if X is a connected edge set in G. The class of constraint graphs (G, X) that are constraint connected is closed under contracting arbitrary edges and deleting edges not in X. A constraint graph obtained by these operations from (G, X) is a *constraint minor* of (G, X). It is straightforward to show that a 2-connected[1] constraint graph (G, X) is constraint connected if and only if it has no constraint minor isomorphic to the 4-cycle whose constraint consists of two opposite edges. The analogue question for connected graphs is not much more interesting.

However, it turns out that the question gets nontrivial if we restrict our attention to 3-connected graphs.

[1]A constraint graph (G, X) is *k-connected* if G is k-connected.

Theorem 3.3.1. *A 3-connected constraint graph (G, X) is constraint connected if and only if it does not have one of the six (3-connected) constraint minors depicted in Figure 3.1, Figure 3.2 or Figure 3.3.*

It is straightforward to deduce Theorem 3.3.1 from Theorem 3.2.1 above. However the converse is also true as follows.

Proof that Theorem 3.3.1 implies Theorem 3.2.1. Let (M, X) be a constraint matroid. If M is 3-connected, then it is the cycle matroid of a unique graph G by a theorem of Whitney [100]. In this case Theorem 3.2.1 for (M, X) is a restatement of Theorem 3.3.1 for (G, X).

Now let (M, X) be a constraint matroid that has no constraint minor depicted in Figure 3.1, Figure 3.2 or Figure 3.3. It remains to show that (M, X) is realisable. Since a constraint matroid is realisable if and only if each of its 2-connected components is, we may assume that M is 2-connected.

Now we prove by induction that (M, X) is realisable. The base case is that M is 3-connected.

If M is not 3-connected, its Tutte-decomposition [94] has a non-trivial 2-separation (A, B). Let M_1 and M_2 be the two matroids obtained by decomposing M along the 2-separation (A, B). In particular, M_1 and M_2 both contain a virtual element e and the 2-sum[2] of M_1 and M_2 along e is M. Note that the M_i can be obtained from M by contracting elements and replacing a parallel class by the virtual element e. For $i = 1, 2$, let (M_i, X_i) be the constraint matroid, where X_i is $X \cap E(M_i)$ plus possibly e if M_{i+1} contains a circuit o such that $o - e \subseteq X$. It is straightforward to check that the (M_i, X_i) are constraint minors of M. Hence by induction, they are realisable. Let G_i be a graph realising (M_i, X_i).

Let G be the 2-sum of the graphs G_1 and G_2 along the virtual element e. By construction M is the cycle matroid of G. If the virtual element e is in X_1 or X_2, it is straightforward to see that (G, X) is constraint connected. So M is realisable. So we may assume that e is in no X_i. If one of the X_i is empty, then (G, X) is constraint connected. So we may assume that both X_i are nonempty.

Then not only (M_1, X_1) but also $(M_1, X_1 + e)$ is a constraint minor of (M, X). So by induction there is a graph G'_1 realising $(M_1, X_1 + e)$. In G'_1 an element of the set X_1 is incident with an endvertex of e. Similarly, there is a graph G'_2 realising $(M_2, X_2 + e)$, and there is an element of the set X_2 is incident with an endvertex of e. Let G' be the 2-sum of the graphs G'_1 and G'_2. By flipping[3] the 2-separator given by the endvertices of e in G' if necessary, we ensure that X is connected in G'. Put another way, (G', X) is constraint connected witnessing that (M, X) is realisable. □

Hence the rest of this chapter is dedicated to the proof of Theorem 3.3.1, which is purely graph-theoretic. Before jumping into the proof, let us fix a few lines of notation. In this chapter all graphs are simple. In particular, if we

[2] See [76] for a definition.

[3] By a theorem of Whitney, graphs represented by a 2-connected matroid are unique up to flipping 2-separators [100].

contract an edge, we afterwards delete all but one edge from every parallel class. In the context of a constraint graph (G, X), we first delete edges in a parallel classes that are not in X (so that constraint minors on simple graphs preserve constraint connectedness). Throughout this chapter we follow the convention that the empty set is a connected edge set in G. Beyond that we follow the notation of [36]. Let's get started with the proof.

3.4 Deleting and contracting edges outside the constraint

In this section we prove Lemma 3.4.9 below, which is used in the proof of Theorem 3.3.1.

Given a constraint graph (G, X), an edge e not in X is *essential* if neither $(G/e, X)$ nor $(G\backslash e, X)$ has a 3-connected constraint minor (G', X') such that X' is disconnected. Informally, Lemma 3.4.9 below gives a structural description of the constraint graphs (G, X) in which every edge not in X is essential.

Before we can prove Lemma 3.4.9 we need some preparation. Our first aim is to prove the following.

Lemma 3.4.1. *Let (G, X) be a 3-connected constraint graph that is not constraint connected. Assume that every edge not in X is essential. Then $G[X]$ has precisely two connected components or (G, X) is the weird prism (defined in Example 3.4.2).*

First we consider some particular examples that will come up in the proof of Lemma 3.4.1.

Example 3.4.2. The *weird prism* is the pair (P, X), where P is the prism and X consists of the three edges in the complement of the two triangles, see Figure 3.4. Contracting any particular edge in X, gives the constraint wheel.

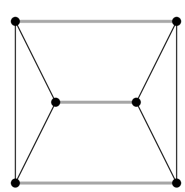

Figure 3.4: The weird prism. The edges in X are depicted grey.

Example 3.4.3. The *constraint Wagner graph* is the pair (W, X), where W is the Wagner graph and X is the set of edges in the complement of one of its six-cycles, see Figure 3.5. If we contract a single edge of X, we get the constraint wheel. If we contract any two opposite edges on the six cycle, then we get a constraint K_4.

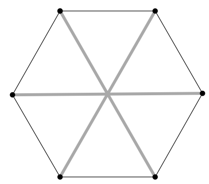

Figure 3.5: The constraint Wagner graph. The edges in X are depicted grey.

Example 3.4.4. The *Wagner prism* is the pair (W', X'), where W' is the prism and X' contains one edge not in the two triangles of the prism. The two other edges in X' are the only two edges of the prism in the triangles that are vertex-disjoint to that edges, see Figure 3.6. There are two opposite edges on the six cycle formed by the edges not in X' whose contraction gives the constraint K_4.

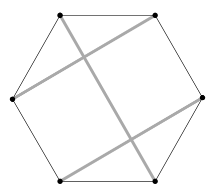

Figure 3.6: The Wagner prism. The edges in X are depicted grey.

Lemma 3.4.5. *Let (G, X) be a 3-connected constraint graph such that $G[X]$ has at least 3-connected components. Assume that G is not the constraint Wagner graph, not the weird prism and not the Wagner prism. Then there is a 3-connected constraint minor (G', X') of (G, X) such that X' is disconnected in G' and such that $E(G') \setminus X'$ is a proper subset of $E(G) \setminus X$.*

Proof that Lemma 3.4.5 implies Lemma 3.4.1. By Example 3.4.3, the constraint Wagner graph has an edge not in X that is not essential. Thus (G, X) is not the constraint Wagner graph. Similarly, (G, X) is not the Wagner prism by Example 3.4.4. Hence by Lemma 3.4.5, $G[X]$ has precisely two connected components or is the weird prism. □

Proof of Lemma 3.4.5. Let e be an arbitrary edge not in X. If the simple graph $G' = G/e$ is 3-connected, then $(G', X \cap E(G'))$ is the desired constraint minor. Otherwise by Bixby's Lemma [76] the graph $G \backslash e$ is 3-connected after suppressing edges of degree 2; note that G cannot be K_4 as the disconnected set X contains at least three edges. Let G' be the graph obtained from $G \setminus e$ by contracting all but one edge from every serial class.

By construction, any vertex of degree 2 of $G \setminus e$ must be an endvertex of e. Hence every nontrivial serial class has size two and there are at most two of them. If a serial class contains an edge in X and an edge not in X, we contract the edge of X in the construction of G'. This construction ensures that we never contract all edges of a path in G that connects two components of $G[X]$. We let $X' = X \cap E(G')$. Hence the components of $G'[X']$ come from those components of $G[X]$ such that not all their edges got contracted.

Thus $G'[X']$ is disconnected unless $G[X]$ has precisely three components and two of these components just consist of a single edge. Furthermore both endvertices of e have degree 3 and each of them is incident with one of these components consisting of a single edge. In this case we say that the edge e *H-shaped*.

Since the edge e was arbitrary, we find the desired constraint graph (G', X') unless every edge of G not in X is H-shaped. Since G is connected, every component of $G[X]$ is incident with an edge not in X. Hence $G[X]$ has precisely three components and they all consist of single edges. Furthermore every vertex of G is incident with one edge in X and two edges not in X. Thus G has precisely six vertices. The edges not in X form a vertex-disjoint union of cycles. So as G is a simple graph, they either form two vertex-disjoint triangles or a 6-cycle. In the first case it is straightforward to check that (G, X) is the weird prism. In the second case it is straightforward to check that (G, X) is isomorphic to the constraint Wagner graph or the Wagner prism. □

This completes the proof of Lemma 3.4.1. Our next step is to prove the following.

Lemma 3.4.6. *Let G be a 3-connected graph and let X be an edge set of G such that $G[X]$ has precisely two components. Let $e \in E(G) \setminus X$ be essential. Then one of the following holds.*

1. *e joins the two components of $G[X]$; or*
2. *there is a component C of $G[X]$ that consists only of a single edge and e has an endvertex v of degree three that is incident with that edge and the third edge incident with v joins the two components of $G[X]$; or*
3. *there is a component C of $G[X]$ that consists of precisely two edges, which form a triangle together with e. The two endvertices of e have degree 3 and are each incident with an edge that joins the two components of $G[X]$.*

Proof. We assume that e does not join the two components of $G[X]$, in particular G is not K_4. If the simple graph $G' = G/e$ is 3-connected, then $(G', X \cap E(G'))$ is a 3-connected constraint minor such that $X \cap E(G')$ is disconnected. Since e is essential this is impossible. Hence by Bixby's Lemma [76] the graph $G \setminus e$ is 3-connected after suppressing edges of degree 2. Let G' be the graph obtained from $G \setminus e$ by contracting all but one edge from every serial class.

By construction, any vertex of degree 2 of $G \setminus e$ must be an endvertex of e. Hence every nontrivial serial class has size two and there are at most two of them. If a serial class contains an edge in X and an edge not in X, we contract the edge of X in the construction of G'. This construction ensures that we never contract all edges of a path in G that connects two components of $G[X]$. We let $X' = X \cap E(G')$. Hence the components of $G'[X']$ come from those components of $G[X]$ such that not all their edges got contracted. Since G' is 3-connected and e is essential, the graph $G'[X']$ is connected.

Hence there must be a component C of $G[X]$ such that all its edges got contracted. Hence C has at most two edges. We split into two cases.

Case 1: C has only a single edge f. Then e has an endvertex v of degree 3 that is incident with f. In this case we shall show that we have outcome 2; that is, the third edge g incident with v joins the two components of $G[X]$. Indeed, we construct G'' like G' but instead of f we contract g. Since G'' is isomorphic to G', it is 3-connected. As e is essential, it must be that $G''[X'+f]$ is connected. Since the component of $G[X]$ different from C does not contain a vertex incident with e, the edge g joins the two components of $G[X]$.

Case 2: C has two edges f_1 and f_2. Then e has two endvertices v_1 and v_2 of degree three such that v_i is incident with f_i. Since G is a simple graph and C is connected, the three edges e, f_1 and f_2 form a triangle. Similar as in Case 1 we prove for each i that the third edge incident with v_i joins the two components of $G[X]$. So we have outcome 3 in this case. □

The following lemma deals with outcome 2 of Lemma 3.4.6.

Lemma 3.4.7. *Let G be a 3-connected graph and X a disconnected edge set of G. Assume that every edge not in X is essential. Assume that a component C of $G[X]$ consists only of a single edge and that there is an edge vw such that v is a vertex of C and w is not in $G[X]$. Then (G, X) is the constraint wheel.*

Proof. The constraint graph (G, X) is not the weird prism; indeed the weird prisms has no edge vw as required in the assumptions. Hence by Lemma 3.4.1, $G[X]$ has only one connected component C' aside from C. The endvertex w of e that is not in C is not incident with any edge of X. Since G is 3-connected, w is incident with at least two edges f_1 and f_2 aside from e. By Lemma 3.4.6 the endvertex of each f_i different from w must be in C or C'. Since C has only one vertex aside from v, one of the f_i must have an endvertex in C'. By symmetry, we may assume that this is true for f_1. Since f_1 has an endvertex that is in neither C nor C', we can apply Lemma 3.4.6 to deduce that C' also consists of a single edge.

Sublemma 3.4.8. *The vertex set of G is $(C \cup C') + w$.*

Proof. By Lemma 3.4.6, each vertex of $C \cup C'$ that has a neighbour outside that set has degree three and at most one neighbour outside that set. Let W be the set of vertices of $C \cup C'$ that have a neighbour outside the set $(C \cup C') + w$. Since w has at least three neighbours in $C \cup C'$, the set W contains at most one vertex. The set W together with w separates G if there are vertices not in $(C \cup C') + w$. Since G is 3-connected, this is not true. Hence $(C \cup C') + w$ is the vertex set of G. □

Since w is adjacent to at least three vertices in $C \cup C'$, at least three vertices of $C \cup C'$ have precisely two neighbours in $C \cup C'$. Hence the graph $G[C \cup C']$ is a 4-cycle. Since G is 3-connected, each of its vertices has degree at least three. Hence by 3-connectivity every vertex of $C \cup C'$ is adjacent to w. Thus G is the constraint wheel. $\square$

Given an edge set Z, by $V(Z)$ we denote the set of endvertices of edges in Z. Summing up, we have the following.

Lemma 3.4.9. *Let (G, X) be a 3-connected constraint graph such that X is disconnected. Assume that every edge not in X is essential and that (G, X) is neither the constraint wheel nor the weird prism. Then $G[X]$ has precisely two connected components C_1 and C_2. All edges not in X have both their endvertices in $V(X)$.*

Proof. By assumption and by Lemma 3.4.1, $G[X]$ has precisely two connected components, C_1 and C_2. By Lemma 3.4.6 and Lemma 3.4.7, every edge not in X has both its endvertices in $V(X)$.

$\square$

3.5 Contracting edges in the constraint

In this section we prove Theorem 3.3.1.

First we need some preparation. Given a bond d in a graph G, then $G - d$ has two connected components which we call the *sides of* d. If we want to specify them, we call them the *left side* and the *right side*.

Given a graph G and a bond d of G, we say that G is *3-connected along* d if G is 2-connected and there does not exist a separator consisting of two vertices from either side of d.

For the rest of this section we fix a graph Q and a bond d of Q so that Q is 3-connected along d. We denote the set of edges on the left side of d by L, and the set of edges on the right side of d by R. We assume throughout that L and R are nonempty. A *special contraction minor* of (Q, d) is a pair (Q', d'), where Q' is obtained from Q by contracting edges not in d, and $d' = d \cap E(Q')$. Note that d' and d need not be equal as contractions might force us to delete edges in parallel classes. Since any parallel class containing one edge of d is a subset of d, the set d' is independent of the choice of the deleted edges.

Example 3.5.1. The following pairs (Q, d) will be of particular interest in this chapter. For any two bonds of K_4 with both sides nonempty, there is an isomorphism of K_4 that induces a bijection between these two bonds. The *special* K_4 is the pair consisting of the graph K_4 and a bond of size 4. The *special prism* is the pair consisting of the prism and a bond whose complement consists of the two triangles of the prism, see Figure 3.7.

Our aim in this section is to prove the following.

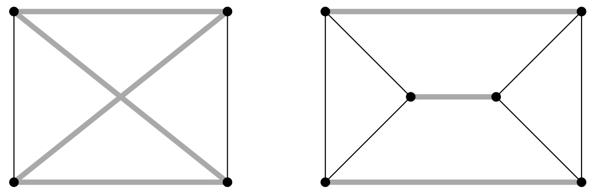

(a) the special K_4 (b) the special prism

Figure 3.7: The edges in the bond d are coloured grey.

Lemma 3.5.2. *Let Q be a graph 3-connected along a bond d such that the two sides of d contain edges. Then (Q, d) has a special contraction minor that is the special K_4 or the special prism.*

Proof that Lemma 3.5.2 implies Theorem 3.3.1. Let (G, X) be a 3-connected constraint graph such that X is disconnected. Our aim is to show that (G, X) has the constraint K_4, the constraint wheel or a constraint prism as a constraint minor. By picking (G, X) minimal, we may assume that every edge not in X is essential. By Example 3.4.2 we may assume that (G, X) is not the weird prism. We may also assume that it is not the constraint wheel. Thus by Lemma 3.4.9, $G[X]$ has precisely two connected components C_1 and C_2. And all edges not in X have both their endvertices in $V(X)$. We take $Q = G$ and d to be the bond consisting of those edges with one endvertex in C_1 and the other in C_2. Note that each C_i contains at least one edge. Since Q is 3-connected, (Q, d) is 3-connected along d.

By Lemma 3.5.2, (Q, d) has a special contraction minor (Q', d') that is the special K_4 or the special prism. Put another way, we can contract edges not in d such that G is K_4 or the prism. Let $X' = X \cap E(Q')$. We recall that if contractions force us to delete edges from a parallel class we first delete edges not in X. Hence since X spans the two sides of d in G, also X' spans the two sides of d' in Q'. Thus if (Q', d') is a special K_4, then (G, X) has the constraint K_4 as a constraint minor. Otherwise (Q', d') is the special prism. It is straightforward to check that in this case (G, X) has a constraint prism as a constraint minor. □

The rest of this section is dedicated to the proof of Lemma 3.5.2. A pair (Q, d) is *irreducible* if Q is 3-connected along d but there does not exist a proper[4] special contraction minor (Q', d') such that both sides of d' contain edges and Q' is 3-connected along d'. The first step in the proof of Lemma 3.5.2 will be to show that the set of irreducible (Q, d) is bounded. Later we examine this bounded set.

Given an edge set Z of Q, by $Q[Z]$ we denote the subgraph of Q whose vertices are those with at least one endvertex in Z and whose edges are those in Z.

[4]non-identical

Lemma 3.5.3. *If the graph $Q[L]$ is not 2-connected and has at least two edges, then (Q, d) is not irreducible.*

Proof. We consider the block-cutvertex-tree of $Q[L]$ and take a leaf block b. Recall that b is a 2-connected subgraph of $Q[L]$ or a single edge attached at a cutvertex $v \in b$ to the rest of $Q[L]$. We obtain Q_1 from Q by contracting all edges of $Q[L]$ not in b. Since by assumption there is an edge in $Q[L]$ that is not in b, Q_1 is a nontrivial contraction of Q.

Next we consider the block-cutvertex-tree of $Q[R]$. Note that unlike that for $Q[L]$ this may consist of just a single node. We obtain Q_2 from Q_1 by successively contracting leaf blocks b' attached at a cutvertex v' onto v' if there is no edge between $b - v$ and $b' - v'$.

In a slight abuse of notation we denote the contraction vertex of Q_2 containing v by v. Similarly after contracting a leaf part on the right side, we denote the contraction vertex containing v' by v'. We let $d_2 = d \cap E(Q_2)$. We denote the edges on the left of d_2 by L_2 and the edges on the right of d_2 by R_2.

Our aim is to show that Q_2 is 3-connected along d_2. By construction L_2 is nonempty.

Sublemma 3.5.4. *The edge set R_2 is nonempty.*

Proof. In the construction of Q we only contract a leaf block b' on the right side attached with cutvertex v' if there is no edge between $b - v$ and $b' - v'$. In particular by contraction we never identify two vertices of $Q[R]$ that have neighbours in $b - v$.

If there was only a single vertex z in $Q[R]$ that has a neighbour in $b - v$, then $Q - v - z$ would be disconnected, contrary to our assumption that Q is 3-connected along d. Hence there are at least two vertices in $Q[R]$ that have neighbours in $b - v$. Thus as explained above, the connected graph $Q_2[R_2]$ contains at least two vertices. Hence R_2 contains an edge. $\square$

Sublemma 3.5.5. *The graph Q_2 is 2-connected.*

Proof. Let x be an arbitrary vertex of Q_2. We distinguish two cases.

Case 1: $x = v$.

By Sublemma 3.5.4, the connected graph $Q_2[R_2]$ has a neighbour in the connected set $b - v$. Hence $Q_2 - x$ is connected.

Case 2: $x \neq v$. If x is not a contraction vertex, then $Q_2 - x$ is connected as $Q - x$ is connected. So x is a vertex of $Q_2[R_2]$. Let K be a component of the graph $Q_2[R_2] - x$. Let K' be the component of $Q[R] - x$ containing K. Since $Q - x$ is connected, there is an edge from K' to $Q[L]$. Hence there is an edge from K to b in Q_2. Hence every component of the graph $Q_2[R_2] - x$ sends an edge to the connected set b. Hence $Q_2 - x$ is connected. $\square$

Sublemma 3.5.6. *For any two vertices $x \in Q_2[L_2]$ and $y \in Q_2[R_2]$ the graph $Q_2 - x - y$ is connected.*

Proof. We distinguish two cases.

Case 1: $x = v$.

Let K be a component of the graph $Q_2[R_2] - y$. Since K did not get contracted, it has a neighbour in $b - v$. Thus every component of $Q_2[R_2] - y$ has a neighbour in the connected set $b - v$. Hence $Q_2 - x - y$ is connected.

Case 2: $x \neq v$.

Let K be a component of the graph $Q_2[R_2] - y$. Let K' be the component of $Q[R] - y$ containing K. Since $Q - x - y$ is connected, there is an edge from K' to $Q[L] - x$. Hence there is an edge from K to $b - x$ in Q_2. Hence every component of the graph $Q_2[R_2] - y$ sends an edge to the connected set $b - x$. Hence $Q_2 - x - y$ is connected. □

By Sublemma 3.5.5 and Sublemma 3.5.6, Q_2 is 3-connected along d_2. By construction Q_2 is obtained from Q by contracting at least one edge. By Sublemma 3.5.4, the edge sets L_2 and R_2 are nonempty. Hence (Q_2, d_2) witnesses that (Q, d) is not irreducible.

□

Lemma 3.5.7. *If the graph $Q[L]$ is 2-connected but not a triangle and the graph $Q[R]$ is 2-connected or consists of a single edge, then (Q, d) is not irreducible.*

In the proof of Lemma 3.5.7 we shall use the following lemma. An edge e in a 2-connected graph G is *contractible* if G/e is 2-connected.

Lemma 3.5.8. *If G is a 2-connected graph that is not a triangle, then it has four contractible edges, two of which do not share an endvertex.*

Proof. If G is 3-connected or a cycle of length at least 4, every edge is contractible and the lemma is true in this case. Hence the Tutte-decomposition [94] of G has at least two leaf parts. The torsos of these parts are cycles or 3-connected. Let v be a vertex in a leaf part that is not in the separator. Then any edge incident with v is contractible. Since there are at least two leaf parts, we can pick vertices v in one of each. Each such vertex is incident with at least two edges and no edge is incident with both these vertices. So there are at least four contractible edges, and there are two of them that do not share an endvertex. □

Proof of Lemma 3.5.7. Suppose for a contradiction that (Q, d) is irreducible. Let vw be a contractible edge of $Q[L]$ (which exists by Lemma 3.5.8).

Sublemma 3.5.9. *Q/vw is 2-connected.*

Proof. As Q is 2-connected and Q/vw is a contraction, it suffices to show that $Q - v - w$ is connected. Since vw is a contractible edge of $Q[L]$, the set $Q[L] - v - w$ is connected. So either $Q - v - w$ is connected or else the connected set $Q[R]$ can only have v or w as neighbours in $Q[L]$.

Hence we may assume that we have the second outcome. Our aim is to derive a contradiction in that case. More precisely, we show that (Q, d) is not irreducible. We obtain $\hat{Q}$ from Q by contracting a spanning tree of $Q[L] - v - w$

and an edge from that set to one of v or w. Note that $\hat{Q}$ is isomorphic to the graph obtained from Q by deleting $Q[L] - v - w$. In our notation we suppress this bijection and just say things like 'v and w are vertices of $\hat{Q}$'.

Our aim is to show that $\hat{Q}$ is 3-connected along d. Suppose not for a contradiction. Then there is a separating set S witnessing that. Let a and b be two vertices in different components of $\hat{Q} - S$. Let P be a path in $Q - S$ joining a and b. If P contains a vertex of $Q[L] - v - w$, we can shortcut it by the edge vw. Hence we may assume that P contains no vertex of $Q[L] - v - w$. So P is a path in $\hat{Q} - S$. This is a contradiction to the assumption that a and b are separated by S. Hence $\hat{Q}$ is 3-connected along d. As both sides of d in $\hat{Q}$ contain edges, $(\hat{Q}, d)$ witnesses that (Q, d) is not irreducible. This is the desired contradiction. □

We abbreviate $Q' = Q/vw$. Let $d' = d \cap E(Q')$. Let L' be the left side of d'. The right side of d' is R.

Sublemma 3.5.10. *If Q' is not 3-connected along d', there is a vertex z of $Q[R]$ such that $Q[L] - v - w$ can only have z as a neighbour in $Q[R]$.*

Proof. By Sublemma 3.5.9, there are vertices y of $Q'[L']$ and z of $Q'[R]$ such that $Q' - y - z$ is disconnected. Since Q is 3-connected along d and Q' is a contraction of Q, it must be that y or z is a contraction vertex. Hence y is the vertex vw. Hence $Q - v - w - z$ is disconnected. Since vw is contractible, $Q[L]-v-w$ is connected. By assumption $Q[R]-z$ is connected. So $Q[L]-v-w$ has no neighbour in $Q[R] - z$. □

By Lemma 3.5.8, $Q[L]$ has three contractible edges a_1a_2, b_1b_2 and c_1c_2 such that a_1, a_2, b_1 and b_2 are distinct vertices. Applying Sublemma 3.5.10 to a_1a_2 and b_1b_2 yields that there are at most two vertices of $Q[R]$ that have neighbours in $Q[L]$. There have to be two such vertices as Q is 2-connected. Call these vertices z_1 and z_2. Sublemma 3.5.10 gives the further information that one of them, say z_1, can only be incident to a_1 or a_2 and z_2 can only be to b_1 or b_2. Now we apply Sublemma 3.5.10 to c_1c_2. Since c_1c_2 is distinct from a_1a_2 and b_1b_2, there have to be vertices on these edges not in c_1c_2. By symmetry, we may assume that a_1 and b_1 are not in c_1c_2. Applying Sublemma 3.5.10 to c_1c_2 yields that there is a single z_i such that a_1 and b_1 can only have z_i as a neighbour in $Q[R]$. By symmetry, we may assume that z_i is equal to z_1. Hence z_2 can only have the neighbour b_2 in $Q[L]$. Hence $Q - z_1 - b_2$ is disconnected. This is a contradiction to the assumption that Q is 3-connected along d. Thus (Q, d) is not irreducible. □

Lemma 3.5.11. *If both graphs $Q[L]$ and $Q[R]$ consist of a single edge, then (Q, d) is the special K_4.*

Proof. Since every vertex is in L or R, the graph Q has precisely four vertices. Since no two vertices from different sides of d separate, Q must contain all four edges joining the endvertices of these edges. Hence Q is a the special K_4. □

Lemma 3.5.12. *If both graphs $Q[L]$ and $Q[R]$ are triangles, then (Q, d) is the special prism or has a (proper) special K_4 as a special contraction minor.*

Proof. If Q has only three edges between $Q[L]$ and $Q[R]$, then as Q is 3-connected along d, these edges must form a matching. So (Q, d) is the special prism.

Thus we may assume that Q has at least four edges between $Q[L]$ and $Q[R]$. So $Q[L]$ and $Q[R]$ each contain a vertex that has at least two neighbours on the other side. Call these vertices ℓ and r. Since ℓ and r do not separate, there is an edge $\ell'r'$ between $Q[L]$ and $Q[R]$ that is not incident with ℓ and r. By symmetry, we may assume that ℓ and ℓ' are in $Q[L]$, and r and r' are in $Q[R]$. As r has two neighbours in $Q[L]$, we can contract a single edge of $Q[L]$ different from $\ell\ell'$ such that r is adjacent to the two remaining vertices of $Q[L]$. Similarly, we contract an edge of $Q[R]$ different from rr' such that the vertex of ℓ is adjacent to the two remaining vertices of $Q[R]$. The resulting contraction is a special K_4. □

Lemma 3.5.13. *If $Q[L]$ is a single edge and $Q[R]$ is a triangle, then (Q, d) has a special K_4 as a (proper) special contraction minor.*

Proof. We denote the edge in $Q[L]$ by vw. Since Q is 2-connected, each of v and w has a neighbour in $Q[R]$. If one of them has only a single neighbour in $Q[R]$, then that neighbour together with the other endvertex of vw is 2-separator. This is impossible as Q is 3-connected along d.

Hence v and w have each at least two neighbours in $Q[R]$. So there is a vertex x in $Q[R]$ adjacent to v and w. Contracting the edge not incident with x to a single vertex, yields a special K_4 as a special contraction minor. □

Proof of Lemma 3.5.2. By taking (Q, d) contraction-minimal, we may assume that it is irreducible. We will show that (Q, d) is a special prism or a special K_4. If both graphs $Q[L]$ and $Q[R]$ are 2-connected, then by Lemma 3.5.7 (and the same lemma applied with the roles of 'L' and 'R' interchanged) both of them are triangles. In this case, by Lemma 3.5.12 (Q, d) is a special prism.

Otherwise one of $Q[L]$ or $Q[R]$ is not 2-connected. By Lemma 3.5.3 (and the same lemma applied with the roles of 'L' and 'R' interchanged) it consists of a single edge. Hence we may assume that one of the two graphs $Q[L]$ and $Q[R]$ must be a single edge. By combining Lemma 3.5.3 with Lemma 3.5.7, we deduce that the other graph must be a single edge or a triangle. It cannot be a triangle by Lemma 3.5.13. Hence (Q, d) is a special K_4 by Lemma 3.5.11 in this case. □

Proof of Theorem 3.3.1. We have just finished the proof of Lemma 3.5.2. And just after the statement of that lemma we showed that it implies Theorem 3.3.1. □

3.6 Concluding remarks

There are various ways how Theorem 3.3.1 might be extended. First, can we replace 'constraint connectedness' by the property that the set X has at most k connected components for some natural number k? More precisely, a constraint graph (G, X) has at *most k islands* if $G[X]$ has at most k connected components. Clearly, the class of constraint graph with at most k islands is closed under taking constraint minors.

Conjecture 3.6.1. *Let $k > 1$. The class of 3-connected constraint graphs with at most k islands is characterised by a finite list of excluded constraint minors.*

Can you explicitly compute the list of excluded minors in Conjecture 3.6.1?

Another extension is as follows. A *double-constraint matroid* (M, X, Y) consists of a matroid M and two sets X and Y of its elements. It is *realisable* if M is the cycle matroid of a graph G such that both X and Y are connected in G. Can you extend Theorem 3.2.1 from constraint matroids to double-constraint matroids? Put another way: is a double-constraint matroid realisable if and only if it does not have one of finitely many excluded double-constraint minors? Although for 3-connected matroids, the answer to this question follows from Theorem 3.2.1, for matroids that are not 3-connected new obstructions arise, see Figure 3.8

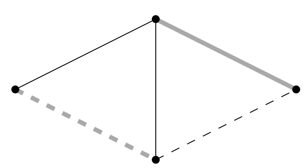

Figure 3.8: The constraint X is depicted in grey, the constraint Y is dashed. Although the matroid represented by this graph is realisable for each of X or Y, it is not realisable for both of them at the same time.

Chapter 4

Dual matroids

4.1 Abstract

We introduce dual matroids of 2-dimensional simplicial complexes. Under certain necessary conditions, duals matroids are used to characterise embeddability in 3-space in a way analogous to Whitney's planarity criterion.

We further use dual matroids to extend a 3-dimensional analogue of Kuratowski's theorem to the class of 2-dimensional simplicial complexes obtained from simply connected ones by identifying vertices or edges.

4.2 Introduction

A well-known characterisation of planarity of graphs is Whitney's theorem from 1932. It states that a graph can be embedded in the plane if and only if its dual matroid is *graphic* (that is, it is the cycle matroid of a graph) [99].

In this chapter we define dual matroids of (2-dimensional) simplicial complexes. We prove under certain necessary assumptions an analogue of Whitney's characterisation for embedding simplicial complexes in 3-space. More precisely, under these assumptions a simplicial complex can be embedded in 3-space if and only if its dual matroid is graphic.

Our definition of dual matroid is inspired by the following fact.

Theorem 4.2.1. *Let C be a directed 2-dimensional simplicial complex embedded into $\mathbb{S}^3$. Then the edge/face incidence matrix of C represents over the integers*[1] *a matroid M which is equal to the cycle matroid of the dual graph of the embedding.*

Indeed, we define[2]. the *dual matroid* of a simplicial complex C to be the matroid represented by the edge/face incidence matrix of C over the finite field $\mathbb{F}_3$.

[1] See Section 4.3 for a definition.

[2] The choice of $\mathbb{F}_3$ is a bit arbitrary. Indeed any other field $\mathbb{F}_p$ with p a prime different from 2 works.

Although the cone over K_5 does not embed in 3-space[3], its dual matroid just consists of a bunch of loops, and thus is graphic. In order to exclude examples like the cone over K_5 we restrict our attention to simplicial complexes C whose dual matroid captures the local structure at all vertices of C. We call such dual matroids *local*, see Section 4.4 for a precise definition. Examples of simplicial complex whose dual matroid is local are those where every edge is incident with precisely three faces and the dual matroid has no loops. Another example is the 3-dimensional grid whose faces are the 4-cycles.

Furthermore matroids (of graphs and also of simplicial complexes) do not depend on the orderings of edges on cycles. Hence it can be shown that dual matroids cannot distinguish triangulations of homology spheres[4] from triangulations of the 3-sphere. While the later ones are always embeddable, this is not true for triangulations of homology spheres in general. Thus we restrict our attention to simply connected simplicial complexes. Under these necessary restrictions we obtain the following 3-dimensional analogue of Whitney's theorem.

Theorem 4.2.2. *Let C be a simply connected 2-dimensional simplicial complex whose dual matroid M is local.*

Then C is embeddable in 3-space if and only if M is graphic.

Tutte's characterisation of graphic matroids [93] yields the following consequence.

Corollary 4.2.3. *Let C be a simply connected simplicial complex whose dual matroid M is local.*

Then C is embeddable in 3-space if and only if M has no minor isomorphic to U_4^2, the fano plane, the dual of the fano plane or the duals of either $M(K_5)$ or $M(K_{3,3})$. $\square$

We further apply dual matroids to study embeddings in 3-space of – not necessarily simply connected – simplicial complexes with locally small separators as follows.

Given a 2-dimensional simplicial complex C, the *link graph*, denoted by $L(v)$, at a vertex v of C is the graph whose vertices are the edges incident with v and whose edges are the faces incident with v and their incidence relation is as in C. If the link graph at v is not connected, we can split v into one vertex for each connected component. There is a similar splitting operation at edges of C. It can be shown that no matter in which order one does all these splittings, one always ends up with the same simplicial complex, *the split complex of C*.

It can be shown that if a simplicial complex embeds topologically into $\mathbb{S}^3$, then so does its split complexes. However, the converse is not true. For an example see Figure 4.1. Here we give a characterisation of when certain simplicial complexes embed, where one of the conditions is that the split complex embeds.

[3]See for example Chapter 1.

[4]These are compact connected 3-manifolds whose homology groups are trivial. Unlike in the 2-dimensional case, this does not imply that the fundamental group is trivial.

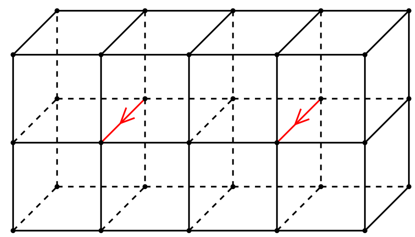

Figure 4.1: The $4 \times 2 \times 1$-grid whose faces are the 4-cycles. It can be shown that the complex obtained by identifying the two edges coloured red cannot be embedded in 3-space.

Theorem 4.2.4. *Let C be a globally 3-connected simplicial complex and $\hat{C}$ be its split complex. Then C embeds into $\mathbb{S}^3$ if and only if $\hat{C}$ embeds into $\mathbb{S}^3$ and the dual matroid of C is the cycle matroid of a graph G and for any vertex or edge of C the set of faces incident with it is a connected edge set of G.*

Here a simplicial complex C is *globally 3-connected*[5] if its dual matroid is 3-connected. For an extension of Theorem 4.2.4 to simplicial complexes that are not globally 3-connected, see Theorem 4.5.19 below.

The condition that a given set of elements of the dual matroid is connected (in some graph representing that matroid) can be characterised by a finite list of obstructions as follows. Given a matroid M and a set X of its elements, a *constraint minor* of (M, X) is obtained by contracting arbitrary elements or deleting elements not in X. In Chapter 3, we prove for any 3-connected graphic matroid M (that is a 3-connected graph) with an edge set X that X is connected in M if and only if (M, X) has no constraint minor from the finite list depicted in Figure 4.2.

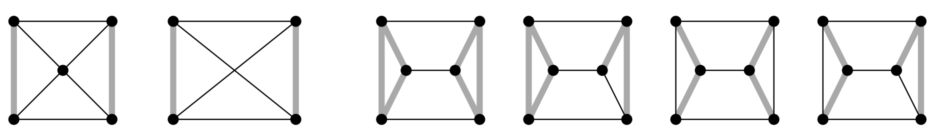

Figure 4.2: The six obstructions characterising connectedness of X. In these graphs we depicted the edge set X in grey.

In Chapter 1, we introduced *space minors* of simplicial complexes and proved that a simply connected locally 3-connected simplicial complex C embeds in 3-space if and only if it does not have a space minor from a finite list $\mathcal{L}$ of obstructions. Using Theorem 4.2.4 we can further extend this characterisation from simply connected simplicial complexes to those whose split complex is simply connected.

[5]In Section 4.7 we give an equivalent definition directly in terms of C.

Theorem 4.2.5. *Let C be a globally 3-connected simplicial complex such that the split complex is simply connected and locally 3-connected*[6]*. Then C embeds into $\mathbb{S}^3$ if and only if its split complex has no space minor from $\mathcal{L}$ and the dual matroid has no constraint minor from the list of Figure 4.2.*

If we do not require global 3-connectivity in Theorem 4.2.5, there are infinitely many obstructions to embeddability, see Section 4.6. We remark that Theorem 4.2.2 can be extended from simply connected simplicial complexes to those whose split complex is simply connected.

The chapter is structured as follows. In Section 4.3 we prove Theorem 4.2.1, which is used in the proof of Theorem 4.2.2 and Theorem 4.2.4. In Section 4.4 we prove Theorem 4.2.2. In Section 2.5 we prove Theorem 4.2.4 and Theorem 4.2.5. Finally in Section 4.6 we construct infinitely many obstructions to embeddability in 3-space (inside the class of simplicial complexes with a simply connected and locally 3-connected split complex).

For graph we follow the notations of [36] and for matroids [76]. Beyond that we rely on some definitions of Chapter 2. The idea of using dual graphs to study embeddability in 3-space also appears in [41].

4.3 Dual matroids

In this section we prove Theorem 4.2.1 and the fact that a simplicial complex and its split complexes have the same dual matroid, which are used in the proofs of Theorem 4.2.2 and Theorem 4.2.4.

A *directed simplicial complex* is a simplicial complex C together with an assignment of a direction to each edge of C and together with an assignment of a cyclic orientation to each face of C. A *signed incidence vector* of an edge e of C has one entry for every face f; this entry is zero if e is not incident with f, it is plus one if f traverses e positively and minus one otherwise.

The matrix given by all signed incidence vectors is called the *(signed) edge/face incidence matrix*. The *dual matroid* of a simplicial complex is the matroid represented by the edge/face incidence matrix of C over the finite field $\mathbb{F}_3$.

Although in this chapter we work with directed simplicial complexes, dual matroids do not depend on the chosen directions. Indeed, changing a direction of an edge or of a face of C changes the linear representation of the dual matroid but not the matroid itself.

A matrix A is a *regular representation* (or *representation over the integers*) of a matroid M if all its entries are integers and the columns are indexed with the elements of M. Furthermore for every circuit o of M there is a $\{0, -1, +1\}$-valued vector[7] v_o in the span over $\mathbb{Z}$ of the rows of A whose support is o. And

[6] In Chapter 5 we discuss how this result can be extended to simplicial complexes whose split complexes are not local 3-connected.

[7] A *vector* is an element of a vector space k^S, where k is a field and S is a set. In a slight abuse of notation, in this chapter we also call elements of modules of the form $\mathbb{Z}^S$ vectors.

the vectors v_o span over $\mathbb{Z}$ all row vectors of A.

4.3.1 Proof of Theorem 4.2.1

Let C be a directed simplicial complex embedded into $\mathbb{S}^3$, the *dual digraph* of the embedding is the following. Its vertex set is the set of components of $\mathbb{S}^3 \setminus C$. It has one edge for every face of C. This face touches one or two components of $\mathbb{S}^3 \setminus C$. If it touches two components, the edge for that face joins the vertices for these two components. The edge is directed from the vertex whose complement touches the chosen orientation of the face to the other component. If the face touches just one component, its edge is a loop attached at the vertex corresponding to that component.

Let $(\sigma(e)|e \in E(C))$ be the planar rotation system of C induced by the topological embedding of C. It is not hard to check that $\sigma(e)$ is a closed trail[8] in the dual graph. The *dual complex* of the embedding is the directed simplicial complex obtained from the dual digraph by adding for each edge of C the cyclic orderings of the cyclic orientations $\sigma(e)$ as faces and we choose their orientations to be $\sigma(e)$.

Observation 4.3.1. *Let C be a connected and locally connected[9] simplicial complex embedded in $\mathbb{S}^3$ with induced planar rotation system Σ. Then the dual complex of the embedding is equal to the dual complex of (C, Σ).*

Proof. By Lemma 2.4.4, the local surfaces for (C, Σ) agree with the local surfaces of the embedding[10]. Hence these two complexes have the vertex set. As they also have the same incidence relations between edges and vertices and edges and faces, they must coincide. □

By Observation 4.3.1 and the definition of 'generated over the integers' and by Theorem 4.8.6, in order to prove Theorem 4.2.1 it suffices to show that the dual complex for (C, Σ) is nullhomologous[11].

First we prove this in the special case when C is nullhomologous and locally connected.

Lemma 4.3.2. *Let C be a nullhomologous locally connected simplicial complex together with a planar rotation system Σ such that local surfaces for (C, Σ) are spheres.[12] Then the dual complex D of (C, Σ) is nullhomologous.*

[8] A *trail* is sequence $(e_i|i \leq n)$ of distinct edges such that the endvertex of e_i is the starting vertex of e_{i+1} for all $i < n$. A trail is *closed* if the starting vertex of e_1 is equal to the endvertex of e_n.

[9] A simplicial complex C is *locally connected* if all its link graphs are connected.

[10] Local surfaces of embeddings are defined in Chapter 2.

[11] A simplicial complex C is *nullhomologous* if the face boundaries of C generate all cycles over the integers. This is equivalent to the condition that the face boundaries of C generate all cycles over the field $\mathbb{F}_p$ for every prime p.

[12] This last property follows from the first two if we additionally assume that Σ is induced by a topological embedding in $\mathbb{S}^3$ by Theorem 2.7.1.

Proof. By Lemma 2.7.5, Lemma 2.7.7 and Lemma 2.7.3 the complexes C and D satisfy euler's formula, that is:

$$|V(C)| - |E| + |F| - |V(D)| = 0$$

Hence we deduce that D nullhomologous by applying the 'Moreover'-part of Lemma 2.7.3 for every prime p. $\square$

Next we shall extend Lemma 4.3.2 to simplicial complexes that are only locally connected.

Lemma 4.3.3. *Let C be a locally connected simplicial complex together with a planar rotation system Σ that is induced by a topological embedding ι in $\mathbb{S}^3$. Then the dual complex D of (C, Σ) is nullhomologous.*

Proof. By Theorem 2.8.1 there is a simplicial complex C' that is obtained from C by subdividing edges, baricentric subdivisions of faces and adding faces along closed trails. And C' is nullhomotopic and has an embedding ι' into $\mathbb{S}^3$ that induces[13] ι. Let D' be the dual of ι'. By Lemma 4.3.2, D' is nullhomologous.

We shall deduce that D is nullhomologous by showing that reversing each of the operations in the construction of C' from C preserves being nullhomologous in the dual. We call such an operation *preserving*.

Sublemma 4.3.4. *Subdividing an edge is preserving.*

Proof. Subdividing an edge in the primal corresponds to adding a copy of a face in the dual. Clearly, the deletion of the copy preserves being nullhomologous for the dual. $\square$

Sublemma 4.3.5. *A baricentric subdivision of a face is preserving.*

Proof. It suffices to show that the subdivision by a single edge is preserving. Subdividing a face by an edge in the primal corresponds to replacing an edge in the dual by two edges in parallel and adding a face containing precisely these two edges. Reversing this operation preserves being nullhomologous. $\square$

Sublemma 4.3.6. *Adding a face is preserving.*

Proof. Adding a face in the primal corresponds to coadding[14] an edge in the dual. Contracting that edge preserves being nullhomologous. $\square$

By Sublemma 4.3.4, Sublemma 4.3.5 and Sublemma 4.3.6, the fact that D' is nullhomologous implies that D is nullhomologous. $\square$

[13] This means that we obtain ι from ι' by deleting the newly added faces, contracting the newly added subdivision edges and undoing the baricentric subdivisions.

[14] A complex A is obtained from a complex A' by *coadding* an edge e if A' is obtained from A by contracting the edge e.

It remains to prove Theorem 4.2.1 for simplicial complexes C that are not locally connected. First we need some preparation.

Given a simplicial complex C, its *vertical split complex* is obtained from C by replacing each vertex v by one vertex for each connected component of $L(v)$, where the edges and faces incident with that vertex are those in its connected component. We refer to these new vertices as the *clones* of v.

Observation 4.3.7. *The vertical split complex of any simplicial complex is locally connected.* □

Observation 4.3.8. *A simplicial complex and its vertical split complex have the same dual matroid.*

Proof. A simplicial complex and its vertical split complex have the same edge/face incidence matrix. □

Given an embedding ι of a simplicial complex C into $\mathbb{S}^3$, we will define what an *induced* embedding of the vertical split complex is.

For that we need some preparation. Let v be a vertex of C whose link graph is not connected. By changing ι a little bit locally (but not its induced planar rotation system) if necessary, we may assume that there is a 2-ball B of small radius around v such that firstly v is the only vertex of C contained in the inside of B. And secondly its boundary ∂B intersects each edge incident with v in a point and each face incident with v in a line. In other words, the intersection of C with the boundary is the link graph at v. As the link graph is disconnected, there is a circle (homeomorphic image of $\mathbb{S}^1$) γ in the boundary such that the two components of $B \setminus \gamma$ both contain vertices of the link graph, see Figure 4.3.

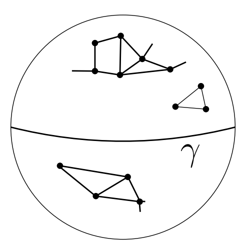

Figure 4.3: The link graph at v embedded into ∂B.

The simplicial complex C_γ is obtained from C by replacing the vertex v by two vertices, one for each connected component of $\partial B \setminus \gamma$ that is incident with the edges and faces whose vertices and edges, respectively, are in that connected component.

The embedding ι *induces*[15] the following embedding ι_γ of C_γ into $\mathbb{S}^3$. We pick a disc contained in B with boundary γ that intersects C only in v. We

[15] The construction of ι_C depends on the choice of B. Still we use the term 'induced' in this context since in this chapter we consider topological embeddings equivalent if they have the same planar rotation system.

replace v by its two clones – both with tiny distance from v and one above that disc and the other below. We only need to change faces and edges incident with v in a tiny neighbourhood around v. Faces and edges above and below do not interfere.

It is easy to see that ι and ι_C have the same planar rotation system and that C and C_γ have the same vertical split complex.

A topological embedding of the vertical split complex of C into $\mathbb{S}^3$ is *(vertically) induced* by ι if it is obtained by applying the above procedure iteratively until C_γ is equal to the vertical split complex of C. It is clear that if ι is a topological embedding of a simplicial complex C into $\mathbb{S}^3$, then its vertical split complex has a topological embedding into $\mathbb{S}^3$ that is induced by ι.

Observation 4.3.9. *Let ι be an embedding of a simplicial complex into $\mathbb{S}^3$ and let ι' be an induced embedding of ι of the vertical split complex. Then ι and ι' have the same dual complex.*

Proof. In both embeddings, the incidence relation between the local surfaces and the faces is the same. Hence both dual complexes have the same vertex/edge incidence relation. They also have the same sets of faces as ι and ι' have the same rotation system. $\square$

A set S of vertices in a simplicial complex C is a *vertex separator* if C can be obtained from two disjoint simplicial complexes that each have at least one face by gluing them together at the vertex set S. As the empty set might also be a vertex separator, any simplicial complex with no vertex separator is connected.

Lemma 4.3.10. *Let C be a simplicial complex without a vertex separator. Assume that C has an embedding ι into $\mathbb{S}^3$. Then the dual complex D of ι is nullhomologous.*

Proof. Let C' be the vertical split complex of C. By Observation 4.3.7, C' is locally connected. By assumption C has no vertex separator. Thus C' is connected. Let ι' be the embedding of C' induced by ι. Let Σ' be the planar rotation system induced by ι'.

By Lemma 4.3.3, the dual D' for (C', Σ') is nullhomologous. By Observation 4.3.1, D' is the dual complex of ι'. By Observation 4.3.9, D' is equal to D. So D is nullhomologous. $\square$

Lemma 4.3.11. *Let C be a simplicial complex embedded into $\mathbb{S}^3$ that is obtained from two simplicial complexes C_1 and C_2 by gluing them together at a set of vertices. Assume that C_2 has no separating vertex set. Let G_i be the dual graph of the embedding restricted to C_i for $i = 1, 2$. Then the dual graph of the embedding of C is equal to a graph obtained by gluing together G_1 and G_2 at a single vertex.*

Proof. We denote the embedding of C into $\mathbb{S}^3$ by ι and the restricted embedding of C_1 by ι_1. Suppose for a contradiction that ι maps interior points of faces of C_2 to interior points of different local surfaces of ι_1. Let ℓ be a local surface of ι_1 to which an interior point of a face of C_2 is mapped by ι. Let C_2' be the subcomplex of C_2 that contains all faces whose interior points are mapped to interior points of ℓ. Its edges and vertices are those of C_2 that are incident with these faces. Note that if one interior point of a face is mapped to ℓ, then all are. Hence the subcomplex C_2'' that contains all other faces and their incident vertices and edges contains a face. The subcomplexes C_2' and C_2'' of C_2 can only intersect in points of C_1. Hence they only can intersect in vertices. Thus C_2' and C_2'' witness that C_2 has a separating vertex set contrary to our assumption.

Thus there is a single local surface of ι_1 to which all interior points of faces of C_2 are mapped by ι. Hence the dual graph of ι is equal to the graph obtained by gluing together G_1 and G_2 at that vertex. □

Proof of Theorem 4.2.1. By applying Lemma 4.3.11 recursively, we may assume that C has no separating vertex set. Recall that the dual graph of the embedding is the 1-skeleton of the dual complex of the embedding. By Lemma 4.3.10, the edge/face incidence matrix is a representation over the integers of the cycle matroid of the dual graph of the embedding. □

4.3.2 Split complexes

A naive way to define splittings of edges might be to consider the incidences at one of their endvertices and split according to that. We shall show that when using this notion of splitting, split complexes will not have all nice properties we want them to have, see Section 4.7. A more refined definition takes into account the incidences at both endvertices, defined as follows.

Given a simplicial complex C and an edge e with two endvertices v and w, two faces incident with e are *v-related* if - when considered as edges of e, they have endvertices in the same connected component of the link graph $L(v) - e$ with the vertex e removed. Analogously, we define *w-related*. Two faces f_1 and f_2 incident with e are in the *same connected component at e* if there is a chain of faces incident with e from f_1 to f_2 such that adjacent faces in the chain are v-related or w-related. Note that 'being in the same connected component at e' is the equivalence relation generated from the union of 'v-related' and 'w'-related.

The simplicial complex obtained from C by *splitting* the edge e is obtained from C by replacing the edge e by one copy e_X for every connected component X at e. The faces incident with e_X are those in X.

We refer to the edges e_X as the *clones* of e. If we apply several splittings, we extend the notion of cloning iteratively so that each edge of the resulting simplicial complex is cloned from a unique edge of C.

If we split an edge in a nontrivial way, then the resulting simplicial complex has the same number of faces but at least one edge more. As in a simplicial complex every edge is incident with a face, we can only split edges a bounded number of times. A simplicial complex obtained from C by splitting edges such that for every edge there is only one component at e is called an *edge split complex* of C. As explained above, every simplicial complex has an edge split complex.

Since splitting edges, does not change the 2-blocks of the link graphs, splittings of edges commute. In particular, edge split complexes are unique. In the following we will talk about 'the edge split complex'.

The *split complex* of a simplicial complex C is the vertical split complex of its edge split complex. Clearly, splitting a vertex does not change the edge split complex.

Example 4.3.12. A simplicial complex, its vertical split complex and its edge split complex have the same split complex. Locally 2-connected[16] simplicial complexes are equal to their split complex.

Lemma 4.3.13. *A simplicial complex and its edge split complex have the same dual matroid.*

Proof. We shall show that a simplicial complex C and a simplicial complex C' have the same dual matroid, where we obtain C' from C by splitting an edge e. Once this is shown, the lemma follows inductively as an edge split complex is obtained by a sequence of edge splittings.

Clearly, C and C' have the same set of faces. Hence their dual matroids have the same ground sets.

The vectors indexed by clones of the edge e of the edge/face incidence matrix A' of C' sum up to the vector indexed by e of the edge/face incidence matrix A of C. Hence the vectors indexed by edges of A' generate the vectors indexed by edges of A. So it remains to show that any vector indexed by a clone e' of e of A' is generated by the vectors indexed by edges of A.

Let v be an endvertex of e. Let K be the connected component of the link graph $L(v)$ of C at v that contains e. Let Y be the union of the components Y' of $K - e$ such that faces incident with e' – when considered as edges of $L(v)$ – have an endvertex in Y'. The sum over all vectors indexed by edges $y \in V(Y)$ of A is the vector indexed by e' of A'. Since e' was an arbitrary clone, the vectors indexed by edges of A generate the vectors indexed by edges of A'.

We have shown that splitting a single edge preserves the dual matroid. Since the edge split complex is obtained by splitting edges, it must have the same dual matroid as the original complex. □

Corollary 4.3.14. *A simplicial complex and its split complex have the same dual matroid.*

Proof. A simplicial complex and its vertical split complex have the same incidence relations between edges and faces. Hence this is a consequence of Lemma 4.3.13. □

4.4 A Whitney type theorem

In this section we prove Theorem 4.2.2.

[16] A simplicial complex is *locally 2-connected* if its link graphs are connected and have no cutvertices.

In general the dual matroid of a simplicial complex C does not contain enough information to decide whether C is embeddable in 3-space. For example, the dual matroid of the cone over K_5 consists of a bunch of loops. So it cannot distinguish this non-embeddable simplicial complex from other embeddable ones. The following fact gives an explanation of this phenomenon (in the notation of that fact: from the graph G we can in general not reconstruct the matroid $M[v]$). Given a vertex v of a simplicial complex, we denote the dual matroid of the link graph at v by $M[v]$.

Fact 4.4.1. *Let C be a simplicial complex embedded in $\mathbb{S}^3$. Then the dual matroid M restricted to the faces incident with v is represented by a graph G. Moreover, G can be obtained from some graph representing $M[v]$ by identifying vertices.*

Proof. By Theorem 4.2.1 M is the cycle matroid of the dual graph of the embedding of C. So G is the restriction of that graph to the faces incident with v.

By G' be denote the 'local dual graph' of C at v. This is defined as the 'dual graph' but with '$\mathbb{S}^3$' replaced by 'a small neighbourhood U around v' in the embedding. Clearly, G' represents $M[v]$. We obtain the vertices of G from those of G' by identifying those vertices for components of $U \setminus C$ that lie in the same component of $\mathbb{S}^3 \setminus C$. The 'Moreover'-part follows. □

To exclude the phenomenon described in Fact 4.4.1 we restrict our attention to simplicial complexes C whose dual matroid captures the local structure at all vertices of C, defined as follows. Given a simplicial complex C with dual matroid M, we say that M is *local* if for every vertex v the matroid $M[v]$ is equal to M restricted to the faces incident with v.

Furthermore matroids (of graphs and also of simplicial complexes) do not depend on the orderings of edges on cycles. Hence it can be shown that dual matroids cannot distinguish triangulations of homology spheres[17] from triangulations of the 3-sphere. While the later ones are always embeddable, this is not true for triangulations of homology spheres. Thus we restrict our attention to simply connected simplicial complexes.

If we exclude these two phenomenons, Theorem 4.2.2, stated in the Introduction, characterises when a simplicial complex is embeddable just in terms of its dual matroid.

Remark 4.4.2. The assumptions of Theorem 4.2.2 can be interpreted as some face maximality assumption. By Theorem 2.8.1 this is true for being simply connected. For locality, let C be any embeddable simplicial complex embeddable. By Fact 4.4.1 we can add faces until for every vertex v the matroid $M[v]$ is equal to M restricted to the faces incident with v. This preserves being simply connected.

Now we prepare for the proof of Theorem 4.2.2.

[17] These are compact connected 3-manifolds whose homology groups are trivial. Unlike in the 2-dimensional case, this does not imply that the fundamental group is trivial.

Lemma 4.4.3. *Let H be a graph whose cycle matroid is the dual matroid M of a simplicial complex C. There is a directed graph $\vec{H}$ with underlying graph H such that for all edges e of C the signed vectors are 3-flows*[18].

Proof. First we consider the case when H is 2-connected. We start with an arbitrarily directed graph $\vec{H}$ with underlying graph H some of whose directions of the edges we might reverse later on in the argument. Since H is 2-connected, the set of edges incident with a vertex is a bond of H, which is called the *atomic bond* of v. By elementary properties of representations, there is a vector b_v with all entries -1, $+1$ or 0 that has the same support[19] as the atomic bond at v.

Given an edge e of H and one of its endvertices v, we say that e is *effectively directed towards* v with respect to a vector b with entries in $\mathbb{Z}$ if $\vec{e}$ is directed towards v and $b(e)$ is positive or $\vec{e}$ is directed away from v and $b(e)$ is negative. First we shall prove that we can modify the directions of the edges of $\vec{H}$ such that all edges e of H are directed such that for some endvertex v they are effectively directed towards v with respect to the at b_v.

Let T be a spanning tree of H. Since T does not contain any cycle, we can pick the b_v such that if vw is an edge of T, then $b_v(vw) = -b_w(vw)$. Hence an edge vw of T is effectively directed towards v with respect to b_v if and only if it is effectively directed towards w with respect to b_w. So by reversing the direction of an edge if necessary[20], we may assume that every edge vw of T is effectively directed towards v with respect to b_v and also effectively directed towards w with respect to b_w.

Next let xy be an edge not in T. By reversing the direction of xy if necessary we may assume that xy is effectively directed towards x with respect to b_x. Our aim is to show that xy is effectively directed towards y with respect to b_y. Let C be the fundamental circuit of xy with respect to T. By elementary properties of representations, there is a vector v_C with support C that is orthogonal over $\mathbb{F}_3$ to all the vectors b_z for vertices z on C. At all vertices z of C except possibly y, the two edges on C incident with z are effectively directed towards z with respect to the vector b_z. Hence for v_C to be orthogonal, precisely one of these edges must be effectively directed towards z with respect to v_C. Using this property inductively along C, we deduce that of the two edges on C incident with y also precisely one is effectively directed towards y with respect to v_C. Since b_y is orthogonal to v_C and the edge incident with y that is on T and C is effectively directed towards y with respect to b_y, also xy must be effectively directed towards y with respect to b_y.

Hence our final directed graph $\vec{H}$ has the property that all edges e of H are effectively directed towards any of their endvertices v with respect to b_v. Since signed vectors of edges e of C are orthogonal at to b_v, it follows that it accumulates 0 (mod 3) at all vertices v. So the signed vectors of C are 3-flows

[18] A *3-flow* in a directed graph $\vec{H}$ is an assignment of integers to the edges of $\vec{H}$ that satisfies Kirchhoff's first law modulo three at every vertex of $\vec{H}$.

[19] The *support* of a vector is the set of coordinates with nonzero values.

[20] To be very formal, we delete the edge from the graph and glue it back the other way round. Note that we do not change the director.

for $\vec{H}$. This completes the proof if H is 2-connected. If H is not 2-connected, we do the same construction independently in every 2-connected component and the result follows. □

First we prove Theorem 4.2.2 under the additional assumption that C is locally 2-connected:

Lemma 4.4.4. *Let C be a simply connected locally 2-connected simplicial complex whose dual matroid is local.*

Then C is embeddable in 3-space if and only if M is graphic.

Proof. Assume that C is embeddable and let D by its dual complex. Then by Theorem 4.2.1 M is equal to the cycle matroid of the 1-skeleton of D. In particular M is graphic.

Now conversely assume that C is a simply connected simplicial complex such for every vertex v the matroid $M[v]$ is equal to dual matroid M restricted to the faces incident with v; and that there is a graph G whose cycle matroid is M. We pick an arbitrary direction at each edge of C and an arbitrary orientation at each face of C. Our aim is to construct a planar rotation system Σ of C and apply Theorem 2.2.1 to deduce that C is embeddable.

By Lemma 4.4.3 there is a direction $\vec{G}$ of G such that the signed incidence vector v_e for each edge e of C is a 3-flow in $\vec{G}$. As the link graph $L(v)$ at each vertex v is 2-connected, none of its vertices e is a cutvertex. Hence the edges incident with e in $L(v)$ form a bond. So they form a circuit in the dual matroid $M[v]$. Thus by assumption the support of v_e is a circuit in the matroid M. By the construction of $\vec{G}$, the signed vector v_e is a directed cycle[21] in $\vec{G}$. This directed cycle defines a cyclic orientation $\sigma(e)$. In terms of C this is a cyclic orientation of the oriented faces incident with the directed edge $\vec{e}$. Put another way $\Sigma = (\sigma(e)|e \in E(C))$ is a rotation system.

Our aim is to prove that Σ is planar. So let v be a vertex of C and let Σ_v be the rotation system of the link graph $L(v)$ induced by Σ. This rotation system of $L(v)$ defines an embedding of $L(v)$ in a 2-dimensional oriented surface S_v in the sense of [71][22]. It remains to show the following.

Sublemma 4.4.5. *S_v is a sphere.*

Proof. As the graph $L = L(v)$ is connected, S_v is connected. Thus it suffices to show that it has Euler genus two, that is:

$$V_L - E_L + F_L = 2 \tag{4.1}$$

Here we abbreviate: $|V(L)| = V_L$, $|E(L)| = E_L$ and F_L denotes the faces of the embedding of $L(v)$ in S_v.

We denote the dual graph of the embedding of L in S_v by H. Our aim is to show that H is equal to the restriction R of G to the faces incident with v. We

[21] A vector v whose entries are in $\{0, +1, -1\}$ is a *directed cycle* if its support is a cycle and it satisfies Kirchhoff's first law at every vertex, see [36].

[22] This is explained in more detail in Chapter 2.

obtain S' from R by gluing on each directed cycle v_e the face $\sigma(e)$. Similarly as in Chapter 2 we use the Edmonds-Hefter-Ringel rotation principle [71, Theorem 3.2.4] to deduce that L is the surface dual of R with respect to the embedding into S'. In particular $S' = S$ and R is equal to H.

Having shown that R is the surface dual of L, we conclude our proof of Equation 4.1 as follows. We denote the dimension of the cycle space of L by d. We have $V_L - E_L = -d + 1$ and $F_L = V_R$ (where V_R is the number of vertices of V_R). Hence in order to prove Equation 4.1 it suffices to show that $d = V_R - 1$. This follows from the assumption that the cycle matroid of R is the dual of the cycle matroid of L. Indeed, the cycle matroid of L is 2-connected by assumption. $\square$

$\square$

Proof of Theorem 4.2.2. As in the proof of Lemma 4.4.4, by Theorem 4.2.1 it suffices to show that any simply connected simplicial complex C whose dual matroid M is graphic and local can be embedded in 3-space.

We prove this in two steps. First we prove it for locally connected simplicial complexes. We prove this by induction. The base case is when C is locally 2-connected and this is dealt with in Lemma 4.4.4. So now we assume that C has a vertex v such that the link graph $L(v)$ has a cut vertex[23]; and that we proved the statement for every simplicial complex as above such that it has a fewer number of cutvertices – summed over all link graphs. Let e be an edge of C that is a cutvertex in $L(v)$.

Sublemma 4.4.6. *The simplicial complex C is obtained from a simplicial complex C' by identifying two vertex-disjoint edges e_1 and e_2 onto e.*

Proof. In the link graph $L(v)$, let f_1 and f_2 be two edges incident with e that are in different 2-blocks of $L(v)$. Hence $L(v)$ has a 1-separation (X_1, X_2) with cutvertex e such that f_i is in the side X_i for $i = 1, 2$.

Let w be the endvertex of e in C different from v. Our aim is to construct a 1-separation (Y_1, Y_2) with cutvertex e of $L(w)$ such that X_i and Y_i agree when restricted to the edges incident with e for $i = 1, 2$. For that we have to show that if two such edges are in different X_i then they do not lie in the same 2-block of $L(w)$. That is, in the matroid $M[w]$ they do not lie in a common circuit consisting of edges incident with e. By the assumption, this property is true in $M[w]$ if and only if it is true in M if and only if it is true in $M[v]$, which it is not true as (X_1, X_2) is a 1-separation.

We obtain C' from C by replacing v by two new vertices v_1 and v_2 and w by two new vertices w_1 and w_2. A face or edge incident with v is in v_i if and only if it is in X_i. Similarly, a face or edge incident with w is incident with w_i if and only if it is in Y_i. Thus every edge or face incident with v is incident with precisely one of v_1 and v_2 except for the edge e for which we introduce two

[23] A vertex v of a graph is a *cut vertex* if the component of the graph containing v with v removed is disconnected.

copies, which we denote by e_1 and e_2. The same is holds with 'w' in place of 'v'. Clearly, the edge e_i joins v_i and w_i. Hence C' has the desired properties. $\square$

Sublemma 4.4.7. *The edges e_1 and e_2 lie in different connected components of C'.*

Proof. The simplicial complex C/e is simply connected and obtained from $C'/\{e_1, e_2\}$ by identifying the vertices e_1 and e_2 onto e. Since C/e is not locally connected at e we can apply Lemma 2.6.1 to deduce that e has to be a cutvertex of C/e.

Since the link graph $L(e)$ of C/e is a disjoint union of the connected link graphs $L(e_1)$ and $L(e_2)$ of $C'/\{e_1, e_2\}$, two faces incident with the same edge e_i in C' cannot be cut off by e in C/e. Hence the only way e can cut C/e is that e_1 and e_2 are cut off from one another. Put another way, e_1 and e_2 lie in different connected components of C'. $\square$

For $i = 1, 2$, let C_i be the component of C' containing e_i and M_i the dual matroid of C_i. We may assume that C is connected. Hence C' is the disjoint union of the C_i. By Sublemma 4.4.7 and Lemma 4.3.13, the dual matroid M of S is the disjoint union of the matroids M_i. So we can apply the induction hypothesis to each simplicial complex C_i. So all C_i are embeddable. Analogously to Lemma 2.6.2 one proves that C is embeddable in 3-space[24].

Finally, we prove the statement for arbitrary simplicial complexes. Again, we prove it by induction. This time the locally connected case is the base case. So now we assume that C has a vertex v such that the link graph $L(v)$ is disconnected; and that we proved the statement for every simplicial complex as above such that the number of components of link graphs minus the total number of link graphs is smaller. As C is simply connected, by Lemma 2.6.1 the vertex v is a cutvertex of C. That is, C is obtained from gluing together two simplicial complexes C' and C'' at the vertex v. Since splitting vertices preserves dual matroids, the dual matroid of C is the disjoint union of the dual matroid of C' and the dual matroid of C''. Thus the simplicial complexes C' and C'' are embeddable in $\mathbb{S}^3$ by induction. Hence by Lemma 2.6.2 C is embeddable. $\square$

Remark 4.4.8. The proof of Theorem 4.2.2 works also if we change the definition of dual matroid in that we replace '$\mathbb{F}_3$' by '$\mathbb{F}_p$ with p prime and $p > 2$'. By Theorem 4.2.1, if C is embeddable, the signed incidence vectors of the edges of C generate the same matroid over any field $\mathbb{F}_p$ with p prime. So if C is embeddable all these definitions of dual matroids coincide.

The special role of $p = 2$ is visible in Corollary 4.2.3, where we have to exclude the matroid $U_{2,4}$, which is representable over any field $\mathbb{F}_p$ with p prime and $p > 2$ but not over $\mathbb{F}_2$.

[24] An alternative is the following: it is easy to see that a simplicial complex S is embeddable if and only if S/e is embeddable for some nonloop e. So the C_i/e_i are embeddable. Then by Lemma 2.6.2 C/e is embeddable. So C is embeddable.

4.5 Constructing embeddings from embeddings of split complexes

In this section we prove Theorem 4.2.4. We subdivide this proof in four subsections.

4.5.1 Constructing embeddings from vertical split complexes

Lemma 4.5.1. *Let C be a simplicial complex obtained from a simplicial complex C' by identifying two vertices v and w. Let ι' be a topological embedding of C' into $\mathbb{S}^3$. Assume that there is a local surface of ι' that contains both v and w. Then there is a topological embedding of C into $\mathbb{S}^3$ that has the same dual graph as ι'.*

Proof. We join v and w by a copy of the unit interval I inside the local surface of ι' that contains them both. We may assume that there is an open cylinder around I that does not intersect C'. We obtain a topological embedding ι of C from ι' by moving v along I to w. We do this in such a way that we change the edges and faces incident with v only inside the small cylinder. It is clear that ι' and ι have the same dual graph. □

Lemma 4.5.2. *Let x be a vertex or edge of a simplicial complex C embedded into $\mathbb{S}^3$. The set of faces incident with x is a connected edge set of the dual graph of the embedding.*

Proof. If x is an edge, then the set of faces incident with x is a closed trail, and hence connected. Hence it remains to consider the case that x is a vertex. Let H_x be the dual graph of the link graph at x with respect to the embedding in the 2-sphere given by the embedding of C. The restriction R_x of the dual graph of the embedding of C to the faces incident with x is obtained from H_x by identifying vertices. Since H_x is connected, also R_x is connected. This completes the proof. □

Given a simplicial complex C and a topological embedding ι of its vertical split complex into $\mathbb{S}^3$, we say that ι satisfies the *vertical dual graph connectivity constraints* if for any vertex x of C, the set of faces incident with x is a connected edge set of the dual graph of ι.

Theorem 4.5.3. *Let C be a simplicial complex. Then C embeds into $\mathbb{S}^3$ if and only if its vertical split complex $\hat{C}$ has an embedding into $\mathbb{S}^3$ that satisfies the vertical dual graph connectivity constraints.*

Proof. First assume that C has a topological embedding ι in $\mathbb{S}^3$. Let ι' be the embedding induced by ι of $\hat{C}$. By Observation 4.3.9, ι and ι' have the same dual graph. Hence by Lemma 4.5.2, ι' satisfies the vertical dual graph connectivity constraints.

Now conversely assume that ι' is an embedding into $\mathbb{S}^3$ of $\hat{C}$ that satisfies the vertical dual graph connectivity constraints. Let G be the dual graph of ι'. We shall recursively construct a sequence (C_n) of simplicial complexes by identifying vertices that belong to the same vertex of C that all have the vertical split complex $\hat{C}$ and topological embeddings ι_n of C_n into $\mathbb{S}^3$ that all have the same dual graph G.

If $C_n = C$, we stop and are done. So there is a vertex v of C such that C_n has at least two vertices cloned from v. The set of faces incident with v is a connected edge set of G. So there are two distinct vertices v_1 and v_2 of C_n cloned from v whose incident faces share a vertex when considered as edge sets of G. Hence there is a local surface of ι_n that contains v_1 and v_2. We obtain C_{n+1} from C_n by identifying v_1 and v_2. The existence of a suitable embedding ι_{n+1} follows from Lemma 4.5.1.

Since this recursion cannot continue forever, we must eventually have that $C_n = C$. Then ι_n is the desired embedding of C and we are done. □

4.5.2 Constructing embeddings from edge split complexes

Our next step is to prove the following lemma analogously to one of the implications of Theorem 4.5.3. Given a simplicial complex C and a topological embedding ι into $\mathbb{S}^3$ of any of its split complex $\hat{C}$ into $\mathbb{S}^3$, we say that ι satisfies the *dual graph connectivity constraints* (with respect to C) if for any vertex or edge x of C, the set of faces incident with x is a connected edge set of the dual graph of ι.

Lemma 4.5.4. *Let C be a locally connected simplicial complex. Assume that the split complex of C has an embedding ι' into $\mathbb{S}^3$ that satisfies the dual graph connectivity constraints. Then C has an embedding in $\mathbb{S}^3$ that has the same dual graph as ι'.*

Working with a strip instead of a unit interval, one shows the following analoguously to Lemma 4.5.1.

Lemma 4.5.5. *Let C be a simplicial complex obtained from a simplicial complex C' by identifying two edges e and e' with disjoint sets of endvertices. Let ι' be a topological embedding of C' into $\mathbb{S}^3$. Assume that there is a local surface of ι' that contains both e and e'. Then there is a topological embedding of C into $\mathbb{S}^3$ that has the same dual graph as ι'.* □

Proof of Lemma 4.5.4. Since the split complex is independent of the ordering in which we do splittings, the split complex C' of C is obtained by a sequence of the following operations: first we split an edge. Then we split the two endvertices of that edge. After that the complex is again locally connected. So we eventually derive at the split complex.

We make an inductive argument similary as in the proof of Theorem 4.5.3. Thus it suffices to show that if a complex embeds and satisfies the dual graph connectivity constraints at the clones of some edge, we can reverse the splitting at that edge within the embedding.

After such a splitting operation the original edge is split into a set of vertex-disjoint edges. By the dual graph connectivity constraints, there are two of these edges in a common local surface of the embedding. So we can apply Lemma 4.5.5 to identify them. Arguing inductively, we can identify them all recursively. This shows why one such splitting can be reversed. Hence we can argue inductively as in the proof of Theorem 4.5.3 to complete the proof. □

4.5.3 Embeddings induce embeddings of split complexes

The goal of this subsection is to prove the following.

Lemma 4.5.6. *Let C be a locally connected simplicial complex with an embedding ι in $\mathbb{S}^3$. Then its split complex has an embedding into $\mathbb{S}^3$ that satisfies the dual graph connectivity constraints and has the same dual graph as ι.*

Before we can prove this, we need some preparation. We start with the following lemma very similar to Lemma 4.5.5. We define 'detemined' and reveal the definition in the proof of the next lemma.

Lemma 4.5.7. *Let C be a simplicial complex obtained from a simplicial complex C' by identifying two edges e and e' that only share the vertex v. Let ι' be a topological embedding of C' into $\mathbb{S}^3$. Assume that the embedding of $L(v)$ in the plane induced by ι' has a region[25] that contains both e and e'. Then there is a topological embedding of C into $\mathbb{S}^3$ that has the same dual graph as ι'. The cyclic orientation at the new edge is determined.*

Proof. We image that the link graph at v is embedded in a small ball around v. Then the region R containing e and e' is included in a unique local surface of ι'. We call that local surface ℓ. We obtain $\bar{C}$ from C' by adding a face f at the edges e, e' and one new edge. The embedding ι induces an embedding of $\bar{C}$ as follows. We embed C' as prescribed by ι' and embed f in ℓ. It remains to specify the faces just before or just after f at e and e'. The face f' just before f at e corresponds to some edge of $L(v)$ that has the region R on its left, when directed towards e. Similarly, the face f'' just after f at e' corresponds to some edge of $L(v)$ that has the region R on its right, when directed towards e'. This embedding of $\bar{C}$ induces some embedding of C by first contracting the third edge of f, the one not equal to e or e' and then contracting the face f, that is, we identify e and e' along f. Clearly this embedding has the same dual graph as ι'.

It remains to show that the cyclic orientation of the incident faces induced by the embedding at the new edge is determined. For that we reveal the definition of determined. It means that the cyclic ordering at the new edge is obtained by concatenating the cyclic orientations of e and e' induced by ι' so that f' is followed by f''. □

For the rest of this subsection we fix a topological embedding ι of a locally connected simplicial complex C into $\mathbb{S}^3$. Our aim is to explain how ι gives rise

[25] Component of $\mathbb{S}^2$ without $L(v)$

to an embedding of any split complex of C. First we need some preparation. Let $\Sigma = (\sigma(e)|e \in E(C))$ be the combinatorial embedding induced by ι.

Let e be an edge of C and I a subinterval of $\sigma(e)$. Let $\bar{C}$ be the simplicial complex obtained from C by replacing e by two edges, one that is incident with the faces in I and the other that is incident with the faces incident with e but not in I. We call $\bar{C}$ the simplicial complex obtained from C by *opening the edge e along I*. We refer to the two new edges as the *opening clones* of e. If we apply several openings, we extend the notion of opening cloning iteratively so that each edge of the resulting simplicial complex is opening cloned from a unique edge of C.

Let C' be a simplicial complex obtained from a simplicial complex C by splitting edges. Given a rotation system Σ of C, we obtain the *induced* rotation system of C' by restricting for each e' of C' cloned from an edge e of C the cyclic ordering $\sigma(e)$ to the faces incident with e'. We define also an induced rotation system if C' is obtained from C by opening edges. This is as above with 'clone' replaced by 'opening clone'.

Let $\bar{C}$ be a simplicial complex obtained from C by opening an edge and let $\bar{\Sigma}$ be the rotation system induced by Σ.

Lemma 4.5.8. *The simplicial complex $\bar{C}$ has a topological embedding $\bar{\iota}$ into $\mathbb{S}^3$ whose induced planar rotation system is $\bar{\Sigma}$.*

The dual graph of $\bar{\iota}$ is obtained from the dual graph G of ι by identifying the two endvertices of I when considered as a trail in G.

In particular, if I is a closed trail in G, then G is the dual graph of $\bar{\iota}$.

Proof of Lemma 4.5.8. We can modify the embedding of C such that there is an open cylinder around e that does not intersect any edge except for e or any face not incident with e. And all faces in I intersect that cylinder only in the left half of the cylinder and the others only in the right half. Now we replace e by two copies - one in the left half, the other in the right half. It is straightforward to check that the dual graph of the embedding has the desired property. $\square$

We fix an edge e of C with endvertices v and w.

Lemma 4.5.9. *There is an embedding ι' of C in $\mathbb{S}^3$ that has the same dual graph as ι such that there is some connected component X at e that is a subinterval of the cyclic orientation $\sigma'(e)$, where $\Sigma' = (\sigma'(e)|e \in E(C))$ is the induced rotation system of ι'.*

Example 4.5.10. The following example demonstrates that in Lemma 4.5.9 we cannot always pick $\iota' = \iota$. In the embedding in 3-space indicated in Figure 4.4 no component at the edge e is a subinterval of the cyclic orientation of the faces incident with e induced by the embedding.

Before we can prove Lemma 4.5.9, we need some preparation.

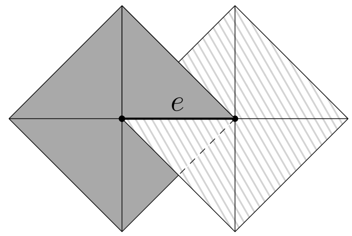

Figure 4.4: This complex is obtained by gluing together two discs, each with four faces, at the edge e.

Given a cyclic orientation σ and a subset X, we say that two elements y_1 and y_2 of σ *separate* X for σ if they are both not in X and the two intervals[26] $y_1\sigma y_2$ and $y_2\sigma y_1$ both contain elements of X.

Lemma 4.5.11. *Let σ be a cyclic orientation and $(P_i|i \in [n])$ be a partition of the elements of σ such that no two elements of the same P_i separate some other P_j. Then there is some P_k that is a subinterval of σ.*

Proof. We pick an arbitrary element a of P_1. We may assume that a partition class P_2 exists. For any P_i not containing a, we define its *first element* to first element of P_i after a in σ, and its *last element* to first element of P_i before a in σ. The *closure* of P_i consists of those elements of σ between its first and last element (including the first and the last one). We denote the closure of P_i by $\overline{P_i}$.

By assumption any two such closures $\overline{P_i}$ and $\overline{P_j}$ are either disjoint or contained in one another, that is, $\overline{P_i} \subseteq \overline{P_j}$ or vice versa. Let P_k be such that its closure is inclusion-wise minimal. Then P_k is equal to its closure and hence a subinterval of σ. $\square$

Given $e \in \sigma$, we denote the element just before e by $e - 1$ and the element just after e by $e + 1$. Given a cyclic orientation σ and four of its elements x_1, x_2, x_3, x_4 such that $(x_1x_2x_3x_4)$ is a cyclic subordering of σ, the *exchange* of σ with respect to x_1, x_2, x_3, x_4 is the following cyclic orientation on the same elements as σ. We concatenate the two cyclic orientations obtained from σ by deleting $x_1\sigma x_3 - x_1 - x_3$ and $x_3\sigma x_1$ such that the immediate successor of x_4 is x_2; see Figure 4.5, formally, it is

$$x_3\sigma x_4[x_2\sigma(x_3 - 1)][(x_1 + 1)\sigma(x_2 - 1)][(x_4 + 1)\sigma x_1]x_3$$

Let $(P_i|i \in I)$ be a partition of the elements of σ, the *fluctuation* of σ with respect to $(P_i|i \in I)$ is the number of adjacent elements of σ in different P_i. Given a partition $\mathcal{P} = (P_i|i \in I)$ of σ, an exchange is *$\mathcal{P}$-improving* if x_2 and x_4 are in the same P_i but none of the following four pairs is in the same P_i: $(x_4, x_4 + 1)$, $(x_2, x_2 - 1)$, $(x_1, x_1 + 1)$, $(x_3, x_3 - 1)$.

[26]By $y_1\sigma y_2$ we denote the subinterval of σ starting at y_1 and ending with y_2.

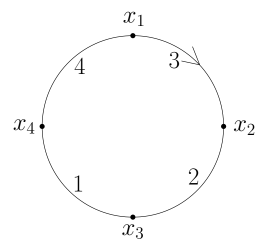

Figure 4.5: The cyclic orientation σ is depicted as a cycle. The four segments between the element x_i are labelled with the elements of $\mathbb{Z}_4$. This describes the ordering in which these segments are travered by the exchanged cyclic orientation.

Lemma 4.5.12. *A cyclic orientation σ' obtained from σ by an exchange that is $\mathcal{P}$-improving has strictly smaller fluctuation.*

Proof. The adjacent elements of σ and σ' are the same except for four pairs involving x_1, x_2, x_3, x_4. For σ these pairs are those mentioned in the definition of '$\mathcal{P}$-improving'. All these four pairs contribute to the fluctuation by the definition of $\mathcal{P}$-improving. For σ' the pair (x_4, x_2) does not contribute to the fluctuation. □

One way to partition $\sigma(e)$ is to put two elements of $\sigma(e)$ in the same class if – when considered as edges of $L(v)$ – they have endvertices in the same component of $L(v) - e$. An exchange is *v-improving* for $\sigma(e)$ if it is $\mathcal{P}$-improving for that particular partition.

For the next lemma we fix the following notation. Let X be the set of edges between e and a connected component of $L(v) - e$. Let Y be the set of edges between e and a connected component of $L(w) - e$. Assume that no connected component at e includes both X and Y.

Lemma 4.5.13. *Assume that two elements of Y separate X in the cyclic orientation $\sigma(e)$. Then there is an embedding ι' of C in $\mathbb{S}^3$ that has the same dual graph as ι such that $\sigma'(e)$ is obtained from $\sigma(e)$ by a v-improving exchange, where $\Sigma' = (\sigma'(e) | e \in E(C))$ is the induced rotation system of ι'.*

Proof. We abbreviate $\sigma(e)$ by σ. We denote the connected component[27] at e including Y by $c(Y)$.

Sublemma 4.5.14. *There are edges f_1 and f_3 of $c(Y)$ that separate X such that the region of $L(w)$ just after f_1 is equal to the region just before f_3. And $f_1 + 1$ and $f_3 - 1$ are not in $c(Y)$.*

Proof. Let f_1' and f_3' be two elements of Y that separate X. We fix two elements x_1 and x_2 of X such that x_1 is in $f_1'\sigma f_3'$ and x_2 is in $f_3'\sigma f_1'$. By choosing f_1' and

[27]See Subsection 4.3.2 above for a definition.

f_3' as near to x_1 as possible, we ensure that the region just after f_1' is equal to the region just before f_3'. We denote this region by R. Let Y' be the set of edges between e and a connected component of $L(w)-e$ that is included in $c(Y)$. The set of all such Y' is denoted by $\mathcal{Y}$. By replacing Y by any $Y' \in \mathcal{Y}$ if necessary, we may assume that no set $Y' \in \mathcal{Y}$ contains elements both before and after x_1 on $f_1'\sigma f_3'$; indeed, by any such replacement $f_1'\sigma f_3'$ strictly decreases.

Sublemma 4.5.15. *The interval $f_1'\sigma x_1$ contains some $f_1 \in c(Y)$ such that the region just after f_1 is R and $f_1 + 1$ is not in $c(Y)$.*

Proof. We recursively define a sequence f_1^n of elements of $f_1'\sigma x_1$. They are strictly increasing and contained in $c(Y)$. We start with $f_1^1 = f_1'$. Assume that we already constructed f_1^n. If $f_1^n + 1$ is not in $c(Y)$ we stop and let $f_1 = f_1^n$. Otherwise $f_1^n + 1$ is in $c(Y)$. Let $Y' \in \mathcal{Y}$ so that $f_1^n + 1 \in Y'$.

We prove inductively during this construction that any set $Y'' \in \mathcal{Y}$ that contains an element of $(f_1^n + 1)\sigma x_1$ contains no element of $f_1'\sigma f_1^n$.

By the induction hypothesis, Y' is a subset of $(f_1^n+1)\sigma x_1$. Let f_1^{n+1} be the maximal element of Y' in $(f_1^n+1)\sigma x_1$. By construction $f_1^{n+1} \in c(Y)$ and f_1^{n+1} is strictly larger than f_1^n. The region R is just before $f_1^n + 1$, the first element of Y'. Thus the region after f_1^{n+1}, the last element of Y', must also be R. The induction step follows from the planarity of $L(w)$ as there is a component of $L(w) - e$ that is adjacent to the set Y', and the induction hypothesis.

This process has to stop as $f_1'\sigma x_1$ is finite and the f_1^n are strictly increasing. Thus we eventually find an f_1. □

Similarly as Sublemma 4.5.15 one shows that the interval $x_1\sigma f_3'$ contains some $f_3 \in c(Y)$ such that the region just before f_3 is R and $f_3 - 1$ is not in $c(Y)$. So f_1 and f_3 have the desired properties. □

We obtain C_1 from C by opening the edge e at the subinterval $f_1\sigma f_3$ of σ. By ι_1 we denote the embedding of C_1 induced by ι. By the choice of f_1 and f_3, the local surface just after f_1 is equal to the local surface just before f_3. Hence by Lemma 4.5.8 the embeddings ι_1 and ι have the same dual graph.

By Sublemma 4.5.14, the link graph at w of C_1 has two connected components. We obtain C_2 from C_1 by splitting the vertex w. By ι_2 we denote the embedding of C_2 induced by ι_1. As splitting vertices does not change the dual graph by Observation 4.3.9, the embeddings ι_2 and ι_1 have the same dual graph. Summing up, ι_2 and ι have the same dual graph.

We denote the copy of e incident with f_1 by e' and the other copy by e''. Since e' and e'' are both incident with edges of X, the component of $L(v) - e$ adjacent to the edges of X has in the link graph of C_2 the two vertices e' and e'' in the neighbourhood. Thus the vertices e' and e'' share a face in the link graph at v of C_2.

By Lemma 4.5.7 ι_2 induces an embedding ι' of C in $\mathbb{S}^3$ that has the same dual graph as ι_2. Let $\Sigma' = (\sigma'(e)|e \in E(C))$ be the induced rotation system of ι'. We denote the element of X in $f_1\sigma f_3$ nearest to f_1 by f_2. Similarly, by f_4 we denote the element of X in $f_3\sigma f_1$ nearest to f_1. As $\sigma'(e)$ is determined by

Lemma 4.5.7, it is obtained by concatenating the cyclic orientations at e' and e'' so that f_4 is followed by f_2. That is, $\sigma'(e)$ is obtained from $\sigma(e)$ by exchanging with respect to f_1, f_2, f_3, f_4.

It remains to check that this exchange is v-improving. Both f_2 and f_4 are in X. On the other hand f_1 and f_3 are in $c(Y)$ but $f_1 + 1$ and $f_3 - 1$ are not in $c(Y)$. In particular, they are in different P_i. Whilst f_2 and f_4 are in X, the two elements $f_2 - 1$ and $f_4 + 1$ are not in X. Thus this exchange is v-improving. □

Proof of Lemma 4.5.9. By $(R_k | k \in K)$ we denote the partition of the faces incident with e into the connected components at e. If no two elements of the same R_a separate some other R_b, then by Lemma 4.5.11 there is some R_a that is a subinterval of $\sigma(e)$. In this case we can just pick $\iota' = \iota$ and are done.

We define the partition $(P_i | i \in I)$ of the faces incident with e as follows. Two faces incident with e are in the same partition if – when considered as edges of $L(v)$ – they have endvertices in the same component of $L(v) - e$. We define the partition $(Q_j | j \in J)$ the same with 'w' in place of 'v'. If some P_i contains two elements separating some Q_j for the cyclic orientation at e, we can apply Lemma 4.5.13 to construct a new embedding of C. We do this until there are no longer such pairs (P_i, Q_j). This has to stop after finitely many steps as by Lemma 4.5.12 the fluctuation – which is a non-negative constant only defined in terms of $(P_i | i \in I)$ – of the cyclic orientation at e strictly decreases in each step. So there is an embedding ι' of C in $\mathbb{S}^3$ such that no P_i contains two elements separating some Q_j for the cyclic orientation $\sigma'(e)$ and such that ι' has the same dual graph as ι; here we denote by $\Sigma' = (\sigma'(e) | e \in E(C))$ is the induced rotation system of ι'. Hence by applying Lemma 4.5.11, it suffices to prove the following.

Sublemma 4.5.16. *For $\sigma'(e)$, either there is some P_i containing two elements separating some Q_j or no two elements of the same R_a separate some other R_b.*

Proof. We assume that there is some R_a that contains two elements r_1 and r_2 that separate some other R_b. The set R_b is a disjoint union of sets P_i. Either r_1 and r_2 separate one of these P_i or by the definition of connected component at e, there is some Q_j included in R_b that contains elements of different P_i, one included in $r_1 \sigma'(e) r_2$ and the other in $r_2 \sigma'(e) r_1$. Summing up there is some P_i or Q_j included in R_b that is separated by r_1 and r_2.

First we consider the case that there is a set P_i. So two elements of that set P_i separate R_a. By an argument as above we conclude that there is some P_m or Q_n included in R_a that is separated by two elements of P_i.

Since the sets P_m are defined from components of $L(v) - e$ and Σ' induces an embedding of $L(v)$ in the plane, these components cannot attach at e in a 'crossing way', that no two elements of some P_i can separate some other P_m. Thus there has to be such a set Q_n.

Summing up, if there is a set P_i separated by r_1 and r_2, then it contains two elements separating some Q_n. Analogously one shows that otherwise the set Q_j separated by r_1 and r_2 contains two elements separating some P_n. But then two elements of P_n separate Q_j. This completes the proof. □

By Sublemma 4.5.16 and the construction of ι', no two elements of the same R_a separate some other R_b for $\sigma'(e)$. Then by Lemma 4.5.11 there is some R_a that is a subinterval of $\sigma'(e)$, as desired. □

Let C' be a simplicial complex obtained from the locally connected simplicial complex C by splitting the edge e.

Lemma 4.5.17. *There is a topological embedding ι' of C' whose induced planar rotation system is the rotation system induced by Σ.*

Moreover ι and ι' have the same dual graph.

Proof. We denote the dual graph of ι by G. We prove this lemma by induction on the number of connected components at e. If there is only one such component, then $C' = C$ and the lemma is trivially true. So we may assume that there are at least two components. By changing the embedding if necessary, by Lemma 4.5.9 we may assume that there is a component J at e that is a subinterval of $\sigma(e)$. As J is a subinterval of the closed trail $\sigma(e)$ of G, it is a trail in G. Next we show that it is a closed one:

Sublemma 4.5.18. *The interval J is a closed trail in G.*

Proof. We are to show that the local surface of the embedding just before the first face f_1 of J is the same as the local surface just after the last edge f_2 of J. For that it suffices to show that in the embedding of the link graph $L(v)$ of v induced by Σ, the region just before the edge f_1 is the same as the region just after the edge f_2. This follows from the fact that J is the set of edges out of a set of connected components of $L(v) - e$. Indeed, the first and last edge out of every component are always in the same region. □

We obtain $\bar{C}$ from C by opening the edge e along J. By Lemma 4.5.8, $\bar{C}$ has a topological embedding $\bar{\iota}$ into $\mathbb{S}^3$ whose induced planar rotation system is induced by Σ. By Sublemma 4.5.18 and Lemma 4.5.8, the dual graph of $\bar{\iota}$ is G.

We observe that C' is obtained from $\bar{C}$ by splitting the clone of e that corresponds to the subinterval $\sigma(e) \setminus J$. Thus the lemma follows by applying induction on $\bar{C}$ and $\bar{\iota}$. □

Proof of Lemma 4.5.6. The split complex of C is obtained from C by a sequence of edge splittings and vertex splittings. By changing the order of the splittings if necessary, we may assume that the complex is always locally connected before we perform an edge splitting. Hence we can apply Lemma 4.5.17 and Theorem 4.5.3 recursively to construct an embedding of the split complex. Since in each splitting step the dual graph is preserved, it satisfies the dual graph connectivity constraints by Lemma 4.5.2 applied to the dual graph of ι. □

4.5.4 Proof of Theorem 4.2.4

We summarise the results of the earlier subsections in the following.

Theorem 4.5.19. *Let C be a simplicial complex and $\hat{C}$ be its split complex. Then C embeds into $\mathbb{S}^3$ if and only if $\hat{C}$ has an embedding into $\mathbb{S}^3$ that satisfies the dual graph connectivity constraints.*

Proof. Assume that C embeds into $\mathbb{S}^3$. Then by Theorem 4.5.3 its vertical split complex embeds into $\mathbb{S}^3$ and satisfies the vertical graph connectivity constraints. Since the vertical split complex is locally connected, we can apply Lemma 4.5.6 to get the desired embedding of the split complex. Note that this embedding has the same dual graph as the vertical split complex. Hence it also satisfies the connectivity constraints for the vertices.

Now conversely assume that the split complex has an embedding ι' that satisfies the dual graph connectivity constraints. By Lemma 4.5.4 the vertical split complex has an embedding in $\mathbb{S}^3$. As this embedding has the same dual graph as ι', it satisfies the vertical dual graph connectivity constraints. So we can apply Theorem 4.5.3. This completes the proof. □

Now we show how Theorem 4.5.19 implies Theorem 4.2.4.

Proof of Theorem 4.2.4. Let C be a globally 3-connected simplicial complex and let $\hat{C}$ be its split complex. If C embeds into $\mathbb{S}^3$, then $\hat{C}$ has an embedding into $\mathbb{S}^3$ whose dual graph G satisfies the dual graph connectivity constraints by Theorem 4.5.19. By Corollary 4.3.14, the two simplicial complexes C and $\hat{C}$ have the same dual matroid. So by Theorem 4.2.1 the cycle matroid of G is the dual matroid of C. This completes the proof of the 'only if'-implication.

Conversely assume that a split complex $\hat{C}$ of a simplicial complex C has an embedding $\hat{\iota}$ into $\mathbb{S}^3$ and the dual matroid M of C is the cycle matroid of a graph G and the set of faces incident with any vertex or edge of C is a connected edge set of G. By Corollary 4.3.14 M is the dual matroid of $\hat{C}$. Let G' be the dual graph of the embedding $\hat{\iota}$ of $\hat{C}$. By Theorem 4.2.1 the cycle matroid of G' is equal to M. Since M is 3-connected by assumption, by a theorem of Whitney [100], the graphs G and G' are identical. Hence G' satisfies the connectivity constraints. So we can apply the 'if'-implication of Theorem 4.5.19 to deduce the 'if'-implication of Theorem 4.2.4. □

Proof of Theorem 4.2.5. By Chapter 1, it suffices to show that a simplicial complex C whose split complex is embeddable has an embedding if and only if its dual matroid has no constraint minor in the list of Figure 4.2. Since the split complex is embeddable, its dual matroid is the cycle matroid of a graph G. By Corollary 4.3.14 the dual matroid of C is the cycle matroid of G. By Theorem 4.2.4, C is embeddable if and only if G satisfies the graph connectivity constraints. The later is true if and only if there is no vertex or edge such that the set X of incident faces is disconnected in G. By the main result of Chapter 3, X is disconnected in G if and only if (G, X) has a constraint minor in the list of Figure 4.2. □

4.6 Infinitely many obstructions to embeddability into 3-space

In this section we construct an infinite sequence $(A_n|n \in \mathbb{N})$ of minimal obstructions to embeddability. More precisely, A_n will have the property that its split complex is simply connected and embeddable, its dual matroid M_n is the cycle matroid of a graph but no such graph will satisfy the connectivity constraints. However, if we remove a constraint or contract or delete an element from the dual matroid, then there is such a graph.

The dual matroid M_n of A_n will be the disjoint union of a cycle C_n of length n and a loop ℓ, see Figure 4.6.

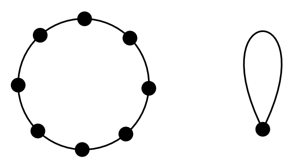

Figure 4.6: The matroid M_8. For each of the eight vertices on the cycle, there is a connectivity constraint forbidding that the loop is attached at that vertex.

The connectivity constraints are as follows. Fix a cyclic orientation $\{e_i|i \in \mathbb{Z}_i\}$ of the edges on C_n. We have a connectivity constraint for every $i \in [n]$, namely that $X[i,n] = C_n - e_i - e_{i+1} + \ell$ is a connected set.

Fact 4.6.1. *There is no graph whose cycle matroid is M_n that meets all the connectivity constraints $X[i,n]$.*

Proof. By $\overline{C_n}$ we denote the graph that is a cycle of length n whose edges have the cyclic ordering $\{e_i|i \in \mathbb{Z}_i\}$. It is straightforward to see that $\overline{C_n}$ is the unique graph whose cycle matroid is C_n that meets all the connectivity constraints $X[i,n] - e$.

Now suppose for a contradiction that there is a graph G whose cycle matroid is M_n that meets all the connectivity constraints $X[i,n]$. Then G is obtained from $\overline{C_n}$ by attaching a loop. Since each $X[i,n]$ contains e, we have to attach the loop at some vertex of $\overline{C_n}$. The connectivity constraint $X[i,n]$, however, forbids us to attach the loop at the vertex incident with e_i and e_{i+1}. Hence G does not exist. $\square$

A careful analysis of this proof yields the following simple facts.

Fact 4.6.2. *1. There is a graph whose cycle matroid is M_n that meets all the connectivity constraints $X[i,n]$ but one.*

2. for every element e, there is a graph whose cycle matroid is $M_n - e$ that meets all the connectivity constraints $X[i,n] - e$;

3. *for every element e, there is a graph whose cycle matroid is M_n/e that meets all the connectivity constraints $X[i,n] - e$.*

□

Hence it remains to construct A_n such that its dual matroids is M_n and so that the nontrivial connectivity constraints are the $X[i,n]$. We remark that we allow the faces of A_n to be arbitrary closed walks. (One obtains a simplicial complexes from A_n by applying baricentric subdivisions to the faces.)

We start the construction of A_n with a cycle C of length n. We attach n faces, which we call $e_1, ..., e_n$. For each e_i, and each vertex v_k of C except for the i-th vertex v_i, we attach $n-1$ edges and let e_i traverse them in between the two edges incident with v_k. We denote the endvertices of the new edges not on C by $x(i,k,j)$ where $(k, j \leq n; k, j \neq i)$, see Figure 4.7.

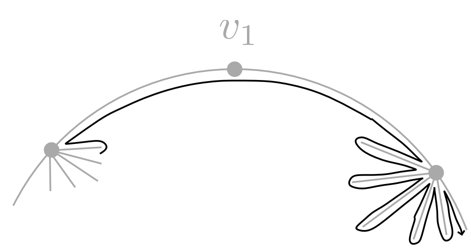

Figure 4.7: In grey, we indicate the cycle C with the new edges. In black we sketched the traversal of the face e_1 after addition of the new edges.

Next we disjointly add a copy of the original cycle C and only attach a single face to it which we denote by ℓ. Call the resulting walk-complex[28] A'_n. We finally obtain A_n from A'_n by identifying for each $i \in [n]$ the $i-th$ vertex v_i on the new copy of C with all vertices $x(i', i, i)$ with $i' \neq i$.

By construction, the split complex of A_n is A'_n. Hence by Corollary 4.3.14 above, the dual matroid of A_n is M_n. By construction, the nontrivial connectivity constraints are the $X[i,n]$. Clearly, the split complex A'_n is simply connected and embeddable.

This completes the construction of the A_n. By Fact 4.6.1 and Fact 4.6.2 they have the desired properties.

4.7 Appendix I

First we give a definition of 'globally 3-connected' directly in terms of the simplicial complex without referring to its dual matroid. Given a simplicial complex C, its edge/face incidence matrix A and a subset L of the faces of C, we denote by $r(L)$ the rank over $\mathbb{F}_3$ of the submatrix of A induced by the vectors

[28]A *walk-complex* is a graph together with a family of closed walks, which we call its faces. Every simplicial complex is a walk-complex. Conversely, from every walk complex we can build a simplicial complex by attaching at each face a cone over that walk.

whose faces are in L. A *2-separation* of a simplicial complex C is a partition of its set F of faces into two sets L and R both of size at least two such that $r(L) + r(R) \leq r(F) + 1$. It is straightforward it check that a simplicial complex is globally 3-connected if and only if it has no 2-separation.

When defining 'edge split complexes', we mentioned a related more naive definition. Here we give this definition. In Example 4.7.1 and Example 4.7.2 we show that this notion lacks two important features of edge split complexes. *Splitting an edge e at an endvertex v* is defined like 'splitting e' but with 'in the same connected component at e' replaced by 'v-related'. A *lazy edge split complex* is defined as 'edge split complex' but with 'for every edge there is only one component at e' replaced by 'it is locally 2-connected'. *lazy split complex* is defined like 'split complex' with 'lazy edge split complex' in place of 'edge split complex'.

Example 4.7.1. In this example we construct a simplicial complex C that has two distinct lazy edge split complexes. We will construct C such that it has two vertices v and w; these vertices are joined by five edges e, e_1, e_2, e_3 and e_4. The edge e is a cut vertex in the link graphs at v and w. And splitting e at one endvertex will make the link graph at the other endvertex 2-connected, see Figure 4.8.

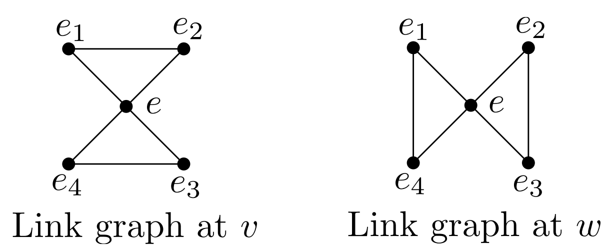

Figure 4.8: If we split one of these link graphs at e, the other becomes a six-cycle.

Next we construct C with the above properties. We obtained C from four triangular faces f_1, f_2, f_3 and f_4 glued together at a single edge e. Let v and w be the two endvertices of that edge. Let $e_i[v]$ be the edge of f_i incident with v different from e. Let $e_i[w]$ be the edge of f_i incident with w different from e. Let v_i be the vertex incident with f_i that is not incident with e. We add the edges e_k between v_k and v_{k+1} for any $k \in \mathbb{Z}_4$. We add the four faces: $e_1[v]e_1e_2[v]$, $e_3[v]e_3e_4[v]$, $e_2[w]e_3e_3[w]$ and $e_4[w]e_4e_1[w]$. This completes the construction of C.

Example 4.7.2. In this example we show that Theorem 4.5.19 with 'split complex' replaced by 'lasy split complex' is false. Let H be a planar graph with vertices v and w such that the graph H' obtained from H by identifying the vertices v and w is not planar. Let C be the cone over H. We obtain C' from C by identifying the two edges corresponding to v and w. Whilst the link at the top of C is H, the complex C' has the link H' and is hence not embeddable. By choosing v and w far apart in H, one ensures that C' is a simplicial complex.

The lazy split complex of C' is unique and equal to C. Unlike C', the simplicial complex C is embeddable. The dual graph of every embedding consists of a single vertex, and so trivially satisfies the graph connectivity constraints. This completes the example.

Concerning Theorem 4.2.4, it is straightforward to modify the example to make the dual graph of the embedding 3-connected.

4.8 Appendix II: Matrices representing matroids over the integers

Matroids representable over the integers are well-studied [76]. In this appendix, we study something very related but slightly different, namely matrices that represent matroids over the integers. Our aim in this appendix is to prove Theorem 4.8.6 below, which is a characterisation of certain matrices representing matroids over the integers.

A matrix A is a *representation of a matroid M over a field k* if all its entries are in k and the columns are indexed with the elements of M. Furthermore for every circuit o of M there is a vector v_o in the span over k of the rows of A whose support is o. And the vectors v_o span over k all row vectors of A.

The following is well-known.

Lemma 4.8.1. *Let A be a matrix representing a matroid M over some field k. Let I an element set that is independent in M. Then the matrix obtained from A by deleting all columns belonging to elements of I represents the matroid M/I over k.* □

A matrix A is a *regular representation* (or *representation over the integers*) of a matroid M if all its entries are integers and the columns are indexed with the elements of M. Furthermore for every circuit o of M there is a $\{0, -1, +1\}$-valued vector[29] v_o in the span over $\mathbb{Z}$ of the rows of A whose support is o. And the vectors v_o span over $\mathbb{Z}$ all row vectors of A. The following is well-known.

Lemma 4.8.2. *Assume that a matrix A regularly represents a matroid M. Then for every cocircuit d of M, there is a $\{0, -1, +1\}$-valued vector w_d whose support is equal to d that is orthogonal*[30] *over $\mathbb{Z}$ to all row vectors of A. These vectors w_d generate over $\mathbb{Z}$ all vectors that are orthogonal over $\mathbb{Z}$ to every row vector.* □

The following is well-known.

Lemma 4.8.3. *Let M be a matroid regularly represented by a matrix A. Let v be a sum of row vectors of A with integer coefficients. If the support of v is nonempty, then it includes a circuit of M.* □

[29] A *vector* is an element of a vector space k^S, where k is a field and S is a set. In a slight abuse of notation, in this chapter we also call elements of modules of the form $\mathbb{Z}^S$ vectors.

[30] Two vectors a and b in k^S are *orthogonal* if $\sum_{s \in S} a(s) \cdot b(s)$ is identically zero over k.

Example 4.8.4. A matrix is *unimodular* if it is $\{0, -1, +1\}$-valued and the determinant of every quadratic submatrix is $\{0, -1, +1\}$-valued[31]. Every unimodular matrix is a regular representation of some matroid, see for example [91]. For example, the vertex/edge incidence matrix of a graph G is a regular representation of the graphic matroid of G.

There also exist regular representations that are not totally unimodular:

Example 4.8.5.

$$\begin{pmatrix} 1 & 1 \\ 1 & -1 \\ 1 & 0 \\ 0 & -1 \end{pmatrix}$$

This matrix is a regular representation of the matroid consisting of two elements in parallel but it is not totally unimodular.

A matroid is *regular* if it can be regularly represented by some matrix. The class of regular matroids has many equivalent characterisations [76]. For example, a matroid has a regular representation (in fact a totally unimodular one) if and only if it has a representation over every field. In this chapter, we need the following related fact, which focuses on the matrices instead of the matroids:

Theorem 4.8.6. *Let A be a matrix whose entries are -1, $+1$ or 0. Then A regularly represents a matroid if and only if there is a single matroid M such that A represents M over any field.*

Whilst the 'only if'-implication is immediate, the other implication is less obvious. To prove it we rely on the following.

Lemma 4.8.7. *Let $(v_i|i \in I)$ be a family of integer valued vectors of $\mathbb{Z}^S$, where S is a finite set. Assume that the family $(v_i|i \in I)$ considered as vectors of the vector space $\mathbb{Q}^S$ spans the whole of $\mathbb{Q}^S$ over $\mathbb{Q}$. Additionally, assume that for every prime number p, the same assumption is true with the finite field '$\mathbb{F}_p$' in place of '$\mathbb{Q}$'. Then the family $(v_i|i \in I)$ spans over $\mathbb{Z}$ all integer valued vectors in $\mathbb{Z}^S$.*

Proof that Lemma 4.8.7 implies Theorem 4.8.6. Assume that A is an integer valued matrix that represents the matroid M over $\mathbb{Q}$ and over all finite fields $\mathbb{F}_p$ for every prime number p, when we interpret[32] the entries of A as elements of the appropriate field. Our aim is to show that A regularly represents the matroid M.

Let b be a base of M. Let A' be the matrix obtained from A by deleting all columns belonging to elements of b. We denote by M' the matroid M/b, in which every element is a loop. By Lemma 4.8.1, A' represents the matroid M' over $\mathbb{Q}$ and over all finite fields $\mathbb{F}_p$. Let $(v_i|i \in I)$ be the family of row vectors of

[31] Here we evaluate the determinate over $\mathbb{Z}$

[32] Here in $\mathbb{F}_p$ we interpret the integer m as its remainder after division by p.

A'. Since every element of M' is a loop, we can apply Lemma 4.8.7 and deduce that the family $(v_i|i \in I)$ spans over $\mathbb{Z}$ all integer valued vectors in $\mathbb{Z}^{E'}$, where E' is the set of elements of M'.

Let v be any integer valued vector that is generated by the rows of A over $\mathbb{Q}$. We show that v is also generated by the rows of A with integer coefficients. By the above, there is a vector w generated from the row vectors of A over $\mathbb{Z}$ that agrees with v in all coordinates of E'. Hence $v - w$ is generated by the row vectors over $\mathbb{Q}$. So if $v - w$ is nonzero, its support must contain a circuit of M by Lemma 4.8.3. Since the support of $v - w$ is contained in the base b, the support does not contains a circuit of M. Hence v must be equal to w. Thus v is in the span of the row vectors with coefficients in $\mathbb{Z}$.

Now let o be a circuit of M. Since A is a regular representation of M over $\mathbb{Q}$, there is a vector v_o with entries in $\mathbb{Q}$ generated by the row vectors of A over $\mathbb{Q}$ whose support is o. We multiplying all entries with a suitable rational number if necessary, we may assume that additionally all entries of v_o are integers and that the greatest common divisor of the entries is one. By the above v_o is in the span of the row vectors with coefficients in $\mathbb{Z}$.

Next we show that all entries of v_o are zero, plus one, or minus one. Suppose for a contradiction that there is some prime number p that divides some entry of v_o. If we interpret the entries of v_o as elements of $\mathbb{F}_p$, then v_o is also in the span of the row vectors with coefficients in $\mathbb{F}_p$. Indeed, the coefficients are just the integer coefficients we have in the representation over $\mathbb{Z}$ interpreted as elements of $\mathbb{F}_p$. Since the greatest common divisor of the entries of v_o is one, v_o when interpreted over $\mathbb{F}_p$ is nonzero but its support is properly contained in o. Since in M the circuit o does not include another circuit, we get a contraction to the assumption that A represents M over $\mathbb{F}_p$. Thus all entries of v_o are zero, plus one, or minus one.

It remains to show that the set of vectors v_o where o is a fundamental circuit of b generates every row vector x of A. Since for every element not in b, there is a unique v_o which takes the value plus one or minus one at that element and zero at every other elements not in b, there is a vector x' generated over $\mathbb{Z}$ by the v_o that agrees with x when restricted to E'. As above we deduce that $x' = x$, and hence x is generated by the v_o over $\mathbb{Z}$. Thus A regularly represents M. □

In order to prove Lemma 4.8.7, we rely on the following well-known lemma.

Lemma 4.8.8. *Let m and n be integer and let d be their greatest common divisor. Then there are integers α and β such that $\alpha \cdot m - \beta \cdot n = d$.* □

Proof of Lemma 4.8.7. Let $s \in S$ be arbitrary. By e_s we denote the vector which in coordinate s has the entry one and otherwise the entry zero. Since the family $(v_i|i \in I)$ spans e_s over $\mathbb{Q}$, there is some positive natural number γ_s so that the family $(v_i|i \in I)$ spans $\gamma_s \cdot e_s$ over $\mathbb{Z}$. Let δ_s be the least possible value for γ_s. Our aim is to show that all δ_s are equal to one. Suppose not for a contradiction. Then there is some prime number p that divides some δ_s. Let $\bar{s}$ be the index so that in the factorisation of $\delta_{\bar{s}}$ the prime number p has the highest multiplicity, say k.

Sublemma 4.8.9. *There is some nonzero integer ϵ such that p has the multiplicity at most $k-1$ in the factorisation of ϵ and such that $\epsilon \cdot e_{\bar{s}}$ is spanned by the family $(v_i|i \in I)$ over $\mathbb{Z}$.*

Let us first see how we finish the proof assuming Sublemma 4.8.9. By Lemma 4.8.8, there are α and β such that $\alpha \cdot \delta_{\bar{s}} - \beta \cdot \epsilon$ is equal to the greatest common divisor D of $\delta_{\bar{s}}$ and ϵ. Hence by Sublemma 4.8.9 $D \cdot e_{\bar{s}}$ is generated by the family $(v_i|i \in I)$ over $\mathbb{Z}$. Since p has the multiplicity at most $k-1$ in the factorisation of D, the number D is strictly smaller than $\delta_{\bar{s}}$. This contradicts the choice of $\delta_{\bar{s}}$. Hence all δ_s are equal to one. It remains so show that the following.

Proof of Sublemma 4.8.9. Since the family $(v_i|i \in I)$ spans $e_{\bar{s}}$ over $\mathbb{F}_p$, there is an integer valued vector w such that the family $(v_i|i \in I)$ spans $e_{\bar{s}} + p \cdot w$ over $\mathbb{Z}$. For a subset T of S we denote by w_T the vector which takes the value $w(s)$ in coordinate s if $s \in T$ and zero otherwise. We denote the multiplicity of p in the factorisation of an integer n by $\sharp_p(n)$.

We shall show inductively for every subset T of S that there is some nonzero natural number ϵ_T with $\sharp_p(\epsilon_T) \leq k-1$ such that $\epsilon_T \cdot (e_{\bar{s}} + p \cdot w_T)$ is spanned by the family $(v_i|i \in I)$ over $\mathbb{Z}$. We start the induction with $T = S$ and $\epsilon_T = 1$ and so $w_T = w$. Assume that we already proved the induction hypothesis for a nonempty subset T of S. Let $t \in T$ be arbitrary. We denote the greatest common divisor of $\epsilon_T \cdot p \cdot w(t)$ and δ_t by d_t. We let $\epsilon_{T-t} = \epsilon_T \cdot \frac{\delta_t}{d_t}$. We have

$$\sharp_p(\epsilon_{T-t}) = \sharp_p(\epsilon_T) + \sharp_p(\delta_t) - \sharp_p(d_t) \leq \sharp_p(\epsilon_T) + \sharp_p(\delta_t) - min\{\sharp_p(\epsilon_T)+1, \sharp_p(\delta_t)\} =$$

$$= max\{\sharp_p(\delta_t), \sharp_p(\epsilon_T) - 1\}$$

Hence by the choice of $\bar{t}$ and by induction $\sharp_p(\epsilon_{T-t}) \leq k-1$. Furthermore:

$$\frac{\delta_t}{d_t} \cdot \epsilon_T \cdot (e_{\bar{t}} + p \cdot w_T) - \frac{\epsilon_T \cdot p \cdot w(t)}{d_t} \cdot \delta_t e_t = \epsilon_{T-t} \cdot (e_{\bar{t}} + p \cdot w_{T-t})$$

Note that all fractions in the above equation are integers. This completes the induction step. Hence the vector $\epsilon_\emptyset \cdot e_{\bar{t}}$ is spanned by the family $(v_i|i \in I)$ over $\mathbb{Z}$, which completes the proof. □

□

Chapter 5

A refined Kuratowski-type characterisation

5.1 Abstract

Building on earlier chapters, we prove an analogue of Kuratowski's characterisation of graph planarity for three dimensions.

More precisely, a simply connected 2-dimensional simplicial complex embeds in 3-space if and only if it has no obstruction from an explicit list. This list of obstructions is finite except for one infinite family.

5.2 Introduction

We assume that the reader is familiar with Chapter 1. In that chapter we prove that a locally 3-connected simply connected 2-dimensional simplicial complex has a topological embedding into 3-space if and only if it has no space minor from a finite explicit list $\mathcal{Z}$ of obstructions. The purpose of this chapter is to extend that theorem beyond locally 3-connected (2-dimensional) simplicial complexes to simply connected simplicial complexes in general.

The first question one might ask in this direction is whether the assumption of local 3-connectedness could simply be dropped from the result of Chapter 1. Unfortunately this is not true. One new obstruction can be constructed from the Möbius-strip as follows.

Consider the central cycle of the Möbius-strip, see Figure 5.1. Now attach a disc at that central cycle. In a few lines we explain why this topological space X cannot be embedded in 3-space. Any triangulation of X gives an obstruction to embeddability. It can be shown that such triangulations have no space minor in the finite list $\mathcal{Z}$.

Why can X not be embedded in 3-space? To answer this, consider a small torus around the central cycle. The disc and the Möbius-strip each intersect

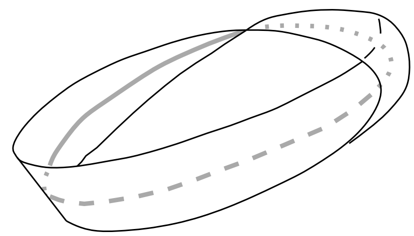

Figure 5.1: The Möbius-strip. The central cycle is depicted in grey.

that torus in a circle. These circles however have a different homotopy class in the torus. Since any two circles in the torus of a different homotopy class intersect[1], the space X cannot be embedded in 3-space without intersections of the disc and the Möbius-strip. Obstructions of this type we call *torus crossing obstructions*. A precise definition is given in Section 5.3.

A refined question might now be whether the result of Chapter 1 extends to simply connected simplicial complex if we add the list $\mathcal{T}$ of torus crossing obstructions to the list $\mathcal{Z}$ of obstructions. The answer to this question is 'almost yes'. Indeed, we just need to add to the space minor operation the operations of stretching defined in Section 5.6. These operations are illustrated in Figure 5.5, Figure 5.9 and Figure 5.11.

It is not hard to show that stretching preserves embeddability. The main result of this chapter is the following.

Theorem 5.2.1. *Let C be a simply connected simplicial complex. The following are equivalent.*

- *C has a topological embedding in 3-space;*
- *C has no stretching that has a space minor in $\mathcal{Z} \cup \mathcal{T}$.*

We deduce Theorem 5.2.1 from the results of Chapter 1 in two steps as follows. The notion of 'local almost 3-connectedness and stretched out' is slightly more general and more technical than 'local 3-connectedness', see Section 5.3 for a definition. First we extend the results of Chapter 1 to locally almost 3-connected and stretched out simply connected simplicial complexes, see Theorem 5.3.4 below. We conclude the proof by showing that any simplicial complex can be stretched to a locally almost 3-connected and stretched out one. More precisely:

Theorem 5.2.2. *For any simplicial complex C, there is a simplicial complex C' obtained from C by stretching so that C' is locally almost 3-connected and stretched out or C' has a non-planar link.*

Moreover C has a planar rotation system if and only if C' has a planar rotation system.

[1]A simple way to see this is to note that the torus with a circle removed is an annulus.

The chapter is organised as follows. In Section 5.3 we prove the extension of the results of Chapter 1 to locally almost 3-connected and stretched out simplicial complexes, Theorem 5.3.4. In Section 5.4, Section 5.5 and Section 5.6 we prove Theorem 5.2.2. Then we prove Theorem 5.2.1. Finally we describe algorithmic consequences.

For graph theoretic definitions we refer the reader to [36].

5.3 A Kuratowski theorem for locally almost 3-connected simply connected simplicial complexes

In this section we prove Theorem 5.3.4, which is used in the proof of the main theorem. First we define the list $\mathcal{T}$ of torus crossing obstructions.

Given a simplicial complex C, a *mega face* $F = (f_i | i \in \mathbb{Z}_n)$ is a cyclic orientation of faces f_i of C together with for every $i \in \mathbb{Z}_n$ an edge e_i of C that is only incident with f_i and f_{i+1} such that the e_i and f_i are locally distinct, that is, $e_i \neq e_{i+1}$ and $f_i \neq f_{i+1}$ for all $i \in \mathbb{Z}_n$. We remark that since in a simplicial complex any two faces can share at most one edge, the edges e_i are implicitly given by the faces f_i. A *boundary component* of a mega face F is a connected component of the 1-skeleton of C restricted to the faces f_i after we delete the edges e_i. Given a cycle o that is a boundary component of a mega face F, we say that F is *locally monotone* at o if for every edge e of o and each face f_i containing e, the next face of F after f_i that contains an edge of o contains the unique edge of o that has an endvertex in common with e and e_{i+1}. Under these assumptions for each edge e of o the number of indices i such that e is incident with f_i is the same. This number is called the *winding number* of F at o.

A *torus crossing obstruction* is a simplicial complex C with a cycle o (called the *base cycle*) whose faces can be partitioned into two mega faces that both have o has a boundary component and are locally monotone at o but with different winding numbers. We denote the set of torus crossing obstructions by $\mathcal{T}$.

Remark 5.3.1. The set of torus crossing obstructions is infinite. Indeed, it contains at least one member for every pair of distinct winding numbers. So it is not possible to reduce it to a finite set. However one can further reduce torus crossing obstruction as follows. First, by working with the class of 3-bounded 2-complexes as defined in Chapter 1 instead of simplicial complexes, one may assume that the cycle o is a loop. Secondly, one may introduce the further operation of gluing two faces along an edge if that edge is only incident with these two faces. This way one can glue the two mega faces into single faces. Thirdly, one can enlarge the holes of the mega faces to make them into one big hole (after contracting edges one may assume that this single hole is bounded by a loop). After all these steps we only have one torus crossing obstruction left for any pair of distinct winding numbers. This obstruction consists of three vertex-disjoint loops and two faces, each incident with two loops. The loop contained by both faces is the base cycle o. Here the faces may have winding number greater than one. The faces have winding number precisely one at the other loops.

A *parallel graph* consists of two vertices, called the *branch vertices*, and a set of disjoint paths between them. Put another way, start with a graph with only two vertices and all edges going between these two vertices, now subdivide these edges arbitrarily, see Figure 5.2.

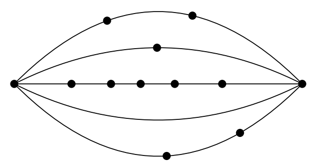

Figure 5.2: A parallel graph with five paths.

For example, parallel graphs where the branch vertices have degree two are cycles.

Given a simplicial complex C and a cycle o of C, we say that o is a *para-cycle* if all link graphs at the vertices of o are parallel graphs.

Lemma 5.3.2. *Let C be a simplicial complex. Assume that C has a para-cycle o such that for some edge e of o the link graph L of the contraction $C/(o-e)$ at the vertex $o-e$ is not loop planar. Then a torus crossing obstruction can be obtained from C by deleting faces.*

Proof. Our aim is to define a torus crossing obstruction with base cycle o. For that we define a set of possible mega faces as follows.

The complex $C/(o-e)$ has only one loop and that is e. We denote the two vertices of L corresponding to e by ℓ_1 and ℓ_2. Since o is a para-cycle, the link graph L is (isomorphic to) a parallel graph with branching vertices ℓ_1 and ℓ_2. We shall define mega faces such that every edge of the parallel graph incident with ℓ_1 is a face of precisely one of these mega faces. We define these mega faces recursively. So let f be an edge of the parallel graph incident with ℓ_1 that is not already assigned to a mega face. Let P be the path of the parallel graph between ℓ_1 and ℓ_2 that contains f. The edges on that path after f are its consecutives in its mega face. The last edge of that path is incident with ℓ_2 and hence it also corresponds to an edge incident with ℓ_1. If that face is equal to f we stop. Otherwise we continue with that face as we did with f, see Figure 5.3.

Eventually, we will come back to the face f. This completes the definition of the mega face containing f. This defines a mega face as all interior vertices of these paths have degree two. It is clear from this definition that the mega faces partition the edges of the link graph. Since o is a para-cycle, these mega-faces are also mega-faces of C and the cycle o is a boundary component of each of them. It is straightforward to check that these mega-faces are monotone at o.

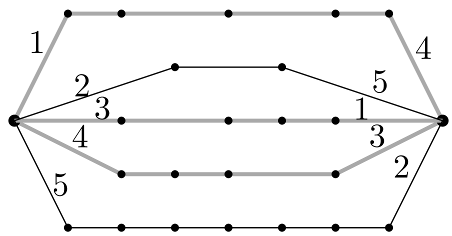

Figure 5.3: The construction of a mega face in a subdivision of B_5. The bijection between the edges incident with ℓ_1 and ℓ_2 is indicated by numbers. In grey we marked a set of the edges whose faces form a mega face.

It suffices to show that two of these mega faces have distinct winding number at o. Suppose not for a contradiction. Then all mega faces have the same winding number.

We enumerate the mega faces and let K be their total number. The winding number of a mega face is equal to the number of its traversals of the edge e, that is, its number of faces that – when considered as edges of the link graph – are incident with ℓ_1. So by our assumption, there is a constant W such that all our mega faces contain precisely W faces incident with e. We enumerate these faces in a subordering of the mega face. More precisely, by $f[k,w]$ we denote the w-th face incident with e on the k-th mega face, where k and w are in the cyclic groups $\mathbb{Z}_K$ and $\mathbb{Z}_W$, respectively.

We will derive a contradiction by constructing a rotation system of the link graph L that is loop planar. We embed it in the plane such that the rotation system at ℓ_1 is $f[1,1]$, $f[2,1]$, $\ldots$, $f[K,1]$, $f[1,2]$, $f[2,2]$, $\ldots$, $f[K,2]$, $f[1,3]$, $\ldots$,$\ldots$, $f[K,W]$, $f[1,1]$.

Then the rotation system at ℓ_2 is obtained from the that of ℓ_1 by replacing each face $f[k,w]$ by $f[k,w+1]$ and then reversing. Since this shift operation keeps this particular cyclic ordering invariant, the rotation systems at ℓ_1 and ℓ_2 are reverse. So this defines a loop planar embedding of the link graph. Hence L has a loop planar rotation system. This is the desired contradiction to our assumption. Hence two mega faces must have a different winding number. So C contains a torus crossing obstruction. $\square$

Next we define 'stretched out'. This is a technical condition, which is used only twice in the argument, namely in the proof of Lemma 5.3.3 and Lemma 5.3.5 below. A simplicial complex is *stretched out* if every edge incident with only two faces has an endvertex x such that the link graph at x is not a subdivision of a 3-connected graph and not a parallel graph whose branching vertices have degree at least three.

Next we define 'para-paths', which are similar to para-cycles and analyse them. A path in a simplicial complex C is a *para-path* if

1. the link graphs at all interior vertices of P are parallel graphs, where the

branching vertices have degree at least three;

2. the link graphs at the two endvertices of P are subdivisions of 3-connected graphs.

Lemma 5.3.3. *Let C be a stretched out simplicial complex*[2] *with a para-path P. Then the complex C' obtained from C by contracting all edges of the path P has at most one loop.*

Proof. Let v and w be the endvertices of the path P. Since C is a simplicial complex, there is at most one edge between v and w. We will show that no other edge of C becomes a loop in C'.

So let e be an edge of C that has an endvertex u on the path P different from v and w. Thus the vertex u is an interior vertex of P, so $L(u)$ is a parallel graph whose branching vertices have degree at least three. As C is stretched out, the other endvertex x of e has a link graph different from all link graphs at vertices on the para-path P. Thus x does not lie on the para-path P. Thus the edge e is not a loop in the simplicial complex C'. □

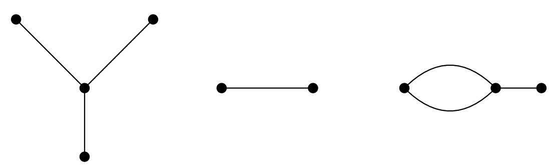

Figure 5.4: A 3-star, an edge and a 2-cycle with an attached leaf. Free-graphs are subdivisions of these graphs.

A *free-graph* is a subdivision of a 3-star, a path or a cycle with an attached path, see Figure 5.4. These graphs are 'free' in the sense that any rotation system on them defines an embedding in the plane. A graph is *almost 3-connected* if it is a subdivision of a 3-connected graph, a parallel graph or a free-graph. A simplicial complex is *locally almost 3-connected* if all its link graphs are almost 3-connected.

Theorem 5.3.4. *Let C be a simplicial complex that is locally almost 3-connected and stretched out. The following are equivalent.*

- *C has a planar rotation system;*
- *C has no space minor in $\mathcal{Z} \cup \mathcal{T}$.*

As a preparation for the proof of Theorem 5.3.4, we prove the following analogue of Lemma 1.3.1. Recall that an edge e is a *chord* of a cycle o in a simplicial complex if e is not in o but joins two vertices of o.

[2] In this chapter we follow the convention that every edge of a simplicial complex is incident with some face.

Lemma 5.3.5. *Let C be a simplicial complex that is locally almost 3-connected and stretched out. Then C has a planar rotation system unless*

1. *C is not locally planar;*
2. *there is a para-path P such that C/P is not locally planar at the vertex P;*
3. *the contraction $C/(o-e)$ is not locally planar, where o is a chordless cycle and e is an edge of o and o contains an edge aside from e.*

Proof. We obtain H from the 1-skeleton of C by deleting all edges of C that are incident with precisely two faces or such that the link graph at one endvertex is a free-graph. Let H' be a connected component of H. We say that a rotation system of C is *planar at H'* if it is planar at all vertices of H'. In order to show that C has a planar rotation system, it suffices to construct for each connected component H' of H a rotation system of C that is planar at H'. Indeed, since the rotators at vertices of degree two are unique, we can combine these rotation systems for the different components of H to a planar rotation system of C. And if the link graph at one endvertex v of an edge e is a free-graph, we can just change the rotator at e in $L(v)$ so as to be the reverse of the rotator at the link graph at e in the other link graph containing e.

First assume that H' just consists of a single vertex. Either C has a rotation system that is planar at H' or the link graph of C at the single vertex of H' is not loop planar. That is, we have the first outcome of the lemma.

Note that vertices whose link graphs are free-graphs are included in this case, as they do not have any outgoing edges in H.

Next assume that all link graphs at vertices of H' are parallel-graphs. Since we may assume that H' contains at least two vertices, each branching vertex of such a parallel graph has degree at least three; and each vertex of H' is incident with precisely two edges (which are the branching vertices in its link graph). So the connected graph H' is a cycle o. In fact, it is a para-cycle.

Our aim is to show that there is a rotation system planar at H' or we get outcome 3 from the lemma. For that we contract the edges of o one by one until a single edge e remains. After each contraction one gets a para-cycle with one fewer vertex. Similarly as Lemma 1.1.2 one proves that there is a rotation system planar at H' before the contraction of a single edge if and only if there is such a planar rotation system after the contraction. Thus C has a rotation system planar at H' or the 2-complex $C/(o-e)$ is not loop planar at $o-e$. That is, we have the third outcome of the lemma, as the cycle o of the stretched out simplicial complex C is chordless.

Thus it suffices to consider the case that H' contains a vertex whose link graph is a subdivision of a 3-connected graph. Let X be the set of those vertices of H' where the link graph is a subdivision of a 3-connected graph.

Remark 5.3.6. If all vertices of H' are in X, the proof of Lemma 1.3.1 extends almost verbatim to this case. We reduce the general case to Lemma 1.3.1 as follows.

Vertices of H' not in X have parallel graphs at their links. These vertices have degree one or two in H'. (Indeed, as they are in H' they must be parallel graphs whose (two) branching vertices have degree at least three.)

Thus the graph H' consists of the set X together with some paths between these vertices, or cycles and paths attached at single vertices of X. These paths starting at a single vertex of X must have a deleted edge d of degree three incident with their last vertex. So the link graph at the other endvertex of d is a free graph. We obtain H'' from H' by deleting all the paths attached at a single vertex of X; here we stress that we do not delete their starting vertex in X and we do not delete attached cycles. Note that there is a rotation system that is planar for H' if and only if there is a rotation system that is planar for H''.

Let C' be the simplicial complex obtained from C contracting all but one edge from every path of H'' between two vertices of X or every cycle of H'' containing precisely one vertex of X. Then C' is a 3-bounded 2-complex such that the link graph at every vertex of X is a subdivision of a 3-connected graph. As C is stretched out, no edge of degree two is a loop in C'. The 2-complex C' has one loop for every cycle of H'' containing a single vertex of H, and no further loops by construction.

As contractions of non-loops and their inverse operations preserve the existence of planar rotation systems in locally 2-connected 2-complexes by Lemma 1.1.2 and local 2-connectivity is preserved by contraction by Lemma 1.2.4, there is a planar rotation system for H'' in C if and only if there is a planar rotation system for X in C'. The same proof of as that of Lemma 1.3.1 gives that there is a planar rotation system for X in C' unless one of the following occurs.

1. C' is not locally planar;
2. there is a non-loop e of C' such that C/e is not locally planar at the vertex e;
3. the contraction $C'/(o-e)$ is not locally planar, where o is a chordless cycle and e is an edge of o and o contains an edge aside from e.

In the first case, let v' be the vertex of C' whose link graph is not loop-planar. Let v be the unique vertex of X contracted onto v'. If the link graph at v of C is not planar, we have outcome 1 of Lemma 5.3.5. Hence we may assume that this is not the case, and so the link graph $L(v)$ is planar, and so also $L(v')$ is planar – but not loop-planar. As the link graph $L(v')$ is a subdivision of a 3-connected graph, there is a single loop e of C' that witnesses that the link graph $L(v')$ is not loop planar. Let o be the unique cycle of H'' containing e. Then the 2-complex $C/(o-e)$ is not loop planar at the contraction vertex. Hence we have outcome 3 of Lemma 5.3.5, as the cycle o of the stretched out simplicial complex C is chordless.

In the second case, the edge e is a path of C all whose interior vertices have parallel graphs at their links. Hence we get a para-path as in outcome 2 of Lemma 5.3.5.

In the third case, each edge of the cycle o is a path of C, and all these paths together form a cycle of C. This cycle has no chord as o has no chord. We pick an arbitrary edge on the path for e, and we get outcome 3 of Lemma 5.3.5. $\square$

Proof of Theorem 5.3.4. By Lemma 5.3.2 we may assume that C has no para-cycle o such that for some edge e of o the contraction $C/(o-e)$ is not loop planar at the vertex $o-e$.

Next we treat the case that C has a para-path P such that the link graph $L(P)$ of C/P at P is not loop planar. The link graph $L(P)$ is the vertex-sum of the link graphs at the vertices of P. Thus it is a subdivision of a 3-connected graph by Lemma 1.2.4. By Lemma 5.3.3, C/P has at most one loop, which is incident with $L(P)$. By Lemma 1.5.6 or Lemma 1.5.12 C' has a space minor that is a generalised cone or a looped generalised cone that is not loop planar at its top, respectively. In the first case we deduce by Lemma 1.5.11 that C' has a space minor in $\mathcal{Z}_1$. In the second case we deduce similarly as in the last paragraph of the proof of Theorem 1.5.16 that C' has a space minor in $\mathcal{Z}_2$.

Having treated the above cases the rest of the proof of Theorem 5.3.4 is analogue to the proof of Theorem 1.5.16 except that we refer to Lemma 5.3.5 instead of Lemma 1.3.1. $\square$

5.4 Streching local 2-separators

In this section we define stretching at local 2-separators and prove basic properties of this operation. This operation is necessary for Theorem 5.2.2.

A *2-separator* in a 2-connected graph[3] L is a pair of vertices (a,b) such that $L-a-b$ has at least two connected components.

Given a simplicial complex C with a vertex v such that its link graph $L(v)$ is 2-connected and has a 2-separator (a,b), the simplicial complex C_2 obtained from C by *stretching* $\{a,b\}$ *at* v is defined as follows, see Figure 5.5.

We denote by Δ_2 the simplicial complex obtained from two disjoint faces of size three by gluing them together at an edge, see Figure 5.6. Let Δ_n^+ be the simplicial complex obtained by gluing n copies of Δ_2 together at a path of length two whose endvertices have degree two in Δ_2 (this is uniquely defined up to isomorphism), see Figure 5.7.

Informally, we obtain C_2 from C by replacing the edges a and b by Δ_n^+, where n is the number of components of $L(v)-a-b$. More precisely, the simplicial complex C_2 is defined as follows. Let n be the number of components of $L(v)-a-b$. We denote the gluing edges of Δ_n^+ by $\bar{a}$ and $\bar{b}$. We label the vertices of Δ_n^+ incident with neither $\bar{a}$ nor $\bar{b}$ by the components of $L(v)-a-b$.

In our notation we suppress a bijection between vertices of C and Δ_n^+ as follows. We label the common vertex of the edges of $\bar{a}$ and $\bar{b}$ by v. We denote the

[3]In this chapter we will only consider 2-separators of link graphs of simplicial complexes; such link graphs do not have parallel edges or loops. For multigraphs, it seems suitable to also consider (a,b) a 2-separator if there are two parallel edges between them and $L-a-b$ is not empty or a and b have three parallel edges in between.

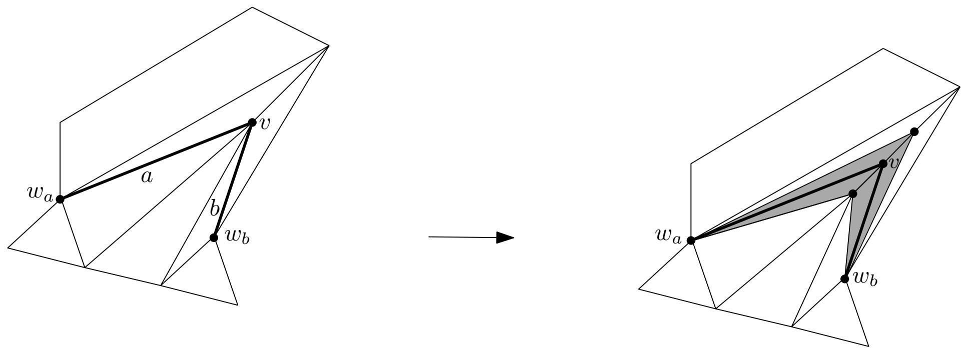

Figure 5.5: If we stretch the highlighted pair of edges in the simplicial complex on the left, we obtain the one on the right. The newly added faces are depicted in grey.

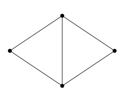

Figure 5.6: The simplicial complex Δ_2.

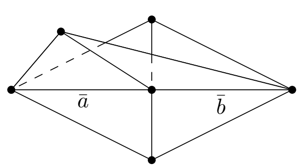

Figure 5.7: The simplicial complex Δ_3^+ with the gluing edges labelled $\bar{a}$ and $\bar{b}$.

endvertex of the edge a in C different from v by w_a; and we label the endvertex of the edge $\bar{a}$ different from v by w_a. Similarly, we denote the endvertex of the edge b in C different from v by w_b; and we label the endvertex of the edge $\bar{b}$ different from v by w_b.

- The vertex set of C_2 is union of the vertex set of C together with the vertex set of Δ_n^+, in formulas: $V(C_2) = V(C) \cup V(\Delta_n^+)$. We stress that the sets $V(C)$ and $V(\Delta_n^+)$ share the vertices v, w_a and w_b and hence these vertices appear in $V(C_2)$ only once as $V(C_2)$ is just a set and not a multiset;

- the edge set of C_2 is (in bijection with) the edge set of C with the edges a and b replaced by the set of edges of Δ_n^+, in formulas: $E(C_2) = (E(C) - a - b) \cup E(\Delta_n^+)$. The incidences between vertices and edges are as in C or Δ_n^+, except for those edges of C that have the vertex v as an endvertex. This defines all incidences of edges except those of C that have the endvertex v. Given an edge x of C incident with v, and denote its other endvertex by x'. Then its corresponding edge of C_2 has the endvertices x' and the vertex of Δ_n^+ that is the component of $L(v) - a - b$ containing x. This completes the definition of the edges of C_2. We stress that the vertex w_a of C_2 is incident with those edges of $C - a - b$ with endvertex w_a and those edges of Δ_n^+ with endvertex w_a;

- the faces of C_2 are the faces of C together with the faces of Δ_n^+; in formulas: $F(C_2) = F(C) \cup F(\Delta_n^+)$. We stress that the sets $F(C)$ and $F(\Delta_n^+)$ are disjoint. The incidences between edges and faces are as in C or Δ_n^+, where defined. This defines all incidences of faces except for those faces f of C incident with the edges a or b, which are defined as follows. There are three cases:

  – if f is a face of C incident with both the edges a and b, then in C_2 these incidences are replaced by incidences with the edges $\bar{a}$ and $\bar{b}$;

  – if f is a face of C incident with the edge a but not b, then in C_2 the incidence of f with a is replaced with an incidence with the edge $w_a x$ of Δ_n^+; where x is the component of $L(v) - a - b$ such that in $L(v)$ the edge f joins a with a vertex of x;

  – similarly, if f is a face of C incident with the edge b but not a, then in C_2 the incidence of f with b is replaced with an incidence with the edge $w_b x$ of Δ_n^+; where x is the component of $L(v) - a - b$ such that in $L(v)$ the edge f joins b with a vertex of x.

This completes the definition of stretching a 2-separator at a vertex.

We refer to the vertices of C_2 that are not in $V(C) - v$ as the *new vertices*, other vertices of C_2 are called *old*.

The link graph at w_a of C is obtained from the link graph at w_a in C_2 by contracting all edges incident with the vertex $\bar{a}$. Note that w_a cannot be incident with b as C is a simplicial complex.

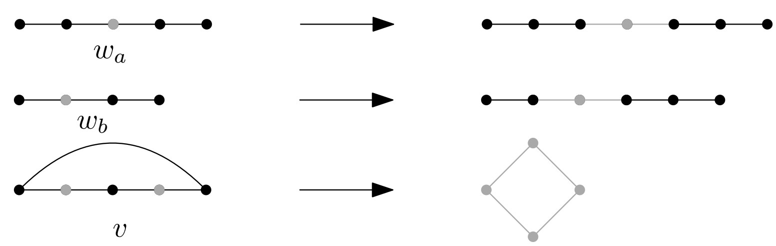

Figure 5.8: On the left we see the link graphs at the vertices w_a, w_b and v of the simplicial complex in Figure 5.5. On the right we see the link graphs after the stretching at (a, b). The vertices a and b and the new vertices and edges in the link graphs are depicted in grey.

Example 5.4.1. In Figure 5.8 we explain how the link graphs of Figure 5.5 change.

The *(abbreviated) degree-sequence* of a graph is the sequences of degrees of its vertices, ordered by size, where we leave out the degrees which are at most two. We compare degree-sequences in the lexicographical order.

Lemma 5.4.2. *Let C_2 be a simplicial complex obtained from C by stretching the 2-separator (a, b) at v. Then at each vertex of C aside from v, the degree-sequence at its link in C_2 is at most the degree-sequence at its link in C. At all new vertices of C_2 the degree-sequence at the link is strictly smaller than the degree-sequence of the link graph at v in C – unless the link graph at v in C is a parallel graph or $L(v) - a - b$ has two components and one is a path.*

Proof. Coadding a star at a vertex cannot increase the abbreviated degree-sequence, hence the lemma is true at old vertices of C_2. So it remains to prove the lemma for the new vertices of C_2 as it is obvious at the others. As the link graph $L(v)$ at v in C is not a parallel graph, the degree-sequence at the link at v in C_2 is strictly smaller than that in C. Now let X be a component of $L(v) - a - b$. If $L(v) - a - b$ has at least three components, then the degree-sequence at the link at X in C_2 is strictly smaller the degree-sequence of the link graph at v in C. This is also true if $L(v) - a - b$ has only two components and the other component has a vertex of degree greater than two; that is, is not a path. This completes the proof of the lemma. □

The *degree-parameter* of a 2-complex is the sequence of degree-sequences of all its link graphs, ordered by size. We compare degree-parameters in the lexicographical order.

Lemma 5.4.3. *Let C be a simplicial complex such that all link graphs are 2-connected or free-graphs. Then we can apply stretchings at local 2-separators of 2-connected link graphs such that the resulting simplicial complex is locally almost 3-connected.*

Before we prove this, we need a definition. A 2-separator (x, y) in a graph G is *proper* unless $G - x - y$ has precisely two components and one of them is a path and xy is not an edge.

Proof. If C has a 2-connected link graph that is not a parallel graph or a subdivision of a 3-connected graph, it contains a proper 2-separator and we stretch at that 2-separator. Link graphs at other vertices remain 2-connected or free graphs, respectively. By Lemma 5.4.2 the degree-parameter goes down and hence this process has to stop after finitely many steps – with the desired simplicial complex. □

Until the rest of this section we fix a simplicial complex C with a vertex v such that the link $L(v)$ is 2-connected and let (a, b) be a 2-separator of $L(v)$. We denote the simplicial complex obtained from C by stretching (a, b) at v by C_2.

Remark 5.4.4. C can be obtained from C_2 as follows. First we contract the edges incident with the vertex v except for $\bar{a}$ and $\bar{b}$. We relabel $\bar{a}$ by a and $\bar{b}$ by b. We obtain some faces of size two, we refer to these faces as *tiny* faces. Then we contract all these tiny faces. This gives C.

We say that an operation, such as contracting an edge, is an *equivalence* for a property, such as the existence of a planar rotation systems, if a simplicial complex has that property if and only if the simplicial complex after applying this operation has this property.

In Lemma 1.1.2 it is shown that contracting a non-loop edge where the link graph at both endvertices are 2-connected is an equivalence for the existence of planar rotation systems. Contracting a face of size two is not always an equivalence for the existence of planar rotation systems but here the contracted faces have the following additional property.

A face f incident with only two edges e_1 and e_2 is *redundant* if there is a vertex v incident with f such that in C/f in any planar rotation system of the link graph $L(v)$ at the rotator at f, the edges incident with e_1 in the link at v for C form an interval. (This implies that also the edges incident with e_2 in the link at v for C form an interval.)

The following is obvious.

Observation 5.4.5. *Let C' be obtained from C by contracting a redundant face. If C' has a planar rotation system, then C has a planar rotation system.* □

Observation 5.4.6. *Tiny faces (as defined in Remark 5.4.4) are redundant.*

Proof. Let X be a component of $L(v) - a - b$. Since $L(v)$ is 2-connected, the edges between a and X form an interval in any rotator at a for any embedding of $L(v)$ in the plane. The same is true for 'b' in place of 'a'. □

Lemma 5.4.7. *The simplicial complex C has a planar rotation system if and only if the simplicial complex C_2 has a planar rotation system.*

Proof. It is shown in Lemma 1.1.2 that contracting a non-loop edge where the link graph at both endvertices are 2-connected is an equivalence for the existence of planar rotation systems.

By Lemma 1.5.4 contracting a face of size two preserves the existence of planar rotation systems. So contracting tiny faces is an equivalence for the existence of planar rotation systems by Observation 5.4.5 and Observation 5.4.6.

Hence all the operation that transform the simplicial complex C_2 to C as described in Remark 5.4.4 are equivalences. Thus stretching at local 2-separators is an equivalence for planar rotation systems. $\square$

The following is geometrically clear, see Figure 5.5, and we will not use it in our proofs.

Lemma 5.4.8. *If C embeds in 3-space, then also C_2 embeds in 3-space.* $\square$

Remark 5.4.9. Also the converse of Lemma 5.4.8 is true.

5.5 Stretching a local branch

In this section we define stretching local branches and prove basic properties of this operation. This operation is necessary for Theorem 5.2.2.

Given a connected graph G with a cut-vertex v, a *branch* at v is a connected component X of $G-v$ together with the vertex v (and all edges between X and v). A *branch of* G is a branch at some cut-vertex of G. For any branch B, there is a unique vertex v such that B is a branch at v; we refer to that vertex v as *the cut-vertex* of the branch B.

Given a 2-complex C with a vertex v such that the link graph $L(v)$ at v is connected and a branch B of $L(v)$, the complex $C[B]$ obtained from C by *pre-stretching* B is defined as follows, see Figure 5.9. We denote the cut-vertex of the branch B by e; and remark that e is an edge of the simplicial complex C.

- The vertex set of $C[B]$ is that of C together with one new vertex, which we denote by $v[B]$, in formulas: $V(C_1) = V(C) \cup \{v[B]\}$;

- the edge set of $C[B]$ is (in bijection with) the edge set of C together with one additional edge, which we denote by $e[B]$, in formulas: $E(C[B]) = E(C) \cup \{e[B]\}$; The incidences between edges and vertices are as in C except for those edges $z \neq e$ of C that are vertices of the branch B. Such edges are incident with the new vertex $v[B]$ in place of v, the other endvertex is not changed. The edge $e[B]$ has the endvertices v and $v[B]$;

- the faces of $C[B]$ are (in bijection with) the faces of C; in formulas: $F(C[B]) = F(C)$. The incidences between faces and edges are as in C except for those faces of C that are incident with the edge e and are in the link graph $L(v)$ edges of the branch B. These faces now have size four. They are now additionally incident with the edge $e[B]$.

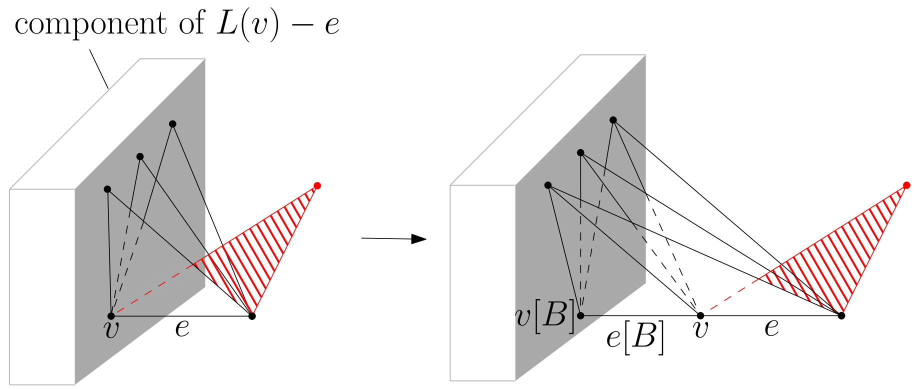

Figure 5.9: The 2-complex on the right is obtained from the 2-complex on the left by stretching the branch B, which consists of the grey box together with the three faces attaching at the grey box. Stretching is defined like pre-stretching but there we additionally subdivide faces to make them all have size three.

This completes the definition of pre-stretching the branch B at v. *Stretching* the branch B is defined the same way except that we additionally subdivide each face f of size four once. Namely we add a subdivision-edge between the vertex v and the unique vertex of the face that is not in the edge e and different from $v[B]$. Hence for any simplicial complex C any stretching at a branch is again a simplicial complex.

See Figure 5.10 for an example illustrating how the link changes at the vertex v and how the link looks like at the vertex $v[B]$.

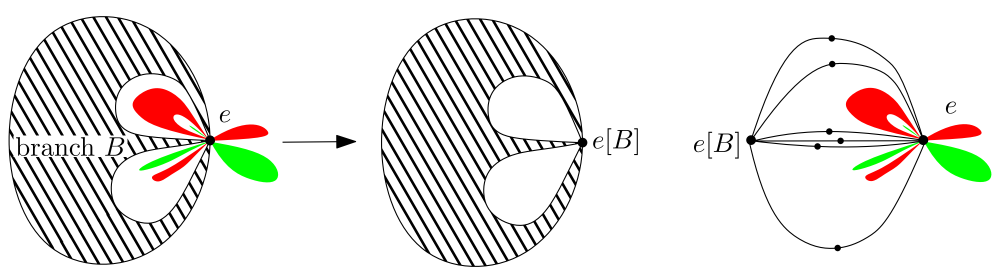

Link at v before. Link at $v[B]$. Link at v afterwards.

Figure 5.10: On the left we see the link graph of a vertex v with a cut-vertex e and a branch B. On the right we see the link graphs of the two vertices obtained from the link graph of v by stretching B.

Until the rest of this section we fix a simplicial complex C with a vertex v such that the link graph $L(v)$ is connected and let B be a branch of that link.

We denote the simplicial complex obtained from C by stretching B by $C[B]$.

Remark 5.5.1. The simplicial complex C can be obtained from $C[B]$ as follows. First we contract the edge $e[B]$. This makes the faces incident with $e[B]$ in $C[B]$ have size two. Then we contract these faces. This gives C. The contracted faces are incident with an edge that is incident with only one other face.

Lemma 5.5.2. *The simplicial complex $C[B]$ has a planar rotation system if and only if C has a planar rotation system.*

Proof. Let Σ be a planar rotation system of the simplicial complex C. We define[4] a rotation system Σ' of the simplicial complex $C[B]$ by taking the same rotator as Σ at every edge except for $e[B]$ and other new edges, which are incident with two faces. At the edges incident with two faces we take the unique cyclic ordering of size two. The rotator at the edge $e[B]$ is constructed from the rotator at the edge e for Σ by restricting it to the faces incident with the edge $e[B]$.

This rotation system is obviously planar at all vertices of $C[B]$ except for the vertex v and $v[B]$. We denote by Π the rotation system induced by Σ of the link graph of C at the vertex v. By the construction given directly after Lemma 1.4.3, Π induces a planar rotation system Π_1 at the branch B, and Π induces a planar rotation system Π_2 at the minor of the link graph at v in C obtained by contracting $B - e$ to a single vertex. It is immediate that the rotation system induced by Σ' at $v[B]$ is Π_1, and the rotation system induced by Σ' at v is Π_2. Hence the rotation system Σ' is planar for the simplicial complex $C[B]$.

By Lemma 1.1.2 contracting an edge preserves the existence of planar rotation systems, and by Lemma 1.5.4 contracting a face of size two preserves the existence of planar rotation systems. Hence by Remark 5.5.1 if $C[B]$ has a planar rotation system, then C has a planar rotation system. $\square$

The following is geometrically clear, see Figure 5.9, and we will not use it in our proofs.

Lemma 5.5.3. *If C embeds in 3-space, then also $C[B]$ embeds in 3-space.* $\square$

Remark 5.5.4. Also the converse of Lemma 5.5.3 is true.

5.6 Increasing local connectivity

In the first three subsections of this section we define stretchings and prove basic properties; these are necessary for Theorem 5.2.2. The forth subsection is a preparation for the last subsection, in which we prove Theorem 5.2.2, and Theorem 5.2.1.

[4] Faces incident with the edge e in C correspond in $C[B]$ either to a single face of size three or two faces of size three obtained from a face of size four by subdivision. This induces a bijective map from the faces incident e in C to the faces incident with e in $C[B]$, and an injective partial map from the faces incident e in C to the faces incident with $e[B]$ in $C[B]$. In order to simplify the presentation of the definition we suppress these two maps.

5.6.1 The operation of stretching edges

Let C be a 2-complex and let e be an edge of C incident with two faces f_1 and f_2. Assume that there is an endvertex v of the edge e such that in any planar rotation system of the link graph $L(v)$ at v the edges f_1 and f_2 are adjacent in the rotator at e. The complex C' obtained from C by *pre-stretching the edge e in the direction of f_1 and f_2* is obtained from C as follows, see Figure 5.11. We replace the edge e by two edges new edges e_1 and e_2, both with the same endvertices as e. We add a face of size two only incident with e_1 and e_2. The faces f_1 and f_2 are incident with e_1 instead of e, all other faces incident with e in C are incident with e_2 instead. This completes the definition of pre-stretching an edge. *Stretching* an edge is defined the same way except that additionally we subdivide the new face of size two to obtain a simplicial complex, see Figure 5.12.

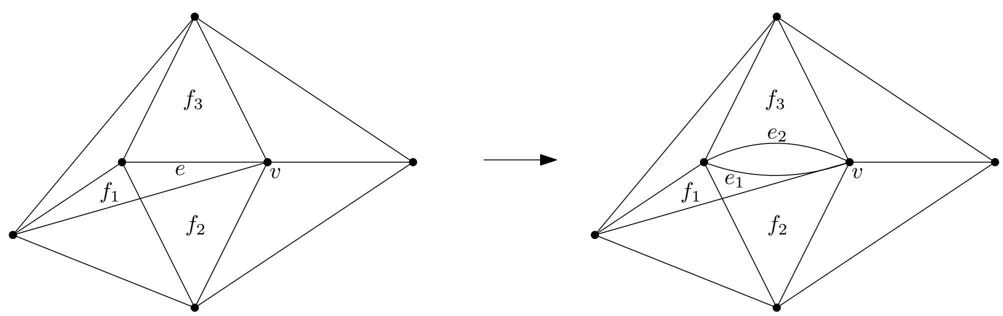

Figure 5.11: The graph on the left defines a simplicial complex by adding faces on all cycles of size three. We obtain the 2-complex on the right by pre-stretching the edge e in the direction of the faces f_1 and f_2. Its faces are all triangles of the graph on the left except that the edge e_1 is only incident with the two faces f_1 and f_2 and the new face $\{e_1, e_2\}$ and the edge e_2 is only incident with the face f_3 and the new face $\{e_1, e_2\}$.

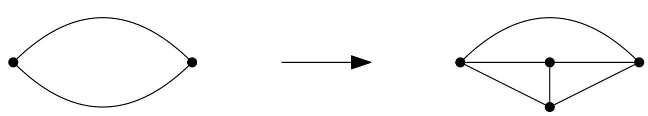

Figure 5.12: Subdivision of a face of size two to a simplicial complex.

Example 5.6.1. The assumption for stretching an edge e is particularly easy to verify if the link graph $L(v)$ is 3-connected. Indeed, then by a theorem of Whitney, we just need to check whether for a particular embedding of the link graph $L(v)$ the edges f_1 and f_2 are adjacent.

Remark 5.6.2. The inverse operation of pre-stretching an edge e to a face $\{e_1, e_2\}$ is contracting the face $\{e_1, e_2\}$ to the edge e as defined in Chapter 1.

Lemma 5.6.3. *Let C' be obtained from C by pre-stretching an edge e. Then C' has a planar rotation system if and only if C has a planar rotation system.*

Proof. Let Σ be a planar rotation system of the 2-complex C. We denote the two new edges of the 2-complex C' by e_1 and e_2. We obtain a rotation system Σ' of the 2-complex C' from Σ by taking the same rotators at all edges of the 2-complex C' except for e_1 and e_2. By assumption, the faces f_1 and f_2 along which we pre-stretch the edge e are adjacent in the rotator at the edge e. We define the new rotator at the edge e_1 to be the rotator of the edge e restricted to the adjacent faces f_1 and f_2 and we add the new face $\{e_1, e_2\}$ in place of the interval formed by the deleted faces. Similarly, we define a rotator at the edge e_2: we delete from the rotator at e the faces f_1 and f_2 and add the face $\{e_1, e_2\}$ in the interval formed by the two deleted faces. It remains to check that the rotation system Σ' is planar. This is immediate at all vertices except for the two endvertices of the edge e. For the two endvertices, note that pre-stretching the edge e has the effect on the link graph as coadding an edge at the vertex e. As the edges f_1 and f_2 of the link graphs are adjacent, the coaddition can be done within the embeddings of the link graphs given by Σ.

By Lemma 1.5.4, contracting a face of size two preserves the existence of planar rotation systems. Hence by Remark 5.6.2 if C' has a planar rotation system, then C has a planar rotation system. □

The following is geometrically clear, see Figure 5.11, and we will not use it in our proofs.

Lemma 5.6.4. *Let C' be obtained from C by stretching an edge e. If C embeds in 3-space, then also C' embeds in 3-space.* □

Remark 5.6.5. Also the converse of Lemma 5.6.4 is true.

5.6.2 The operation of contracting edges

An edge e in a 2-complex C is *reversible* if the 2-complex C has a planar rotation system if and only if the 2-complex C/e has a planar rotation system.

A *para-star* is a graph obtained from a family of disjoint parallel graphs by gluing them together at a single vertex.

Lemma 5.6.6. *Let e be a non-loop edge with endvertices v and w of a 2-complex C such that the link graphs $L(v)$ and $L(w)$ are para-stars and the vertex e is a maximum degree vertex in both of them. Then the edge e is reversible.*

Proof. By Lemma 1.1.2, it suffices to show how any planar rotation system Σ on the 2-complex C/e induces a planar rotation system Σ' on the 2-complex C. Letting Σ' to be equal to Σ at all edges of C not incident with v or w, it suffices to show the following.

Sublemma 5.6.7. *Let $L(v)$ and $L(w)$ be para-stars and let the vertex e have maximal degree in both of them. Let $L(e)$ be the vertex sum of $L(v)$ and $L(w)$ along e. For any planar rotation system Π of the graph $L(e)$, there are planar* *rotation systems of the graphs $L(v)$ and $L(w)$ that are reverse of one another at the vertex e, and otherwise agree with the rotation system* Π.

Proof. Throughout we assume in the graphs $L(v)$ and $L(w)$, the vertex e is adjacent to any other vertex. This easily implies the general case by suppressing suitable degree two vertices as rotators at such vertices are unique.

We prove this by induction on the number of branches of the graph $L(v)$. The base case is that the graph $L(v)$ is a parallel graph. Then the graph $L(e)$ is isomorphic to the graph $L(w)$. So a planar rotation system on the graph $L(e)$ induces a planar rotation system on the graph $L(w)$. And there is a unique planar rotation system on the graph $L(v)$ whose rotator at e is reverse to the rotator at e in that planar rotation system of $L(w)$.

So we may assume that the graph $L(v)$ has at least two branches. We split into two cases.

Case 1: the graph $L(e)$ is disconnected. We consider $L(e)$ as a bipartite graph with the vertex set of $L(v) - e$ on the left and the vertex set of $L(w) - e$ on the right. As every vertex of $L(e)$ is incident with an edge, there are two vertices of $L(v) - e$ in different connected components of the bipartite graph $L(e)$. Denote these two vertices by y and z. We obtain $L(v)'$ from $L(v)$ by identifying the vertices y and z into a single vertex. Denote that new vertex by u. We denote the vertex sum of $L(v)'$ and $L(w)$ along e by $L(e)'$. The graph $L(e)'$ is equal to the graph obtained from $L(e)$ by identifying the vertices y and z. Thus any planar rotation system of the graph $L(e)$ induces a planar rotation system of the graph $L(e)'$ by sticking the rotation systems at the vertices y and z together so that the rotator at the new vertex u contains the edges incident with the vertex y or z, respectively, as a subinterval. By induction such a rotation system induces planar rotation systems at the graphs $L(v)'$ and $L(w)$. This planar rotation system at the graph $L(v)'$ induces a rotation system on the graph $L(v)$ by splitting the rotator at u into the two subintervals for the vertices y and z. This induced rotation system is planar for $L(v)$ as the rotators for y and z are subintervals of the rotator for u.

Case 1: not Case 1, so the graph $L(e)$ is connected. As above, we consider $L(e)$ as a bipartite graph, and let a planar rotation system of the graph $L(e)$ be given. Since the left side has at least two vertices, there is a vertex on the right of the connected bipartite graph $L(e)$ that has two neighbours on the left. Pick such a vertex x. We pick neighbours y and z of x in $L(v) - e$ such that there are edges e_y between y and x and e_z between z and x such that these two edges are incident in the rotator at the vertex x. Let $L(v)'$ be the graph obtained from $L(v)$ by identifying the vertices y and z to a single vertex. Call that new vertex u. We denote the vertex sum of $L(v)'$ and $L(w)$ along e by $L(e)'$. The graph $L(e)'$ is equal to the graph obtained from $L(e)$ by identifying the vertices y and z. The chosen planar rotation system of the graph $L(e)$ induces a rotation system for the graph $L(e)'$ by sticking the rotation systems at the vertices y and z together so that the rotator at the new vertex u contains the edges incident with the vertex y or z, respectively, as a subinterval. By the choice of y and z this rotation system is planar. By induction this planar rotation system on $L(e)'$ induces planar rotation systems on the graphs $L(v)'$ and $L(w)$. This planar rotation system at the graph $L(v)'$ induces a rotation system on the graph $L(v)$ by splitting the rotator at u into the two subintervals for the vertices y and z. This induced rotation system is planar for $L(v)$ as the rotators for y and z are subintervals of the rotator for u. $\square$

To summarise the proof of Lemma 5.6.6, we define the planar rotation system Σ' for the 2-complex C as indicated above, and we choose the rotators at the edges incident with the vertices v or w as induced in the sense of Sublemma 5.6.7 by the rotation system of the link graph $L(e)$ at the vertex e of the 2-complex C/e. $\square$

5.6.3 The definition of stretching

We say that a simplicial complex $\tilde{C}$ is obtained from a simplicial complex C by *stretching*, if it is obtained from C by applying successively operations of the following types:

1. stretching local branches at connected link graphs;
2. 2-stretching at local 2-separators of 2-connected link graphs;
3. stretching edges;
4. contracting reversible edges that are not loops;
5. splitting vertices.

We also call $\tilde{C}$ a *stretching* of C.

Lemma 5.6.8. *Assume C' is a stretching of C. Then C has a planar rotation system if and only if C' has a planar rotation system.*

Proof. In the language introduced above we are to show that all five stretching operations are equivalences for the property 'existence of planar rotation systems'. For the first operation it is proved in Lemma 5.5.2, for the second it is proved in Lemma 5.4.7, and for the third it is proved in Lemma 5.6.3 for pre-stretchings of edges, and so the result for stretchings follows. For the forth operation it is true by the definition of reversible. Splitting vertices is clearly an equivalence for the existence of planar rotation systems. $\square$

5.6.4 Increasing the local connectivity a bit

A 2-complex is *locally almost 2-connected* if all its link graphs are 2-connected or free graphs. The following is a key step towards Theorem 5.2.2.

Theorem 5.6.9. *Any simplicial complex C has a stretching C' that is a simplicial complex so that C' is locally almost 2-connected or C' has a non-planar link graph.*

Before we prove Theorem 5.6.9, we need some preparation. A *star of parallel graphs* is a graph that is not 2-connected and is obtained from a set of disjoint parallel graphs by gluing them together at a single vertex.

Example 5.6.10. The only parallel graphs that are stars of parallel graphs are paths. Stars of parallel graphs are 2-connected para-stars.

Lemma 5.6.11. *Let C be a simplicial complex that is locally connected. Then there is a simplicial complex $\tilde{C}$ that is obtained from C by stretching local branches such that every link graph of $\tilde{C}$ is 2-connected or a star of parallel graphs.*

Proof. We will prove this by induction. The base case is that every link graph is 2-connected or a star of parallel graphs. Next we consider the case that each link graph has at most one cut-vertex. Let v be a vertex of the simplicial complex C such that its link graph has a cut-vertex e. Then all branches of e are 2-connected graphs. We stretch all branches of e that are not parallel graphs, one after the other. Then the link at v becomes a star of parallel graphs and all other new link graphs are 2-connected. The old link graphs, those at vertices of C aside from v, do not change except possibly for subdividing edges[5]. We apply this recursively to all link graphs with cut-vertices, and so reduce this case to the base case.

Next suppose that there is a vertex v such that its link graph has at least two cut-vertices. Let e_1 be an arbitrary cut-vertex of that link graph. And let B be a branch of e_1 containing another cut-vertex e_2. Then we stretch B. All link graphs at vertices of C aside from v are not changed (except for possibly subdividing edges). The vertex v is replaced by two new vertices. Each cut-vertex of the link graph of v is in precisely one of the two new link graphs, and e_1 and e_2 are in different link graphs. Hence both new link graphs have strictly less cut-vertices than the link graph of v. Hence we can apply induction (on the sequence of numbers of cut-vertices of link graphs, ordered by size and compared in lexicographical order).

□

Proof of Theorem 5.6.9. The *cutvertex-degree* of a simplicial complex C is the maximal degree of a cutvertex of a link graph of the simplicial complex C. We prove Theorem 5.6.9 by induction on the cutvertex-degree. So let C be a simplicial complex with cutvertex-degree a.

We obtain C_1 from C by splitting all vertices whose link graphs are disconnected. The simplicial complex C_1 is locally connected. By Lemma 5.6.11 there is a stretching C_2 of the simplicial complex C_1 such that all its link graphs are 2-connected or stars of parallel graphs.

If the cutvertex-degree a is at most three, then all link graphs of C_2 are 2-connected or stars of parallel graphs where the unique cut-vertex has degree at most three. Graphs of the second type are always free, see Figure 5.4. This

[5]The vertices v' of C where those subdivisions occur are those such that there is an edge e' between v and v' such that e' is a vertex of one of the branches we stretch.

completes the proof if the cutvertex-degree is at most three, so from now on let the cutvertex-degree a be at least four.

We say that a simplicial complex C is *a-nice* if all its link graphs are 2-connected, or stars of parallel graphs whose cutvertex has degree precisely a or else have maximum degree strictly less than a. For example the simplicial complex C_2 is a-nice. We say that a simplicial complex C is *a-structured* if all its link graphs are parallel graphs, or stars of parallel graphs whose cutvertex has degree precisely a or else have maximum degree strictly less than a. For example, every a-structured simplicial complex is a-nice.

Sublemma 5.6.12. *Assume C_2 is a-nice. There is a stretching C_3 of C_2 that is a-structured or else has a non-planar link.*

Proof. We prove this sublemma by induction on the degree-parameter as defined in Section 5.4. So let C_2 be an a-nice simplicial complex such that all a-nice simplicial complexes with strictly smaller degree-parameter have a stretching that is a-structured or has a non-planar link.

We may assume that the simplicial complex C_2 is not a-structured; that is, it has a vertex v such that the link graph $L(v)$ is 2-connected but no parallel graph and $L(v)$ has vertex e of degree at least a. Also we may assume that $L(v)$ is planar.

Case 1: the vertex e is not contained in a proper 2-separator of $L(v)$. Since embeddings of 2-connected graphs in the plane are unique up to flipping at 2-separators by a theorem of Whitney, any embedding of the graph $L(v)$ in the plane has the same rotator at the vertex e (up to reversing). Take two edges f_1 and f_2 incident with the vertex e that are adjacent in the rotator. Now we stretch the edge e of C_2 in the direction of the faces corresponding to f_1 and f_2. The link graphs at all vertices of C_2 except for v and the other endvertex w of the edge e of C_2 do not change. In the link graphs for v and w the vertex e is replaced by two new vertices (and a path of length two joining them), each of strictly smaller degree than e, as its degree a is at least four. This new simplicial complex C_3 has strictly smaller degree-parameter than C_2.

In order to be able to apply induction, we need to show that C_3 is a-nice. The link graph at v is still 2-connected in C_3. If the link graph at w in C_2 is 2-connected, this is still true in C_3. Hence it remains to consider the case that it is a star of parallel graphs. In this case the vertex e must be the cutvertex of $L(w)$ by the choice of a. Then in the simplicial complex C_3, the link graph at w has maximum degree less than a. Thus C_3 is a-nice and we can apply the induction hypothesis. So C_3 has a stretching of the desired type, and C_3 is a stretching of C_2. This completes the induction step in this case.

Case 2: not Case 1. Then the vertex e is contained in a proper 2-separator of $L(v)$. Let x be the other vertex in that 2-separator. We obtain C_3 from C_2 by stretching at the 2-separator $\{e, x\}$. The simplicial complex C_3 has strictly smaller degree parameter than C_2 by Lemma 5.4.2. We verify that C_3 is a-nice. All link graphs at new vertices are still 2-connected in C_3. Let w and w' be the endvertices of the edges e and x aside from v, respectively. Hence it remains so show that the link graphs at w and w' in C_3 are 2-connected, stars of parallel graphs whose cutvertex has degree a or have maximal degree less than a. If the link graph at w in C_2 is 2-connected, it is also 2-connected in C_3 (as coadding a star preserves 2-connectedness). Hence we may assume that the link graph at w in C_2 is a star of parallel graphs. and e is its cutvertex by the choice of a. Then either $L(w)$ is still a star of parallel graphs in C_3 or else it has maximum degree less than a. The same analysis applies to the vertex 'w'' in place of 'w'. Thus C_3 is a-nice. So by induction there is a stretching of C_3 of the desired type, and C_3 is a stretching of C_2. This completes the induction step, and hence the proof of this sublemma. □

Let C_3 be stretching of the simplicial complex C_2 as in Sublemma 5.6.12. If C_3 has a non-planar link, we are done. Hence we may assume that the simplicial complex C_3 is a-structured. If C_3 has cutvertex-degree less than a, we can apply the induction hypothesis. Hence we may assume that C_3 has a vertex v such that the link graph at v is a star of parallel graphs whose cutvertex e has degree precisely a. Now we show how the property 'a-structured' implies the existence of certain paths, which can then be contracted to reduce the cutvertex-degree.

Sublemma 5.6.13. *There is a path P_n from the vertex v starting with e to another vertex w_n whose link graph is a star of parallel graphs. All link graphs at internal vertices of the path are parallel graphs and all edges of the path have the same face-degree.*

Proof. We build the path $P_n = w_0e_1w_1...e_nw_n$ recursively as follows. We start with $e_1 = e$ and $w_0 = v$ and let w_1 be the endvertex of the edge e aside from v. Assume we already constructed $w_0e_1w_1...e_iw_i$. If the link graph $L(w_i)$ is a star of parallel graphs we stop and let $i = n$ and $w_i = w_n$. Otherwise by assumption, the link graph $L(w_i)$ must be 2-connected. As the edge e_i has degree precisely a the link graph $L(w_i)$ of the a-structured simplicial complex C_3 is a parallel graph. So the link graph $L(w_i)$ contains a unique vertex except from e_i that has degree larger than two, and this vertex has the same degree as the vertex e_i. We pick this vertex for e_{i+1}. Note that e_{i+1} is an edge of the simplicial complex C. We let w_{i+1} be the endvertex of e_{i+1} different from w_i. Note that all edges e_i have the same face-degree by construction. Since any path[6] in C must be finite, it suffices to prove the following:

Fact 5.6.14. *For all i the walk P_i is a path.*

Proof. We prove this by induction on i. The base case is that $i = 1$. Suppose for a contradiction there is some $j < i$ such that $w_i = w_j$.

Case 1: $j = 0$: then e_i must be equal to the only vertex of $L(v)$ of the same degree; that is, e_i is equal to the edge e_1. But then the endvertex w_{i-1} of the edge e_i is equal to the vertex w_1. This is a contradiction to the induction hypothesis. Hence w_i cannot be equal to w_0.

[6] A *path* in a graph is a sequence alternating between vertices and edges such that adjacent members are incident, and all vertices (and edges) are distinct.

Case 2: $j \geq 1$: as in the link graph $L(w_j)$ the only two vertices with the same degree as the vertex e_i are e_j and e_{j+1}, it must be that the edge e_i is equal to one of these two edges; that is, the endvertex w_{i-1} of e_i must be equal to w_{j-1} or w_{j+1}. The vertex w_{j-1} cannot be an option by the induction hypothesis. Similarly, the vertex w_{j+1} cannot be an option by the induction hypothesis if $j+1 < i-1$. So $j+2 \geq i$, so $j = i-2$ or $j = i-1$. Since the simplicial complex C has no loops or parallel edges any three consecutive vertices on P_i, such as w_{i-2}, w_{i-1} and w_i, are distinct. Hence neither $j = i-2$ nor $j = i-1$ are possible. Thus we have also reached a contradiction in this case. Hence the vertex w_i is distinct from all previous vertices on the walk P_i. □

□

Given a path P_n with endvertex w_n as in Sublemma 5.6.13, whose link graph $L(w_n)$ at w_n is a star of parallel graphs, denote the (unique) cut-vertex of the link graph $L(w_n)$ by x.

Sublemma 5.6.15. *The cut-vertex x is equal to the last edge e_n on the path P_n.*

Proof. We denote the degree of the cut-vertex x by a'. By the definition of a, we have, $a' \leq a$.

On the other hand by Sublemma 5.6.13 the vertices e and e_n have the same degree in the graphs $L(v)$ and $L(w_n)$, and this degree is equal to a by the choice of the vertex v. As $L(w_n)$ is a star of parallel graphs with a cut-vertex, the degree of the cut-vertex x is strictly larger than the degree of any other vertex of $L(w_n)$. Hence it must be that $a = a'$ and $x = e_n$. □

By Sublemma 5.6.13 and Sublemma 5.6.15, there is a set of vertex-disjoint paths in C_3 such that any of their endvertices has a link graph that is a star of parallel graphs whose cutvertex has degree a. All internal vertices of these paths are parallel graphs. And by taking this collection maximal, we ensure that any vertex whose link graph is a star of parallel graphs whose cutvertex has degree a is an endvertex of one of these paths. We denote the set of these paths by $\mathcal{P}$. We obtain the 2-complex C_4 from C_3 by contracting all edges on these paths of $\mathcal{P}$. Contracting the edges on the paths recursively, we note at each step that these edges are reversible by Lemma 5.6.6. At all vertices of C except for those vertices on the paths, the two 2-complexes C_3 and C_4 have the same link graphs. In addition, C_4 has the contraction vertices, one for each of the vertex-disjoint paths. These link graphs are the vertex-sum of the link graphs at the vertices on its path, see Section 1.2 for background on vertex-sums. So the link graph at a new contraction vertex is (isomorphic to) the vertex sum of the link graphs at the two endvertices plus various subdivision vertices coming from the parallel graphs at internal vertices of the path. By Sublemma 5.6.15, each of these vertices in the link graph has degree strictly less than a. Hence all new contraction vertices have maximum degree less than a. Hence the cutvertex-degree of C_4 is strictly smaller than a. So the 2-complex

C_4 satisfies all the conditions to apply the induction hypothesis except that it may not be a simplicial complex as it may have edges that are loops or parallel edges.

Now we show how we can stretch local branches of C_3 to get a simplicial complex C_3' so that the simplicial complex C_4' obtained from C_3' by contracting all the paths in $\mathcal{P}$ is a simplicial complex. We obtain C_3' from C_3 by stretching at each endvertex of a path in $\mathcal{P}$ all the branches and at each interior vertex of a path in $\mathcal{P}$ we stretch at the 2-separator consisting of the two branching vertices of its parallel graph. We obtain C_4' from C_3' by contracting the above defined family of paths $\mathcal{P}$. It is straightforward to check that C_4' is a simplicial complex – and is a stretching of C with smaller cutvertex-degree. This completes the induction step, and hence this proof. $\square$

5.6.5 Proofs of Theorem 5.2.1 and Theorem 5.2.2

We conclude this section by proving the following theorems mentioned in the Introduction.

Proof of Theorem 5.2.2. Let C be a simplicial complex. Recall that Theorem 5.2.2 says there is a simplicial complex C''' obtained from C by stretching so that C''' is locally almost 3-connected and stretched out or C''' has a non-planar link; moreover C has a planar rotation system if and only if C''' has a planar rotation system.

By Theorem 5.6.9 there is a stretching C' of C that is a simplicial complex that is locally almost 2-connected or has a non-planar link. As we are done otherwise, we may assume that C' is locally almost 2-connected. By Lemma 5.4.3 there is a stretching C'' of C' that is a locally almost 3-connected simplicial complex.

Sublemma 5.6.16. *Let C'' be a locally almost 3-connected simplicial complex. Then there is a stretching C''' of C'' that has additionally the property that it is stretched out.*

Proof. We say that an edge of face-degree two is *stretched out* if it has one endvertex that is not a subdivision of a 3-connected graph or a parallel graph whose branch vertices have degree at least three. Note that a simplicial complex in which every edge of degree two is stretched out is stretched out itself. We prove this sublemma by induction on the number of edges of degree two that are not stretched out. So assume there is an edge e that is not stretched out. Let v be one of its endvertices.

Case 1: the link graph at v is a parallel graph whose two branch vertices x_1 and x_2 have degree at least three. Then we stretch at the 2-separator (x_1, x_2) at v. This gives a simplicial complex $\tilde{C}$ that in addition to the vertex v has also one new vertex for every component of $L(v) - x_1 - x_2$. The link graphs at these new vertices are cycles. Hence every edge of degree two incident with these new vertices is stretched out. Thus $\tilde{C}$ has strictly less edges of degree two that are not stretched out.

Case 2: the link graph at v is a subdivision of a 3-connected graph. Then the vertex e of $L(v)$ is contained in a subdivided edge. Let P be the path of that subdivided edge, and let x_1 and x_2 be its endvertices. Then we stretch at the 2-separator (x_1, x_2) at v. The rest of the analysis is analogue to Case 1. This completes the proof of the sublemma. $\square$

By Sublemma 5.6.16 we may assume that C has a stretching C''' that is a locally almost 3-connected and stretched out simplicial complex. The 'Moreover'-part follows from the fact that C''' is a stretching of C as shown in Lemma 5.6.8. This completes the proof. $\square$

Proof of Theorem 5.2.1. Let C be a simply connected simplicial complex. Recall that Theorem 5.2.1 says that C has an embedding in 3-space if and only if C has no stretching that has a space minor in $\mathcal{Z} \cup \mathcal{T}$. By Theorem 2.2.1 C is embeddable in 3-space if and only if it has a planar rotation system.

By Theorem 5.2.2 there is a simplicial complex C' that is a stretching of C. Moreover C has a planar rotation system if and only if C' has a planar rotation system. By that theorem either the simplicial complex C' has a non-planar link or it is locally almost 3-connected and stretched out. In the first case, by Kuratowski's theorem, Lemma 1.5.6 and Lemma 1.5.11, the simplicial complex C' has a minor in the finite list $\mathcal{Z}$ – so the theorem is true in this case. In the second case by Theorem 5.3.4 C' has a planar rotation system if and only if it has no space minor in $\mathcal{Z} \cup \mathcal{T}$. This completes the proof. $\square$

5.7 Algorithmic consequences

Our proofs give a quadratic algorithm that verifies whether a given 2-dimensional simplicial complex has a planar rotation system. This gives a quadratic algorithm that checks whether a given 2-dimensional simplicial complex has an embedding in a (compact) 3-manifold by Lemma 2.5.1. In particular, for simply connected 2-complexes this gives a quadratic algorithm that tests embeddability in 3-space by Perelman's theorem. The algorithm has several components.

1. The locally almost 3-connected and stretched out case. The corresponding fact in the paper is Lemma 5.3.5. This clearly has a linear time algorithm.

2. Reduction of the locally almost 3-connected case to the locally almost 3-connected and stretched out case. The corresponding fact in the paper is Sublemma 5.6.16. This clearly has a linear time algorithm.

3. Reduction of the locally almost 2-connected case to the locally almost 3-connected case. The corresponding fact in the paper is Lemma 5.4.3. To analyse the running time, we do this step slightly differently than in the paper. First we compute a Tutte-decomposition at every 2-connected link graph. This tells us precisely how we can stretch that vertex along 2-separators. Doing these stretchings at different vertices may affect the link graphs at other vertices. However, once a link graph is a subdivision

of a 3-connected graph, it will stay that. So the vertices we may have to look at multiple times are vertices where the link graphs are parallel graphs. But if we need to stretch there again, the maximum degree goes down. It is straightforward to show that this step can be done in linear time.

4. Reduction of the general case to the locally almost 2-connected case. The corresponding fact in the paper is Theorem 5.6.9.

   This is done by recursion on the cutvertex-degree a. So let us analyse the step from a to $a-1$ in detail. The input is a simplicial complex C and we measure its size by $\sum_{e\in E(C)}(deg(e)-2)$; subtracting minus two per edge is okay as edges of degree two have unique rotation systems, and we may assume that C has no edge of degree less than two.

   The stretching related to Lemma 5.6.11 can be done in linear time as computing the block-cutvertex-tree of link graphs can be done in linear time. For the part corresponding to Sublemma 5.6.12 we compute the stretching via Tutte-decompositions as in step 3 explained above, and then we check for planarity for each 3-connected link graph. If it is planar, we remember a planar rotation system and if we stretch later an edge incident with that vertex at the other endvertex we check whether this stretching is compatible with the chosen planar rotation system. This can be done in linear time. The construction of the set $\mathcal{P}$ of paths can clearly be done in linear time. Hence the whole recursion step from a to $a-1$ just takes linear time. The output is the simplicial complex C_4'.

   However, with the current argument, the size of C_4' might be bigger than the size of the input C. Indeed, stretching a local branch may increase the size. Hence here we explain how we modify the construction of the simplicial complex C_3' so that this does not occur. First note that none of the stretching operations except for stretching a branch increases the size. We obtain C'' from C_3' by contracting all edges $e[B]$ of degree at least three that were added by stretching a local branch. It is easy to see that C'' is a simplicial complex. Each path $P\in\mathcal{P}$ of C_3' contracts onto a closed trail of C''. We obtain C_3'' from C'' by stretching branches for each closed trail $P\in\mathcal{P}$ as follows.

   Case 1: in the simplicial complex C'' the trail P has at least one internal vertex. Denote the edge of P incident with v by e_1 and the edge of P incident with w by e_2. Then we stretch at the link graphs of v and w all branches at e_1 and e_2, respectively.

   Case 2: not Case 1. Note that P must consist of at least one edge by Sublemma 5.6.13. And that edge is not a loop as C'' is a simplicial complex. So P consists of a single edge e. Then in the simplicial complex C'' there is no edge in parallel to e. We stretch all branches of the link graph at v at the vertex e (but not for w).

   We obtain C_4'' from C_3'' by contracting all $P\in\mathcal{P}$. It is straightforward to check that C_4'' is a simplicial complex and that the size of C_4'' is at most that of C. So C_4'' is a suitable output of the recursion step. As each recursion step takes linear time, all of them together take at most quadratic time.

This completes the description of the algorithm.

Acknowledgement

I thank Radoslav Fulek for pointing out an error in an earlier version of this paper. I thank Nathan Bowler and Reinhard Diestel for useful discussions on this topic.

Part II

Infinite graphs

Chapter 6

Edge-disjoint double rays in infinite graphs: a Halin type result

6.1 Abstract

We show that any graph that contains k edge-disjoint double rays for any $k \in \mathbb{N}$ contains also infinitely many edge-disjoint double rays. This was conjectured by Andreae in 1981.

6.2 Introduction

We say a graph G has *arbitrarily many vertex-disjoint H* if for every $k \in \mathbb{N}$ there is a family of k vertex-disjoint subgraphs of G each of which is isomorphic to H. Halin's Theorem says that every graph that has arbitrarily many vertex-disjoint rays, also has infinitely many vertex-disjoint rays [44]. In 1970 he extended this result to vertex-disjoint double rays [47]. Jung proved a strengthening of Halin's Theorem where the initial vertices of the rays are constrained to a certain vertex set [62].

We look at the same questions with 'edge-disjoint' replacing 'vertex-disjoint'. Consider first the statement corresponding to Halin's Theorem. It suffices to prove this statement in locally finite graphs, as each graph with arbitrarily many edge-disjoint rays contains a locally finite union of tails of these rays. But the statement for locally finite graphs follows from Halin's original Theorem applied to the line-graph.

This reduction to locally finite graphs does not work for Jung's Theorem or for Halin's statement about double rays. Andreae proved an analog of Jung's Theorem for edge-disjoint rays in 1981, and conjectured that a Halin-type Theorem would be true for edge-disjoint double rays [7]. Our aim in the current chapter is to prove this conjecture.

More precisely, we say a graph G has *arbitrarily many edge-disjoint H* if for every $k \in \mathbb{N}$ there is a family of k edge-disjoint subgraphs of G each of which is isomorphic to H, and our main result is the following.

Theorem 6.2.1. *Any graph that has arbitrarily many edge-disjoint double rays has infinitely many edge-disjoint double rays.*

Even for locally finite graphs this theorem does not follow from Halin's analogous result for vertex-disjoint double rays applied to the line graph. For example a double ray in the line graph may correspond, in the original graph, to a configuration as in Figure 6.1.

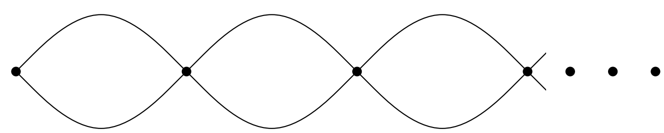

Figure 6.1: A graph that does not include a double ray but whose line graph does.

A related notion is that of ubiquity. A graph H is *ubiquitous* with respect to a graph relation $\leq$ if $nH \leq G$ for all $n \in \mathbb{N}$ implies $\aleph_0 H \leq G$, where nH denotes the disjoint union of n copies of H. For example, Halin's Theorem says that rays are ubiquitous with respect to the subgraph relation. It is known that not every graph is ubiquitous with respect to the minor relation [8], nor is every locally finite graph ubiquitous with respect to the subgraph relation [66, 102], or even the topological minor relation [8, 9]. However, Andreae has conjectured that every locally finite graph is ubiquitous with respect to the minor relation [8]. For more details see [9]. In Section 6.7 (the outlook) we introduce a notion closely related to ubiquity.

The proof is organised as follows. In Section 6.4 we explain how to deal with the cases that the graph has infinitely many ends, or an end with infinite vertex-degree. In Section 6.5 we consider the 'two ended' case: That in which there are two ends ω and ω' both of finite vertex-degree, and arbitrarily many edge-disjoint double rays from ω to ω'.

The only remaining case is the 'one ended' case: That in which there is a single end ω of finite vertex-degree and arbitrarily many edge-disjoint double rays from ω to ω. One central idea in the proof of this case is to consider 2-rays instead of double rays. Here a 2-ray is a pair of vertex-disjoint rays. For example, from each double ray one obtains a 2-ray by removing a finite path. The remainder of the proof is subdivided into two parts: In Subsection 6.6.3 we show that if there are arbitrarily many edge-disjoint 2-rays into ω, then there are infinitely many such 2-rays. In Subsection 6.6.2 we show that if there are infinitely many edge-disjoint 2-rays into ω, then there are infinitely many edge-disjoint double rays from ω to ω.

We finish by discussing the outlook and mentioning some open problems.

6.3 Preliminaries

All our basic notation for graphs is taken from [36]. In particular, two rays in a graph are equivalent if no finite set separates them. The equivalence classes of this relation are called the *ends* of G. We say that a ray in an end ω *converges* to ω. A double ray *converges* to all the ends of which it includes a ray.

6.3.1 The structure of a thin end

It follows from Halin's Theorem that if there are arbitrarily many vertex-disjoint rays in an end of G, then there are infinitely many such rays. This fact motivated the central definition of the *vertex-degree* of an end ω: the maximal cardinality of a set of vertex-disjoint rays in ω.

An end is *thin* if its vertex-degree is finite, and otherwise it is *thick*. A pair (A, B) of edge-disjoint subgraphs of G is a *separation* of G if $A \cup B = G$. The number of vertices of $A \cap B$ is called the *order* of the separation.

Definition 6.3.1. Let G be a locally finite graph and ω a thin end of G. A countable infinite sequence $((A_i, B_i))_{i\in\mathbb{N}}$ of separations of G *captures* ω if for all $i \in \mathbb{N}$

- $A_i \cap B_{i+1} = \emptyset$,
- $A_{i+1} \cap B_i$ is connected,
- $\bigcup_{i\in\mathbb{N}} A_i = G$,
- the order of (A_i, B_i) is the vertex-degree of ω, and
- each B_i contains a ray from ω.

Lemma 6.3.2. *Let G be a locally finite graph with a thin end ω. Then there is a sequence that captures ω.*

Proof. Without loss of generality G is connected, and so is countable. Let $v_1, v_2, \dots$ be an enumeration of the vertices of G. Let k be the vertex-degree of ω. Let $\mathcal{R} = \{R_1, \dots, R_k\}$ be a set of vertex-disjoint rays in ω and let S be the set of their start vertices. We pick a sequence $((A_i, B_i))_{i\in\mathbb{N}}$ of separations and a sequence (T_i) of connected subgraphs recursively as follows. We pick (A_i, B_i) such that S is included in A_i, such that there is a ray from ω included in B_i, and such that B_i does not meet $\bigcup_{j<i} T_j$ or $\{v_j \mid j \le i\}$: subject to this we minimise the size of the set X_i of vertices in $A_i \cap B_i$. Because of this minimization B_i is

connected and X_i is finite. We take T_i to be a finite connected subgraph of B_i including X_i. Note that any ray that meets all of the B_i must be in ω.

By Menger's Theorem [36] we get for each $i \in \mathbb{N}$ a set $\mathcal{P}_i$ of vertex-disjoint paths from X_i to X_{i+1} of size $|X_i|$. From these, for each i we get a set of $|X_i|$ vertex-disjoint rays in ω. Thus the size of X_i is at most k. On the other hand it is at least k as each ray R_j meets each set X_i.

Assume for contradiction that there is a vertex $v \in A_i \cap B_{i+1}$. Let R be a ray from v to ω inside B_{i+1}. Then R must meet X_i, contradicting the definition of B_{i+1}. Thus $A_i \cap B_{i+1}$ is empty.

Observe that $\bigcup \mathcal{P}_i \cup T_i$ is a connected subgraph of $A_{i+1} \cap B_i$ containing all vertices of X_i and X_{i+1}. For any vertex $v \in A_{i+1} \cap B_i$ there is a v–X_{i+1} path P in B_i. P meets B_{i+1} only in X_{i+1}. So P is included in $A_{i+1} \cap B_i$. Thus $A_{i+1} \cap B_i$ is connected. The remaining conditions are clear. □

Remark 6.3.3. Every infinite subsequence of a sequence capturing ω also captures ω. □

The following is obvious:

Remark 6.3.4. Let G be a graph and $v, w \in V(G)$ If G contains arbitrarily many edge-disjoint v–w paths, then it contains infinitely many edge-disjoint v–w paths. □

We will need the following special case of the theorem of Andreae mentioned in the Introduction.

Theorem 6.3.5 (Andreae [7])**.** *Let G be a graph and $v \in V(G)$. If there are arbitrarily many edge-disjoint rays all starting at v, then there are infinitely many edge-disjoint rays all starting at v.*

6.4 Known cases

Many special cases of Theorem 6.2.1 are already known or easy to prove. For example Halin showed the following.

Theorem 6.4.1 (Halin)**.** *Let G be a graph and ω an end of G. If ω contains arbitrarily many vertex-disjoint rays, then G has a half-grid as a minor.*

Corollary 6.4.2. *Any graph with an end of infinite vertex-degree has infinitely many edge-disjoint double rays.* □

Another simple case is the case where the graph has infinitely many ends.

Lemma 6.4.3. *A tree with infinitely many ends contains infinitely many edge-disjoint double rays.*

Proof. It suffices to show that every tree T with infinitely many ends contains a double ray such that removing its edges leaves a component containing infinitely many ends, since then one can pick those double rays recursively.

There is a vertex $v \in V(T)$ such that $T - v$ has at least 3 components C_1, C_2, C_3 that each have at least one end, as T contains more than 2 ends. Let e_i be the edge vw_i with $w_i \in C_i$ for $i \in \{1, 2, 3\}$. The graph $T \setminus \{e_1, e_2, e_3\}$ has precisely 4 components (C_1, C_2, C_3 and the one containing v), one of which, D say, has infinitely many ends. By symmetry we may assume that D is neither C_1 nor C_2. There is a double ray R all whose edges are contained in $C_1 \cup C_2 \cup \{e_1, e_2\}$. Removing the edges of R leaves the component D, which has infinitely many ends. □

Corollary 6.4.4. *Any connected graph with infinitely many ends has infinitely many edge-disjoint double rays.* □

6.5 The 'two ended' case

Using the results of Section 6.4 it is enough to show that any graph with only finitely many ends, each of which is thin, has infinitely many edge-disjoint double rays as soon as it has arbitrarily many edge-disjoint double rays. Any double ray in such a graph has to join a pair of ends (not necessarily distinct), and there are only finitely many such pairs. So if there are arbitrarily many edge-disjoint double rays, then there is a pair of ends such that there are arbitrarily many edge-disjoint double rays joining those two ends. In this section we deal with the case where these two ends are different, and in Section 6.6 we deal with the case that they are the same. We start with two preparatory lemmas.

Lemma 6.5.1. *Let G be a graph with a thin end ω, and let $\mathcal{R} \subseteq \omega$ be an infinite set. Then there is an infinite subset of $\mathcal{R}$ such that any two of its members intersect in infinitely many vertices.*

Proof. We define an auxilliary graph H with $V(H) = \mathcal{R}$ and an edge between two rays if and only if they intersect in infinitely many vertices. By Ramsey's Theorem either H contains an infinite clique or an infinite independent set of vertices. Let us show that there cannot be an infinite independent set in H. Let k be the vertex-degree of ω: we shall show that H does not have an independent set of size $k + 1$. Suppose for a contradiction that $X \subseteq \mathcal{R}$ is a set of $k + 1$ rays that is independent in H. Since any two rays in X meet in only finitely many vertices, each ray in X contains a tail that is disjoint to all the other rays in X. The set of these $k+1$ vertex-disjoint tails witnesses that ω has vertex-degree at least $k + 1$, a contradiction. Thus there is an infinite clique $K \subseteq H$, which is the desired infinite subset. □

Lemma 6.5.2. *Let G be a graph consisting of the union of a set $\mathcal{R}$ of infinitely many edge-disjoint rays of which any pair intersect in infinitely many vertices. Let $X \subseteq V(G)$ be an infinite set of vertices, then there are infinitely many edge-disjoint rays in G all starting in different vertices of X.*

Proof. If there are infinitely many rays in $\mathcal{R}$ each of which contains a different vertex from X, then suitable tails of these rays give the desired rays. Otherwise there is a ray $R \in \mathcal{R}$ meeting X infinitely often. In this case, we choose the desired rays recursively such that each contains a tail from some ray in $\mathcal{R} - R$. Having chosen finitely many such rays, we can always pick another: we start at some point in X on R which is beyond all the (finitely many) edges on R used so far. We follow R until we reach a vertex of some ray R' in $\mathcal{R} - R$ whose tail has not been used yet, then we follow R'. $\square$

Lemma 6.5.3. *Let G be a graph with only finitely many ends, all of which are thin. Let ω_1, ω_2 be distinct ends of G. If G contains arbitrarily many edge-disjoint double rays each of which converges to both ω_1 and ω_2, then G contains infinitely many edge-disjoint double rays each of which converges to both ω_1 and ω_2.*

Proof. For each pair of ends, there is a finite set separating them. The finite union of these finite sets is a finite set $S \subseteq V(G)$ separating any two ends of G. For $i = 1, 2$ let C_i be the component of $G - S$ containing ω_i.

There are arbitrarily many edge-disjoint double rays from ω_1 to ω_2 that have a common last vertex v_1 in S before staying in C_1 and also a common last vertex v_2 in S before staying in C_2. Note that v_1 may be equal to v_2. There are arbitrarily many edge-disjoint rays in $C_1 + v_1$ all starting in v_1. By Theorem 6.3.5 there is a countable infinite set $\mathcal{R}_1 = \{R_1^i \mid i \in \mathbb{N}\}$ of edge-disjoint rays each included in $C_1 + v_1$ and starting in v_1. By replacing $\mathcal{R}_1$ with an infinite subset of itself, if necessary, we may assume by Lemma 6.5.1 that any two members of $\mathcal{R}_1$ intersect in infinitely many vertices. Similarly, there is a countable infinite set $\mathcal{R}_2 = \{R_2^i \mid i \in \mathbb{N}\}$ of edge-disjoint rays each included in $C_2 + v_2$ and starting in v_2 such that any two members of $\mathcal{R}_2$ intersect in infinitely many vertices.

Let us subdivide all edges in $\bigcup \mathcal{R}_1$ and call the set of subdivision vertices X_1. Similarly, we subdivide all edges in $\bigcup \mathcal{R}_2$ and call the set of subdivision vertices X_2. Below we shall find double rays in the subdivided graph, which immediately give rise to the desired double rays in G.

Suppose for a contradiction that there is a finite set F of edges separating X_1 from X_2. Then v_i has to be on the same side of that separation as X_i as there are infinitely many v_i–X_i edges. So F separates v_1 from v_2, which contradicts the fact that there are arbitrarily many edge-disjoint double rays containing both v_1 and v_2. By Remark 6.3.4 there is a set $\mathcal{P}$ of infinitely many edge-disjoint X_1–X_2 paths. As all vertices in X_1 and X_2 have degree 2, and by taking an infinite subset if necessary, we may assume that each end-vertex of a path in $\mathcal{P}$ lies on no other path in $\mathcal{P}$.

By Lemma 6.5.2 there is an infinite set Y_1 of start-vertices of paths in $\mathcal{P}$ together with an infinite set $\mathcal{R}_1'$ of edge-disjoint rays with distinct start-vertices whose set of start-vertices is precisely Y_1. Moreover, we can ensure that each ray in $\mathcal{R}_1'$ is included in $\bigcup \mathcal{R}_1$. Let Y_2 be the set of end-vertices in X_2 of those paths in $\mathcal{P}$ that start in Y_1. Applying Lemma 6.5.2 again, we obtain an infinite set $Z_2 \subseteq Y_2$ together with an infinite set $\mathcal{R}_2'$ of edge-disjoint rays included in $\bigcup \mathcal{R}_2$ with distinct start-vertices whose set of start-vertices is precisely Z_2.

For each path P in $\mathcal{P}$ ending in Z_2, there is a double ray in the union of P and the two rays from $\mathcal{R}_1'$ and $\mathcal{R}_2'$ that P meets in its end-vertices. By construction, all these infinitely many double rays are edge-disjoint. Each of those double rays converges to both ω_1 and ω_2, since each ω_i is the only end in C_i. $\square$

Remark 6.5.4. Instead of subdividing edges we also could have worked in the line graph of G. Indeed, there are infinitely many vertex-disjoint paths in the line graph from $\bigcup \mathcal{R}_1$ to $\bigcup \mathcal{R}_2$.

6.6 The 'one ended' case

We are now going to look at graphs G that contain a thin end ω such that there are arbitrarily many edge-disjoint double rays converging only to the end ω. The aim of this section is to prove the following lemma, and to deduce Theorem 6.2.1.

Lemma 6.6.1. *Let G be a countable graph and let ω be a thin end of G. Assume there are arbitrarily many edge-disjoint double rays all of whose rays converge to ω. Then G has infinitely many edge-disjoint double rays.*

We promise that the assumption of countability will not cause problems later.

6.6.1 Reduction to the locally finite case

A key notion for this section is that of a 2-ray. A *2-ray* is a pair of vertex-disjoint rays. For example, from each double ray one obtains a 2-ray by removing a finite path.

In order to deduce that G has infinitely many edge-disjoint double rays, we will only need that G has arbitrarily many edge-disjoint 2-rays. In this subsection, we illustrate one advantage of 2-rays, namely that we may reduce to the case where G is locally finite.

Lemma 6.6.2. *Let G be a countable graph with a thin end ω. Assume there is a countable infinite set $\mathcal{R}$ of rays all of which converge to ω.*

Then there is a locally finite subgraph H of G with a single end which is thin such that the graph H includes a tail of any $R \in \mathcal{R}$.

Proof. Let $(R_i \mid i \in \mathbb{N})$ be an enumeration of $\mathcal{R}$. Let $(v_i \mid i \in \mathbb{N})$ be an enumeration of the vertices of G. Let U_i be the unique component of $G \setminus \{v_1, \ldots, v_i\}$ including a tail of each ray in ω.

For $i \in \mathbb{N}$, we pick a tail R'_i of R_i in U_i. Let $H_1 = \bigcup_{i \in \mathbb{N}} R'_i$. Making use of H_1, we shall construct the desired subgraph H. Before that, we shall collect some properties of H_1.

As every vertex of G lies in only finitely many of the U_i, the graph H_1 is locally finite. Each ray in H_1 converges to ω in G since $H_1 \setminus U_i$ is finite for every $i \in \mathbb{N}$. Let Ψ be the set of ends of H_1. Since ω is thin, Ψ has to be finite: $\Psi = \{\omega_1, \ldots, \omega_n\}$. For each $i \leq n$, we pick a ray $S_i \subseteq H_1$ converging to ω_i.

Now we are in a position to construct H. For any $i > 1$, the rays S_1 and S_i are joined by an infinite set $\mathcal{P}_i$ of vertex-disjoint paths in G. We obtain H from H_1 by adding all paths in the sets $\mathcal{P}_i$. Since H_1 is locally finite, H is locally finite.

It remains to show that every ray R in H is equivalent to S_1. If R contains infinitely many edges from the $\mathcal{P}_i$, then there is a single $\mathcal{P}_i$ which R meets infinitely, and thus R is equivalent to S_1. Thus we may assume that a tail of R is a ray in H_1. So it converges to some $\omega_i \in \Psi$. Since S_i and S_1 are equivalent, R and S_1 are equivalent, which completes the proof. $\square$

Corollary 6.6.3. *Let G be a countable graph with a thin end ω and arbitrarily many edge-disjoint 2-rays of which all the constituent rays converge to ω. Then there is a locally finite subgraph H of G with a single end, which is thin, such that H has arbitrarily many edge-disjoint 2-rays.*

Proof. By Lemma 6.6.2 there is a locally finite graph $H \subseteq G$ with a single end such that a tail of each of the constituent rays of the arbitrarily many 2-rays is included in H. $\square$

6.6.2 Double rays versus 2-rays

A connected subgraph of a graph G including a vertex set $S \subseteq V(G)$ is a *connector* of S in G.

Lemma 6.6.4. *Let G be a connected graph and S a finite set of vertices of G. Let $\mathcal{H}$ be a set of edge-disjoint subgraphs H of G such that each connected component of H meets S. Then there is a finite connector T of S, such that at most $2|S| - 2$ graphs from $\mathcal{H}$ contain edges of T.*

Proof. By replacing $\mathcal{H}$ with the set of connected components of graphs in $\mathcal{H}$, if necessary, we may assume that each member of $\mathcal{H}$ is connected. We construct graphs T_i recursively for $0 \leq i < |S|$ such that each T_i is finite and has at most $|S| - i$ components, at most $2i$ graphs from $\mathcal{H}$ contain edges of T_i, and each component of T_i meets S. Let $T_0 = (S, \emptyset)$ be the graph with vertex set S and no edges. Assume that T_i has been defined.

If T_i is connected let $T_{i+1} = T_i$. For a component C of T_i, let C' be the graph obtained from C by adding all graphs from $\mathcal{H}$ that meet C.

As G is connected, there is a path P (possibly trivial) in G joining two of these subgraphs C_1' and C_2' say. And by taking the length of P minimal, we may assume that P does not contain any edge from any $H \in \mathcal{H}$. Then we can extend P to a C_1–C_2 path Q by adding edges from at most two subgraphs from $\mathcal{H}$ — one included in C_1' and the other in C_2'. We obtain T_{i+1} from T_i by adding Q.

$T = T_{|S|-1}$ has at most one component and thus is connected. And at most $2|S| - 2$ many graphs from $\mathcal{H}$ contain edges of T. Thus T is as desired. $\square$

Let d, d' be 2-rays. d is a *tail* of d' if each ray of d is a tail of a ray of d'. A set D' is a *tailor* of a set D of 2-rays if each element of D' is a tail of some element of D but no 2-ray in D includes more than one 2-ray in D'.

Lemma 6.6.5. *Let G be a locally finite graph with a single end ω, which is thin. Assume that G contains an infinite set $D = \{d_1, d_2, \dots\}$ of edge-disjoint 2-rays.*

Then G contains an infinite tailor D' of D and a sequence $((A_i, B_i))_{i\in\mathbb{N}}$ capturing ω (see Definition 6.3.1) such that there is a family of vertex-disjoint connectors T_i of $A_i \cap B_i$ contained in $A_{i+1} \cap B_i$, each of which is edge-disjoint from each member of D'.

Proof. Let k be the vertex-degree of ω. By Lemma 6.3.2 there is a sequence $((A'_i, B'_i))_{i\in\mathbb{N}}$ capturing ω. By replacing each 2-ray in D with a tail of itself if necessary, we may assume that for all $(r, s) \in D$ and $i \in \mathbb{N}$ either both r and s meet A'_i or none meets A'_i. By Lemma 6.6.4 there is a finite connector T'_i of $A'_i \cap B'_i$ in the connected graph B'_i which meets in an edge at most $2k-2$ of the 2-rays of D that have a vertex in A'_i.

Thus, there are at most $2k-2$ 2-rays in D that meet all but finitely many of the T'_i in an edge. By throwing away these finitely many 2-rays in D we may assume that each 2-ray in D is edge-disjoint from infinitely many of the T'_i. So we can recursively build a sequence $N_1, N_2, \dots$ of infinite sets of natural numbers such that $N_i \supseteq N_{i+1}$, the first i elements of N_i are all contained in N_{i+1}, and d_i only meets finitely many of the T'_j with $j \in N_i$ in an edge. Then $N = \bigcap_{i\in\mathbb{N}} N_i$ is infinite and has the property that each d_i only meets finitely many of the T'_j with $j \in N$ in an edge. Thus there is an infinite tailor D' of D such that no 2-ray from D' meets any T'_j for $j \in N$ in an edge.

We recursively define a sequence $n_1, n_2, \dots$ of natural numbers by taking $n_i \in N$ sufficiently large that B'_{n_i} does not meet T'_{n_j} for any $j < i$. Taking $(A_i, B_i) = (A'_{n_i}, B'_{n_i})$ and $T_i = T'_{n_i}$ gives the desired sequences. $\square$

Lemma 6.6.6. *If a locally finite graph G with a single end ω which is thin contains infinitely many edge-disjoint 2-rays, then G contains infinitely many edge-disjoint double rays.*

Proof. Applying Lemma 6.6.5 we get an infinite set D of edge-disjoint 2-rays, a sequence $((A_i, B_i))_{i\in\mathbb{N}}$ capturing ω, and connectors T_i of $A_i \cap B_i$ for each $i \in \mathbb{N}$ such that the T_i are vertex-disjoint from each other and edge-disjoint from all members of D.

We shall construct the desired set of infinitely many edge-disjoint double rays as a nested union of sets D_i. We construct the D_i recursively. Assume that a set D_i of i edge-disjoint double rays has been defined such that each of its members is included in the union of a single 2-ray from D and one connector T_j. Let $d_{i+1} \in D$ be a 2-ray distinct from the finitely many 2-rays used so far. Let C_{i+1} be one of the infinitely many connectors that is different from all the finitely many connectors used so far and that meets both rays of d_{i+1}. Clearly, $d_{i+1} \cup C_{i+1}$ includes a double ray R_{i+1}. Let $D_{i+1} = D_i \cup \{R_{i+1}\}$. The union $\bigcup_{i\in\mathbb{N}} D_i$ is an infinite set of edge-disjoint double rays as desired. $\square$

6.6.3 Shapes and allowed shapes

Let G be a graph and (A, B) a separation of G. A *shape* for (A, B) is a word $v_1x_1v_2x_2 \dots x_{n-1}v_n$ with $v_i \in A \cap B$ and $x_i \in \{l, r\}$ such that no vertex appears twice. We call the v_i the *vertices* of the shape. Every ray R *induces* a shape $\sigma = \sigma_R(A, B)$ on every separation (A, B) of finite order in the following way: Let $<_R$ be the *natural order* on $V(R)$ induced by the ray, where $v <_R w$ if w lies in the unique infinite component of $R - v$. The vertices of σ are those vertices of R that lie in $A \cap B$ and they appear in σ in the order given by $<_R$. For v_i, v_{i+1} the path v_iRv_{i+1} has edges only in A or only in B but not in both. In the first case we put l between v_i and v_{i+1} and in the second case we put r between v_i and v_{i+1}.

Let $(A_1, B_1), (A_2, B_2)$ be separations with $A_1 \cap B_2 = \emptyset$ and thus also $A_1 \subseteq A_2$ and $B_2 \subseteq B_1$. Let σ_i be a nonempty shape for (A_i, B_i). The word $\tau = v_1x_1v_2 \dots x_{n-1}v_n$ is an *allowed shape linking* σ_1 *to* σ_2 with *vertices* $v_1 \dots v_n$ if the following holds.

- v is a vertex of τ if and only if it is a vertex of σ_1 or σ_2,
- if v appears before w in σ_i, then v appears before w in τ,
- v_1 is the initial vertex of σ_1 and v_n is the terminal vertex of σ_2,
- $x_i \in \{l, m, r\}$,
- the subword vlw appears in τ if and only if it appears in σ_1,
- the subword vrw appears in τ if and only if it appears in σ_2,
- $v_i \neq v_j$ for $i \neq j$.

Each ray R defines a word $\tau = \tau_R[(A_1, B_1), (A_2, B_2)] = v_1x_1v_2 \dots x_{n-1}v_n$ with vertices v_i and $x_i \in \{l, m, r\}$ as follows. The vertices of τ are those vertices of R that lie in $A_1 \cap B_1$ or $A_2 \cap B_2$ and they appear in τ in the order given by $<_R$. For v_i, v_{i+1} the path v_iRv_{i+1} has edges either only in A_1, only in $A_2 \cap B_1$, or only in B_2. In the first case we set $x_i = l$ and τ contains the subword v_ilv_{i+1}. In the second case we set $x_i = m$ and τ contains the subword v_imv_{i+1}. In the third case we set $x_i = r$ and τ contains the subword v_irv_{i+1}.

For a ray R to induce an allowed shape $\tau_R[(A_1, B_1), (A_2, B_2)]$ we need at least that R starts in A_2. However, each ray in ω has a tail such that whenever it meets an A_i it also starts in that A_i. Let us call such rays *lefty*. A 2-ray is *lefty* if both its rays are.

Remark 6.6.7. Let (A_1, B_1), and (A_2, B_2) be two separations of finite order with $A_1 \subseteq A_2$, and $B_2 \subseteq B_1$. For every lefty ray R meeting A_1, the word $\tau_R[(A_1, B_1), (A_2, B_2)]$ is an allowed shape linking $\sigma_R(A_1, B_1)$ and $\sigma_R(A_2, B_2)$. $\square$

From now on let us fix a locally finite graph G with a thin end ω of vertex-degree k. And let $((A_i, B_i))_{i\in\mathbb{N}}$ be a sequence capturing ω such that each member has order k.

A 2-*shape* for a separation (A, B) is a pair of shapes for (A, B). Every 2-ray induces a 2-shape coordinatewise in the obvious way. Similarly, an *allowed* 2-*shape* is a pair of allowed shapes.

Clearly, there is a global constant $c_1 \in \mathbb{N}$ depending only on k such that there are at most c_1 distinct 2-shapes for each separation (A_i, B_i). Similarly, there is a global constant $c_2 \in \mathbb{N}$ depending only on k such that for all $i, j \in \mathbb{N}$ there are at most c_2 distinct allowed 2-shapes linking a 2-shape for (A_i, B_i) with a 2-shape for (A_j, B_j).

For most of the remainder of this subsection we assume that for every $i \in \mathbb{N}$ there is a set D_i consisting of at least $c_1 \cdot c_2 \cdot i$ edge-disjoint 2-rays in G. Our aim will be to show that in these circumstances there must be infinitely many edge-disjoint 2-rays.

By taking a tailor if necessary, we may assume that every 2-ray in each D_i is lefty.

Lemma 6.6.8. *There is an infinite set $J \subseteq \mathbb{N}$ and, for each $i \in \mathbb{N}$, a tailor D'_i of D_i of cardinality $c_2 \cdot i$ such that for all $i \in \mathbb{N}$ and $j \in J$ all 2-rays in D'_i induce the same 2-shape $\sigma[i, j]$ on (A_j, B_j).*

Proof. We recursively build infinite sets $J_i \subseteq \mathbb{N}$ and tailors D'_i of D_i such that for all $k \leq i$ and $j \in J_i$ all 2-rays in D'_k induce the same 2-shape on (A_j, B_j). For all $i \geq 1$, we shall ensure that J_i is an infinite subset of J_{i-1} and that the $i - 1$ smallest members of J_i and J_{i-1} are the same. We shall take J to be the intersection of all the J_i.

Let $J_0 = \mathbb{N}$ and let D'_0 be the empty set. Now, for some $i \geq 1$, assume that sets J_k and D'_k have been defined for all $k < i$. By replacing 2-rays in D_i by their tails, if necessary, we may assume that each 2-ray in D_i avoids A_ℓ, where ℓ is the $(i - 1)$st smallest value of J_{i-1}. As D_i contains $c_1 \cdot c_2 \cdot i$ many 2-rays, for each $j \in J_{i-1}$ there is a set $S_j \subseteq D_i$ of size at least $c_2 \cdot i$ such that each 2-ray in S_j induces the same 2-shape on (A_j, B_j). As there are only finitely many possible choices for S_j, there is an infinite subset J_i of J_{i-1} on which S_j is constant. For D'_i we pick this value of S_j. Since each $d \in D'_i$ induces the empty 2-shape on each (A_k, B_k) with $k \leq \ell$ we may assume that the first $i - 1$ elements of J_{i-1} are also included in J_i.

It is immediate that the set $J = \bigcap_{i\in\mathbb{N}} J_i$ and the D'_i have the desired property. $\square$

Lemma 6.6.9. *There are two strictly increasing sequences $(n_i)_{i\in\mathbb{N}}$ and $(j_i)_{i\in\mathbb{N}}$ with $n_i \in \mathbb{N}$ and $j_i \in J$ for all $i \in \mathbb{N}$ such that $\sigma[n_i, j_i] = \sigma[n_{i+1}, j_i]$ and $\sigma[n_i, j_i]$ is not empty.*

Proof. Let H be the graph on $\mathbb{N}$ with an edge $vw \in E(H)$ if and only if there are infinitely many elements $j \in J$ such that $\sigma[v, j] = \sigma[w, j]$.

As there are at most c_1 distinct 2-shapes for any separator (A_i, B_i), there is no independent set of size $c_1 + 1$ in H and thus no infinite one. Thus, by

Ramsey's theorem, there is an infinite clique in H. We may assume without loss of generality that H itself is a clique by moving to a subsequence of the D'_i if necessary. With this assumption we simply pick $n_i = i$.

Now we pick the j_i recursively. Assume that j_i has been chosen. As i and $i+1$ are adjacent in H, there are infinitely many indicies $\ell \in \mathbb{N}$ such that $\sigma[i,\ell] = \sigma[i+1,\ell]$. In particular, there is such an $\ell > j_i$ such that $\sigma[i+1,\ell]$ is not empty. We pick j_{i+1} to be one of those ℓ.

Clearly, $(j_i)_{i\in\mathbb{N}}$ is an increasing sequence and $\sigma[i,j_i] = \sigma[i+1,j_i]$ as well as $\sigma[i,j_i]$ is non-empty for all $i \in \mathbb{N}$, which completes the proof. $\square$

By moving to a subsequence of (D'_i) and $((A_j,B_j))$, if necessary, we may assume by Lemma 6.6.8 and Lemma 6.6.9 that for all $i,j \in \mathbb{N}$ all $d \in D'_i$ induce the same 2-shape $\sigma[i,j]$ on (A_j,B_j), and that $\sigma[i,i] = \sigma[i+1,i]$, and that $\sigma[i,i]$ is non-empty.

Lemma 6.6.10. *For all $i \in \mathbb{N}$ there is $D''_i \subseteq D'_i$ such that $|D''_i| = i$, and all $d \in D''_i$ induce the same allowed 2-shape $\tau[i]$ that links $\sigma[i,i]$ and $\sigma[i,i+1]$.*

Proof. Note that it is in this proof that we need all the 2-rays in D''_i to be lefty as they need to induce an allowed 2-shape that links $\sigma[i,i]$ and $\sigma[i,i+1]$ as soon as they contain a vertex from A_i. As $|D'_i| \geq i \cdot c_2$ and as there are at most c_2 many distinct allowed 2-shapes that link $\sigma[i,i]$ and $\sigma[i,i+1]$ there is $D''_i \subseteq D'_i$ with $|D''_i| = i$ such that all $d \in D''_i$ induce the same allowed 2-shape. $\square$

We enumerate the elements of D''_j as follows: $d^j_1, d^j_2, \ldots, d^j_j$. Let (s^j_i, t^j_i) be a representation of d^j_i. Let $S^j_i = s^j_i \cap A_{j+1} \cap B_j$, and let $\mathcal{S}_i = \bigcup_{j\geq i} S^j_i$. Similarly, let $T^j_i = t^j_i \cap A_{j+1} \cap B_j$, and let $\mathcal{T}_i = \bigcup_{j\geq i} T^j_i$.

Clearly, $\mathcal{S}_i$ and $\mathcal{T}_i$ are vertex-disjoint and any two graphs in $\bigcup_{i\in\mathbb{N}}\{\mathcal{S}_i, \mathcal{T}_i\}$ are edge-disjoint. We shall find a ray R_i in each of the $\mathcal{S}_i$ and a ray R'_i in each of the $\mathcal{T}_i$. The infinitely many pairs (R_i, R'_i) will then be edge-disjoint 2-rays, as desired.

Lemma 6.6.11. *Each vertex v of $\mathcal{S}_i$ has degree at most 2. If v has degree 1 it is contained in $A_i \cap B_i$.*

Proof. Clearly, each vertex v of $\mathcal{S}_i$ that does not lie in any separator $A_j \cap B_j$ has degree 2, as it is contained in precisely one S^j_i, and all the leaves of S^j_i lie in $A_j \cap B_j$ and $A_{j+1} \cap B_{j+1}$ as d^j_i is lefty. Indeed, in S^j_i it is an inner vertex of a path and thus has degree 2 in there. If v lies in $A_i \cap B_i$ it has degree at most 2, as it is only a vertex of S^j_i for one value of j, namely $j = i$.

Hence, we may assume that $v \in A_j \cap B_j$ for some $j > i$. Thus, $\sigma[j,j]$ contains v and $l : \sigma[j,j] : r$ contains precisely one of the four following subwords:

$$lvl, lvr, rvl, rvr$$

(Here we use the notation $p : q$ to denote the concatenation of the word p with the word q.) In the first case $\tau[j-1]$ contains mvm as a subword and $\tau[j]$ has no m adjacent to v. Then S^{j-1}_i contains precisely 2 edges adjacent to v and S^j_i

has no such edge. The fourth case is the first one with l and r and j and $j-1$ interchanged.

In the second and third cases, each of $\tau[j-1]$ and $\tau[j]$ has precisely one m adjacent to v. So both S_i^{j-1} and S_i^j contain precisely 1 edge adjacent to v.

As v appears only as a vertex of S_i^ℓ for $\ell = j$ or $\ell = j-1$, the degree of v in $\mathcal{S}_i$ is 2. □

Lemma 6.6.12. *There are an odd number of vertices in $\mathcal{S}_i$ of degree* 1*.*

Proof. By Lemma 6.6.11 we have that each vertex of degree 1 lies in $A_i \cap B_i$. Let v be a vertex in $A_i \cap B_i$. Then, $\sigma[i,i]$ contains v and $l : \sigma[i,i] : r$ contains precisely one of the four following subwords:

$$lvl, lvr, rvl, rvr$$

In the first and fourth case v has even degree. It has degree 1 otherwise. As $l : \sigma[i,i] : r$ starts with l and ends with r, the word lvr appear precisely once more than the word rvl. Indeed, between two occurrences of lvr there must be one of rvl and vice versa. Thus, there are an odd number of vertices with degree 1 in $\mathcal{S}_i$. □

Lemma 6.6.13. *$\mathcal{S}_i$ includes a ray.*

Proof. By Lemma 6.6.11 every vertex of $\mathcal{S}_i$ has degree at most 2 and thus every component of $\mathcal{S}_i$ has at most two vertices of degree 1. By Lemma 6.6.12 $\mathcal{S}_i$ has a component C that contains an odd number of vertices with degree 1. Thus C has precisely one vertex of degree 1 and all its other vertices have degree 2, thus C is a ray. □

Corollary 6.6.14. *G contains infinitely many edge-disjoint* 2*-rays.*

Proof. By symmetry, Lemma 6.6.13 is also true with $\mathcal{T}_i$ in place of $\mathcal{S}_i$. Thus $\mathcal{S}_i \cup \mathcal{T}_i$ includes a 2-ray X_i. The X_i are edge-disjoint by construction. □

Recall that Lemma 6.6.1 states that a countable graph with a thin end ω and arbitrarily many edge-disjoint double rays all whose subrays converge to ω, also has infinitely many edge-disjoint double rays. We are now in a position to prove this lemma.

Proof of Lemma 6.6.1. By Lemma 6.6.6 it suffices to show that G contains a subgraph H with a single end which is thin such that H has infinitely many edge-disjoint 2-rays. By Corollary 6.6.3, G has a subgraph H with a single end which is thin such that H has arbitrarily many edge-disjoint 2-rays. But then by the argument above H contains infinitely many edge-disjoint 2-rays, as required. □

With these tools at hand, the remaining proof of Theorem 6.2.1 is easy. Let us collect the results proved so far to show that each graph with arbitrarily many edge-disjoint double rays also has infinitely many edge-disjoint double rays.

Proof of Theorem 6.2.1. Let G be a graph that has a set D_i of i edge-disjoint double rays for each $i \in \mathbb{N}$. Clearly, G has infinitely many edge-disjoint double rays if its subgraph $\bigcup_{i \in \mathbb{N}} D_i$ does, and thus we may assume without loss of generality that $G = \bigcup_{i \in \mathbb{N}} D_i$. In particular, G is countable.

By Corollary 6.4.4 we may assume that each connected component of G includes only finitely many ends. As each component includes a double ray we may assume that G has only finitely many components. Thus, there is one component containing arbitrarily many edge-disjoint double rays, and thus we may assume that G is connected.

By Corollary 6.4.2 we may assume that all ends of G are thin. Thus, as mentioned at the start of Section 6.5, there is a pair of ends (ω, ω') of G (not necessarily distinct) such that G contains arbitrarily many edge-disjoint double rays each of which converges precisely to ω and ω'. This completes the proof as, by Lemma 6.5.3 G has infinitely many edge-disjoint double rays if ω and ω' are distinct and by Lemma 6.6.1 G has infinitely many edge-disjoint double rays if $\omega = \omega'$. □

6.7 Outlook and open problems

We will say that a graph H is *edge-ubiquitous* if every graph having arbitrarily many edge-disjoint H also has infinitely many edge-disjoint H.

Thus Theorem 6.2.1 can be stated as follows: the double ray is edge-ubiquitous. Andreae's Theorem implies that the ray is edge-ubiquitous. And clearly, every finite graph is edge-ubiquitous.

We could ask which other graphs are edge-ubiquitous. It follows from our result that the 2-ray is edge-ubiquitous. Let G be a graph in which there are arbitrarily many edge-disjoint 2-rays. Let $v * G$ be the graph obtained from G by adding a vertex v adjacent to all vertices of G. Then $v * G$ has arbitrarily many edge-disjoint double rays, and thus infinitely many edge-disjoint double rays. Each of these double rays uses v at most once and thus includes a 2-ray of G.

The vertex-disjoint union of k rays is called a *k-ray*. The k-ray is edge-ubiquitous. This can be proved with an argument similar to that for Theorem 6.2.1: Let G be a graph with arbitrarily many edge-disjoint k-rays. The same argument as in Corollaries 6.4.4 and 6.4.2 shows that we may assume that G has only finitely many ends, each of which is thin. By removing a finite set of vertices if necessary we may assume that each component of G has at most one end, which is thin. Now we can find numbers k_C indexed by the components C of G and summing to k such that each component C has arbitrarily many edge-disjoint k_C-rays. Hence, we may assume that G has only a single end, which is thin. By Lemma 6.6.2 we may assume that G is locally finite.

In this case, we use an argument as in Subsection 6.6.3. It is necessary to use k-shapes instead of 2-shapes but other than that we can use the same combinatorial principle. If C_1 and C_2 are finite sets, a (C_1, C_2)-*shaping* is a pair (c_1, c_2) where c_1 is a partial colouring of $\mathbb{N}$ with colours from C_1 which is defined at all but finitely many numbers and c_2 is a colouring of $\mathbb{N}^{(2)}$ with colours from C_2 (in our argument above, C_1 would be the set of all k-shapes and C_2 would be the set of all allowed k-shapes for all pairs of k-shapes).

Lemma 6.7.1. *Let $D_1, D_2, \dots$ be a sequence of sets of (C_1, C_2)-shapings where D_i has size i. Then there are strictly increasing sequences $i_1, i_2, \dots$ and $j_1, j_2, \dots$ and subsets $S_n \subseteq D_{i_n}$ with $|S_n| \geq n$ such that*

- *for any $n \in \mathbb{N}$ all the values of $c_1(j_n)$ for the shapings $(c_1, c_2) \in S_{n-1} \cup S_n$ are equal (in particular, they are all defined).*
- *for any $n \in N$, all the values of $c_2(j_n, j_{n+1})$ for the shapings $(c_1, c_2) \in S_n$ are equal.*

Lemma 6.7.1 can be proved by the same method with which we constructed the sets D_i'' from the sets D_i. The advantage of Lemma 6.7.1 is that it can not only be applied to 2-rays but also to more complicated graphs like k-rays.

A *talon* is a tree with a single vertex of degree 3 where all the other vertices have degree 2. An argument as in Subsection 6.6.2 can be used to deduce that talons are edge-ubiquitous from the fact that 3-rays are. However, we do not know whether the graph in Figure 6.2 is edge-ubiquitous.

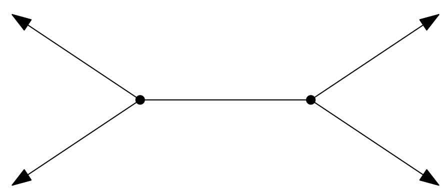

Figure 6.2: A graph obtained from 2 disjoint double rays, joined by a single edge. Is this graph edge-ubiquitous?

We finish with the following open problem.

Question 6.7.2. *Is the directed analogue of Theorem 6.2.1 true? More precisely: Is it true that if a directed graph has arbitrarily many edge-disjoint directed double rays, then it has infinitely many edge-disjoint directed double rays?*

It should be noted that if true the directed analogue would be a common generalization of Theorem 6.2.1 and the fact that double rays are ubiquitous with respect to the subgraph relation.

Chapter 7

The colouring number of infinite graphs

7.1 Abstract

We show that, given an infinite cardinal μ, a graph has colouring number at most μ if and only if it contains neither of two types of subgraph. We also show that every graph with infinite colouring number has a well-ordering of its vertices that simultaneously witnesses its colouring number and its cardinality.

7.2 Introduction

Our point of departure is a recent article by Péter Komjáth [63] one of whose results addresses infinite graphs with infinite colouring number. Recall

Definition 7.2.1. The *colouring number* $\mathrm{col}(G)$ of a graph $G = (V, E)$ is the smallest cardinal κ such that there exists a well-ordering $<^*$ of V with

$$|N(v) \cap \{w \mid w <^* v\}| < \kappa \qquad \text{for all } v \in V,$$

where $N(v)$ is the set of neighbours of v. We call such well-orderings *good*.

This notion was introduced by Erdős and Hajnal in [40].

What Komjáth proved in [63] is that if the colouring number of a graph G is bigger than some infinite cardinal μ, then G contains either a K_μ, i.e., μ mutually adjacent vertices, or G contains for each positive integer k an induced copy of the complete bipartite graph $K_{k,k}$. This condition is not a characterisation: there are graphs, such as $K_{\omega,\omega}$, which have small colouring number but nevertheless include an induced $K_{k,k}$ for each k.

Since having colouring number $\leq \mu$ is closed not only under taking induced subgraphs but even under taking subgraphs, it seems easier to look first for a characterisation in terms of forbidden subgraphs. When playing with the ideas appearing in Komjáth's proof, we realized that they can be used to give just such a transparent characterization of "having colouring number $\leq \mu$" in terms of forbidden subgraphs. For some explicit graphs called *μ-obstructions*, to be introduced in Definition 7.3.1 below, we shall prove

Theorem 7.2.2. *Let G be a graph and let μ denote some infinite cardinal. Then the statement $\operatorname{col}(G) > \mu$ is equivalent to G containing some μ-obstruction as a subgraph.*

The proof we describe has an interesting consequence:

Theorem 7.2.3. *Every graph G whose colouring number is infinite possesses a good well-ordering of length $|V(G)|$.*

It is not hard to re-obtain the result of Komjáth mentioned above from our characterisation 7.2.2 by inspecting whether the μ-obstructions satisfy it. In fact, one can easily obtain the following strengthening:

Theorem 7.2.4. *If G is a graph with $\operatorname{col}(G) > \mu$, where μ denotes some infinite cardinal, then G contains either a K_μ or, for each positive integer k, an induced $K_{k,\omega}$.*

We will also give an example in Section 2 demonstrating that the conclusion cannot be improved further to the presence of an induced $K_{\omega,\omega}$. Which complete bipartite graphs exactly one gets by this approach depends on which properties the relevant cardinals have in the partition calculus.

For standard set-theoretical background we refer to Kunen's textbook [65].

7.3 Obstructions

Throughout this section, we fix an infinite cardinal μ. There are two kinds of μ-obstructions relevant for the condition $\operatorname{col}(G) > \mu$ in Theorem 7.2.2. They are introduced next.

Definition 7.3.1. (1) A *μ-obstruction of type I* is a bipartite graph H with bipartition (A, B) such that for some cardinal $\lambda \geq \mu$ we have

- $|A| = \lambda$, $|B| = \lambda^+$,
- every vertex of B has at least μ neighbours in A, and
- every vertex of A has λ^+ neighbours in B.

(2) Let $\kappa > \mu$ be regular, and let G be a graph with $V(G) = \kappa$. Define T_G to be the set of those $\alpha \in \kappa$ with the following properties:

- $\operatorname{cf}(\alpha) = \operatorname{cf}(\mu)$
- The order type of $N(\alpha) \cap \alpha$ is μ.

- The supremum of $N(\alpha) \cap \alpha$ is α.

If T_G is stationary in κ, then G is a *μ-obstruction of type II.* We also call graphs isomorphic to such graphs μ-obstructions of type II.

Now we can directly proceed to the easier direction of Theorem 7.2.2.

Proposition 7.3.2. *If a graph G has a μ-obstruction of either type as a subgraph, then* $\mathrm{col}(G) > \mu$.

Proof. Suppose first that G contains a μ-obstruction of type I, say with bipartition (A, B) as in Definition 7.3.1 above, and $|A| = \lambda \geq \mu$. Assume for a contradiction that there is a good well-ordering of G. Thus every $b \in B$ has a neighbour in A above it in that well-ordering. For $a \in A$, we denote by X_a the set of those neighbours of a that are below a in the well-ordering. Hence $B = \bigcup_{a \in A} X_a$. Since all the X_a have size less than μ, we deduce that $|B| \leq \lambda$, which is the desired contradiction.

In the second case, we may without loss of generality assume that G itself is an obstruction of type II. Again we suppose for a contradiction that there is a good well-ordering $<^*$ of $V(G)$. Notice that each $\alpha \in T_G$ has a neighbour $\beta < \alpha$ such that $\alpha <^* \beta$. Let $f \colon T_G \longrightarrow \kappa$ be a function sending each α to some such β. By Fodor's Lemma, there must be some $\beta < \kappa$ such that

$$T = \{\alpha \in T_G \mid f(\alpha) = \beta\}$$

is stationary. Now every element of T is a neighbour of β, and β comes after T in the ordering $<^*$, which in view of $|T| = \kappa > \mu$ contradicts our assumption that this ordering is good. □

We say that a graph is *μ-unobstructed* if it contains no μ-obstruction of either type. To complete the proof of Theorem 7.2.2 we still need to show that every μ-unobstructed graph G satisfies $\mathrm{col}(G) \leq \mu$. This will be the objective of Sections 3 and 4.

In the remainder of this section, we prove two results asserting that in order to find an obstruction in a given graph G it suffices to find something weaker.

Definition 7.3.3. A *μ-barricade* is bipartite graph with bipartition (A, B) such that

- $|A| < |B|$,
- and every vertex of B has at least μ neighbours in A.

Lemma 7.3.4. *If G has a μ-barricade as a subgraph, then it also has a μ-obstruction of type I as a subgraph.*

Proof. Let H with bipartition (A, B) be a barricade which is a subgraph of G, chosen so that $\lambda = |A|$ is minimal. By deleting some vertices of B if necessary, we may assume that B has cardinality λ^+. Let A' be the set of $a \in A$ for which $N_B(a)$ is of size λ^+, and let B' be the set of elements of B with no neighbour in

$A \setminus A'$. By the definition of A', there are at most λ edges ab with $a \in A \setminus A'$ and $b \in B$. So $B \setminus B'$ is of size at most λ. It follows that B' has cardinality λ^+. In particular, the subgraph H' of H on (A', B') is a barricade, so by minimality of $|A|$ we have $|A'| = \lambda$. Since by construction every vertex of A' has λ^+ neighbours in B and hence in B', the subgraph H' is a μ-obstruction of type I. □

Definition 7.3.5. Let $\kappa > \mu$ be regular. A graph G with set of vertices κ is said to be a *μ-ladder* if there is a stationary set T such that each $\alpha \in T$ has at least μ neighbours in α. Also, every graph isomorphic to such a graph is regarded as a μ-ladder.

Lemma 7.3.6. *Every graph containing a μ-ladder is μ-obstructed.*

Proof. It suffices to prove that every μ-ladder is μ-obstructed. So let G with $V(G) = \kappa$ be as described in the previous definition. For each $\alpha \in T$ we let the sequence $\langle \alpha_i \,|\, i < \mu \rangle$ enumerate the μ smallest neighbours of α in increasing order and denote the limit point of this sequence by $f(\alpha)$. Clearly we have $f(\alpha) \leq \alpha$ and $\mathrm{cf}\big(f(\alpha)\big) = \mathrm{cf}(\mu)$ for all $\alpha \in T$.

Let us first suppose that the set

$$T' = \{\alpha \in T \,|\, f(\alpha) < \alpha\}$$

is stationary in κ. Then for some $\gamma < \kappa$ the set

$$B = \{\alpha \in T' \,|\, f(\alpha) = \gamma\}$$

is stationary and as $|\gamma| < \kappa = |B|$ the pair (γ, B) is a μ-barricade in G. Due to Lemma 7.3.4 it follows that G contains a μ-obstruction of type I.

So it remains to consider the case that

$$T'' = \{\alpha \in T \,|\, f(\alpha) = \alpha\}$$

is stationary in κ. In that case we have $N(\alpha) \cap \alpha = \{\alpha_i \,|\, i < \mu\}$ for all $\alpha \in T''$. So T_G is a superset of T'' and thus stationary, meaning that G is a μ-obstruction of type II. □

7.4 The regular case

In this and the next section we shall prove the harder part of Theorem 7.2.2, in such a way that Theorem 7.2.3 is also immediate. To this end we shall show

Theorem 7.4.1. *Let G denote an infinite graph of order κ and let μ be an infinite cardinal. Then at least one of the following three cases occurs:*

- *G has a subgraph H with $|V(H)| < |V(G)|$ and $\mathrm{col}(H) > \mu$.*
- *G is μ-obstructed.*
- *G has a good well-ordering of length κ exemplifying $\mathrm{col}(G) \leq \mu$.*

Suppose for a moment that we know this. To deduce Theorem 7.2.2 we consider any graph with $\operatorname{col}(G) > \mu$. Let G^* be subgraph of G with $\operatorname{col}(G^*) > \mu$ and subject to this with $|V(G^*)|$ as small as possible. Then G^* is still infinite and when we apply Theorem 7.4.1 to G^* the first and third outcome are impossible, so the second one most occur. Thus G^* and hence G contains a μ-obstruction, as desired. To obtain Theorem 7.2.3 we apply Theorem 7.4.1 to G with $\mu = \operatorname{col}(G)$.

The proof of Theorem 7.4.1 itself is divided into two cases according to whether κ is regular or singular. The former case will be treated immediately and the latter case is deferred to the next section.

Proof of Theorem 7.4.1 when κ is regular. Let $V(G) = \kappa$ and consider the set

$$T = \{\alpha < \kappa \,|\, \text{Some } \beta \geq \alpha \text{ has at least } \mu \text{ neighbours in } \alpha\}\,.$$

First Case: T is not stationary in κ.

Let $\langle \delta_i \,|\, i < \kappa \rangle$ be a strictly increasing continuous sequence of ordinals with limit κ such that $\delta_i \notin T$ holds for all $i < \kappa$. If for some $i < \kappa$ the restriction G_i of G to the half-open interval $[\delta_i, \delta_{i+1})$ has colouring number $> \mu$, then the first alternative holds. Otherwise we may fix for each $i < \kappa$ a well-ordering $<_i$ of $V(G_i)$ that exemplifies $\operatorname{col}(G_i) \leq \mu$. The concatenation $<^*$ of all these well-orderings has length κ, so it suffices to verify that it demonstrates $\operatorname{col}(G) \leq \mu$.

To this end, we consider any vertex x of G. Let $i < \kappa$ be the ordinal with $x \in G_i$. The neighbours of x preceding it in the sense of $<^*$ are either in δ_i or they belong to G_i and precede x under $<_i$. Since $x \geq \delta_i$ and $\delta_i \notin T$, there are less than μ neighbours of x in δ_i. Also, by our choice of $<_i$, there are less than μ such neighbours in G_i.

Second Case: T is stationary in κ.

Let us fix for each $\alpha \in T$ an ordinal $\beta_\alpha \geq \alpha$ with $|N(\beta_\alpha) \cap \alpha| \geq \mu$. A standard argument shows that the set

$$E = \{\delta < \kappa \,|\, \text{If } \alpha \in T \cap \delta, \text{ then } \beta_\alpha < \delta\}$$

is club in κ. Thus $T \cap E$ is unbounded in κ. Let the sequence $\langle \eta_i \,|\, i < \kappa \rangle$ enumerate the members of this set in increasing order. Then for each $i < \kappa$ the ordinal $\xi_i = \beta_{\eta_i}$ is at least η_i and smaller than η_{i+1}, because the latter ordinal belongs to E. In particular, each of the equations $\eta_i = \xi_j$ and $\xi_i = \xi_j$ is only possible if $i = j$. Thus it makes sense to define

$$v_\alpha = \begin{cases} \alpha & \text{if } \alpha \neq \eta_i, \xi_i \text{ for all } i < \kappa\,, \\ \xi_i & \text{if } \alpha = \eta_i \text{ for some } i < \kappa\,, \\ \eta_i & \text{if } \alpha = \xi_i \text{ for some } i < \kappa\,. \end{cases}$$

The map π sending each $\alpha < \kappa$ to v_α is a permutation of κ. If α belongs to the stationary set $T \cap E$, then $v_\alpha = \xi_i$ for some $i < \kappa$ and therefore v_α has at least μ neighbours in η_i and all of these are of the form v_β with $\beta < \alpha$. So π gives an isomomorphism between G and a μ-ladder, and in the light of Lemma 7.3.6 we are done. $\square$

7.5 The singular case

Next we consider the case that κ is a singular cardinal. The form of our argument will be recognisable to anyone who is familiar with Shelah's singular compactness theorem (see for instance [88]). We will not, however, assume such familiarity.

We will refer to sets of size at least μ as *big* and sets of size less than μ as *small*

We will often consider $\subseteq$-increasing families $(X_i)_{i<\gamma}$ of sets for which each $N_{X_i}(v)$ is small. In such cases we would like to conclude that also $N_{\bigcup_{i<\gamma} X_i}(v)$ is small. We can do this as long as γ and μ have different cofinalities. So we fix the notation ϖ for the rest of the argument to mean the least infinite cardinal whose cofinality is not equal to $\mathrm{cf}(\mu)$. Thus ϖ is either ω or ω_1.

Definition 7.5.1. A set X of vertices of a graph G is *robust* if for any $v \in V(G) \setminus X$ the neighbourhood $N_X(v)$ is small.

Remark 7.5.2. Let $(X_i)_{i<\varpi}$ be a $\subseteq$-increasing family of robust sets. Then $\bigcup_{i<\varpi} X_i$ is also robust.

Lemma 7.5.3. *Let G be a μ-unobstructed graph and let X be an uncountable set of vertices of G. Then there is a robust set Y of vertices of G which includes X and is of the same cardinality.*

Proof. Let λ be the cardinality of X.

We build a $\subseteq$-increasing family $(X_i)_{i<\varpi}$ of sets recursively by letting $X_0 = X$, taking $X_{i+1} = X_i \cup \{v \in V(G) : N_{X_i}(v) \text{ is big}\}$ and $X_l = \bigcup_{i<l} X_i$ for l a limit ordinal. We take $Y = \bigcup_{i<\varpi} X_i$. Since by construction Y is robust and includes X, it remains to prove that $|Y| = \lambda$.

To do this, we prove by induction on i that each X_i is of size λ. The cases where i is 0 or a limit are clear, so suppose $i = j+1$. By the induction hypothesis, $|X_j| = \lambda$. If $|X_{j+1}|$ were greater than λ then the induced bipartite subgraph on (X_j, X_{j+1}) would be a barricade, which is impossible by Lemma 7.3.4. Thus $|X_{j+1}| = \lambda$, as required.

□

Remark 7.5.4. Lemma 7.5.3 also holds when X is countably infinite, but the proof is more involved and so we have omitted it (unlike in the above proof, we need that there are no type II obstructions).

Proof of Theorem 7.4.1 when κ is singular. If G is μ-obstructed then we are done, so we suppose that it isn't.

Let $(v_i)_{i<\kappa}$ be an enumeration of the set of vertices. Let $(\kappa_i)_{i<\mathrm{cf}(\kappa)}$ be a continuous cofinal sequence for κ, where $\kappa_0 > \mathrm{cf}(\kappa)$ is uncountable. We begin by building a family $(X_{i,j})_{i<\mathrm{cf}(\kappa),j<\varpi}$ of robust sets of vertices of G, with $X_{i,j}$ of size κ_i, together with a family of enumerations $((x^k_{i,j})_{k<\kappa_i})_{i<\mathrm{cf}(\kappa),j<\varpi}$ of these sets. These enumerations will be chosen arbitrarily. We choose the sets in such a way that they satisfy the following conditions:

1. $X_{i',j'} \subseteq X_{i,j}$ for $i' \leq i$ and $j' \leq j$.

2. $v_k \in X_{i,0}$ for $k < \kappa_i$
3. $x^k_{i',j} \in X_{i,j+1}$ for $k < \kappa_i$ (this ensures that for any limit ordinal l all elements of $X_{l,j}$ also appear in some $X_{i,j+1}$ with $i < l$).

We do this by nested recursion on i and j. When we come to choose $X_{i,j}$, we have already chosen all $X_{i',j'}$ with $j' < j$ or both $j' = j$ and $i' \le i$. The three conditions above specify some collection of κ_i-many vertices which must appear in $X_{i,j}$. We can extend this collection to a robust set of the same size (which we take as $X_{i,j}$) by Lemma 7.5.3.

Now for $i < \mathrm{cf}(\kappa)$ let $X_i = \bigcup_{j<\varpi} X_{i,j}$, which is robust by Remark 7.5.2. We claim that for any limit ordinal l we have $X_l = \bigcup_{i<l} X_i$. That each X_i with $i < l$ is a subset of X_l is clear by condition 1 above. On the other hand, for any $x \in X_l$ there must be some $j < \varpi$ with $x \in X_{l,j}$, say $x = x^k_{l,j}$. But then as $k < \kappa_l$ it follows from the continuity of the κ_i that there is some $i < l$ with $k < \kappa_i$. Thus by condition 3 above we have $x \in X_{i,j+1} \subseteq X_i$, so that $x \in \bigcup_{i<l} X_i$.

Each vertex must lie in some set X_i by condition 2 above, and it follows from what we have just shown that the least such i can never be a limit. That is, X_0 together with all the sets $X_{i+1} \setminus X_i$ gives a partition of the vertex set. If the induced subgraph of G on any of these sets has colouring number $> \mu$ then the first alternative of Theorem 7.4.1 holds. Otherwise all of these induced subgraphs have good well-orderings. Since each X_i is robust, the well-ordering obtained by concatenating all of these well-orderings is also good, so that the third alternative of Theorem 7.4.1 holds. □

7.6 A necessary condition

In this section we derive Theorem 7.2.4 from Theorem 7.2.2. For that we shall rely on the following.

Theorem 7.6.1 (Dushnik, Erdős, and Miller, [39]). *For each infinite cardinal λ we have $\lambda \longrightarrow (\lambda, \omega)$. This means that if the edges of a complete graph on λ vertices are coloured red and green, then there is either a red clique of size λ, or a green clique of size ω.*

By restricting ones attention to the red graph, one realises that this means that every infinite graph G either contains a clique of size $|V(G)|$ or an infinite independent set. When used in this formulation, we refer to the above as DEM.

Proof of Theorem 7.2.4. By Theorem 7.2.2 it remains to show that every graph with an obstruction of type I or II has a K_μ subgraph or an induced $K_{k,\omega}$.

First we check this for obstructions (A, B) of type I. By DEM, we may assume that the neighbourhood $N(b)$ of every $b \in B$ contains an independent set Y_b of size k. Let f be the function mapping b to Y_b. There must be a finite subset Y of A such that $|f^{-1}(Y)| = |B|$. By DEM, we may assume that $f^{-1}(Y)$ contains an infinite independent set B'. Then $G[B' \cup Y]$ is isomorphic to $K_{k,\omega}$.

Hence it remains to show that every obstruction G of type II has a K_μ subgraph or an induced $K_{k,\omega}$. For every $\alpha \in T_G$, we may assume by DEM that $N(\alpha) \cap \alpha$ contains an independent set Y_α of size k. For each i with $1 \leq i \leq k$, let $f_i\colon T \to \kappa$ be the function mapping α to the i-th smallest element of Y_α. By Fodor's Lemma, there is some stationary $T' \subseteq T_G$ at which f_1 is constant, and some stationary $T'' \subseteq T'$ at which f_2 is constant. Proceeding like this, we find some stationary $S \subseteq T_G$ at which all the f_i are constant. Let X be the set of these k constants. By DEM, we may assume that S contains an infinite independent set I. Then $G[X \cup I]$ is isomorphic to $K_{k,\omega}$. $\square$

In the following example, we show that if we replace '$K_{k,\omega}$' by '$K_{\omega,\omega}$' in Theorem 7.2.4, then it becomes false.

Example 7.6.2. Let A be the set of finite 0-1-sequences, and let B be the set of infinite 0-1-sequences. We define a bipartite graph G with vertex set $A \cup B$ by adding for each $a \in A$ and $b \in B$ the edge ab if a is an initial segment of b. Since G is bipartite, it cannot contain a K_ω. It cannot contain a $K_{\omega,\omega}$ either since any two vertices in B have only finitely many neighbours in common.

On the other hand, $\mathrm{col}(G) > \aleph_0$ since G is an $\aleph_0$-barricade.

Remark 7.6.3. The proof of Theorem 7.2.4 actually shows something slightly stronger: in order to have $\mathrm{col}(G) \leq \mu$ it is enough to forbid K_μ and a K_{k,μ^+}-subgraph where the k vertices on the left are independent. If $\mu = \omega$, then DEM implies it is enough to forbid K_μ and an *induced* K_{k,μ^+}. On the other hand if $\kappa = 2^\omega$ and $\mu = \omega_1$, it may happen that the bipartite graph contains neither a K_μ nor an induced K_{k,ω_1} by Sierpiski's theorem, which says that

$$2^\omega \not\longrightarrow (\omega_1)^2_2\,.$$

Chapter 8

On tree-decompositions of one-ended graphs

8.1 Abstract

We prove that one-ended graphs whose end is undominated and has finite vertex degree have tree-decompositions that display the end and that are invariant under the group of automorphisms.

This can be applied to prove a conjecture of Halin from 2000 and solves a recent problem of Boutin and Imrich. Furthermore, it implies for every transitive one-ended graph that its end must have infinite vertex degree.

8.2 Introduction

In [38], Dunwoody and Krön constructed tree-decompositions invariant under the group of automorphisms that are non-trivial for graphs with at least two ends. In the same paper, they applied them to obtain a combinatorial proof of generalization of Stalling's theorem of groups with at least two ends. This tree-decomposition method has multifarious applications, as demonstrated by Hamann in [53] and Hamann and Hundertmark in [54]. For graphs with only a single end, however, these tree-decompositions may be trivial. Hence such a structural understanding of this class of graphs remains elusive.

For many one-ended graphs, such as the 2-dimensional grid, such tree-decompositions cannot exist. Indeed, it is necessary for existence that the end has *finite vertex degree*; that is, there is no infinite set of pairwise vertex-disjoint rays belonging to that end. Already in 1965 Halin [46] knew that one-ended graphs whose end has finite vertex degree have tree-decompositions displaying the end (a precise definition can be found towards the end of Section 8.4). Nevertheless, for these tree-decompositions to be of any use for applications as above, one needs them to have the additional property that they are invariant under the group of automorphisms. Unfortunately such tree-decomposition do not exist for all graphs in question, see Example 8.4.10 below, but in the example there is a vertex dominating the end. In this chapter we construct such tree-decompositions if the end is not dominated.

Theorem 8.2.1. *Every one-ended graph whose end is undominated and has finite vertex degree has a tree-decomposition that displays its end and that is invariant under the group of automorphisms.*

This better structural understanding leads to applications similar to those for graphs with more than one end. Indeed, below we deduce from Theorem 8.2.1 a conjecture of Halin from 2000, and answer a recent question of Boutin and Imrich. A further application was pointed out by Hamann.

For graphs like the one in Figure 8.1, the tree-decompositions of Theorem 8.2.1 can be constructed using the methods of Dunwoody and Krön. Namely, if the graphs in question contain 'highly connected tangles' aside from the end. In general such tangles need not exist, for an example see Figure 8.2. It is the essence of Theorem 8.2.1 to provide a construction that is invariant under the group of automorphisms that decomposes graphs as those in Firgure 8.2 in a tree-like way.

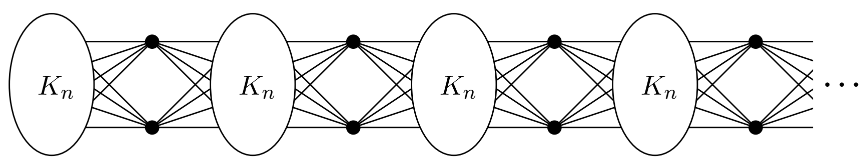

Figure 8.1: Complete graphs glued together at separators of size two along a ray. The method of Dunwoody and Krön gives a tree-decomposition of this graph along an end whose separators have size two.

Applications. In [52] Halin showed that one-ended graphs with vertex degree equal to one cannot have countably infinite automorphism group. Not completely satisfied with his result, he conjectured that this extends to one-ended graphs with finite vertex degree. Theorem 8.2.1 implies this conjecture.

Theorem 8.2.2. *Given a graph with one end which has finite vertex degree, its automorphism group is either finite or has at least* $2^{\aleph_0}$ *many elements.*

Theorem 8.2.2 can be further applied to answer a question posed by Boutin and Imrich, who asked in [13] whether there is a graph with linear growth and countably infinite automorphism group. Theorem 8.2.2 implies a negative answer to this question as well as strengthenings of further results of Boutin and Imrich, see Section 8.5 for details.

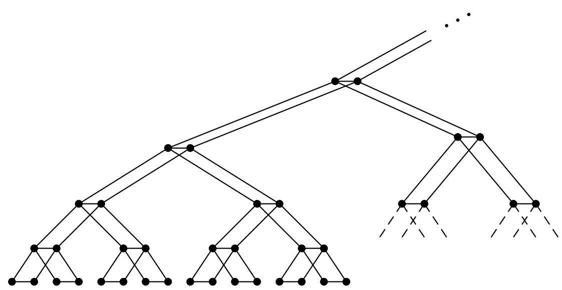

Figure 8.2: The product of the canopy tree with K_1. This graph has a tree decomposition whose decomposition tree is the canopy tree.

Finally, Matthias Hamann[1] pointed out the following consequence of Theorem 8.2.1.

Theorem 8.2.3. *Ends of transitive one-ended graphs must have infinite vertex degree.*

We actually prove a stronger version of Theorem 8.2.3 with 'quasi-transitive'[2] in place of 'transitive'.

The rest of this chapter is structured as follows: in Section 8.3 we set up all necessary notations and definitions. As explained in [31], there is a close relation between tree-decompositions and nested sets of separations. In this chapter we work mainly with nested sets of separations. In Section 8.4 we prove Theorem 8.2.1, and Section 8.5 is devoted to the proof of Theorem 8.2.2, and its implications on the work of Boutin and Imrich. Finally, in Section 8.6 we prove Theorem 8.2.3.

Many of the lemmas we apply in this work were first proved by Halin. Since in some cases we need slight variants of the original results and also since Halin's original papers might not be easily accessible, proofs of some of these results are included in appendices.

8.3 Preliminarlies

Throughout this chapter $V(G)$ and $E(G)$ denote the sets of vertices and edges of a graph G, respectively. We refer to [36] for all graph theoretic notions which are not explicitly defined.

[1]personal communication

[2]Here a graph is *quasi-transitive*, if there are only finitely many orbits of vertices under the automorphism group.

8.3.1 Separations, rays and ends

A *separator* in a graph G is a subset $S \subseteq V(G)$ such that $G-S$ is not connected. We say that a separator S *separates vertices u and v* if u and v are in different components of $G-S$. Given two vertices u and v, a separator S separates u and v *minimally* if it separates u and v and the components of $G-S$ containing u and v both have the whole of S in their neighbourhood. The following lemma can be found in Halin's 1965 paper [49, Statement 2.4], and also in his later paper [50, Corollary 1] and then with a different proof.

Lemma 8.3.1. *Given vertices u and v and $k \in \mathbb{N}$, there are only finitely many distinct separators of size at most k separating u and v minimally.*

A *separation* is a pair (A, B) of subsets of $V(G)$ such that $A \cup B = V(G)$ and there is no edge connecting $A \setminus B$ to $B \setminus A$. This immediately implies that if u and v are adjacent vertices in G then u and v are both contained in either A or B. The sets A and B are called the *sides* of the separation (A, B). A separation (A, B) is said to be *proper* if both $A \setminus B$ to $B \setminus A$ are non-empty and then $A \cap B$ is a separator. A separation (A, B) is *tight* if every vertex in $A \cap B$ has neighbours in both $A \setminus B$ and $B \setminus A$. The *order* of a separation is the number of vertices in $A \cap B$. Throughout this chapter we will only consider separations of finite order. The following is well-known.

Lemma 8.3.2. (See [30, Lemma 2.1]) *Given any two separations (A, B) and (C, D) of G then the sum of the orders of the separations $(A \cap C, C \cup D)$ and $(B \cap D, A \cup C)$ is equal to the sum of the orders of the separations (A, B) and (C, D). In particular if the orders of (A, B) and (C, D) are both equal to k the the sum of the orders of $(A \cap C, C \cup D)$ and $(B \cap D, A \cup C)$ is equal to $2k$.* □

The separations (A, B) and (C, D) are *strongly nested* if $A \subseteq C$ and $D \subseteq B$. They are *nested* if they are strongly nested after possibly exchanging '(A, B)' by '(B, A)' or '(C, D)' by '(D, C)'. That is, (A, B) and (C, D) are nested if one of the following holds:

- $A \subseteq C$ and $D \subseteq B$,
- $A \subseteq D$ and $C \subseteq B$,
- $B \subseteq C$ and $D \subseteq A$,
- $B \subseteq D$ and $C \subseteq A$.

We say a set $\mathcal{S}$ of separations is *nested*, if any two separations in it are nested.

A *ray* in a graph G is a one-sided infinite path $v_0, v_1, \dots$ in G. The sub-rays of a ray are called its *tails*. Given a finite separator S of G, there is for every ray γ a unique component of $G-S$ that contains all but finitely many vertices of γ. We say that γ *lies* in that component of $G-S$. Given a separation (A, B) of finite order one can similarly say that γ lies in one of the sides of the separation. Two rays are *in the same end* if they lie in the same component of

$G-S$ for every finite separator of G. Clearly, this is an equivalence relation. An equivalence class is called a *(vertex) end*[3]. An alternative way to define ends is to say that two rays R_1 and R_2 are in the same end if there are infinitely many pairwise disjoint $R_1 - R_2$ paths. (Given subsets X and Y of the vertex set, an $X - Y$ path is a path that has its initial vertex in X and terminal vertex in Y and every other vertex is neither in X nor Y. In the case where $X = \{x\}$ then we speak of $x - Y$ paths instead of $X - Y$ paths and if $Y = \{y\}$ we speak of $x-y$ paths.) An end ω *lies* in a component C of $G-S$ if every ray that belongs to ω lies in C. Clearly, every end lies in a unique component of $G-S$ for every finite separator S and if (A,B) is a separation of finite order then an end either lies in A or B.

A vertex $v \in V(G)$ *dominates* an end ω of G, if there is no separation (A,B) of finite order such that $v \in A \setminus B$ and ω lies in B. Equivalently, v dominates ω if for every ray R in ω there are infinitely many paths connecting v to R such that any two of them only intersect in v.

The *vertex degree* of an end ω is equal to a natural number k if the maximal cardinality of a family of pairwise disjoint rays belonging to the end is k. If no such number k exists then we say that the vertex-degree of the end is infinite. Halin [46] (see also [35, Theorem 8.2.5]) proved that if the vertex-degree of an end is infinite then there is an infinite family of pairwise disjoint rays belonging to the end. Ends with finite vertex degree are sometimes called *thin* and those with infinite vertex degree are called *thick*.

The following lemma is well-known. A proof can be found in Appendix A.

Lemma 8.3.3. (Cf. [52, Section 3]) *Let G be a connected graph and ω an end of G having a finite vertex degree. Then there are only finitely many vertices in G that dominate the end ω.*

In this chapter we are focusing on 1-ended graphs where the end ω has vertex degree k. In the following definition we pick out a class of separations that are relevant in this case.

Definition 8.3.4. Let G be an arbitrary graph. If ω is an end of G that has vertex degree k then say that a separation (A,B) is *ω-relevant* if it has the following properties

- the order of (A,B) is exactly k,
- $A \setminus B$ is connected,
- every vertex in $A \cap B$ has a neighbour in $A \setminus B$,
- ω lives in B, and
- there is no separation (C,D) of order $< k$ such that $A \subseteq C$ and ω lives in D.

[3] A notion related to 'vertex ends' are 'topological ends'. In this chapter we are mostly interested in graphs where no vertex dominates a vertex end. In this context the two notions of end agree.

Define $\mathcal{S}_\omega$ as the set of all ω-relevant separations.

The following characterization of ω-relevant separations is uses a Menger type result. A proof based on [45] and [46] is contained in Appendix A.

Lemma 8.3.5. *Let G be an arbitrary graph. Suppose ω is an end of G with vertex degree k.*

1. *If (A, B) is an ω-relevant separation then there is a family of k pairwise disjoint rays in ω such that each of them has its initial vertex in $A \cap B$.*
2. *Conversely, if (A, B) is a separation of order k such that $A \setminus B$ is connected, every vertex in $A \cap B$ has a neighbour in $A \setminus B$, the end ω lies in B and there is a family of k disjoint rays in ω such that each of these rays has its initial vertex in $A \cap B$ then the separation (A, B) is ω-relevant.*

In particular, for $(A, B) \in \mathcal{S}_\omega$ the component of $G - (A \cap B)$ in which ω lives has the whole of $A \cap B$ in its neighbourhood and hence every separation in $\mathcal{S}_\omega$ is tight. Note that the set $A \setminus B$ completely determines the ω-relevant separation (A, B).

The relation

$$(A, B) \leq (C, D) :\iff A \subseteq C \text{ and } B \supseteq D$$

defines a partial order on the set of all separations, so in particular on the set $\mathcal{S}_\omega$. Since (C, D) is a tight separation, the condition $A \subseteq C$ implies that $D \subseteq B$. This is shown in [31, (7) on p. 17] and the argument goes as follows: Suppose that $D \not\subseteq B$ and $x \in D \setminus B$. Then $x \in A \subseteq C$ so $x \in (C \cap D) \setminus B$. Because (C, D) is a tight separation, x has a neighbour $y \in D \setminus C$. But $x \in A \setminus B$ and hence y must also be in A. But $y \notin C$, contradicting the assumption that $A \subseteq C$. Hence $D \subseteq B$ and $(A, B) \leq (C, D) \iff A \subseteq C$.

The next result follow from results of Halin in [46]. These results are in turn proved by using Menger's Theorem. For the convenience of the reader a detailed proof is provided in Appendix A.

Theorem 8.3.6. *Let G be a connected 1-ended graph such that the end ω is undominated and has finite vertex degree k. Then there is a sequence $\{(A_n, B_n)\}_{n \geq 0}$ of ω-relevant separations, such that the sequence of sets B_n is strictly decreasing and for every finite set of vertices F there is a number n such that $F \subseteq A_n \setminus B_n$.*

We will not use the following in our proof.

Remark 8.3.7. Theorem 8.3.6 is also true if we leave out the assumption that G is one-ended (and replace 'the end ω' by 'there exists an end ω that').

8.3.2 Automorphism groups

An *automorphism* of a graph $G = (V, E)$ is a bijective function $\gamma : V \to V$ that preserves adjacency and whose inverse also preserves adjacency. Clearly an automorphism γ also induces a bijection $E \to E$ which by abuse of notation we will also call γ. The *automorphism group of* G, i.e. the group of all automorphisms of G, will be denoted by $\mathrm{Aut}(G)$.

Let Γ be a subgroup of $\mathrm{Aut}(G)$. For a set $D \subseteq V(G)$ we define the *setwise stabiliser* of D as the subgroup $\Gamma_{\{D\}} = \{\gamma \in \Gamma \mid \gamma(D) = D\}$ and the *pointwise stabiliser* of D is defined as $\Gamma_{(D)} = \{\gamma \in \Gamma \mid \gamma(d) = d \text{ for all } d \in D\}$. The setwise stabiliser is the subgroup of all elements in Γ that leave the set D invariant and the pointwise stabiliser is the subgroup of all those elements in Γ that fix every vertex in D. If $D \subseteq V(G)$ is invariant under Γ then we use Γ^D to denote the permutation group on D induced by Γ, i.e. Γ^D is the group of all permutation σ of D such that there is some element $\gamma \in \Gamma$ such that the restriction of γ to D is equal to σ. Note that $\Gamma_{(D)}$ is a normal subgroup of $\Gamma_{\{D\}}$ and the index $\Gamma_{(D)}$ in $\Gamma_{\{D\}}$ is equal to the number of elements in $(\Gamma_{\{D\}})^D$.

The full automorphism group of a graph has a special property relating to separations. Suppose γ is an automorphism of a graph G and that γ leaves both sides of a separation (A, B) invariant and fixes every vertex in the separator $A \cap B$. Then the full automorphism group contains automorphisms σ_A and σ_B such that σ_A like γ on A fixes every vertex in B and *vice versa* for σ_B. Informally one can describe this property by saying that the pointwise stabiliser (in the full automorphism group) of a set D of vertices acts indpendently on the components of $G - D$. We will refer to this property as *the independence property.*

There is a natural topology on $\mathrm{Aut}(G)$, called the *permutation topology*: endow the vertex set with the discrete topology and consider the topology of pointwise convergence on $\mathrm{Aut}(G)$. Clearly, the permutation topology also makes sense for any group of permutations of a set. The following lemma is a special case of a result in [24, (2.6) on p. 28]. In particular it tells us that the limit of a sequence of automorphisms again is an automorphism. This fact will be central to the proof of Theorem 8.2.2.

Lemma 8.3.8. *The automorphism group of a graph is closed in the set of all permutations of the vertex set endowed with the topology of pointwise convergence.*

The next result is also a special case of a result from Cameron's book refered to above. This time we look at [24, (2.2) on p. 28].

Lemma 8.3.9. *The automorphism group of a countable graph is finite, countably infinite or has at least $2^{\aleph_0}$ elements.*

8.4 Invariant nested sets

In this section we will prove Theorem 8.4.8. The following two facts about sequences of nested separations will be useful at several points in the proof.

Lemma 8.4.1. *Let G be a connected graph. Assume that $(A_i, B_i)_{i\in\mathbb{N}}$ is a sequence of proper separations of order at most some fixed natural number k.*

Assume also that $A_i \subsetneq A_{i-1}$, every $A_i \setminus B_i$ is connected, and every vertex in $A_i \cap B_i$ has a neighbour in $A_i \setminus B_i$. Define X as the set of vertices contained in infinitely many A_i. Then

1. *$X \subseteq B_i$ for all but finitely many i,*

2. *there is a unique end μ which lies in every A_i, and*

3. *$x \in X$ if and only if x dominates μ.*

Proof. First observe that $X = \bigcap_{i \in \mathbb{N}} A_i$ because the sequence A_i is decreasing. Let X' be the set of vertices in X with a neighbour outside of X. For every $x \in X'$ we can find a neighbour y of x and $i_0 \in \mathbb{N}$ such that $y \notin A_i$ for every $i \geq i_0$. Since the edge xy must be contained in either A_i or B_i we conclude that $x \in B_i$ and thus $x \in A_i \cap B_i$ for $i \geq i_0$.

Hence there is $i_1 \in \mathbb{N}$ such that $X' \subseteq A_i \cap B_i$ for every $i \geq i_1$. The order of each separation is at most k, so X' contains at most k vertices. Now for $i \geq i_1$ every path from $X \setminus B_i$ to $A_i \setminus (X \cup B_i)$ must pass through X' and thus through B_i. Since $A_i \setminus B_i$ is connected this means that one of the two sets must be empty, i.e., either $X \setminus B_i = \emptyset$ or $X \setminus B_i = A_i \setminus B_i$. Assume that the latter is the case. Then A_i contains at most k vertices which are not contained in X and the same is clearly true for every A_j for $j > i$. This contradicts the fact that the sequence A_i was assumed to be infinite and strictly decreasing. We conclude that $X \subseteq B_i$ for $i \geq i_1$. Note that this implies that $X = X'$ because if $i \geq i_1$ then $X \subseteq A_i \cap B_i$ and every vertex in $A_i \cap B_i$ has an neighbour in $A_i \setminus B_i$.

To see that there is an end μ which lies in every A_i we construct a ray which has a tail in each A_i. For this purpose pick for $i \geq i_1$ a vertex $v_i \in A_i \setminus X$ and paths P_i connecting v_i to v_{i+1} in $A_i \setminus X$. This is possible because $A_i \setminus X$ contains $A_i \setminus B_i$ and is connected ($A_i \setminus B_i$ is connected and every vertex in $B_i \cap A_i$ has a neighbour in $A_i \setminus B_i$). No vertex lies on infinitely many paths P_i because no vertex is contained in infinitely many sets $A_i \setminus X$. Hence the union of the paths P_i is an infinite, locally finite graph and thus contains a ray. This ray belongs to an end μ which lies in every A_i.

Finally we need to show that every vertex in X dominates the end μ. Without loss of generality we can assume that $X \subseteq B_i$ for all i. So, let R be a ray in μ and $x \in X$. We will inductively construct infinitely many paths from x to R which only intersect in x. Assume that we already constructed some finite number of such paths. Since all of them have finite length, there is an index i such that $A_i \setminus B_i$ doesn't contain any vertex in their union. The ray R has a tail contained in $A_i \setminus B_i$ and since $x \in A_i \cap B_i$ we know that x has a neighbour in $A_i \setminus B_i$. Finally $A_i \setminus B_i$ is connected, so we can find a path connecting x to the tail of R which intersects the previously constructed paths only in x. Proceeding inductively we obtain infinitely many paths connecting x to R which pairwise only intersect in x completing the proof of the Lemma. $\square$

We would now like to construct a subset of the set $\mathcal{S}_\omega$ of ω-relevant separations that is both nested and invariant under all automorphisms and from that set we construct a tree. The following two lemmas give us important properties of nestedness when we restrict to ω-relevant separations.

Lemma 8.4.2. *Two separations $(A,B),(C,D)$ in $\mathcal{S}_\omega$ are nested if and only if they are either comparable with respect to $\leq$, or $A \subseteq D$.*

Proof. First assume that the two separations are nested. It is impossible that $B \subseteq C$ and $D \subseteq A$ since the end ω lies in B and D, but not in C and A. Hence, if the two separations are not comparable, then we know that $A \subseteq D$ and $C \subseteq B$.

For the converse implication first consider the case that $A \subseteq D$. We want to show that $C \subseteq B$. Assume for a contradiction that there is a vertex x in $C \setminus B$. This vertex must be contained in $A \subseteq D$ and hence in the separator $C \cap D$. By the definition of $\mathcal{S}_\omega$ the vertex x must have a neighbour y in $C \setminus D$. Then $y \notin A$ and $x \notin B$, contradicting the fact that the edge xy must lie in either A or B, as (A,B) is a separation.

Finally, note that any two separations in $\mathcal{S}_\omega$ that are comparable with respect to $\leq$ are obviously nested. □

Lemma 8.4.3. (Analogies with [38, Lemma 4.2]) For each $(A,B) \in \mathcal{S}_\omega$ there are only finitely many $(C,D) \in \mathcal{S}_\omega$ not nested with (A,B).

Proof. The first step is to show that if (C,D) is not nested with (A,B) then (C,D) separates some vertices v and w in $A \cap B$. Then we show that we may assume that the separation is minimal. Since $A \cap B$ is finite there are only finitely many possibilities for the pair v,w and we can apply Lemma 8.3.1 to deduce the result.

First suppose for a contradiction that $(C \setminus D) \cap (A \cap B)$ is empty. Since $C \setminus D$ is connected, it must be a subset of $A \setminus B$ or $B \setminus A$. As every vertex in $C \cap D$ has a neighbour in $C \setminus D$ it follows that $C \subseteq A$ in the first case, whilst $C \subseteq B$ in the second. In both cases (A,B) and (C,D) are nested by Lemma 8.4.2, contrary to our assumption. Hence there exists a vertex $v \in (C \setminus D) \cap (A \cap B)$. Note that by letting the separations (A,B) and (C,D) switch roles we see that $(A \setminus B) \cap (C \cap D)$ is also non-empty.

Since the separation (C,D) is in $\mathcal{S}_\omega$ there is by Lemma 8.3.5 a family of k disjoint rays that all have their initial vertices in $C \cap D$. Because ω lives in D, all vertices in these rays, except their initial vertices, are contained in the component of $D \setminus C$ that contains ω. Pick a vertex v' from $(A \setminus B) \cap (C \cap D)$. This vertex v' is the initial vertex of one of the rays mentioned above. Since ω lives in B this rays must contain a vertex w from $A \cap B$ and as mentioned above w is contained in the component of $D \setminus C$ that contains ω. Now we have shown that (C,D) separates the two vertices v and w. This separation is minimal because v is in $C \setminus D$ and $C \setminus D$ is connected and has $C \cap D$ as it neighbourhood, and w is contained in the component of $G - (C \cap D)$ that contains ω and that component has the whole of $C \cap D$ as its neighbourhood. □

Let G be a one-ended graph whose end ω is undominated and has finite vertex degree k. Recall that by Lemma 8.4.1 there are no infinite decreasing chains in

$\mathcal{S}_\omega$—such a chain would define an end $\mu \neq \omega$, contradicting the assumption that G has only one end. In particular, $\mathcal{S}_\omega$ has minimal elements. Assign recursively an ordinal $\alpha(A, B)$ to each $(A, B) \in \mathcal{S}_\omega$ by the following method: if (A, B) is minimal (with respect to $\leq$ in $\mathcal{S}_\omega$) then set $\alpha(A, B) = 0$; otherwise define $\alpha(A, B)$ as the smallest ordinal β such that $\alpha(C, D) < \beta$ for all separations $(C, D) \in \mathcal{S}_\omega$ such that $(C, D) < (A, B)$. For $v \in V(G)$, let $\mathcal{S}_\omega(v)$ be the set of those separations (A, B) in $\mathcal{S}_\omega$ with $v \in A \cap B$. Now set

$$\alpha(v) = \sup\{\alpha(A, B) \mid (A, B) \in \mathcal{S}_\omega(v)\}.$$

If it so happens that $\mathcal{S}_\omega(v)$ is empty then $\alpha(v) = 0$. For a vertex set S, we let $\alpha(S)$ be the supremum over all $\alpha(v)$ with $v \in S$. Note that the functions $\alpha(A, B)$ and $\alpha(v)$ are both invariant under the action of the automorphism group of G.

Example 8.4.4. Below is a construction of a graph where α takes ordinal values that are not natural numbers. However, it is not difficult to show that for a locally finite connected graph the α-values are always natural numbers.

We construct a graph G at which α takes values that are not natural numbers. Let $P_n = v_0^n, \ldots, v_n^n$ be a path of length n. We obtain G by taking a ray and identifying its starting vertex r with the vertices v_n^n for each $n \geq 0$. This graph has only one end μ and its vertex degree is 1. For $0 \leq k \leq n-1$ the separation $(\{v_0^n, \ldots, v_k^n\}, VG \setminus \{v_0^n, \ldots, v_{k-1}^n\})$ is μ-relevant and its α-value is k. Hence any separation (A, B) with r (and all the attached paths) in A has α-value at least the ordinal ω.

Lemma 8.4.5. *Let G be a graph with only one end ω. Assume that ω is undominated and has vertex degree k. Let (C, D) be in $\mathcal{S}_\omega$. Then for all but finitely many vertices v in C, we have $\alpha(v) \leq \alpha(C, D)$.*

Proof. By Lemma 8.4.3, there are only finitely many separations in $\mathcal{S}_\omega$ that are not nested with (C, D). Let C' the set of those vertices in $C \setminus D$ that are not in any separator of these finitely many separations. It suffices to show that if $v \in C'$ and (A, B) in $\mathcal{S}_\omega(v)$ then $\alpha(A, B) < \alpha(C, D)$. Note that the result is trivially true if $\mathcal{S}_\omega(v)$ is empty. By the choice of v, the separations (A, B) and (C, D) are nested. Since v is in $(C \setminus D) \cap (A \cap B)$, it is not true that $A \subseteq D$ or $B \subseteq D$. Since the end ω does not lie in the sides A and C, it does not lie in the side $A \cup C$ of the separation $(A \cup C, B \cap D)$. Hence it lies in the side $B \cap D$. In particular $B \cap (D \setminus C)$ is nonempty. Thus it is not true that $B \subseteq C$. Looking at the definition of nestedness we see that $A \subseteq C$. Hence $(A, B) < (C, D)$ and thus $\alpha(A, B) < \alpha(C, D)$ and the result follows. $\square$

Lemma 8.4.6. *Let G be a graph with only one end ω. Assume that ω is undominated and has vertex degree k. For every separation (C, D) in $\mathcal{S}_\omega$, there is a separation $(A, B) \in \mathcal{S}_\omega$ such that $C \subseteq A$ and $\alpha(C) < \alpha(A, B)$.*

Proof. Let $\{(A_n, B_n)\}_{n \geq 0}$ be a sequence of ω-relevant separations as described in Theorem 8.3.6. Find a separation (A, B) in this sequence such that $C \cap D \subseteq A \setminus B$. Suppose for a contradiction that $C \setminus D$ contains a vertex x from $A \cap B$.

There is a ray R that has x as a starting vertex and every other vertex is contained in $B \setminus A$. Because $C \cap D$ contains no vertex from B we see that this ray would be contained in $C \setminus D$, contradicting the assumption that the end ω lies in D. Hence, $C \setminus D$ does not intersect $A \cap B$ and then, since $C \setminus D$ is connected, we conclude that $C \subseteq A$. Thus $\alpha(C, D) \leq \alpha(A, B)$.

By the previous Lemma there are at most finitely many vertices v in C such that $\alpha(v) > \alpha(C, D)$. Suppose for a contradiction that v is such a vertex and there is no value of n such that $\alpha(v) < \alpha(A_n, B_n)$. Then we can find a sequence $\{(C_n, D_n)\}_{n \geq 0}$ of separations in $\mathcal{S}_\omega(v)$ such that $\alpha(C_1, D_1) < \alpha(C_2, D_2) < \cdots$ and for every n there is a number r_n such $\alpha(A_n, B_n) < \alpha(C_{n_r}, B_{n_r})$. By Lemma 8.4.3 we may assume that for all values of n and m the separations (C_n, D_n) and (C_m, B_m) are nested. Say that a pair of separations $\{(C_n, D_n), (C_m, D_m)\}$ is blue if the separations are comparable with respect to $\leq$ and red otherwise. By Ramsey's Theorem, see e.g. [24, (1.9) on p. 16], there is an infinite set of separations such that all pairs from that set have the same colour. If all pairs from that set were blue then we could find an infinite increasing or a decreasing chain. By Lemma 8.4.1(2) there cannot be an infinite descending chain of separations and if there was an infinite increasing chain in $\mathcal{S}_\omega(v)$ then, by Lemma 8.4.1(3) with the roles of the A_i's and the B_i's reversed, v would be a dominating vertex for the end ω, contrary to assumptions. Hence all pairs from that infinite set must be red and we can conclude that there is an infinite set of separations in the family $\{(C_n, D_n)\}_{n \geq 0}$ such that no two of them are comparable with respect to ordering. We may assume that if n and m are distinct then (C_n, D_n) and (C_m, D_m) are not comparable and then $C_n \setminus D_n$ and $C_m \setminus D_m$ are disjoint. Start by choosing n such that $v \in A_n \setminus B_n$ and then choose m such that none of the vertices in $A_n \cap B_n$ is in $C_m \setminus D_m$. There must be some vertex u that belongs both to B_n and $C_m \setminus D_m$. The set $(C_m \setminus D_m) \cup \{v\}$ is connected and thus it contains a $v - u$ path P. But $v \in A_n \setminus B_n$ and $u \in B_n \setminus A_n$ and the path P contains no vertices from $A_n \cap B_n$. We have reached a contradiction. Hence our original assumption must be wrong. $\square$

Let X be a connected set of vertices which cannot be separated from the end ω by a separation of order less than k. A separation $(A, B) \in \mathcal{S}_\omega$ is called *X-nice*, if for every $v \in A \cap B$ we have $\alpha(v) > \alpha(X)$ and there is some $\varphi \in \mathrm{Aut}(G)$ such that $\varphi(X) \subseteq A$ (then we must have $\varphi(X) \subseteq A \setminus B$). Let $\mathcal{N}(X)$ be the set of all X-nice separations in $\mathcal{S}_\omega$ which are minimal with respect to $\leq$, i.e. $\mathcal{N}(X)$ contains all X-nice separations $(A, B) \in \mathcal{S}_\omega$ such that A is minimal with respect to inclusion.

Lemma 8.4.7. *Let G be a graph with only one end ω. Assume that ω is undominated and has vertex degree k.*

Suppose $(X, Y) \in \mathcal{S}_\omega$. Then $\mathcal{N}(X)$ is non-empty. For each automorphism φ of G there is a unique element (A, B) in $\mathcal{N}(X)$ such that $\varphi(X) \subseteq A$. If (A, B) and (C, D) are not equal and in $\mathcal{N}(X)$, then $A \subseteq D$ and $C \subseteq B$. Furthermore, any two elements of $\mathcal{N}(X)$ can be mapped onto each other by an automorphism.

Proof. The existence of an X-nice separation follows from Lemma 8.4.6. Minimal such separations exist because by Lemma 8.4.1 an infinite descending chain would imply that G had another end $\mu \neq \omega$.

Let (A, B) and (C, D) be elements of $\mathcal{N}(X)$. Suppose $\varphi(X) \subseteq A$ and $\psi(X) \subseteq C$, where $\varphi, \psi \in \mathrm{Aut}(G)$. Note that $\varphi(X)$ is disjoint from $C \cap D$ because $\alpha(\varphi(X)) = \alpha(X)$, which is strictly less than $\alpha(v)$ for any $v \in C \cap D$. Hence it is a subset of either $C \setminus D$ or $D \setminus C$. We next prove that if (A, B) and (C, D) are not equal, then $A \subseteq D$ and $C \subseteq B$.

First we consider the case that $\varphi(X)$ is a subset of $C \setminus D$. Our aim is to show that (A, B) and (C, D) are equal. This also implies that (A, B) is the unique element in $\mathcal{N}(X)$ such that $\varphi(X) \subseteq A$. Our strategy will be to construct a X-nice separation that is $\leq$ to both of them and by minimality of (A, B) and (C, D) we will conclude that it must be equal to both of them. Note that $\varphi(X)$ is included in $(C \setminus D) \cap (A \setminus B)$. Let A' be the connected component of $(C \setminus D) \cap (A \setminus B)$ that contains the connected set $\varphi(X)$ together with the separator of $(A \cap C, B \cup D)$. Let B' be the union of $B \cup D$ with the other components of $(C \setminus D) \cap (A \setminus B)$.

Next we show that the separation (A', B') is in $\mathcal{N}(X)$. Since the end ω lies in $B \cap D$, this vertex set is infinite. Because (A, B) is in $\mathcal{S}_\omega$, the separation $(A \cup C, B \cap D)$ has order at least k. Hence by Lemma 8.3.2, the separation $(A \cap C, B \cup D)$ has order at most k. The property that X cannot be separated from ω by fewer than k vertices implies that the separation (A', B') has order precisely k. Also, every vertex of the separator of (A', B') has a neighbour in $A' \setminus B'$ and in $B' \setminus A'$. Clearly ω lies in B' and there is no separation (C', D') of order less than k such that $A' \subseteq C'$ and ω lies in D' as $(X, Y) \in \mathcal{S}_\omega$. Hence (A', B') is in $\mathcal{S}_\omega$ and thus it is in $\mathcal{N}(X)$ as $A' \subseteq A$. Since $A' \subseteq A$, it must be that $A' = A$ by the minimality of (A, B). Similarly, $A' = C$. Thus $A = C$ and so $(A, B) = (C, D)$. This completes the case when $\varphi(X)$ is a subset of $C \setminus D$.

So we may assume that $\varphi(X) \subseteq D \setminus C$, and by symmetry that $\psi(X) \subseteq B \setminus A$. Consider the separations $(A \cap D, B \cup C)$ and $(B \cap C, A \cup D)$. They must have order at least k because $\varphi(X) \subseteq A \cap D$, $\omega \in B \cup C$ and $\psi(X) \subseteq B \cap C$, $\omega \in A \cup D$. So they must have order precisely k by Lemma 8.3.2. Let A' be the component of $G - (B \cup C)$ that contains $\varphi(X)$ together with the separator of $(A \cap D, B \cup C)$. Let B' be the union of $B \cup C$ with the other components. Similar as in the last case we show that (A', B') is in $\mathcal{N}(X)$. By the minimality of (A, B) it must be that $A \subseteq D$. The above argument with the separation $(B \cap C, A \cup D)$ in place of $(A \cap D, B \cup C)$ yields that $C \subseteq B$. This completes the proof that if (A, B) and (C, D) are not equal and in $\mathcal{N}(X)$, then (A, B) and (C, D) are nested.

By the above there is for each $\varphi \in \mathrm{Aut}(G)$ a unique separation $(A_\varphi, B_\varphi) \in \mathcal{N}(X)$ such that $\varphi(X) \subseteq A_\varphi$. If we apply φ^{-1} to this separation we must obtain the unique separation $(A, B) \in \mathcal{N}(X)$ such that $X \subseteq A$. Hence any separation of $\mathcal{N}(X)$ can be mapped by an automorphism to every other separation in $\mathcal{N}(X)$. $\square$

Theorem 8.4.8. *Let G be a connected graph with only one end ω, which is undominated and has finite vertex degree k. Then there is a nested set $\mathcal{S}$ of ω-relevant separations of G that is* $\mathrm{Aut}(G)$*-invariant. And there is a 1-ended* *tree T and a bijection between the edge set of T and $\mathcal{S}$ such that the natural action of* $\mathrm{Aut}(G)$ *on S induces an action on T by automorphisms.*

Proof. Pick some ω-relevant separation (A_0, B_0). Define a sequence (A_n, B_n) of separations as follows. For $n \in \mathbb{N}_{>0}$ pick $(A_n, B_n) \in \mathcal{N}(A_{n-1})$ such that $A_{n-1} \subsetneq A_n$, which is possible by Lemma 8.4.7. Observe that the sequence of separations (A_n, B_n) has the same properties as the sequence in Theorem 8.3.6.

Now let

$$\mathcal{S} = \{(\varphi(A_n), \varphi(B_n)) \mid n \in \mathbb{N}_{>0}, \varphi \in \mathrm{Aut}(G)\}.$$

Note that (A_0, B_0) is not an element in $\mathcal{S}$.

First we prove that $\mathcal{S}$ is nested. Let $(\varphi(A_n), \varphi(B_n))$ and $(\psi(A_m), \psi(B_m))$ be two different elements of $\mathcal{S}$ (here φ and ψ are automorphisms of G). If $m = n$ then they are nested by Lemma 8.4.7, since they both are elements of $\mathcal{N}(A_{n-1})$. Hence assume without loss of generality that $n < m$. If $\varphi(A_m) = \psi(A_m)$ then $\varphi(A_n) \subseteq \varphi(A_m) = \psi(A_m)$ which implies that the two separations are nested. Otherwise by Lemma 8.4.7 we have $\varphi(A_n) \subseteq \varphi(A_m) \subseteq \psi(B_m)$, also showing nestedness, by Lemma 8.4.2.

Next we construct a directed graph T_+. We define T_+ as follows. Its vertex set is $\mathcal{S}$. We add a directed edge from $(\varphi(A_n), \varphi(B_n))$ to $(\psi(A_{n+1}), \psi(B_{n+1}))$ if $\varphi(A_n)$ is a subset of $\psi(A_{n+1})$. By Lemma 8.4.7, each vertex has outdegree at most one. And by the construction of $\mathcal{S}$ it has outdegree at least one.

The next step is to show that the graph is connected. Let $(C, D) = \varphi(A_n, B_n)$ be a vertex in T_+. Find an m such that $C \subseteq A_m \setminus B_m$. Suppose for a contradiction that $(\varphi(A_m), \varphi(B_m)) \neq (A_m, B_m)$. Both $(\varphi(A_m), \varphi(B_m))$ and (A_m, B_m) are in $\mathcal{N}(X)$. By Lemma 8.4.7 $\varphi(A_m) \subseteq B_m$. Thus $\varphi(A_m)$ is empty. This is a contradiction to the assumption that (A_m, B_m) is a proper separation. Now we see that

$$(A_m, B_m) = (\varphi(A_m), \varphi(B_m)), (\varphi(A_{m-1}), \varphi(B_{m-1})), \ldots, (\varphi(A_n), \varphi(B_n)) = (C, D)$$

is a path in T_+ from (A_m, B_m) to (C, D). Thus every vertex in T_+ is in the same connected component as some vertex (A_m, B_m) and since they all belong to the same component we deduce that T_+ is connected. Hence the corresponding undirected graph T is a tree.

The map that sends $(\varphi(A_n), \varphi(B_n))$ to the edge with endvertices $(\varphi(A_n), \varphi(B_n))$ and $(\psi(A_{n+1}), \psi(B_{n+1}))$ is clearly a bijection. If the ray $(A_1, B_1), (A_2, B_2), \ldots$ is removed from T then what remains of T is clearly rayless and thus the tree T is one-ended.

The statement about the action of $\mathrm{Aut}(G)$ on T follows easily since the properties used to define T are invariant under $\mathrm{Aut}(G)$. $\square$

A *tree-decomposition* of a graph G consists of a tree T and a family $(P_t)_{t \in V(T)}$ of subsets of $V(G)$, one for each vertex of T such that

(T1) $V(G) = \bigcup_{t \in V(T)} P_t$,

(T2) for every edge $e \in E(G)$ there is $t \in V(T)$ such that both endpoints of e lie in P_t, and

(T3) $P_{t_1} \cap P_{t_3} \subseteq P_{t_2}$ whenever t_2 lies on the unique path connecting t_1 and t_3 in T.

The tree T is called *decomposition tree*, the sets P_t are called the *parts* of the tree-decomposition.

We associate to an edge $e = st$ of the decomposition tree a separation of G as follows. Removing e from T yields two components T_s and T_t. Let $X_s = \bigcup_{u\in T_s} P_u$ and $X_t = \bigcup_{u\in T_t} P_u$. If $X_s \setminus X_t$ and $X_t \setminus X_s$ are non-empty (this will be the case for all tree-decompositions considered in this chapter), then (X_s, X_t) is a proper separation of G. Clearly, the set of all separations associated to edges of a decomposition tree is nested.

The separators $A \cap B$ of the separations associated to edges of a decomposition tree are called *adhesion sets*. The supremum of the sizes of adhesion sets is called the *adhesion* of the tree-decomposition. The tree-decompositions constructed in this chapter all have finite adhesion.

Given a graph G with only one end ω and a tree-decomposition $(T, P_t \mid t \in V(T))$ of G of finite adhesion, then $(T, P_t \mid t \in V(T))$ *displays* ω if firstly the decomposition tree T has only one end; call it μ. And secondly for any edge st of T with μ in T_t, the associated separation (X_s, X_t) has the property that ω lies in X_t.

A tree-decomposition is $\mathrm{Aut}(G)$*-invariant* if the set S of separations associated to it is closed by the natural action of $\mathrm{Aut}(G)$ on S. The following implies Theorem 8.2.1.

Theorem 8.4.9. *Let G be a connected graph with only one end ω, which is undominated and has finite vertex degree k. Then G has a tree-decomposition $(T, P_t \mid t \in V(T))$ of adhesion k that displays ω and is* $\mathrm{Aut}(G)$*-invariant.*

Proof. We follow the notation of the proof of Theorem 8.4.8.

Given a vertex t of T_+, the *inward neighbourhood* of t, denoted by $N_+(t)$, is the set of vertices u of T_+ such that there is a directed edge from u to t in T_+. Recall that the vertices of T_+ are (in bijection with) separations; we refer to the separation associated to the vertex t by (A_t, B_t). Given a vertex t, we let $P_t = A_t \setminus \bigcup_{u\in N_+(t)}(A_u \setminus B_u)$.

It is straightforward that $(T, P_t \mid t \in V(T))$ is a tree-decomposition of adhesion k (whose set of associated separations is $\mathcal{S} \cup \{(B, A) \mid (A, B) \in \mathcal{S}\}$). It is not hard to see that $(T, P_t \mid t \in V(T))$ displays ω and is $\mathrm{Aut}(G)$-invariant. $\square$

Example 8.4.10. In this example we construct a one-ended graph G whose end is dominated and has vertex degree 1, but the graph G has no tree-decomposition of finite adhesion that is invariant under the group of automorphisms and whose decomposition tree is one-ended. We obtain G from the canopy tree by adding a new vertex adjacent to all the leaves of the canopy tree. Then we add infinitely many vertices of degree one only incident to that new vertex, see Figure 8.3.

Suppose for a contradiction that G has a tree-decomposition $(T, P_t \mid t \in V(T))$ of finite adhesion that is invariant under the group of automorphisms and such that T is one-ended.

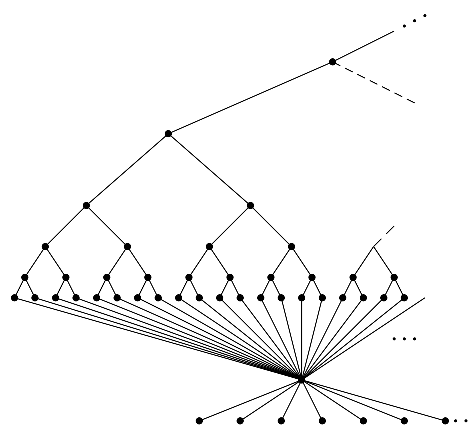

Figure 8.3: A graph with no $\mathrm{Aut}(G)$-invariant tree-decomposition of finite adhesion.

There cannot be a single part P_t that contains a ray of the canopy tree. To see that first note that there cannot be two such parts by the assumption of finite adhesion. Hence any such part would contain all vertices of the canopy tree from a certain level onwards. This is not possible by finite adhesion.

Having shown that there cannot be a single part P_t that contains a ray of the canopy tree, it must be that every part P_t with t near enough to the end of T contains a vertex of the canopy tree.

Our aim is to show that any vertex u of degree 1 is in all parts. Suppose not for a contradiction. Then since T is one-ended, there is a vertex t of T such that t separates in T all vertices s with $u \in P_s$ from the end of T. We pick t high enough in T such that there is a vertex v of the canopy tree in P_t. If P_t contained all vertices of the orbit of v, then P_t together with all parts P_s, where s has some fixed bounded distance from t in T, would contain a ray. This is impossible; the proof is similar as that that P_t cannot contain a ray. Hence there is a vertex v' in the orbit of v that is not in P_t. Take an automorphism of G that fixes u and moves v to v'. As the tree-decomposition is $\mathrm{Aut}(G)$-invariant, T has a vertex s such that $u, v' \in P_s$ but $v \notin P_s$. Since T is $\mathrm{Aut}(G)$-invariant and one-ended, t does not separate s from the end of T. This is a contradiction as $u \in P_s$.

Hence u must be in all parts. As u was arbitrary, every vertex of degree one must be in every part. So the tree-decomposition does not have finite adhesion. This is the desired contradiction. Hence such a tree-decomposition does not exist.

8.5 A dichotomy result for automorphism groups

Before we turn to a proof of Theorem 8.2.2, we state a few helpful auxiliary results. The following lemma can be seen as a consequence of [52, Lemma 7], but for completeness a direct proof is provided in Appendix B.

Lemma 8.5.1. *If T is a one-ended tree and R is a ray in T, then every automorphism of T fixes some tail of R pointwise.*

The next result is Lemma 3 in [52]. For completeness a proof is included in Appendix C.

Lemma 8.5.2. *The pointwise (and hence also the setwise) stabiliser of a finite set of vertices in the automorphism group of a rayless graph is either finite or contains at least $2^{\aleph_0}$ many elements.*

The next result is an extension of Lemma 8.5.2 to one-ended graphs where the end has finite vertex degree.

Lemma 8.5.3. *Let G be a graph with only one end ω. Assume that ω has finite vertex degree k. Let X be a finite set of vertices in G that contains all the vertices that dominate the end. If the graph $G - X$ is connected then the pointwise stabiliser of X in $\mathrm{Aut}(G)$ is either finite or contains at least $2^{\aleph_0}$ many elements.*

Proof. Denote by Γ the pointwise stabiliser of X in $\mathrm{Aut}(G)$. If Γ is finite, then there is nothing to show, hence assume that Γ is infinite.

Consider a nested $\mathrm{Aut}(G-X)$-invariant set of ω-relevant separations of $G-X$ as in Theorem 8.4.8 and a tree T built from this set in the way described. Clearly Γ gives rise to a subgroup of $\mathrm{Aut}(G - X)$ whence this nested set is Γ-invariant. Adding X to both sides of every separation in $\mathcal{S}$ gives rise to a new Γ invariant set $\mathcal{S}$ of nested separations such that each separation has order $k + |X|$. The tree we get from $\mathcal{S}$ is the same as T. From now on we will work with $\mathcal{S}$.

Every element $\gamma \in \Gamma$ induces an automorphism of T. Note that this canonical action of Γ on T is in general not faithful, i.e. it is possible that different elements of Γ induce the same automorphism of T.

Let R be a ray in T and let $(e_n)_{n\in\mathbb{N}}$ be the family of edges of R (in the order in which they appear on R). Let (A_n, B_n) be the separation of G corresponding to e_n. Denote by Γ_n the stabiliser of e_n in Γ. By Lemma 8.5.1 every automorphism of T (and hence also every $\gamma \in \Gamma$) fixes some tail of R, so Γ_n is non-trivial for large enough n. Furthermore, Γ_n is a subgroup of Γ_m whenever $n \leq m$.

We claim that for all but finitely many n, we have at least one non-trivial γ in the pointwise stabiliser of B_n. To see this, let $\gamma_1, \ldots, \gamma_{(k+|X|)!+1}$ be a set of $(k + |X|)! + 1$ different non-trivial automorphisms in Γ. Choose n large enough such that they all are contained in Γ_n and act differently on A_n. By a simple pigeon hole argument, at least two of them, γ_1 and γ_2 say, have the same action on $A_n \cap B_n$. Then $\gamma_1 \circ \gamma_2^{-1}$ is an automorphism which fixes $A_n \cap B_n$ pointwise, and fixes A_n setwise but not pointwise. Now, using the *independence property* from Section 8.3.2 we can define an automorphism

$$\gamma(x) = \begin{cases} \gamma_1 \circ \gamma_2^{-1}(x) & \text{if } x \in A_n \setminus B_n \\ x & \text{if } x \in B_n \end{cases}$$

with the desired properties.

Note that the subgroup leaving A_n invariant in the pointwise stabiliser of B_n in Γ induces the same permutation group on the rayless graph induced by A_n in G as does the subgroup leaving A_n invariant in the pointwise stabiliser of $A_n \cap B_n$. Hence, if there is $n \in \mathbb{N}$ such that the pointwise stabiliser of B_n in Γ is infinite, then this stabiliser contains at least $2^{\aleph_0}$ many elements by Lemma 8.5.2.

So (by passing to a tail of R) we may assume that the pointwise stabiliser of B_n is a finite but non-trivial subgroup of Γ for every $n \in N$.

Next we claim that for every n there is a non-trivial automorphism in the pointwise stabiliser of A_n. If not, then Γ_n is finite and we choose $\sigma \in \Gamma \setminus \Gamma_n$. For an edge e of T, denote by T_e the component of $T - e$ which does not contain the end of T. Clearly $\sigma(T_e) = T_{\sigma(e)}$ for every edge e. In particular, if $e = e_m$ is the last edge of R which is not fixed by σ, then clearly $\sigma(T_e) \subseteq T - T_e$. Furthermore $n < m$, so $A_n \subseteq A_m$, and $B_m \subseteq B_n$. Hence $\sigma(A_n) \subseteq \sigma(A_m) \subseteq B_m \subseteq B_n$. Now let γ be a nontrivial automorphism in the pointwise stabiliser of B_n. Then $\sigma^{-1} \circ \gamma \circ \sigma$ is easily seen to be a nontrivial element of the pointwise stabiliser of A_n: for $a \in A_n$ we have

$$\sigma^{-1} \circ \gamma \circ \sigma(a) = \sigma^{-1} \circ \sigma(a) = a$$

since $\sigma(a) \in B_n$ is fixed by γ.

Now define an infinite sequence $(\gamma_k)_{k \in \mathbb{N}}$ of elements of Γ as follows. Pick a nontrivial γ_1 in the pointwise stabiliser of A_1. Assume that γ_i has been defined for $i < k$, then let n_k be such that γ_i acts non-trivially on A_{n_k} for all $i < k$ and pick a nontrivial element γ_k in the pointwise stabiliser of A_{n_k}. For an infinite 0-1-sequence $(r_j)_{j \geq 1}$, define

$$\psi_i = \gamma_i^{r_i} \circ \gamma_{i-1}^{r_{i-1}} \circ \cdots \circ \gamma_1^{r_1},$$

in other words, ψ_n is the composition of all γ_j with $j \leq n$ and $r_j = 1$. Finally define ψ to be the limit of the ψ_n in the topology of pointwise convergence. This limit exists, because for $j > i$ the restriction ψ_i and ψ_j to A_{n_i} coincide, and the A_{n_i} exhaust $V(G)$. By Lemma 8.3.8, ψ is contained in $\mathrm{Aut}(G)$ and is also in $\Gamma \subseteq \mathrm{Aut}(G)$ because every ψ_i stabilises X pointwise.

Finally assume that we have two different 0-1-sequences $(r_j)_{j \geq 1}$ and $(r'_j)_{j \geq 1}$ and let $(\psi_j)_{j \geq 1}$ and $(\psi'_j)_{j \geq 1}$ be the corresponding sequences of automorphisms. If l is the first index such that $r_l \neq r'_l$ then the restrictions of ψ_l and ψ'_l (and hence also of ψ_i and ψ'_i for $i > l$) to A_{n_l} differ. Hence different 0-1-sequences give different elements of Γ and Γ contains at least $2^{\aleph_0}$ many elements. □

Theorem 8.2.2. *Let G be a graph with one end which has finite vertex degree. Then* $\mathrm{Aut}(G)$ *is either finite or has at least* $2^{\aleph_0}$ *many elements.*

Proof. Let X be the set of vertices which dominate ω. This set is possibly empty and by Lemma 8.3.3 it is finite. Every automorphism stabilises X setwise. Therefore the pointwise stabiliser of X is a normal subgroup of $\mathrm{Aut}(G)$ with finite index. So it suffices to show that the conclusion of Theorem 8.2.2 holds for the stabiliser Γ of X.

For every component C of $G - X$ let Γ_C be the pointwise stabiliser of X in $\mathrm{Aut}(C \cup X)$. Then Γ_C is either finite or contains at least $2^{\aleph_0}$ many elements by Lemma 8.5.2 and Lemma 8.5.3. If $|\Gamma_C| = 2^{\aleph_0}$ for some component C then we need do no more. So assume that all the groups Γ_C are finite. The same argument as used towards the end of the proof of Lemma 8.5.2 (see Appendix C) now shows that either Γ is finite or has at least cardinality $2^{\aleph_0}$. $\square$

As a corollary we can answer a question posed by Boutin and Imrich in [13]. In order to state this question, we first need some notation. For a vertex v in a graph G we define $B_v(n)$, *the ball of radius n centered at v*, as the set of all vertices in G in distance at most n from v. We also define $S_v(n)$, *the sphere of radius n centered at v*, as the set of all vertices in G in distance exactly n from v. A connected locally finite graph is said to have *linear growth* if there is a constant c such that $|B_v(n)| \leq cn$ for all $n = 1, 2, \ldots$. It is an easy exercise to show that the property of having linear growth does not depend on the choice of the vertex v.

In relation to their work on the distinguishing cost of graphs Boutin and Imrich [13] ask whether there exist one-ended locally finite graphs that has linear growth and countably infinite automorphism group.

If G is a locally finite graph with linear growth and v is a vertex in G then there is a constant k such that $|S_v(n)| = k$ for infinitely many values of n. (This is observed by Boutin and Imrich in their paper [13, Fact 2 in the proof of Proposition 13].) From this we deduce that the vertex-degree of an end of G is at most equal to k, since each ray in G must pass through all but finitely many of the spheres $S_v(n)$. Using Theorem 8.2.2 one can now give a negative answer to the above question.

Theorem 8.5.4. *If G is a connected locally finite graph with one end and linear growth, then the automorphism group of G is either finite or contains exactly $2^{\aleph_0}$ many elements.*

Proof. Since G is locally finite and connected, the graph G is countable. Hence the automorphism group cannot contain more than $2^{\aleph_0}$ many elements. Furthermore linear growth implies that all ends must have finite vertex degree, hence we can apply Theorem 8.2.2. $\square$

In particular a connected graph with linear growth and a countably infinite autormorphism group cannot have one end. Thus one can strengthen [13, Theorem 22] and get:

Theorem 8.5.5. (Cf. [13, Theorem 22]) *Every locally finite connected graph with linear growth and countably infinite automorphism group has 2 ends.*

Furthermore one can in [13, Theorem 18] remove the assumption that the graph is 2-ended, since it is implied by the other assumptions.

8.6 Ends of quasi-transitive graphs

Finally, another application was pointed out to the authors by Matthias Hamann. Recall that a graph is called *transitive*, if all vertices lie in the same orbit under the automorphism group, and *quasi-transitive* (or *almost-transitive*), if there are only finitely many orbits on the vertices.

The groundwork for the study of automorphisms of infinite graphs was laid in the 1973 paper of Halin [48]. Among the results there is a classification of automorphisms of a connected infinite graph, see [48, Sections 5, 6 and 7]. *Type 1* automorphisms, to use Halin's terminology, leave a finite set of vertices invariant. An automorphism is said to be of *type 2* if it is not of type 1. Type 2 automorphism are of two kinds, the first kind fixes precisely one end which is then thick (i.e. has infinite vertex degree) and the second kind fixes precisely two ends which are then both thin (i.e. have finite vertex degrees). In Halin's paper these results are stated with the additional assumption that the graph is locally finite but the classification remains true without this assumption.

It is a well known fact that a connected, transitive graph has either 1, 2, or infinitely many ends (follows for locally finite graphs from Halin's paper [44, Satz 2] and for the general case see [34, Corollary 4]). It is a consequence of a result of Jung [61] that if such a graph has more than one end then there is a type 2 automorphism that fixes precisely two ends and thus the graph has at least two thin ends. In particular, in the two-ended case both of the ends must be thin. Contrary to this, we deduce from Theorem 8.4.8 that the end of a one-ended transitive graph is always thick. This even holds in the more general case of quasi-transitive graphs. This was proved for locally finite graphs by Thomassen [89, Proposition 5.6]. A variant of this result for *metric ends* was proved by Krön and Möller in [64, Theorem 4.6].

Theorem 8.6.1. *If G is a one-ended, quasi-transitive graph, then the unique end is thick.*

For the proof we need the following auxiliary result.

Proposition 8.6.2. *There is no one-ended quasi-transitive tree.*

Proof. Assume that T is a quasi-transitive tree and that R is a ray in T. Then there is an edge-orbit under $\mathrm{Aut}(T)$ containing infinitely many edges of R. Contract all edges not in this orbit to obtain a tree T' whose automorphism group acts transitively on edges. Clearly, every end of T' corresponds to an end of T (there may be more ends of T which we contracted). But edge transitive trees must be either regular, or bi-regular. Hence T', and thus also T, has at least 2 ends. $\square$

Proof of Theorem 8.6.1. Assume for a contradiction that G is a quasi-transitive, one-ended graph whose end is thin.

If the end ω is dominated, then remove all vertices which dominate it and only keep the component C in which ω lies. The resulting graph is still quasi-transitive since C must be stabilised setwise by every automorphism. Furthermore, the degree of ω does not increase by deleting parts of the graph. Hence we can without loss of generality assume that the end of the counterexample G is undominated.

Now apply Theorem 8.4.8 to G. This gives a nested set $\mathcal{S}$ of separations which is invariant under automorphisms—in particular, there are only finitely many orbits of $\mathcal{S}$ under the action of $\mathrm{Aut}(G)$. Theorem 8.4.8 further tells us that there is a bijection between $\mathcal{S}$ and the edges of a one-ended tree T such that the action of $\mathrm{Aut}(G)$ on $\mathcal{S}$ induces an action on T by automorphisms. Hence T is a quasi-transitive one-ended tree, which contradicts Proposition 8.6.2. $\square$

8.7 Appendix A

We say that a vertex v *dominates* a ray L if there are infinitely many $v-L$ paths, any two only having v as a common vertex. It follows from the definition of an end that if a vertex domintes one ray belonging to an end then it dominates every ray belonging to that end and dominates the end.

Proof of Lemma 8.3.3. Assume that the set X of dominating vertices is infinite. By the above we can assume that there is a ray R and infinitely many vertices $x_1, x_2, \dots$ that dominate R in G. We show that G must then contain a subdivision of the complete graph on $x_1, x_2, \dots$. Start by taking vertices v_1 and v_2 on R_1 such that there are disjoint $x_1 - v_1$ and $x_2 - v_2$ paths. Then we find vertices w_1 and w_2 furher along the ray R_1 such that there are disjoint $x_1 - w_1$ and $x_3 - w_3$ paths and still further along we find vertices u_2 and u_3 such that there are disjoint $x_2 - u_2$ and $x_3 - u_3$ paths. Adding the relevant segments of R we find $x_1 - x_2$, $x_1 - x_3$ and $x_2 - x_3$ paths having at most their endvertices in common. The subgraph of G consisting of these three paths is thus a subdivision of the complete graph on three vertices. Using induction we can find an increasing sequence of subgraphs H_n of G that contains the vertices $x_1, x_2, \dots, x_n$ and also paths P_{ij} linking x_i and x_j such that any two such paths have at most their end vertices in common. The subgraph H_n is a subdivision of the complete graph on n-vertices. The subgraph $H = \bigcup_{i=1}^{\infty} H_i$ is a subdivision of the complete graph on (countably) infinite set of vertices and contains an infinite family of pairwise disjoint rays that all belong to the end ω. This contradicts our assumptions and we conclude that T must be finite. $\square$

A *ray decomposition*[4] of *adhesion* m of a graph G consists of subgraphs $G_1, G_2, \dots$ such that:

1. $G = \bigcup_{i=1}^{\infty} G_i$;

[4] Halin used the German term 'schwach m-fach kettenförmig'.

2. if $T_{n+1} = \big(\bigcup_{i=1}^{n} G_i\big) \cap G_{n+1}$ then $|T_{n+1}| = m$ and $T_{n+1} \subseteq G_n \setminus \big(\bigcup_{i=1}^{n-1} G_i\big)$ for $n = 1, 2, \ldots$;

3. for each value of $n = 1, 2, \ldots$ there are m pairwise disjoint paths in G_{n+1} that have their initial vertices in T_{n+1} and teminal vertices in T_{n+2};

4. none of the subgraphs G_i contains a ray.

The following Menger-type result is used by Halin in his proof of [46, Satz 2]. In the proof we also use ideas from another one of Halin's papers [45, Proof of Satz 3].

Theorem 8.7.1. *Let G be a locally finite connected graph with the property that G contains a family of m pairwise disjoint rays but there is no such family of $m+1$ pairwise disjoint rays. Then there is in G a family of pairwise disjoint separators $T_1, T_2, \ldots$ such that each contains precisely m vertices and a ray in G must for some n_0 intersects all the sets T_n for $n \geq n_0$.*

Proof. Fix a reference vertex v_0 in G. Let E_j denote the set of vertices in distance precisely j from v_0. Define also B_i as the set of vertices in distance at most i from v_0. For numbers i and j such that $i+1 < j$ we construct a new graph H_{ij} such that we start with the subgraph of G induces by B_j, then we remove B_i but add a new vertex a that has as its neighbourhood the set ∂B_i (for a set C of vertices ∂C denotes the set of vertices that are not in C but are adjacent to some vertex in C) and we also add a new vertex b that has every vertex in $\partial(G \setminus B_j)$ as its neighbour. Since G is assumed to be locally finite the graph H_{ij} is finite. (By abuse of notation we do not distinguish the additional vertices a and b in different graphs H_{ij}.)

Suppose that, for a fixed value of i, there are always for j big enough at least k distinct $a-b$ paths in H_{ij} such that any two of them interesect only in the vertices a and b. Then one can use the same argument as in the proof of König's Infinity Lemma to show that then G contains a family of k pairwise disjoint rays. Because G does not contain a family of $m+1$ pairwise disjoint rays there are for each i a number j_i such that for every $j \geq j_i$ there are at most m disjoint $a-b$ paths in H_{ij_i}. Since a and b are not adjacent in H_{ij_i} then the Menger Theorem says that minimum number of a vertices in an $a-b$ separator is equal to the maximal number of $a-b$ paths such that any two of the paths have no inner vertices in common. Whence there is in $H_{ij_i} \setminus \{a, b\}$ a set T and $a-b$ separator with precisely m vertices. This set is also an separator in G and every ray in G that has its initial vertex in B_i must intersect T. From this information we can easily construct our sequence of separators $T_1, T_2, \ldots$.

We can also clearly assume that if i_j is the smallest number such that T_j is in B_{i_j} then $T_k \cap B_{i_j} = \emptyset$ for all $k > j$. $\square$

Corollary 8.7.2. *Let G be a connected locally finite graph. Suppose ω is an end of G and ω has finite vertex degree m. Then there is a sequence $T_1, T_2, \ldots$ of separators each containing precisely m vertices such that if C_i denotes the component of $G - T_i$ that ω belongs to then $C_1 \supseteq C_2 \supseteq \ldots$ and $\bigcap_{i=1}^{\infty} C_i = \emptyset$.*

Proof. We use exactly the same argument as above except that when we construct the H_{ij} we only put in edges from b to those vertices in E_j that are in the boundary of the component of $G \setminus B_j$ that ω lies in. □

Proof of Lemma 8.3.5. The first part of the Lemma about the existence of a family of k pairwise disjoint rays in ω with their initial vertices in $A \cap B$ follows directly from the above.

For the second part, the only thing we need to show is that there cannot exist a separation (C, D) of order $< k$ such that $A \subseteq C$ and ω lies in D. Such a separation cannot exist because the k pairwise disjoint rays that have their initial vertices in $A \cap B$ and belong to ω would all have to pass through $C \cap D$. □

Theorem 8.7.3. ([46, Satz 2]) *Let G be a graph with the property that it contains a family of m pairwise disjoint rays but no family of $m+1$ pairwise disjoint rays. Let X denote the set of vertices in G that dominate some ray. Then the set X is finite and the graph $G - X$ has a ray decomposition of adhesion m.*

Proof. Let $R_1, \ldots, R_m$ denote a family of pairwise disjoint rays. Set $R = R_1 \cup \cdots \cup R_m$.

Any ray in G must intersect the set R in infinitely many vertices and thus intersects one of the rays $R_1, \ldots, R_m$ in infinitely many vertices. From this we conclude that every ray in G is in the same end as one of the rays $R_1, \ldots, R_m$. Thus a vertex that dominates some ray in G must dominate one of the rays $R_1, \ldots, R_m$.

In Lemma 8.3.3 we have already shown that the set of vertices dominating an end of finite vertex degree is finite. Note also that if a vertex in R is in infinitely many distinct sets of the type ∂C where C is a component of $G \setminus R$ then x would be a dominating vertex of some ray R_i. Thus there can only be finitely many vertices in R with this property.

We will now show that $G - X$ has a ray decomposition of adhesion m. To simplify the notation we will in the rest of the proof assume that X is empty.

Assume now that there is a component C of $G - R$ such that ∂C is infinite. Take a spanning tree of C and then adjoin the vertices in ∂C to this tree using edges in G. Now we have a tree with infinitely many leafs. It is now apparent that either the tree contains a ray that does not intersect R or there is a vertex in C that dominates a ray in G. Both possibilities are contrary to our assumptions and we can conclude that ∂C is finite for every component C of $G \setminus R$.

For every set S in R of such that $S = \partial C$ for some component C in $G \setminus R$ we find a locally finite connected subgraph C_S of $C \cup S$ containing S. The graph G' that is the union of R and all the subgraphs C_S is a locally finite graph. The original graph G has a ray decomposition of adhesion m if and only if G' has a ray decomposition of adhesion m.

At this point we apply Theorem 8.7.1. From Theorem 8.7.1 we have the sequence $T_2, T_3, \ldots$ of separators. We choose T_2 such that all the rays $R_1, \ldots, R_m$ intersect T_2. We start by defining G_i for $i \geq 2$ as the union of T_i and all those components of $G - T_i$ that contain the tail of some ray R_i. Finally, set $G_1 = G \setminus (G_2 \setminus T_2)$. Note that none of the subgraphs G_i can contain a ray and our family of rays provides a family of m pairwise disjoint $T_i - T_{i+1}$ paths. Now we have shown that G has a ray decomposition of adhesion m. □

Finally, we are now ready to show how Halin's result above implies Theorem 8.3.6 that concerns ω-relevant separations.

Proof of Theorem 8.3.6. We continue with the notation in the proof of Theorem 8.7.3. Recall that there are infinitely many pairwise disjoint paths connecting a ray R_i to a ray R_j. Thus we may assume that the initial vertices of the rays $R_1, \ldots R_k$ all belong to the same component of $G - T_2$. We set A_n as the union of the component of $G - T_{n+1}$ that contains these initial vertices with T_{n+1}. Then set $B_n = (G \setminus A_n) \cup T_{n+1}$. Now it is trivial to check that the sequence (A_n, B_n) of separations satisfies the conditions. □

8.8 Appendix B

Proof of Lemma 8.5.1. Let σ be an automorphism of T. In cite [90, Proposition 3.2] Tits proved that there are three types of automorphisms of a tree: (i) those that fix some vertex, (ii) those that fix no vertex but leave an edge invariant and (iii) those that leave some double-ray $\ldots, v_{-1}, v_0, v_1, v_2, \ldots$ invariant and act as non-trivial translations on that double-ray. (Similar results were proved independently by Halin in [48].) Since T is one-ended it contains no double-ray and thus (iii) is impossible. Suppose now that σ fixes no vertex in T but leaves the edge e invariant. The end of T lives in one of the components of $T - e$ and σ swaps the two components of $T - e$. This is impossible, because T has only one end and this end must belong to one of the components of $T - e$. Hence σ must fix some vertex v. There is a unique ray R' in T with v as an initial vertex and this ray is fixed pointwise by σ. The two rays R and R' intersect in a ray that is a tail of R and this tail of R is fixed pointwise by σ. □

8.9 Appendix C

In this Appendix we prove Lemma 8.5.2 which is a slightly sharpened version of Lemma 3 from Halin's paper [52]. The change is that 'uncountable' in Halin's results is replaced by 'at least $2^{\aleph_0}$ elements'.

First there is an auxilliary result that corresponds to Lemma 2 in [52].

Lemma 8.9.1. *Let G be a connected graph and $\Gamma = \mathrm{Aut}(G)$. Suppose D is a subset of the vertex set of G. Let $\{C_i\}_{i \in I}$ denote the family of components of $G - D$. Define G_i as the subgraph spanned by $C_i \cup \partial C_i$. Set $\Gamma_i = \mathrm{Aut}(G_i)_{(\partial C_i)}$. Suppose that Γ_i is either finite or has at least $2^{\aleph_0}$ elements for all i. Then $\Gamma_{(D)}$ is either finite or has at least $2^{\aleph_0}$ elements.*

Proof. If one of the groups γ_i has at least $2^{\aleph_0}$ elements then there is nothing more to do. So, we assume that all these groups are finite.

Now there are two situations where it is possible that $\Gamma_{(D)}$ is infinite. The first is when infinitely many of the groups Γ_i are non-trivial. For any family $\{\sigma_i\}_{i\in I}$ such that $\sigma_i \in \Gamma_i$ we can find an automorphism $\sigma \in \Gamma_{(G\setminus C_i)} \subseteq \Gamma_{(D)}$ such that the restriction to C_i equals σ_i for all i. If infinitely many of the groups Γ_{C_i} are nontrivial, then there are at least $2^{\aleph_0}$ such families $\{\sigma_i\}_{i\in I}$ and $\Gamma_{(D)}$ must have at least $2^{\aleph_0}$ elements.

We say that two components C_i and C_j are equivalent if $\partial C_i = \partial C_j$ and there is an isomorphism φ_{ij} from the subgraph G_i to the subgraph G_j fixing every vertex in $\partial C_i = \partial C_j$. Clearly there is an automorpism σ_{ij} of G that fixes every vertex that is neither in C_i nor C_j such that $\sigma_{ij}(v) = \varphi_{ij}(v)$ for $v \in C_i$ and $\sigma_{ij}(v) = \varphi_{ij}^{-1}(v)$ for $v \in C_j$. If there are infinitely many disjoint ordered pairs of equivalent components we can for any subset of these pairs find an automorphism $\sigma \in \Gamma_{(D)}$ such that if (C_i, C_j) is in our subset then the restriction of σ to $C_i \cup C_j$ is equal to the restriction of σ_{ij}. There are at least $2^{\aleph_0}$ such sets and thus $\Gamma_{(D)}$ has at least $2^{\aleph_0}$ elements.

If neither of the two cases above occurs then $\Gamma_{(D)}$ is clearly finite. $\square$

Proof of Lemma 8.5.2. Following Schmidt [85] (see also Halin's paper [51, Section 3]) we define, using induction, for each ordinal λ a class of graphs $A(\lambda)$. The class $A(0)$ is the class of finite graphs. Suppose $\lambda > 0$ and $A(\mu)$ has already been defined for all $\mu < \lambda$. A graph G is in the class $A(\lambda)$ if and only if it contains a finite set F of vertices such that each component of $G - F$ is in $A(\mu)$ for some $\mu < \lambda$. It is shown in the papers referred to above that if G belongs to $A(\lambda)$ for some ordinal λ then G is rayless and, conversely, every rayless graph belongs to $A(\lambda)$ for some ordinal λ. For a rayless graph G we define $o(G)$ as the smallest ordinal λ such that G is in $A(\lambda)$.

The Lemma is proved by induction over $o(G)$. If $o(G) = 0$ then the graph G is finite and the automorphism group is also finite.

Assume that the result is true for all rayless graphs H such that $o(H) < o(G)$. Find a finite set F of vertices such that each of the components of $G - F$ has a smaller order than G. Denote the family of components of $G - F$ with $\{C_i\}_{i\in I}$. Denote with G_i the subgraph induced by $C_i \cup \partial C_i$. By induction hypothesis the pointwise stabiliser of ∂C_i in $\mathrm{Aut}(G_i)$ is either finite or has at least $2^{\aleph_0}$ elements. Lemma 8.9.1 above implies that $\mathrm{Aut}(G)_{(D)}$ is either finite or has at least $2^{\aleph_0}$ elements. $\square$

Part III

Infinite matroids

Chapter 9

Matroid intersection, base packing and base covering for infinite matroids

9.1 Abstract

As part of the recent developments in infinite matroid theory, there have been a number of conjectures about how standard theorems of finite matroid theory might extend to the infinite setting. These include base packing, base covering, and matroid intersection and union. We show that several of these conjectures are equivalent, so that each gives a perspective on the same central problem of infinite matroid theory. For finite matroids, these equivalences give new and simpler proofs for the finite theorems corresponding to these conjectures.

This new point of view also allows us to extend, and simplify the proofs of, some cases where these conjectures were known to be true.

9.2 Introduction

The well-known finite matroid intersection theorem of Edmonds states that for any two finite matroids M and N the size of a biggest common independent set is equal to the minimum of the rank sum $r_M(E_M) + r_N(E_N)$, where the minimum is taken over all partitions $E = E_M \dot\cup E_N$. The same statement for infinite matroids is true, but for a silly reason [32], which suggests that more care is needed in extending this statement to the infinite case.

Nash-Williams [4] proposed the following for finitary matroids.

Conjecture 9.2.1 (The Matroid Intersection Conjecture). *Any two matroids M and N on a common ground set E have a common independent set I admitting a partition $I = J_M \cup J_N$ such that $\mathrm{Cl}_M(J_M) \cup \mathrm{Cl}_N(J_N) = E$.*

For finite matroids this is easily seen to be equivalent to the intersection theorem, which is why we refer to Conjecture 9.2.1 as the Matroid Intersection Conjecture. If for a pair of matroids M and N on a common ground set there are sets I, J_M and J_N as in Conjecture 9.2.1, we say that M and N have the *Intersection property*, and that I, J_M and J_N *witness* this.

In [6], it was shown that this conjecture implies the celebrated Aharoni-Berger-Theorem [2], also known as the Erdős-Menger-Conjecture. Call a matroid *finitary* if all its circuits are finite and *co-finitary* if its dual is finitary. The conjecture is true in the cases where M is finitary and N is co-finitary [6].[1] Aharoni and Ziv [4] proved the conjecture for one matroid finitary and the other a countable direct sum of finite rank matroids.

In this chapter we will demonstrate that the Matroid Intersection Conjecture is a natural formulation by showing that it is equivalent to several other new conjectures in unexpectedly different parts of infinite matroid theory.

Suppose we have a family of matroids $(M_k|k \in K)$ on the same ground set E. A *packing* for this family consists of a spanning set S_k for each M_k such that the S_k are all disjoint. Note that not all families of matroids have a packing. More precisely, the well-known finite base packing theorem states that if E is finite then the family has a packing if and only if for every subset $Y \subseteq E$ the following holds.

$$\sum_{k \in K} r_{M_k.Y}(Y) \leq |Y|$$

The Aharoni-Thomassen graphs [3, 36] show that this theorem does not extend verbatim to finitary matroids. However, the base packing theorem extends to finite families of co-finitary matroids [5]. This implies the topological tree packing theorems of Diestel and Tutte. Independently from our main result, we close the gap in between by showing that the base packing theorem extends to arbitrary families of co-finitary matroids (for example, topological cycle matroids).

Similar to packings are coverings: a *covering* for the family $(M_k|k \in K)$ consists of an independent set I_k for each M_k such that the I_k cover E. And analogously to the base packing theorem, there is a base covering theorem characterising the finite families of finite matroids admitting a covering.

We are now in a position to state our main conjecture, which we will show is equivalent to the intersection conjecture. Roughly, the finite base packing theorem says that a family has a packing if it is very dense. Similarly, the finite base covering theorem says roughly that a family has a covering if it is very sparse. Although not every family of matroids has a packing and not every family has a covering, we could ask: is it always possible to divide the ground set into a "dense" part, which has a packing, and a "sparse" part, which has a covering?

Definition 9.2.2. We say that a family of matroids $(M_k|k \in K)$ on a common ground set E, has the *Packing/Covering* property if E admits a partition $E = P\dot{\cup}C$ such that $(M_k\restriction_P|k \in K)$ has a packing and $(M_k.C|k \in K)$ has a covering.

[1] In fact in [6] the conjecture was proved for a slightly larger class.

Conjecture 9.2.3. *Any family of matroids on a common ground set has the Packing/Covering property.*

Here $M_k\restriction_P$ is the restriction of M_k to P and $M_k.C$ is the contraction of M_k onto C. Note that if $(M_k\restriction_P|k \in K)$ has a packing, then $(M_k.P|k \in K)$ has a packing, so we get a stronger statement by taking the restriction here. Similarly, we get a stronger statement by contracting to get the family which should have a covering than we would get by restricting.

For finite matroids, we show that this new conjecture is true and implies the base packing and base covering theorems. So the finite version of Conjecture 9.2.3 unifies the base packing and the base covering theorem into one theorem.

For infinite matroids, we show that Conjecture 9.2.3 and the intersection conjecture are equivalent, and that both are equivalent to Conjecture 9.2.3 for pairs of matroids. In fact, for pairs of matroids, we show that (M, N) has the Packing/Covering property if and only if M and N^* have the Intersection property. As the Packing/Covering property is preserved under duality for pairs of matroids, this shows the less obvious fact that the Intersection property is also preserved under duality:

Corollary 9.2.4. *If M and N are matroids on the same ground set then M and N have the intersection property if and only if M^* and N^* do.*

Conjecture 9.2.3 also suggests a base packing conjecture and a base covering conjecture which we show are equivalent to the intersection conjecture but not to the above mentioned rank formula formulation of base packing for infinite matroids.

The various results about when intersection is true transfer via these equivalences to give results showing that these new conjectures also hold in the corresponding special cases. For example, while the rank-formulation of the covering theorem is not true for all families of co-finitary matroids, the new covering conjecture is true in that case. This yields a base covering theorem for the algebraic cycle matroid of any locally finite graph and the topological cycle matroid of any graph. Similarly, we immediately obtain in this way that the new packing and covering conjectures are true for finite families of finitary matroids. Thus we get packing and covering theorems for the finite cycle matroid of any graph.

For finite matroids, the proofs of the equivalences of these conjectures simplify the proofs of the corresponding finite theorems.

We show that Conjecture 9.2.3 might be seen as the infinite analogue of the rank formula of the matroid union theorem. It should be noted that there are two matroids whose union is not a matroid [5], so there is no infinite analogue of the finite matroid union theorem as a whole.

This new point of view also allows us to give a simplified account of the special cases of the intersection conjecture and even to extend the results a little bit. Our result includes the following:

Theorem 9.2.5. *Any family of matroids $(M_k|k \in K)$ on the same ground set*

E for which there are only countably many sets appearing as circuits of matroids in the family has the Packing/Covering property.

This chapter is organised as follows: In Section 2, we recall some basic matroid theory and introduce a key idea, that of exchange chains. After this, in Section 3, we restate our main conjecture and look at its relation to the infinite matroid intersection conjecture. In Section 4, we prove a special case of our main conjecture. In the next two sections, we consider base coverings and base packings of infinite matroids. In the final section, Section 7, we give an overview over the various equivalences we have proved.

9.3 Preliminaries

9.3.1 Basic matroid theory

Throughout, notation and terminology for graphs are that of [36], for matroids that of [76, 22], and for topology that of [10]. M always denotes a matroid and $E(M)$, $\mathcal{I}(M)$, $\mathcal{B}(M)$, $\mathcal{C}(M)$ and $\mathcal{S}(M)$ denote its ground set and its sets of independent sets, bases, circuits and spanning sets, respectively.

Recall that the set $\mathcal{I}(M)$ is required to satisfy the following *independence axioms* [22]:

(I1) $\emptyset \in \mathcal{I}(M)$.

(I2) $\mathcal{I}(M)$ is closed under taking subsets.

(I3) Whenever $I, I' \in \mathcal{I}(M)$ with I' maximal and I not maximal, there exists an $x \in I' \setminus I$ such that $I + x \in \mathcal{I}(M)$.

(IM) Whenever $I \subseteq X \subseteq E$ and $I \in \mathcal{I}(M)$, the set $\{I' \in \mathcal{I}(M) \mid I \subseteq I' \subseteq X\}$ has a maximal element.

The axiom (IM) for the dual M^* of M is equivalent to the following:

(IM*) Whenever $Y \subseteq S \subseteq E$ and $S \in \mathcal{S}(M)$, the set $\{S' \in \mathcal{S}(M) \mid Y \subseteq S' \subseteq S\}$ has a minimal element.

As the dual of any matroid is also a matroid, every matroid satisfies this. We need the following facts about circuits, the first of which is commonly referred to as the infinite circuit elimination axiom [22]:

(C3) Whenever $X \subseteq C \in \mathcal{C}(M)$ and $\{C_x \mid x \in X\} \subseteq \mathcal{C}(M)$ satisfies $x \in C_y \Leftrightarrow x = y$ for all $x, y \in X$, then for every $z \in C \setminus \left(\bigcup_{x \in X} C_x\right)$ there exists a $C' \in \mathcal{C}(M)$ such that $z \in C' \subseteq \left(C \cup \bigcup_{x \in X} C_x\right) \setminus X$.

(C4) Every dependent set contains a circuit.

A matroid is called *finitary* if every circuit is finite.

Lemma 9.3.1. *A set S is M-spanning iff it meets every M-cocircuit.*

Proof. We prove the dual version where $I := E(M) \setminus S$.

$$\text{A set } I \text{ is } M^*\text{-independent iff it does not contain an } M^*\text{-circuit.} \tag{9.1}$$

Clearly, if I contains a circuit, then it is not independent. Conversely, if I is not independent, then by (C4) it also contains a circuit. □

Let 2^X denote the power set of X. If $M = (E, \mathcal{I})$ is a matroid, then for every $X \subseteq E$ there are matroids $M{\restriction}_X := (X, \mathcal{I} \cap 2^X)$ (called the *restriction* of M to X), $M \backslash X := M{\restriction}_{E\setminus X}$ (which we say is obtained from M by *deleting* X)[2], $M.X := (M^*{\restriction}_X)^*$ (which we say is obtained by *contracting onto* X) and $M/X := M.(E \setminus X)$ (which we say is obtained by *contracting* X). For $e \in E$, we will also denote $M/\{e\}$ by M/e and $M\backslash\{e\}$ by $M\backslash e$.

Given a base B of X (that is, a maximal independent subset of X), the independent sets of M/X can be characterised as those subsets I of $E \setminus X$ for which $B \cup I$ is independent in M.

Lemma 9.3.2. *Let M be a matroid with ground set $E = C \dot\cup X \dot\cup D$ and let o' be a circuit of $M' = M/C\backslash D$. Then there is an M-circuit o with $o' \subseteq o \subseteq o' \cup C$.*

Proof. Let s be any M-base of C. Then $s \cup o'$ is M-dependent since o' is M'-dependent. On the other hand, $s \cup o' - e$ is M-independent whenever $e \in o'$ since $o' - e$ is M'-independent. Putting this together yields that $s \cup o'$ contains an M-circuit o, and this circuit must not avoid any $e \in o'$, as desired. □

For a family $(M_k | k \in K)$ of matroids, where M_k has ground set E_k, the *direct sum* $\bigoplus_{k\in K} M_k$ is the matroid with ground set $\bigcup_{k\in K} E_k \times \{k\}$, with independent sets the sets of the form $\bigcup_{k\in K} I_k \times \{k\}$ where for each k the set I_k is independent in M_k. Contraction and deletion commute with direct sums, in the sense that for a family $(X_k \subseteq E_k | k \in K)$ we have $\bigoplus_{k\in K}(M_k/X_k) = (\bigoplus_{k\in K} M_k)/(\bigcup_{k\in K} X_k \times \{k\})$ and $\bigoplus_{k\in K}(M_k\backslash X_k) = (\bigoplus_{k\in K} M_k)\backslash(\bigcup_{k\in K} X_k \times \{k\})$

Lemma 9.3.3. *Let M be a matroid and $X \subseteq E(M)$. If $S_1 \subseteq X$ spans $M{\restriction}_X$ and $S_2 \subseteq E \setminus X$ spans M/X, then $S_1 \cup S_2$ spans M.*

Proof. Let B be a maximal independent subset of S_1. Then B spans S_1 and S_1 spans X, so B spans X. Thus B is a base of X. Now let $e \in M \setminus X \setminus S_2$. Since $e \in \mathrm{Cl}_{M/X}(S_2)$ there is a set $I \subseteq E \setminus X$ such that I is M/X-independent but $I + e$ is not. Then $B \cup I$ is M-independent but $B \cup I + e$ is not, so that $e \in \mathrm{Cl}_M(S_1 + S_2)$, as witnessed by the set $B + I$. Any other element of E is either in S_2 or is in $X \subseteq \mathrm{Cl}_M(S_1)$, and so is in the span of $S_1 \cup S_2$. □

Lemma 9.3.4 ([23], Lemma 5)**.** *Let M be a matroid with a circuit C and a co-circuit D, then $|C \cap D| \neq 1$.*

[2] We use the notation $M{\restriction}_X$ rather than the conventional notation $M|X$ to avoid confusion with our notation $(M_k|k \in K)$ for families of matroids.

A particular class of matroids we shall employ is the *uniform* matroids $U_{n,E}$ on a ground set E, in which the bases are the subsets of E of size n. In fact, the matroids we will use are those of the form $U^*_{1,E}$, in which the bases are all those sets obtained by removing a single element from E. Such a matroid is said to consist of a single circuit, because $\mathcal{C}(U^*_{1,E}) = \{E\}$. A subset is independent iff it isn't the whole of E. Note that for a subset X of E, $U^*_{1,E}\restriction_X$ is free (every subset is independent) unless X is the whole of E, and $U^*_{1,E}.X = U^*_{1,X}$ unless X is empty.

9.3.2 Exchange chains

Below, we will need a modification of the concept of exchange chains introduced in [5]. The only modification is that we need not only exchange chains for families with two members but more generally exchange chains for arbitrary families, which we define as follows: Let $(M_k|k \in K)$ be a family of matroids and let $B_k \in \mathcal{I}(M_k)$. A $(B_k|k \in K)$*-exchange chain (from y_0 to y_n)* is a tuple $(y_0, k_0; y_1, k_1; \ldots; y_n)$ where $B_{k_l} + y_l$ includes an M_{k_l}-circuit containing y_l and y_{l+1}. A $(B_k|k \in K)$-exchange chain from y_0 to y_n is called *shortest* if there is no $(B_k|k \in K)$-exchange chain $(y'_0, k'_0; y'_1, k'_1; \ldots; y'_m)$ with $y'_0 = y_0$, $y'_m = y_n$ and $m < n$. A typical exchange chain is shown in Figure 9.1.

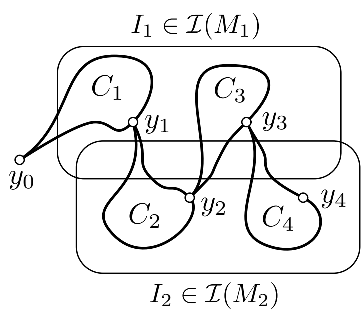

(a) Before the exchange

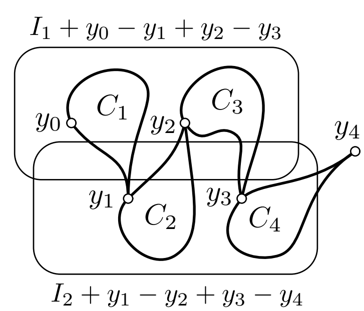

(b) After the exchange

Figure 9.1: An (I_1, I_2)-exchange chain of length 4.

Lemma 9.3.5. *Let $(M_k|k \in K)$ be a family of matroids and let $B_k \in \mathcal{I}(M_k)$. If $(y_0, k_0; y_1, k_1; \ldots; y_n)$ is a shortest $(B_k|k \in K)$-exchange chain from y_0 to y_n, then $B'_k \in \mathcal{I}(M_k)$ for every k, where*

$$B'_k := B_k \cup \{y_l|k_l = k\} \setminus \{y_{l+1}|k_l = k\}$$

Moreover, $\mathrm{Cl}_{M_k} B_k = \mathrm{Cl}_{M_k} B'_k$.

Proof (Sketch). The proof that the B'_k are independent is done by induction on n and is that of Lemma 4.5 in [5]. To see the second assertion, first note that

$\{y_l | k_l = k\} \subseteq \mathrm{Cl}_{M_k} B_k$ and thus $B'_k \subseteq \mathrm{Cl}_{M_k} B_k$. Thus it suffices to show that $B_k \subseteq \mathrm{Cl}_{M_k} B'_k$. For this, note that the reverse tuple $(y_n, k_{n-1}; y_{n-1}, k_{n-2}; \dots; y_0)$ is a B'_k-exchange chain giving back the original B_k, so we can apply the preceding argument again. □

Lemma 9.3.6. *Let M be a matroid and $I, B \in \mathcal{I}(M)$ with B maximal and $B \setminus I$ finite. Then $|I \setminus B| \leq |B \setminus I|$.*

Lemma 9.3.7. *Let $(M_k | k \in K)$ be a family of matroids, let $B_k \in \mathcal{I}(M_k)$ and let C be a circuit for some M_{k_0} such that $C \setminus B_{k_0}$ only contains one element, e. If there is a $(B_k | k \in K)$-exchange chain from x_0 to e, then for every $c \in C$, there is a $(B_k | k \in K)$-exchange chain from x_0 to c.*

Proof. Let $(y_0 = x_0, k_0; y_1, k_1; \dots; y_n = e)$ be an exchange chain from x_0 to e. Then $(y_0 = x_0, k_0; y_1, k_1; \dots; y_n = e, k_0; c)$ is the desired exchange chain. □

9.4 The Packing/Covering conjecture

The matroid union theorem is a basic result in the theory of finite matroids. It gives a way to produce a new matroid $M = \bigvee_{k \in K} M_k$ from a finite family $(M_k | k \in K)$ of finite matroids on the same ground set E. We take a subset I of E to be M-independent iff it is a union $\bigcup_{k \in K} I_k$ with each I_k independent in the corresponding matroid M_k. The fact that this gives a matroid is interesting, but a great deal of the power of the theorem comes from the fact that it gives an explicit formula for the ranks of sets in this matroid:

$$r_M(X) = \min_{X = P \dot\cup C} \sum_{k \in K} r_{M_k}(P) + |C| \tag{9.2}$$

Here the minimisation is over those pairs (P, C) of subsets of X which partition X.

For infinite matroids, or infinite families of matroids, this theorem is no longer true [5], in that M is no longer a matroid. However, it turns out, as we shall now show, that we may conjecture a natural extension of the rank formula to infinite families of infinite matroids.

First, we state the formula in a way which does not rely on the assumption that M is a matroid:

$$\max_{I_k \in \mathcal{I}(M_k)} \left| \bigcup_{k \in K} I_k \right| = \min_{E = P \dot\cup C} \sum_{k \in K} r_{M_k}(P) + |C| \tag{9.3}$$

Note that this is really only the special case of (9.2) with $X = E$. However, it is easy to deduce the more general version by applying (9.3) to the family $(M_k \restriction_X | k \in K)$.

Note also that no value $|\bigcup_{k \in K} I_k|$ appearing on the left is bigger than any value $\sum_{k \in K} r_{M_k}(P) + |C|$ appearing on the right. To see this, note that $|\bigcup_{k \in K} (I_k \cap P)| \leq \sum_{k \in K} r_{M_k}(P)$ and $\bigcup_{k \in K} (I_k \cap C) \subseteq C$. So the formula is equivalent to the statement that we can find $(I_k|k \in K)$ and P and C with $P \dot{\cup} C = E$ so that

$$\left| \bigcup_{k \in K} I_k \right| = \sum_{k \in K} r_{M_k}(P) + |C| \,. \tag{9.4}$$

For this, what we need is to have equality in the two inequalities above, so we get

$$\left| \bigcup_{k \in K} (I_k \cap P) \right| = \sum_{k \in K} r_{M_k}(P) \text{ and } \bigcup_{k \in K} (I_k \cap C) = C \,. \tag{9.5}$$

The equation on the left can be broken down a bit further: it states that each $I_k \cap P$ is spanning (and so a base) in the appropriate matroid $M_k \restriction_P$, and that all these sets are disjoint. This is the familiar notion of a packing:

Definition 9.4.1. Let $(M_k|k \in K)$ be a family of matroids on the same ground set E. A *packing* for this family consists of a spanning set S_k for each M_k such that the S_k are all disjoint.

So the $I_k \cap P$ form a packing for the family $(M_k \restriction_P |k \in K)$. In fact, in this case, each $I_k \cap P$ is a base in the corresponding matroid. In Definition 9.4.1, we do not require the S_k to be bases, but of course if we have a packing we can take a base for each S_k and so obtain a packing employing only bases.

Dually, the right hand equation in (9.5) corresponds to the presence of a covering of C:

Definition 9.4.2. Let $(M_k|k \in K)$ be a family of matroids on the same ground set E. A *covering* for this family consists of an independent set I_k for each M_k such that the I_k cover E.

It is immediate that the sets $I_k \cap C$ form a covering for the family $(M_k \restriction_C |k \in K)$. In fact we get the stronger statement that they form a covering for the family $(M_k.C|k \in K)$ where we contract instead of restricting, since for each k we have that $I_k \cap P$ is an M_k-base for P, and we also have that I_k, which is the union of $I_k \cap C$ with $I_k \cap P$, is M_k-independent.

Putting all of this together, we get the following self-dual notion:

Definition 9.4.3. Let $(M_k|k \in K)$ be a family of matroids on the same ground set E. We say this family has the *Packing/Covering property* iff there is a partition of E into two parts P (called the *packing side*) and C (called the *covering side*) such that $(M_k \restriction_P |k \in K)$ has a packing, and $(M_k.C|k \in K)$ has a covering.

We have established above that this property follows from the rank formula for union, but the argument can easily be reversed to show that in fact Packing/Covering is equivalent to the rank formula, where that formula makes sense. However, Packing/Covering also makes sense for infinite matroids, where the rank formula is no longer useful. We are therefore led to the following conjecture:

Conjecture 9.2.3. *Every family of matroids on the same ground set has the Packing/Covering property.*

Because of this link to the rank formula, we immediately get a special case of this conjecture:

Theorem 9.4.4. *Every finite family of finite matroids on the same ground set has the Packing/Covering property.*

Packing/Covering for pairs of matroids is closely related to another property which is conjectured to hold for all pairs of matroids.

Definition 9.4.5. A pair (M, N) of matroids on the same ground set E has the *Intersection property* iff there is a subset J of E, independent in both matroids, and a partition of J into two parts J^M and J^N such that

$$\mathrm{Cl}_M(J^M) \cup \mathrm{Cl}_N(J^N) = E\,.$$

Conjecture 9.2.1. *Every pair of matroids on the same ground set has the Intersection property.*

We begin by demonstrating a link between Packing/Covering for pairs of matroids and Intersection.

Proposition 9.4.6. *Let M and N be matroids on the same ground set E. Then M and N have the Intersection property iff (M, N^*) has the Packing/Covering property.*

Proof. Suppose first of all that (M, N^*) has the Packing/Covering property, with packing side P decomposed as $S^M \dot\cup S^{N^*}$ and covering side C decomposed as $I^M \dot\cup I^{N^*}$. Let J^M be an M-base of S^M, and J^N an N-base of $C \setminus I^{N^*}$. $J = J^M \cup J^N$ is independent in M since $J^N \subseteq I^M$ is independent in $M.C$ and J^M is independent in $M\restriction_P$. Similarly J is independent in N since $J^M \subseteq P \setminus S^{N^*}$ is independent in $N.P$ and J^N is independent in $N\restriction_C$. But also

$$\mathrm{Cl}_M(J^M) \cup \mathrm{Cl}_N(J^N) = \mathrm{Cl}_M(S^M) \cup \mathrm{Cl}_N(C \setminus I^{N^*}) \supseteq P \cup C = E\,.$$

Now suppose instead that M and N have the Intersection property, as witnessed by $J = J^M \dot\cup J^N$. Let $J^M \subseteq P \subseteq \mathrm{Cl}_M(J^M)$ and $J^N \subseteq C \subseteq \mathrm{Cl}_N(J^N)$ be a partition of E (this is possible since $\mathrm{Cl}_M(J^M) \cup \mathrm{Cl}_N(J^N) = E$). We shall show first of all that $M\restriction_P$ and $N^*\restriction_P$ have a packing, with the spanning sets given by $S^M = J^M$ and $S^{N^*} = P \setminus J^M$. J^M is spanning in $M\restriction_P$ since $P \subseteq \mathrm{Cl}_M(J^M)$, so it is enough to check that $P \setminus J^M$ is spanning in $N^*\restriction_P$, or equivalently that J^M is independent in $N.P$. But this is true since J^N is an N-base of C and $J^M \cup J^N$ is N-independent.

Similarly, J^N is independent in $M.C$, and since $C \subseteq \mathrm{Cl}_N(J^N)$ J^N is spanning in $N\restriction_C$ and so $C \setminus J^N$ is independent in $N^*.C$. Thus the sets $I^M = J^N$ and $I^{N^*} = C \setminus J^N$ form a covering for $(M.C, N^*.C)$. □

Corollary 9.4.7. *If M and N are matroids on the same ground set then (M,N) has the Packing/Covering property iff (M^*, N^*) does.* $\square$

This corollary is not too hard to see directly. However, the following similar corollary is less trivial.

Corollary 9.2.4. *If M and N are matroids on the same ground set then M and N have the Intersection property iff M^* and N^* do.* $\square$

Proposition 9.4.6 shows that Conjecture 9.2.1 follows from Conjecture 9.2.3, but so far we would only be able to use it to deduce that any pair of matroids has the Packing/Covering property from Conjecture 9.2.1. However, this turns out to be enough to give the whole of Conjecture 9.2.3.

Proposition 9.4.8. *Let $(M_k|k \in K)$ be a family of matroids on the same ground set E, and let $M = \bigoplus_{k\in K} M_k$, on the ground set $E \times K$. Let N be the matroid on the same ground set given by $\bigoplus_{e\in E} U^*_{1,K}$. Then the M_k have the Packing/Covering property iff M and N do.*

Proof. First of all, suppose that the M_k have the Packing/Covering property and let P, C, S_k and I_k be as in Definition 9.4.3. We can partition $E \times K$ into $P' = P \times K$ and $C' = C \times K$. Let $S^M = \bigcup_{k\in K} S_k \times \{k\}$, and let $S^N = P' \setminus S^M$. S^M is spanning in $M\restriction_{P'}$ by definition, and since the sets S_k are disjoint, there is for each $e \in P$ at most one $k \in K$ with $(e,k) \notin S^N$. Thus S^N is spanning in $N\restriction_{P'}$. Similarly, let $I^M = \bigcup_{k\in K} I_k \times \{k\}$ and let $I^N = C' \setminus I^M$. I^M is independent in $M.C'$ by definition, and since the sets I_k cover C there is for each $e \in E$ at least one $k \in K$ with $(e,k) \notin I^N$. Thus I^N is independent in $N.C'$.

Now suppose instead that M and N have the Packing/Covering property, with packing side P decomposed as $S^M \dot\cup S^N$ and covering side C decomposed as $I^M \dot\cup I^N$. First we modify these sets a little so that the packing and covering sides are given by $\overline{P} \times K$ and $\overline{C} \times K$ for some sets $\overline{P}$ and $\overline{C}$. To this end, we let $\overline{P} = \{e \in E|(\forall k \in K)(e,k) \in P\}$, and $\overline{C} = \{e \in E|(\exists k \in K)(e,k) \in C\}$, so that $\overline{P}$ and $\overline{C}$ form a partition of E. Let $\overline{S}^N = S^N \cap (\overline{P} \times K)$ and $\overline{I}^N = I^N \cup ((\overline{C} \times K) \setminus C)$. We shall show that $(S^M, \overline{S}^N)$ is a packing for $(M\restriction_{\overline{P}\times K}, N\restriction_{\overline{P}\times K})$ and $(I^M, \overline{I}^N)$ is a covering for $(M.(\overline{C} \times K), N.(\overline{C} \times K))$.

For any $e \in \overline{C}$, the restriction of the corresponding copy of $U^*_{1,K}$ to $P \cap (\{e\} \times K)$ is free, and so since the intersection of S^N with this set is spanning there, it must contain the whole of $P \cap (\{e\} \times K)$. So since $S^M \subseteq P$ is disjoint from S^N, it can't contain any (e,k) with $e \in \overline{C}$. That is, $S^M \subseteq \overline{P} \times K$. It also spans $\overline{P} \times K$ in M, since it spans the larger set P. For each $e \in \overline{P}$, $\overline{S}^N \cap (\{e\} \times K) = S^N \cap (\{e\} \times K)$ N-spans $\{e\} \times K$. Thus $\overline{S}^N$ N-spans $\overline{P} \times K$, so $(S^M, \overline{S}^N)$ is a packing for $(M\restriction_{\overline{P}\times K}, N\restriction_{\overline{P}\times K})$.

To show that $(I^M, \overline{I}^N)$ is a covering for $(M.(\overline{C} \times K), N.(\overline{C} \times K))$, it suffices to show that $\overline{I}^N$ is $N.(\overline{C} \times K)$-independent. For each $e \in \overline{C}$, the set $C \cap (\{e\} \times K)$

is nonempty, so the contraction of the corresponding copy of $U^*_{1,K}$ to this set consists of a single circuit, so there is some point in this set but not in I^N. Then that same point is also not in $\overline{I}^N$, and so $\overline{I}^N \cap (\{e\} \times K)$ is independent in the corresponding copy of $U^*_{1,K}$, so $\overline{I}^N$ is indeed $N.(\overline{C} \times P)$-independent.

Now that we have shown that $\overline{P} \times K$, $\overline{C} \times K$, $(S^M, \overline{S}^N)$ and $(I^M, \overline{I}^N)$ also witness that M and N have the Packing/Covering property, we show how we can construct a packing and a covering for $(M_k \restriction_{\overline{P}} | k \in K)$ and $(M_k.\overline{C} | k \in K)$ respectively.

For each $k \in K$ let $I_k = \{e \in E | (e,k) \in I^M\}$. Since, as we saw above, I^M meets each of the sets $\{e\} \times K$ with $e \in \overline{C}$, the union of the I_k is $\overline{C}$. Since also each I_k is independent in $M_k.\overline{C}$, they form a covering for $(M_k.\overline{C} | k \in K)$. Similarly, let $S_k = \{e \in E | (e,k) \in S^M\}$. Since the intersection of $\overline{S}^N$ with $\{e\} \times K$ is spanning in the corresponding copy of $U^*_{1,k}$ for any $e \in \overline{P}$, it follows that for such e it misses at most one point of this set, so that there can be at most one point in $S^M \cap (\{e\} \times K)$, so the S_k are disjoint. Thus they form a packing of $(M_k \restriction_{\overline{P}} | k \in K)$. □

Corollary 9.4.9. *The following are equivalent:*

(a) Any two matroids have the Intersection property (Conjecture 9.2.1).

(a) Any two matroids in which the second is a direct sum of copies of $U_{1,2}$ have the Intersection property.

(a) Any pair of matroids has the Packing/Covering property.

(a) Any pair of matroids in which the second is a direct sum of copies of $U_{1,2}$ has the Packing/Covering property.

(a) Any family of matroids has the Packing/Covering property (Conjecture 9.2.3).

Proof. We shall prove the following equivalences.

$$\begin{tikzcd}
(b) \arrow[r, leftrightarrow] & (d) \arrow[d, leftrightarrow] & \\
(a) \arrow[r, leftrightarrow] & (c) \arrow[r, leftrightarrow] & (e)
\end{tikzcd}$$

The equivalences of (a) with (c) and (b) with (d) both follow from Proposition 9.4.6. (c) evidently implies (d), but we can also get (c) from (d) by applying Proposition 9.4.8. Similarly, (e) evidently implies (c) and we can get (e) from (c) by applying Proposition 9.4.8. □

9.5 A special case of the Packing/Covering conjecture

In [4], Aharoni and Ziv prove a special case of the intersection conjecture. Here we employ a simplified form of their argument to prove a special case of the

Packing/Covering conjecture. Our simplification also yields a slight strengthening of their theorem.

Key to the argument is the notion of a wave.

Definition 9.5.1. Let $(M_k|k \in K)$ be a family of matroids all on the ground set E. A *wave* for this family is a subset P of E together with a packing $(S_k|k \in K)$ of $(M_k\restriction_P|k \in K)$. In a slight abuse of notation, we shall sometimes refer to the wave just as P or say that elements of P are in the wave. A wave is a *hindrance* if the S_k don't completely cover P. The family is *unhindered* if there is no hindrance, and *loose* if the only wave is the empty wave.

Remark 9.5.2. Those familiar with Aharoni and Ziv's notion of wave should observe that if $(P, (S_1, S_2))$ is a wave as above and we let F be an M_2-base of S_2 then F is not only M_2-independent but also $M_1^*.P$-independent, since $S_1 \subseteq P \setminus F$ is $M_1\restriction_P$-spanning. Now since $P \subseteq \mathrm{Cl}_{M_2}(F)$, we get that F is also $M_1^*.\mathrm{Cl}_{M_2}(F)$-independent. Thus F is a wave in the sense of Aharoni and Ziv for the matroids M_1^* and M_2. There is a similar correspondence of the other notions defined above.

Similarly, they say that the pair (M_1, M_2) is *matchable* iff there is a set which is M_1-spanning and M_2-independent. Those interested in translating between the two contexts should note that there is a covering for (M_1, M_2) iff (M_1^*, M_2) is matchable.

We define a partial order on waves by $(P, (S_k|k \in K)) \leq (P', (S'_k|k \in K))$ iff $P \subseteq P'$ and for each $k \in K$ we have $S_k \subseteq S'_k$. We say a wave is *maximal* iff it is maximal with respect to this partial order.

Lemma 9.5.3. *For any wave P there is a maximal wave $P_{\max} \geq P$.*

Proof. This follows from Zorn's Lemma since for any chain $((P_i, (S^i_k|k \in K))|i \in I)$ the union $(\bigcup_{i\in I} P_i, (\bigcup_{i\in I} S^i_k|k \in K))$ is a wave. □

Lemma 9.5.4. *Let $(M_k|k \in K)$ be a family of matroids on the same ground set E, and let $(P, (S_k|k \in K))$ and $(P', (S'_k|k \in K))$ be two waves. Then $(P \cup P', (S_k \cup (S'_k \setminus P)|k \in K))$ is a wave.*

Proof. Clearly, the $S_k \cup (S'_k \setminus P)$ are disjoint and $cl_{M_k} S_k$ includes $S'_k \cap P$ and hence $cl_{M_k}(S_k \cup (S'_k \setminus P))$ includes $P \cup P'$, as desired. □

Corollary 9.5.5. *If $P_{\max}$ is a maximal wave then anything in any wave P is in $P_{\max}$.*

Proof. We apply Lemma 9.5.4 to the pair $(P_{\max}, P)$. □

Lemma 9.5.6. *For any $e \in E$ and $k \in K$, any maximal wave P satisfies $e \in \mathrm{Cl}_{M_k} P$ whenever there is any wave P' with $e \in \mathrm{Cl}_{M_k} P'$.*

In particular, if e is not contained in any wave, there are at least two k such that, for every wave P', we have $e \notin \mathrm{Cl}_{M_k} P'$.

Proof. Let $(P, (S_k|k \in K))$ be a maximal wave. By Corollary 9.5.5 for any wave $(P', (S'_k|k \in K))$ we have $S'_k \subseteq \mathrm{Cl}_{M_k} S_k$. Thus $e \in \mathrm{Cl}_{M_k} P' = \mathrm{Cl}_{M_k} S'_k$ implies $e \in \mathrm{Cl}_{M_k} P$, as desired.

For the second assertion, assume toward contradiction that there is at most one k_0 such that, for every wave P', $e \notin \mathrm{Cl}_{M_{k_0}} P'$. Then $e \in \mathrm{Cl}_{M_k} P$ for all $k \neq k_0$. But then the following is a wave and contains e:
$X := (P + e, (\overline{S}_k|k \in K))$ where $\overline{S}_{k_0} = S_{k_0} + e$ and $\overline{S}_k = S_k$ for other values of k. This is a contradiction. □

Lemma 9.5.7. *Let $(P, (S_k|k \in K))$ be a wave for a family $(M_k|k \in K)$ of matroids. Let $(P', (S'_k|k \in K))$ be a wave for the family $(M_k/P|k \in K)$. Then $(P \cup P', (S_k \cup S'_k|k \in K))$ is a wave for the family $(M_k|k \in K)$. If either P or P' is a hindrance then so is $P \cup P'$.*

Remark 9.5.8. In fact, though we will not need this, a similar statement can be shown for an ordinal indexed family of waves P^β, with P^β a wave for the family $(M_k/\bigcup_{\gamma<\beta} P^\gamma|k \in K)$.

Proof. For each k, the set S'_k is spanning in $M_k\restriction_{P\cup P'}/P$ and S_k is spanning in $M_k\restriction_{P\cup P'}\restriction P$, so by Lemma 9.3.3 each set $S_k \cup S'_k$ spans $P \cup P'$, and they are clearly disjoint. If the S_k don't cover some point of P then the $S_k \cup S'_k$ also don't cover that point, and the argument in the case where P' is a hindrance is similar. □

Corollary 9.5.9. *For any maximal wave $P_{\max}$, the family $(M_k/P_{\max}|k \in K)$ is loose.*

We are now in a position to present another Conjecture equivalent to the Packing/Covering Conjecture. It is for this new form that we shall present our partial proof.

Conjecture 9.5.10. *Any unhindered family of matroids has a covering.*

Proposition 9.5.11. *Conjecture 9.5.10 and Conjecture 9.2.3 are equivalent.*

Proof. First of all, suppose that Conjecture 9.2.3 holds, and that we have an unhindered family $(M_k|k \in K)$ of matroids. Using Conjecture 9.2.3, we get P, C, S_k and I_k as in Definition 9.4.3. Then $(P, (S_k|k \in K))$ is a wave, and since it can't be a hindrance the sets S_k cover P. They must also all be independent, since otherwise we could remove a point from one of them to obtain a hindrance. So the sets $S_k \cup I_k$ give a covering for $(M_k|k \in K)$.

Now suppose instead that Conjecture 9.5.10 holds, and let $(M_k|k \in K)$ be any family of matroids on the ground set E. Then let $(P, (S_k|k \in K))$ be a maximal wave. By Corollary 9.5.9, $(M_k/P|k \in K)$ is loose, and so in particular this family is unhindered. So it has a covering $(I_k|k \in K)$. Taking covering side $C = E \setminus P$, this means that the M_k have the Packing/Covering property. □

Lemma 9.5.12. *Suppose that we have an unhindered family $(M_k|k \in K)$ of matroids on a ground set E. Let $e \in E$ and $k_0 \in K$ such that for every wave P*

we have $e \notin \mathrm{Cl}_{M_{k_0}} P$. Then the family $(M'_k|k \in K)$ on the ground set $E - e$ is also unhindered, where $M'_{k_0} = M_{k_0}/e$ but $M'_k = M_k \backslash e$ for other values of k.

Proof. Suppose not, for a contradiction, and let $(P, (S_k|k \in K))$ be a hindrance for $(M'_k|k \in K)$. Without loss of generality, we assume that the S_k are bases of P. Let $\overline{S}_k$ be given by $\overline{S}_{k_0} = S_{k_0} + e$ and $\overline{S}_k = S_k$ for other values of k. Note that $\overline{S}_{k_0}$ is independent because otherwise, by the M_{k_0}/e-independence of S_{k_0}, we must have $e \in \mathrm{Cl}_{M_{k_0}}(S_{k_0})$ (in fact, $\{e\}$ must be an M_{k_0}-circuit), so that $P \subseteq \mathrm{Cl}_{M_{k_0}}(S_{k_0})$, and thus $(P, (S_k|k \in K))$ is a wave for the M_k with $e \in \mathrm{Cl}_{M_{k_0}} P$. Let P' be the set of $x \in P$ such that there is no $(\overline{S}_k|k \in K)$-exchange chain from x to e.

Let $x_0 \in P \setminus \bigcup_{k \in K} S_k$. If $x_0 \in P'$, then we will show that $(P', (P' \cap \overline{S}_k|k \in K))$ is a wave containing x_0. This contradicts the assumption that $(M_k|k \in K)$ is unhindered. We must show for every k that every $x \in P' \setminus P' \cap \overline{S}_k$ is M_k-spanned by $P' \cap \overline{S}_k$. Since $e \notin P'$ we cannot have $x = e$. Let C be the unique circuit contained in $x + \overline{S}_k$. If $x \in P'$, then $C \subseteq P'$ by Lemma 9.3.7, so $x \in \mathrm{Cl}_{M_k}(P' \cap \overline{S}_k)$, as desired.

If $x_0 \notin P'$, there is a shortest $(\overline{S}_k|k \in K)$-exchange chain

$$(y_0 = x_0, k_0; y_1, k_1; \ldots; y_n = e)$$

from x_0 to e. Let $\overline{S}'_k := \overline{S}_k \cup \{y_l|k_l = k\} \setminus \{y_{l+1}|k_l = k\}$. By Lemma 9.3.5, $\overline{S}'_k$ is M_k-independent and $\mathrm{Cl}_{M_k} \overline{S_k} = \mathrm{Cl}_{M_k} \overline{S_k}'$ for all $k \in K$. Thus each $\overline{S}'_k$ M_k-spans P but avoids e, in other words: $(P, (\overline{S}'_k|k \in K))$ is an $(M_k|k \in K)$-wave. But also $e \in \mathrm{Cl}_{M_{k_0}} P$ since $e \in \overline{S}_{k_0}$, a contradiction. □

We will now discuss those partial versions of Conjecture 9.5.10 which we can prove. We would like to produce a covering of the ground set by independent sets - and that means that we don't want any of the sets in the covering to include any circuits for the corresponding matroid. First of all, we show that we can at least avoid *some* circuits. In fact, we'll prove a slightly stronger theorem here, showing that we can specify a countable family of sets, which are to be avoided whenever they are dependent. In all our applications, the dependent sets we care about will be circuits.

Theorem 9.5.13. *Let $(M_k|k \in K)$ be an unhindered family of matroids on the same ground set E. Suppose that we have a sequence of subsets o_n of E. Then there is a family $(I_k|k \in K)$ whose union is E and such that for no $k \in K$ and $n \in \mathbb{N}$ do we have both $o_n \subseteq I_k$ and o_n dependent in M_k.*

Proof. If some wave includes the whole ground set, then as the family is unhindered, this wave would yield the desired covering. Unfortunately, we may not assume this. Instead, we recursively build a family $(J_k|k \in K)$ of disjoint sets such that some wave $(P, (S_k|k \in K))$ for the $M_k/J_k \backslash \bigcup_{l \neq k} J_l$ includes enough of $E \setminus \bigcup_k J_k$ that any family $(I_k|k \in K)$ whose union is E and with $I_k \cap (P \cup \bigcup_{k \in K} J_k) = S_k \cup J_k$ will work.

We construct J_k as the nested union of some $(J_k^n|n \in \mathbb{N} \cup \{0\})$ with the following properties. Abbreviate $M_k^n := M_k/J_k^n \backslash \bigcup_{l \neq k} J_l^n$.

(a) J_k^n is independent in M_k.

(a) For different k, the sets J_k^n are disjoint.

(a) $(M_k^n|k \in K)$ is unhindered.

(a) Either the set $o_n \setminus \bigcup_{k \in K} J_k^n$ is included in some $(M_k^n|k \in K)$-wave or there are distinct l, l' such that there is some $e \in o_n \cap J_l^n$ and some $e' \in o_n \cap J_{l'}^n$.

Put $J_k^0 := \emptyset$ for all k. These satisfy ((a))-((a)), and ((a)) is vacuous since there is no term o_0 (we are following the convention that 0 is not a natural number). Assume that we have already constructed J_k^n satisfying ((a))-((a)).

If ((a)) with o_{n+1} in place of o_n is already satisfied by the $(J_k^n|k \in K)$ we can simply take $J_k^{n+1} := J_k^n$ for all k.

Otherwise, if we let P_{max} be a maximal wave, there is some $e \in o_{n+1} \setminus \bigcup_{k \in K} J_k^n$ not in P_{max} and so not in any $(M_k^n|k \in K)$-wave. By Lemma 9.5.6, there are at least two $k \in K$ such that $e \notin \mathrm{Cl}_{M_k^n} P'$ for every wave P'. In particular, e is not a loop ($\{e\}$ is independent) in M_k^n for those two k. Let l be one of these two values of k. Now let $\overline{J_l^{n+1}} := J_l^n + e$ and $\overline{J_k^{n+1}} := J_k^n$ for $k \neq l$. Then the $\overline{J_k^{n+1}}$ satisfy ((a)) and ((a)). By Lemma 9.5.12 and the choice of e, we also have ((a)).

If the $\overline{J_k^{n+1}}$ already satisfy ((a)), then we are done. Else, to obtain ((a)), repeat the induction step so far and find $e' \in o_{n+1} \setminus \bigcup_{k \in K} \overline{J_k^{n+1}}$ not in any $(\overline{M_k^n}|k \in K)$-wave. Here $\overline{M_k^n}$ is M_k^n/e if $k = l$ and $M_k^n \backslash e$ otherwise. Further we find, $l' \neq l$ such that $\{e'\}$ is independent in $\overline{M_{l'}^n}$ and $e' \notin \mathrm{Cl}_{M_l} P'$ for every wave P'. Now let $J_{l'}^{n+1} := \overline{J_{l'}^{n+1}} + e'$ and $J_k^{n+1} := \overline{J_k^{n+1}}$ for $k \neq l'$. Then the J_k^{n+1} satisfy ((a)) and ((a)) and now also ((a)). By Lemma 9.5.12 and the choice of e', we also have ((a)).

We now define a new family of matroids by $M_k' := M_k/J_k \backslash \bigcup_{l \neq k} J_l$, and we construct an $(M_k'|k \in K)$-wave $(P, (S_k|k \in K))$. We once more do this by taking the union of a recursively constructed nested family. Explicitly, we take $S_k = \bigcup_{n \in \mathbb{N}} S_k^n$ and $P = \bigcup_{n \in \mathbb{N}} P^n$, where for each n the wave $W^n = (P^n, (S_k^n|k \in K))$ is a maximal wave for $(M_k^n|k \in K)$ and the S_k^n are nested. We can find such waves using Lemma 9.5.3: for each n we have that W^n is also a wave for $(M_k^{n+1}|k \in K)$ since in our construction we never contract or delete anything which is in a wave.

Now let $(I_k|k \in K)$ be chosen so that $\bigcup I_k = E$ and for each $k_0 \in K$ we have $I_{k_0} \cap (P \cup \bigcup_{k \in K} J_k) = S_{k_0} \cup J_{k_0}$. Suppose for a contradiction that for some pair (k_0, n) we have $o_n \subseteq I_{k_0}$ and o_n is dependent in M_{k_0}. Then by ((a)), either the set $o_n \setminus \bigcup_{k \in K} J_k^n$ is included in some $(M_k^n|k \in K)$-wave or there are distinct l, l' such that there is some $e \in o_n \cap J_l^n$ and some $e' \in o_n \cap J_{l'}^n$. In the second case, clearly $o_n \nsubseteq I_{k_0}$.

In the first case, we will find a hindrance for $(M_k^n|k \in K)$, which contradicts ((a)). It suffices to show that $S_{k_0}^n$ is dependent in $M_{k_0}^n$, since then we can obtain a hindrance by removing a point from $S_{k_0}^n$ in W^n. Let $o = o_n \backslash \bigcup_{k \in K} J_k^n = o_n \backslash J_{k_0}^n$. Note that o is dependent in $M_{k_0}^n$, since o_n is dependent in $M_{k_0}^n$ but $J_{k_0}^n$ is not by ((a)). By assumption, $o \subseteq P^n$, and so since also $o \subseteq o_n \subseteq I_{k_0}$ we have $o \subseteq I_{k_0} \cap P^n = S^n_{k_0}$, so that $S^n_{k_0}$ is $M^n_{k_0}$-dependent as required.

□

Note that, in particular, if we have a countable family of matroids each with only countably many circuits then Theorem 9.5.13 applies in order to prove Conjecture 9.2.3 in that special case. Requiring only countably many circuits might seem quite restrictive, but there are many cases where it holds:

Proposition 9.5.14. *A matroid of any of the following types on a countable ground set has only countably many circuits:*

(a) A finitary matroid.

(a) A matroid whose dual has finite rank.

(a) A direct sum of matroids each with only countably many circuits.

Proof. ((a)) follows from the fact that the countable ground set has only countably many finite subsets. For ((a)), since every base B has finite complement, there are only countably many bases. As every circuit is a fundamental circuit for some base, there can only be countably many circuits, as desired. For ((a)), there can only be countably many nontrivial summands in the direct sum since the ground set is countable, and the result follows. □

In particular, Theorem 9.5.13 applies to any countable family of matroids each of which is a direct sum of matroids that are finitary or whose duals have finite rank. This includes the main result of Aharoni and Ziv in [4], if the ground set E is countable, by Proposition 9.4.6.

If we have a family of sets $(I_k|k \in K)$ which does not form a covering, because some elements aren't independent, how might we tweak it to make them more independent? Suppose that the reason why I_k is dependent is that it contains a circuit o of M_k, but that o also includes a cocircuit for another matroid $M_{k'}$ from our family. Then we could move some point from I_k into $I_{k'}$ to remove this dependence without making $I_{k'}$ any more dependent.[3] We are therefore not so worried about circuits including cocircuits in this way as we are about other sorts of circuits. Therefore we now consider cases where most circuits do include such cocircuits:

Definition 9.5.15. Let $(M_k|k \in K)$ be a family of matroids on the same ground set E. For each $k \in K$ we let W_k be the set of all M_k-circuits that do not contain an $M_{k'}$-cocircuit with $k' \neq k$. Call the family $(M_k|k \in K)$ of matroids *at most countably weird* if $\bigcup W_k$ is at most countable.

Note that if E is countable then $(M_k|k \in K)$ is at most countably weird if and only if $\bigcup W^\infty_k$ is countable where W^∞_k is the subset of W_k consisting only of the infinite circuits in W_k.

[3]We may assume that the I_k are disjoint. Then any new circuits in $I_{k'}$ would have to meet the cocircuit in just one point, which is impossible by Lemma 9.3.4.

Theorem 9.5.16. *Any unhindered and at most countably weird family* $(M_k|k \in K)$ *of matroids has a covering.*

Proof. Apply Theorem 9.5.13 to $(M_k|k \in K)$ where the o_n enumerate $\bigcup W_k$ where the W_k are defined as in Definition 9.5.15.

So far $(I_k|k \in K)$ is not necessarily a covering since each I_k might still contain circuits. But by the choice of the family of circuits each circuit contained in I_k contains an $M_{k'}$-cocircuit with $k' \neq k$.

In the following, we tweak $(I_k|k \in K)$ to obtain a covering $(L_k|k \in K)$. First extend I_k into a minimal M_k-spanning set B_k by $(IM)^*$. We obtain L_k from B_k by removing all elements in $I_k \cap \bigcup_{l \neq k} B_l$. We can suppose without loss of generality $(I_k|k \in K)$ was a partition of E, and so the family $(L_k|k \in K)$ covers E. It remains to show that L_k is independent. For this, assume for a contradiction that L_k contains an M_k-circuit C. By the choice of B_k, the circuit C is contained in I_k. In particular, C contains an M_l-cocircuit X for some $l \neq k$. By construction B_l meets X and thus C. As $C \subseteq I_k$, the circuit C is not contained in L_k, a contradiction. So $(L_k|k \in K)$ is the desired covering. □

Theorem 9.5.17. *Any at most countably weird family* $(M_k|k \in K)$ *of matroids has the Packing/Covering property.*

Proof. For each $k \in K$, let W_k be the set of all M_k-circuits that do not contain an $M_{k'}$-cocircuit with $k' \neq k$. Let $(P, (S_k|k \in K))$ be a maximal wave. We may assume that each S_k is a base of P. It suffices to show that the family $(M_k/P|k \in K)$ has a covering.

By Theorem 9.5.16, it suffices to show that the family $(M_k/P|k \in K)$ is at most countably weird . Let $\overline{W}_k$ be the set of M_k/P-circuits that do not include some $M_{k'}/P$-cocircuit for some $k' \neq k$. By Lemma 9.3.2, for each $o \in \overline{W}_k$, there is an M_k-circuit $\hat{o}$ included in $o \cup S_k$ with $o \subseteq \hat{o}$.

Next we show that if $\hat{o}$ includes some $M_{k'}$-cocircuit b, then $b \subseteq o$. In particular o includes some $M_{k'}/P$-cocircuit. Indeed, otherwise $b \cap P$ is nonempty and includes some $M_{k'}\restriction_P$-cocircuit. This cocircuit would be included in S_k, which is impossible since $S_{k'}$ spans P, and is disjoint from S_k. Thus if $\hat{o}$ is in W_k, then o is in $\overline{W}_k$.

For each $o \in \bigcup \overline{W}_k$, we pick some $k \in K$ such that $o \in \overline{W}_k$, and let $\iota(o) = \hat{o}$. Then $\iota : \bigcup \overline{W}_k \to \bigcup W_k$ is an injection since if $\iota(o) = \iota(q)$, then $o = \iota(o) \setminus P = \iota(q) \setminus q = q$. Thus $(M_k/P|k \in K)$ is at most countably weird and so $(M_k/P|k \in K)$ has a covering by Theorem 9.5.16, which completes the proof. □

However, there are still some important open questions here.

Definition 9.5.18 ([6])**.** The *finitarisation of a matroid* M is the matroid M^{fin} whose circuits are precisely the finite circuits of M.[4] A matroid is called *nearly finitary* if every base misses at most finitely many elements of some base of the finitarisation.

[4] It is easy to check that M^{fin} is indeed a matroid [6].

From Proposition 9.4.6 and the corresponding case of Matroid Intersection [6] we obtain the following:

Corollary 9.5.19. *Any pair of nearly finitary matroids has the Packing/Covering property.*

By Proposition 9.4.8 Corollary 9.5.19 implies that any finite family of nearly finitary matroids has the Packing/Covering property. Since every countable set has only countably many finite subsets, any family of finitary matroids supported on a countable ground set is at most countably weird, and thus has the Packing/Covering property by Theorem 9.5.17. On the other hand any family of two cofinitary matroids has the Packing/Covering by Corollary 9.5.19 since the pairwise Packing/Covering Property is self-dual. By Proposition 9.4.8, this implies that any family of cofinitary matroids has the Packing/Covering property. We sum up these results in the following table.

Type of family	cofinitary	finitary	nearly finitary
finite	✓	✓	✓
countable ground set	✓	✓	?
arbitrary	✓	?	?

In particular, we do not know the answer to the following open questions.

Open Question 9.5.20. *Must every family of nearly finitary matroids on a countable common ground set have the Packing/Covering property?*

Open Question 9.5.21. *Must every family of finitary matroids have the Packing/Covering property?*

9.6 Base covering

The well-known base covering theorem reads as follows.

Theorem 9.6.1. *Any family of finite matroids $(M_k|k \in K)$ on a finite common ground set E has a covering if and only if for every finite set $X \subseteq E$ the following holds.*

$$\sum_{k \in K} r_{M_k}(X) \geq |X|$$

Taking the family to contain only one matroid, consisting of one infinite circuit, we see that this theorem does not extend verbatim to infinite matroids. However, Theorem 9.6.1 extends verbatim to finite families of finitary matroids by compactness [5].[5] The requirement that the family is finite is necessary as $(U_k = U_{1,\mathbb{R}}|k \in \mathbb{N})$ satisfies the rank formula but does not have a covering.

In the following, we conjecture an extension of the finite base covering theorem to arbitrary infinite matroids. Our approach is to replace the rank formula

[5]The argument in [5] is only made in the case where all M_k are the same but it easily extends to finite families of arbitrary finitary matroids.

by a condition that for finite sets X is implied by the rank formula but is still meaningful for infinite sets. A first attempt might be the following:

$$\text{Any packing for the family } (M_k\restriction_X | k \in K) \text{ is already a covering.} \tag{9.6}$$

Indeed, for finite X, if $(M_k\restriction_X | k \in K)$ has a packing and there is an element of X not covered by the spanning sets of this packing, then this violates the rank formula. However, there are infinite matroids that violate (9.6) and still have a covering, see Figure 9.2.

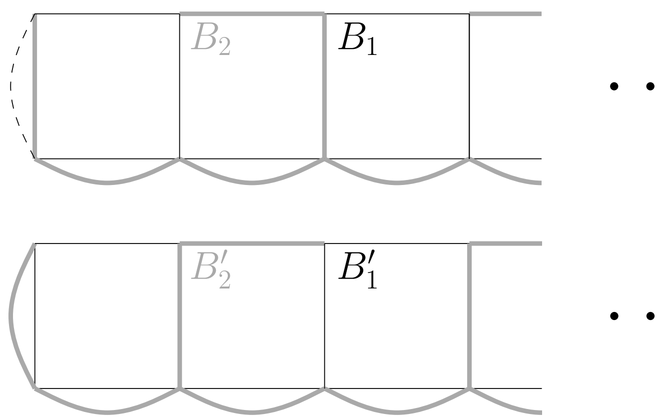

Figure 9.2: Above is a base packing which isn't a base covering. Below that is a base covering for the same matroids, namely the finite cycle matroid for the graph, taken twice.

We propose to use instead the following weakening of (9.6).

$$\text{If } (M_k\restriction_X | k \in K) \text{ has a packing, then it also has a covering.} \tag{9.7}$$

To see that (9.7) does not imply the rank formula for some finite X, consider the family (M, M), where M is the finite cycle matroid of the graph

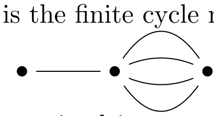

This graph has an edge not contained in any cycle (so that (M, M) does not have a packing) but enough parallel edges to make the rank formula false.

Using (9.7), we obtain the following:

Conjecture 9.6.2 (Covering Conjecture)**.** *A family of matroids $(M_k | k \in K)$ on the same ground set E has a covering if and only if (9.7) is true for every $X \subseteq E$.*

Proposition 9.6.3. *Conjecture 9.2.3 and Conjecture 9.6.2 are equivalent.*

Proof. For the "only if" direction, note that Conjecture 9.6.2 implies Conjecture 9.5.10, which by Proposition 9.5.11 implies Conjecture 9.2.3.

For the "if" direction, note that by assumption we have a partition $E = P\dot{\cup}C$ such that there exist disjoint $M_k\restriction_P$-spanning sets S_k and $M_k.C$-independent sets I_k whose union is C. By (9.7), $(M_k\restriction_P|k \in K)$ has a covering with sets B_k, where $B_k \in \mathcal{I}(M_k\restriction_P)$. As $I_k \cup B_k \in \mathcal{I}(M_k)$, the sets $I_k \cup B_k$ form the desired covering. □

As Packing/Covering is true for finite matroids, Proposition 9.6.3 implies the non-trivial direction of Theorem 9.6.1. By Theorem 9.5.17 we obtain the following applications.

Corollary 9.6.4. *Any at most countably weird family of matroids $(M_k|k \in K)$ has a covering if and only if (9.7) is true for every $X \subseteq E$.*

Let us now specialise to graphs. A good introduction to the algebraic and the topological cycle matroids of infinite graphs is [21]. We rely on the fact that the algebraic cycle matroid of any locally finite graph and the topological cycle matroid of any graph are co-finitary.

Definition 9.6.5. The bases of the topological cycle matroid are called *topological trees* and the bases of the algebraic cycle matroid are called *algebraic trees.* Using this we define *topological tree-packing, topological tree-covering, algebraic tree-packing, algebraic tree-covering.*

Corollary 9.6.6 (Base covering for the topological cycle matroids)**.** *A family of multigraphs $(G_k|k \in K)$ on a common ground set E has a topological tree-covering if and only if the following is true for every $X \subseteq E$.*

If $(G_k[X]|k \in K)$ has a topological tree-packing, then it also has a topological tree-covering. (9.8)

Corollary 9.6.7 (Base covering for the algebraic cycle matroids of locally finite graphs)**.** *A family of locally finite multigraphs $(G_k|k \in K)$ on a common ground set E has an algebraic tree-covering if and only if the following is true for every $X \subseteq E$.*

If $(G_k[X]|k \in K)$ has an algebraic tree-packing, then it also has an algebraic tree-covering. (9.9)

9.7 Base packing

The well-known base packing theorem reads as follows.

Theorem 9.7.1. *Any family of finite matroids $(M_k|k \in K)$ on a finite common ground set E has a packing if and only if for every finite set $Y \subseteq E$ the following holds.*

$$\sum_{k\in K} r_{M_k.Y}(Y) \leq |Y|$$

Aigner-Horey, Carmesin and Fröhlich [5] extended this theorem to families consisting of finitely many copies of the same co-finitary matroid. We extend this to arbitrary co-finitary families.

Theorem 9.7.2. *Any family of co-finitary matroids $(M_k|k \in K)$ on a common ground set E has a packing if and only if for every finite set $Y \subseteq E$ the following holds.*

$$\sum_{k\in K} r_{M_k.Y}(Y) \leq |Y|$$

Proof by a compactness argument. We will think of partitions of the ground set E as functions from E to K - such a function f corresponds to a partition $(S_k^f|k \in K)$, given by $S_k^f = \{e \in E|f(e) = k\}$. Endow K with the co-finite topology where a set is closed iff it is finite or the whole of K. Then endow K^E with the product topology, which is compact since the topology on K is compact.

By Lemma 9.3.1 a set S is spanning for a matroid M iff it meets every cocircuit of that matroid. So we would like a function f contained in each of the sets $C_{k,B} = \{f|S_k^f \cap B \neq \emptyset\}$, where B is a cocircuit for the matroid M_k. We will prove the existence of such a function by a compactness argument: we need to show that each $C_{k,B}$ is closed in the topology given above and that any finite intersection of them is nonempty.

To show that $C_{k,B}$ is closed, we rewrite it as $\bigcup_{e\in B}\{f|f(e) = k\}$. Each of the sets $\{f|f(e) = k\}$ is closed since their complements are basic open sets, and the union is finite since M_k is co-finitary.

Now let $(k_i|1 \leq i \leq n)$ and $(B_i|1 \leq i \leq n)$ be finite families with each B_i a cocircuit in M_{k_i}. We need to show that $\bigcap_{1\leq i\leq n} C_{k_i,B_i}$ is nonempty. Let $X = \bigcup_{1\leq i\leq n} B_i$. Since the rank formula holds for each subset of X, we have by the finite version of the base packing Theorem a packing $(S_k|k \in K)$ of $(M_k.X|k \in K)$. Now any f such that $f(e) = k$ for $e \in S_k$ will be in $\bigcap_{1\leq i\leq n} C_{k_i,B_i}$ by Lemma 9.3.1, since each B_i is an $M_{k_i}.X$-cocircuit. This completes the proof.

□

Theorem 9.7.1 does not extend verbatim to arbitrary infinite matroids. Indeed, for every integer k there exists a finitary matroid M on a ground set E with no three disjoint bases yet satisfying $|Y| \geq kr_{M.Y}(Y)$ for every finite $Y \subseteq E$ [3, 36].

In the following we conjecture an extension of the finite base packing theorem to arbitrary infinite matroids. This extension uses the following condition, which for finite sets Y is implied by the rank formula of the base packing theorem but is still meaningful for infinite sets:

$$\text{If } (M_k.Y|k \in K) \text{ has a covering, then it also has a packing.} \tag{9.10}$$

Indeed, if $(M_k.Y|k \in K)$ has a covering and there is an element of Y contained in several of the corresponding independent sets, then this violates the rank formula.

Using our new condition, we obtain the following:

Conjecture 9.7.3 (Packing Conjecture)**.** *A family of matroids $(M_k|k \in K)$ on the same ground set E has a packing if and only if (9.10) is true for every $Y \subseteq E$.*

Proposition 9.7.4. *Conjecture 9.2.3 and Conjecture 9.7.3 are equivalent.*

Proof. Since by Lemma 9.3.1 condition (9.10) for a pair of matroids is equivalent to (9.7) for the duals of those matroids and a pair of matroids have a packing if and only if their duals have a covering, Conjecture 9.7.3 implies that any pair of matroids satisfying (9.7) has a covering, and in particular that any unhindered pair of matroids has a covering. As in the proof of (9.5.11), this implies that any pair of matroids has the Packing/Covering property, which implies Conjecture 9.2.3 by Corollary 9.4.9.

The converse is proved as in the proof of Proposition 9.6.3. □

As Packing/Covering is true for finite matroids, Proposition 9.7.4 implies the non-trivial direction of Theorem 9.7.1. By Theorem 9.5.17 we obtain the following applications.

Corollary 9.7.5. *Any at most countably weird family of matroids on ground set E has a packing if and only if (9.10) is true for every $Y \subseteq E$.*

Now let us specialise to graphs. The question if there is a packing theorem for the finite cycle matroid of an infinite graph was raised by Nash-Williams in 1967 [74] , who suggested that a countable graph G has k edge-disjoint spanning trees if and if $k \cdot r_{M.Y}(Y) \leq |Y|$ for every finite edge set Y. Here M is the finite cycle matroid of G. However, Aharoni and Thomassen constructed a counterexample in 1989 [3]. Our approach gives the following two packing theorems for finite cycle matroids of infinite graphs. We rely on the fact that the finite cycle matroid of any graph is finitary.

Corollary 9.7.6 (Base packing theorem for the finite cycle matroid)**.** *Any family of countable multigraphs $(G_k|k \in K)$ with a common edge set E has a tree-packing if and only if (9.11) is true for every $Y \subseteq E$.*

If $(M_k.Y|k \in K)$ has a tree-covering, then it also has a tree-packing. (9.11)

Corollary 9.7.7 (Base packing theorem for the finite cycle matroid)**.** *Any finite family of multigraphs $(G_k|k \in K)$ with common edge set E has a tree-packing if and only if (9.11) is true for every $Y \subseteq E$.*

A similar result was obtained by Aharoni and Ziv [4]. However, their argument is different and they have the additional assumption that the ground set is countable.

Note that the covering conjecture for arbitrary finitary families is still open and equivalent to Open Question 9.5.21.

9.8 Overview

We have shown that a great many natural conjectures are equivalent, which we will review in this section. We are indebted to a reviewer for pointing out the importance of the fact that many of the equivalences we have proved specialise to smaller classes than the class of all matroids. We therefore consider the following conjectures, each of which could be made relative to a class $\mathcal{M}$ of matroids.

The Intersection conjecture: Any two matroids in $\mathcal{M}$ on the same ground set have the Intersection property

The pairwise Packing/Covering conjecture: Any pair of matroids from $\mathcal{M}$ on the same ground set has the Packing/Covering property

The Packing/Covering conjecture: Any family of matroids from $\mathcal{M}$ on the same ground set has the Packing/Covering property

The Packing conjecture: A family of matroids $(M_k \in \mathcal{M} | k \in K)$ on the same ground set E has a packing if and only if the following condition is true for every $Y \subseteq E$:

If $(M_k.Y | k \in K)$ has a covering, then it also has a packing.

The Covering conjecture: A family of matroids $(M_k \in \mathcal{M} | k \in K)$ on the same ground set E has a covering if and only if the following condition is true for every $Y \subseteq E$:

If $(M_k\restriction_Y | k \in K)$ has a packing, then it also has a covering.

Most crudely, if $\mathcal{M}$ is a class of matroids containing all matroids $U^*_{1,K}$ and closed under duality, minors and direct sums then all of the above conjectures are equivalent to each other, with proofs exactly as in this chapter. However, particular equivalences only depend on weaker conditions on the class $\mathcal{M}$. For the equivalence of the Intersection conjecture with the pairwise Packing/Covering conjecture, both relative to $\mathcal{M}$, we just need that $\mathcal{M}$ is closed under duality. For the equivalence of the pairwise Packing/Covering conjecture with the Packing/Covering conjecture, we just need that $\mathcal{M}$ contains all the matroids $U^*_{1,K}$ and is closed under direct sums. This equivalence also holds for classes of matroids of bounded size:

Lemma 9.8.1. *Let $\mathcal{M}_{<\kappa}$ be the class of all matroids on ground sets of cardinality less than κ for some regular*[6] *cardinal κ. Then the pairwise Packing/Covering conjecture for $\mathcal{M}_\kappa$ is equivalent to the Packing/Covering conjecture for $\mathcal{M}_\kappa$.*

[6]Recall that an infinite cardinal κ is *regular* if and only if no set of cardinality κ can be expressed as a union of fewer than κ sets, all of cardinality less than κ.

Proof (assuming the axiom of choice). It is clear that the pairwise Packing/Covering conjecture follows from the Packing/Covering conjecture. For the converse, suppose the pairwise Packing/Covering conjecture holds, and let $(M_k|k \in K)$ be a family of matroids on the same ground set E of cardinality less than κ. For each $e \in E$, let K_e be the set of $k \in K$ for which $\{e\}$ is independent in M_k. Let $E' = \{e \in E|\#(K_e) < \kappa\}$, and let $K' = \bigcup_{e \in E'} K_e$. Then K' has cardinality less than κ, so by Proposition 9.4.8 the family $(M_k\restriction_{E'}|k \in K')$ has the Packing/Covering property: call the packing side P and the covering side C, and let the packing and the covering be $(I_k|k \in K')$ and $(S_k|k \in K')$.

Let $C' = E \setminus P$, and for any $k \in K \setminus K'$ let $S_k = \emptyset$, which is spanning in $M_k\restriction E'$ by the definition of K'. Using some well-ordering of $E \setminus E'$, we can choose recursively for each $e \in E \setminus E'$ an element $k(e)$ of K_e such that all of the $k(e)$ are distinct. For each $k \in K \setminus K'$, we now set $I_k = \{e \in E \setminus E'|k(e) = k\}$, which is either empty or has size 1 and is independent in M_k. Then the S_k form a packing of P and the I_k form a covering of C', so $(M_k|k \in K)$ has the Packing/Covering property. □

For the equivalence of the Packing/Covering conjecture with the Covering conjecture, both relative to $\mathcal{M}$, we just need that $\mathcal{M}$ is closed under contraction. For the equivalence of the Packing/Covering conjecture with the Packing conjecture, both relative to $\mathcal{M}$, we just need that $\mathcal{M}$ is closed under deletion. To see this, it is not enough to use the argument in the proof of Proposition 9.7.4, for that argument goes via the pairwise Packing/Covering conjecture. Instead, an argument dual to that for the Covering conjecture must be used, relying on the existence of maximal cowaves, where a cowave is a pair $(C, (I_k|k \in K))$ with the I_k forming a covering of $(M_k.C|k \in K)$. The existence of maximal cowaves can be demonstrated by an argument dual to that for Lemma 9.5.3.

Chapter 10

On the intersection of infinite matroids

10.1 Abstract

We show that the *infinite matroid intersection conjecture* of Nash-Williams implies the infinite Menger theorem proved by Aharoni and Berger in 2009.

We prove that this conjecture is true whenever one matroid is nearly finitary and the second is the dual of a nearly finitary matroid, where the nearly finitary matroids form a superclass of the finitary matroids.

In particular, this proves the infinite matroid intersection conjecture for finite-cycle matroids of 2-connected, locally finite graphs with only a finite number of vertex-disjoint rays.

10.2 Introduction

The infinite Menger theorem[1] was conjectured by Erdős in the 1960s and proved recently by Aharoni and Berger [2]. It states that for any two sets of vertices S and T in a connected graph, there is a set of vertex-disjoint S-T-paths whose maximality is witnessed by an S-T-separator picking exactly one vertex form each of these paths.

The complexity of the only known proof of this theorem and the fact that the finite Menger theorem has a short matroidal proof, make it natural to ask whether a matroidal proof of the infinite Menger theorem exists. In this chapter, we propose a way to approach this problem by proving that a conjecture of Nash-Williams regarding infinite matroids implies the infinite Menger theorem.

Building on earlier work of Higgs and Oxley, recently, Bruhn, Diestel, Kriesell, Pendavingh and Wollan [22] found axioms for infinite matroids in terms of independent sets, bases, circuits, closure and (relative) rank. These axioms allow for

[1]see Theorem 10.4.1 below.

duality of infinite matroids as known from finite matroid theory, which settled an old problem of Rado. With these new axioms it is now possible to study which theorems of finite matroid theory have infinite analogues.

Here, we shall look at Edmonds' *matroid intersection theorem*, which is a classical result in finite matroid theory [75]. It asserts that *the maximum size of a common independent set of two matroids M_1 and M_2 on a common ground set E is given by*

$$\min_{X\subseteq E} \mathrm{rk}_{M_1}(X) + \mathrm{rk}_{M_2}(E \setminus X), \tag{10.1}$$

where rk_{M_i} denotes the rank function of the matroid M_i.

In this chapter, we consider the following conjecture of Nash-Williams, which first appeared in [4][2] and serves as an infinite analogue to the finite matroid intersection theorem[3].

Conjecture 10.2.1. [The infinite matroid intersection conjecture]
Any two matroids M_1 and M_2 on a common ground set E have a common independent set I admitting a partition $I = J_1 \cup J_2$ such that $\mathrm{cl}_{M_1}(J_1)\cup \mathrm{cl}_{M_2}(J_2) = E$.

Here, $\mathrm{cl}_M(X)$ denotes the *closure* of a set X in a matroid M; it consists of X and the elements spanned by X in M (see [75]).

10.2.1 Our results

Aharoni and Ziv [4] proved that Conjecture 10.2.1 implies the infinite analogues of Knig's and Hall's theorems. We strengthen this by showing that this conjecture implies the celebrated *infinite Menger theorem* (in the undirected version as stated in Theorem 10.4.1 below), which is known to imply the infinite analogues of Knig's and Hall's theorems [36].

Theorem 10.2.2. *The infinite matroid intersection conjecture for finitary matroids implies the infinite Menger theorem.*

We are able to prove new instances of Conjecture 10.2.1.[4], see Theorem 10.2.5 below. Before we can state this theorem, we need to introduce the class of 'nearly finitary matroids'. For any matroid M, taking as circuits only the finite circuits of M defines a (finitary) matroid with the same ground set as M. This matroid is called the *finitarization* of M and denoted by M^{fin}.

It is not hard to show that every basis B of M extends to a basis B^{fin} of M^{fin}, and conversely every basis B^{fin} of M^{fin} contains a basis B of M. Whether or not $B^{\mathrm{fin}} \setminus B$ is finite will in general depend on the choices for B and B^{fin}, but given a choice for one of the two, it will no longer depend on the choice for the second one.

[2] Historical note: in [4], Nash-Williams's Conjecture is only made for *finitary matroids*, those all of whose circuits are finite.

[3] An alternative notion of infinite matroid intersection was recently proposed by Christian [32].

[4] The methods of this chapter are refined in [16], which was submitted to the arxiv half a year after this chapter.

We call a matroid M *nearly finitary* if every base of its finitarization contains a base of M such that their difference is finite.

Next, let us look at some examples of nearly finitary matroids. There are three natural extensions to the notion of a finite graphic matroid in an infinite context [22]; each with ground set $E(G)$. The most studied one is the *finite-cycle matroid*, denoted $M_{FC}(G)$, whose circuits are the finite cycles of G. This is a finitary matroid, and hence is also nearly finitary.

The second extension is the *algebraic-cycle matroid*, denoted $M_A(G)$, whose circuits are the finite cycles and double rays of G [22, 21][5].

Proposition 10.2.3. *$M_A(G)$ is a nearly finitary matroid if and only if G has only a finite number of vertex-disjoint rays.*

The third extension is the *topological-cycle matroid*, denoted $M_C(G)$[6], whose circuits are the topological cycles of G (Thus $M_C^{\text{fin}}(G) = M_{FC}(G)$ for any finitely separable graph G; see Subsection 10.7.2 or [21] for definitions).

Proposition 10.2.4. *Suppose that G is 2-connected and locally finite. Then, $M_C(G)$ is a nearly finitary matroid if and only if G has only a finite number of vertex-disjoint rays.*

Here we prove the following.

Theorem 10.2.5. *Conjecture 10.2.1 holds for M_1 and M_2 whenever M_1 is nearly finitary and M_2 is the dual of a nearly finitary matroid.*

Aharoni and Ziv [4] proved that the infinite matroid intersection conjecture is true whenever one matroid is finitary and the other is a countable direct sum of finite-rank matroids. Note that Theorem 10.2.5 does not imply this result of [4] nor is it implied by it.

Proposition 10.2.4 and Theorem 10.2.5 can be used to prove the following.

Corollary 10.2.6. *Suppose that G and H are 2-connected, locally finite graphs with only a finite number of vertex-disjoint rays. Then their finite-cycle matroids $M_{FC}(G)$ and $M_{FC}(H)$ satisfy the intersection conjecture.*

Similar results are true for *the algebraic-cycle matroid*, *the topological-cycle matroid*, and their duals.

10.2.2 An overview of the proof of Theorem 10.2.5

In finite matroid theory, an exceptionally short proof of the matroid intersection theorem employing the well-known *finite matroid union theorem* [75, 86] is

[5] $M_A(G)$ is not necessarily a matroid for any G; see [58].
[6] $M_C(G)$ is a matroid for any G; see [21].

known. The latter theorem asserts[7] that for two finite matroids $M_1 = (E_1, \mathcal{I}_1)$ and $M_2 = (E_2, \mathcal{I}_2)$ the set system

$$\mathcal{I}(M_1 \vee M_2) = \{I_1 \cup I_2 \mid I_1 \in \mathcal{I}_1,\ I_2 \in \mathcal{I}_2\} \tag{10.2}$$

forms the set of independent sets of their *union matroid* $M_1 \vee M_2$. Throughout, M^* denotes the dual of a matroid M. We prove that this strategy of proof extends to infinite matroids.

Theorem 10.2.7. *If M_1 and M_2 are matroids on a common ground set E and $M_1 \vee M_2^*$ is a matroid, then Conjecture 10.2.1 holds for M_1 and M_2.*

Thus in order to prove Conjecture 10.2.1, it would be enough to prove that the union of any two matroids is a matroid. Unfortunately, this is not true.[8] We provide examples in Section 10.8. However, we can prove that the union of two nearly finitary matroids is a matroid.

Theorem 10.2.8. *If M_1 and M_2 are nearly finitary matroids, then $M_1 \vee M_2$ is a nearly finitary matroid.*

Hence Theorem 10.2.5 follows from combining Theorem 10.2.8 and Theorem 10.2.7.

This chapter is organized as follows. Additional notation, terminology, and basic lemmas are given in Section 10.3. In Section 10.4 we prove Theorem 10.2.2. In Section 10.5 we prove Theorem 10.2.8. In Section 10.6 we prove Theorem 10.2.7, and in Section 10.7 we prove Propositions 10.2.3 and 10.2.4 and Corollary 10.2.6. In Section 10.8, we construct matroids whose union is not a matroid.

10.3 Preliminaries

Notation and terminology for graphs are that of [36], and for matroids that of [22, 75].

Throughout, G always denotes a graph where $V(G)$ and $E(G)$ denote its vertex and edge sets, respectively. We write M to denote a matroid and write $E(M)$, $\mathcal{I}(M)$, $\mathcal{B}(M)$, and $\mathcal{C}(M)$ to denote its ground set, independent sets, bases, and circuits, respectively.

It will be convenient to have a similar notation for set systems. That is, for a set system $\mathcal{I}$ over some ground set E, an element of $\mathcal{I}$ is called *independent*, a maximal element of $\mathcal{I}$ is called a *base* of $\mathcal{I}$, and a minimal element of $\mathcal{P}(E) \setminus \mathcal{I}$ is called *circuit* of $\mathcal{I}$. A set system is *finitary* if an infinite set belongs to the

[7] Often the matroid union theorem is complemented by a formula for the rank function of the union. This, however, is implied by the fact that the union is a matroid (as follows from Theorem 10.2.7 below and results of [16]). This rank formula and its relation to Conjecture 10.2.1 is studied in [16].

[8] This is not that surprsing as the methods of this chapter are much more elementary than those developed by Aharoni and Berger in [2].

system provided each of its finite subsets does; with this terminology, M is finitary provided that $\mathcal{I}(M)$ is finitary.

We review the definition of a matroid as this is given in [22]. A set system $\mathcal{I}$ is the set of independent sets of a matroid if it satisfies the following *independence axioms*:

(I1) $\emptyset \in \mathcal{I}$.

(I2) $\lceil \mathcal{I} \rceil = \mathcal{I}$, that is, $\mathcal{I}$ is closed under taking subsets.

(I3) Whenever $I, I' \in \mathcal{I}$ with I' maximal and I not maximal, there exists an $x \in I' \setminus I$ such that $I + x \in \mathcal{I}$.

(IM) Whenever $I \subseteq X \subseteq E$ and $I \in \mathcal{I}$, the set $\{I' \in \mathcal{I} \mid I \subseteq I' \subseteq X\}$ has a maximal element.

In [22], an equivalent axiom system to the independence axioms is provided and is called the *circuit axioms system*; this axiom system characterises a matroid in terms of its circuits. Of these circuit axioms, we shall make frequent use of the so called *(infinite) circuit elimination axiom* phrased here for a matroid M:

(C) Whenever $X \subseteq C \in \mathcal{C}(M)$ and $\{C_x \mid x \in X\} \subseteq \mathcal{C}(M)$ satisfies $x \in C_y \Leftrightarrow x = y$ for all $x, y \in X$, then for every $z \in C \setminus \left(\bigcup_{x \in X} C_x\right)$ there exists a $C' \in \mathcal{C}(M)$ such that $z \in C' \subseteq \left(C \cup \bigcup_{x \in X} C_x\right) \setminus X$.

The following is a well-known fact for finite matroids (see, e.g., [75]), which generalizes easily to infinite matroids.

Lemma 10.3.1. [22, Lemma 3.11]
Let M be a matroid. Then, $|C \cap C^*| \neq 1$, whenever $C \in \mathcal{C}(M)$ and $C^* \in \mathcal{C}(M^*)$.

10.4 From infinite matroid intersection to the infinite Menger theorem

In this section, we prove Theorem 10.2.2; asserting that the infinite matroid intersection conjecture implies the infinite Menger theorem.

Given a graph G and $S, T \subseteq V(G)$, a set $X \subseteq V(G)$ is called an S–T *separator* if $G - X$ contains no S–T path. The infinite Menger theorem reads as follows.

Theorem 10.4.1 (Aharoni and Berger [2])**.** *Let G be a connected graph. Then for any $S, T \subseteq V(G)$ there is a set $\mathcal{L}$ of vertex-disjoint S–T paths and an S–T separator $X \subseteq \bigcup_{P \in \mathcal{L}} V(P)$ satisfying $|X \cap V(P)| = 1$ for each $P \in \mathcal{L}$.*

Infinite matroid union cannot be used in order to obtain the infinite Menger Theorem directly via Theorem 10.2.7 and Theorem 10.2.2. Indeed, in Section 10.8 we construct a finitary matroid M and a co-finitary matroid N such that their union is not a matroid. Consequently, one cannot apply Theorem 10.2.7 to the finitary matroids M and N^* in order to obtain Conjecture 10.2.1 for them. However, it is easy to see that Conjecture 10.2.1 is true for these particular M and N^*.

Next, we prove Theorem 10.2.2.

Proof of Theorem 10.2.2. Let G be a connected graph and let $S, T \subseteq V(G)$ be as in Theorem 10.4.1. We may assume that $G[S]$ and $G[T]$ are both connected. Indeed, an S–T separator with $G[S]$ and $G[T]$ connected gives rise to an S–T separator when these are not necessarily connected. Abbreviate $E(S) := E(G[S])$ and $E(T) := E(G[T])$, let M be the finite-cycle matroid $M_F(G)$, and put $M_S := M/E(S) - E(T)$ and $M_T := M/E(T) - E(S)$; all three matroids are clearly finitary.

Assuming that the infinite matroid intersection conjecture holds for M_S and M_T, there exists a set $I \in \mathcal{I}(M_S) \cap \mathcal{I}(M_T)$ which admits a partition $I = J_S \cup J_T$ satisfying

$$\mathrm{cl}_{M_S}(J_S) \cup \mathrm{cl}_{M_T}(J_T) = E,$$

where $E = E(M_S) = E(M_T)$. We regard I as a subset of $E(G)$.

For the components of $G[I]$ we observe two useful properties. As I is disjoint from $E(S)$ and $E(T)$, the edges of a cycle in a component of $G[I]$ form a circuit in both, M_S and M_T, contradicting the independence of I in either. Consequently,

$$\text{the components of } G[I] \text{ are trees.} \tag{10.3}$$

Next, an S-path[9] or a T-path in a component of $G[I]$ gives rise to a circuit of M_S or M_T in I, respectively. Hence,

$$|V(C) \cap S| \leq 1 \text{ and } |V(C) \cap T| \leq 1 \text{ for each component } C \text{ of } G[I]. \tag{10.4}$$

Let $\mathcal{C}$ denote the components of $G[I]$ meeting both of S and T. Then by (10.3) and (10.4) each member of $\mathcal{C}$ contains a unique S–T path and we denote the set of all these paths by $\mathcal{L}$. Clearly, the paths in $\mathcal{L}$ are vertex-disjoint.

In what follows, we find a set X comprised of one vertex from each $P \in \mathcal{L}$ to serve as the required S–T separator. To that end, we show that one may alter the partition $I = J_S \cup J_T$ to yield a partition

$$I = K_S \cup K_T \text{ satisfying } cl_{M_S}(K_S) \cup cl_{M_T}(K_T) = E \text{ and (Y.1-4)}, \tag{10.5}$$

where (Y.1-4) are as follows.

(Y.1) Each component C of $G[I]$ contains a vertex of $S \cup T$.

(Y.1) Each component C of $G[I]$ meeting S but not T satisfies $E(C) \subseteq K_S$.

(Y.1) Each component C of $G[I]$ meeting T but not S satisfies $E(C) \subseteq K_T$.

[9]A non-trivial path meeting $G[S]$ exactly in its end vertices.

(Y.1) Each component C of $G[I]$ meeting both, S and T, contains at most one vertex which at the same time

(a) lies in S or is incident with an edge of K_S, and

(a) lies in T or is incident with an edge of K_T.

Postponing the proof of (10.5), we first show how to deduce the existence of the required S–T separator from (10.5). Define a pair of sets of vertices (V_S, V_T) of $V(G)$ by letting V_S consist of those vertices contained in S or incident with an edge of K_S and defining V_T in a similar manner. Then $V_S \cap V_T$ may serve as the required S–T separator. To see this, we verify below that (V_S, V_T) satisfies all of the terms (Z.1-4) stated next.

(Z.1) $V_S \cup V_T = V(G)$;

(Z.1) for every edge e of G either $e \subseteq V_S$ or $e \subseteq V_T$;

(Z.1) every vertex in $V_S \cap V_T$ lies on a path from $\mathcal{L}$; and

(Z.1) every member of $\mathcal{L}$ meets $V_S \cap V_T$ at most once.

To see (Z.(Z.1)), suppose v is a vertex not in $S \cup T$. As G is connected, such a vertex is incident with some edge $e \notin E(T) \cup E(S)$. The edge e is spanned by K_T or K_S; say K_T. Thus, $K_T + e$ contains a circle containing e or $G[K_T + e]$ has a T-path containing e. In either case v is incident with an edge of K_T and thus in V_T, as desired.

To see (Z.(Z.1)), let $e \in cl_{M_T}(K_T) \backslash K_T$; so that $K_T + e$ has a circle containing e or $G[K_T + e]$ has T-path containing e; in either case both end vertices of e are in V_T, as desired. The treatment of the case $e \in cl_{M_S}(K_S)$ is similar.

To see (Z.(Z.1)), let $v \in V_S \cap V_T$; such is in S or is incident with an edge of K_S, and in T or is incident with an edge in K_T. Let C be the component of $G[I]$ containing v. By (Y.1-4), $C \in \mathcal{C}$, i.e. it meets both, S and T and therefore contains an S–T path $P \in \mathcal{L}$. Recall that every edge of C is either in K_S or K_T and consider the last vertex w of a maximal initial segment of P in $C - K_T$. Then w satisfies (Y.(a)), as well as (Y.(a)), implying $v = w$; so that v lies on a path from $\mathcal{L}$.

To see (Z.(Z.1)), we restate (Y.(Y.1)) in terms of V_S and V_T: each component of $\mathcal{C}$ contains at most one vertex of $V_S \cap V_T$. This clearly also holds for the path from $\mathcal{L}$ which is contained in C.

It remains to prove (10.5). To this end, we show that any component C of $G[I]$ contains a vertex of $S \cup T$. Suppose not. Let e be the first edge on a $V(C)$–S path Q which exists by the connectedness of G. Then $e \notin I$ but without loss of generality we may assume that $e \in \mathrm{cl}_{M_S}(J_S)$. So in $G[I] + e$ there must be a cycle or an S-path. The latter implies that C contains a vertex of S and the former means that Q was not internally disjoint to $V(C)$, yielding contradictions in both cases.

We define the sets K_S and K_T as follows. Let C be a component of $G[I]$.

1. If C meets S but not T, then include its edges into K_S.

2. If C meets T but not S, then include its edges into K_T.

3. Otherwise (C meets both of S and T) there is a path P from $\mathcal{L}$ in C. Denote by v_C the last vertex of a maximal initial segment of P in $C - J_T$. As C is a tree, each component C' of $C - v_C$ is a tree and there is a unique edge e between v_C and C'. For every such component C', include the edges of $C' + e$ in K_S if $e \in J_S$ and in K_T otherwise, i.e. if $e \in J_T$.

Note that, by choice of v_C, either v_C is the last vertex of P or the next edge of P belongs to J_T. This ensures that K_S and K_T satisfy (Y.(Y.1)). Moreover, they form a partition of I which satisfies (Y.(Y.1)-(Y.1)) by construction. It remains to show that $\mathrm{cl}_{M_S}(K_S) \cup \mathrm{cl}_{M_T}(K_T) = E$.

As $K_S \cup K_T = I$, it suffices to show that any $e \in E \setminus I$ is spanned by K_S in M_S or by K_T in M_T. Suppose $e \in \mathrm{cl}_{M_S}(J_S)$, i.e. $J_S + e$ contains a circuit of M_S. Hence, $G[J_S]$ either contains an e-path R or two disjoint e–S paths R_1 and R_2. We show that $E(R) \subseteq K_S$ or $E(R) \subseteq K_T$ in the former case and $E(R_1) \cup E(R_2) \subseteq K_S$ in the latter.

The path R is contained in some component C of $G[I]$. Suppose $C \in \mathcal{C}$ and v_C is an inner vertex of R. By assumption, the edges preceding and succeeding v_C on R are both in J_S and hence the edges of both components of $C - v_C$ which are met by R plus their edges to v_C got included into K_S, showing $E(R) \subseteq K_S$. Otherwise $C \notin \mathcal{C}$ or $C \in \mathcal{C}$ but v_C is no inner vertex of R. In both cases the whole set $E(R)$ got included into K_S or K_T.

We apply the same argument to R_1 and R_2 except for one difference. If $C \notin \mathcal{C}$ or $C \in \mathcal{C}$ but v_C is no inner vertex of R_i, then $E(R_i)$ got included into K_S as R_i meets S.

Although the definitions of K_S and K_T are not symmetrical, a similar argument shows $e \in \mathrm{cl}_{M_S}(K_S) \cup \mathrm{cl}_{M_T}(K_T)$ if e is spanned by J_T in M_T. □

Note that the above proof requires only that Conjecture 10.2.1 holds for finite-cycle matroids.

10.5 Union

In this section, we prove Theorem 10.2.8. The main difficulty in proving this theorem is the need to verify that given two nearly finitary matroids M_1 and M_2, that the set system $\mathcal{I}(M_1 \vee M_2)$ satisfies the axioms (IM) and (I3).

To verify the (IM) axiom for the union of two nearly finitary matroids we shall require the following theorem, proved below in Subsection 10.5.2.

Proposition 10.5.1. *If M_1 and M_2 are finitary matroids, then $M_1 \vee M_2$ is a finitary matroid.*

To verify (IM) for the union of finitary matroids we use a compactness argument (see Subsection 10.5.2). More specifically, we will show that $\mathcal{I}(M_1 \vee M_2)$ is a finitary set system whenever M_1 and M_2 are finitary matroids. It is then an easy consequence of Zorn's lemma that all finitary set systems satisfy (IM).

The verification of axiom (I3) is dealt in a joint manner for both matroid families. In the next section we prove the following.

Proposition 10.5.2. *The set system $\mathcal{I}(M_1 \vee M_2)$ satisfies* (I3) *for any two matroids M_1 and M_2.*

Indeed, for finitary matroids, Proposition 10.5.2 is fairly simple to prove. We, however, require this proposition to hold for nearly finitary matroids as well. Consequently, we prove this proposition in its full generality, i.e., for any pair of matroids. In fact, it is interesting to note that the union of infinitely many matroids satisfies (I3); though the axiom (IM) might be violated as seen in Observation 10.5.10).

At this point it is insightful to note a certain difference between the union of finite matroids to that of finitary matroids in a more precise manner. By the finite matroid union theorem if M admits two disjoint bases, then the union of these bases forms a base of $M \vee M$. For finitary matroids the same assertion is false.

Claim 10.5.3. *There exists an infinite finitary matroid M with two disjoint bases whose union is not a base of the matroid $M \vee M$ as it is properly contained in the union of some other two bases.*

Proof. Consider the infinite one-sided ladder with every edge doubled, say H, and recall that the bases of $M_F(H)$ are the ordinary spanning trees of H. In Figure 10.1, (B_1, B_2) and (B_3, B_4) are both pairs of disjoint bases of $M_F(H)$. However, $B_3 \cup B_4$ properly covers $B_1 \cup B_2$ as it additionally contains the leftmost edge of H $\square$

Clearly, a direct sum of infinitely many copies of H gives rise to an infinite sequence of unions of disjoint bases, each properly containing the previous one. In fact, one can construct a (single) matroid formed as the union of two nearly finitary matroids that admits an infinite properly nested sequence of unions of disjoint bases.

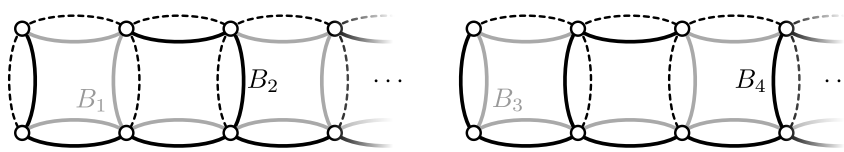

Figure 10.1: The disjoint bases B_1 and B_2 on the left are properly covered by the bases B_3 and B_4 on the right.

10.5.1 Exchange chains and the verification of axiom (I3)

In this section, we prove Proposition 10.5.2. Throughout this section M_1 and M_2 are matroids. It will be useful to show that the following variant of (I3) is satisfied.

Proposition 10.5.4. *The set $\mathcal{I} = \mathcal{I}(M_1 \vee M_2)$ satisfies the following.*

(I3') For all $I, B \in \mathcal{I}$ where B is maximal and all $x \in I \backslash B$ there exists $y \in B \backslash I$ such that $(I + y) - x \in \mathcal{I}$.

Observe that unlike in (I3), the set I in (I3') may be maximal.

We begin by showing that Proposition 10.5.4 implies Proposition 10.5.2.

Proof of Proposition 10.5.2 from Proposition 10.5.4. Let $I \in \mathcal{I}$ be non-maximal and $B \in \mathcal{I}$ be maximal. As I is non-maximal there is an $x \in E \setminus I$ such that $I + x \in \mathcal{I}$. We may assume $x \notin B$ or the assertion follows by (I2). By (I3'), applied to $I + x$, B, and $x \in (I + x) \setminus B$ there is $y \in B \setminus (I + x)$ such that $I + y \in \mathcal{I}$. □

We proceed to prove Proposition 10.5.4. The following notation and terminology will be convenient. A circuit of M which contains a given set $X \subseteq E(M)$ is called an *X-circuit*.

By a *representation of* a set $I \in \mathcal{I}(M_1 \vee M_2)$, we mean a pair (I_1, I_2) where $I_1 \in \mathcal{I}(M_1)$ and $I_2 \in \mathcal{I}(M_2)$ such that $I = I_1 \cup I_2$.

For sets $I_1 \in \mathcal{I}(M_1)$ and $I_2 \in \mathcal{I}(M_2)$, and elements $x \in I_1 \cup I_2$ and $y \in E(M_1) \cup E(M_2)$ (possibly in $I_1 \cup I_2$), a tuple $Y = (y_0 = y, \ldots, y_n = x)$ with $y_i \neq y_{i+1}$ for all i is called an *even* (I_1, I_2, y, x)*-exchange chain*[10] (or *even* (I_1, I_2, y, x)*-chain*) of *length* n if the following terms are satisfied.

(X1) For an even i, there exists a $\{y_i, y_{i+1}\}$-circuit $C_i \subseteq I_1 + y_i$ of M_1.

(X1) For an odd i, there exists a $\{y_i, y_{i+1}\}$-circuit $C_i \subseteq I_2 + y_i$ of M_2.

If $n \geq 1$, then (X1) and (X2) imply that $y_0 \notin I_1$ and that, starting with $y_1 \in I_1 \setminus I_2$, the elements y_i alternate between $I_1 \setminus I_2$ and $I_2 \setminus I_1$; the single exception being y_n which can lie in $I_1 \cap I_2$.

By an *odd exchange chain* (or *odd chain*) we mean an even chain with the words 'even' and 'odd' interchanged in the definition. Consequently, we say *exchange chain* (or *chain*) to refer to either of these notions. Furthermore, a subchain of a chain is also a chain; that is, given an (I_1, I_2, y_0, y_n)-chain $(y_0, \ldots, y_n)$, the tuple $(y_k, \ldots, y_l)$ is an (I_1, I_2, y_k, y_l)-chain for $0 \leq k \leq l \leq n$.

Lemma 10.5.5. *If there exists an (I_1, I_2, y, x)-chain, then $(I+y)-x \in \mathcal{I}(M_1 \vee M_2)$ where $I := I_1 \cup I_2$. Moreover, if $x \in I_1 \cap I_2$, then $I + y \in \mathcal{I}(M_1 \vee M_2)$.*

Remark. In the proof of Lemma 10.5.5 chains are used in order to alter the sets I_1 and I_2; the change is in a single element. Nevertheless, to accomplish this

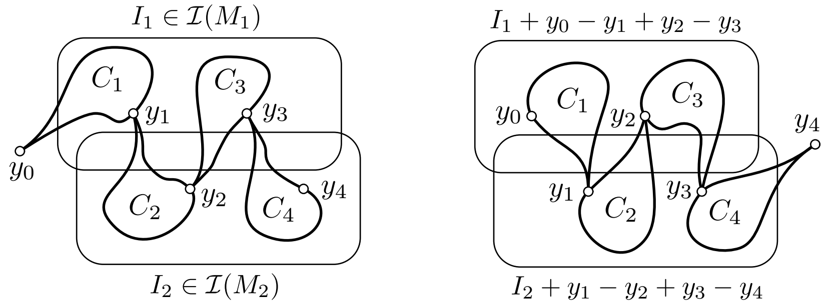

(a) the initial representation

(b) the obtained representation

Figure 10.2: An even exchange chain of length 4.

change, exchange chain of arbitrary length may be required; for instance, a chain of length four is needed to handle the configuration depicted in Figure 10.2.

Next, we prove Lemma 10.5.5.

Proof of Lemma 10.5.5. The proof is by induction on the length of the chain. The statement is trivial for chains of length 0. Assume $n \geq 1$ and that $Y = (y_0, \ldots, y_n)$ is a shortest (I_1, I_2, y, x)-chain. Without loss of generality, let Y be an even chain. If $Y' := (y_1, \ldots, y_n)$ is an (odd) (I_1', I_2, y_1, x)-chain where $I_1' := (I_1 + y_0) - y_1$, then $((I_1' \cup I_2) + y_1) - x \in \mathcal{I}(M_1 \vee M_2)$ by the induction hypothesis and the assertion follows, since $(I_1' \cup I_2) + y_1 = (I_1 \cup I_2) + y_0$. If also $x \in I_1 \cap I_2$, then either $x \in I_1' \cap I_2$ or $y_1 = x$ and hence $n = 1$. In the former case $I + y \in \mathcal{I}(M_1 \vee M_2)$ follows from the induction hypothesis and in the latter case $I + y = I_1' \cup I_2 \in \mathcal{I}(M_1 \vee M_2)$ as $x \in I_2$.

Since I_2 has not changed, (X2) still holds for Y', so to verify that Y' is an (I_1', I_2, y_1, x)-chain, it remains to show $I_1' \in \mathcal{I}(M_1)$ and to check (X1). To this end, let C_i be a $\{y_i, y_{i+1}\}$-circuit of M_1 in $I_1 + y_i$ for all even i. Such exist by (X1) for Y. Notice that any circuit of M_1 in $I_1 + y_0$ has to contain y_0 since $I_1 \in \mathcal{I}(M_1)$. On the other hand, two distinct circuits in $I_1 + y_0$ would give rise to a circuit contained in I_1 by the circuit elimination axiom applied to these two circuits, eliminating y_0. Hence C_0 is the unique circuit of M_1 in $I_1 + y_0$ and $y_1 \in C_0$ ensures $I_1' = (I_1 + y_0) - y_1 \in \mathcal{I}(M_1)$.

To see (X1), we show that there is a $\{y_i, y_{i+1}\}$-circuit C_i' of M_1 in $I_1' + y_i$ for every even $i \geq 2$. Indeed, if $C_i \subseteq I_1' + y_i$, then set $C_i' := C_i$; else, C_i contains an element of $I_1 \setminus I_1' = \{y_1\}$. Furthermore, $y_{i+1} \in C_i \setminus C_0$; otherwise $(y_0, y_{i+1}, \ldots, y_n)$ is a shorter (I_1, I_2, y, x)-chain for, contradicting the choice of Y. Applying the circuit elimination axiom to C_0 and C_i, eliminating y_1 and fixing y_{i+1}, yields a circuit $C_i' \subseteq (C_0 \cup C_i) - y_1$ of M_1 containing y_{i+1}. Finally, as I_1' is independent and $C_i' \setminus I_1' \subseteq \{y_i\}$ it follows that $y_i \in C_i'$. $\square$

[10]Some authors call them *augmenting paths*

We shall require the following. For $I_1 \in \mathcal{I}(M_1)$, $I_2 \in \mathcal{I}(M_2)$, and $x \in I_1 \cup I_2$, let

$$A(I_1, I_2, x) := \{a \mid \text{there exists an } (I_1, I_2, a, x)\text{-chain}\}.$$

This has the property that

$$\text{for every } y \notin A, \text{ either } I_1 + y \in \mathcal{I}(M_1) \text{ or the unique circuit } C_y \text{ of } M_1 \text{ in } I_1 + y \text{ is disjoint from } A. \tag{10.6}$$

To see this, suppose $I_1 + y \notin \mathcal{I}(M_1)$. Then there is a unique circuit C_y of M_1 in $I_1 + y$. If $C_y \cap A = \emptyset$, then the assertion holds so we may assume that $C_y \cap A$ contains an element, a say. Hence there is an (I_1, I_2, a, x)-chain $(y_0 = a, y_1, \ldots, y_{n-1}, y_n = x)$. As $a \in I_1$ this chain must be odd or have length 0, that is, $a = x$. Clearly, $(y, a, y_1, \ldots, y_{n-1}, x)$ is an even (I_1, I_2, y, x)-chain, contradicting the assumption that $y \notin A$.

Next, we prove Proposition 10.5.4.

Proof of Proposition 10.5.4. Let $B \in \mathcal{I}(M_1 \vee M_2)$ maximal, $I \in \mathcal{I}(M_1 \vee M_2)$, and $x \in I \setminus B$. Recall that we seek a $y \in B \setminus I$ such that $(I+y)-x \in \mathcal{I}(M_1 \vee M_2)$. Let (I_1, I_2) and (B_1, B_2) be representations of I and B, respectively. We may assume $I_1 \in \mathcal{B}(M_1|I)$ and $I_2 \in \mathcal{B}(M_2|I)$. We may further assume that for all $y \in B \setminus I$ the sets I_1+y and I_2+y are dependent in M_1 and M_2, respectively, for otherwise it holds that $I+y \in \mathcal{I}(M_1 \vee M_2)$ so that the assertion follows. Hence, for every $y \in (B \cup I) \setminus I_1$ there is a circuit $C_y \subseteq I_1 + y$ of M_1; such contains y and is unique since otherwise the circuit elimination axiom applied to these two circuits eliminating y yields a circuit contained in I_1, a contradiction.

If $A := A(I_1, I_2, x)$ intersects $B \setminus I$, then the assertion follows from Lemma 10.5.5. Else, $A \cap (B \setminus I) = \emptyset$, in which case we derive a contradiction to the maximality of B. To this end, set (Figure 10.3)

$$B'_1 := (B_1 \setminus b_1) \cup i_1 \qquad \text{where} \quad b_1 := B_1 \cap A \quad \text{and} \quad i_1 := I_1 \cap A$$

$$B'_2 := (B_2 \setminus b_2) \cup i_2 \qquad \text{where} \quad b_2 := B_2 \cap A \quad \text{and} \quad i_2 := I_2 \cap A$$

Since A contains x but is disjoint from $B \setminus I$, it holds that $(b_1 \cup b_2)+x \subseteq i_1 \cup i_2$ and thus $B + x \subseteq B'_1 \cup B'_2$. It remains to verify the independence of B'_1 and B'_2 in M_1 and M_2, respectively.

Without loss of generality it is sufficient to show $B'_1 \in \mathcal{I}(M_1)$. For the remainder of the proof 'independent' and 'circuit' refer to the matroid M_1. Suppose for a contradiction that the set B'_1 is dependent, that is, it contains a circuit C. Since i_1 and $B_1 \setminus b_1$ are independent, neither of these contain C. Hence there is an element $a \in C \cap i_1 \subseteq A$. But $C \setminus I_1 \subseteq B_1 \setminus A$ and therefore no C_y with $y \in C \setminus I_1$ contains a by (10.6). Thus, applying the circuit elimination axiom on C eliminating all $y \in C \setminus I_1$ via C_y fixing a, yields a circuit in I_1, a contradiction. $\square$

Since in the proof of Proposition 10.5.4 the maximality of B is only used in order to avoid the case that $B + x \in \mathcal{I}(M_1 \vee M_2)$, one may prove the following slightly stronger statement.

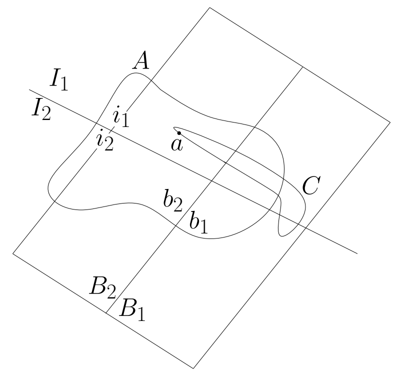

Figure 10.3: The independent sets I_1, at the top, and I_2, at the bottom, the bases B_1, on the right, and B_2, on the left, and their intersection with A.

Corollary 10.5.6. *For all $I, J \in \mathcal{I}(M_1 \vee M_2)$ and $x \in I \setminus J$, if $J + x \notin \mathcal{I}(M_1 \vee M_2)$, then there exists $y \in J \setminus I$ such that $(I + y) - x \in \mathcal{I}(M_1 \vee M_2)$.*

Next, the proof of Proposition 10.5.4, shows that for any maximal representation (I_1, I_2) of I there is $y \in B \setminus I$ such that exchanging finitely many elements of I_1 and I_2 gives a representation of $(I + y) - x$.

For subsequent arguments, it will be useful to note the following corollary. Above we used chains whose last element is fixed. One may clearly use chains whose first element is fixed. If so, then one arrives at the following.

Corollary 10.5.7. *For all $I, J \in \mathcal{I}(M_1 \vee M_2)$ and $y \in J \setminus I$, if $I + y \notin \mathcal{I}(M_1 \vee M_2)$, then there exists $x \in I \setminus J$ such that $(I + y) - x \in \mathcal{I}(M_1 \vee M_2)$.*

10.5.2 Finitary matroid union

In this section, we prove Proposition 10.5.1. In view of Proposition 10.5.2, it remains to show that $\mathcal{I}(M_1 \vee M_2)$ satisfies (IM) whenever M_1 and M_2 are finitary matroids.

The verification of (IM) for countable finitary matroids can be done using König's infinity lemma. Here, in order to capture matroids on any infinite ground set, we employ a topological approach. See [10] for the required topological background needed here.

We recall the definition of the product topology on $\mathcal{P}(E)$. The usual base of this topology is formed by the system of all sets

$$C(A, B) := \{X \subseteq E \mid A \subseteq X, B \cap X = \emptyset\},$$

where $A, B \subseteq E$ are finite and disjoint. Note that these sets are closed as well. Throughout this section, $\mathcal{P}(E)$ is endowed with the product topology and *closed* is used in the topological sense only.

We show that Proposition 10.5.1 can easily be deduced from Proposition 10.5.8 and Lemma 10.5.9, presented next.

Proposition 10.5.8. *Let $\mathcal{I} = \lceil \mathcal{I} \rceil \subseteq \mathcal{P}(E)$. The following are equivalent.*

10.5.8.1. $\mathcal{I}$ is finitary;

10.5.8.1. $\mathcal{I}$ is compact, in the subspace topology of $\mathcal{P}(E)$.

A standard compactness argument can be used in order to prove 10.5.8.10.5.8.1.. Here, we employ a slightly less standard argument to prove 10.5.8.10.5.8.1. as well. Note that as $\mathcal{P}(E)$ is a compact Hausdorff space, assertion 10.5.8.10.5.8.1. is equivalent to the assumption that $\mathcal{I}$ is closed in $\mathcal{P}(E)$, which we use quite often in the following proofs.

Proof of Proposition 10.5.8. To deduce 10.5.8.10.5.8.1. from 10.5.8.10.5.8.1., we show that $\mathcal{I}$ is closed. Let $X \notin \mathcal{I}$. Since $\mathcal{I}$ is finitary, X has a finite subset $Y \notin \mathcal{I}$ and no superset of Y is in $\mathcal{I}$ as $\mathcal{I} = \lceil \mathcal{I} \rceil$. Therefore, $C(Y, \emptyset)$ is an open set containing X and avoiding $\mathcal{I}$ and hence $\mathcal{I}$ is closed.

For the converse direction, assume that $\mathcal{I}$ is compact and let X be a set such that all finite subsets of X are in $\mathcal{I}$. We show $X \in \mathcal{I}$ using the finite intersection property[11] of $\mathcal{P}(E)$. Consider the family $\mathcal{K}$ of pairs (A, B) where $A \subseteq X$ and $B \subseteq E \setminus X$ are both finite. The set $C(A, B) \cap \mathcal{I}$ is closed for every $(A, B) \in \mathcal{K}$, as $C(A, B)$ and $\mathcal{I}$ are closed. If $\mathcal{L}$ is a finite subfamily of $\mathcal{K}$, then

$$\bigcup_{(A,B)\in\mathcal{L}} A \in \bigcap_{(A,B)\in\mathcal{L}} \left(C(A, B) \cap \mathcal{I}\right).$$

As $\mathcal{P}(E)$ is compact, the finite intersection property yields

$$\left(\bigcap_{(A,B)\in\mathcal{K}} C(A, B) \right) \cap \mathcal{I} = \bigcap_{(A,B)\in\mathcal{K}} \left(C(A, B) \cap \mathcal{I}\right) \neq \emptyset.$$

However, $\bigcap_{(A,B)\in\mathcal{K}} C(A, B) = \{X\}$. Consequently, $X \in \mathcal{I}$, as desired. □

Lemma 10.5.9. *If $\mathcal{I}$ and $\mathcal{J}$ are closed in $\mathcal{P}(E)$, then so is $\mathcal{I} \vee \mathcal{J}$.*

Proof. Equipping $\mathcal{P}(E) \times \mathcal{P}(E)$ with the product topology, yields that Cartesian products of closed sets in $\mathcal{P}(E)$ are closed in $\mathcal{P}(E) \times \mathcal{P}(E)$. In particular, $\mathcal{I} \times \mathcal{J}$ is closed in $\mathcal{P}(E) \times \mathcal{P}(E)$. In order to prove that $\mathcal{I} \vee \mathcal{J}$ is closed, we note that $\mathcal{I} \vee \mathcal{J}$ is exactly the image of $\mathcal{I} \times \mathcal{J}$ under the union map

$$f : \mathcal{P}(E) \times \mathcal{P}(E) \to \mathcal{P}(E), \quad f(A, B) = A \cup B.$$

It remains to check that f maps closed sets to closed sets; which is equivalent to showing that f maps compact sets to compact sets as $\mathcal{P}(E)$ is a compact

[11] *The finite intersection property* ensures that an intersection over a family $\mathcal{C}$ of closed sets is non-empty if every intersection of finitely many members of $\mathcal{C}$ is.

Hausdorff space. As continuous images of compact spaces are compact, it suffices to prove that f is continuous, that is, to check that the pre-images of subbase sets $C(\{a\},\emptyset)$ and $C(\emptyset,\{b\})$ are open as can be seen here:

$$f^{-1}(C(\{a\},\emptyset)) = (C(\{a\},\emptyset)\times\mathcal{P}(E)) \cup (\mathcal{P}(E)\times C(\{a\},\emptyset))$$
$$f^{-1}(C(\emptyset,\{b\})) = C(\emptyset,\{b\})\times C(\emptyset,\{b\}).$$

□

Next, we prove Proposition 10.5.1.

Proof of Proposition 10.5.1. By Proposition 10.5.2 it remains to show that the union $\mathcal{I}(M_1)\vee\mathcal{I}(M_2)$ satisfies (IM). As all finitary set systems satisfy (IM), by Zorn's lemma, it is sufficient to show that $\mathcal{I}(M_1\vee M_2)$ is finitary. By Proposition 10.5.8, $\mathcal{I}(M_1)$ and $\mathcal{I}(M_2)$ are both compact and thus closed in $\mathcal{P}(E)$, yielding, by Lemma 10.5.9, that $\mathcal{I}(M_1)\vee\mathcal{I}(M_2)$ is closed in $\mathcal{P}(E)$, and thus compact. As $\mathcal{I}(M_1)\vee\mathcal{I}(M_2) = \lceil\mathcal{I}(M_1)\vee\mathcal{I}(M_2)\rceil$, Proposition 10.5.8 asserts that $\mathcal{I}(M_1)\vee\mathcal{I}(M_2)$ is finitary, as desired. □

We conclude this section with the following observation.

Observation 10.5.10. *A countable union of finitary matroids need not be a matroid.*

Proof. We show that for any integer $k\geq 1$, the set system

$$\mathcal{I} := \bigvee_{n\in\mathbb{N}} U_{k,\mathcal{R}}$$

is not a matroid, where here $U_{k,\mathcal{R}}$ denotes the k-uniform matroid with ground set $\mathcal{R}$.

Since a countable union of finite sets is countable, we have that the members of $\mathcal{I}$ are the countable subsets of $\mathcal{R}$. Consequently, the system $\mathcal{I}$ violates the (IM) axiom for $I=\emptyset$ and $X=\mathcal{R}$. □

Above, we used the fact that the members of $\mathcal{I}$ are countable and that the ground set is uncountable. One can have the following more subtle example, showing that a countable union of finite matroids need not be a matroid.

Let $A=\{a_1,a_2,\ldots\}$ and $B=\{b_1,b_2,\ldots\}$ be disjoint countable sets, and for $n\in\mathbb{N}$, set $E_n := \{a_1,\ldots,a_n\}\cup\{b_n\}$. Then $\bigvee_{n\in\mathbb{N}} U_{1,E_n}$ is an infinite union of finite matroids and fails to satisfy (IM) for $I=A$ and $X=A\cup B=E(M)$.

10.5.3 Nearly finitary matroid union

In this section, we prove Theorem 10.2.8.

For a matroid M, let $\mathcal{I}^{\text{fin}}(M)$ denote the set of subsets of $E(M)$ containing no finite circuit of M, or equivalently, the set of subsets of $E(M)$ which have all their finite subsets in $\mathcal{I}(M)$. We call $M^{\text{fin}} = (E(M),\mathcal{I}^{\text{fin}}(M))$ the *finitarization* of M. With this notation, a matroid M is *nearly finitary* if it has the property that

$$\text{for each } J \in \mathcal{I}(M^{\text{fin}}) \text{ there exists an } I \in \mathcal{I}(M) \text{ such that } |J \setminus I| < \infty. \quad (10.7)$$

For a set system $\mathcal{I}$ (not necessarily the independent sets of a matroid) we call a maximal member of $\mathcal{I}$ a *base* and a minimal member subject to not being in $\mathcal{I}$ a *circuit*. With these conventions, the notions of *finitarization* and *nearly finitary* carry over to set systems.

Let $\mathcal{I} = \lceil \mathcal{I} \rceil$. The finitarization $\mathcal{I}^{\text{fin}}$ of $\mathcal{I}$ has the following properties.

1. $\mathcal{I} \subseteq \mathcal{I}^{\text{fin}}$ with equality if and only if $\mathcal{I}$ is finitary.
2. $\mathcal{I}^{\text{fin}}$ is finitary and its circuits are exactly the finite circuits of $\mathcal{I}$.
3. $(\mathcal{I}|X)^{\text{fin}} = \mathcal{I}^{\text{fin}}|X$, in particular $\mathcal{I}|X$ is nearly finitary if $\mathcal{I}$ is.

The first two statements are obvious. To see the third, assume that $\mathcal{I}$ is nearly finitary and that $J \in \mathcal{I}^{\text{fin}}|X \subseteq \mathcal{I}^{\text{fin}}$. By definition there is $I \in \mathcal{I}$ such that $J \setminus I$ is finite. As $J \subseteq X$ we also have that $J \setminus (I \cap X)$ is finite and clearly $I \cap X \in \mathcal{I}|X$.

Proposition 10.5.11. *The pair $M^{\text{fin}} = (E, \mathcal{I}^{\text{fin}}(M))$ is a finitary matroid, whenever M is a matroid.*

Proof. By construction, the set system $\mathcal{I}^{\text{fin}} = \mathcal{I}(M^{\text{fin}})$ satisfies the axioms (I1) and (I2) and is finitary, implying that it also satisfies (IM).

It remains to show that $\mathcal{I}^{\text{fin}}$ satisfies $(I3)$. By definition, a set $X \subseteq E(M)$ is not in $\mathcal{I}^{\text{fin}}$ if and only if it contains a finite circuit of M.

Let $B, I \in \mathcal{I}^{\text{fin}}$ where B is maximal and I is not, and let $y \in E(M) \setminus I$ such that $I + y \in \mathcal{I}^{\text{fin}}$. If $I + x \in \mathcal{I}^{\text{fin}}$ for any $x \in B \setminus I$, then we are done.

Assuming the contrary, then $y \notin B$ and for any $x \in B \setminus I$ there exists a finite circuit C_x of M in $I + x$ containing x. By maximality of B, there exists a finite circuit C of M in $B + y$ containing y. By the circuit elimination axiom (in M) applied to the circuits C and $\{C_x\}_{x \in X}$ where $X := C \cap (B \setminus I)$, there exists a circuit

$$D \subseteq \left(C \cup \bigcup_{x \in X} C_x \right) \setminus X \subseteq I + y$$

of M containing $y \in C \setminus \bigcup_{x \in X} C_x$. The circuit D is finite, since the circuits C and $\{C_x\}$ are; this contradicts $I + y \in \mathcal{I}^{\text{fin}}$. $\square$

Proposition 10.5.12. *For arbitrary matroids M_1 and M_2 it holds that*

$$\mathcal{I}(M_1^{\text{fin}} \vee M_2^{\text{fin}}) = \mathcal{I}(M_1^{\text{fin}} \vee M_2^{\text{fin}})^{\text{fin}} = \mathcal{I}(M_1 \vee M_2)^{\text{fin}}.$$

Proof. By Proposition 10.5.11, the matroids M_1^{fin} and M_2^{fin} are finitary and therefore $M_1^{\mathrm{fin}} \vee M_2^{\mathrm{fin}}$ is a finitary as well, by Proposition 10.5.1. This establishes the first equality.

The second equality follows from the definition of finitarization provided we show that the finite members of $\mathcal{I}(M_1^{\mathrm{fin}} \vee M_2^{\mathrm{fin}})$ and $\mathcal{I}(M_1 \vee M_2)$ are the same.

Since $\mathcal{I}(M_1) \subseteq \mathcal{I}(M_1^{\mathrm{fin}})$ and $\mathcal{I}(M_2) \subseteq \mathcal{I}(M_2^{\mathrm{fin}})$ it holds that $\mathcal{I}(M_1^{\mathrm{fin}} \vee M_2^{\mathrm{fin}}) \supseteq \mathcal{I}(M_1 \vee M_2)$. On the other hand, a finite set $I \in \mathcal{I}(M_1^{\mathrm{fin}} \vee M_2^{\mathrm{fin}})$ can be written as $I = I_1 \cup I_2$ with $I_1 \in \mathcal{I}(M_1^{\mathrm{fin}})$ and $I_2 \in \mathcal{I}(M_2^{\mathrm{fin}})$ finite. As I_1 and I_2 are finite, $I_1 \in \mathcal{I}(M_1)$ and $I_2 \in \mathcal{I}(M_2)$, implying that $I \in \mathcal{I}(M_1 \vee M_2)$. $\square$

With the above notation a matroid M is nearly finitary if each base of M^{fin} contains a base of M such that their difference is finite. The following is probably the most natural manner to construct nearly finitary matroids (that are not finitary) from finitary matroids.

For a matroid M and an integer $k \geq 0$, set $M[k] := (E(M), \mathcal{I}[k])$, where

$$\mathcal{I}[k] := \{I \in \mathcal{I}(M) \mid \exists J \in \mathcal{I}(M) \text{ such that } I \subseteq J \text{ and } |J \setminus I| = k\}.$$

Proposition 10.5.13. *If* $\mathrm{rk}(M) \geq k$*, then* $M[k]$ *is a matroid.*

Proof. The axiom (I1) holds as $\mathrm{rk}(M) \geq k$; the axiom (I2) holds as it does in M. For (I3) let $I', I \in \mathcal{I}(M[k])$ such that I' is maximal and I is not. There is a set $F' \subseteq E(M) \setminus I'$ of size k such that, in M, the set $I' \cup F'$ is not only independent but, by maximality of I', also a base. Similarly, there is a set $F \subseteq E(M) \setminus I$ of size k such that $I \cup F \in \mathcal{I}(M)$.

We claim that $I \cup F$ is non-maximal in $\mathcal{I}(M)$ for any such F. Suppose not and $I \cup F$ is maximal for some F as above. By assumption, I is contained in some larger set of $\mathcal{I}(M[k])$. Hence there is a set $F^+ \subseteq E(M) \setminus I$ of size $k+1$ such that $I \cup F^+$ is independent in M. Clearly $(I \cup F) \setminus (I \cup F^+) = F \setminus F^+$ is finite, so Lemma 10.5.14 implies that

$$\left|F^+ \setminus F\right| = \left|(I \cup F^+) \setminus (I \cup F)\right| \leq \left|(I \cup F) \setminus (I \cup F^+)\right| = \left|F \setminus F^+\right|.$$

In particular, $k+1 = |F^+| \leq |F| = k$, a contradiction.

Hence we can pick F such that $F \cap F'$ is maximal and, as $I \cup F$ is non-maximal in $\mathcal{I}(M)$, apply (I3) in M to obtain a $x \in (I' \cup F') \setminus (I \cup F)$ such that $(I \cup F) + x \in \mathcal{I}(M)$. This means $I + x \in \mathcal{I}(M[k])$. And $x \in I' \setminus I$ follows, as $x \notin F'$ by our choice of F.

To show (IM), let $I \subseteq X \subseteq E(M)$ with $I \in \mathcal{I}(M[k])$ be given. By (IM) for M, there is a $B \in \mathcal{I}(M)$ which is maximal subject to $I \subseteq B \subseteq X$. We may assume that $F := B \setminus I$ has at most k elements; for otherwise there is a superset $I' \subseteq B$ of I such that $|B \setminus I'| = k$ and it suffices to find a maximal set containing $I' \in \mathcal{I}(M[k])$ instead of I.

We claim that for any $F^+ \subseteq X \setminus I$ of size $k+1$ the set $I \cup F^+$ is not in $\mathcal{I}(M[k])$. For a contradiction, suppose it is. Then in $M|X$, the set $B = I \cup F$ is a base and $I \cup F^+$ is independent and as $(I \cup F) \setminus (I \cup F^+) \subseteq F \setminus F^+$ is finite, Lemma 10.5.14 implies

$$\left|F^+ \setminus F\right| = \left|(I \cup F^+) \setminus (I \cup F)\right| \leq \left|(I \cup F) \setminus (I \cup F^+)\right| = \left|F \setminus F^+\right|.$$

This means $k+1 = |F^+| \leq |F| = k$, a contradiction. So by successively adding single elements of $X \setminus I$ to I as long as the obtained set is still in $\mathcal{I}(M[k])$ we arrive at the wanted maximal element after at most k steps. $\square$

We conclude this section with a proof of Theorem 10.2.8. To this end, we shall require following two lemmas.

Lemma 10.5.14. *Let M be a matroid and $I, B \in \mathcal{I}(M)$ with B maximal and $B \setminus I$ finite. Then, $|I \setminus B| \leq |B \setminus I|$.*

Proof. The proof is by induction on $|B \setminus I|$. For $|B \setminus I| = 0$ we have $B \subseteq I$ and hence $B = I$ by maximality of B. Now suppose there is $y \in B \setminus I$. If $I + y \in \mathcal{I}$ then by induction

$$|I \setminus B| = |(I+y) \setminus B| \leq |B \setminus (I+y)| = |B \setminus I| - 1$$

and hence $|I \setminus B| < |B \setminus I|$. Otherwise there exists a unique circuit C of M in $I + y$. Clearly C cannot be contained in B and therefore has an element $x \in I \setminus B$. Then $(I+y) - x$ is independent, so by induction

$$|I \setminus B| - 1 = |((I+y) - x) \setminus B| \leq |B \setminus ((I+y) - x)| = |B \setminus I| - 1,$$

and hence $|I \setminus B| \leq |B \setminus I|$. $\square$

Lemma 10.5.15. *Let $\mathcal{I} \subseteq \mathcal{P}(E)$ be a nearly finitary set system satisfying (I1), (I2), and the following variant of (I3):*

- *(*) For all $I, J \in \mathcal{I}$ and all $y \in I \setminus J$ with $J + y \notin \mathcal{I}$ there exists $x \in J \setminus I$ such that $(J+y) - x \in \mathcal{I}$.*

Then $\mathcal{I}$ satisfies (IM).

Proof. Let $I \subseteq X \subseteq E$ with $I \in \mathcal{I}$. As $\mathcal{I}^{\text{fin}}$ satisfies (IM) there is a set $B^{\text{fin}} \in \mathcal{I}^{\text{fin}}$ which is maximal subject to $I \subseteq B^{\text{fin}} \subseteq X$ and being in $\mathcal{I}^{\text{fin}}$. As $\mathcal{I}$ is nearly finitary, there is $J \in \mathcal{I}$ such that $B^{\text{fin}} \setminus J$ is finite and we may assume that $J \subseteq X$. Then, $I \setminus J \subseteq B^{\text{fin}} \setminus J$ is finite so that we may choose a J minimizing $|I \setminus J|$. If there is a $y \in I \setminus J$, then by (*) we have $J + y \in \mathcal{I}$ or there is an $x \in J \setminus I$ such that $(J+y) - x \in \mathcal{I}$. Both outcomes give a set containing more elements of I and hence contradicting the choice of J.

It remains to show that J can be extended to a maximal set B of $\mathcal{I}$ in X. For any superset $J' \in \mathcal{I}$ of J, we have $J' \in \mathcal{I}^{\text{fin}}$ and $B^{\text{fin}} \setminus J'$ is finite as it is a subset of $B^{\text{fin}} \setminus J$. As $\mathcal{I}^{\text{fin}}$ is a matroid, Lemma 10.5.14 implies

$$|J' \setminus B^{\text{fin}}| \leq |B^{\text{fin}} \setminus J'| \leq |B^{\text{fin}} \setminus J|.$$

Hence, $|J' \setminus J| \leq 2|B^{\text{fin}} \setminus J| < \infty$. Thus, we can greedily add elements of X to J to obtain the wanted set B after finitely many steps. $\square$

Next, we prove Theorem 10.2.8.

Proof of Theorem 10.2.8. By Proposition 10.5.4, in order to prove that $M_1 \vee M_2$ is a matroid, it is sufficient to prove that $\mathcal{I}(M_1 \vee M_2)$ satisfies (IM). By Corollary 10.5.7 and Lemma 10.5.15 it remains to show that $\mathcal{I}(M_1 \vee M_2)$ is nearly finitary.

So let $J \in \mathcal{I}(M_1 \vee M_2)^{\text{fin}}$. By Proposition 10.5.12 we may assume that $J = J_1 \cup J_2$ with $J_1 \in \mathcal{I}(M_1^{\text{fin}})$ and $J_2 \in \mathcal{I}(M_2^{\text{fin}})$. By assumption there are $I_1 \in \mathcal{I}(M_1)$ and $I_2 \in \mathcal{I}(M_2)$ such that $J_1 \setminus I_1$ and $J_2 \setminus I_2$ are finite. Then $I = I_1 \cup I_2 \in \mathcal{I}(M_1 \vee M_2)$ and the assertion follows as $J \setminus (I_1 \cup I_2) \subseteq (J_1 \setminus I_1) \cup (J_2 \setminus I_2)$ is finite. $\square$

10.6 From infinite matroid union to infinite matroid intersection

In this section, we prove Theorem 10.2.7.

Proof of Theorem 10.2.7. Our starting point is the well-known proof from finite matroid theory that matroid union implies a solution to the matroid intersection problem. With that said, let $B_1 \cup B_2^* \in \mathcal{B}(M_1 \vee M_2^*)$ where $B_1 \in \mathcal{B}(M_1)$ and $B_2^* \in \mathcal{B}(M_2^*)$, and let $B_2 = E \setminus B_2^* \in \mathcal{B}(M_2)$. Then, put $I = B_1 \cap B_2$ and note that $I \in \mathcal{I}(M_1) \cap \mathcal{I}(M_2)$. We show that I admits the required partition.

For an element $x \notin B_i$, $i = 1, 2$, we write $C_i(x)$ to denote the fundamental circuit of x into B_i in M_i. For an element $x \notin B_2^*$, we write $C_2^*(x)$ to denote the fundamental circuit of x into B_2^* in M_2^*. Put $X = B_1 \cap B_2^*$, $Y = B_2 \setminus I$, and $Z = B_2^* \setminus X$, see Figure 10.4.

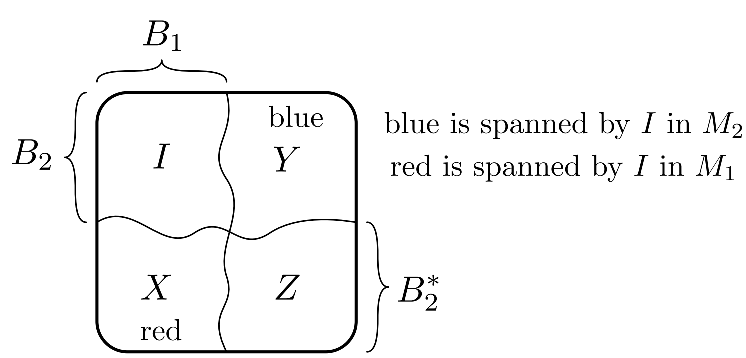

Figure 10.4: The sets X, Y, and Z and their colorings.

Observe that

$$cl_{M_1}(I) \cup cl_{M_2}(I) = E = I \cup X \cup Y \cup Z. \tag{10.8}$$

To see (10.8), note first that

$$X \subseteq cl_{M_2}(I). \tag{10.9}$$

Clearly, no member of X is spanned by I in M_1. Assume then that $x \in X$ is not spanned by I in M_2 so that there exists a $y \in C_2(x) \cap Y$. Then, $x \in C_2^*(y)$, by Lemma 10.3.1. Consequently, $B_1 \cup B_2^* \subsetneq B_1 \cup (B_2^* + y - x) \in \mathcal{I}(M_1 \vee M_2^*)$; contradiction to the maximality of $B_1 \cup B_2^*$, implying (10.9).

By a similar argument, it holds that

$$Y \subseteq cl_{M_1}(I). \tag{10.10}$$

To see that

$$Z \subseteq cl_{M_1}(I) \cup cl_{M_2}(I), \tag{10.11}$$

assume, towards contradiction, that some $z \in Z$ is not spanned by I neither in M_1 nor in M_2 so that there exist an $x \in C_1(z) \cap X$ and a $y \in C_2(z) \cap Y$. Then $B_1 - x + z$ and $B_2 - y + z$ are bases and thus $B_1 \cup B_2^* \subsetneq (B_1 - x + z) \cup (B_2^* - z + y)$; contradiction to the maximality of $B_1 \cup B_2^*$. Assertion (10.8) is proved.

The problem of finding a suitable partition $I = J_1 \cup J_2$ can be phrased as a (directed) graph coloring problem. By (10.8), each $x \in E \setminus I$ satisfies $C_1(x) - x \subseteq I$ or $C_2(x) - x \subseteq I$. Define $G = (V, E)$ to be the directed graph whose vertex set is $V = E \setminus I$ and whose edge set is given by

$$E = \{(x, y) : C_1(x) \cap C_2(y) \cap I \neq \emptyset\}. \tag{10.12}$$

Recall that a *source* is a vertex with no incoming edges and a *sink* is a vertex with no outgoing edges. As $C_1(x)$ does not exist for any $x \in X$ and $C_2(y)$ does not exist for any $y \in Y$, it follows that

$$\text{the members of } X \text{ are sinks and those of } Y \text{ are sources in } G. \tag{10.13}$$

A 2-coloring of the vertices of G, by say blue and red, is called *divisive* if it satisfies the following:

(D.1) I spans all the blue elements in M_1;

(D.1) I spans all the red elements in M_2; and

(D.1) $J_1 \cap J_2 = \emptyset$ where $J_1 := (\bigcup_{x \text{ blue}} C_1(x)) \cap I$ and $J_2 := (\bigcup_{x \text{ red}} C_2(x)) \cap I$.

Clearly, if G has a divisive coloring, then I admits the required partition.

We show then that G admits a divisive coloring. Color with blue all the sources. These are the vertices that can only be spanned by I in M_1. Color with red all the sinks, that is, all the vertices that can only be spanned by I in M_2. This defines a partial coloring of G in which all members of X are red and those of Y are blue. Such a partial coloring can clearly be extended into a divisive coloring of G provided that

$$G \text{ has no } (y, x)\text{-path with } y \text{ blue and } x \text{ red.} \tag{10.14}$$

Indeed, given (10.14) and (10.13), color all vertices reachable by a path from a blue vertex with the color blue, color all vertices from which a red vertex is reachable by a path with red, and color all remaining vertices with, say, blue. The resulting coloring is divisive.

It remains to prove (10.14). We show that the existence of a path as in (10.14) contradicts the following property:

Suppose that M and N are matroids and $B \cup B'$ is maximal in $\mathcal{I}(M \vee N)$. Let $y \notin B \cup B'$ and let $x \in B \cap B'$. Then, (by Lemma 10.5.5)

$$\textit{there exists no } (B, B', y, x)\textit{-chain;} \tag{10.15}$$

(in fact, the contradiction in the proofs of (10.9),(10.10), and (10.11) arose from simple instances of such forbidden chains).

Assume, towards contradiction, that P is a (y, x)-path with y blue and x red; the intermediate vertices of such a path are not colored since they are not a sink nor a source. In what follows we use P to construct a $(B_1, B_2^*, y_0, y_{2|P|})$-chain $(y_0, y_1, \ldots, y_{2|P|})$ such that $y_0 \in Y$, $y_{2|P|} \in X$, all odd indexed members of the chain are in $V(P) \cap Z$, and all even indexed elements of the chain other than y_0 and $y_{2|P|}$ are in I. Existence of such a chain would contradict (10.15).

Definition of $\boldsymbol{y_0}$. As y is pre-colored blue then either $y \in Y$ or $C_2(y) \cap Y \neq \emptyset$. In the former case set $y_0 = y$ and in the latter choose $y_0 \in C_2(y) \cap Y$.

Definition of $\boldsymbol{y_{2|P|}}$. In a similar manner, x is pre-colored red since either $x \in X$ or $C_1(x) \cap X \neq \emptyset$. In the former case, set $y_{2|P|} = x$ and in the latter case choose $y_{2|P|} \in C_1(x) \cap X$.

The remainder of the chain. Enumerate $V(P) \cap Z = \{y_1, y_3, \ldots, y_{2|P|-1}\}$ where the enumeration is with respect to the order of the vertices defined by P. Next, for an edge $(y_{2i-1}, y_{2i+1}) \in E(P)$, let $y_{2i} \in C_1(y_{2i-1}) \cap C_2(y_{2i+1}) \cap I$; such exists by the assumption that $(y_{2i-1}, y_{2i+1}) \in E$. As $y_{2i+1} \in C_2^*(y_{2i})$ for all relevant i, by Lemma 10.3.1, the sequence $(y_0, y_1, y_2, \ldots, y_{2|P|})$ is a $(B_1, B_2^*, y_0, y_{2|P|})$-chain in $\mathcal{I}(M_1 \vee M_2^*)$.

This completes our proof of Theorem 10.2.7. □

Note that in the above proof, we do not use the assumption that $M_1 \vee M_2^*$ is a matroid; in fact, we only need that $\mathcal{I}(M_1 \vee M_2^*)$ has a maximal element.

10.7 The graphic nearly finitary matroids

In this section we prove Propositions 10.2.3 and 10.2.4 yielding a characterization of the graphic nearly finitary matroids.

For a connected graph G, a maximal set of edges containing no finite cycles is called an *ordinary spanning tree.* A maximal set of edges containing no finite cycles nor any double ray is called an *algebraic spanning tree.* These are the bases of $M_F(G)$ and $M_A(G)$, respectively. We postpone the discussion about $M_C(G)$ to Subsection 10.7.2.

To prove Propositions 10.2.3 and 10.2.4, we require the following theorem of Halin [35, Theorem 8.2.5].

Theorem 10.7.1 (Halin 1965)**.** *If an infinite graph G contains k disjoint rays for every $k \in \mathbb{N}$, then G contains infinitely many disjoint rays.*

10.7.1 The nearly finitary algebraic-cycle matroids

The purpose of this subsection is to prove Proposition 10.2.3.

Proof of Proposition 10.2.3. Suppose that G has k disjoint rays for every integer k; so that G has a set $\mathcal{R}$ of infinitely many disjoint rays by Theorem 10.7.1. We show that $M_A(G)$ is not nearly finitary.

The edge set of $\bigcup \mathcal{R} = \bigcup_{R \in \mathcal{R}} R$ is independent in $M_A(G)^{\text{fin}}$ as it induces no finite cycle of G. Therefore there is a base of $M_A(G)^{\text{fin}}$ containing it; such induces an ordinary spanning tree, say T, of G. We show that

$$T - F \text{ contains a double ray for any finite edge set } F \subseteq E(T). \tag{10.16}$$

This implies that $E(T) \setminus I$ is infinite for every independent set I of $M_A(G)$ and hence $M_A(G)$ is not nearly finitary. To see (10.16), note that $T - F$ has $|F| + 1$ components for any finite edge set $F \subseteq E(T)$ as T is a tree and successively removing edges always splits one component into two. So one of these components contains infinitely many disjoint rays from $\mathcal{R}$ and consequently a double ray.

Suppose next, that G has at most k disjoint rays for some integer k and let T be an ordinary spanning tree of G, that is, $E(T)$ is maximal in $M_A(G)^{\text{fin}}$. To prove that $M_A(G)$ is nearly finitary, we need to find a finite set $F \subseteq E(T)$ such that $E(T) \setminus F$ is independent in $M_A(G)$, i.e. it induces no double ray of G. Let $\mathcal{R}$ be a maximal set of disjoint rays in T; such exists by assumption and $|\mathcal{R}| \leq k$. As T is a tree and the rays of $\mathcal{R}$ are vertex-disjoint, it is easy to see that T contains a set F of $|\mathcal{R}| - 1$ edges such that $T - F$ has $|\mathcal{R}|$ components which each contain one ray of $\mathcal{R}$. By maximality of $\mathcal{R}$ no component of $T - F$ contains two disjoint rays, or equivalently, a double ray. $\square$

10.7.2 The nearly finitary topological-cycle matroids

In this section we prove Proposition 10.2.4 that characterizes the nearly finitary topological-cycle matroids. Prior to that, we first define these matroids. To that end we shall require some additional notation and terminology on which more details can be found in [21].

An *end of G* is an equivalence class of rays, where two rays are *equivalent* if they cannot be separated by a finite edge set. In particular, two rays meeting infinitely often are equivalent. Let the *degree* of an end ω be the size of a maximal set of vertex-disjoint rays belonging to ω, which is well-defined [36]. We say that a double ray *belongs to* an end if the two rays which arise from the removal of one edge from the double ray belong to that end; this does not depend on the choice of the edge. Such a double ray is an example of a *topological cycle*[12]

For a graph G the topological-cycle matroid of G, namely $M_C(G)$, has $E(G)$ as its ground set and its set of circuits consists of the finite and topological cycles. In fact, every infinite circuit of $M_C(G)$ induces at least one double ray; provided that G is locally finite [36].

A graph G has only finitely many disjoint rays if and only if G has only finitely many ends, each with finite degree. Also, note that

$$\text{every end of a 2-connected locally finite graph has degree at least 2.} \tag{10.17}$$

Indeed, applying Menger's theorem inductively, it is easy to construct in any k-connected graph for any end ω a set of k disjoint rays of ω.

Now we are in a position to start the proof of Proposition 10.2.4.

Proof of Proposition 10.2.4. If G has only a finite number of vertex-disjoint rays then $M_A(G)$ is nearly finitary by Proposition 10.2.3. Since $M_A(G)^{\text{fin}} = M_C(G)^{\text{fin}}$ and $\mathcal{I}(M_A(G)) \subseteq \mathcal{I}(M_C(G))$, we can conclude that $M_C(G)$ is nearly finitary as well.

Now, suppose that G contains k vertex-disjoint rays for every $k \in \mathbb{N}$. If G has an end ω of infinite degree, then there is an infinite set $\mathcal{R}$ of vertex-disjoint rays belonging to ω. As any double ray containing two rays of $\mathcal{R}$ forms a circuit of $M_C(G)$, the argument from the proof of Proposition 10.2.3 shows that $M_C(G)$ is not nearly finitary.

Assume then that G has no end of infinite degree. There are infinitely many disjoint rays, by Theorem 10.7.1. Hence, there is a countable set of ends $\Omega = \{\omega_1, \omega_2, \dots\}$.

We inductively construct a set $\mathcal{R}$ of infinitely many vertex-disjoint double rays, one belonging to each end of Ω. Suppose that for any integer $n \geq 0$ we have constructed a set $\mathcal{R}_n$ of n disjoint double rays, one belonging to each of the ends $\omega_1, \dots, \omega_n$. Different ends can be separated by finitely many vertices so there is a finite set S of vertices such that $\bigcup \mathcal{R}_n$ has no vertex in the component C of $G - S$ which contains ω_{n+1}. Since ω_{n+1} has degree 2 by (10.17), there are two disjoint rays from ω_{n+1} in C an thus also a double ray D belonging to ω_{n+1}. Set $\mathcal{R}_{n+1} := \mathcal{R}_n \cup \{D\}$ and $\mathcal{R} := \bigcup_{n \in \mathbb{N}} \mathcal{R}_n$.

As $\bigcup \mathcal{R}$ contains no finite cycle of G, it can be extended to an ordinary spanning tree of G. Removing finitely many edges from this tree clearly leaves an element of $\mathcal{R}$ intact. Hence, the edge set of the resulting graph still contains a circuit of $M_C(G)$. Thus, $M_C(G)$ is not nearly finitary in this case as well. $\square$

In the following we shall propose a possible extension of Theorem 10.7.1 to matroids. We call a matroid M *k-nearly finitary* if every base of its finitarization contains a base of M such that their difference has size at most k. Note that

[12] Formally, the topological cycles of G are those subgraphs of G which are homeomorphic images of S^1 in the Freudenthal compactification $|G|$ of G. However, the given example is the only type of topological cycle which shall be needed for the proof.

saying 'at most k' is not equivalent to saying 'equal to k', consider for example the algebraic-cycle matroid of the infinite ladder. We conjecture the following.

Conjecture 10.7.2. *Every nearly finitary matroid is k-nearly finitary for some k.*

We remark that Propositions 10.2.3 and 10.2.4 above are special cases of this conjecture. In the proof of Proposition 10.2.4 we used Theorem 10.7.1. In fact it is not difficult to show that Proposition 10.2.4 and Theorem 10.7.1 are equivalent. In particular, Conjecture 10.7.2 implies Theorem 10.7.1.

10.7.3 Graphic matroids and the intersection conjecture

By Theorem 10.2.5, the intersection conjecture is true for $M_C(G)$ and $M_{FC}(H)$ for any two graphs G and H since the first is co-finitary and the second is finitary. Using also Proposition 10.2.4, we obtain the following.

Corollary 10.7.3. *Suppose that G and H are 2-connected, locally finite graphs with only a finite number of vertex-disjoint rays. Then, $M_C(G)$ and $M_C(H)$ satisfy the intersection conjecture.* □

Using Proposition 10.2.3 instead of Proposition 10.2.4, we obtain the following.

Corollary 10.7.4. *Suppose that G and H are graphs with only a finite number of vertex-disjoint rays. Then, $M_A(G)$ and $M_A(H)$ satisfy the intersection conjecture if both are matroids.* □

With a little more work, the same is also true for $M_{FC}(G)$, see Corollary 10.2.6.

Proof of Corollary 10.2.6. First we show that $(((M_C(G)^{\text{fin}})^*)^{\text{fin}})^* = M_C(G)$ if G is locally finite. Indeed, then $M_C(G)^{\text{fin}} = M_{FC}(G)$, $M_{FC}(G)^*$ is the matroids whose circuits are the finite and infinite bonds of G, and its finitarization has as its circuits the finite bonds of G. And the dual of this matroid is $M_C(G)$, see [22] for example.

Having showed that $(((M_C(G)^{\text{fin}})^*)^{\text{fin}})^* = M_C(G)$ if G is locally finite, we next show that if $M_C(G)$ is nearly finitary, then so is $M_{FC}(G)^*$. For this let B be a base of $M_{FC}(G)^*$ and B' be a base of $(M_{FC}(G)^*)^{\text{fin}}$. Then $B' \setminus B = (E \setminus B) \setminus (E \setminus B')$. Now $E \setminus B$ is a base of $M_{FC}(G) = M_C(G)^{\text{fin}}$ and by the above $E \setminus B'$ is a base of $M_C(G)$. Since $M_C(G)$ is nearly finitary, $B' \setminus B$ is finite, yielding that $M_{FC}(G)^*$ is nearly finitary.

As $M_{FC}(G)^*$ is nearly finitary and $M_{FC}(H)$ is finitary, $M_{FC}(H)$ and $M_{FC}(G)$ satisfy the intersection conjecture by Theorem 10.2.5. □

A similar argument shows that if G and H are are 2-connected, locally finite graphs with only a finite number of vertex-disjoint rays, then one can also prove that $M_{FC}(G)^*$ and $M_{FC}(H)^*$ satisfy the intersection conjecture. Similar results are true for $M_C(G)^*$ or $M_A(G)^*$ in place of $M_{FC}(G)^*$.

10.8 Union of arbitrary infinite matroids

In this section, we show that there exists infinite matroids M and N whose union is not a matroid.

In Claim 10.8.1, we treat the relatively simpler case in which M is finitary and N is co-finitary and both have uncountable ground sets. Second, then, in Claim 10.8.2, we refine the argument as to have M both finitary and co-finitary and N co-finitary and both on countable ground sets.

Claim 10.8.1. *There exists a finitary matroid M and a co-finitary matroid N such that $\mathcal{I}(M \vee N)$ is not a matroid.*

Proof. Set $E = E(M) = E(N) = \mathbb{N} \times \mathcal{R}$. Next, put $M := \bigoplus_{n\in\mathbb{N}} M_n$, where $M_n := U_{1,\{n\}\times\mathcal{R}}$. The matroid M is finitary as it is a direct sum of 1-uniform matroids. For $r \in \mathcal{R}$, let N_r be the circuit matroid on $\mathbb{N} \times \{r\}$; set $N := \bigoplus_{r\in\mathcal{R}} N_r$. As N is a direct sum of circuits, it is co-finitary. (see Figure 10.5).

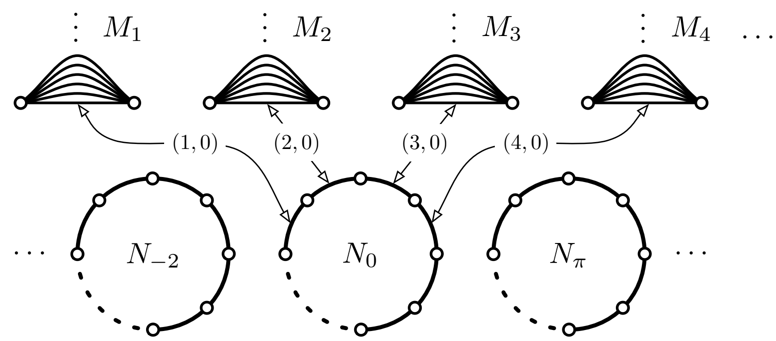

Figure 10.5: $M = \bigoplus_{n\in\mathbb{N}} M_n$ and $N = \bigoplus_{r\in\mathcal{R}} N_r$.

We show that $\mathcal{I}(M \vee N)$ violates the axiom (IM) for $I = \emptyset$ and $X = E$; so that $\mathcal{I}(M \vee N)$ has no maximal elements. It is sufficient to show that a set $J \subseteq E$ belongs to $\mathcal{I}(M \vee N)$ if and only if it contains at most countably many circuits of N. For if so, then for any $J \in \mathcal{I}(M \vee N)$ and any circuit $C = \mathbb{N} \times \{r\}$ of N with $C \nsubseteq J$ (such a circuit exists) we have $J \cup C \in \mathcal{I}(M \vee N)$.

The point to observe here is that every independent set of M is countable, (since every such set meets at most one element of M_n for each $n \in \mathbb{N}$), and that every independent set of N misses uncountably many elements of E (as any such set must miss at least one element of N_r for each $r \in \mathcal{R}$).

Suppose $J \subseteq E$ contains uncountably many circuits of N. Since each independent set of N misses uncountably many elements of E, every set $D = J \setminus J_N$ is uncountable whenever $J_N \in \mathcal{I}(J)$. On the other hand, since each independent set of M is countable, we have that $D \notin \mathcal{I}(M)$. Consequently, $J \notin \mathcal{I}(M \vee N)$, as required.

We may assume then that $J \subseteq E$ contains only countably many circuits of N, namely, $\{C_{r_1}, C_{r_2}, \ldots\}$. Now the set $J_M = \{(i, r_i) : i \in \mathbb{N}\}$ is independent in

M; consequently, $J \setminus J_M$ is independent in N; completing the proof. $\square$

We proceed with matroids on countable ground sets.

Claim 10.8.2. *There exist a matroid M that is both finitary and co-finitray, and a co-finitary matroid N whose common ground is countable such that $\mathcal{I}(M \vee N)$ is not a matroid.*

Proof. For the common ground set we take $E = (\mathbb{N} \times \mathbb{N}) \cup L$ where $L = \{\ell_1, \ell_2, \dots\}$ is countable and disjoint to $\mathbb{N} \times \mathbb{N}$. The matroids N and M are defined as follows. For $r \in \mathbb{N}$, let N_r be the circuit matroid on $\mathbb{N} \times \{r\}$. Set N to be the matroid on E obtained by adding the elements of L to the matroid $\bigoplus_{r \in \mathbb{N}} N_r$ as loops. Next, for $n \in \mathbb{N}$, let M_n be the 1-uniform matroid on $(\{n\} \times \{1, 2, \dots, n\}) \cup \{\ell_n\}$. Let M be the matroid obtained by adding to the matroid $\bigoplus_{n \in \mathbb{N}} M_n$ all the members of $E \setminus E(\bigoplus_{n \in \mathbb{N}} M_n)$ as loops

We show that $\mathcal{I}(M \vee N)$ violates the axiom (IM) for $I = \mathbb{N} \times \mathbb{N}$ and $X = E$. It is sufficient to show that

(a) $I \in \mathcal{I}(M \vee N)$; and that

(b) every set J satisfying $I \subset J \subseteq E$ is in $\mathcal{I}(M \vee N)$ if and only if it misses infinitely many elements of L.

To see that $I \in \mathcal{I}(M \vee N)$, note that the set $I_M = \{(n, n) \mid n \in \mathbb{N}\}$ is independent in M and meets each circuit $\mathbb{N} \times \{r\}$ of N. In particular, the set $I_N := (\mathbb{N} \times \mathbb{N}) \setminus I_M$ is independent in N, and therefore $I = I_M \cup I_N \in \mathcal{I}(M \vee N)$.

Let then J be a set satisfying $I \subseteq J \subseteq E$, and suppose, first, that $J \in \mathcal{I}(M \vee N)$. We show that J misses infinitely many elements of L.

There are sets $J_M \in \mathcal{I}(M)$ and $J_N \in \mathcal{I}(N)$ such that $J = J_M \cup J_N$. As J_N misses at least one element from each of the disjoint circuits of N in I, the set $D := I \setminus J_N$ is infinite. Moreover, we have that $D \subseteq J_M$, since $I \subseteq J$. In particular, there is an infinite subset $L' \subseteq L$ such that $D + l$ contains a circuit of M for every $\ell \in L'$. Indeed, for every $e \in D$ is contained in some M_{n_e}; let then $L' = \{\ell_{n_e} : e \in D\}$ and note that $L' \cap J = \emptyset$. This shows that J_M and L' are disjoint and thus J and L' are disjoint as well, and the assertion follows.

Suppose, second, that there exists a sequence $i_1 < i_2 < \dots$ such that J is disjoint from $L' = \{\ell_{i_r} : r \in \mathbb{N}\}$. We show that the superset $E \setminus L'$ of J is in $\mathcal{I}(M \vee N)$. To this end, set $D := \{(i_r, r) \mid r \in \mathbb{N}\}$. Then, D meets every circuit $\mathbb{N} \times \{r\}$ of N in I, so that the set $J_N := \mathbb{N} \times \mathbb{N} \setminus D$ is independent in N. On the other hand, D contains a single element from each M_n with $n \in L'$. Consequently, $J_M := (L \setminus L') \cup D \in \mathcal{I}(M)$ and therefore $E \setminus L' = J_M \cup J_N \in \mathcal{I}(M \vee N)$. $\square$

While the union of two finitary matroids is a matroid, by Proposition 10.5.1, the same is not true for two co-finitary matroids.

Corollary 10.8.3. *The union of two co-finitary matroids is not necessarily a matroid.*

Since two matroids M and N^* satisfy Conjecture 10.2.1 by Theorem 10.2.7 if the union of M and N is a matroid, it seems worth investigating where the boundaries of this approach are. In particular, we have the following question. Is the class of nearly finitary matroids the largest class containing the fintary matroids that is closed under taking (finite) unions in the following sense?

Question 10.8.4. *Is there for every non-nearly finitary matroid M a finitary matroid N such that the union of M and N is not a matroid?*

In [5] we prove that this conjecture is true for any matroid M such that the finitarization of M has an independent set I containing only countably many M-circuits such that I has no finite subset meeting all of these circuits.

Chapter 11

An excluded minors method for infinite matroids

11.1 Abstract

The notion of thin sums matroids was invented to extend the notion of representability to non-finitary matroids. A matroid is tame if every circuit-cocircuit intersection is finite. We prove that a tame matroid is a thin sums matroid over a finite field k if and only if all its finite minors are representable over k.

11.2 Introduction

Given a family of vectors in a vector space over some field k, there is a matroid structure on that family whose independent sets are given by the linearly independent subsets of the family. Matroids arising in this way are called *representable* matroids over k. A classical theorem of Tutte [92] states that a finite matroid is binary (that is, representable over $\mathbb{F}_2$) if and only if it does not have $U_{2,4}$ as a minor. In the same spirit, a key aim of finite matroid theory has been to determine such 'forbidden minor' characterisations for the classes of matroids representable over other finite fields. For example Bixby and Seymour [12, 87] characterized the finite ternary matroids (those representable over $\mathbb{F}_3$) by forbidden minors, and more recently there is a forbidden minors characterisation for the finite matroids representable over $\mathbb{F}_4$, due to Geelen, Gerards and Kapoor [42]. In 1971 Rota conjectured that for any finite field the class of finite matroids representable over that field is characterised by finitely many forbidden minors. A proof of this conjecture has been announced by Geelen, Gerards and Whittle. An outline of the proof has already appeared in [43]. In this chapter we develop a method which makes it possible to extend the above excluded minor characterisations from finite to infinite matroids.

It is clear that any representable matroid is *finitary*, that is, all its circuits are finite, and so many interesting examples of infinite matroids are not representable. However, since the construction of many standard examples, including the algebraic cycle matroids of infinite graphs, is suggestively similar to that of representable matroids, the notion of *thin sums matroids* was introduced in [21]: it is a generalisation of representability which captures these infinite examples. We will work with thin sums matroids rather than with representable matroids.

In [1] it was shown that the class of tame thin sums matroids over a fixed field is closed under duality, where a matroid is *tame* if any circuit-cocircuit intersection is finite. On the other hand, there are thin sums matroids whose dual is not a thin sums matroid [15] - such counterexamples cannot be tame. A simple consequence of this closure under duality is that the class of tame thin sums matroids over a fixed field is closed under taking minors, and so we may consider the forbidden minors for this class.

Minor closed classes may have infinite 'minimal' forbidden minors. For example the class of finitary matroids has the infinite circuit $U^*_{1,\mathbb{N}}$ as a forbidden minor. Similarly, the class of tame thin sums matroids over $\mathbb{R}$ has $U_{2,\mathcal{P}(\mathbb{R})}$ as a forbidden minor. However, our main result states that the class of tame thin sums matroids over a fixed *finite* field has only finite minimal forbidden minors.

Theorem 11.2.1. *Let M be a tame matroid and k be a finite field. Then M is a thin sums matroid over k if and only if every finite minor of M is k-representable.*

The proof is by a compactness argument. All previous compactness proofs in infinite matroid theory known to the authors use only that either all finite restrictions or all finite contractions have a certain property to conclude that the matroid itself has the desired property. For our purposes, arguments of this kind must fail because there is a tame matroid all of whose finite restrictions and finite contractions are binary but which is not a thin sums matroid over $\mathbb{F}_2$ - in fact, it has a $U_{2,4}$-minor. We shall briefly sketch how to construct such a matroid. Start with $U_{2,4}$, and add infinitely many elements parallel to one of its elements. This ensures that every finite contraction is binary. If we also add infinitely many elements which are parallel in the dual to some other element then we guarantee in addition that all finite restrictions are binary, but the matroid itself has a $U_{2,4}$ minor.

Theorem 11.2.1 implies that each of the excluded minor characterisations for finite representable matroids mentioned in the first paragraph extends to tame matroids. Thus, for example, a tame matroid is a thin sums matroid over $\mathbb{F}_2$ if and only if it has no $U_{2,4}$ minor. Any future excluded minor characterisations for finite matroids representable over a fixed finite field will also immediately extend to tame matroids by this theorem.

Our approach makes it possible to lift many other standard theorems about finite matroids representable over a finite field to theorems about tame thin sums matroids over the same field. For example, the same method shows that a tame matroid is regular (that is, a thin sums matroid over every field) if and only if all its finite minors are, and that regularity is equivalent to signability for tame matroids (see Section 3 for a definition). Our method applies to excluded minor characterisations of properties other than representability. In [17] the same method is employed to show that the tame matroids all of whose finite minors are graphic are precisely those matroids that arise from some *graph-like space*, in the sense that the circuits are given by topological circles and the cocircuits by topological bonds.

The proof of Theorem 11.2.1 will appear in Section 4, but first we must introduce some basic preliminary results for those without a background in infinite matroid theory. In Section 3 we treat the binary case separately. This is simpler but many ideas can alredy be seen there. In Section 5 we apply these methods to related representations such as regular ones.

11.3 Preliminaries

11.3.1 Basics

Throughout, notation and terminology for graphs are those of [36], and for matroids those of [76, 22]. And M always denotes a matroid and $E(M)$ (or just E), $\mathcal{I}(M)$ and $\mathcal{C}(M)$ denote its ground set and its sets of independent sets and circuits, respectively.

A set system $\mathcal{I} \subseteq \mathcal{P}(E)$ is the set of independent sets of a matroid if and only if it satisfies the following *independence axioms* [22].

(I1) $\emptyset \in \mathcal{I}(M)$.

(I2) $\mathcal{I}(M)$ is closed under taking subsets.

(I3) Whenever $I, I' \in \mathcal{I}(M)$ with I' maximal and I not maximal, there exists an $x \in I' \setminus I$ such that $I + x \in \mathcal{I}(M)$.

(IM) Whenever $I \subseteq X \subseteq E$ and $I \in \mathcal{I}(M)$, the set $\{I' \in \mathcal{I}(M) \mid I \subseteq I' \subseteq X\}$ has a maximal element.

A set system $\mathcal{C} \subseteq \mathcal{P}(E)$ is the set of circuits of a matroid if and only if it satisfies the following *circuit axioms* [22].

(C1) $\emptyset \notin \mathcal{C}$.

(C2) No element of $\mathcal{C}$ is a subset of another.

(C3) (Circuit elimination) Whenever $X \subseteq o \in \mathcal{C}(M)$ and $\{o_x \mid x \in X\} \subseteq \mathcal{C}(M)$ satisfies $x \in o_y \Leftrightarrow x = y$ for all $x, y \in X$, then for every $z \in o \setminus \left(\bigcup_{x \in X} o_x\right)$ there exists a $o' \in \mathcal{C}(M)$ such that $z \in o' \subseteq \left(o \cup \bigcup_{x \in X} o_x\right) \setminus X$.

(CM) $\mathcal{I}$ satisfies (IM), where $\mathcal{I}$ is the set of those subsets of E not including an element of $\mathcal{C}$.

We will rely on the following straightforward lemmas, which may be proved for infinite matroids in essentially the same way as for finite matroids. Fix some matroid M.

Lemma 11.3.1. *Let M be a matroid and s be a base. Let o_e and b_f a fundamental circuit and a fundamental cocircuit with respect to s, then*

1. *$o_e \cap b_f$ is empty or $o_e \cap b_f = \{e, f\}$ and*

2. *$f \in o_e$ iff $e \in b_f$.*

Proof. To see the first note that $o_e \subseteq s + e$ and $b_f \subseteq (E \setminus s) + f$. So $o_e \cap b_f \subseteq \{e, f\}$. As a circuit and a cocircuit can never meet in only one edge, the assertion follows.

To see the second, first let $f \in o_e$. Then $f \in o_e \cap b_f$, so by (1) $o_e \cap b_f = \{e, f\}$ and so $e \in b_f$. The converse implication is the dual statement of the above implication. □

Lemma 11.3.2. *For any circuit o containing two edges e and f, there is a cocircuit b such that $o \cap b = \{e, f\}$.*

Proof. As $o - e$ is independent, there is a base including $o - e$. By Lemma 11.3.1, the fundamental cocircuit of f of this base intersects o in e and f, as desired. □

Lemma 11.3.3. *Let M be a matroid with ground set $E = C \dot\cup X \dot\cup D$ and let o' be a circuit of $M' = M/C \backslash D$. Then there is an M-circuit o with $o' \subseteq o \subseteq o' \cup C$.*

Proof. Let s be any M-base of C. Then $s \cup o'$ is M-dependent since o' is M'-dependent. On the other hand, $s \cup o' - e$ is M-independent whenever $e \in o'$ since $o' - e$ is M'-independent. Putting this together yields that $s \cup o'$ contains an M-circuit o, and this circuit must not avoid any $e \in o'$, as desired. □

Corollary 11.3.4. *Let M' be a minor of M. Further let o' be an M'-circuit and b' be an M'-cocircuit. Then there is an M-circuit $o \subseteq o' \cup (E(M) \setminus E(M'))$ and an M-cocircuit $b \subseteq b' \cup (E(M) \setminus E(M'))$ such that $o \cap b = o' \cap b'$.*

A *scrawl* is a union of circuits. For any matroid M, M can be recovered from its set of scrawls since the circuits are precisely the minimal nonempty scrawls.

Lemma 11.3.5. *Let M be a matroid, and let $w \subseteq E$. The following are equivalent:*

1. *w is a scrawl of M.*

2. *w never meets a cocircuit of M just once.*

Corollary 11.3.6. *Let M be a matroid with ground set $E = C \dot\cup X \dot\cup D$, and let $w' \subseteq X$. Then w' is a scrawl of $M' = M/C \backslash D$ if and only if there is a scrawl w of M with $w' \subseteq w \subseteq w' \cup C$.*

11.3.2 Thin sums matroids

Throughout the whole chapter, we will follow the convention that if we write that a sum equals zero then this implicitly includes the statement that this sum is well-defined, that is, that only finitely many summands are nonzero.

Definition 11.3.7. Let A be a set, and k a field. Let $f = (f_e | e \in E)$ be a family of functions from A to k, and let $\lambda = (\lambda_e | e \in E)$ be a family of elements of k. We say that λ is a *thin dependence* of f if and only if for each $a \in A$ we have

$$\sum_{e \in E} \lambda_e f_e(a) = 0 \,,$$

We say that a subset I of E is *thinly independent* for f if and only if the only thin dependence of f which is 0 everywhere outside I is $(0|e \in E)$. The *thin sums system* M_f of f is the set of such thinly independent sets. This isn't always the set of independent sets of a matroid [22], but when it is we say that this matroid is the *thin sums matroid* of f, and that it is *thinly represented* by f over k.

This definition is deceptively similar to the definition of the representable matroid corresponding to f considered as a family of vectors in the k-vector space k^A. The difference is in the more liberal definition of dependence: it is possible for λ to be a thin dependence even if there are infinitely many $e \in E$ with $\lambda_e \neq 0$, provided that for each $a \in A$ there are only finitely many $e \in E$ such that *both* $\lambda_e \neq 0$ and $f_e(a) \neq 0$.

Indeed, the notion of thin sums matroid was introduced as a generalisation of the notion of representable matroid: every representable matroid is finitary, but this restriction does not apply to thin sums matroids.

There are many natural examples of thinly representable matroids: for example, finite, topological and algebraic cycle matroids of graphs are always thinly representable over every field [1]. A finitary matroid is thinly representable over k if and only if it is representable in the usual sense [1].

The following connection between scrawls and thin dependences will turn out to be useful.

Lemma 11.3.8 ([1])**.** *Let M_f be a thinly representable matroid, and let c be a linear dependence for f. Then the support of c is a scrawl.*

Let k be a field and let k^* denote the set of nonzero elements of k. A *k-painting for the matroid M* is a choice of a function $c_o \colon o \to k^*$ for each circuit o of M and a function $d_b \colon b \to k^*$ for each cocircuit b of M such that for any circuit o and cocircuit b we have

$$\sum_{e \in o \cap b} c_o(e) d_b(e) = 0 \,. \tag{11.1}$$

A matroid is *k-paintable* if it has a k-painting.

Lemma 11.3.9 ([1])**.** *Let M be a tame matroid. Then M is a thin sums matroid over the field k if and only if M is k-paintable.*

By symmetry of the definition, it is clear that a matroid is k-paintable if and only if its dual is. Each k-painting $((c_o|o \in \mathcal{C}(M)), (d_b|b \in \mathcal{C}(M^*)))$ of M induces at least one k-painting $((c'_o|o \in \mathcal{C}(N)), (d'_b|b \in \mathcal{C}(N^*)))$ for each minor $N = M/C\backslash D$ in the following way: By Lemma 11.3.3, for each circuit o of N we can pick a circuit $\bar{o}$ of M such that $o \subseteq \bar{o} \subseteq o \cup C$. Similarly, for each cocircuit b of N we can pick a cocircuit $\bar{b}$ of M such that $b \subseteq \bar{b} \subseteq b \cup D$. Let $c'_o = c_{\bar{o}}\restriction_o$ and $d'_b = d_{\bar{b}}\restriction_b$. Then $((c'_o|o \in \mathcal{C}(N)), (d'_b|b \in \mathcal{C}(N^*)))$ is a k-painting of N.

In particular, if a matroid is k-paintable then all its finite minors are.

11.4 Binary matroids

Theorem 11.4.1. *Let M be a tame matroid. Then the following are equivalent:*

1. *M is a binary thin sums matroid.*
2. *For any circuit o and cocircuit b of M, $|o \cap b|$ is even.*
3. *For any circuit o and cocircuit b of M, $|o \cap b| \neq 3$*
4. *M has no minor isomorphic to $U_{2,4}$.*
5. *If o_1, o_2 are circuits then $o_1 \triangle o_2$ is empty or includes a circuit.*
6. *If o_1, o_2 are circuits then $o_1 \triangle o_2$ is a disjoint union of circuits.*
7. *If $(o_i|i \in I)$ is a finite family of circuits then $\triangle_{i \in I} o_i$ is empty or includes a circuit.*
8. *If $(o_i|i \in I)$ is a finite family of circuits then $\triangle_{i \in I} o_i$ is a disjoint union of circuits.*
9. *For any base s of M, and any circuit o of M, $o = \triangle_{e \in o \backslash s} o_e$, where o_e is the fundamental circuit of e with respect to s.*

Proof. We shall prove the following implications:

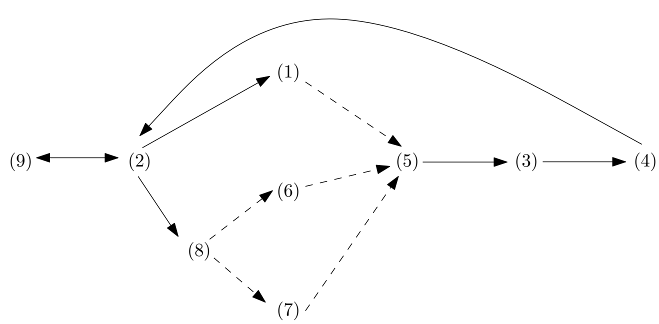

Those implications indicated by dotted arrows are clear. We shall prove the remaining implications.

(2) implies (1): We need to find a suitable thin sums system. Let A be the set of cocircuits of M, and let $E \xrightarrow{f} \mathbb{F}_2^A$ be the map sending e to the function which sends $b \in A$ to 1 if $e \in b$ and 0 otherwise.

We are to show that the thin sums matroid M_{ts} defined by f is M. Since the characteristic function of any M-circuit is a thin dependence for f with support equal to that circuit by (2), any M-dependent set is also $M_{ts}(f)$-dependent.

It remains to show that the support of every non-zero thin dependence is M-dependent. By Lemma 11.3.5 the support of every non-zero thin dependence includes a circuit, as desired.

(2) implies (8): Let $(o_i|i \in I)$ be a finite family of circuits. By Zorn's Lemma, we can choose a maximal family $(o_j|j \in J)$ of disjoint circuits such that $\bigcup_{j\in J} o_j \subseteq \triangle_{i\in I} o_i$, and let $w = \triangle_{i\in I} o_i \setminus \bigcup_{j\in J} o_j$. Let b be any cocircuit of M, so that $|b \cap o_i|$ is even for each $i \in I$. Then $|b \cap \triangle_{i\in I} o_i|$ is also even, and in particular finite. Since the o_j are disjoint, there can only be finitely many of them that meet $b \cap \triangle_{i\in I} o_i$, and since for each such j we have that $|b \cap o_j|$ is even, it follows that $|b \cap w|$ is even. In particular, $b \cap w$ doesn't have just one element. Since b was arbitrary, by Lemma 11.3.5 w is a scrawl of M and so if it is nonempty it includes a circuit. But in that case, we could add that circuit to the family $(o_j|j \in J)$, contradicting the maximality of that family. Thus w is empty, and $\triangle_{i\in I} o_i = \bigcup_{j\in J} o_j$ is a disjoint union of circuits.

(5) implies (3): Suppose, for a contradiction, that (5) holds but (3) fails, and choose a circuit o and a cocircuit b with $o \cap b = \{x, y, z\}$ of size 3. Pick a base s of $(E \setminus b) + x$ including $o - y - z$, which exists by (IM). As b is a cocircuit, $b - x$ avoids some M-base, thus $(E \setminus b) + x$ is spanning and thus s is spanning, as well. Let o_y and o_z be the fundamental circuits of y and z with respect to s.

It suffices to show that $o_y \triangle o_z \subseteq o - x$. Indeed, since $y, z \in o_y \triangle o_z$, (5) then yields a circuit properly included in o, which is impossible. By Lemma 11.3.2 we can't have $o_y \cap b = \{y\}$ so we must have $x \in o_y$. Similarly, $x \in o_z$, and so $x \notin o_y \triangle o_z$. So it is sufficient to show that o_y and o_z agree outside o, in other words: $o_y \subseteq o_z \cup o$ and $o_z \subseteq o_y \cup o$.

To see this, first note that by uniqueness of the fundamental circuit of y it suffices to show that y is spanned by $(o_z - z) \cup (o - y - z)$. As z is spanned by $(o_z - z)$, $o - y$ is spanned by $(o_z - z) \cup (o - y - z)$. Since o is a circuit, y is also spanned by $(o_z - z) \cup (o - y - z)$, as desired. A similar argument yields $o_z \subseteq o_y \cup o$, completing the proof.

(3) implies (4): Since any subset of the ground set of $U_{2,4}$ of size 3 is both a circuit and a cocircuit, it is easy to find a circuit and cocircuit in $U_{2,4}$ whose intersection has size 3. So we simply apply Corollary 11.3.4.

(4) implies (2): Suppose for a contradiction that (4) holds but (2) does not. Then let o be a circuit and b a cocircuit such that $|o \cap b| = k$ is odd. By contracting $o \setminus b$ and deleting $b \setminus o$, we obtain a minor M' of M in which $o \cap b$ is both a circuit and a cocircuit. Let s be a minimal spanning set containing $o \cap b$, which exists by (IM^*). Then in the minor M'' of M' obtained by contracting

$s \setminus (o \cap b)$, $(o \cap b)$ is spanning, and is still both a circuit and a cocircuit. By a similar removal, we can find a minor M''' of M'' in which $o \cap b$ is a circuit and a cocircuit and is both spanning and cospanning. Let $x \in o \cap b$. Then $o \cap b - x$ is both a base and a cobase of M''', and it is finite (it has size $k-1$). As $o \cap b - x$ is a base and a cobase, the complement of $o \cap b - x$ is also a base and a cobase. Thus the ground set of M''' is also finite (it has size $2k-2$). Applying the finite version of the theorem, then, M''' contains a $U_{2,4}$ minor, which is also a minor of M, giving the desired contradiction.

(9) implies (2): first we will show that the following implies (2):

> For any base s of M, any circuit o meets every fundamental cocircuit of s in an even number of edges. $(\diamond)$

To see that $(\diamond)$ implies (2), it suffices to show that every cocircuit b is fundamental cocircuit of some base s. Let $e \in b$. Then as b is a cocircuit, $E \setminus (b - e)$ is spanning. Thus by (IM) there is a base s of $E \setminus (b - e)$, which clearly has b as fundamental cocircuit.

So it remains to see that (9) implies $(\diamond)$. By (9), $o = \triangle_{e \in o \setminus s} o_e$. Let b_f be some fundamental cocircuit of s for some $f \in s$. By Lemma 11.3.1 $o_e \cap b_f$ is empty or $o_e \cap b_f = \{e, f\}$. So it suffices to show that every f is in only finitely many o_e, which follows from the fact that $o = \triangle_{e \in o \setminus s} o_e$ is well defined at f. This completes the proof.

(2) implies (9): we have to show for every edge f that it is contained in only finitely many o_e and that $f \in o \iff f \in \triangle_{e \in o \setminus s} o_e(f)$. If $f \notin s$, this is easy, so let $f \in s$. By Lemma 11.3.1 $f \in o_e$ iff $e \in b_f$. As M is tame $|o \cap b_f|$ is finite, so there are only finitely many such e. By (2), $|o \cap b_f|$ is even. If $f \notin o$, all such e are not contained in s, so $f \notin \triangle_{e \in o \setminus s} o_e$. If $f \in o$, all such e but f are not contained in s, so $f \in \triangle_{e \in o \setminus s} o_e$. This completes the proof. $\square$

We remark that we might also put the duals of the statements in the list onto the list. It might be worth noting that (7) becomes false if we also allow I to be infinite. To see this, consider the finite cycle matroid of the graph obtained from a ray by adding a vertex that is adjacent to every vertex on the ray. Indeed, the symmetric difference of all 3-cycles is a ray starting at this new vertex. This set is not empty, and nor does it include a circuit, so the infinite version of (7) fails.

More generally, the finite cycles of a locally finite graph generate the cycle space, which may contain infinite cycles [36].

We offer the following related open questions. Let (10) be the statement like (9) but for only one base of M. For finite matroids, (10) is equivalent to (9). Is the same true for tame matroids?

The following simple question also remains open:

In Theorem 11.4.1, we assumed that M is tame. Without this assumption, the theorem is no longer true. For example, in [18] there is an example of a wild matroid satisfying (2-6) and (10), but not (1) or (7-9). However, this matroid is not a binary thin sums matroid. In fact, we still do not know the answer to the following:

Open Question 11.4.2. *Is every binary thin sums matroid tame?*

In a binary tame matroid, it is easy to see that any set meeting every cocircuit not in an odd number of edges is a disjoint union of circuits provided that the set is either countable or does not meet any cocircuit infinitely. A well-known result of Nash-Williams says that the above is also true if the matroid is the finite cycle matroid of some graph. Does this extend to all binary tame matroids?

Open Question 11.4.3. *Let M be a binary tame matroid and let X be a set that meets no cocircuit in an odd number of edges. Must X be a disjoint union of circuits?*

11.5 Excluded minors of representable matroids

In this section, we will prove the main result, Theorem 11.2.1. The proof will be by a compactness argument, but because we wish to prove the result for tame matroids rather than just finitary ones, we will need to go beyond the usual compactness arguments for finitary matroids in two ways. First, we need the characterisation in Lemma 11.3.9, since the definition of thin sums matroids is not suited to compactness arguments. Second, we need the following lemma, which allows us to move to a finite minor whilst preserving a finite amount of complexity in a tame matroid.

Lemma 11.5.1. *Let M be a tame matroid, O a finite set of circuits of M and B a finite set of cocircuits of M. Then there exists a finite minor N of M and functions $f\colon O \to \mathcal{C}(N)$ and $g\colon B \to \mathcal{C}(N^*)$ such that for any $o \in O$ and $b \in B$ we have $f(o) \cap g(b) = o \cap b$.*

Proof. We pick an element $e_o \in o$ for each $o \in O$ and an element $e_b \in b$ for each $b \in B$. Let $F = \bigcup_{o\in O} \bigcup_{b\in B} o \cap b \cup \{e_o | o \in O\} \cup \{e_b | b \in B\}$. Since M is tame, F is finite. Next, for each $o \in O$ and each $e \in o \cap F - e_o$ we pick a cocircuit $b_{o,e}$ with $o \cap b_{o,e} = \{e_o, e\}$ (this is possible by Lemma 11.3.2). Let B' be the set of all cocircuits picked in this way or contained in B. Note that B' is finite. Similarly, we pick for each $b \in B$ and each element e_b (which by construction is in $F \cap b$) a circuit $o_{b,e}$ with $o_{b,e} \cap b = \{e_b, e\}$ for each $e \in F \cap b - e_b$, and we collect all of these, together with all circuits contained in O, in a finite set O'.

Let $F' = F \cup (\bigcup O' \cap \bigcup B') = F \cup \bigcup_{o\in O'} \bigcup_{b\in B'} o \cap b$. Note that F' is also finite. Let $C = \bigcup O' \setminus F'$, and let $D = E \setminus (C \cup F')$. Thus $E = C \dot\cup F' \dot\cup D$. Let N be the finite minor of M with ground set F' that is given by $M/C\backslash D$. For each $o \in O$, $o \setminus F' \subseteq C$ and so $o \cap F'$ is a scrawl of N by Corollary 11.3.6. Let $f(o)$ be a circuit of N with $e_o \in f(o) \subseteq o \cap F'$. Then for each $e \in o \cap F - e_o$ we know that $F' \cap b_{o,e}$ is a scrawl of N^*, again by Corollary 11.3.6, so it can't meet $f(o)$ in just one point. But $e_o \in f(o) \cap F' \cap b_{o,e} \subseteq \{e_o, e\}$ so we must have $f(o) \cap F' \cap b_{o,e} = \{e_o, e\}$ and we conclude that $e \in f(o)$. Since e was arbitrary, this implies that $o \cap F \subseteq f(o)$.

Similarly, for each $b \in B$, we find a cocircuit $g(b)$ of N such that $g(b) \subseteq F' \cap b$ but $b \cap F \subseteq g(b)$. Thus for $o \in O$ and $b \in B$ we have $f(o) \cap g(b) = o \cap b$, as required. $\square$

By Lemma 11.3.9, Theorem 11.2.1 is equivalent to the following result.

Theorem 11.5.2. *Let M be a tame matroid and k be a finite field. Then M is k-paintable if and only if all of its finite minors are k-paintable.*

Proof. The 'only if' part was established in Section 11.3. For the 'if' part, we begin by defining the topological space whose compactness we will use. We would like an element of this space to correspond to a choice of functions $c_o \colon o \to k^*$ and $d_b \colon b \to k^*$ for each $o \in \mathcal{C}(M)$ and $b \in \mathcal{C}(M^*)$, so we take

$$H = \left(\bigcup_{o \in \mathcal{C}(M)} \{o\} \times o \right) \amalg \left(\bigcup_{b \in \ \mathcal{C}(M^*)} \{b\} \times b \right)$$

and take the underlying set of our space to be $X = (k^*)^H$ - the compact topology on X that we will use is given by the product of $|H|$ copies of the discrete topology on k^*.

For each circuit o and cocircuit b of M, the set

$$C_{o,b} = \left\{ c \in (k^*)^H \,\middle|\, \sum_{e \in o \cap b} c(o,e)c(b,e) = 0 \right\}$$

is closed because $o \cap b$ is finite. We shall now show that any finite intersection of such sets is nonempty. That is, we shall show that $\bigcap_{(o,b) \in K} C_{o,b} \neq \emptyset$ for every finite subset K of $\mathcal{C}(M) \times \mathcal{C}(M^*)$.

Let O be the set of circuits appearing as first components of elements of K, and let B be the set of cocircuits appearing as second components of elements of K. By Lemma 11.5.1, there are a finite minor N of M and functions $f \colon O \to \mathcal{C}(N)$ and $g \colon B \to \mathcal{C}(N^*)$ such that for any $o \in O$ and $b \in B$ we have $f(o) \cap g(b) = o \cap b$.

Since N is finite, it is k-paintable. So we can find functions $c_{f(o)} \colon f(o) \to k^*$ and $d_{g(b)} \colon g(b) \to k^*$ for all $o \in O$ and $b \in B$ such that $\sum_{e \in o \cap b} c_{f(o)}(e) d_{g(b)}(e) = 0$ for each such o and b. Let $c \in (k^*)^H$ be chosen so that, for each $o \in O$, $b \in B$ and $e \in o \cap b$ we have $c(o,e) = c_{f(o)}(e)$ and $c(b,e) = c_{g(b)}(e)$. These choices ensure that $c \in \bigcap_{(o,b) \in K} C_{o,b}$.

Since $(k^*)^H$ is compact, and any finite intersection of the $C_{o,b}$ is nonempty, we have that $\bigcap_{(o,b) \in \mathcal{C}(M) \times \mathcal{C}(M^*)} C_{o,b}$ is nonempty. As any element in the intersection gives a k-painting, this completes the proof. □

We note that this gives a uniform way to extend excluded minor characterisations of representability from finite to infinite matroids. For example, we may immediately extend the result of [12, 87] as follows:

Corollary 11.5.3. *A tame matroid M is a thin sums matroid over $GF(3)$ if and only if it has no minor isomorphic to $U_{2,5}$, $U_{3,5}$, F_7 or F_7^*.*

11.6 Other applications of the method

11.6.1 Regular matroids

A key definition to prove Theorem 11.5.2 was that of a k-painting. The corresponding notion for regular matroids is as follows.

A *signing* for a matroid M is a choice of a function $c_o \colon o \to \{1, -1\}$ for each circuit o of M and a function $d_b \colon b \to \{1, -1\}$ for each cocircuit b of M such that for any circuit o and cocircuit b we have

$$\sum_{e \in o \cap b} c_o(e) d_b(e) = 0\,,$$

where the sum is evaluated over $\mathbb{Z}$. A matroid is *signable* if it has a signing.

Lemma 11.6.1. *[[76, Proposition 13.4.5],[98]] Let M be a finite matroid. Then M is regular if and only if M is signable.*

Using similar ideas to those in the proof of Theorem 11.5.2, we obtain the following.

Theorem 11.6.2. *Let M be a tame matroid. Then the following are equivalent.*

1. *M is a thin sums matroid over every field.*
2. *M is signable*
3. *Every finite minor of M is regular.*

Proof. (2) implies that M is k-paintable for every field k, and so implies (1). (1) implies that every finite minor of M is representable over every field, and so is regular, which gives (3). (3) implies that every finite minor of M is signable, by Lemma 11.6.1. We may then deduce (2) by a compactness argument like that in the proof of Theorem 11.5.2. □

Motivated by this theorem, we call a tame matroid *regular* if any of these equivalent conditions hold.

11.6.2 Partial fields

Theorem 11.6.2 is a special case of a more general result extending characterisations of simultaneous representations over multiple fields using partial fields to tame infinite matroids. For some background on partial fields, see [95].

A *partial field* consists of a pair (R, S), where R is a ring and S is a subgroup of the group of units of R under multiplication, such that $-1 \in S$. In this context, an (R, S)*-painting* for a matroid M is a choice of a function $c_o \colon o \to S$ for each circuit o of M and a function $d_b \colon b \to S$ for each cocircuit b of M such that for any circuit o and cocircuit b we have

$$\sum_{e \in o \cap b} c_o(e) d_b(e) = 0\,. \tag{11.2}$$

For example, for any field k a matroid M is k-paintable if and only if it is (k, k^*)-paintable, and M is signable if and only if it is $(\mathbb{Z}, \{-1, 1\})$-paintable. It is clear that the class of (R, S)-paintable matroids is closed under duality and under taking minors. In particular, any finite minor of an (R, S)-paintable matroid is (R, S)-paintable. The converse follows from an almost identical compactness argument to that used for Theorem 11.5.2, giving:

Theorem 11.6.3. *Let (R, S) be a partial field with S finite. A tame matroid is (R, S)-paintable if and only if all its finite minors are.*

It follows from the results of [95, Section 2.7] that a finite matroid is (R, S)-paintable if and only if it is (R, S)-representable. For finite matroids it is known that simultaneous representability over sets of fields corresponds to representability over partial fields, and we are now in a position to lift many such results to all tame matroids. For example, we can lift [101, Theorem 1.2] as follows:

Corollary 11.6.4. *A tame matroid M is a thin sums matroid over both $\mathbb{F}_3$ and $\mathbb{F}_4$ if and only if it is $(\mathbb{C}, \{\zeta^i | i \leq 6\})$-paintable for ζ a primitive sixth root of unity.*

11.6.3 Ternary matroids

For finite matroids, a useful property of $\mathbb{F}_3$-representable matroids is the uniqueness of the representations. In this section, we shall prove the corresponding property for tame ternary matroids.

Let M be a k-paintable matroid for some field k. We say that two k-paintings $((c_o | o \in \mathcal{C}(M)), (d_b | b \in \mathcal{C}(M^*)))$ and $((\tilde{c}_o | o \in \mathcal{C}(M)), (\tilde{d}_b | b \in \mathcal{C}(M^*)))$ are *equivalent* if and only if there are constants $x(o)$ for every $o \in \mathcal{C}(M)$, constants $x(b)$ for every $b \in \mathcal{C}(M^*)$, constants $x(e)$ for every edge e and a field automorphism φ such that the following are true:

1. $\tilde{c}_o(e) = \varphi(x(o)x(e)c_o(e))$ for any $e \in o \in \mathcal{C}(M)$.
2. $\tilde{d}_b(e) = \varphi\left(\frac{x(b)d_b(e)}{x(e)}\right)$ for any $e \in b \in \mathcal{C}(M^*)$.

Two signings of the same matroid M are *equivalent* if and only if they induce equivalent $\mathbb{F}_3$-paintings of M.

Via Lemma 11.3.9 for any tame matroid any thin sums representation over k corresponds to a k-painting. For finite matroids, the notions of equivalence for representations and paintings coincide: it is straightforward to check that two representations are equivalent iff the corresponding paintings are. As for finite matroids, we obtain the following.

Theorem 11.6.5. *Any two $\mathbb{F}_3$-paintings of the same matroid M are equivalent.*

Proof. M, being $\mathbb{F}_3$-paintable, must be tame. Without loss of generality we may also assume that M is connected and has more than one edge. Thus any edges

e and f of M lie on a common circuit[1]. We nominate a particular edge g_1, and for each other edge g we nominate a circuit $o(g)$ containing both g_1 and g. We also nominate for each circuit o of M an edge $e(o) \in o$ and for each cocircuit b of M an edge $e(b) \in b$.

We denote the two $\mathbb{F}_3$-paintings $((c_o|o \in \mathcal{C}(M)), (d_b|b \in \mathcal{C}(M^*)))$ and $((\tilde{c}_o|o \in \mathcal{C}(M)), (\tilde{d}_b|b \in \mathcal{C}(M^*)))$. We shall construct witnesses to the equivalence as in the definition above. Since every automorphism of $\mathbb{F}_3$ is trivial, we shall take φ to be the identity.

We now set $x(g) = \frac{\tilde{c}_{o(g)}(g)c_{o(g)}(g_1)}{\tilde{c}_{o(g)}(g_1)c_{o(g)}(g)}$ for each $g \in E$, $x(o) = \frac{\tilde{c}_o(e(o))}{x(e(o))c_o(e(o))}$ for each circuit o of M and $x(b) = \frac{x(e(b))\tilde{d}_b(e(b))}{d_b(e(b))}$ for each cocircuit b of M.

In order to prove that these values satisfy (1) at a particular circuit o and $g \in o$, let $O = \{o, o(g), o(e(o))\}$ and $F = \{g, g_1, e(o)\}$ and use the construction from the proof of Lemma 11.5.1 to obtain a finite minor $M' = M/C \backslash D$ such that for every $o \in O$ there is an M'-circuit $o' \subseteq o$ such that $o' \cap F = o \cap F$ and for every $b \in B$ there is an M'-cocircuit $b' \subseteq b$ such that $b' \cap F = b \cap F$.

Let $((c'_o|o \in \mathcal{C}(M')), (d'_b|b \in \mathcal{C}(M'^*))$ be the $\mathbb{F}_3$-painting of M' induced by $((c_o|o \in \mathcal{C}(M)), (d_b|b \in \mathcal{C}(M^*))$, and $((\tilde{c}'_o|o \in \mathcal{C}(M')), (\tilde{d}'_b|b \in \mathcal{C}(M'^*))$ that induced by $((\tilde{c}_o|o \in \mathcal{C}(M)), (\tilde{d}_b|b \in \mathcal{C}(M^*)))$.

By uniqueness of representation for finite matroids, we can find constants $x'(o')$ for every $o' \in \mathcal{C}(M')$, constants $x'(b')$ for every $b' \in \mathcal{C}(M'^*)$ and constants $x'(g)$ for every $g \in X$ such that

3. $\tilde{c}'_{o'}(g) = x'(o')x'(g)c'_{o'}(g)$ for any $g \in o' \in \mathcal{C}(M')$.
4. $\tilde{d}'_{b'}(g) = \frac{x'(b')d'_{b'}(g)}{x'(g)}$ for any $g \in b' \in \mathcal{C}(M'^*)$.

Lemma 11.6.6. *For each $o \in O$ there is $\lambda_o \in k^*$ such that*

5. $c_o\restriction_F = \lambda_o c'_{o'}\restriction_F$

Proof. As part of the construction of M', we picked a canonical element e_o of o'. Let $\lambda = \frac{c_o(e_o)}{c'_{o'}(e_o)}$. For any other $e \in o' \cap F$, there is by construction a cocircuit $b_{o,e}$ of M with $o \cap b_{o,e} = \{e_o, e\}$. Then by the dual of Corollary 11.3.6 $b_{o,e} \cap E(M')$ is a coscrawl of M', and so there is a cocircuit b' of M' with $e_o \in b' \subseteq b_{o,e}$, and so $e_o \in o' \cap b' \subseteq \{e_o, e\}$. Since o' and b' can't meet in only one element, $e \in b'$. Since the painting of M' is induced from that of M, there is a cocircuit b of M such that $d_b(e) = d'_{b'}(e)$ for all $e \in E(M')$ and $b' \subseteq b \subseteq b' \cup D$, and so $o \cap b = \{e_o, e\}$. Using the identities in the definition of painting, we deduce that

$$c_o(e_o)d_b(e_o) + c_o(e)d_b(e) = 0 \quad \text{and} \quad c'_{o'}(e_o)d'_{b'}(e_o) + c'_{o'}(e)d'_{b'}(e) = 0$$

and so

$$c_o(e) = -\frac{c_o(e_o)d_b(e_o)}{d_b(e)} = -\frac{\lambda c'_{o'}(e_o)d'_{b'}(e_o)}{d'_{b'}(e)} = \lambda c'_{o'}(e)$$

which gives the desired result, since $e \in o' \cap F$ was arbitrary. □

[1] In Section 3 of [23], it is shown that the relation 'e is in a common circuit with f' is indeed an equivalence relation for infinite matroids.

Similarly, we can find constants $\tilde{\lambda}_o$ for each $o \in O$ such that

6. $\tilde{c}_o\restriction_F = \tilde{\lambda}_o \tilde{c}'_{o'}\restriction_F$

Now we must simply unwind all the algebraic relationships to obtain the desired result.

$$x(g) = \frac{\tilde{c}_{o(g)}(g)c_{o(g)}(g_1)}{\tilde{c}_{o(g)}(g_1)c_{o(g)}(g)} = \frac{\tilde{c}'_{o(g)'}(g)c'_{o(g)'}(g_1)}{\tilde{c}'_{o(g)'}(g_1)c'_{o(g)'}(g)} = \frac{x'(o(g)')x'(g)}{x'(o(g)')x'(g_1)} = \frac{x'(g)}{x'(g_1)}$$

where the first equation follows from the definitions, the second from (5) and (6) and the third from (3). Similarly, we get:

$$x(o) = \frac{\tilde{c}_o(e(o))}{x(e(o))c_o(e(o))} = \frac{\tilde{\lambda}_o}{\lambda_o}\frac{\tilde{c}'_{o'}(e(o))}{x(e(o))c'_{o'}(e(o))} = \frac{\tilde{\lambda}_o}{\lambda_o}\frac{x'(o')x'(e(o))}{x(e(o))}$$

And finally:

$$x(o)x(g)c_o(g) = \frac{\tilde{\lambda}_o}{\lambda_o}\frac{x'(o')x'(e(o))x'(g_1)}{x'(e(o))}\frac{x'(g)}{x'(g_1)}c_o(g) = \frac{\tilde{\lambda}_o}{\lambda_o}x'(o')x'(g)c_o(g)$$

Now the last term is just $\tilde{c}_o(g)$ by first applying (5) and then (3). This completes the proof of the above assignment satisfies (1). The proof that it also satisfies (2) is similar.

□

As every tame regular matroid is a thin sums matroid over $\mathbb{F}_3$, it also has a unique representation. In particular the finite cycle matroid, the algebraic cycle matroid and the topological cycle matroid of a given graph (and their duals) have a unique signing.

In what follows, we will describe this signing of the finite cycle matroid of a given graph G — the other cases are similar. First direct the edges of G in an arbitrary way. To define the functions c_o, let o be some cycle of G. Pick a cyclic order of o. For $e \in o$, let $c_o(e) = 1$ if e is directed according to the cyclic order of o and -1 otherwise.

Next, let b be some cocircuit. By minimality of the cocircuit, it is contained in a single component of G and its removal separates this component into two components, say $C_1(b)$ and $C_2(b)$. Note that every edge in b has precisely one endvertex in each of these components. For $e \in b$, let $d_b(e) = 1$ if e points to a vertex in C_1 and -1 otherwise.

It remains to check that $\sum_{e\in o\cap b} c_o(e)d_b(e) = 0$ for all circuits o and cocircuits b. As every circuit is finite, the above sum is finite. Since the directions we gave to the edges of G do not influence the values of the products $c_o(e)d_b(e)$, we may assume without loss of generality that in the bond b all edges are directed from $C_1(b)$ to $C_2(b)$. So we get a summand of $+1$ for each edge along which o

traverses b from $C_1(b)$ to $C_2(b)$ and a summand of -1 for each edge along which o traverses b from $C_2(b)$ to $C_1(b)$. Since o must traverse b the same number of times in each direction, the sum evaluates to 0.

Let us look at how to modify the above construction to make it work for the algebraic cycle matroid and the topological cycle matroid instead. Finite circuits in the algebraic cycle matroid may be dealt with as before. To define c_o for a double ray o, we pick an orientation of o and let $c_o(e)$ be 1 if e is directed in agreement with this orientation and -1 otherwise. The above argument still applies: using the tameness of the algebraic cycle matroid, we obtain that a double ray can cross a skew cut only finitely many times, and both tails of the double ray must lie on the same side (as one side is rayless), so the double ray must cross the skew cut the same number of times in each direction.

Using the fact that topological circles are homeomorphic to the unit circle, we get a cyclic order on each circuit of the topological cycle matroid and the above construction again gives us a signing.

Chapter 12

Matroids with an infinite circuit-cocircuit intersection

12.1 Abstract

We construct some matroids that have a circuit and a cocircuit with infinite intersection.

This answers a question of Bruhn, Diestel, Kriesell, Pendavingh and Wollan. It further shows that the axiom system for matroids proposed by Dress in 1986 does not axiomatize all infinite matroids.

We show that one of the matroids we define is a thin sums matroid whose dual is not a thin sums matroid, answering a natural open problem in the theory of thin sums matroids.

12.2 Introduction

In [22], Bruhn, Diestel, Kriesell, Pendavingh and Wollan introduced axioms for infinite matroids in terms of independent sets, bases, circuits, closure and (relative) rank. These axioms allow for duality of infinite matroids as known from finite matroid theory, which settled an old problem of Rado. Unlike the infinite matroids known previously, such matroids can have infinite circuits or infinite cocircuits. Many infinite matroids are *finitary*, that is, every circuit is finite, or *cofinitary*, that is, every cocircuit is finite, but nontrivial matroids with both infinite circuits and infinite cocircuits have been known for some time [22, 58].

However in all the known examples, all intersections of circuit with cocircuit are finite. Moreover, this finiteness seems to be a natural requirement in many theorems [1, 14]. This phenomenon prompted the authors of [22] to ask the following.

Question 12.2.1 ([22])**.** *Is the intersection of a circuit with a cocircuit in an infinite matroid always finite?*

Dress [37] even thought that the very aim to have infinite matroids with duality, as in Rado's problem, would make it necessary that circuit-cocircuit intersection were finite. He therefore proposed axioms for infinite matroids which had the finiteness of circuit-cocircuit intersections built into the definition of a matroid, in order to facilitate duality.

And indeed, it was later shown by Wagowski [97] that the axioms proposed by Dress capture all infinite matroids as axiomatised in [22] if and only if Question 12.2.1 has a positive answer. We prove that the assertion of Question 12.2.1 is false and consequently that the axiom system for matroids proposed by Dress does not capture all matroids.

We call a matroid *tame* if the intersection of any circuit with any cocircuit is finite, and otherwise *wild*.

Theorem 12.2.2. *There exists a wild matroid.*

To construct such matroids M, we use some recent result from an investigation of matroid union [6]. We later became aware that Matthews and Oxley had constructed some matroids similar to ours by means of a more involved construction in [70], though they did not have the distinction between tame and wild matroids in mind.

We hope that the wild matroids we construct here may be sufficiently badly behaved to serve as generic counterexamples also for other open problems. To illustrate this potential, we shall show that we do obtain a counterexample to a natural open question about thin sums matroids, a generalisation of representable matroids.

If we have a family of vectors in a vector space, we get a matroid structure on that family whose independent sets are given by the linearly independent subsets of the family. Matroids arising in this way are called *representable* matroids. Although many interesting finite matroids (eg. all graphic matroids) are representable, it is clear that any representable matroid is finitary and so many interesting examples of infinite matroids are not of this type. However, since the construction of many of these examples, including the algebraic cycle matroids of infinite graphs, is suggestively similar to that of representable matroids, the notion of *thin sums matroids* was introduced in [21]: it is a generalisation of representability which captures these infinite examples.

Since thin sums matroids need not be finitary, and the duals of many thin sums matroids are again thin sums matroids, it is natural to ask whether the class of thin sums matroids itself is closed under duality. It is shown in [1] that the class of tame thin sums matroids is closed under duality, so that any counterexample must be wild. We show below that one of the wild matroids we have constructed does give a counterexample.

Theorem 12.2.3. *There exists a thin sums matroid whose dual is not a thin sums matroid.*

The chapter is organised as follows. In Section 2, we recall some basic matroid theory. After this, in Section 3, we give the first example of a wild matroid. In Section 4, we give a second example, which is obtained by taking the union of a matroid with itself. In Section 5, we show that the class of thin sums matroids is not closed under duality by constructing a suitable wild thin sums matroid whose dual is not a thin sums matroid.

12.3 Preliminaries

Throughout, notation and terminology for graphs are that of [36], for matroids that of [76, 22]. A set system $\mathcal{I}$ is the set of independent sets of a matroid if it satisfies the following *independence axioms* [22].

(I1) $\emptyset \in \mathcal{I}$.

(I2) $\mathcal{I}$ is closed under taking subsets.

(I3) Whenever $I, I' \in \mathcal{I}$ with I' maximal and I not maximal, there exists an $x \in I' \setminus I$ such that $I + x \in \mathcal{I}$.

(IM) Whenever $I \subseteq X \subseteq E$ and $I \in \mathcal{I}$, the set $\{I' \in \mathcal{I} \mid I \subseteq I' \subseteq X\}$ has a maximal element.

M always denotes a matroid and $E(M)$, $\mathcal{I}(M)$, $\mathcal{B}(M)$, $\mathcal{C}(M)$ and $\mathcal{S}(M)$ denote its ground set and its sets of independent sets, bases, circuits and spanning sets, respectively. A matroid is called *finitary* if every circuit is finite.

In our constructions, we will make use of algebraic cycle matroid $M_A(G)$ of a graph G. The circuits of $M_A(G)$ are the edge sets of finite cycles of G and the edge sets of double rays[1]. If G is locally finite, then $M_A(G)$ is *cofinitary*, that is, its dual is finitary [6]. If G is not locally finite, then this is no longer true [22]. Higgs [58] characterized those graphs G that have an algebraic cycle matroid, that is, whose finite circuits and double rays from the circuits of a matroid: G has an algebraic cycle matroid if and only if G does not contain a subdivision of the Bean-graph, see Figure 12.1.

12.4 First construction: the matroid M^+

In this example, we will need the following construction from [6], where it is also shown that this construction gives a matroid:

Definition 12.4.1. (Truncation) Let M be a matroid, in which $\emptyset$ isn't a base. Then the matroid M^-, on the same groundset, is that whose bases are those obtained by removing a point from a base of M. That is, $\mathcal{B}(M^-) = \{B - e | B \in \mathcal{B}(M), e \in B\}$. Dually, if M is a matroid whose ground set E isn't a base, we define M^+ by $\mathcal{B}(M^+) = \{B + e | B \in \mathcal{B}(M), e \in E \setminus B\}$.

[1] A *double ray* is a two sided infinite path

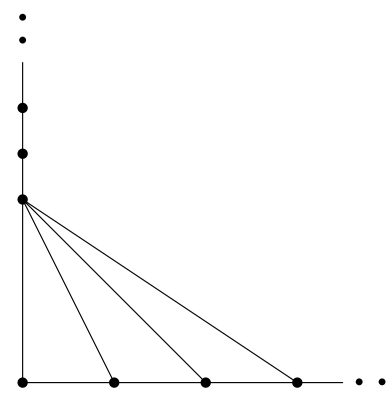

Figure 12.1: The Bean-graph

Thus $(M^+)^* = (M^*)^-$.

We shall show that the matroids constructed in this way are very often wild.

Since M^- is obtained from M by making the bases of M into dependent sets, we may expect that $\mathcal{C}(M^-) = \mathcal{C}(M) \cup \mathcal{B}(M)$: that is, the set of circuits of M^- contains exactly the circuits and the bases of M. This is essentially true, but there is one complication: an M-circuit might include an M-base, which would prevent it being an M^--circuit. Let O be a circuit of M^-. If O is M-independent, it is clear that O must be an M-base. Conversely, any M-base is a circuit of M^-. If O is M-dependent, then since all proper subsets of O are M^--independent and so M-independent, O must be an M-circuit. Conversely, an M-circuit not including an M-base is an M^--circuit.

On the other hand, none of the circuits of M is a circuit of M^+: for any circuit O of M, pick any $e \in O$ and extend $O - e$ to a base B of M. Then $O \subseteq B + e$, so $O \in \mathcal{I}(M^+)$. In fact, a circuit of M^+ is a set minimal with the property that at least two elements must be removed before it becomes M-independent. To see this note that the independent sets of M^+ are those sets from which an M-independent set can be obtained by removing at most one element.

Now we are in a position to construct a wild matroid: let M be the algebraic cycle matroid of the graph in Figure 12.2. Then the dashed edges form a circuit in M^+, and the bold edges form a circuit in $(M^+)^* = (M^*)^-$ (they form a base in M^* since their complement forms a base in M). The intersection, consisting of the dotted bold edges, is evidently infinite.

For the remainder of this section, we will generalize this example to construct a large class of wild matroids. To do so, we first have a closer look at the circuits of M^+. It is clear that if M is the finite cycle matroid of a graph G, then we get as circuits of M^+ any subgraphs which are subdivisions of those in Figure 12.3.

More generally, we can make precise a sense in which every circuit of M^+ is obtained by sticking together two circuits.

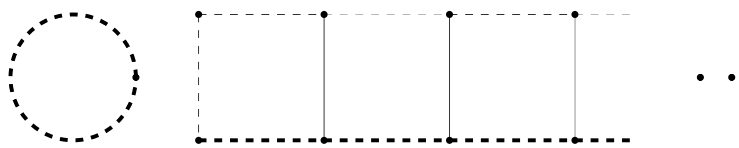

Figure 12.2: A circuit and a cocircuit with infinite intersection

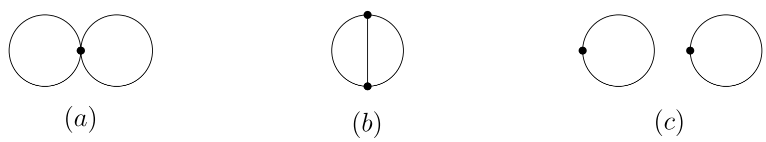

Figure 12.3: Shapes of circuits in M^+, with M a finite cycle matroid

Lemma 12.4.2. *Let O be a circuit of M, and $I \subseteq E(M) \setminus O$. Then $O \cup I$ is M^+-independent iff I is M/O-independent.*

Proof. If: Extend I to a base B of M/O. Pick any $e \in O$. Then $B' = B \cup O - e$ is a basis of M and $O \cup I \subseteq B' + e$.

Only if: Pick B a base of M and $e \in E \setminus B$ such that $O \cup I \subseteq B \cup e$. Since O is dependent, we must have $e \in O$, and so $I \subseteq B \setminus O$. Finally, $B \setminus O$ is a base of M/O, since $B \cap O = O - e$ is a base of O. □

Lemma 12.4.3. *Let O_1 be a circuit of M, and O_2 a circuit of M/O_1. Then $O_1 \cup O_2$ is a circuit of M^+. Every circuit of M^+ arises in this way.*

Proof. $O_1 \cup O_2$ is M^+-dependent by Lemma 12.4.2. Next, we shall show that any set $O_1 \cup O_2 - e$ obtained by removing a single element from $O_1 \cup O_2$ is M^+-independent, and so that $O_1 \cup O_2$ is a *minimal* dependent set (a circuit) in M^+. The case $e \in O_2$ is immediate by Lemma 12.4.2. If $e \in O_1$, then we pick any $e' \in O_2$. Now extend $O_2 - e'$ to a base B of M/O_1. Then $B' = B \cup O_1 - e$ is a base of M and $O_1 \cup O_2 - e \subseteq B' + e'$.

Finally, we need to show that any circuit O of M^+ arises in this way. O must be M-dependent, and so we can find a circuit $O_1 \subseteq O$ of M. Let $O_2 = O \setminus O_1$: O_2 is a circuit of M/O_1 by Lemma 12.4.2. □

Corollary 12.4.4. *Any union of two distinct circuits of M is dependent in M^+.*

It follows from Lemma 12.4.3 that the subgraphs of the types illustrated in Figure 12.3 give all of the circuits of M^+ for M a finite cycle matroid. Similarly, subdivisions of the graphs in Figure 12.3 and Figure 12.4 give circuits in the algebraic cycle matroid of a graph.

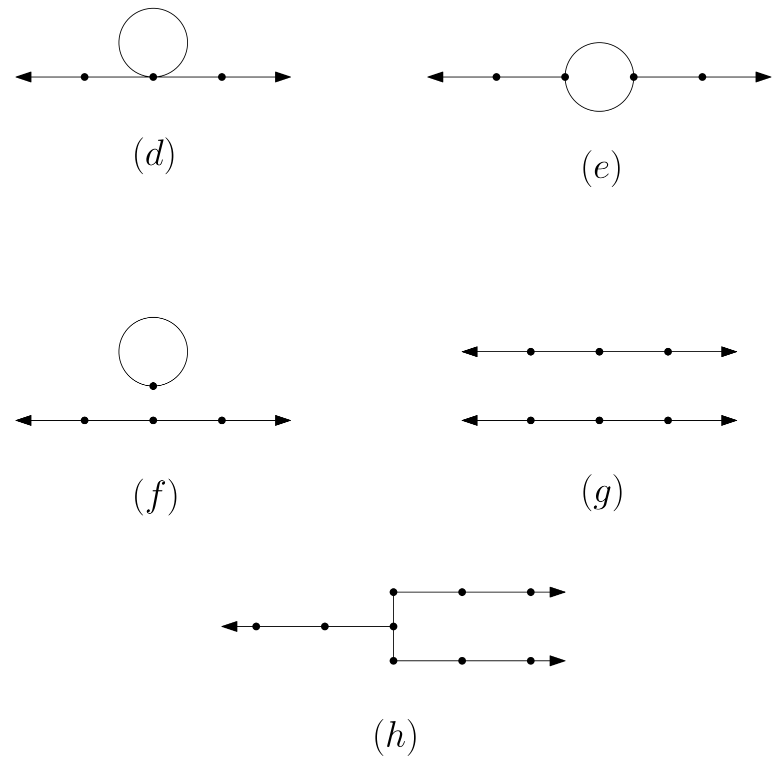

Figure 12.4: Shapes of circuits in M^+, with M an algebraic cycle matroid

Now that we have a good understanding of the circuits of matroids constructed this way, we can find many matroids M such that M^+ is wild.

Theorem 12.4.5. *Let M be a matroid with a base B and a circuit O such that $O \setminus B$ is infinite. Then M^+ is wild.*

Proof. Let $e \in O \setminus B$, and let O' be the fundamental circuit of e with respect to B. As $O' \setminus O$ is dependent in M/O, there is an M/O-circuit O'' included in O'. By Lemma 12.4.3, $O \cup O''$ is an M^+-circuit.

Since $E \backslash B$ is an M^*-base, it is a circuit of $(M^*)^- = (M^+)^*$. Now $(O \cup O'') \cap (E \backslash B)$ includes $O \setminus B$ and so it is infinite.

□

12.5 Second construction: matroid union

The *union of two matroids* $M_1 = (E_1, \mathcal{I}_1)$ *and* $M_2 = (E_2, \mathcal{I}_2)$ is the pair $(E_1 \cup E_2, \mathcal{I}_1 \vee \mathcal{I}_2)$, where

$$\mathcal{I}_1 \vee \mathcal{I}_2 := \{I_1 \cup I_2 \mid I_1 \in \mathcal{I}_1,\ I_2 \in \mathcal{I}_2\}$$

The *finitarization* M^{fin} *of a matroid* M is the matroid whose circuits are precisely the finite circuits of M. In [6] it is shown that M^{fin} is always a matroid. Note that every base of M^{fin} contains some base of M and conversely every base of M is contained in some base of M^{fin}. A matroid M is called *nearly finitary* if for every base of M, it suffices to add finitely many elements to that base to obtain some base of M^{fin}. It is easy to show that M is nearly finitary if and only if for every base of M^{fin} it suffices to delete finitely many elements from that base to obtain some base of M.

The main tool for this example is the following theorem.

Theorem 12.5.1 ([6]). *The union of two nearly finitary matroids is a matroid, and in fact nearly finitary.*

Note that there are two matroids whose union is not a matroid [5].

One can also define M^+ using matroid union: $M^+ = M \vee U_{1,E(M)}$. Here $U_{1,E(M)}$ is the matroid with groundset $E(M)$, whose bases are the 1-element subsets of $E(M)$. In this section, we will obtain a wild matroid as union of some non-wild matroid M with itself.

Let us start constructing M. We obtain the graph H from the infinite one-sided ladder L by doubling every edge, see Figure 12.5.

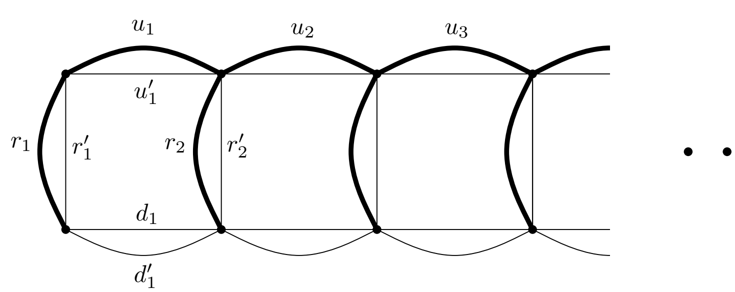

Figure 12.5: The graph H

As in the figure, we fix the following notation for the edges of H: In L, call the edges on the upper side of the ladder $u_1, u_2, \ldots$, the edges on the lower side $d_1, d_2, \ldots$ and the rungs $r_1, r_2, \ldots$. For every edge e of L, call its clone e'.

Let $M_A(H)$ be the algebraic cycle matroid of H. Note that $M_A(H)$ is a matroid by the results mentioned in the Preliminaries. Now we define M as the union of $M_A(H)$ with itself. To show that M is a matroid, by Theorem 12.5.1 it suffices to show the following.

Lemma 12.5.2. *$M_A(H)$ is nearly finitary.*

Proof. First note that the finitarization of $M_A(H)$ is the finite cycle matroid $M_F(H)$, whose circuits are the finite cycles of H. To see that $M_A(H)$ is nearly finitary, it suffices to show that each base B of $M_F(H)$ contains at most one double ray. It is easy to see that a double ray R of H contains precisely one rung r_i or r'_i. From this rung onwards, R contains precisely one of u_j or u'_j for $j \geq i$ and one of d_j or d'_j for $j \geq i$. Let R and S be two distinct double rays with unique rung edges e_R and e_S. Wlog assume that the index of e_R is less or equal than the index of e_S. Then already $R + e_S$ contains a finite circuit, which consists of e_R, e_S and all edges of R with smaller index than that of e_S. So each base B of $M_F(H)$ contains at most one double ray, proving the assumption. $\square$

Having proved that $M \vee M$ is a matroid, we next prove that it is wild.

Theorem 12.5.3. *The matroid $M \vee M$ is wild.*

To prove this, we will construct a circuit C and a cocircuit D with infinite intersection. Let us start with C, which we define as the set of all horizontal edges in Figure 12.5 together with the rung r_1.

Lemma 12.5.4. *$C := \{u_i, u'_i, d_i, d'_i | i = 1, 2, \ldots\} + r_1$ is a circuit of M.*

Proof. First, we show that C is dependent. To this end, it suffices to show that $C - r_1 = \{u_i, u'_i, d_i, d'_i | i = 1, 2, \ldots\}$ is a basis of M. As $I_1 = \{u_i, d_i | i = 1, 2, \ldots\}$ and $I_2 = \{u'_i, d'_i | i = 1, 2, \ldots\}$ are both independent in $M_A(H)$, their union $C - r_1$ is independent in M. All other representations $C - r_1 = I_1 \cup I_2$ with $I_1, I_2 \in \mathcal{I}(M_A(H))$ are the upper one up to exchanging parallel edges since from u_i and u'_i precisely one is I_1 and the other is in I_2. Similarly, the same is true for d_i and d'_i. So $C - r_1$ is a base and C is dependent, as desired.

It remains to show that $C - e$ is independent for every $e \in C$. The case $e = r_1$, was already consider above. By symmetry, we may else assume that $e = u_i$. Then $C - u_i = I_1 \cup I_2$ where $I_1 = \{u_i, d_i | i = 1, 2, \ldots\} - u_i + r_1$ and $I_2 = \{u'_i, d'_i | i = 1, 2, \ldots\}$ and I_1 and I_2 are both independent in $M_A(H)$, proving the assumption. $\square$

Next we turn to D, drawn bold in Figure 12.5.

Lemma 12.5.5. *$D := \{u_i, r_i | i = 1, 2, \ldots\}$ is a cocircuit of M.*

Proof. To this end, we show that $E \setminus D$ is a hyperplane, that is, $E \setminus D$ is non-spanning and $E \setminus D$ together with any edge is spanning in M. To see that $E \setminus D$ is non-spanning, we properly cover it by the following two bases B_1 and B_2 of $M_A(H)$, see Figure 12.6. Formally, $B_1 := \{d_i | i = 1, 2, \ldots\} \cup \{r'_i | i \text{ odd}\} \cup \{u'_i | i \text{ odd}\}$, $B_2 := \{d'_i | i = 1, 2, \ldots\} \cup \{r'_i | i \text{ even}\} \cup \{u'_i | i \text{ even}\} + r_1$.

To see that $E \setminus D$ together with any edge is spanning in M, we even show that $E \setminus D$ together with any edge is a base of M. This is done in two steps: first we show that $E \setminus D$ together with any edge e is independent in M and then that $E \setminus D$ together with any two edges is dependent in M. Concerning the first

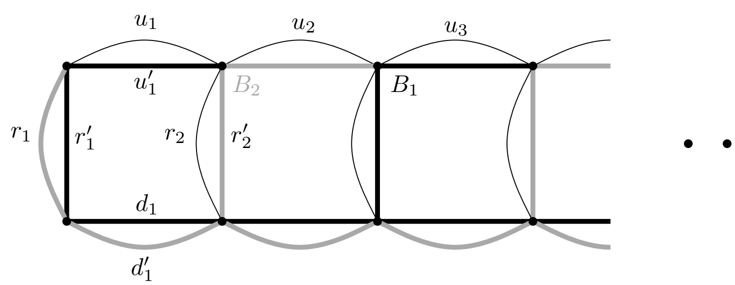

Figure 12.6: The two bases B_1 and B_2 properly cover $E \setminus D$

assertion, we distinguish between the cases $e = u_n$ for some n and $e = r_n$ for some n. In both cases we assume that n is odd. If n is even, then the argument is similar. In both cases we will cover $E \setminus D + e$ with two bases of $M_A(H)$, which arise from a slight modification of B_1 and B_2, see Figures 12.7 and 12.8.

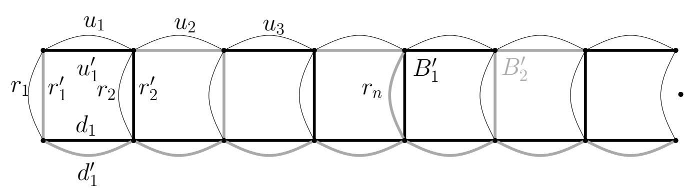

Figure 12.7: The two bases B'_1 and B'_2 cover $E \setminus D + r_n$

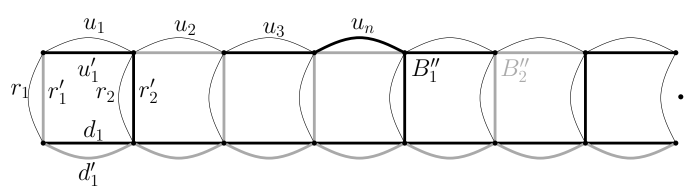

Figure 12.8: The two bases B''_1 and B''_2 cover $E \setminus D + u_n$

In the first case the bases are

$$B'_1 := B_1 \setminus \{r'_i | i < n \text{ and odd}\} \cup \{r'_i | i < n \text{ and even}\},$$

$$B'_2 := B_2 \cup \{r'_i | i < n \text{ and odd}\} \setminus \{r'_i | i < n \text{ and even}\} - r_1 + r_n$$

In the second case, the bases arise from B'_1 and B'_2 as follows:

$$B''_1 := B'_1 + u_n - r'_{n-1}, \quad B''_1 := B'_1 + r'_{n-1} - r_n$$

Having shown that $E \setminus D + e$ is independent for every $e \in D$, it remains to show for any two $e_1, e_2 \in D$ that $E \setminus D + e_1 + e_2$ cannot be covered by two bases of $M_A(H)$. In fact we prove the slightly stronger fact that $E \setminus D + e_1 + e_2$ cannot be covered by two bases of $M_F(H)$, that is by two spanning trees T_1 and T_2 of H.

Let H_n be the subgraph of H consisting of those $2n$ vertices that have the least distance to r_1. Choose n large enough so that $e_1, e_2 \in H_n$. An induction argument shows that $E \setminus D$ has $4n - 3$ edges in H_n since $E \setminus D$ has 1 edge in H_1 and $E \setminus D$ has 4 edges is $H_n \setminus H_{n-1}$. On the other hand, $T_1 \cup T_2$ can have at most $2(2n-1)$ edges in H_n since H_n has $2n$ vertices. This shows that T_1 and T_2 cannot cover $E \setminus D + e_1 + e_2$ because they cannot cover $H_n \setminus D + e_1 + e_2$. So for any $e \in D$ the set $E \setminus D + e$ is a base of M, proving the assumption. $\square$

As $|C \cap D| = \infty$, this completes the proof of Theorem 12.5.3.

In the previous section, we were able to generalise our example and give a necessary condition under which M^+ is wild. Here, we do not see a way to do this, because the description of C and D made heavy use of the structure of M. It would be nice to have a large class of matroids M, as in the previous section, such that $M \vee M$ is wild.

Open Question 12.5.6. *For which matroids M is $M \vee M$ wild?*

12.6 A thin sums matroid whose dual isn't a thin sums matroid

The constructions introduced so far give us examples of matroids which are wild, and so badly behaved. We therefore believe they will be a fruitful source of counterexamples in matroid theory. In this section, we shall illustrate this by giving a counterexample for a very natural question.

First we recall the notion of a thin sums matroid.

Definition 12.6.1. Let A be a set, and k a field. Let $f = (f_e | e \in E)$ be a family of functions from A to k, and let $\lambda = (\lambda_e | e \in E)$ be a family of elements of k. We say that λ is a *thin dependence* of f iff for each $a \in A$ we have

$$\sum_{e \in E} \lambda_e f_e(a) = 0\,,$$

where the equation is taken implicitly to include the claim that the sum on the left is well defined, that is, that there are only finitely many $e \in E$ with $\lambda_e f_e(a) \neq 0$.

We say that a subset I of E is *thinly independent* for f iff the only thin dependence of f which is 0 everywhere outside I is $(0|e \in E)$. The *thin sums system* M_f of f is the set of such thinly independent sets. This isn't always the set of independent sets of a matroid [22], but when it is we call it the *thin sums matroid* of f.

This definition is deceptively similar to the definition of the representable matroid corresponding to f considered as a family of vectors in the k-vector space k^A. The difference is in the more liberal definition of dependence: it is possible for λ to be a thin dependence even if there are infinitely many $e \in E$ with $\lambda_e \neq 0$, provided that for each $a \in A$ there are only finitely many $e \in E$ such that *both* $\lambda_e \neq 0$ and $f_e(a) \neq 0$.

Indeed, the notion of thin sums matroid was introduced as a generalisation of the notion of representable matroid: every representable matroid is finitary, but this restriction does not apply to thin sums matroids. Thus, although it is clear that the class of representable matroids isn't closed under duality, the question of whether the class of thin sums matroids is closed under duality remained open. It is shown in [1] that the class of tame thin sums matroids is closed under duality, so that any counterexample must be wild. We show below that one of the wild matroids we have constructed does give a counterexample.

There are many natural examples of thin sums matroids: for example, the algebraic cycle matroid of any graph not including a subdivision of the Bean graph is a thin sums matroid, as follows:

Definition 12.6.2. Let G be a graph with vertex set V and edge set E, and k a field. We can pick a direction for each edge e, calling one of its ends its *source* $s(e)$ and the other its *target* $t(e)$. Then the family $f^G = (f^G_e | e \in E)$ of functions from V to k is given by $f_e = \chi_{t(e)} - \chi_{s(e)}$, where for any vertex v the function χ_v takes the value 1 at v and 0 elsewhere.

Theorem 12.6.3. *Let G be a graph not including any subdivision of the Bean graph. Then M_{f^G} is the algebraic cycle matroid of G.*

This theorem, which motivated the defintion of M_f, is proved in [1].

For the rest of this section, M will denote the algebraic cycle matroid for the graph G in Figure 12.9, in which we have assigned directions to all the edges and labelled them for future reference. We showed in the Section 12.4 that M^+ is wild. We shall devote the rest of this Section 12.4 to showing that in fact it gives an example of a thin sums matroid whose dual isn't a thin sums matroid.

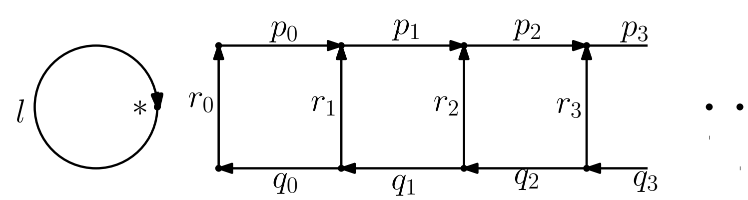

Figure 12.9: The graph G

As usual, we denote the vertex set of G by V and the edge set by E. We call the unique vertex lying on the loop at the left $*$.

Theorem 12.6.4. *M^+ is a thin sums matroid over the field $\mathbb{Q}$.*

Proof. We begin by specifying the family $(f_e|e \in E)$ of functions from V to $\mathbb{Q}$ for which $M^+ = M_f$. We take f_e to be f_e^G as in Definition 12.6.2 if e is one of the p_i or q_i, to be χ_* if $e = l$, and to be $f_e^G + i \cdot \chi_*$ if $e = r_i$.

First, we have to show that every circuit of M^+ is dependent in M_f. There are a variety of possible circuit types: in fact, types (b), (c), (e) and (f) from Figures 12.3 and 12.4 can arise. We shall only consider type (f): the proofs for the other types are very similar. Figure 12.10 shows the two ways a circuit of type (f) can arise.

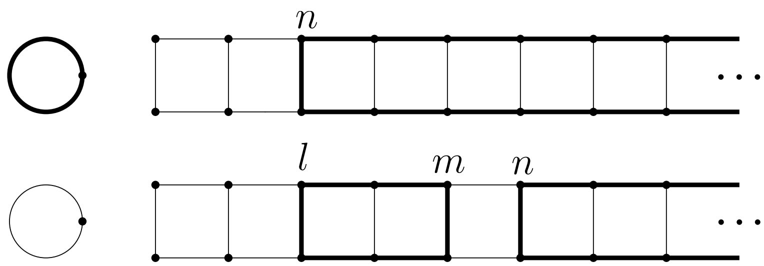

Figure 12.10: The two ways of obtaining a circuit of type (f)

The first includes the edge l, together with r_n for some n and all those p_i and q_i with $i \geq n$. We seek a thin dependence λ such that λ is nonzero on precisely these edges.

We shall take $\lambda_{r_n} = 1$. We can satisfy the equations $\sum_{e \in E} \lambda_e f_e(v)$ with $v \neq *$ by taking $\lambda_{p_i} = \lambda_{q_i} = 1$ for all $i \geq n$. The equation $\sum_{e \in E} \lambda_e f_e(*) = 0$ reduces to $\lambda_* + n\lambda_{r_n} = 0$, which we can satisfy by taking $\lambda_* = -n$. It is immediate that this gives a thin dependence of f.

The second way a circuit of type (f) can arise includes the edges r_l, r_m and r_n, together with those p_i and q_i with either $l \leq i < m$ or $n \leq i$. We seek a thin dependence λ such that λ is nonzero on precisely these edges.

The equations $\sum_{e \in E} \lambda_e f_e(v)$ with $v \neq *$ may be satisfied by taking $\lambda_{p_i} = \lambda_{q_i} = \lambda_{r_l} = -\lambda_{r_m}$ for $l \leq i \leq m$ and $\lambda_{p_i} = \lambda_{q_i} = \lambda_{r_n}$ for $i \geq n$. The equation $\sum_{e \in E} \lambda_e f_e(*) = 0$ reduces to $l\lambda_{r_l} + m\lambda_{r_m} + n\lambda_{r_n} = 0$, which since $\lambda_{r_m} = -\lambda_{r_l}$ reduces further to $(m - l)\lambda_{r_m} = n\lambda_{r_n}$. We can satisfy this equation by taking $\lambda_{r_m} = n$ and $\lambda_{r_n} = m - l$. Taking the remaining λ_e to be given as above then gives a thin dependence of f. Note that $\lambda \neq 0$ since $m \neq l$ and thus $\lambda_{r_n} \neq 0$.

Next, we need to show that every dependent set of M_f is also dependent in M^+, completing the proof. Let D be such a dependent set, as witnessed by a nonzero thin dependence λ of f which is 0 outside D. Let $D' = \{e|\lambda_e \neq 0\}$, the *support* of λ. Using the equations $\sum_{e \in E} \lambda_e f_e(v)$ with $v \neq *$, we may deduce that the degree of D' at each vertex (except possibly $*$) is either 0 or at least 2. Therefore any edge (except possibly l) contained in D' is contained in some circuit of M included in D'. Since $\{l\}$ is already a circuit of M, we can even drop the qualification 'except possibly l'.

Since D' is nonempty, it must include some circuit O of M. Suppose first of all for a contradiction that $D' = O$. The intersection of D' with the set $\{l\} \cup \{r_i | i \in \mathbb{N}_0\}$ is nonempty, so by the equation $\sum_{e \in E} \lambda_e f_e(*) = 0$ this intersection must have at least 2 elements. The only way this can happen with D' a circuit is if there are $m < n$ such that D' consists of r_m, r_n, and the p_i and q_i with $m \leq i < n$. We now deduce, since λ is a thin dependence, that $\lambda_{p_i} = \lambda_{q_i} = \lambda_{r_m} = -\lambda_{r_n}$ for $m \leq i \leq n$. In particular, the equation $\sum_{e \in E} \lambda_e f_e(*) = 0$ reduces to $(m-n)\lambda_{r_m} = 0$, which is the desired contradiction as by assumption $\lambda_{r_m} \neq 0$ and $m < n$. Thus $D' \neq O$, and we can pick some $e \in D \setminus O$. As above, D' includes some M-circuit O' containing e. Then the union $O \cup O' \subseteq D$ is M^+-dependent by Corollary 12.4.4. $\square$

Theorem 12.6.5. *$(M^+)^*$ is not a thin sums matroid over any field.*

Proof. Suppose for a contradiction that it is a thin sums matroid M_f, with f a family of functions $A \to k$. For each circuit O of $(M^+)^*$, we can find a nonzero thin dependence λ of f which is nonzero only on O - it must be nonzero on the whole of O by minimality of O.

The circuits of $(M^+)^* = (M^*)^-$ are precisely the circuits and the bases of M^*, the dual of the algebraic cycle matroid of G, since no circuit in M^* includes a base. This dual M^*, called the *skew cuts* matroid of G, is known to have as its circuits those cuts of G which are minimal subject to the condition that one side contains no rays [21].

Thus since $\{r_0, q_0\}$ is a skew cut, we can find a thin dependence λ^0 which is nonzero precisely at r_0 and q_0. Similarly, for each $i > 0$ we can find a thin dependence λ^i which is nonzero precisely at q_{i-1}, r_i and q_i. Since the set of bold edges in Figure 12.2 is also a circuit of $(M^+)^*$, there is a thin dependence λ which is nonzero on precisely those edges.

To obtain a contradiction, we will show that $\{r_i | i \in \mathbb{N}\}$ is dependent in M_f. The idea behind the following calculations is to consider $\{r_i | i \in \mathbb{N}\}$ as the limit of the M_f-circuits $\{r_i | 0 \leq i \leq n\} \cup \{q_n\}$ and then to use the properties of thin sum representations to show that the "limit" $\{r_i | i \in \mathbb{N}\}$ inherits the dependence.

Now define the sequences $(\mu_i | i \in \mathbb{N})$ and $(\nu_i | i \in \mathbb{N})$ inductively by $\nu_0 = 1$, $\nu_i = -(\lambda^i_{q_i}/\lambda^i_{q_{i-1}})\nu_{i-1}$ for $i > 0$ and $\mu_i = -(\lambda^i_{r_i}/\lambda^i_{q_i})\nu_i$. Pick any $a \in A$. Then we have $0 = \sum_{e \in E} \lambda^0_e f_e(a) = \lambda^0_{r_0} f_{r_0}(a) + \lambda^0_{q_0} f_{q_0}(a)$, and rearranging gives

$$\nu_0 f_{q_0}(a) = \mu_0 f_{r_0}(a)\,.$$

Similarly, $0 = \sum_{e \in E} \lambda^i_e f_e(a) = \lambda^i_{q_{i-1}} f_{q_{i-1}}(a) + \lambda^i_{r_i} f_{r_i}(a) + \lambda^i_{q_i} f_{q_i}(a)$, and rearranging gives

$$\nu_i f_{q_i}(a) = \nu_{i-1} f_{q_{i-1}}(a) + \mu_i f_{r_i}(a)\,.$$

So by induction on i we get the formula

$$\nu_i f_{q_i}(a) = \sum_{j=0}^{i} \mu_j f_{r_j}(a)\,.$$

The formula $\sum_{e \in E} \lambda_e f_e(a) = 0$ implicitly includes the statement that the sum is well defined, so only finitely many summands can be nonzero. In particular, there can only be finitely many i for which $f_{q_i}(a) \neq 0$. It then follows by the formula above that there are only finitely many i such that $f_{r_i}(a)$ is nonzero, since if $f_{r_i}(a) \neq 0$, then as $\mu_i \neq 0$ we have $\nu_i f_{q_i}(a) \neq \nu_{i-1} f_{q_{i-1}}(a)$. So as $\nu_i \neq 0$ and $\nu_{i-1} \neq 0$, one of $f_{q_i}(a)$ or $f_{q_{i-1}}(a)$ is not equal to zero. Therefore all but finitely many $f_{r_i}(a)$ are zero since all but finitely many $f_{q_i}(a)$ are zero. So the following sum is well defined and evaluates to zero.

$$\sum_{i=0}^{\infty} \mu_i f_{r_i}(a) = 0\,.$$

Therefore, if we define a family $(\lambda'_e | e \in E)$ by $\lambda'_{r_i} = \mu_i$ and $\lambda'_e = 0$ for other values of e, then we have

$$\sum_{e \in E} \lambda'_e f_e(a) = 0\,.$$

Since $a \in A$ was arbitrary, this implies that λ' is a thin dependence of f. Note that $\lambda' \neq 0$ since $\lambda'_{r_0} \neq 0$. Thus the set $\{r_i | i \in \mathbb{N}\}$ is dependent in $M_f = (M^*)^-$. But it is also an $(M^*)^-$basis, since adding l gives a basis of M^*. This is the desired contradiction. $\square$

Bibliography

[1] H. Afzali and N. Bowler. Thin sums matroids and duality. *Adv. Math,*, 271:1–29, 2015.

[2] R. Aharoni and E. Berger. Menger's theorem for infinite graphs. *Invent. math.*, 176:1–62, 2009.

[3] R. Aharoni and C. Thomassen. Infinite, highly connected digraphs with no two arc-disjoint spanning trees. *J. Graph Theory*, 13:71–74, 1989.

[4] R. Aharoni and R. Ziv. The intersection of two infinite matroids. *J. London Math. Soc.*, 58:513–525, 1998.

[5] Elad Aigner-Horev, Johannes Carmesin, and Jan-Oliver Fröhlich. Infinite matroid union. Preprint (2011), current version available at http://arxiv.org/abs/1111.0602.

[6] Elad Aigner-Horev, Johannes Carmesin, and Jan-Oliver Fröhlich. On the intersection of infinite matroids. Discrete Mathematics, to appear, current version available at http://arxiv.org/abs/1111.0606.

[7] Th. Andreae. Über maximale Systeme von kantendisjunkten unendlichen Wegen in Graphen. *Math. Nachr.*, 101:219–228, 1981.

[8] Thomas Andreae. On disjoint configurations in infinite graphs. *Journal of Graph Theory*, 39(4):222–229, 2002.

[9] Thomas Andreae. Classes of locally finite ubiquitous graphs. *Journal of Combinatorial Theory, Series B*, 103(2):274 – 290, 2013.

[10] M.A. Armstrong. *Basic Topology.* Springer-Verlag, 1983.

[11] R. H. Bing. An alternative proof that 3-manifolds can be triangulated. *Ann. Math.(2)*, 69:37–65, 1959.

[12] Robert E Bixby. On Reid's characterization of the ternary matroids. *Journal of Combinatorial Theory, Series B*, 26(2):174 – 204, 1979.

[13] Debrah Boutin and Wilfried Imrich. The cost of distinguishing graphs. In Tullio Ceccherini-Silberstein, Maura Salvatori, and Ecaterina Sava-Huss, editors, *Groups, Graphs and Random Walks*, London Mathematical Society Lecture Note Series. Cambridge University Press, publication planned for April 2017.

[14] Nathan Bowler and Johannes Carmesin. An excluded minors method for infinite matroids. Journal of Combi. Theory (Series B), to appear, current version available at http://arxiv.org/pdf/1212.3939v1.

[15] Nathan Bowler and Johannes Carmesin. Matroids with an infinite circuit-cocircuit intersection. *J. Combin. Theory (Series B)*, 107:78–91, 2014.

[16] Nathan Bowler and Johannes Carmesin. Matroid intersection, base packing and base covering for infinite matroids. *Combinatorica*, 35(2):153–180, 2015.

[17] Nathan Bowler, Johannes Carmesin, and Robin Christian. Infinite graphic matroids. To appear in Combinatorica, available at `http://arxiv.org/abs/1309.3735`.

[18] Nathan Bowler, Johannes Carmesin, and Luke Postle. Reconstruction of infinite matroids from their 3-connected minors. to appear in European J. Combin.

[19] Nathan Bowler, Johannes Carmesin, and Julian Pott. Edge-disjoint double rays in infinite graphs: a Halin type result. *J. Combin. Theory Ser. B*, 111:1–16, 2015.

[20] Ulrich Brehm. A nonpolyhedral triangulated mobius strip. *Proc. Amer. Math. Soc.*, 89(3):519–522, 1983.

[21] Henning Bruhn and Reinhard Diestel. Infinite matroids in graphs. *Discrete Math.*, 311(15):1461–1471, 2011.

[22] Henning Bruhn, Reinhard Diestel, Matthias Kriesell, Rudi Pendavingh, and Paul Wollan. Axioms for infinite matroids. *Adv. Math.*, 239:18–46, 2013.

[23] Henning Bruhn and Paul Wollan. Finite connectivity in infinite matroids. *European J. Combin.*, 33(8):1900–1912, 2012.

[24] Peter J. Cameron. *Oligomorphic permutation groups.* Cambridge: Cambridge University Press, 1990.

[25] Johannes Carmesin. Embedding simply connected 2-complexes in 3-space – I. A Kuratowski-type characterisation. Preprint 2017, available at ”https://arxiv.org/pdf/1709.04642.pdf”.

[26] Johannes Carmesin. Embedding simply connected 2-complexes in 3-space – II. Rotation systems. Preprint 2017, available at "https://arxiv.org/pdf/1709.04643.pdf".

[27] Johannes Carmesin. Embedding simply connected 2-complexes in 3-space – III. Constraint minors. Preprint 2017, available at "https://arxiv.org/pdf/1709.04645.pdf".

[28] Johannes Carmesin. Embedding simply connected 2-complexes in 3-space – IV. Dual matroids. Preprint 2017, available at "https://arxiv.org/pdf/1709.04652.pdf".

[29] Johannes Carmesin. Embedding simply connected 2-complexes in 3-space – V. A refined Kuratowski-type characterisation. Preprint 2017, available at "https://arxiv.org/pdf/1709.04659.pdf".

[30] Johannes Carmesin, Reinhard Diestel, Matthias Hamann, and Fabian Hundertmark. k-blocks: a connectivity invariant for graphs. *SIAM J. Discrete Math.*, 28(4):1876–1891, 2014.

[31] Johannes Carmesin, Reinhard Diestel, Fabian Hundertmark, and Maya Stein. Connectivity and tree structure in finite graphs. *Combinatorica*, 34(1):11–45, 2014.

[32] R. Christian. *Infinite graphs, graph-like spaces and B-matroids.* PhD thesis, University of Waterloo, 2010.

[33] Arnaud de Mesmay, Yo'av Rieck, Eric Sedgwick, and Martin Tancer. Embeddability in $\mathbb{R}^3$ is NP-hard. Preprint 2017, available at: "https://arxiv.org/pdf/1708.07734".

[34] R. Diestel, H. A. Jung, and R. G. Möller. On vertex transitive graphs of infinite degree. *Arch. Math.*, 60(6):591–600, 1993.

[35] Reinhard Diestel. *Graph Theory* (4th edition). Springer-Verlag, 2010. Electronic edition available at: `http://diestel-graph-theory.com/index.html`.

[36] Reinhard Diestel. *Graph Theory* (5th edition). Springer-Verlag, 2016. Electronic edition available at: `http://diestel-graph-theory.com/index.html`.

[37] A. Dress. Duality theory for finite and infinite matroids with coefficients. *Advances in Mathematics*, 59:97–123, 1986.

[38] M. J. Dunwoody and B. Krön. Vertex cuts. *J. Graph Theory*, 80(2):136–171, 2015.

[39] B. Dushnik and E. W. Miller. Partially ordered sets. *American Journal of Mathematics*, 63, 1941.

[40] P. Erdős and A. Hajnal. On chromatic number of graphs and set-systems. *Acta Mathematica Academiae Scientarum Hungarica*, 17:61–99, 1966.

[41] Kazufumi Eto, Shosaku Matsuzaki, and Makoto Ozawa. An obstruction to embedding 2-dimensional complexes into the 3-sphere. *Topology Appl.*, 198:117–125, 2016.

[42] J. Geelen, A. Gerards, and A. Kapoor. The excluded minors for GF(4)-representable matroids. *J. Combin. Theory (Series B)*, 79:247–299, 2000.

[43] Jim Geelen, Bert Gerards, and Geoff Whittle. Solving Rota's conjecture. *Notices Amer. Math. Soc.*, 61(7):736–743, 2014.

[44] R. Halin. Über unendliche Wege in Graphen. *Math. Annalen*, 157:125–137, 1964.

[45] R. Halin. Über trennende Eckenmengen in Graphen und den Mengerschen Satz. *Math. Ann.*, 157:34–41, 1964.

[46] R. Halin. Über die Maximalzahl fremder unendlicher Wege in Graphen. *Math. Nachr.*, 30:63–85, 1965.

[47] R. Halin. Die Maximalzahl fremder zweiseitig unendlicher Wege in Graphen. *Math. Nachr.*, 44:119–127, 1970.

[48] R. Halin. Automorphisms and endomorphisms of infinite locally finite graphs. *Abh. Math. Semin. Univ. Hamb.*, 39:251–283, 1973.

[49] R. Halin. Lattices of cuts in graphs. *Abh. Math. Semin. Univ. Hamb.*, 61:217–230, 1991.

[50] R. Halin. Some finiteness results concerning separation in graphs. *Discrete Math.*, 101(1-3):97–106, 1992.

[51] R. Halin. The structure of rayless graphs. *Abh. Math. Semin. Univ. Hamb.*, 68:225–253, 1998.

[52] R. Halin. A note on graphs with countable automorphism group. *Abh. Math. Semin. Univ. Hamb.*, 70:259–264, 2000.

[53] Matthias Hamann. End-transitive graphs. *Isr. J. Math.*, 189:437–459, 2012.

[54] Matthias Hamann and Fabian Hundertmark. The classification of connected-homogeneous digraphs with more than one end. *Trans. Am. Math. Soc.*, 365(1):531–553, 2013.

[55] Hatcher. Notes on basic 3-manifold topology. available at "http://www.math.cornell.edu/ hatcher/3M/3Mfds.pdf".

[56] A. Hatcher. *Algebraic Topology.* Cambridge Univ. Press, 2002.

[57] Michael Heusner and Raphael Zentner. A new algorithm for 3-sphere recognition. Preprint 2016, available at arXiv:1610.04092.

[58] D.A. Higgs. Infinite graphs and matroids. Recent Progress in Combinatorics, Proceedings Third Waterloo Conference on Combinatorics, Academic Press, 1969, pp. 245–53.

[59] S. V. Ivanov. Recognizing the 3-sphere. *Illinois J. Math.*, 45(4):1073–1117, 2001.

[60] R. Möller J. Carmesin, F. Lehner. On tree-decompositions of one-ended graphs. Preprint 2017.

[61] H. A. Jung. A note on fragments of infinite graphs. *Combinatorica*, 1:285–288, 1981.

[62] H.A. Jung. Wurzelbäume und unendliche Wege in Graphen. *Math. Nachr.*, 41:1–22, 1969.

[63] Pter Komjth. A note on uncountable chordal graphs. Discrete Mathematics, to appear.

[64] Bernhard Krön and Rögnvaldur G. Möller. Quasi-isometries between graphs and trees. *J. Comb. Theory, Ser. B*, 98(5):994–1013, 2008.

[65] Kenneth Kunen. *Set Theory.* Studies in Logic 34. London: College Publications 2011. 1990.

[66] John Lake. A problem concerning infinite graphs. *Discrete Mathematics*, 14(4):343 – 345, 1976.

[67] László Lovász. Graph minor theory. *Bull. Amer. Math. Soc. (N.S.)*, 43(1):75–86, 2006.

[68] Jiri Matousek, Eric Sedgwick, Martin Tancer, and Uli Wagner. Embeddability in the 3-sphere is decidable. In *Computational geometry (SoCG'14)*, pages 78–84. ACM, New York, 2014. Extended version available at "https://arxiv.org/pdf/1402.0815".

[69] Jiri Matousek, Martin Tancer, and Uli Wagner. Hardness of embedding simplicial complexes in $\mathbb{R}^d$. *J. Eur. Math. Soc. (JEMS)*, 13(2):259–295, 2011.

[70] Laurence R. Matthews and James G. Oxley. Infinite graphs and bicircular matroids. *Discrete Math.*, 19(1):61–65, 1977.

[71] B. Mohar and C. Thomassen. *Graphs on Surfaces.* Johns Hopkins, 2001.

[72] Edwin E. Moise. Affine structures in 3-manifolds. V. The triangulation theorem and Hauptvermutung. *Annals of Mathematics. Second Series*, 56:96–114, 1952.

[73] P. Komjáth C. Reiher N. Bowler, J. Carmesin. The colouring number of infinite graphs. Preprint 2015.

[74] C.St.J.A. Nash-Williams. Infinite graphs—a survey. *J. Combin. Theory*, 3:286–301, 1967.

[75] James Oxley. *Matroid Theory*. Oxford University Press, 1992.

[76] James Oxley. *Matroid Theory* (2nd edition). Oxford University Press, 2011.

[77] C. Papakyriakopoulos. A new proof for the invariance of the homology groups of a complex (in greek). *Bull. Soc. Math. Grece*, 22:1–154, 1946.

[78] C. D. Papakyriakopoulos. On Dehn's Lemma and the Asphericity of Knots. *Ann. Math.*, 66:1–26, 1957.

[79] G. Perelman. The entropy formula for ricci flow and its geometric applications. 2002, available at "https://arxiv.org/abs/math.DG/0211159".

[80] G. Perelman. Finite extinction time for the solutions to the ricci flow on certain three-manifolds. 2003, available at "https://arxiv.org/abs/math.DG/0307245".

[81] G. Perelman. Ricci flow with surgery on three-manifolds. 2003, available at "https://arxiv.org/abs/math.DG/0303109".

[82] J. L. Ramírez Alfonsín. Knots and links in spatial graphs: a survey. *Discrete Math.*, 302(1-3):225–242, 2005.

[83] Neil Robertson and P. D. Seymour. Graph minors. XX. Wagner's conjecture. *J. Combin. Theory Ser. B*, 92(2):325–357, 2004.

[84] Saul Schleimer. Sphere recognition lies in np. *In Michael Usher, editor, Low- dimensional and Symplectic Topology. American Mathematical Society.*, 82:183–214, 2011.

[85] Ruediger Schmidt. Ein Ordnungsbegriff für Graphen ohne unendliche Wege mit einer Anwendung auf n-fach zusammenhaengende Graphen. *Arch. Math.*, 40:283–288, 1983.

[86] A. Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency - Volume B*. Springer-Verlag, 2003.

[87] P.D Seymour. Matroid representation over GF(3). *Journal of Combinatorial Theory, Series B*, 26(2):159 – 173, 1979.

[88] Saharon Shelah. A compactness theorem for singular cardinals, free algebras, whitehead problem and transversals. *Israel Journal of Mathematics*, 21:319–349, 1975.

[89] Carsten Thomassen. The Hadwiger number of infinite vertex-transitive graphs. *Combinatorica*, 12(4):481–491, 1992.

[90] Jacques Tits. Sur le groupe des automorphismes d'un arbre. Essays Topol. Relat. Top., Mém. dédiés à Georges de Rham, 188-211 (1970)., 1970.

[91] K. Truemper. *Matroid Decompositions (Revised Version)*. Leibnitz company, 2017.

[92] W. T. Tutte. A Homotopy Theorem for Matroids, II. *Transactions of the American Mathematical Society*, 88(1):pp. 161–174, 1958.

[93] W. T. Tutte. Lectures on matroids. *J. Res. Nat. Bur. Standards Sect. B*, 69B:1–47, 1965.

[94] W. T. Tutte. *Graph theory*, volume 21 of *Encyclopedia of Mathematics and its Applications*. Addison-Wesley Publishing Company, Advanced Book Program, Reading, MA, 1984. With a foreword by C. St. J. A. Nash-Williams.

[95] S.H.M. van Zwam. *Partial fields in matroid theory*. PhD thesis, Eindhoven University of Technology, 1999.

[96] Uli Wagner. Minors, embeddability, and extremal problems for hypergraphs. pages 569–607, 2013.

[97] M. Wagowski. Strong duality property for matroids with coefficients. *Europ. J. Comb.*, 15:293–302, 1994.

[98] N. White. Unimodular matroids. In *Combinatorial geometries*, volume 29 of *Encyclopedia Math. Appl.*, pages 40–52. Cambridge Univ. Press, Cambridge, 1987.

[99] H. Whitney. Non-separable and planar graphs. *Trans. Am. Math. Soc.*, 34:339–362, 1932.

[100] H. Whitney. 2-Isomorphic Graphs. *Amer. J. Math.*, 55:245–254, 1933.

[101] Geoff Whittle. On matroids representable over GF(3) and other fields. *Transactions of the American Mathematical Society*, 349(2):pp. 579–603, 1997.

[102] D.R Woodall. A note on a problem of Halin's. *Journal of Combinatorial Theory, Series B*, 21(2):132 – 134, 1976.

[103] Raphael Zentner. Integer homology 3-spheres admit irreducible representations in SL(2,C). Preprint 2016, available at arXiv:1605.08530.