Graphs and their Circuits

—From Finite to Infinite—

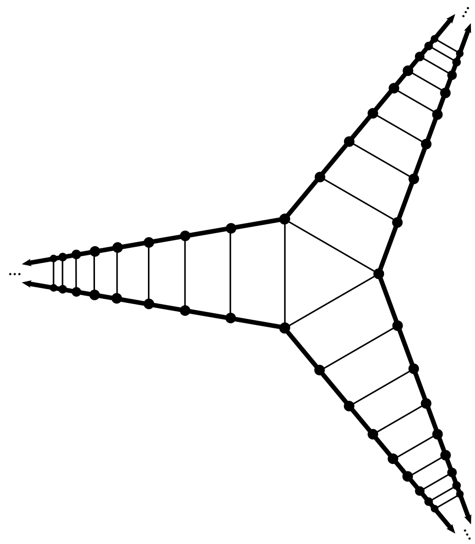

Habilitationsschrift
Department Mathematik
der Universität Hamburg

vorgelegt von Henning Bruhn

Hamburg 2009

Contents

最愛のなほへ

Acknowledgement

I am deeply indebted to Reinhard Diestel for long years of advice and support, and for teaching me most I know about graphs. I am very grateful to András Sebő for his kind hospitality in Grenoble and for introducing me to Combinatoral Optimisation. Moreover, I'd like to thank Ron Aharoni and Eli Berger, and Robin Christian and Bruce Richter for memorable and productive stays in Haifa and Waterloo, respectively. Finally, I am very happy to have Maya Stein as my long-time collaborator and as a friend.

Chapter 1

Infinite circuits

All our graphs are allowed to have loops and parallel edges, with the exception of 2-connected graphs, which we require to be loopless, and of 3-connected graphs, which, in addition, cannot have parallel egdes. Most of our graphs will be locally finite, that is, graphs in which all the vertices have finite degree. In general, our notation follows Diestel [31].

1.1 The main theme

Often, easy problems in finite graphs become challenging, and sometimes very interesting indeed, in infinite graphs. This observation is the main theme of this work. A good example is Menger's theorem:

Theorem 1.1 (Menger)**.** *For any two vertex sets A and B in a finite graph it holds that the maximal cardinality of a set of disjoint A–B paths equals the minimal cardinality of an A–B separator.*

While Menger's theorem is not trivial to prove it is also not very hard and can be done in about half a page or less. In contrast, the analogue of Menger's theorem in infinite graphs, the Erdős-Menger conjecture, is fiendishly hard.

Conjecture 1.2 (Erdős-Menger)**.** *For any two vertex sets A and B in a graph there exists a set $\mathcal{P}$ of pairwise disjoint A–B paths and an A–B separator that contains precisely one vertex from each path in $\mathcal{P}$.*

A proof of the Erdős-Menger conjecture was announced by Aharoni and Berger [2]. Their proof, which runs for some 50 pages, is full of intricate details and complex arguments. Surely, a good amount of this is due to difficulties inherent in handling uncountable sets. We will be spared these additional complexities as we will mostly be concerned with countable graphs. Nevertheless, even the countable case of Conjecture 1.2 is still considerably more difficult than Menger's theorem. We will encounter more examples of this phenomenon throughout this thesis.

Menger's theorem is instructive in another respect, too. When we extend a result about finite graphs to infinite graphs it depends heavily on the precise formulation whether we obtain a useful result or a triviality. Here is a simple variant of Menger's theorem in infinite graphs.

Theorem 1.3. *For any two vertex sets A and B in a graph it holds that the maximal cardinality of a set of disjoint A–B paths equals the minimal cardinality of an A–B separator.*

On first glance, the theorem might appear a more direct analogue of Menger's theorem. This, however, is not the case. In fact, the theorem, which is easy to prove, misses a hidden but central point of Theorem 1.1. Besides determining the cardinality of a minimal separator Menger's theorem yields, in addition, an insight in the structure of the separator. This insight is captured in the Erdős-Menger conjecture but lost in Theorem 1.3. As a result, the case of Theorem 1.3 when there is an infinite set $\mathcal{P}$ of disjoint A–B paths becomes easy to prove. Choosing a maximal such $\mathcal{P}$ we can simply choose $\bigcup_{P\in\mathcal{P}} V(P)$ as the separator of the desired size. In contrast, the hard case of Conjecture 1.2 is precisely when A and B are infinitely connected.

The two observations, that easy problems become hard in infinite graphs and the heavy dependence on the precise formulation, are a salient feature of the cycle space in an infinite graph, the second theme of this thesis.

The cycle space is the set of all symmetric differences of circuits, ie the edge sets of cycles. In a finite graph, the elements of the cycle space are easily characterised in terms of degrees: an edge set Z is an element of the cycle space if and only if every vertex is incident with an even number of edges in Z. In infinite graphs this fact is either trivially false, the edge set of a double ray may serve as a counterexample, or trivially true—provided we require additionally that Z is a finite set.

Tutte's planarity criterion, which we will discuss a bit more in Section 1.4, provides a second example. It states that a finite 3-connected graph is planar if and only if every edge lies in at most two induced non-separating cycles. As we will see, the criterion becomes outright false in infinite graphs.

In both cases, in the characterisation of cycle space elements by degrees and in Tutte's planarity criterion, the naive formulation is to blame. Indeed, in the setting of the *topological cycle space*, which allows infinite circuits, both results become challenging, non-trivial and true. We will describe and define the topological cycle space in the next two sections.

1.2 Sources of infinite circuits

The topological cycle space has been developed by Diestel and Kühn, see [33, 34, 35]. Its main feature is that it allows infinite circuits. Apart from that it seems natural to have infinite circuits in infinite graphs there are three good reasons that motivate the definition of infinite circuits.

Firstly, in a 2-connected infinite plane graph infinite circuits arise from the face boundaries. In Figure 1.1, the boundary of the outer face consists of the union of three disjoint rays. The fact that face boundaries in a finite 2-connected graph are cycles suggests that we should view the edge set of the infinite face boundary as a circuit, too. We should point out, though, that even in the case of a 3-connected graph the infinite face boundaries depend strongly on the very nature of the drawing.

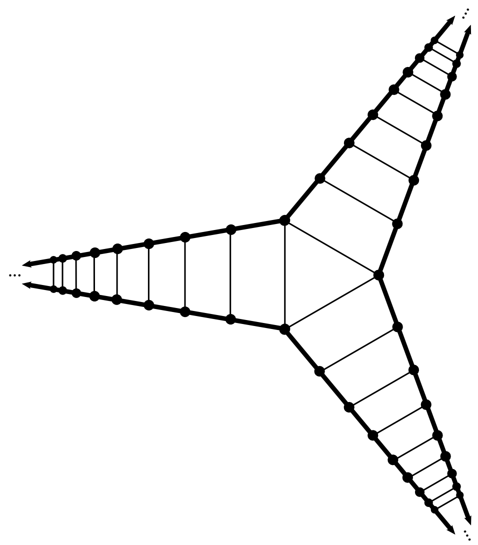

Figure 1.1: An infinite face boundary

Secondly, the matroids associated with a graph lead to infinite circuits, too. If, for a locally finite graph G, we declare the minimal finite cuts to be the cocircuits we obtain a matroid M on the edges of G. If G is finite then the circuits of M are precisely the circuits of G. On the other hand, if G is infinite then M may have infinite circuits. For example, if G is the graph in Figure 1.1 then M will have the edge set of the infinite face boundary as a circuit. Clearly, these considerations suppose that it is possible to define a well-behaved matroid on an infinite ground set in the first place. While it might not be widely known, this is possible, and we shall encounter infinite matroids in the last chapter, where we will also discuss the relationship between matroids and the infinite circuits of a graph.

Finally, and this is the most compelling reason for them, infinite circuits make previously unsolvable problems solvable. A number of theorems about circuits in finite graphs are false in infinite graphs but become true if infinite circuits are considered. Below and throughout the rest of the thesis we will discuss a number of examples. But first, let us give the formal definition of infinite circuits, which is based on a topology on the graph together with its ends.

1.3 The topological cycle space

Let us now describe the topology with which we endow an infinite graph. In most cases the graph in question will be locally finite. Then the resulting space, called the Freudenthal compactification, will be obtained by compactifying the graph by its ends. Sometimes, most notably when we will treat dual graphs, we will find it more useful to work within a slightly large class of graphs. This class consists of the graphs in which

$$\textit{any two vertices can be separated by finitely many edges.} \tag{1.1}$$

The reason for considering the class of these graphs will become apparent in Chapter 2.

Let $G = (V, E)$ be a fixed graph in this section. A 1-way infinite path is called a *ray*, a 2-way infinite path is a *double ray*. A subray of a ray or a double ray is a *tail*. Two rays are called *equivalent* if there are infinitely many disjoint paths between them. The equivalence classes of rays are the *ends of* G, we denote the set of these by $\Omega(G)$.

We define a topological space $|G|$ on G together with its ends. In order to do so we view every edge as homeomorphic to the unit interval and pick a fixed homeomorphism for each edge. Next, consider a vertex v. The basic open neighbourhoods of v are then the sets consisting of v together with all points of distance $< \frac{1}{n}$ from v on incident edges for $n \in \mathbb{N}$ (where the distance is measured by the fixed homeomorphism for that edge). It remains to describe the basic open neighbourhoods of an end, ω say. Pick a finite vertex set S, and denote the component of $G - S$ that contains a ray of ω (and thus a subray for every ray in ω) by $C(S, \omega)$. We say that ω *belongs to* $C(S, \omega)$. A basic open neighbourhood of ω now consists of $C(S, \omega)$, all ends that have a ray in $C(S, \omega)$ and the union of all interior points of edges between S and $C(S, \omega)$.

The resulting space is denoted by $|G|$. For more on this space and subtle variants of it, we refer the reader to Diestel [32]. If G is locally finite then $|G|$ is Hausdorff and compact. On the other hand, if G is not locally finite then $|G|$ fails to be Hausdorff and in fact is generally not well suited for our purpose. Therefore, we shall define a quotient space $\tilde{G}$ of $|G|$.

We say that a vertex v *dominates* an end ω if there are infinitely many paths from v to a ray R in ω that meet pairwise only in v. The end ω is then a *dominated end*. The quotient space $\tilde{G}$ is obtained from $|G|$ by identifying every dominated end with its dominating vertex. We remark that if G satisfies (1.1) then no two vertices will be identified (and in fact we will never use $\tilde{G}$ otherwise), and if G is in addition 2-connected then $\tilde{G}$ is a compact space.

Having defined the underlying topologies, we can now state precisely what circuits should be. For this, G is assumed to be locally finite or to at least satisfy (1.1). While we formulate our definition in terms of $\tilde{G}$ the reader should bear in mind that $\tilde{G} = |G|$ for locally finite G.

A *circle* is the homeomorphic image of the unit circle in $\tilde{G}$, and a *circuit* is its edge set, i.e. the set of all edges contained in it. For a circle C the subgraph $C \cap G$ is called a *cycle*. We note that every edge of which a circle contains an interior point lies completely in the circle, see Diestel and Kühn [35]. This definition not only includes the traditional, finite circuits but also allows infinite ones. As an example consider Figure 1.2. There the (edge set of the) double ray D is a circuit, since both tails of D are in the same end to the right. On the other hand, double ray D' is not a circuit. Yet, the union of D' and D'' is a circuit. Circles may become fairly complex. It is not hard to construct a circle consisting of countably infinite many disjoint double rays and uncountably many ends, see Diestel [30].

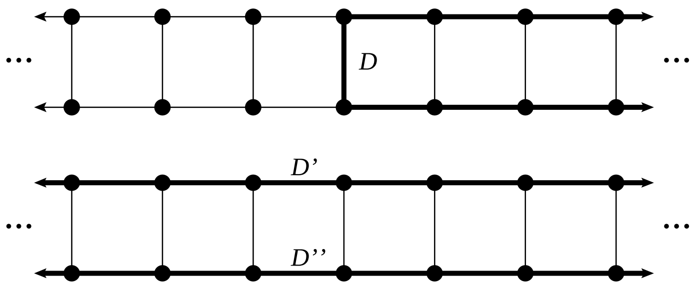

Figure 1.2: Circuits in the double ladder

The image of a continuous mapping $[0, 1] \to \tilde{G}$ is called a *topological path*. The homeomorphic image of the unit interval $[0, 1]$ in $\tilde{G}$ is an *arc*. Observe that circuits as well as arcs must contain edges.

The topological cycle space consists of all sets of all finite or (well-defined) infinite sums of circuits. That infinite sums are allowed is the second main feature of the topological cycle space. We omit here the discussion of why this is not only natural but also necessary, see Diestel [30] and [15]. Nevertheless, we will return to infinite sums in Chapter 6, where we will investigate them in a more abstract setting.

Call a family $\mathcal{F}$ of edge sets *thin* if every edge appears in at most finitely many members of $\mathcal{F}$. The *thin sum* of a thin family $\mathcal{F}$ is then defined as the set of edges that occur in precisely an odd number of the members of $\mathcal{F}$ and is denoted by $\sum_{F\in\mathcal{F}} F$. All sums of edge sets in this thesis are considered to be thin sums. Now, the *topological cycle space* $\mathcal{C}(G)$ is the set of all thin sums of families of circuits.

1.4 Infinite circuits solve problems

Many facts about circuits in finite graphs are false or trivial in infinite graphs but become true once infinite circuits are considered. A prime example is Tutte's planarity criterion. Call a circuit peripheral it is chordless and if deletion of all vertices incident with the edges in the circuit does not separate the graph.

Theorem 1.4 (Tutte[1])**.** *A 3-connected finite graph is planar if and only if every edge lies in at most two peripheral circuits.*

To see why the theorem fails in infinite graphs consider the graph G in Figure 1.3. The graph is certainly non-planar as it is comprised of $K_{3,3}$ (in bold) to which three disjoint infinite 3-ladders are added. On the other hand,

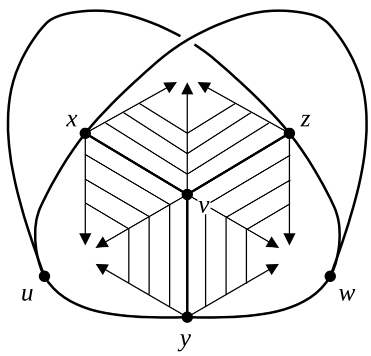

Figure 1.3: Tutte's criterion fails without infinite circuits

[1] It is not entirely clear, whether 'Tutte's planarity criterion' should rightly be attributed to Tutte. It certainly follows readily from his generating theorem proved in [83] and MacLane's planarity criterion. The criterion was, indepedently, discovered by Kelmans [56].

it is not hard to check that no edge lies in three finite peripheral circuits. Thus, if only finite circuits are considered Tutte's planarity criterion would indicate that the graph is planar, which it is not. With infinite circuits, however, the criterion becomes valid. The edge xw, for instance, can be seen to be contained in two finite peripheral circuits and, in addition, in one infinite peripheral circuit. The edge thereby serves as a certificate for non-planarity. In [23] it is proved that not only the example in Figure 1.3 can be repaired by allowing infinite circuits but that Tutte's planarity criterion is valid in all locally finite graphs.

Infinite circuits and the topological cycle space have been almost surprisingly successful in generalising a large number of other theorems, too. In most cases, the topological permits verbatim extensions of results about circuits and the cycle space in a finite graph. Examples are:

- MacLane's planarity criterion, see Chapter 5 as well as [23];
- Tutte's generating theorem [15];
- Gallai's partition theorem [19];
- geodesic circuits generate the cycle space (Georgakopoulos and Sprüssel [47]);
- Whitney's planarity criterion and dual graphs, see next chapter, as well as [16];
- tree packing and arboricity (Stein [76]); and
- fundamental circuits generate the cycle space, see next section and Diestel and Kühn [33].

Moreover, since the circuits of Diestel and Kühn may be infinite, it becomes possible to consider Hamilton circuits in locally finite graphs, see Georgakopoulos [45], Cui, Wang and Yu [29] and [26]. Coming from a more algebraic viewpoint, Diestel and Sprüssel [36] develop a homology that captures the topological cycle space. A more general approach has been pursued by Vella and Richter [87], who define cycle spaces for different compactifications of a graph. This work has been followed up in Casteels and Richter [27] and in Richter and Rooney [70].

1.5 Topological spanning trees

Let us consider one more example in more depth, as it will play a role later.

The starting point of the research into the topological cycle space was a question of Richter. Richter asked how the fact that the fundamental circuits of a spanning tree generate the cycle space could be extended in a meaningful way to infinite graphs. The first observation to make here is that what is called a circuit determines what a tree is. Indeed, a tree is a connected subgraph that does not contain circuits. Now, an ordinary tree might very well contain an infinite circuit, which means we have to amend our definition of a tree. Moreover, as we work within the topological space $\tilde{G}$, for some graph G, it appears natural to consider a subspace rather than a subgraph. If we, finally, then interprete 'connectivity' in a topological way we have all ingredients for the infinite analogon of a spanning tree.

To sum up, given a graph G satisfying (1.1) we say that a subspace T of $\tilde{G}$ is a *topological spanning tree*, or TST for short, if it is closed, contains all the vertices but no circle, and which contains every edge of which it contains an interior point.

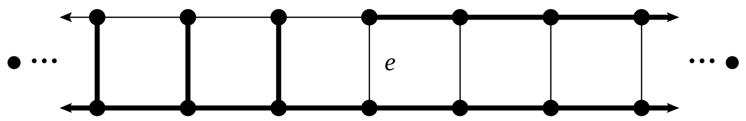

Figure 1.4: A topological spanning tree

The bold subspace in Figure 1.4 of $\tilde{G}$ (or rather $|G|$ as the double ladder is locally finite) constitutes a TST but the subgraph contained in it fails to be a spanning tree as it is not (graph-theoreticially) connected. The addition of the edge e makes (the subgraph of) this subspace into a spanning tree. That tree, however, contains an infinite circuit, namely the edge set of the double ray contained in it with both tails in the right end.

The following theorem answers the original question of Richter. Let e be an edge outside a TST T. The *fundamental circuit of e (with respect to T)* is then the unique circuit contained in $E(T) \cup \{e\}$.

Theorem 1.5 (Diestel and Kühn [35]). *In a graph satisfying* (1.1) *the topological cycle space is generated by the fundamental circuits of any topological spanning tree. Furthermore, the fundamental circuits of a TST form a thin family.*

Spanning trees are the bases of the cycle matroid of a connected graph, and the finite circuits of the graph are the circuits of that matroid. In the same way, a matroid can be associated with a graph in which TSTs are the bases and all circuits (finite or infinite) of the graph are the circuits of the matroid. See Chapter 8.

Topological spanning trees occur naturally in the tree packing problem for infinite graphs. In finite graphs the problem is solved by the following theorem, independently discovered by Nash-Williams and by Tutte. In a partition of the vertex set of a graph an edge is called *cross-edge* if its endvertices lie in distinct partition classes.

Theorem 1.6 (Nash-Williams [63], Tutte [82])**.** *Let G be a finite graph, and let $k \in \mathbb{N}$. Then G has k edge-disjoint spanning trees if and only if for every partition of $V(G)$ with ℓ partition classes there are $\geq k(\ell - 1)$ cross-edges.*

Tutte [82] extended the result to infinite locally finite graphs. The statement he was able to prove, however, does not ascertain the existence of spanning trees but of objects he called 'semi-connected'. The result, at first, appears slightly inelegant and unconvincing. Yet, phrased with TSTs, it becomes a verbatim extension of Theorem 1.6.

Theorem 1.7 (Tutte [82], Stein [76])**.** *Let G be a locally finite graph, and let $k \in \mathbb{N}$. Then G has k edge-disjoint topological spanning trees if and only if for every partition of $V(G)$ with ℓ partition classes there are $\geq k(\ell - 1)$ cross-edges.*

The TSTs in the theorem are necessary—the result becomes false with ordinary spanning trees, see Oxley [66]. Finding a necessary and sufficient condition for the existence of k edge-disjoint spanning trees seems hard. We will briefly encounter this problem again in Section 8.9.

1.6 Basic properties of $\mathcal{C}(G)$

The plethora of results listed above indicate that the topological cycle space is the right setting to overcome the deficits of ordinary finite circuits in infinite graphs. Moreover, the most basic properties of its finite counterpart are shared by the topological cycle space. These properties are the following, which are almost too trivial to necessitate a proof in finite graphs. Let $G = (V, E)$ be a finite graph. Then

- every element of the cycle space is the (edge-)disjoint union of circuits;
- an edge set $F \subseteq E$ lies in the cycle space if and only if every vertex in (V, F) has even degree;
- an edge set $F \subseteq E$ lies in the cycle space if and only if F meets every cut of G in an even number of edges; and

- the cycle space is closed under taking sums.

The last property is so obvious that it is normally not even considered. In infinite graphs, however, it needs proof. Diestel and Kühn [35] demonstrate that the topological cycle space is closed under taking thin sums, and in Chapter 6 we shall revisit this issue in a more general setting.

The first two items on the list become non-trivial theorems in infinite graphs. Moreover, they turn out to be very valuable when working with infinite circuits. Especially Theorem 1.9 is used over and over in many arguments as it allows to dispense with the underlying topology. Whenever it becomes necessary to identify some edge set as a circuit or an element of the cycle space it is much easier to simply test whether it meets every finite cut evenly than to construct a homeomorphism to the unit circle.

Theorem 1.8 (Diestel and Kühn [35])**.** *Every element of the cycle space of a graph satisfying* (1.1) *is the (edge-)disjoint union of circuits.*

Theorem 1.9 (Diestel and Kühn [35])**.** *Let Z be a set of edges in a graph G satisfying* (1.1)*. Then $Z \in \mathcal{C}(\tilde{G})$ if and only if Z meets every finite cut of G in an even number of edges.*

Let us consider the remaining property on the list above. In locally finite graphs every vertex is incident with an even number of edges in any element of the topological cycle space. The converse, however, is no longer true. That is, an edge set that induces an even graph does not necessarily lie in the topological cycle space. Consider, for instance, the edge set of a double ray contained in some graph. Depending on whether the two disjoint tails of the double ray lie in the same end or in distinct ends the edge set lies in the topological cycle space or not. Yet the degrees of the vertices in the double ray are always two and, thus, cannot distinguish between the two cases. What is needed is some measure on the ends, an *end degree*, that together with the vertex degrees decides whether a given edge set is an element of the topological cycle space. We will address this problem in Chapter 3.

Chapter 2

Duality of ends

2.1 Duality in infinite graphs

In 1932 Whitney [92] introduced the concept of dual graphs: a graph G^* is a dual of a finite graph G if there exists a bijection $* : E(G) \to E(G^*)$ so that a set $F \subseteq E(G)$ is a circuit of G precisely when F^* is a bond in G^*. (A *bond* is a minimal non-empty cut.) Dual graphs allow to formulate a criterion for planarity.

Theorem 2.1 (Whitney [92])**.** *A finite graph G has a dual if and only if it is planar.*

Two further key properties of dual graphs are symmetry and uniqueness. That is, a graph is the dual of its dual (symmetry), and a planar graph has exactly one dual, provided the graph is 3-connected.

Duality for infinite graphs has first been explored by Thomassen [79]. Faced with the incongruity that an infinite graph may have infinite cuts as well as finite ones but (in the traditional definition) only finite circuits he chose to ignore infinite cuts. Consequently, G^* is a dual of G, in the sense of Thomassen, if for all *finite* sets $F \subseteq E(G)$, F is a circuit precisely when F^* is a bond. This concept allowed him to prove an infinite version of Whitney's planarity criterion: a 2-connected graph G has a (Thomassen-)dual if and only if it is planar and satisfies (1.1), i.e. that every two vertices of G can be separated by finitely many edges.

Thomassen's definition is not completely satisfactory, as the symmetry in taking duals is lost, as well as the uniqueness of the duals of 3-connected graphs. These deficits are ultimately due to the disregard of infinite cuts. Consider the graph G, the half-grid in solid lines, in Figure 2.1. Geometrically, the dotted graph G^* should be its dual, and indeed it is a dual in

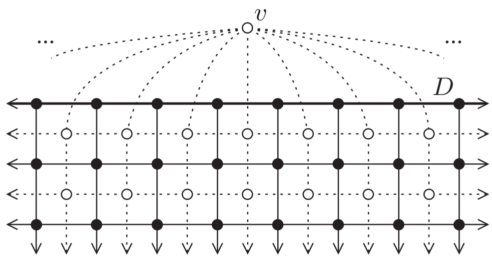

Figure 2.1: A pair of dual graphs

Thomassen's sense. We can obtain a second Thomassen-dual by splitting the vertex v into two vertices u and w and making each of these adjacent with infinitely neighbours of v so that every neighbour is adjacent to exactly one of u and w. Such a Thomassen-dual H violates the uniqueness of the dual for a 3-connected graph, in this case the half-grid. Furthermore, symmetry is broken too, since G is no longer a dual of H. In fact, H might not even be planar depending on how we split up the neighbours of v between u and w.

By considering infinite as well as finite circuits, however, we can restore uniqueness. In our example, consider the edge set F of the double ray D in G. In G^*, its dual set F^* (the set of edges incident with v) is a bond. But F^* is not a bond in H, because it contains the edges incident with u (say) as a proper subset. Thus, if F counts as a circuit, then G^* will be a dual of G but H will not, as should be our aim. Taking the circuits of G in $|G|$ achieves this.

Figure 2.1 underlines a second point. Namely, the class of locally finite graphs is not closed under taking dual graphs. Fortunately, as Thomassen showed for his weak definition of dual graphs, a necessary condition to have a dual is (1.1). Moreover, we shall see below that for dual graphs (considering infinite circuits) condition (1.1) holds too.

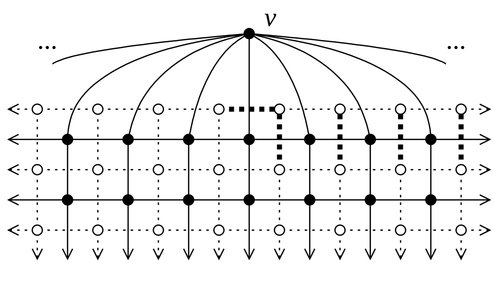

Figure 2.2: Rays starting in v have to be circuits

In Chapter 1, we have introduced two topological spaces, $|G|$ and $\tilde{G}$, for a graph G satisfying (1.1). Which of those should we choose as basis for our definition of duality? In Figure 2.2 the roles of the graphs in Figure 2.1 are reversed. We start with the graph G drawn in solid lines and want to convince ourselves that the half-grid is its dual. Now, the bold edges F^* in the half-grid form a bond, so the corresponding edges in G should yield a circuit. Clearly, as F as the edge set of a ray starting in v it cannot be a circuit with respect to $|G|$. In $\tilde{G}$, however, F is a circuit as v dominates the end of the ray. In conclusion, the example indicates that we have to work in $\tilde{G}$.

Let G be a graph satisfying (1.1). Let G^* be another graph, with a bijection ${}^*\colon E(G) \to E(G^*)$. Call G^* a *dual* of G if the following holds for every set $F \subseteq E(G)$, finite or infinite: F is a circuit in $\tilde{G}$ if and only if F^* is a bond in G^*.

Requiring correspondence between all bonds and all circuits (finite or infinite) restores symmetry and uniqueness of duals:

Theorem 2.2. **[16]** *Let G be a countable graph satisfying* (1.1)*.*

(i) *G has a dual if and only if G is planar.*

(ii) *If G^* is a dual of G, then G^* satisfies* (1.1)*, G is a dual of G^*, and this is witnessed by the inverse bijection of *.*

(iii) *If G is* 3*-connected then it has at most one dual, up to isomorphism.*

Another well-known feature of dual graphs is the duality of spanning trees. Given a pair of finite connected dual graphs G and G^*, a set D is the edge set of a spanning tree of G, if and only if $E(G^*) \setminus D^*$ is the edge set of a spanning tree in G^*. Moreover, if a pair of graphs has this kind of correspondence of spanning trees then they are dual to each other.

With ordinary spanning trees the duality of trees fails in infinite graphs. Indeed, the edge set of a spanning tree is allowed to contain an infinite circuit C. Then C^* is a bond, and $G^* - C^*$ therefore disconnected and cannot contain any spanning tree of G^*. In Chapter 1 we introduced topological spanning trees, which are better suited for dealing with infinite circuits. Unfortunately, a duality of TSTs is impossible too. The main issue here is that a TST (or rather the subgraph consisting of the vertices and edges of the TST) does not need to be connected in a graph-theoretical sense, it merely needs to be topologically connected. Thus, a TST with edge set T might miss a (then necesssarily infinite) bond C. The complement $E(G^*) \setminus T^*$ of T in the dual graph will then contain the circuit C^*, but a TST cannot contain a circuit.

Instead of working with spanning trees or with TSTs we will formulate our duality of trees with trees that marry the properties of both. We call an (ordinary) spanning tree of a graph G *acirclic* if it does not contain any (infinite) circuit of $\tilde{G}$. Observe that the closure of an acirclic spanning tree is a TST of $\tilde{G}$.

Theorem 2.3. **[16]** *Let $G = (V, E)$ and $G^* = (V^*, E^*)$ be connected graphs satisfying* (1.1)*, and let* $^*\colon E \to E^*$ *be a bijection. Then the following two assertions are equivalent:*

(i) *G and G^* are duals of each other, and this is witnessed by the map * and its inverse.*

(ii) *Given a set $F \subseteq E$, the graph (V, F) is an acirclic spanning tree of G if and only if $(V^*, E^* \setminus F^*)$ is an acirclic spanning tree of G^* (both in $\tilde{G}$).*

We shall come back to the relation of spanning trees with duals in Chapter 8. Moreover, it will play a brief role in the proofs in Section 2.4.

2.2 Duality of ends

The main aim of this chapter, which is based on [22], is the study of the relation between the end space of a graph and the end space of its dual. Our first result states that there exists a homeomorphism between these two spaces that arises in a natural way from the bijection $*$ on the edges.

More precisely, we will demonstrate that, given a pair G, G^* of (infinite) duals, the endvertices of a set $F \subseteq E(G)$ converge towards an end ω of G if and only if the endvertices of F^* converge towards the dual end ω^*. This is the content of Theorem 2.7, which we will discuss below.

Thick ends, those that contain an infinite set of disjoint rays, play an important role in the study of the automorphism group of a graph, see for instance Halin [51]. As our second result, we will prove that thickness is preserved in the dual end:

Theorem 2.4. **[22]** *Let G, G^* be a pair of dual graphs, and let ω be an end of G. Then ω is thick if and only if ω^* is thick.*

In order to prove Theorem 2.4, we make use of a notion of connectivity, introduced by Tutte [85], that coincides with the matroid connectivity of the cycle-matroid of the graph. As a by-product we obtain a generalisation to infinite graphs of the following classical result:

Theorem 2.5 (Tutte [85]). *Let G and G^* be a pair of finite dual graphs, and let $k \geq 2$. Then G is k-Tutte-connected if and only if G^* is k-Tutte-connected.*

We will define Tutte-connectivity in Section 2.4 (all other definitions can be found in the next section), but let us remark here that a graph is 3-Tutte-connected if and only if it is 3-connected. Therefore, Theorem 2.16, our extension of Tutte's theorem, has the following consequence:

Corollary 2.6 (Thomassen [79]). *Let G and G^* be a pair of dual graphs. Then G is 3-connected if and only if G^* is 3-connected.*

Let us make our aims for this chapter more precise. We start with the bijection we wish to define between the end spaces of dual graphs. Our mapping will be an extension of the bijection $* : E(G) \to E(G^*)$ on the edges (and we will therefore, slightly abusing notation, denote it with $*$ as well). More precisely, we aim at a bijection $*$ between $\Omega(G)$ and $\Omega(G^*)$, so that for all $F \subseteq E(G)$, the endvertices of F converge against an end ω of G if and only if the endvertices of F^* converge against ω^*.

In the space $\tilde{G}$, which is instrumental in the definition of duality, the accumulation points of vertex sets are the identification classes of ends. Recall that any two ends that cannot be separated by finitely many edges, are identified, giving rise to larger equivalence classes of rays called *edge-ends* by some authors (e.g. Hahn, Laviolette and Širáň [50]). So, should we not search for a bijection of the edge-ends rather than of the ends?

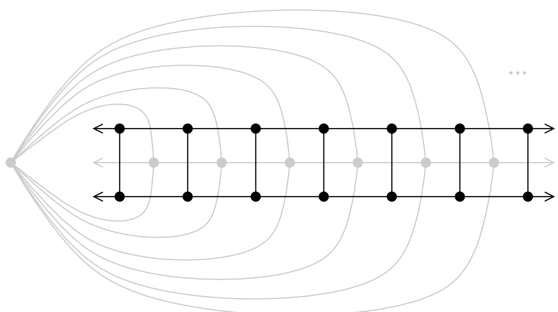

Figure 2.3: No correspondence between edge-ends of duals

Figure 2.3 demonstrates that there is no hope for a bijection between edge-ends (even without any structural requirements). The double ladder has two edge-ends, while its dual graph has only one edge-end.

The reason that this attempt fails lies in the nature of duals. The existence of finite edge-cuts between (edge-)ends will not be preserved in the dual. In fact, such a (minimal) cut corresponds to a circuit in the dual, which need not separate anything. By contrast, a vertex-separation whose deletion results in two sufficiently large sides does, in some sense, carry over to the dual graph; this is the essence of Theorem 2.16 and will be more explored in Section 2.4.

Our bijection will thus be between the ends of G and G^*. This means that we will work in $|G|$, since any two identified ends cannot be distinguished topologically in $\tilde{G}$. Endowing $\Omega(G)$ resp. $\Omega(G^*)$ with the subspace topology of $|G|$ resp. $|G^*|$, we will show the existence of a bijection $\Omega(G) \to \Omega(G^*)$, which is structure-preserving in the sense above. Moreover, we will see that $*$ is a homoeomorphism:

Theorem 2.7. **[22]** *Let G and G^* be 2-connected dual graphs. Then there is a homeomorphism $* : \Omega(G) \to \Omega(G^*)$, where the two spaces are endowed with the subspace topology of $|G|$ resp. $|G^*|$, so that*

$$\text{for all } F \subseteq E(G) \text{ and ends } \omega \text{ it holds that } \omega \in \overline{F} \text{ if and only if } \omega^* \in \overline{F^*}. \tag{2.1}$$

We remark that the requirement that G and G^* are 2-connected cannot be dropped. This is illustrated by the example of the double ray. Every dual of the double ray is a graph whose edge set is the union of countably many loops, and thus contains no end at all.

We shall prove Theorem 2.7 in the next section.

Let us now turn to our second objective: showing that our bijection $*$ preserves thickness. This will be achieved in Theorem 2.4. Again, we are confronted with the question why focus on preserving (vertex-)thickness instead of "edge-thickness", i.e. the existence of infinitely many edge-disjoint rays in an end.

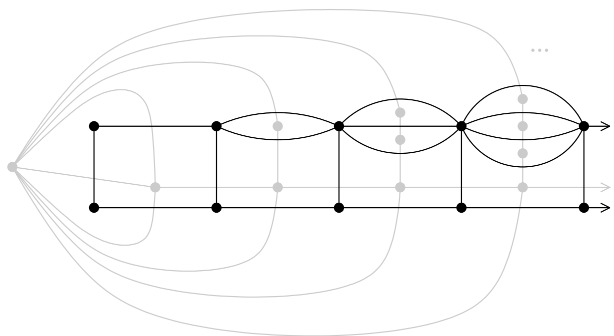

Figure 2.4: Edge-thick end with edge-thin dual end

This is answered by Figure 2.4, which shows a graph that has a single edge-thick end while the unique end of its dual graph does not even possess two edge-disjoint rays. The reason is the same as above: although (or because) the notion of duals is based on edges and operations with edges, the existence of (small) edge-separators is not preserved in the dual.

Since not all vertex-separators are preserved in the dual, connectivity is not an invariant of (finite or infinite) duals, as we have already remarked in the introduction. But, the related notion of *Tutte-connectivity* is. We defer to Section 2.4 for the definition; suffice it to say here that there are two reasons why a graph may have low Tutte-connectivity: Either it has a small vertex-separator or it contains a small circuit. In Section 2.4, we prove that Tutte-connectivity is an invariant of infinite duals, too (Theorem 2.16).

Theorem 2.16 is an important step on our way to proving Theorem 2.4. Our proof of Theorem 2.16 differs from the usual proof of its finite version, Theorem 2.5, which is done in two steps. First, one shows that Tutte-connectivity coincides with the connectivity of the cycle-matroid of the graph. Then one observes that matroid connectivity is invariant under duality.

If we want to use this approach for Theorem 2.16 as well, we first have to answer two questions. Which notion of infinite matroids should we use? And how do we define higher connectivity in a matroid?

The first question is easy to answer. Although it is sometimes claimed that there is no proper concept of an infinite matroid that provides duality and the existence of bases at the same time, B-matroids, as defined by Higgs [54], accomplish that (see also Oxley [67]). Moreover, one can prove that duality in B-matroids is compatible with taking dual graphs. While the second problem, the definition of higher connectivity, can also be overcome in a satisfactory way, its solution together with the introduction of B-matroids would take quite a bit of time and effort. Therefore, we will, in Section 2.4, present a matroid-free proof of Theorem 2.4. We shall return to B-matroids as well as to the relation between matroid-connectivity and graph-connectivity in Chapter 8

2.3 $*$ induces a homeomorphism on the ends

Before we are able to prove Theorem 2.7, we need three lemmas. The proof of the first lemma, which is not hard, can be found in [31, Lemma 8.2.2].

Lemma 2.8. *Let G be a connected graph, and let U be an infinite subset of $V(G)$. Then G contains a ray R with infinitely many disjoint R–U paths or a subdivided star with infinitely many leaves in U.*

Lemma 2.9. [22] *Let G be a* 2*-connected graph satisfying* (1.1)*. If U is an infinite set of vertices then $\overline{U}$ contains an end of G.*

Proof. Suppose otherwise. Then there is no ray R in G with infinitely many disjoint R–U paths. So, an application of Lemma 2.8 yields a subdivided star S that contains an infinite subset U' of U. We delete the centre of S and apply Lemma 2.8 again, this time to U', which yields another subdivided star S' with infinitely many leaves in U'. But then, the centre of S and the centre of S' are infinitely connected, contradicting (1.1). $\square$

As a convenience we will, for a set F of edges, write $V[F]$ to denote the set of endvertices of the edges in F.

Lemma 2.10. **[22]** *Let G be a 2-connected graph, and let X and Y be two sets of edges such that $\overline{X} \cap \overline{Y} \cap \Omega(G) \neq \emptyset$. Then there are infinitely many (edge-)disjoint finite circuits each of which meets both X and Y.*

Proof. Let $\mathcal{Z}$ be an $\subseteq$-maximal set of finite disjoint circuits so that each $C \in \mathcal{Z}$ meets both X and Y, and suppose that $|\mathcal{Z}|$ is finite. Putting $Z := \bigcup \mathcal{Z}$, we pick for every two $x, y \in V[Z]$ for which it is possible an x–y path $P_{x,y}$ that is edge-disjoint from Z. Denote by Z' the union of Z with the edge sets of all these paths, and observe that still $|Z'| < \infty$.

We claim that for every component K of $G - V[Z']$ it holds that

$$\text{for every } v, w \in N(K) \text{ there is a } v\text{–}w \text{ path in } G[Z'] - Z. \tag{2.2}$$

Indeed, by construction, there are $x, y \in V[Z]$ and (possibly trivial) v–x resp. w–y paths Q_v resp. Q_w in $G[Z'] - Z$. Then x and y are connected through $K \cup Q_v \cup Q_w \subseteq G - Z$. Hence in $P_{x,y} \cup Q_v \cup Q_w \subseteq G[Z'] - Z$ we find a v–w path. This proves (2.2).

Now, because $\overline{X} \cap \overline{Y}$ contains an end, there exists a component K of $G - V[Z']$ which contains infinitely many vertices of both $V[X]$ and $V[Y]$. Choose edges $e_X, e_Y \in E(K) \cup E(K, G - K)$ so that $e_X \in X$, and $e_Y \in Y$. Since G is 2-connected, there is a finite circuit C which contains both e_X and e_Y. The maximality of $\mathcal{Z}$ implies that C meets Z in at least one edge. In particular, C contains the edge sets of (possibly identical) $N(K)$-paths P_X and P_Y so that $e_X \in E(P_X)$, and $e_Y \in E(P_Y)$.

Being connected, K contains a $V(P_X)$–$V(P_Y)$ path P. Thus, we find in $P \cup P_X \cup P_Y$ an $N(K)$-path P' with $e_X, e_Y \in E(P')$. By (2.2), there exists a path R in $G[Z'] - Z$ between the endvertices of P'. Now, $E(P') \cup E(R)$ is a circuit that meets both X and Y but is edge-disjoint from Z, a contradiction to the maximality of $\mathcal{Z}$. $\square$

Proof of Theorem 2.7. We start by claiming that for each $F \subseteq E(G)$ and each end ω of G the following is true:

$$\text{if } \overline{F} \cap \Omega(G) = \{\omega\} \text{ then } \overline{F^*} \text{ contains exactly one end.} \tag{2.3}$$

Suppose the claim is not true. By Lemma 2.9, this cannot be because $\overline{F^*}$ fails to contain an end; rather there must be (at least) two ends, α_1 and α_2, in $\overline{F^*}$. Take a finite connected subgraph T of G^* so that $V(T)$ separates α_1 and α_2 in G^*. For $i = 1, 2$, denote by K_i the component of $G^* - T$ to which α_i belongs, and set $X_i^* := (E(K_i) \cup E(K_i, T)) \cap F^*$. Since each of the X_i^* is infinite, it follows from Lemma 2.9 that $\overline{X_i}$ contains an end. As $\overline{X_i} \subseteq \overline{F}$, this end must be ω. Hence, Lemma 2.10 yields disjoint finite circuits $C_1, C_2, \ldots$ in G each of which meets X_1 as well as X_2.

We claim that each of the bonds C_i^* contains an edge of T. Indeed, let M_1 and M_2 be the two components of $G^* - C_i^*$. Since C_i^* meets both X_1^* and X_2^*, each M_j contains a vertex in $K_1 \cup T$ and a vertex in $K_2 \cup T$. As, for $j = 1, 2$, M_j is connected it follows that $V(M_j) \cap V(T) \neq \emptyset$. So, since T is connected, there is an M_1–M_2 edge in $E(T)$, i.e. $C_i^* \cap E(T) \neq \emptyset$, for each $i \in \mathbb{N}$. This yields a contradiction since the C_i^* are disjoint but T is finite. Therefore, Claim (2.3) is established.

Now, we define $* : \Omega(G) \to \Omega(G^*)$. Given an end $\omega \in \Omega(G)$, pick any set $F \subseteq E(G)$ with $\overline{F} \cap \Omega(G) = \{\omega\}$ (choose, for instance, the edge set of a ray in ω). Define $\omega^* = \omega^*(F)$ to be the, by (2.3), unique end in $\overline{F^*}$. To see that this mapping is well-defined, i.e. that it does not depend on the choice of F, consider a second set $D \subseteq E(G)$ as above, and observe that $\omega^*(D) = \omega^*(D \cup F) = \omega^*(F)$. Since G is a dual of G^* (Theorem 2.2 (ii)), we may apply (2.3) to G^* and see that $*$ is a bijection and satisfies (2.1).

Next, we prove that $* : \Omega(G) \to \Omega(G^*)$ is continuous. For this, let an end $\omega^* \in \Omega(G^*)$ and an open neighbourhood $U^* \subseteq \Omega(G^*)$ of ω^* be given. Then there exists a finite vertex set $S \subseteq V(G^*)$, and a component K of $G^* - S$ so that $W^* := K \cap \Omega(G^*) \subseteq U^*$.

Setting $F^* := E(G^*) \setminus (E(K) \cup E(S, K))$, we observe that $W^* = \Omega(G^*) \setminus \overline{F^*}$. Hence, by (2.1), $W = \Omega(G) \setminus \overline{F}$. So, W is an open neighbourhood of ω whose image is contained in U^*. Finally, by interchanging the roles of G and G^* we see that the inverse of $*$ is continuous as well. $\square$

2.4 Tutte-connectivity

In this and in the next section, we are concerned with how (Tutte-) connectivity is preserved in the dual. The main idea underlying our proofs is the duality of spanning trees, to which we already alluded to in Section 2.1. We will use the tree duality implicitly in the key lemma, Lemma 2.14, below. The next two lemmas help to relate the tree duality to vertex separations.

We shall need to work within both spaces $|G|$ and $\tilde{G}$. In order to distinguish between closures of sets $X \subseteq V(G) \cup E(G)$ in the two spaces, we write

$\overline{X}$ for the closure of X in $|G|$, and $\widetilde{X}$ for the closure of X in $\tilde{G}$. This notation will only be valid in this chapter.

Lemma 2.11. **[22]** *Let G be a graph satisfying* (1.1)*, let T be a subgraph that does not contain any circuits, and let $U \subseteq V(T)$ such that $0 < |U| < \infty$. Then there exists a set $F \subseteq E(T)$ of size at most $|U| - 1$ so that every arc in $\widetilde{T}$ between two vertices in U meets F.*

Proof. We use induction on $|U|$. The assertion is trivial for $|U| = 1$, so for the induction step asume that $|U| > 1$. Choose $v \in U$, then by the induction assumption there is a set $D \subseteq E(T)$ such that each vertex w of $U \setminus \{v\}$ lies in a different path-component K_w of $\widetilde{T} - D$. If there is no vertex $w \in U \setminus \{v\}$ such that $v \in K_w$, we are done, so assume there is such a w.

Observe that there exists exactly one v–w arc A in $\widetilde{T} - D$. Indeed, if there were two, then it is easy to see that the edge set of their union would contain a circuit. Now, choose any edge e on A, and set $F := D \cup \{e\}$. Clearly, F is as desired, which completes the proof. $\square$

Lemma 2.12. **[22]** *Let H be a connected graph, let $F \subseteq E(H)$, and let $W \subseteq V(H)$. If every W-path in H meets F then $|F| \geq |W| - 1$.*

Proof. Since no two vertices of W can lie in the same component of $H - F$, we deduce that $H - F$ has at least $|W|$ components. As each deletion of a single edge increases the number of components by at most one, $H - F$ can have at most $|F| + 1$ components. $\square$

Let us now introduce the notion of Tutte-connectivity, see Tutte [85]. For finite graphs, the Tutte-connectivity coincides with the connectivity of the cycle-matroid of the graph. We remark that for $k \in \{2, 3\}$, a graph is k-Tutte-connected if and only if it is k-connected. For greater k the two notions of connectivity are not equivalent.

Definition 2.13. *A* k-Tutte-separation *of a graph G is a partition (X, Y) of $E(G)$ so that $|X|, |Y| \geq k$ and so that at most k vertices of G are incident with edges in both of X and Y.*
We say that a graph G is k-Tutte-connected *if G has no ℓ-Tutte-separation for any $\ell < k$.*

Consider a k-Tutte-separation (X, Y) in a (2-connected) graph G with a dual G^*. To prove that Tutte-connectivity is invariant under taking duals, we would ideally like to see that (X^*, Y^*) is a k-Tutte-separation in G^*. This, however, is not always true—if the two sides of the separation do not induce connected subgraphs of G^*, then the number of vertices in $V[X^*] \cap V[Y^*]$ can be much higher than k. Thus we will strengthen the requirements and lessen our expectations. By demanding $(V[Y], Y) - V[X]$ to be connected, we shall be able to guarantee that at least $(V[Y^*], Y^*)$ is connected. Moreover, we will be content with finding an ℓ-Tutte-separation of G^* for some $\ell \leq k$ that is derived from (X^*, Y^*).

The statement of the next lemma, which accomplishes just that, is a bit more general than we need for Theorem 2.16, as we shall reuse it for Theorem 2.4.

Lemma 2.14. **[22]** *Let G and G^* be a pair of 2-connected dual graphs, and let (X, Y) be a k-Tutte-separation such that $C_Y := (V[Y], Y) - V[X]$ is non-empty and connected, and such that $Y = E(C_Y) \cup E(C_Y, V[X])$. Then*

(i) *there exists a component L of $(V[X^*], X^*)$ so that $(E(L), E(G^*) \setminus E(L))$ is an ℓ-Tutte-separation for some $\ell \leq k$; and*

(ii) *for each component K of $(V[X^*], X^*)$ with $|E(K)| \geq k$ it holds that $(E(K), E(G^*) \setminus E(K))$ is a k-Tutte-separation.*

In order to prove the lemma we need a simple fact that follows easily from the observation that every two edges lie in a common circuit if and only if the graph is 2-connected, which is the case precisely when every two edges lie in a common bond. Variants of this lemma can be found in Thomassen [78] as well as in [16].

Lemma 2.15. *Let G and G^* be a pair of dual graphs. Then G is 2-connected if and only if G^* is 2-connected.*

Proof of Lemma 2.14. First, we prove that

$$\widetilde{Y^*} \text{ is path-connected in } \tilde{G}^*. \tag{2.4}$$

Suppose that this is not the case. Then we can write Y as the disjoint union of two sets Y_1 and Y_2 so that there is no $V[Y_1^*]$–$V[Y_2^*]$ arc in $\tilde{G}^*$ that only uses edges from Y^*.

In particular, there is no circle in $\tilde{G}^*$ that only uses edges from Y^* and meets both Y_1^* and Y_2^*. Equivalently, there is no bond in G that only uses edges from Y, and meets both Y_1 and Y_2.

However, since C_Y is connected and since every edge in Y is incident with a vertex in C_Y, there is a vertex $x \in V(C_Y)$ which is incident with both Y_1 and Y_2. Observe that the cut B_x of G, which consists of all edges incident with x, is a subset of Y. As G is 2-connected, B_x is a bond, which yields the desired contradiction and thus proves (2.4).

Now, set $U := V[X] \cap V[Y]$ and $W := V[X^*] \cap V[Y^*]$. Observe that each vertex in W is incident with both X^* and Y^*. So, if $|W|$ is infinite, then Lemma 2.9 implies that $\overline{X^*} \cap \overline{Y^*}$ contains an end, while $\overline{X} \cap \overline{Y}$ does not (as X and Y are finitely separated by U). This contradicts Theorem 2.7. We have thus shown that

$$|W| \textit{ is finite.} \tag{2.5}$$

Let T_X be the edge set of a maximal topological spanning forest of $\widetilde{X}$, i.e. the union of TSTs of the spaces $\tilde{C}$ corresponding to the components C of $(V[X], X)$. We point out that every circuit of G that lies entirely in X is a circuit of $(V[X], X)$. It follows that T_X does not contain any circuits of G.

Next, we prove that

$$\textit{every } W\textit{-path in } (V[X^*], X^*) \textit{ meets } T_X^*. \tag{2.6}$$

Suppose there is a W-path whose edge set D^* lies in $X^* \setminus T_X^*$. By (2.4), there is a circuit C^* of G^* with $C^* \cap X^* = D^*$. Thus, C is a bond in G, and hence D is a finite cut of $(V[X], X)$. Consequently, D contains a bond B of $(V[X], X)$, which then is completely contained in one component K_B of $(V[X], X)$. As $B \subseteq D \subseteq X \setminus T_X$, the intersection of B with T_X is empty. Thus, B is a finite cut of K_B that is disjoint from T_X but that separates two vertices incident with T_X. Since, on the other hand, $\widetilde{T_X}$ restricted to $\widetilde{K_B}$ is path-connected, we obtain a contradiction. This proves (2.6).

Next, Lemma 2.11 yields a set $F \subseteq T_X$ of at most $|U| - 1$ edges so that every U-arc in $\widetilde{T_X} \subseteq \widetilde{X}$ meets F. This means that every circuit C of G with $C \cap X \subseteq T_X$ meets F. Thus, every bond B^* of G^* with $B^* \cap X^* \subseteq T_X^*$ meets F^*. Hence, denoting by $\mathcal{K}$ the set of components of $(V[X^*], X^*)$, we obtain that

$$\textit{for every } K \in \mathcal{K}\textit{, the graph } H_K := K - (T_X^* \setminus F^*) \textit{ is connected.} \tag{2.7}$$

Now, for every $K \in \mathcal{K}$, observe that by (2.6), every W-path in H_K meets F^*. So, by (2.7), we may apply Lemma 2.12 to H_K. Doing so for each $K \in \mathcal{K}$, we obtain that $|F^*| \geq |W| - |\mathcal{K}|$. On the other hand, $|F^*| = |F| \leq |U| - 1$ by the choice of F, implying that

$$|W| \leq |U| + |\mathcal{K}| - 1. \tag{2.8}$$

Suppose that for every $K \in \mathcal{K}$, it holds that $|V(K) \cap W| > |E(K)|$. Then

$$|W| = \sum_{K \in \mathcal{K}} |V(K) \cap W| \geq \sum_{K \in \mathcal{K}} (|E(K)| + 1) = |X^*| + |\mathcal{K}|.$$

As $|X^*| = |X| \geq |U|$, we obtain that $|W| \geq |U| + |\mathcal{K}|$. This yields a contradiction to (2.8), since by (2.5), $|W|$ is finite. Therefore, there exists an $L \in \mathcal{K}$ with

$$\ell := |V(L) \cap W| \leq |E(L)|.$$

Observe that if we can show now that $\ell \leq k$, then it follows that the edge partition $(E(L), E(G^*) \setminus E(L))$ is an ℓ-Tutte-separation of G^*, as desired for (i). So, in order to prove (i), and (ii), it suffices to prove that for each $K \in \mathcal{K}$ it holds that

$$|V(K) \cap W| \leq |U|.$$

Suppose otherwise. Then there exists an $M \in \mathcal{K}$ such that

$$|W| = \sum_{K \in \mathcal{K}} |V(K) \cap W| \geq (|U| + 1) + \sum_{K \in \mathcal{K},\, K \neq M} |V(K) \cap W|.$$

Because G is 2-connected, so is G^* (Lemma 2.15). Thus $|V(K) \cap W| \geq 1$ for every $K \in \mathcal{K}$, resulting again in $|W| \geq |U| + |\mathcal{K}|$, a contradiction, as desired. □

Theorem 2.16. **[22]** *Let G and G^* be a pair of dual graphs, and let $k \geq 2$. Then G is k-Tutte-connected if and only if G^* is k-Tutte-connected.*

Proof. We show that if G has a k-Tutte-separation (X, Y), then G^* has an ℓ-Tutte-separation for some $\ell \leq k$. By Theorem 2.2 (ii), this is enough to prove the theorem.

First, assume that $V[Y] \setminus V[X] \neq \emptyset$. Let K be a component of $(V[Y], Y) - V[X]$, and set $Z := E(K) \cup E(K, G - K)$. As $E(K, G-K)$ contains at least one edge for each vertex in $N(K)$, it follows that $|Z| \geq |N(K)|$. Thus, $(Z, E(G) \setminus Z)$ is a k'-Tutte-separation of G for $k' := |N(K)| \leq k$. We can now apply Lemma 2.14 (i) to obtain the desired ℓ-Tutte-separation of G^*.

So, we may assume that $V[Y] \setminus V[X] = \emptyset$. Then, since $|Y| \geq k$, there is a circuit C in Y, say of length $\ell \leq k$. Hence, C^* is a bond of size ℓ in G^*; let K_1 and K_2 be the components of $G^* - C^*$. Now,

$$|E(K_1 \cup K_2)| = |X^*| + |Y^*| - |C^*| \geq 2k - \ell.$$

Thus, we can partition C^* into C_1^* and C_2^* so that each $Z_i^* := E(K_i) \cup C_i^*$ has cardinality at least ℓ.

In order to show that (Z_1^*, Z_2^*) is an ℓ-Tutte-separation of G^* it remains to check that $U := V[Z_1^*] \cap V[Z_2^*]$ has cardinality at most ℓ. To this end, consider a vertex $v \in U$, and let j be such that $v \in V(K_j)$. Then v is incident with an edge $e_v^* \in C_{3-j}^*$, whose other endvertex lies in K_{3-j}, because C^* is a cut. This defines an injection from $U \to C^*$, which implies $|U| \leq |C^*| \leq \ell$, as desired. □

2.5 The dual preserves the degrees

In this section we will use Lemma 2.14 in order to prove a quantitative version of Theorem 2.4, that relates the 'degree' of an end ω to the degree of its dual end ω^*.

For an end ω, define $m(\omega)$ to be the supremum of the cardinalities of sets of disjoint rays in ω; Halin [51] showed that this supremum is indeed attained. In [24] and in Stein [77] the number of vertex- (or edge-)disjoint rays in an end has been successfully used to serve as the degree of an end in a locally finite graph (whether vertex- or edge-disjoint rays should be considered depends on the application). This motivates the definition of the degree $d(\omega) := m(\omega)$ of an end ω of a locally finite graph. We will discuss end degrees in more depth in Chapter 3.

Now, if G and G^* are a dual pair of 2-connected locally finite graphs, then it will turn out that $m(\omega) = m(\omega^*)$ for every end ω of G. In non-locally finite graphs we need to be a bit more careful: Figure 2.4 indicates that dominating vertices should be taken into account.

For an end $\omega \in \Omega(G)$ and a finite vertex set S, we say that $U \subseteq V(G)$ *separates S from ω* if U meets every ray in ω that starts in S. We define here the *degree* $d(\omega)$ of an end $\omega \in \Omega(G)$ to be the minimal number k such that for each finite set $S \subseteq V(G)$, we can separate S from ω in G by deleting at most k vertices from G. If there is no such k, we set $d(\omega) := \infty$. Lemma 2.17 will show that this definition is consistent with the one given above for locally finite graphs.

So, denote by $\mathrm{dom}(\omega)$ the number of vertices that dominate an end ω (possibly infinite). Note that the graphs we are interested in, namely those that satisfy (1.1), are such that $\mathrm{dom}(\omega) \in \{0, 1\}$ for every end ω.

Lemma 2.17. **[22]** *Let G be a graph and let $\omega \in \Omega(G)$. Then $d(\omega) = m(\omega) + \mathrm{dom}(\omega)$.*

Proof. It is easy to see that $d(\omega)$ is at least $m(\omega) + \mathrm{dom}(\omega)$. For the other direction, we may assume that $\mathrm{dom}(\omega) < \infty$. Denote by D the set of vertices that dominate ω. As D is a finite set, there is an obvious bijection between the ends of $G - D$ and G, which we will tacitly use.

We observe first that for any finite vertex set T, there exists a finite T–ω separator T' in $G - D$ that is contained in $C_{G-D}(T, \omega)$. Indeed, otherwise, by Menger's theorem[1], $G[T \cup C(T, \omega)] - D$ contains infinitely many paths

[1]We use here, and below, that the cardinality version of Menger's theorem holds in infinite graphs. This can easily be deduced from Menger's theorem for finite graphs, see for instance [31, Section 8.4]

between T and some ray in ω that are pairwise disjoint except possibly in T. As T is finite, this implies that $T \setminus D$ contains a vertex which dominates ω, contradicting our choice of D.

Now, consider an arbitrary finite set $S \subseteq V(G)$. Starting with $S_0 := S \setminus D$ we can choose inductively finite vertex sets S_i so that $S_i \subseteq V(C_{G-D}(S_{i-1}, \omega))$ is an S_{i-1}–ω separator in $G - D$, and has minimal cardinality with that property. Since $S_i \subseteq V(C_{G-D}(S_{i-1}, \omega))$, all the S_i are pairwise disjoint.

Applying Menger's theorem repeatedly between S_{i-1} and S_i we obtain a set $\mathcal{R}$ of disjoint rays in ω of cardinality at least $|S_1|$. As $S_1 \cup D$ separates S from ω in G, we have shown that S can be separated from ω by at most $|S_1| + |D| \leq m(\omega) + \mathrm{dom}(\omega)$ vertices, thus proving the lemma. $\square$

We remark that Lemma 2.17 can be obtained easily from results of Polat [69]; we chose to provide the proof nevertheless since the statement of Polat's results together with the necessary adaptions would have taken about as much time and space.

Theorem 2.18. **[22]** *Let G and G^* be a pair of 2-connected dual graphs, and let ω be an end of G. Then $d_G(\omega) = d_{G^*}(\omega^*)$.*

Proof. First assume that $d(\omega) \leq k$, where $k \in \mathbb{N}$ is a finite number. We wish to show that ω^* has vertex-degree $\leq k$, too.

So, let a finite vertex set $T \subseteq V(G^*)$ be given. Pick a finite edge set F^* of cardinality at least k so that $T \subseteq V[F^*]$ and so that F^* induces a connected graph. Now, since $d(\omega) \leq k$ there is a set $U \subseteq V(G)$ of cardinality at most k that separates (the finite set) $V[F]$ from ω. If C is the component of $G - U$ to which ω belongs then set $Y := E(C) \cup E(C, U)$ and $X := E(G) \setminus Y$. Because $k \geq |U| = |V[X] \cap V[Y]|$, and because $|Y| = \infty$ and $|X| \geq |F| \geq k$, it follows that (X, Y) is a k-Tutte-separation.

Since $(V[F^*], F^*) \subseteq (V[X^*], X^*)$ is connected, there is a component K of $(V[X^*], X^*)$ that contains all of F^*. As $|F^*| \geq k$, Lemma 2.14 (ii) implies that $(E(K), E(G^*) \setminus E(K))$ is a k-Tutte-separation. Moreover, as $\omega \notin \overline{X}$, it follows that $\omega^* \notin \overline{K}$. Thus, $N_{G^*}(G^* - K)$ is a vertex set of cardinality $\leq k$ that separates $T \subseteq V[F^*]$ from ω^*, as desired.

In conclusion, since G is also a dual of G^* (Theorem 2.2 (ii)), it follows that $d(\omega) = d(\omega^*)$ if either of ω and ω^* has finite degree. In the remaining case, we trivially have $d(\omega) = \infty = d(\omega^*)$. $\square$

The theorem in conjunction with Lemma 2.17 immediately yields Theorem 2.4.

Chapter 3

The cycle space and end degrees

3.1 Degrees of ends

As outlined in Chapter 1 the topological cycle space makes it possible to extend facts about circuits in finite graphs to infinite graphs. However, one of the most basic and simple results characterising the cycle space of a finite graph did not so far have an analogue in locally finite graphs. For an edge set Z let us call a vertex Z*-even* if it is incident with an even number of edges in Z.

Proposition 3.1. *Let G be a finite graph. Then an edge set $Z \subseteq E(G)$ is an element of the cycle space if and only if every vertex in G is Z-even.*

Easy examples show that the proposition, as it is, cannot carry over to infinite graphs. In Figure 3.1, all the vertices in the double rays D and D', for instance, have degree 2, yet $E(D)$ is a circuit but $E(D')$ is not.

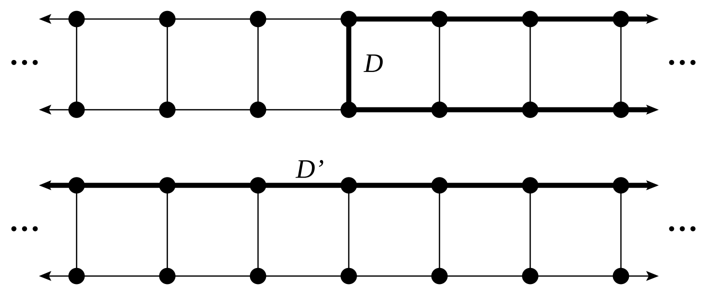

Figure 3.1: The edge set of D but not of D' is a circuit

The key difference between D and D' obviously lies in their behaviour at the ends of the double ladder. In [34] it was proposed to find a suitable definition for the *degree of an end* that captures this behaviour. Such a definition has been offered in [24], where the degree of an end is defined to be the maximal number of edge-disjoint rays in the end. Should this degree be finite, we call an end *even* if the end-degree is an even number, otherwise it is *odd*. For ends of infinite degree it is still possible to assign a parity, even or odd; we defer the details to Section 3.2. The concept of an end-degree allows to prove the following theorem, the main result of [24].

Theorem 3.2. **[24]** *Let G be a locally finite graph. Then $E(G)$ is an element of the topological cycle space of G if and only if every vertex and every end of G has even degree.*

To see whether a subset Z of the egdes of a finite graph $G = (V, E)$ is an element of the cycle space, we evidently need to check the degree a given vertex has in the subgraph (V, Z), whereas the degree in the whole graph is irrelevant. In the same way, if we want to extend Proposition 3.1, we have to measure the degree of an end with respect to Z. Such an end-degree, that classifies ends as Z*-even* or as Z*-odd*, has been introduced in [24], which allowed to formulate the following conjecture.

Conjecture 3.3. **[24]** *Let G be a locally finite graph, and let $Z \subseteq E(G)$. Then Z is an element of the topological cycle space of G if and only if every vertex and every end of G is Z-even.*

The purpose of this chapter, which is based on [11], is to give a proof of the conjecture, which will be achieved over the course of Sections 3.3 and 3.4. In Section 3.2 we briefly discuss and define the notion of an end-degree.

3.2 End-degrees in subgraphs

In this section, let us first give the formal definition of the degree of an end with respect to the whole graph. In a second step we then refine the definition, so that it applies to subgraphs as well. We follow here the exposition of [24], where a more thorough discussion can be found.

In Chapter 3 we have already encountered a notion of an end-degree. There we considered the degree of an end to be the maximal number of disjoint rays of the end. In this context, the definition turns out to be unsuitable as the graph in Figure 3.2 demonstrates. Every vertex has even degree, and both ends admit at most two disjoint rays. Since thus all vertices and ends have even degree we would assume that the whole edge set is contained in its topological cycle space. This, however, is not the case as the graph has an odd cut (see Theorem 1.9).

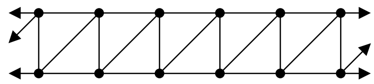

Figure 3.2: Counting disjoint rays does not work

There is an obvious alternative to defining the end-degree by counting disjoint rays, namely by counting *edge-disjoint* rays. This leads to two different definitions of the end-degree, each of which has its merits. Indeed, in Chapter 3 we worked successfully with an end-degree based on disjoint rays, which is also used in Stein [77]. The graph in Figure 3.2 gives an indication that, in order to characterise the elements of the topological cycle space, we need to focus on edge-disjoint rays. While all vertices and ends have even degree if the notion of an end-degree rests on disjoint rays, each of the two ends has three edge-disjoint rays (but not four).

Let ω be an end of a locally finite graph G. In this chapter, the *end-degree of ω* is the supremum (in fact, this is a maximum) over the cardinalities of sets of edge-disjoint rays in ω, and we denote this (possibly infinite) number by $\deg(\omega)$.

For Theorem 3.2 the numerical value $\deg(\omega)$ is not important. Rather, it is essential whether ω can be said to be even or odd. Provided $\deg(\omega)$ is finite then it is obvious that ω should be even if and only if the number $\deg(\omega)$ is even. That raises the question what parity we should assign to an end ω of infinite degree. The graphs in Figure 3.3 demonstrate that we cannot call such an end always even or always odd. In both graphs all the vertices have even degree and all the ends have infinite degree. Yet, as can be easily checked with the help of Theorem 1.9, the edge set of the infinite grid lies in the cycle space, while for the graph H on the right, we have $E(H) \notin \mathcal{C}(H)$. Consequently, for Theorem 3.2 to become a true statement, the single end in the infinite grid has to be even, but the two ends in H should be odd.

The example indicates that we need to distinguish between ends that have infinite degree but are even and ends of odd-infinite degree. This is accomplished by the following definition. We call an end ω *even* if there exists a finite vertex set $S \subseteq V(G)$ so that for all finite vertex sets $S' \supseteq S$ it holds that the maximal number of edge-disjoint rays in ω starting in S' is even. Otherwise, the end is called *odd*.

Let us make two remarks, both of which are discussed in more detail in [24]. First, for an end ω of finite degree, the end is even according to

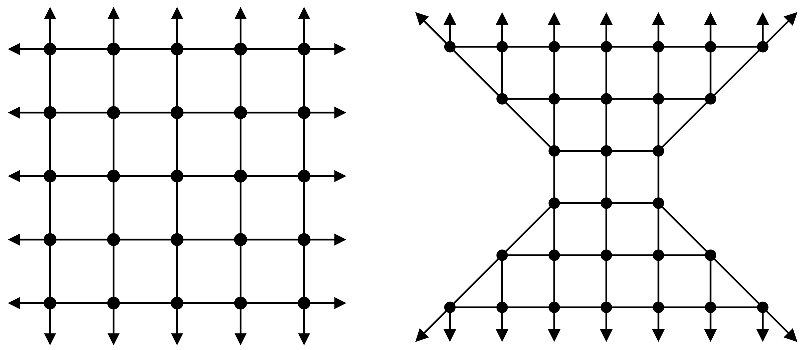

Figure 3.3: Even and odd ends of infinite degree

this definition if and only if $\deg(\omega)$ is an even number. Second, the choice of quantifiers for S and S' might appear arbitrary. Indeed, defining an end to be even with reversed quantifiers seems equally reasonable, i.e. an end ω would be even if *for all* finite vertex sets S *there exists* a finite superset S' such that the maximal number of edge-disjoint rays in ω starting in S' is even. While we suspect that such a definition of *weakly even* ends would still lead to Theorem 3.2 and Conjecture 3.3 to be true, we unfortunately cannot prove much in that respect.

Going back to Figure 3.3, we see that a choice of $S = \emptyset$ is sufficient for the single end of the infinite grid to be even. For any of the two ends of the graph on the right, however, it is not hard to check that as long as S' separates the two ends, the maximal number of edge-disjoint rays in the end starting in S' is odd, and hence the end itself is odd, as desired.

As outlined in Section 3.1, the degree of an end with respect to the whole graph is not of much use to us. To be able to decide whether a given edge set Z lies in the cycle space or not, we require a notion of an end-degree that takes Z into account. The key to adapting the notions introduced above, lies in substituting every occurrence of the word 'ray' by 'arc'.

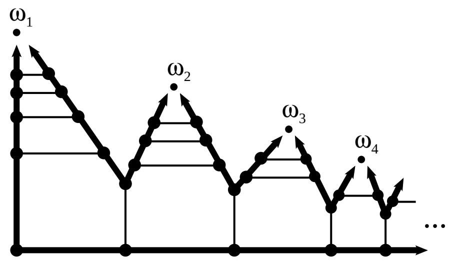

Figure 3.4: In subgraphs we have to count arcs

Figure 3.4 shows that we cannot keep counting rays. The set C of bold edges forms a circuit, yet any two rays of the end to the right contained in C share an edge. However, there are two edge-disjoint arcs in $\overline{C}$.

Denote by $\overline{Z}$ the closure of the point set $\bigcup_{z \in Z} z$ in $|G|$. Moreover, for $S, T \subseteq V(G) \cup \Omega(G)$, we say that A is an *S–T arc* if the first point of A lies in S, the last in T and no interior point in $S \cup T$. For $x \in V(G) \cup \Omega(G)$ we simply speak of *x–T arcs* instead of $\{x\}$–T arcs, and proceed analogously for other combinations of singeltons and sets.

We say that an end ω is *Z-even* if there exists a finite vertex set S so that for all finite vertex sets $S' \supseteq S$ it holds that the maximal number of edge-disjoint S'–ω arcs contained in $\overline{Z}$ is even. If an end is not Z-even, it is *Z-odd*. This definition is consistent with the definition of the end-degree in the whole graph, i.e. an end is $E(G)$-even if it is even. Moreover, if the maximal number N of edge-disjoint arcs contained in $\overline{Z}$ and ending in ω is finite, then ω is Z-even if and only if N is even. For both these facts, see [24].

Consider the double rays D and D' in Figure 3.1 again. There are two (edge-disjoint) arcs contained in $\overline{E(D)}$ that end in the end to the right, so that that end is $E(D)$-even. So, all vertices and all ends are $E(D)$-even, and indeed $E(D)$ is a circuit. In contrast, $E(D')$ is not a circuit and we can easily check that any two arcs contained in $\overline{E(D')}$ terminating in the same end share an edge. Consequently, the ends are $E(D')$-odd and therefore a certificate for $E(D') \notin \mathcal{C}$.

Let us now state the main result of this chapter.

Theorem 3.4. **[11]** *Let G be a locally finite graph, and let $Z \subseteq E(G)$. Then $Z \in \mathcal{C}(G)$ if and only if every vertex and every end of G is Z-even.*

The following lemma will be convenient when we check whether a given end is even or odd. We say that an edge set F *separates* a vertex set S from an end ω if every ray in ω with first vertex in S meets F.

Lemma 3.5. **[11]** *Let ω be an end of a locally finite graph G, let $Z \subseteq E(G)$ and let $S \subseteq V(G)$ be a finite vertex set. Then the maximal number of edge-disjoint S–ω arcs contained in $\overline{Z}$ equals the minimum of $|F \cap Z|$ over all finite cuts F of G that separate S from ω.*

While the proof of the lemma is not overly difficult, it is not very instructive, and similar arguments have been given in [24]. We note that it can also be obtained from a more general result by Thomassen and Vella [81], who prove a Menger-type theorem for graph-like spaces.

With the help of Lemma 3.5 and Theorem 1.9 it becomes easy to prove the forward direction of our main theorem.

Lemma 3.6. [11] *Let G be a locally finite graph, and let $Z \subseteq E(G)$. If $Z \in \mathcal{C}(G)$ then every vertex and every end of G is Z-even.*

Proof. If $Z \in \mathcal{C}(G)$ then, by definition or by Theorem 1.9, every vertex of G is Z-even. To see that any end ω is Z-even too, consider an arbitrary finite vertex set S'. By Lemma 3.5, the maximal number of edge-disjoint S'–ω arcs contained in $\overline{Z}$ equals the minimum $|F \cap Z|$ for all finite cuts F that separate S' from ω. Since Theorem 1.9 implies that for any finite cut F we have that $F \cap Z$ is an even set, we deduce that the number of S'–ω arcs is even, and consequently that ω is Z-even. $\square$

The rest of the chapter will be spent on proving the backward direction. For this, given an edge set Z in a locally finite graph G we assume that every vertex is Z-even but that $Z \notin \mathcal{C}(G)$. Hence, our aim is to find a Z-odd end, which we shall achieve be showing that the conditions of the next lemma are met.

Let $Z \subseteq E(G)$, and let C be some subgraph of G. We write $\partial_Z C$ for the edges in Z with exactly one endvertex in C and exactly one endvertex outside C. We write $\partial_G C$ for $\partial_{E(G)} C$. We call a subgraph R a *region* of G if there is a finite cut F of G so that R is a component of $G - F$. In particular, R is induced and connected. We call a region R *Z-even* if $|\partial_Z R|$ is even, otherwise R is *Z-odd*.

Lemma 3.7. [11] *Let G be a locally finite graph, and let Z be a subset of $E(G)$. Assume there exists a sequence $C_1, C_2, \ldots$ of regions of G with the following properties:*

(i) $|\partial_Z C_n|$ *is odd for all n;*

(ii) $\partial_G C_n \cup E(C_n) \subseteq E(C_{n-1})$ *for all n; and*

(iii) *for every region R of G with $C_k \supseteq R \supseteq C_\ell$ for some $k \leq \ell$ it holds that $|\partial_Z R| \geq |\partial_Z C_k|$.*

Then, G contains a Z-odd end.

Proof. Denote by ω the end of a ray that meets every C_n (that such a ray exists can be seen, for instance, by Lemma 8.2.2 in [31]), and let us show that ω is Z-odd. First, observe that, by (ii), any arc that meets every C_n for large n contains a subarc with ω as endpoint. Now, let a finite vertex set S be given. By (ii), we may pick an N so that S is disjoint from C_N. Put $S' := S \cup N(G - C_N)$ and note that (iii) together with Lemma 3.5 imply that the maximal number of edge-disjoint S'–ω arcs contained in $\overline{Z}$ equals $|\partial_Z C_N|$, which is odd by (i). Thus, ω is Z-odd. $\square$

Let us give a rough outline of the proof of Theorem 3.4. Lemma 3.7 provides us with a recipe for proving the existence of a Z-odd end. But how do we find Z-odd regions C_n as in the lemma? We first note that there is a natural candidate for C_1. By Theorem 1.9, there exists a finite cut of G that meets Z in an odd number of edges. Now, among all finite cuts F so that $|F \cap Z|$ is odd we choose one where $|Z \cap F|$ is minimal. Then, we take C_1 to be a component of $G - F$. This already ensures that for any Z-odd region $R \subseteq C_1$ it holds that $|\partial_Z R| \geq |\partial_Z C_1|$. Furthermore, since every vertex is Z-even the Z-odd cut $F = \partial_G C_1$ propagates into C_1, in the sense that C_1 properly contains Z-odd regions. Finding a suitable region C_2 is more difficult. In a similar way as for C_1, it appears enticing to simply pick among all Z-odd regions C with $\partial_G C \cup E(C) \subseteq E(C_1)$ one so that $|\partial_Z C|$ is minimal. However, since we chose $|\partial_Z C_1|$ to be minimal among all Z-*odd* regions there still could exists a Z-even region R sandwiched between C_1 and C_2 with smaller cut-size in Z than C_1, i.e. with $|\partial_Z C_1| > |\partial_Z R|$.

In order to overcome this problem, we will eliminate all small Z-even cuts before choosing C_2. This will be achieve by contracting certain Z-even regions and obtaining a minor all of whose infinite Z-even regions have large cut-size. In that minor we then choose a region C_2^* so that $\partial_Z C_2^*$ has minimal odd size. By uncontracting we obtain the region C_2 in the original graph. We repeat this procedure. Once again we eliminate all small Z-even cuts, choose C_3^* in the resulting minor and so on. This way we can obtain regions $C_1^*, C_2^*, \ldots$ of different minors of G. We will gain the regions C_n of G by uncontracting the regions C_n^*.

The next section will hand us a tool to eliminate infinite Z-even regions of small cut-size. The main work of constructing a sequence of regions $C_1 \supseteq C_2 \supseteq \ldots$ will be achieved in Section 3.4.

3.3 Elimination of regions with small cutsize

Before we can prove Lemma 3.10, the main tool to eliminate small Z-even cuts, we state two lemmas. The first of which is a standard lemma, that asserts that the function measuring the number of edges leaving a vertex set is submodular.

Lemma 3.8. *Let G be a graph, and let $X, Y \subseteq V(G)$. Then*

$$|\partial X| + |\partial Y| \geq |\partial(X \cap Y)| + |\partial(X \cup Y)|$$

and

$$|\partial X| + |\partial Y| \geq |\partial(X - Y)| + |\partial(Y - X)|.$$

The next lemma is used in the inductive proof of Lemma 3.10. For an edge set Z, we say that D is a *(m, Z)-region* if D is a region with $|\partial_Z D| = m$.

Lemma 3.9. **[11]** *Let G be a locally finite graph, and let $Z \subseteq E(G)$. Let C be a region of G, and denote by m the minimal integer k for which there is an infinite (k, Z)-region R in G with $R \subseteq C$. Assume m to be even. Let $R, S_1, \ldots, S_\ell$ be (m, Z)-regions, where $S_1, \ldots, S_\ell$ are pairwise disjoint and $|R - \bigcup_{i=1}^{\ell} S_i| = \infty$ Then there exist a subgraph K and an (m, Z)-region S satisfying*

(i) *the subgraph K is the union of components of $R - \bigcup_{i=1}^{\ell} S_i$;*

(ii) *the region S is spanned by the union of K with some (possibly none) of $S_1, \ldots, S_\ell$;*

(iii) *$S - S_i$ is connected for every $i = 1, \ldots, \ell$;*

(iv) *K is an infinite subgraph and $|\partial_Z K|$ is even; and*

(v) *if $m = 0$ then each K is connected; and if $m > 0$ then each component of K is incident with an edge in Z.*

Proof. Define $\mathcal{I}$ to be the set of those S_i among $S_1, \ldots, S_\ell$ for which $S_i - R$ is infinite; denote by $\mathcal{J}$ the other ones. Consider an $S_i \in \mathcal{I}$. Observe that by definition of m and since each of $R - S_i$ and $S_i - R$ is an infinite subgraph, Lemma 3.8 implies that $|\partial_Z(R - S_i)| = |\partial_Z(S_i - R)| = m$. Hence, $R - S_i$ contains an infinite (m, Z)-region. In a similar way, we see that for any $S_j \in \mathcal{J}$, the induced subgraph on $R \cup S_j$ contains an infinite (m, Z)-region. As the $S_1, \ldots, S_\ell$ are pairwise disjoint it follows therefore that each infinite component of $G[(R - \bigcup \mathcal{I}) \cup \bigcup \mathcal{J}]$ (and there is at least one) is is an infinite (m, Z)-region. For one of these components, R' say, the subgraph $K' := R' - \bigcup_{i=1}^{\ell} S_i$ will be infinite, so that K' satisfies (i), and (ii) holds for R'. Among all infinite subgraphs K of G and (m, Z)-region S satisfying (i) and (ii), we choose S to be $\subseteq$-minimal.

Let us now show that K satisfies (iv). Indeed, let $\mathcal{T} \subseteq \{S_1, \ldots, S_\ell\}$ so that $S = G[K \cup \bigcup \mathcal{T}]$. Since the S_i are pairwise disjoint it follows that

$$\partial_Z K = \partial_Z(S - \bigcup \mathcal{T}) = \partial_Z S + \sum_{T \in \mathcal{T}} \partial_Z T.$$

(Recall that we consider the sum of edge sets to be their symmetric difference.) Since S as well as all the $T \in \mathcal{T}$ are (m, Z)-regions, we deduce that $|\partial_Z K|$ is even.

Next, assume that $m > 0$ and suppose that K has a component L that is not incident with any edge in Z. Since $m > 0$, L cannot be infinite, which implies that $K - L$ is still infinite. Moreover, as $|\partial_Z(S - L)| = m$, one of the components of $S - L$ is an infinite (m, Z)-region S', which then together with $K' := S' - \bigcup_{i=1}^{\ell} S_i$ constitutes a contradiction to the minimal choice of S. If, on the other hand, $m = 0$ then K cannot be disconnected as each infinite component K' of K would with $S' = K'$ contradict the choice of K and S. Therefore, (v) is proved.

Finally, in order to prove (iii), suppose that there exists a k so that $S - S_k$ is not connected. Since $K \subseteq S - \bigcup_{i=1}^{\ell} S_i$ is infinite one of the components of $S - S_k$, X say, has therefore the property that $K' := X - \bigcup_{i=1}^{\ell} S_i$ is infinite as well. Observe that it follows from (ii) that $S_k \subseteq S$. Setting $Y := S - (S_k \cup X)$, we see that

$$2m = |\partial_Z S| + |\partial_Z S_k| = |\partial_Z(S_k \cup X)| + |\partial_Z(S_k \cup Y)|.$$

As both $S_k \cup X$ and $S_k \cup Y$ are infinite, $S' := G[S_k \cup X] \subseteq S$ is an infinite (m, Z)-region. Again we have, with S' and K', obtained a contradiction to the minimal choice of S. □

In the rough sketch of the proof of Theorem 3.4, we claimed we would construct minors in order to eliminate infinite Z-even regions of small cutsize. This is slightly incorrect. Unfortunately, and this will lead to some technical complications, we are not able to force the contracted branch sets to be connected. Rather, it will sometimes be necessary to contract a disconnected set to a single vertex. Thus, we will not be working with minors but with what we call pseudo-minors.

Let $\mathcal{V}$ be a partition of the vertex set of a graph G. We define a graph H with vertex set $\mathcal{V}$ and edge set $E(H) \subseteq E(G)$, so that e is an edge of H between two distinct vertices U and U' of H if and only if e is an edge of G with one endvertex in U and the other in U'. In particular, we allow H to have parallel edges but no loops. We call such a graph H a *pseudo-minor of* G, denoted by $H \preccurlyeq G$, and define $\mathcal{K}(H, G)$ to be the set of non-singletons in $\mathcal{V}$.

Let D, K be subgraphs of G. We say that D *splits* K if neither $V(K) \subseteq V(D)$ nor $V(D) \cap V(K) = \emptyset$.

Lemma 3.10. **[11]** *Let G be a locally finite graph, and let $Z \subseteq E(G)$. Let C be a region of G, and denote by m the minimal k for which there is an infinite (k, Z)-region R in G with $R \subseteq C$. Assume m to be even. Then there exists a locally finite pseudo-minor G' of G and a set $\mathcal{S}$ of (m, Z)-regions of G so that the following holds:*

(i) *K is infinite for each $K \in \mathcal{K}(G', G)$ and $|\partial_Z K|$ is even;*

(ii) *every region D of G splits at most finitely many $K \in \mathcal{K}(G', G)$;*

(iii) *if D is an infinite (k, Z)-region of G' with $E(D) \subseteq E(C)$ then it follows that $k > m$;*

(iv) *for every $K \in \mathcal{K}(G', G)$ there is an $S \in \mathcal{S}$ with $K \subseteq S \subseteq C$;*

(v) *for every $S \in \mathcal{S}$ there is an $\mathcal{L} \subseteq \mathcal{K}(G', G)$ with $S = G[\bigcup_{L \in \mathcal{L}} L]$; and*

(vi) *if $m = 0$ then each $K \in \mathcal{K}(G', G)$ is connected; and if $m > 0$ then each component of K, $K \in \mathcal{K}(G', G)$, is incident with an edge in Z.*

If the assertions of the lemma are satisfied for a graph G with pseudo-minor G', a region C and a set of regions $\mathcal{S}$ then the tuple $(G', G, C, \mathcal{S}, m)$ will be called a *legal contraction system.*

Proof. We may restrict ourselves to the component of G that contains the region C, and therefore assume that G itself is connected. Let $R_1, R_2, \ldots$ be an enumeration of all infinite (m, Z)-regions of G. (Since G is connected, these are only countably many.)

We shall define inductively subgraphs $K_1, K_2, K_3, \ldots$ and (m, Z)-regions $S_1, S_2, S_3, \ldots \subseteq C$ satisfying

- K_i is infinite for each $i = 1, 2, 3, \ldots$ and $|\partial_Z K_i|$ is even;
- For every $i = 1, 2, 3, \ldots$, the region S_i is spanned by the union of K_i with some (possibly none) of $S_1, \ldots, S_{i-1}$;
- The $K_1, K_2, K_3, \ldots$ are pairwise disjoint;
- if $m = 0$ then each K_i is connected; and if $m > 0$ then each component of each K_i is incident with an edge in Z.

We note that the second and third property imply that

$$\textit{for each } j < i \textit{ either } S_j \subseteq S_i \textit{ or } S_j \cap S_i = \emptyset. \tag{3.1}$$

Taking $\mathcal{S} = \{S_1, S_2, S_3, \ldots\}$ and obtaining G' from G by contracting the K_i, we clearly have (i), (iv)–(vi). A further analysis of the process will yield (ii) and (iii).

We start by setting $K_1 = S_1 = R_1$. Now assume that $K_1, \ldots K_\ell$ and $S_1, \ldots, S_\ell$ are constructed. We denote by n_ℓ the minimal n satisfying $|R_n - \bigcup_{i=1}^{\ell} S_i| = \infty$. (If no such R_n exists, the process terminates.) We then apply

Lemma 3.9 to R_{n_ℓ} and the $\subseteq$-maximal regions among $S_1, \dots, S_\ell$, which are, by (3.1), pairwise disjoint. The resulting K and S found by the lemma will be chosen as $K_{\ell+1}$ and $S_{\ell+1}$ respectively.

In order to see that (ii) is satisfied, let R be a region of G and let $X = (\partial_G R) \cap \bigcup_{i=1}^{\infty} E(S_i)$. Let $\{S_{i_1}, S_{i_2}, \dots, S_{i_k}\}$ be a set of minimal size with $X \subseteq \bigcup_{j=1}^{k} E(S_{i_j})$. We claim that R does not split any K_i with $i > \max(i_1, \dots, i_k)$. To reach a contradiction, suppose that R splits K_n for some $n > \max(i_1, \dots, i_k)$. Then, as K_n is disjoint from all the S_{i_j}, K_n must be disconnected, and have at least one component inside R and at least one component outside R. Since $K_n \subseteq S_n$ and since S_n is connected, S_n must contain some edge from X. Hence it meets some S_{i_j}, S_{i_1} say. As $n > i_1$, we have, by construction, that $S_{i_1} \subseteq S_n$.

Let S_p be the $\subseteq$-maximal region among $S_1, \dots S_{n-1}$ containing S_{i_1}. Recall that we chose S_n and K_n using Lemma 3.9, which states that $S_n - S_p$ is connected. As, furthermore, K_n is disjoint from S_p because it is chosen later, it follows that $S_n - S_p$ contains an edge of X, and thus meets and then contains one of $S_{i_2}, \dots, S_{i_k}$, say S_{i_2}. We thus have $S_{i_2} \subseteq S_n - S_p \subseteq S_n - S_{i_1}$. This, however, leads to $X \subseteq E(S_n) \cup \bigcup_{j=3}^{k} E(S_{i_j})$, which contradicts the minimality of k. This completes the proof of (ii).

Let us finally prove (iii). Note that if $n_{\ell+1} = n_\ell$ then, by Lemma 3.9 (i), $R_{n_\ell} - \bigcup_{i=1}^{\ell+1} S_i$ will have less components than $R_{n_\ell} - \bigcup_{i=1}^{\ell} S_i$. This implies $\lim_{\ell \to \infty} n_\ell = \infty$. Therefore, for every (m, Z)-region R_n it holds that $|R_n - \bigcup_{i=1}^{\ell} S_i| < \infty$ for some ℓ. Now, if D is an infinite (k, Z)-region of G' with $E(D) \subseteq E(C)$ then by uncontracting and (ii) we find a (k', Z)-region R of G with $R \subseteq C$ and $k' \leq k$ so that $E(R) \cap E(G')$ is infinite. By assumption, we have that $k' \geq m$. If $k' = m$ then $R = R_n$ for some n, and consequently $|R_n - \bigcup_{i=1}^{\ell} S_i| < \infty$ for some ℓ, contradicting that $E(R) \cap E(G')$ is infinite. This proves (iii). $\square$

3.4 Proof of main result

We restate and then prove the main result of this chapter.

Theorem 3.4. **[11]** *Let G be a locally finite graph, and let $Z \subseteq E(G)$. Then $Z \in \mathcal{C}(G)$ if and only if every vertex and every end of G is Z-even.*

Proof. In light of Lemma 3.6 we only need to prove the backward direction. In order to do so, assume that every vertex of G is Z-even but that $Z \notin \mathcal{C}(G)$. Our task is to find a Z-odd end in G.

Since $Z \notin \mathcal{C}(G)$ there exists a topologically connected component of $\overline{Z}$ whose edge set is not an element of the cycle space (recall that $\overline{Z}$ is the closure of Z as a subspace of $|G|$). An end of G that is odd with respect to some connected component of $\overline{Z}$ is Z-odd, as each end of G lies in at most one connected component. Thus, we may, by deleting all other edges from Z, assume that $\overline{Z}$ is topologically connected. In particular, this means that

there exists no finite cut of G that avoids Z but separates two edges in Z. (3.2)

Recall that in order to apply Lemma 3.7, we need to find a sequence of nested Z-odd regions $C_1 \supsetneq C_2 \supsetneq \ldots$ of G so that for any region R of G with $C_k \supseteq R \supseteq C_\ell$ for some $k \leq \ell$ it holds that $|\partial_Z C_k| \leq |\partial_Z R|$. We shall do this by first defining a sequence of regions $C_1^*, C_2^*, \ldots$ in certain pseudo-minors of G. We will obtain C_n from the C_n^* by uncontracting.

Let us be more precise. Inductively, we will find Z-odd regions $C_1^*, C_2^* \ldots$ in certain pseudo-minors $G^0 \succcurlyeq G^2 \succcurlyeq \ldots$ of G. We set $m_n := |\partial_Z C_n^*| - 1$ for all $n \geq 1$, and for convenience we put $m_0 = -2$. We start the construction with $G^{-2} = G$. Now, in each step, i.e. for each $n \in \mathbb{N}$, we will first find a region C_n^* of $G^{m_{n-1}}$ and then construct pseudo-minors $G^{m_{n-1}+2} \succcurlyeq \ldots \succcurlyeq G^{m_n}$ of $G^{m_{n-1}}$. We require that for all $n \geq 1$ it holds that

(i) $|\partial_Z C_n^*|$ is odd; and

(ii) $\partial_{G^{m_{n-1}}} C_n^* \cup E(C_n^*) \subseteq E(C_{n-1}^*)$ if $n \geq 2$.

Let us pause for a while before we give a third requirement. Recall that C_n^* is a region in the pseudo-minor $G^{m_{n-1}}$. It should be noted that there might be values of m which are not equal to any m_n, but we will still need to refer to a region of G^m which is naturally obtained from the sequence $C_1^*, C_2^*, C_3^*, \ldots$. For this end, we introduce the following notation:

Let $\ell > m \geq -2$ be two even numbers and let D be an induced subgraph of G^m. Then $\pi_{m,\ell}(D)$ denotes the induced subgraph of G^ℓ on the vertex set $\{X_v : v \in V(D) \text{ and } v \in X_v \in V(G^\ell)\}$ (recall that, as G^ℓ is a pseudo-minor of G^m, its vertex set is a partition of $V(G^m)$). We will also consider the inverse of $\pi_{m,\ell}$, which we denote by $\pi_{\ell,m}$. Given an induced subgraph D' of G^ℓ define $\pi_{\ell,m}(D')$ to be the induced subgraph of G^m on the vertex set $\bigcup_{X \in V(D')} X$. Also for every induced subgraph D of G^m write $\pi_{m,m}(D) = D$. Thus $\pi_{m,\ell}(D)$ is defined for every $m, \ell \in \{-2, 0, 2, 4, \ldots\}$ regardless of the order between them, and for every induced subgraph D of G^m it assigns an induced subgraph of G^ℓ.

With this notation, put $C_n^m := \pi_{m_{n-1},m}(C_n^*)$ for every $m = -2, 0, 2, 4, \ldots$ and $n = 1, 2, 3, \ldots$. In particular, this means that $C_n^* = C_n^{m_{n-1}}$.

We are now ready to state the third requirement. Alongside with the construction of the pseudo-minors G^m and the regions C_n^* we will construct, for every m, sets $\mathcal{S}_m$ of (m, Z)-regions of G^{m-2} so that

(iii) $(G^m, G^{m-2}, C_n^{m-2}, \mathcal{S}_m, m)$ is a legal contraction system, where n is the number satisfying $m_{n-1} < m \leq m_n$.

Thus, by Lemma 3.10 (i) each vertex in G^m remains Z-even. Since it is central to the main idea of the proof we restate another implications of (iii):

$$\text{if } D \subseteq C_n^m \text{ is an infinite } (k, Z)\text{-region of } G^m \text{ for } 0 \leq m \leq m_n \text{ then } k > m. \tag{3.3}$$

The statement follows from Lemma 3.10 (iii) either directly or by induction, depending on whether $m_{n-1} < m \leq m_n$ or not.

Furthermore, we note the following consequence of (iii) or, more specifically, of Lemma 3.10 (iv).

$$C_n^* \text{ does not split any } K \in \mathcal{K}(G^m, G^{m_{n-1}}) \text{ for any even } m \geq m_{n-1}. \tag{3.4}$$

In particular, it follows that the cut $\partial_{G^{m_{n-1}}} C_n^*$ lies in G^m for each even $m \geq m_{n-1}$ and thus is still a cut there.

Let us make one last observation before we finally start with the construction. We claim that

$$\text{if } D' \text{ is an infinite region of } G^m \text{ for some } m \text{ then there exists for each even } \ell \leq m \text{ an infinite region } D \text{ of } G^\ell \text{ with } \partial_{G^\ell} D \subseteq \partial_{G^m} D' \text{ and } \pi_{\ell,m}(D) \subseteq D'. \tag{3.5}$$

Indeed, $\tilde{D} := \pi_{m,\ell}(D')$ is an induced subgraph of G^ℓ with $\partial_{G^\ell} \tilde{D} = \partial_{G^m} D'$. From (iii), resp. Lemma 3.10 (ii), it follows that $\tilde{D}$ has only finitely many components. Hence, one of them is infinite, and this will be the desired region D.

We are now ready to start the construction. We begin with $G^{-2} = G$ and let $C_1 = C_1^*$ be a Z-odd region with the minimal possible value of $|\partial_Z C_1|$. We know that Z-odd regions exist by Theorem 1.9. Recall that we write $m_1 = |\partial_Z C_1| - 1$ We then construct $G^0, \ldots, G^{m_1}$ and $\mathcal{S}_0, \ldots, \mathcal{S}_{m_1}$ using Lemma 3.10 in a way that will be described in more details later on.

For $n > 1$, assume $C_1^*, \ldots, C_{n-1}^*$ and $G^{-2}, G^0, \ldots, G^{m_{n-1}}$ and the corresponding $\mathcal{S}_m$ to be constructed. In order to find a suitable region C_n^*, we first claim that there exist Z-odd regions C satisfying (ii).

For $n > 1$, denote by F the edges in $G^{m_{n-1}}$ incident with the vertices in $N(G^{m_{n-1}} - C_{n-1}^{m_{n-1}})$. Since C_{n-1}^* is a region, F is a finite set and thus $|F \cap Z|$ is even. Furthermore, $\partial_{G^{m_{n-1}}} C_{n-1}^* \subseteq F$ implies that the cut $F \setminus \partial_{G^{m_{n-1}}} C_{n-1}^*$ meets Z in an odd number of edges. One of the components of $G^{m_{n-1}} - (F \setminus \partial_{G^{m_{n-1}}} C_{n-1}^*)$ contained in $C_{n-1}^{m_{n-1}}$ will thus be Z-odd and hence as desired. We now pick C_n^* among all Z-odd regions C of $G^{m_{n-1}}$ satisfying (ii) so that $|\partial_Z C_n^*|$ is minimal.

A consequence of the choice of C_n^* is that for all $n \in \mathbb{N}$:

$$\textit{if } D \textit{ is a } (k,Z)\textit{-region of } G^{m_{n-1}} \textit{ so that } k \textit{ is odd and } D \subseteq C_n^* \textit{ then } k \geq m_n + 1 = |\partial_Z C_n^*|. \tag{3.6}$$

In order to define G^m, for even m with $m_{n-1} < m \leq m_n$, assume G^ℓ for $\ell = -2, 0, 2, 4, \ldots, m-2$ to be already constructed. By (3.3), it holds that the smallest k^* for which there is an infinite (k^*, Z)-region R of G^{m-2} with $R \subseteq C_n^{m-2}$ is at least $m-1$. Now, (3.6) in conjunction with (3.5) shows that $k^* \geq m$. Hence, we may apply Lemma 3.10 to G^{m-2}, Z, C_n^{m-2}, m (in the roles of G, Z, C, m respectively). With the resulting pseudo-minor G^m and set of (m, Z)-regions $\mathcal{S}_m$, the tuple $(G^m, G^{m-2}, C_n^{m-2}, \mathcal{S}_m, m)$ is a legal contraction system, i.e. (iii) is satisfied.

Assume the construction achieved for all n and corresponding m. Set $C_n := \pi_{m_{n-1},-2}(C_n^*) = C_n^{-2}$, i.e. C_n is the induced subgraph of G obtained from C_n^* by uncontracting $\mathcal{K}(G^{m_{n-1}}, G)$. Observe that

$$|\partial_Z C_n| \textit{ is odd and } \partial_G(C_n) \cup E(C_n) \subseteq E(C_{n-1}) \textit{ for all } n \in \mathbb{N}. \tag{3.7}$$

Recall that our aim is to find a sequence of regions C_n of G satisfying the requirements of Lemma 3.7. As (3.7) means that already two of the conditions hold we need only make sure that each C_n is indeed a region, i.e. a connected subgraph, and that for all regions R with $C_k \supseteq R \supseteq C_\ell$ for some $k \leq \ell$ it holds that $|\partial_Z R| \geq |\partial_Z C_k|$. We shall deal with the latter condition first.

Observe that (3.4) implies that

$$C_n \textit{ does not split any } K \in \mathcal{K}(G^{m_n}, G). \tag{3.8}$$

Next, we prove that for every n and $0 \leq m \leq m_n$ it holds that

$$\textit{for every } (k,Z)\textit{-region } R \textit{ of } G \textit{ with } R \subseteq C_n \textit{ and } k \leq m \textit{ it follows that } \pi_{-2,m}(R) \textit{ is a finite subgraph of } G^m. \tag{3.9}$$

Assume the statement to be false for G^m, for some m. For even ℓ, $-2 \leq \ell \leq m$, denote by $\mathcal{D}^\ell$ the set of (k, Z)-regions D of G^ℓ with $k \leq m$, $D \subseteq C_n^\ell$ and so that $\pi_{\ell,m}(D)$ is infinite. Clearly, by assumption we have that $\mathcal{D}^{-2} \neq \emptyset$. On the other hand, it holds that $\mathcal{D}^m = \emptyset$. Indeed, any element in $\mathcal{D}^m$ would contradict (3.3).

Now, choose $\ell \leq m$ to be the maximal even integer so that $\mathcal{D}^{\ell-2} \neq \emptyset$, and among the $D \in \mathcal{D}^{\ell-2}$ pick one, $\tilde{D}$ say, so that $\tilde{D}$ splits a minimum number of elements in $\mathcal{K}(G^\ell, G^{\ell-2})$. (Note that, by Lemma 3.10 (ii), every region splits only finitely many sets in $\mathcal{K}(G^\ell, G^{\ell-2})$.)

Now, since $\mathcal{D}^\ell = \emptyset$, the region $\tilde{D}$ of $G^{\ell-2}$ must split some element K of $\mathcal{K}(G^\ell, G^{\ell-2})$. Let $S \in \mathcal{S}_\ell$ be an (ℓ, Z)-region with $K \subseteq S$ of $G^{\ell-2}$ (see Lemma 3.10 (iv)).

We distinguish two cases. Assume that $S - \tilde{D}$ is infinite. If $|\partial_Z(\tilde{D} - S)| > |\partial_Z \tilde{D}|$ then, by Lemma 3.8, $S - \tilde{D}$ contains an infinite (ℓ', Z)-region of $G^{\ell-2}$ with $\ell' < \ell$, in contradiction to either (3.3) or (3.6) (together with (3.5)). Thus, $\tilde{D} - S$ is a (k', Z)-region with $k' \leq |\partial_Z \tilde{D}| \leq m$. Observe that because $\pi_{\ell-2,\ell}(S)$ is finite by Lemma 3.10 (iii), the subgraph $\pi_{\ell-2,m}(\tilde{D} - S)$ of G^m is still infinite. As, moreover, by Lemma 3.10 (v), $\tilde{D} - S$ splits fewer elements in $\mathcal{K}(G^\ell, G^{\ell-2})$, we obtain a contradiction to the choice of $\tilde{D}$.

So, let $S - \tilde{D}$ be finite. Suppose that $|\partial_Z(S \cup \tilde{D})| > |\partial_Z \tilde{D}|$. Then, by Lemma 3.8, $S \cap \tilde{D}$ contains an infinite (ℓ', Z)-region of $G^{\ell-2}$ with $\ell' < \ell$, in contradiction to either (3.3) or (3.6). Thus, $G[S \cup \tilde{D}]$ is an infinite (k', Z)-region with $k' \leq |\partial_Z \tilde{D}| \leq m$. Since $G[S \cup \tilde{D}]$ splits fewer $K \in \mathcal{K}(G^\ell, G^{\ell-2})$ than $\tilde{D}$, we obtain again a contradiction to the choice of $\tilde{D}$—provided we can show that $S \cup \tilde{D} \subseteq C_n^{\ell-2}$. To do this, observe that, by Lemma 3.10 (v), there is a set $\mathcal{L} \subseteq \mathcal{K}(G^\ell, G^{\ell-2})$ with $S = G[\bigcup_{L \in \mathcal{L}} L]$. Since $S - \tilde{D}$ is finite but all the $L \in \mathcal{L}$ are infinite by Lemma 3.10 (i), it follows that $\tilde{D}$ meets every $L \in \mathcal{L}$. Together with the fact that C_n does not split any elements in $\mathcal{K}(G^\ell, G^{\ell-2})$, by (3.8), and $\tilde{D} \subseteq C_n^{\ell-2}$ it follows that $S \subseteq C_n^{\ell-2}$, and hence, $S \cup \tilde{D} \subseteq C_n^{\ell-2}$. This finishes the proof of (3.9).

In order to prove that the subgraphs $C_1, C_2, \ldots$ of G satisfy Condition (iii) of Lemma 3.7 consider a region R with $C_k \supseteq R \supseteq C_\ell$ for some $\ell \geq k$. Observe that $C_\ell^{m_k}$ is still an infinite subgraph of G^{m_k} since $\partial_Z C_\ell^{m_k}$ is odd but every vertex in G^{m_k} is Z-even. Thus, $\pi_{-2,m_k}(R) \supseteq C_\ell^{m_k}$ is infinite, which with (3.9) implies that $|\partial_Z R| \geq m_k + 1 = |\partial_Z C_k|$, as desired.

For Lemma 3.7 to apply, it remains to show that:

$$C_n \textit{ is a region of } G \textit{ for all } n \in \mathbb{N}. \tag{3.10}$$

For this, it suffices to prove that C_n is connected. If $m_n = 0$ then, by Lemma 3.10 (vi), G^0 is a minor (rather than only a pseudo-minor) of G. As C_n^* is a region in $G^{m_{n-1}}$, which is either G^0 or $G^{-2} = G$, we immediately see that C_n is connected as well.

So, let $m_n > 0$. Since C_n does not split any $K \in \mathcal{K}(G^{m_n}, G)$ and since $\partial_G C_n$ is odd, $C_n^{m_n}$ is infinite. Let C be a component of C_n so that $\pi_{-2,m_n}(C)$ is infinite, and suppose that $R := C_n - C$ is non-empty. If $\partial_Z R \neq \emptyset$ then $|\partial_Z C| \leq m_n$ in contradiction to (3.9). If, on the other hand, $\partial_Z R = \emptyset$ then $\partial_G R$ is a finite cut of G separating two edges in Z, which constitutes a contradiction to that $\overline{Z}$ is topologically connected. Indeed, C contains an edge of Z since $\partial_Z C = \partial_Z C_n$ is an odd set but every vertex is Z-even. To see that $E(R)$ meets Z, recall that C_n^* is a region of $G^{m_{n-1}}$. Thus, there exists an $\ell \leq m_{n-1}$, so that $\pi_{-2,\ell-2}(C)$ splits a $K \in \mathcal{K}(G^\ell, G^{\ell-2})$. By (3.8), the subgraph K of $G^{\ell-2}$ is contained in $C_n^{\ell-2}$ and then K has one component in $\pi_{-2,\ell-2}(C)$ and one in $\pi_{-2,\ell-2}(R)$. Lemma 3.10 (vi) implies that both these components are incident with an edge in Z. As $\partial_Z R = \emptyset$, we obtain $E(R) \cap Z \neq \emptyset$.

In conclusion, the regions C_n satisfy all conditions required in Lemma 3.7, which therefore yields the desired Z-odd end in G. $\square$

Chapter 4

Bicycles and left-right tours

4.1 Bicycles in finite graphs

The graphs in this chapter are assumed to be simple unless otherwise noted. The set of edge sets in a graph G together with symmetric difference as addition forms a $\mathbb{Z}_2$-vector space, the *edge space* $\mathcal{E}(G)$. Two important subspaces of $\mathcal{E}(G)$ are the topological cycle space $\mathcal{C}(G)$ and the *cut space* $\mathcal{C}^*(G)$, the set of all cuts. Although cuts and cycle space elements are, in a sense, orthogonal to each other, it is possible for an edge set to be an element of both, $\mathcal{C}(G)$ and $\mathcal{C}^*(G)$. Such an edge set is called a *bicycle* and the space $\mathcal{B}(G) := \mathcal{C}(G) \cap \mathcal{C}^*(G)$ is the *bicycle space*. See Figure 4.1 for an example.

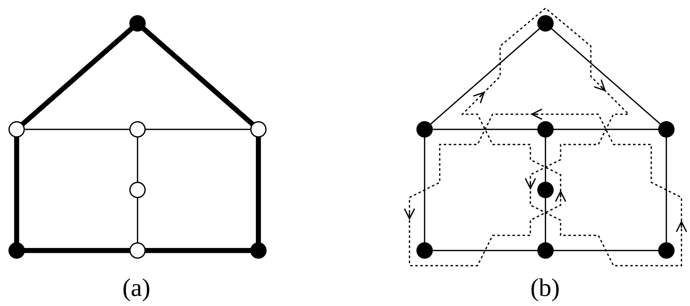

Figure 4.1: (a) a bicycle; (b) a left-right tour

In finite graphs, bicycles have been widely studied and a number of fundamental results involving bicycles are known. The aim of this chapter, which is based on [21], is to extend three of these to locally finite graphs. The first theorem we will extend is Read and Rosenstiehl's tripartition theorem:

Theorem 4.1 (Read and Rosenstiehl [74])**.** *Let e be an edge in a finite graph G. Then exactly one of the following holds:*

(i) *there exists a* $B \in \mathcal{B}(G)$ *with* $e \in B$*; or*

(ii) *there exists a* $Y \in \mathcal{C}(G)$ *with* $e \in Y$ *and* $Y + e \in \mathcal{C}^*(G)$*; or*

(iii) *there exists a* $Z \in \mathcal{C}(G)$ *with* $e \notin Z$ *and* $Z + e \in \mathcal{C}^*(G)$*.*

With the naive definition of $\mathcal{C}(G)$, in which every element of the cycle space is necessarily finite, Theorem 4.1 cannot be expected to carry over to locally finite graphs. The double ladder, depicted in Figure 4.2, constitutes an obvious counterexample: no finite bicycle contains the edge e, yet there is neither a finite Y nor a finite Z as in (ii) or (iii) of the theorem.

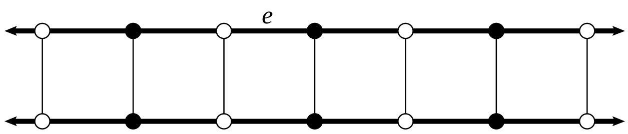

Figure 4.2: There is no finite B, Y or Z as in Theorem 4.1 for e

Again, the almost self-evident solution is that infinite graphs demand infinite circuits. We see that the set of bold edges in the double ladder form an infinite circuit. (The two double rays together with the end to the left and the one to the right are homeomorphic to the unit circle.) Since this edge set is also a cut, we have found an infinite bicycle containing e, and thus the counterexample ceases to be one. More generally, we will prove in Sections 4.3 and 4.4 that the tripartition theorem becomes true for locally finite graphs once infinite circuits are admitted.

In Sections 4.5 and 4.6 we will be concerned with plane graphs. In plane graphs, there is an easy way to find bicycles. Starting with any edge uv, we traverse uv from u to v, and then choose the leftmost edge at v, follow it along, then turn right, again turn left at the next vertex, and we continue alternating between left and right turns until we reach uv again. There we stop, provided we are about to traverse uv again from u to v and provided our turn at v would, again, be a left turn. The closed walk produced in this way is called a *left-right tour*. Its *residue*, the set of edges traversed exactly once, forms a bicycle; see Figure 4.1.

Shank [75] observed that left-right tours not only yield bicycles but that they, moreover, determine already all bicycles in the graph:

Theorem 4.2 (Shank [75]). *In a finite plane graph the residues of the left-right tours generate the bicycle space.*

This is the second of the theorems we shall extend to locally finite graphs. See also Richter and Shank [71] and Lins, Richter and Shank [58].

The third and final result we shall treat in this chapter is a planarity criterion that involves left-right tours and bicycles in a sophisticated way (Section 4.7). For finite graphs this is due to Archdeacon, Bonnington and Little [6].

4.2 Definitions and preliminaries

In the *cut space* $\mathcal{C}^*(G)$, the set of all cuts, a result that is analogous to Theorem 1.9 holds; see the next lemma. A proof of this easy result can, for instance, be found in [16].

Lemma 4.3. *Let F be a set of edges in a graph G. Then F is a cut if and only if it meets every finite circuit in an even number of edges.*

We call the space $\mathcal{B}(G) := \mathcal{C}(G) \cap \mathcal{C}^*(G)$ the *bicycle space* of G; an element of $\mathcal{B}(G)$ is a *bicycle*.[1]

In Sections 4.5 and 4.7 we will be concerned with infinite plane graphs. The usual drawings seem rather insufficient for infinite graphs. Indeed, several of the expected properties may fail. For instance, in a 2-connected graph the face boundaries do not need to be cycles. Moreover, they might even contain only half an edge (for instance, in the drawing there might be vertices converging against an interior point of an edge) or no edges at all. All these problems are overcome when, instead of G, the space $|G|$ is embedded in the sphere. Fortunately, this is not a restriction at all:

Theorem 4.4 (Richter and Thomassen [72])**.** *Let G be a locally finite* 2*-connected planar graph. Then $|G|$ embeds in the sphere.*

While the theorem is formulated for 2-connected graphs, it is not hard to extend it to graphs that are merely connected. And indeed, we will make use of the theorem in graphs that are not necessarily 2-connected.

Assuming $|G|$ to be embedded in the sphere S, we call a connected component of $S \setminus |G|$ a *face* and its boundary a *face boundary*. It can be seen that each face boundary consists of a subgraph of G together with a subset of the ends of G.

[1]There is a certain inconsistency here. Following Diestel [31], we use "cycle" to denote a *subgraph* stemming from a homeomorphic image of S^1. In particular, a finite cycle is a connected subgraph. On the other hand, a finite bicycle, which is an *edge set*, does not need to span a connected graph.

4.3 The tripartition theorem

In this section, we extend Read and Rosenstiehl's tripartition theorem to locally finite graphs. Since the proof is short and because it is worthwhile to see where it breaks down for infinite graphs, we will start by repeating the proof for finite graphs.

For this, let us recall two standard notions. There is a scalar product $*$ defined on $\mathcal{E}(G)$ for a multigraph G as follows: for $X, Y \subseteq E(G)$, we let $X * Y = 0$ if $|X \cap Y|$ is even, and we set $X * Y = 1$ otherwise. With this product, for a set of edge sets $\mathcal{X}$, we can define the orthogonal space $\mathcal{X}^\perp := \{Y \subseteq E(G) : Y * X = 0 \text{ for all } X \in \mathcal{X}\}$. Clearly, this is standard linear algebra and all the usual methods apply. We recall the well-known fact that $\mathcal{C}(G)^\perp = \mathcal{C}^*(G)$, for a finite (multi-)graph G.

Proof of Theorem 4.1. Assume that there is no bicycle containing e. Thus

$$\{e\} \in \mathcal{B}(G)^\perp = (\mathcal{C}(G) \cap \mathcal{C}^*(G))^\perp = \mathcal{C}(G)^\perp + \mathcal{C}^*(G)^\perp = \mathcal{C}^*(G) + \mathcal{C}(G).$$

We omit the easy proof that only one of (i)–(iii) can hold, since these arguments will appear later anyway. □

The first problem we encounter when we apply this proof to infinite graphs concerns the definition of the scalar product: What should the value of $X * Y$ be if the edge sets X, Y have infinite intersection? Fortunately, we will be able to circumvent this issue by only using the scalar product for $X, Y \in \mathcal{E}(G)$ with $|X \cap Y| < \infty$. A proper concept for orthogonal spaces appears to be more difficult, as however defined they seem to lose a number of their usual properties. For this reason, we will make do without them in infinite graphs. We remark that, these problems notwithstanding, Casteels and Richter [27] introduce orthogonal spaces in infinite graphs that still retain many of the usual properties.

Before we state the tripartition theorem for locally finite graphs, let us denote by $\mathcal{C}_{\text{fin}}(G)$ (resp. $\mathcal{C}^*_{\text{fin}}(G)$ or $\mathcal{B}_{\text{fin}}(G)$) the set of all finite edge sets in $\mathcal{C}(G)$ (resp. in $\mathcal{C}^*(G)$ or in $\mathcal{B}(G)$).

Theorem 4.5. **[21]** *Let e be an edge of a locally finite graph G. Then either*

(i) *there exists $B \in \mathcal{B}(G)$ with $e \in B$; or*

(ii) $\{e\} \in \mathcal{C}_{\text{fin}}(G) + \mathcal{C}^*_{\text{fin}}(G)$

but not both.

The reader will have noticed that the theorem only divides the edges into two classes rather than three. We will address this at the end of the section. The proof uses Kőnig's Infinity Lemma, a standard tool in infinite graph theory. For a proof we refer the reader to [31].

Lemma 4.6 (Kőnig's Infinity Lemma)**.** *Let $W_1, W_2, \dots$ be an infinite sequence of disjoint non-empty finite sets, and let H be a graph on their union. For every $n \geq 2$ assume that every vertex in W_n has a neighbour in W_{n-1}. Then H contains a ray $v_1 v_2 \dots$ with $v_n \in W_n$ for all n.*

Proof of Theorem 4.5. We may assume G to be connected and therefore countable. For each $n \in \mathbb{N}$ denote by S_n the set of the first $n+1$ vertices in some fixed enumeration of the vertices of G that starts with the endvertices of e. Define G_n to be the graph $G[S_n]$ together with the edges in $E(S_n, V(G) \setminus S_n)$ and their incident vertices. Let $\widetilde{G}_n$ be the minor of G obtained by contracting the components of $G - S_n$ (where we keep parallel edges but delete loops). Note that $e \in E(G_n) = E(\widetilde{G}_n)$. Put $W_n := \{B \in \mathcal{C}^*(G_n) \cap \mathcal{C}(\widetilde{G}_n) : e \in B\}$.

We distinguish two cases. First, assume there exists an N such that $W_N = \emptyset$. As $e \in E(G_N)$ this means that $\{e\} \in (\mathcal{C}^*(G_N) \cap \mathcal{C}(\widetilde{G}_N))^\perp$ (where we take the orthogonal space with respect to $\mathcal{E}(G_N)$, which is a finite vector space). Since $\mathcal{C}(G_N) \subseteq \mathcal{C}_{\text{fin}}(G)$ and $\mathcal{C}^*(\widetilde{G}_N) \subseteq \mathcal{C}^*_{\text{fin}}(G)$ it follows that

$$\begin{aligned} \{e\} \in (\mathcal{C}^*(G_N) \cap \mathcal{C}(\widetilde{G}_N))^\perp &= \mathcal{C}^*(G_N)^\perp + \mathcal{C}(\widetilde{G}_N)^\perp \\ &= \mathcal{C}(G_N) + \mathcal{C}^*(\widetilde{G}_N) \quad \subseteq \quad \mathcal{C}_{\text{fin}}(G) + \mathcal{C}^*_{\text{fin}}(G) \end{aligned}$$

and hence (ii) holds.

Second, assume $W_n \neq \emptyset$ for all n. It is not hard to check that for each $K \in \mathcal{C}^*(G_{n+1})$ it holds that $K \cap E(G_n) \in \mathcal{C}^*(G_n)$, and that for each $Z \in \mathcal{C}(\widetilde{G}_{n+1})$ the restriction $Z \cap E(\widetilde{G}_n)$ lies in $\mathcal{C}(\widetilde{G}_n)$. It follows that $B \in W_{n+1}$ implies $B \cap E(G_n) \in W_n$. We define a graph on $\bigcup_{n=1}^\infty W_n$ such that $B \in W_{n+1}$ is adjacent to $B' \in W_n$ if and only if $B \cap E(G_n) = B'$. Thus, the conditions for Lemma 4.6 are satisfied, and we obtain for each $n \in \mathbb{N}$ a $B_n \in W_n$ so that $B_{n+1} \cap E(G_n) = B_n$ for all n. Clearly, $B := \bigcup_{n \in \mathbb{N}} B_n$ contains e.

To see that B is a bicycle, consider a finite cut F of G. Choose $N \in \mathbb{N}$ large enough so that $F \subseteq E(\widetilde{G}_N)$—then F is a cut in $\widetilde{G}_N$, too. We get

$$B * F = B * (F \cap E(\widetilde{G}_N)) = (B \cap E(\widetilde{G}_N)) * F = B_N * F = 0,$$

where the last equality follows since $B_N \in \mathcal{C}(\widetilde{G}_N)$. As F was arbitrary, Theorem 1.9 implies that $B \in \mathcal{C}(G)$. In a similar way, but using Lemma 4.3

in G_N instead of Theorem 1.9 in $\widetilde{G}_N$, we see that $B \in \mathcal{C}^*(G)$. Therefore, $B \in \mathcal{B}(G)$ and (i) holds.

Finally, suppose that there is a $B \in \mathcal{B}(G)$ with $e \in B$ and $Z \in \mathcal{C}_{\text{fin}}(G)$, $K \in \mathcal{C}^*_{\text{fin}}(G)$ with $\{e\} = Z + K$. Then, as B is both a cut and an element of the cycle space, we obtain

$$1 = \{e\} * B = (Z + K) * B = Z * B + K * B = 0,$$

which gives a contradiction. $\square$

Casteels and Richter [27] independently proved a complementary result:

Theorem 4.7 (Casteels and Richter [27])**.** *Let e be an edge of a locally finite graph G. Then either*

(i) *there exists $B \in \mathcal{B}_{\text{fin}}(G)$ with $e \in B$; or*

(ii) $\{e\} \in \mathcal{C}(G) + \mathcal{C}^*(G)$

but not both.

It should be noted that Casteels and Richter in fact prove a more general result of which Theorem 4.7 is but a consequence.

Theorems 4.5 and 4.7 look tantalisingly similar. The next lemma sheds some light on their relation.

Lemma 4.8. [21] *Let G be a locally finite graph. If for an edge e of G two of the following conditions hold, then the third one is satisfied, too:*

(i) *there is a $Y \in \mathcal{C}(G)$ with $e \in Y$ and $Y + e \in \mathcal{C}^*(G)$;*

(ii) *there is a $Z \in \mathcal{C}(G)$ with $e \notin Z$ and $Z + e \in \mathcal{C}^*(G)$;*

(iii) *there is a $B \in \mathcal{B}(G)$ with $e \in B$.*

If all of (i)–(iii) *hold for e, then each of Y, Z, B in* (i)–(iii) *is an infinite set.*

The lemma is reminiscent of a theorem by Richter and Shank [71] about (finite) surface duals. In fact, our proof uses similar arguments. We mention, moreover, that all of (i)–(iii) can hold for an edge. In Figure 4.2 we have already seen that e lies in an infinite bicycle, while in Figure 4.3 we witness the other two cases.

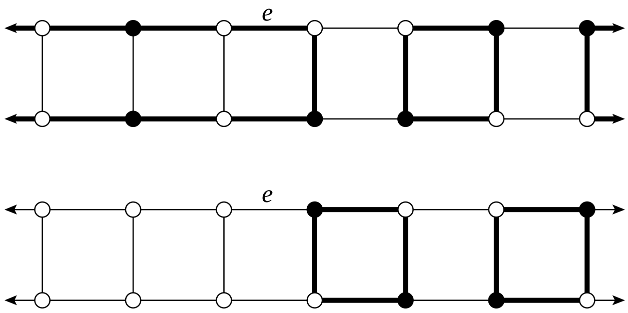

Figure 4.3: (i), (ii) in Lemma 4.8 hold for e

Proof of Lemma 4.8. First, assume (iii) and one of (i),(ii) to hold. Thus, there exist a $B \in \mathcal{B}(G)$ with $e \in B$ and an $X \in \mathcal{C}(G)$ so that $X + e \in \mathcal{C}^*(G)$. Since $X + B \in \mathcal{C}(G)$, we have that if X is as in (i), then $X + B$ satisfies (ii) and if, on the other hand, X is as in (ii), then $X + B$ satisfies (i).

Second, assume that (i) and (ii) hold, and let Y be as in (i) and Z as in (ii). Then $B := Y + Z \in \mathcal{C}(G)$, since $Y, Z \in \mathcal{C}(G)$. From $B = (Y+e)+(Z+e)$ it follows that B is also a cut. Finally, since $e \in Y$ but $e \notin Z$, we have $e \in B$.

For the second part of the lemma, assume that (i)-(iii) hold for e, and let $e \in B \in \mathcal{B}(G)$. By (the trivial part of) Theorem 4.7, it follows that B cannot be finite. On the other hand, Y and Z as in (i) resp. (ii) need to be infinite sets, too, since otherwise this would give a contradiction to Theorem 4.5. $\square$

Read and Rosenstiehl's theorem partitions the edges of a finite graph into three classes. So far, our theorem yields only two classes. So, let us refine Theorem 4.5. For this, we say that an edge e in a locally finite graph G is of *cut-type* if there is a finite cut K containing e so that $K \setminus \{e\} \in \mathcal{C}(G)$. We say that e is of *flow-type* if there is a finite element Z of the cycle space with $e \in Z$ and $Z \setminus \{e\} \in \mathcal{C}^*(G)$. Then, the following immediate corollary of Lemma 4.8 turns Theorem 4.5 into a true tripartition theorem:

Corollary 4.9. **[21]** *No edge in a locally finite graph can be of cut-type and of flow-type at the same time.*

We should point out that to denote by $\mathcal{C}^*(G)$ the set of all cuts is possibly a bit misleading as it might give the impression that it is the dual space of $\mathcal{C}(G)$. That, however, is not the case. Rather, Theorem 1.9 shows that, at least in some sense, $\mathcal{C}(G)$ and $\mathcal{C}^*_{\text{fin}}(G)$ are dual to each other. On the other hand, the dual space of $\mathcal{C}^*(G)$ is $\mathcal{C}_{\text{inf}}(G)$, see for instance [16].

In this respect, our bicycle space $\mathcal{B}(G)$ is situated between these two dualities. Examples as the graph in Figure 4.2 indicate that this is nevertheless justified since in order to make the tripartition theorem work in infinite graphs, whether it is in the form of Theorem 4.5 or in the form of Theorem 4.7, we need both spaces, $\mathcal{C}(G)$ and $\mathcal{C}^*(G)$.

4.4 Principal cuts

Let e be an edge of flow- or of cut-type in a locally finite graph G. Then, by definition, there is a $Z \in \mathcal{C}_{\text{fin}}(G)$ so that $Z + e \in \mathcal{C}^*(G)$. We call Z a *principal flow of* e and $Z + e$ a *principal cut of* e. In this section, we shall demonstrate, partially without proofs, that the properties of principal cuts carry over from finite graphs to locally finite graphs.

As a first notable property, let us see that the principal cuts are unique in a *pedestrian* graph, that is a graph G for which $\mathcal{B}(G) = \{\emptyset\}$. Indeed, let $K, K' \in \mathcal{C}^*(G)$ so that $K + e, K' + e \in \mathcal{C}(G)$. Then $K + K' = (K + e) + (K' + e) \in \mathcal{B}(G)$, which implies that $K = K'$ as $\mathcal{B}(G) = \{\emptyset\}$. For the purpose of this section, given a pedestrian graph let us denote the principal cut of an edge e by K_e and the principal flow by Z_e.

We need the following lemma, which (stated for finite graphs but with exactly the same proof) appears in Read and Rosenstiehl [74]. (We note that the lemma remains true in non-pedestrian graphs; Z_e (resp. Z_f) is then simply any principal flow through e (resp. f), as there is no longer a unique one. And similarly for K_e, K_f.)

Lemma 4.10. *Let e and f be edges in a locally finite pedestrian graph G. Then:*

(i) *$e \in Z_f$ if and only if $f \in Z_e$; and*

(ii) *$e \in K_f$ if and only if $f \in K_e$.*

Proof. To prove (i) consider

$$\begin{aligned}\{e\} * Z_f &= (Z_e + K_e) * Z_f = Z_e * Z_f + K_e * Z_f = Z_e * Z_f \\ &= Z_e * Z_f + Z_e * K_f = Z_e * (Z_f + K_f) = Z_e * \{f\}.\end{aligned}$$

Note that all these scalar products are well-defined since the Z_e and K_e are finite sets. Assertion (ii) is proved analogously. □

Proposition 4.11. **[21]** *In a locally finite pedestrian graph G both of the families $(Z_e)_{e \in E(G)}$ and $(K_e)_{e \in E(G)}$ are thin.*

Proof. Suppose there is an edge e lying in infinitely many Z_f. Since G is a pedestrian graph, e is of flow- or cut-type and Z_e is therefore defined. Thus Lemma 4.10 implies that $f \in Z_e$ for all these infinitely many f, contradicting that Z_e is finite. Thus $(Z_e)_{e \in E(G)}$ is thin. The proof for the principal cuts is the same. □

For an edge e to be of flow- or of cut-type we have required that there is a *finite* $Z \in \mathcal{C}(G)$ with $Z + e \in \mathcal{C}^*(G)$. In the light of Theorem 4.7 one could also quite reasonably relax this, and say that an edge is of flow- or cut-type if there is any such Z, finite or infinite. A pedestrian graph, then, would be one without any *finite* bicycles, since in precisely this case all edges are of flow- or cut-type.

There are several problems with this definition. We have already seen (Figures 4.2 and 4.3) that this would not give a proper tripartition. Furthermore, principal cuts in a pedestrian graph would not necessarily be unique and their family may not be thin. For instance, the cuts in the lower graph in Figure 4.3 would form a non-thin family of principal cuts.

The following corollary lists verbatim extensions of some basic properties of principal flows and cuts. Their proofs for finite graphs (substantially) use the finiteness only in one point, namely that it is allowed to take arbitrary sums of principal cuts. While, clearly, this is never an issue in finite graphs, such sums may be infinite in infinite graphs and then need to be thin in order to be well-defined. But this is exactly what Proposition 4.11 asserts.

Corollary 4.12. **[21]** *Let G be a locally finite pedestrian graph. Then*

(i) $(Z_e)_{e\in E(G)}$ *generates the cycle space; and*

(ii) $(K_e)_{e\in E(G)}$ *generates the cut space; and*

(iii) *the union of all flow-type edges is an element of the cycle space; and*

(iv) *the union of all cut-type edges is a cut.*

Proof. (i) and (ii) can be found in Read and Rosenstiehl [74] and (iii) and (iv) in Godsil and Royle [49]. □

4.5 Left-right tours

What should a left-right tour in an infinite plane graph be? Quite trivially, the name suggests two requirements for a left-right tour. Firstly, it should be "left-right", that is, locally it should consist of alternating left and right turns. And secondly, it should be a "tour", which means it should close up.

The first requirement is fairly simple to guarantee. Just as with left-right tours in finite graphs, we start a walk at an arbitrary edge and then alternately turn left and right. If we reach our starting edge again in this way, we have found a finite left-right tour. Otherwise, we prolong our walk in the other direction from our starting edge, again taking left and right turns. The resulting walk, which we call a *left-right string*, will be two-way infinite; two examples can be seen in Figure 4.5. In general, the two ends of a left-right string will not be identical, and the walk will therefore not be closed. So to achieve that we do not get stuck in an end, it will be necessary to glue together several left-right strings at ends. In this way we shall obtain a topological tour in $|G|$.

Let us start with left-right strings. To define these properly we shall first need to describe what it means to do a left turn followed by a right turn. We follow the treatment of Keir and Richter [55]. Let G be a locally finite graph, and let $|G|$ be embedded in the sphere S. Recall that, by Theorem 4.4, every locally finite planar graph has such an embedding. The interior of an edge of G is homeomorphic to the open unit interval $(0,1)$. For each edge e, we fix a homeomorphism. If η_1 denotes the image of the restriction of this homeomorphism to $(0,\frac{1}{2})$ and η_2 is the image of the restriction to $(\frac{1}{2},1)$ then η_1,η_2 are the *halves of* e. We use the notation $\overline{\eta_1}=\eta_2$ and $\overline{\eta_2}=\eta_1$ to switch back and forth between the two halves of an edge. Furthermore, we fix for e two open, disjoint and connected subsets, σ_1 and σ_2, of $S\setminus|G|$ each of which has e in its boundary. These are the *sides of* e, and as for the halves, we put $\overline{\sigma_1}=\sigma_2$ and $\overline{\sigma_2}=\sigma_1$. A triple (e,η,σ), where $e\in E(G)$, η is a half of e and σ is a side of e, is called a *corner of* $|G|$. We say that $c=(e,\eta,\sigma)$ is a *corner at* e, and it is a *corner at* $v\in V(G)$ if the boundary $\partial\eta$ contains v. Clearly, for each edge e there are four corners at e.

For each $v\in V(G)$ choose an open disc D around v, so that each half of an edge at v intersects ∂D in exactly one point. Then ∂D defines in a natural way a rotation of the halves. We say that two corners (e,η,σ), (e',η',σ') at v are *matched* if η and η' appear consecutively in the local rotation at v, and if the connected component K of $\sigma\cap D$ with $\eta\cap D\subseteq\partial K$ and the connected component K' of $\sigma'\cap D$ with $\eta'\cap D\subseteq\partial K'$ are contained in the same connected component of $D\setminus|G|$. It can be seen that this definition is independent of the actual choice of D. See Figure 4.4 for an illustration.

Corners can be used to describe left-right steps. Formally, this works as follows. Let $W=\ldots(e_{-1},\eta_{-1},\sigma_{-1}),(e_0,\eta_0,\sigma_0),(e_1,\eta_1,\sigma_1)\ldots$ be a (finite, one-way infinite or two-way infinite) sequence of corners satisfying the following properties:

(i) $(e_i,\overline{\eta_i},\overline{\sigma_i})$ and $(e_{i+1},\eta_{i+1},\sigma_{i+1})$ are matched for all i; and

(ii) no corner appears twice in W.

We call such a sequence W a *left-right walk*, which is justified by the fact that the edges $\ldots e_{-1}e_0e_1\ldots$ do indeed form a walk. Moreover, we will sometimes pretend that a left-right walk is in fact a walk, i.e. a sequence of vertices and

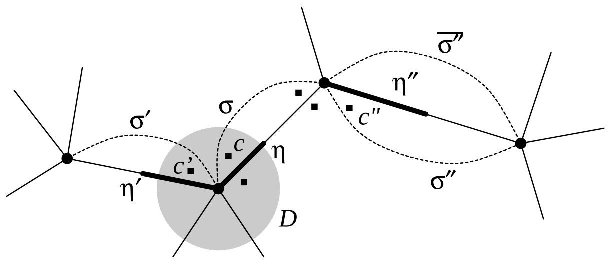

Figure 4.4: We think of a corner $c = (e, \eta, \sigma)$ at $v \in V(G)$ as a point close to v and η, and lying in σ. The corners c and $c' = (e', \eta', \sigma')$ are matched; the corners c and c'' describe a left-right step.

edges, rather than a sequence of corners. The corners c and c'' in Figure 4.4 describe a left-right step as in (i).

We say that S is a *left-right string* (LRS for short) if it is a maximal left-right walk. It is not hard to check that if $S = \ldots (e_{-1}, \eta_{-1}, \sigma_{-1}), (e_0, \eta_0, \sigma_0), (e_1, \eta_1, \sigma_1) \ldots$ then $S' := \ldots (e_1, \overline{\eta_1}, \overline{\sigma_1}), (e_0, \overline{\eta_0}, \overline{\sigma_0}), (e_{-1}, \overline{\eta_{-1}}, \overline{\sigma_{-1}}) \ldots$ is an LRS, too. Clearly, the walks S and S' traverse the same edges, but in opposite directions. Although we will sometimes view S as an oriented walk, we will, in general, not distinguish between S and S' and consider them to be identical. This slight abuse of notation ensures that every edge is covered exactly twice by LRS; see the next lemma. Figure 4.5 gives an example of two different LRS in the double ladder.

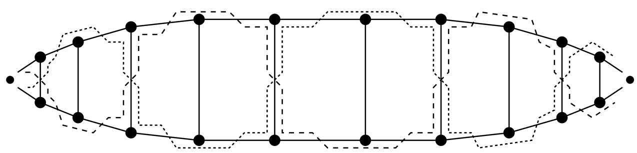

Figure 4.5: Two LRS in the double ladder

A set $\mathcal{W}$ of walks is a *double cover of* G if every edge $e \in E(G)$ is traversed exactly twice by walks in $\mathcal{W}$ (i.e. either once in two walks or twice in one walk). We leave out the proof of the following elementary observation.

Lemma 4.13. **[21]** *For a locally finite graph G, let $|G|$ be embedded in the sphere. Then:*

(i) *No two corners in an LRS are matched.*

(ii) *An LRS is either a closed walk or a two-way infinite walk.*

(iii) *The set of all LRS of G is a double cover of G.*

Observe that because of our somewhat tortuous definition of left-right walks as sequences of corners, (iii) remains true in pathological cases, such as when G is a double ray. Then, there are precisely two (distinct) LRS, which together form a double cover. Both of them traverse the double ray from one end to the other and are as walks indistinguishable. The corner sequences, however, are distinct.

Let G be a locally finite graph (not necessarily planar). We define a *tour T in $|G|$* to be a continuous map $T : S^1 \to |G|$ that is locally injective at every $x \in S^1$ for which $T(x)$ is an interior point of an edge. Note that, therefore, every edge with an interior point in the image of T, denoted by rge T, is completely contained in rge T. We denote the set of all edges that lie in rge T by $E(T)$. The *residue* $\bigtriangledown T$ of a tour T is the set of those edges that are traversed exactly once by T.

Now we can finally extend the definition of left-right tours to infinite graphs. Assume that $|G|$ is embedded in the sphere. Our aim is to give a definition so that an LRT consists of a number of LRS that are glued together at ends so as to constitute a tour in $|G|$. An example would be the two LRS shown in Figure 4.5 together with the two ends of the double ladder.

Formally, we define a *left-right tour L in $|G|$* (LRT for short) to be a tuple $(\mathcal{S}, \tau)$ where $\mathcal{S}$ is a set of LRS of G and $\tau : S^1 \to |G|$ a tour of $|G|$, so that each maximal subwalk of τ (in G, not in $|G|$) corresponds to one $S \in \mathcal{S}$ and vice versa. Usually, however, we will think of L as being a tour in $|G|$, and say that an LRS S *lies in* L if $S \in \mathcal{S}$.

Having defined LRTs, our first task is to prove that the residue of an LRT is indeed a bicycle. In finite graphs, this is due to Shank:

Lemma 4.14 (Shank [75]). *If G is a finite plane graph, then the residue of a left-right tour is a bicycle.*

Lemma 4.14 is proved with the help of plane dual graphs. While abstract dual graphs have been defined in [16], a suitable theory of plane dual graphs that involves infinite cycles has yet to be formulated. This is probably not overly difficult but checking the sometimes tedious geometrical details would take too much space and effort here. Rather, with the help of the next lemma, we will circumvent this obstacle by reducing the problem to finite graphs.

Lemma 4.15. **[21]** *For a locally finite graph G, let $|G|$ be embedded in the sphere. Let H be a finite plane subgraph, and let $L_1, \ldots, L_k$ be a set of LRTs of G so that no LRS of G lies in more than one L_i. Then there exist a finite plane supergraph H' of H and a set $L'_1, \ldots, L'_k$ of LRTs of H' so that for all*

$i = 1, \dots, k$, the LRT L_i traverses precisely the edges $e_1, \dots, e_n$ of H and in this order if and only if L_i' does.

Proof. From the given finite plane subgraph H of G we will construct a finite plane supergraph H' of H (which will not necessarily be a subgraph of G) with the required properties. We may assume H to be induced. Each L_i decomposes in H into a set of walks. Our task is to draw in the faces of H finite graphs so that the subwalks in the set $L_i \cap H$ connect up in the same order as in G (for all i). Since this will be done in the same way in every face, we may assume in what follows that all of $G - H$ is contained in one face.

Denote by F those edges in the cut $E(H, G-H)$ that lie in some L_i, and find in the one face that contains $G - H$ an open disc D so that each edge in F meets ∂D in its interior. For each edge e in F, running along e from H towards $G - H$ we pick the first point, x say, in ∂D and cut off the edge at x. We draw a vertex at x and let the set of these x be X. We denote by H_0 the finite plane graph consisting of H together with the cut-off edges in F (plus the vertices in X). While, technically, F is a subset of $E(G)$, we will view it as a subset of $E(H_0)$, too.

Consider an LRT L, and let $\mathcal{S}$ be the set of LRS that lie in L (here, of the two orientations of an LRS $S \in \mathcal{S}$, we pick the one that is induced by L). We define the set of corners $\mathcal{K}_L$ to be $\bigcup_{S \in \mathcal{S}} S$, and observe that L induces a cyclic ordering on the LRS in $\mathcal{S}$, and therefore also on $\mathcal{K}_L$. Furthermore, we let $\mathcal{M}$ be those of the corners in $\bigcup_{i=1}^{k} \mathcal{K}_{L_i}$ that are corners at edges in F. Clearly, for each corner in $\mathcal{M}$, which is a corner in G, there is a corresponding corner in H_0. For the sake of simplicity, we will not distinguish between these two and, depending on the context, view $\mathcal{M}$ as a set of corners either in G or in H_0. Corners in $\mathcal{M}$ come in two kinds: there are *outgoing* corners, i.e. corners at vertices in $V(H)$, and *ingoing* corners, those at vertices in X.

Next, we will construct a pairing of the corners in $\mathcal{M}$. For each i, we arbitrarily pick an outgoing corner c_1 in $\mathcal{M} \cap \mathcal{K}_{L_i}$. Then, let $c_1, \dots, c_l$ be the corners in $\mathcal{M} \cap \mathcal{K}_{L_i}$ in the cyclic order of $\mathcal{K}_{L_i}$. Since L_i is a tour, l is even and for each odd j the corner c_j is outgoing while c_{j+1} is ingoing. We pair up consecutive corners: $\{c_1, c_2\}, \dots \{c_{l-1}, c_l\} \in \mathcal{P}$. For later use, we note that

$$\textit{if } \{c, c'\} \in \mathcal{P} \textit{ then one of } c, c' \textit{ is outgoing and one ingoing.} \tag{4.1}$$

Our task is to find finite left-right walks between each pair $\{c, c'\} \in \mathcal{P}$. The definition of $\mathcal{P}$ then ensures that for each i the order of the corners in $\mathcal{K}_{L_i}$ within H is maintained.

Define for each $c \in \mathcal{M}$ a left-right walk $K^0(c) := (c)$, i.e. $K^0(c)$ is a walk of length 1, which traverses an edge in F. To simplify the construction in the next steps we will, with the help of a suitable homeomorphism, identify D with $(0,3) \times (0,1) \subseteq \mathbb{R}^2$, where all the vertices in X are assumed to lie in the open segment $\{0\} \times (0,1)$; see Figure 4.6.

Next, we pick $m := |\mathcal{M}|$ distinct points $x_1^1, \dots, x_m^1$ in $\{1\} \times (0,1)$, where we choose the labelling so that x_j^1 has a smaller y-coordinate than x_{j+1}^1 for all j. We consider these points to be vertices and draw non-crossing edges in $(0,1) \times (0,1)$ in order to join each x_j^1 to a vertex w in X so that w receives one edge if its incident edge in F is only traversed once by $L_1, \dots, L_k$; otherwise (when the edge is used twice) we make w adjacent to two of the x_j^1. Clearly, in the resulting plane supergraph H_1 of H_0 each vertex in $x_1^1, \dots, x_m^1$ has degree 1.

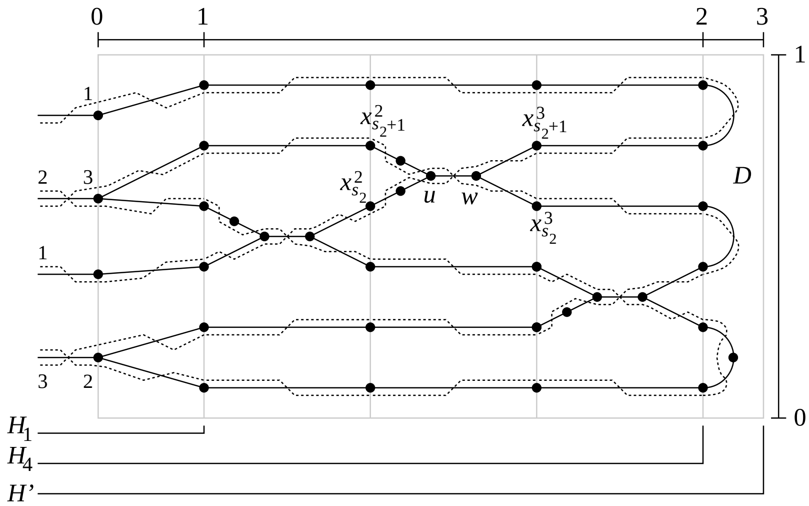

Figure 4.6: The construction of the H_i (not to scale)—corners with the same number are supposed to be paired.

Consider $c = (e, \eta, \sigma) \in \mathcal{M}$. Assume first that c is an ingoing corner. If c is matched with (e', η', σ') (in H_1), we precede the edge e in $K^0(c)$ by e' in order to obtain the left-right walk $K^1(c)$, i.e. we put $K^1(c) := ((e', \overline{\eta'}, \overline{\sigma'}), c)$. (Observe, that in this case, the walk is directed towards H, and hence we have to lengthen it in backward direction.) Second, assume that c is outgoing. If $(e, \overline{\eta}, \overline{\sigma})$ is matched with $c'' := (e'', \eta'', \sigma'')$ (in H_1) we lengthen $K^0(c)$ along the edge e'' to $K^1(c)$, that is, we set $K^1(c) := (c, c'')$. In this way, we define left-right walks $K^1(c)$ for all $c \in \mathcal{M}$, so that each vertex in $x_1^1, \dots, x_m^1$ is used by a unique $K^1(c)$, and this $K^1(c)$ either starts or ends in that vertex.

We will construct supergraphs H_i of H_1 with corresponding left-right walks $K^i(c) \supseteq K^1(c)$, $c \in \mathcal{M}$. More precisely, we will construct finitely many nested plane supergraphs with $H_1 \subset H_2 \subset \ldots \subset H_{t+1}$, where $H_i \setminus H_{i-1}$

is entirely drawn in $(a, b] \times (0,1)$ for some $1 \leq a < b < 3$ (we will determine the respective a and b in a moment). The intersection of H_i with $\{b\} \times (0,1)$ will consist of m vertices; in the order we encounter them on $\{b\} \times (0,1)$ going from $(b, 0)$ to $(b, 1)$ these will be denoted by $x_1^i, \ldots, x_m^i$. For each $j = 1, \ldots, m$ there will then be a unique corner $p_j^i \in \mathcal{M}$ so that the left-right walk $K^i(p_j^i)$ either starts or ends in x_j^i (and is otherwise disjoint from $x_1^i, \ldots, x_m^i$).

Let $(p_1, \ldots, p_m)$ be a permutation of $\mathcal{M}$. For the rest of the proof let us call a *flip at* $s \in \{1, \ldots, m-1\}$ the operation that turns $(p_1, \ldots, p_m)$ into $(p_1, \ldots, p_{s-1}, p_{s+1}, p_s, p_{s+2}, \ldots, p_m)$. Clearly, for some t there is a sequence of t flips at $s_1, \ldots, s_t$ that turns $(p_1^1, \ldots, p_m^1)$ into $(q_1, \ldots, q_m)$ so that for each odd j in $\{1, \ldots, m\}$ it holds that $\{q_j, q_{j+1}\} \in \mathcal{P}$.

Our aim now is to define H_{i+1}, for $i \in \{1, \ldots, t\}$, in such a way that $(p_1^{i+1}, \ldots, p_m^{i+1})$ is obtained from $(p_1^i, \ldots, p_m^i)$ by performing a flip at s_i. Moreover, with the exception of the points $x_1^{i+1}, \ldots, x_m^{i+1}$, we will draw $H_{i+1} \setminus H_i$ in $(1 + \frac{i-1}{t}, 1 + \frac{i}{t}) \times (0,1)$. Assume $H_1, \ldots, H_i$ to be constructed. We put m distinct vertices $x_1^{i+1}, \ldots, x_m^{i+1}$ (in this order) on the segment $\{1 + \frac{i}{t}\} \times (0,1)$. For each $j \in \{1, \ldots, m\}$ with $j \neq s_i, s_i + 1$, draw a straight line between x_j^i and x_j^{i+1}. We extend $K^i(p_j^i)$ to a left-right walk $K^{i+1}(p_j^i)$ along the edge $x_j^i x_j^{i+1}$. Then we draw an edge uw in $(1 + \frac{i-1}{t}, 1 + \frac{i}{t}) \times (0,1)$ so that no crossing edges arise when we connect u to $x_{s_i}^i$ and $x_{s_i+1}^i$, and w to $x_{s_i}^{i+1}$ and $x_{s_i+1}^{i+1}$. If necessary, we subdivide the edge $x_{s_i}^i u$ in order to guarantee the existence of a left-right walk from $x_{s_i}^i$ through uw to $x_{s_i+1}^{i+1}$ (that is disjoint from $x_{s_i+1}^i$). We extend $K^i(p_{s_i}^i)$ by this walk to a left-right walk $K^{i+1}(p_{s_i}^i)$, and proceed in an analogous way for $K^i(p_{s_i+1}^i)$. This ensures that $(p_1^{i+1}, \ldots, p_m^{i+1})$ is obtained from $(p_1^i, \ldots, p_m^i)$ by performing a flip at s_i.

Finally, assume all the H_i up to H_{t+1} to be constructed. For each odd j in $\{1, \ldots, m\}$, we draw an edge in $(2,3) \times (0,1)$ that joins x_j^{t+1} to x_{j+1}^{t+1}. Subdividing $x_j^{t+1} x_{j+1}^{t+1}$ if necessary, we can join $K^{t+1}(p_j^{t+1})$ by this (possibly subdivided) edge to $K^{t+1}(p_{j+1}^{t+1})$, so that the resulting walk is left-right (here, (4.1) ensures that the corner sequences fit with respect to orientation). By construction of the pairing $\mathcal{P}$, we ensure that the resulting LRTs L_i' in the plane graph H' ($:= H_{t+1}$ plus the possibly subdivided edges in $(2,3) \times (0,1)$) behave on H in the same way as the L_i do. $\square$

Lemma 4.16. **[21]** *For a locally finite graph G, let $|G|$ be embedded in the sphere. Then the residue of an LRT in G is an element of the bicycle space.*

Proof. Let F be a finite cut and L an LRT. As a tour, L passes an even number of times through F. Therefore, $| \bigtriangledown L \cap F|$ is even and it follows, by Theorem 1.9, that $\bigtriangledown L$ is an element of the cycle space.

To see that the residue $\bigtriangledown L$ is a cut, consider a finite cycle C. Lemma 4.15 (with $H = C$) yields a finite plane supergraph H' of C and an LRT L' of H' so that $\bigtriangledown L \cap E(C) = \bigtriangledown L' \cap E(C)$. As $\bigtriangledown L'$ is a cut in H' (by Lemma 4.14) and $C \subseteq H'$ a cycle, we have that $\bigtriangledown L' \cap E(C)$ is an even set. Since this implies that $\bigtriangledown L \cap E(C)$ is even, too, it follows from Lemma 4.3 that $\bigtriangledown L \in \mathcal{C}^*(G)$ and hence $\bigtriangledown L \in \mathcal{B}(G)$. $\square$

4.6 LRTs generate the bicycle space

In this section we will prove the analogue of Theorem 4.2 for locally finite graphs.

Let G be a locally finite graph for which $|G|$ is embedded in the plane, and consider a bicycle B of G. Since the cuts of the form $E(v)$ generate the cut space, there is a vertex set X such that $B = \sum_{x \in X} E(x)$. On the other hand, B is also an element of the cycle space. As in finite graphs, $\mathcal{C}(G)$ is generated by the residues of the face boundaries (this is shown in [23]). Thus, there is a set F of face boundaries such that $B = \sum_{f \in F} \bigtriangledown f$. For each bicycle B assume such a pair X, F to be fixed. Following Richter and Shank [71], we say that an LRS S is of *type I* if there is a corner $c = (e, \eta, \sigma)$ in S for which the following statements are either both true or both false:

(i) $\partial\eta$ contains a vertex in X; and

(ii) σ lies in a face whose face boundary is in F.

It is not hard to check that if for one corner in S either both of (i) and (ii) are true or are both false then this holds for every corner in S; see also Richter and Shank [71]. If S is not of type I, then S is of *type II.*

Lemma 4.17. *Let G be a locally finite plane graph, and let B be a bicycle. Then an edge e of G lies in B if and only if it lies in exactly one LRS of type I and in one LRS of type II with respect to B.*

Proof. The proof is identical to the one given for finite graphs in Richter and Shank [71]. $\square$

An LRT L is called *B-uniform* if every two LRS contained in L are of the same type. In finite graphs, Lemma 4.17 is already enough to prove Theorem 4.2: we only need to sum up all LRS (which are identical to LRTs in finite graphs) of type I (or type II, for that matter). By contrast, in locally finite graphs, it is not even clear whether there is a single B-uniform LRT, let alone a set of B-uniform LRTs with the properties as in the last lemma. The next lemma asserts the existence of B-uniform LRTs.

Lemma 4.18. [21] *Let G be a locally finite graph, let $|G|$ be embedded in the sphere, and let B be a bicycle of G. Then there exists a set $\mathcal{L}$ of B-uniform LRTs so that each LRS of G is contained in exactly one $L \in \mathcal{L}$.*

Proof. We may assume G to be connected. Then there is an enumeration $S_1, S_2, \ldots$ of the set of LRS of G, since G is countable.

We construct from G another locally finite graph G' (which, in all likelihood, will not be planar). The vertex set of G' consists of vertices v_p, one for each vertex v of G and for each subwalk p of the form $p = evf$ in each S_i ($e, f \in E(G)$). Such a vertex $v_p \in V(G')$ is called a *clone of* v. The edge set of G' is comprised of two disjoint sets, E' and F'. The set F' contains one edge between each pair of clones v_p and v_q of the same vertex $v \in V(G)$; i.e. the clones of a vertex span a complete graph. Two clones u_p and v_q of distinct vertices $u, v \in V(G)$ are connected by an edge in E' if p and q are subwalks in the same LRS S_i and appear consecutively in S_i, i.e. if $S_i = \ldots e_{-1}v_{-1}e_0v_0e_1v_1e_2\ldots$ then $p = e_{j-1}v_{j-1}e_j$ and $q = e_jv_je_{j+1}$ (or the other way round) for some j. See Figure 4.7 for an illustration.

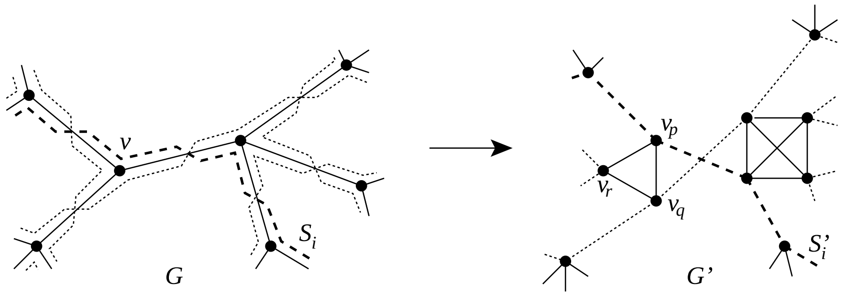

Figure 4.7: Construction of G' in the proof of Lemma 4.18; the edges in E' are dotted.

Let us define a mapping $\phi : V(G') \cup E(G') \to V(G) \cup E(G)$. For each $v \in V(G)$ we map all clones of v and all edges (in F') between two clones of v to v. An edge u_pv_q in E', where u_p is a clone of $u \in V(G)$ and v_q is a clone of $v \neq u$, is mapped to the edge uv of G. Clearly, this map is surjective.

We note, furthermore, that because of Lemma 4.13 (iii),

$$\text{each } e \in E(G) \text{ has exactly two preimages under } \phi, \text{ and these are in } E'. \tag{4.2}$$

For each $S_i = \ldots e_{-1}v_{-1}e_0v_0e_1v_1e_2\ldots$, the map ϕ defines a walk in G'. Indeed, since there is a vertex v_{p_j} in G' for each subwalk $p_j := e_jv_je_{j+1}$, and since each v_{p_j} is linked by an edge e'_{j+1} in E' to $v_{p_{j+1}}$, the sequence $\ldots e'_{-1}v_{p_{-1}}e'_0v_{p_0}e'_1v_{p_1}e'_2\ldots$ is a walk in G', which we denote by S'_i. We claim that for all i it holds that

(i) if $S'_i = \dots e'_{-1} v'_{-1} e'_0 v'_0 e'_1 v'_1 e'_2 \dots$ then
$S_i = \dots \phi(e'_{-1})\phi(v'_{-1})\phi(e'_0)\phi(v'_0)\phi(e'_1)\phi(v'_1)\phi(e'_2)\dots$; and

(ii) each S'_i is either a cycle or a double ray; and

(iii) S'_i and S'_j are disjoint for all $j \neq i$.

Claim (i) is clear by construction, and for (ii) and (iii) simply note that a clone v_p of a vertex $v \in V(G)$ is adjacent to exactly two vertices that are not clones of v.

Denote by $\mathcal{X}_I$ the set of all those S'_i for which S_i is of type I with respect to B, and let $\mathcal{X}_{II}$ be the set of the other S'_i (those for which S_i is of type II). We will show that

$$\textit{both of } X_I := \bigcup_{S' \in \mathcal{X}_I} E(S') \textit{ and } X_{II} := \bigcup_{S' \in \mathcal{X}_{II}} E(S') \textit{ lie in } \mathcal{C}(G'). \tag{4.3}$$

To show that $X_I \in \mathcal{C}(G')$, consider a finite cut K' of G'; by Theorem 1.9, it suffices to prove that $|X_I \cap K'|$ is even.

Fix a vertex a' of G', and for each finite cut $L = E_{G'}(A, B)$ of G' with $a' \in A$ denote by $c(L)$ the number of vertices $w' \in B$ so that there exists a clone $u' \in A$ of the same vertex as w'. Since, by definition, each such w' is adjacent to a vertex in A, the number $c(L)$ is finite.

Now, among all finite cuts L for which $|L \cap X_I|$ has the same parity as $|K' \cap X_I|$ choose one, K say, so that $c(K)$ is minimal. Suppose that $c(K) > 0$, and let $K = E_{G'}(A, B)$ with $a' \in A$. Since $c(K) > 0$ there exist $u' \in A$ and $w' \in B$ that are clones of the same vertex $v \in V(G)$. As $w' = v_p$ for some subwalk p in some S_i, we obtain from (iii) that w' lies in exactly one S'_i, which implies that w' is incident with exactly zero or two edges in X_I, depending on whether S_i is of type II or of type I. Thus, the cut $\widetilde{K} := K + E(w')$ meets X_I in an even number of edges if and only if $|K \cap X_I|$ is even. On the other hand, we have $\widetilde{K} = E_{G'}(A \cup \{w'\}, B \setminus \{w'\})$, which implies $c(\widetilde{K}) < c(K)$, which contradicts the choice of K.

Therefore, it holds that $c(K) = 0$. Since all clones of a vertex are on the same side of K, it follows that $K \subseteq E'$, that $\phi(K)$ is a finite cut of G, and that for each $e \in \phi(K)$ both of the preimages of e under ϕ lie in K. Thus, if we can show that $\phi(K)$ is traversed an even number of times by LRS of type I (with respect to B), then $|X_I \cap K|$ is even, and hence so is $|X_I \cap K'|$.

Lemmas 4.13 (iii) and 4.17 imply that $\phi(K) \setminus B$ is traversed an even number of times by LRS of type I. Since B is an element of the cycle space, the set $B \cap \phi(K)$ is even, by Theorem 1.9. Thus, Lemma 4.17 implies that also $B \cap \phi(K)$ is traversed an even number of times by LRS of type I. With (4.2) we get that $|X_I \cap K|$ is even. The proof for X_{II} is the same.

Next, we use Theorem 1.8 to decompose $X_I + X_{II}$ into a set $\mathcal{D}$ of (edge-)disjoint circuits. We observe that

$$\text{for all } i \text{ and } D \in \mathcal{D} \text{ it holds that if } E(S_i') \cap D \neq \emptyset \text{ then } E(S_i') \subseteq D. \text{ Moreover, for each } D \in \mathcal{D}, \text{ all the } S_i \text{ with } E(S_i') \subseteq D \text{ are of the same type.} \tag{4.4}$$

Indeed, by (ii) and (iii) every vertex of G' is incident with exactly two or zero edges of X_I (resp. X_{II}). Since this also holds for circuits, the assertion follows.

Next, we define a continuous mapping $\phi' : |G'| \to |G|$. On the 1-complex G' we extend ϕ to a continuous mapping ϕ' so that the following holds:

(a) $\phi'(e') = e$ if and only if $\phi(e') = e$ for all $e' \in E(G')$ and $e \in E(G)$ (where, with regard to ϕ' we view e' and e as point sets, while for ϕ we see them as edges of graphs); and

(b) at each interior point of an edge of G', the map ϕ' is locally injective.

To define ϕ' on ends, consider a ray R' in an end ω' of G'. Then $\phi(R')$ is a one-way infinite walk, and thus contains a ray in an end, say ω. We map ω' to ω.

It remains to check that ϕ' is continuous at ends. So, consider an end ω' of G' and let a basic open neighbourhood $C := \hat{C}_G(U, \phi'(\omega'))$ of $\phi'(\omega')$ in $|G|$ be given (recall that U is a finite vertex set). Denoting by U' the set of all clones of vertices in U, we see that $C' := \hat{C}_{G'}(U', \omega')$ is a basic open neighbourhood of ω' in $|G'|$ and that $\phi'(C') \subseteq C$. Therefore, ϕ' is continuous.

Finally, since each $D \in \mathcal{D}$ is a circuit, by definition there exists a homeomorphism $\sigma_D : S^1 \to |G'|$ with image $\overline{D}$. By (b), the continuous mapping $\phi' \circ \sigma_D : S^1 \to |G|$ is locally injective at points $x \in S^1$ that are mapped to interior points of edges. Furthermore, (i) and (a) imply that each maximal subwalk in $\phi' \circ \sigma_D$ is an LRS, and that these are precisely those S_i for which $E(S_i') \subseteq D$. Therefore, each $\phi' \circ \sigma_D$ describes an LRT in $|G|$. By (4.4), each such LRT is B-uniform. We denote the set $\{\phi' \circ \sigma_D : D \in \mathcal{D}\}$ of LRTs by $\mathcal{L}$.

Since for every S_i the set $E(S_i')$ is contained in some $D \in \mathcal{D}$, every S_i occurs in one of the LRTs in $\mathcal{L}$, and on the other hand, since all the $D \in \mathcal{D}$ are (edge-)disjoint, no S_i appears in two elements of $\mathcal{L}$. □

We remark that the LRTs in $\mathcal{L}$ have an additional property, of which we will, however, make no use: each $L \in \mathcal{L}$ is *minimal* in the sense that, if L' is an LRT with $\emptyset \neq E(L') \subseteq E(L)$ then $E(L') = E(L)$. In order to briefly sketch the proof, let $D \in \mathcal{D}$ be the circuit in G' so that $\phi' \circ \sigma_D$ describes the LRT L. Let $\mathcal{Y}$ be the subset of LRS contained in L that also lie in L'. Then it is easy to check that $Y := \bigcup_{S \in \mathcal{Y}} E(S')$ is an element of the cycle space of G'. Since Y is not empty and a subset of the circuit D, it follows that $Y = D$ which implies $E(L) = E(L')$, as claimed.

With Lemma 4.18 we can extend Theorem 4.2 to locally finite graphs using arguments of Richter and Shank [71]. Given a bicycle B, Lemma 4.18 yields a set $\mathcal{M}$ of LRTs, so that every LRS of type I appears in exactly one element of $\mathcal{M}$. Lemma 4.17 assures that $\sum_{M \in \mathcal{M}} \triangledown M = B$. On the other hand, Lemma 4.16 shows that all sums of residues of LRTs are elements of the bicycle space. In conclusion, we have proved:

Theorem 4.19. **[21]** *Let G be a locally finite graph, and let $|G|$ be embedded in the sphere. Then the residues of the left-right tours in $|G|$ generate the bicycle space of G.*

In a finite graph, the set of LRTs is a double cover. In the double ladder, by contrast, we can construct LRTs by glueing together any two of the four LRS, which results in a set of six LRTs that cover all edges more than twice; see Figure 4.5. Moreover, while Lemma 4.18 asserts that there are double covers consisting of LRTs, in the case of the double ladder none of these are sufficient to generate the bicycle space. Indeed, consider a double cover $\mathcal{L}$ of LRTs for the double ladder. Pick an LRT of the double cover and observe that it traverses some edge e twice (in Figure 4.5 this is the case for every second rung). It is easy to check that every edge in the double ladder lies in a bicycle, and hence, no bicycle containing e can be expressed as the sum of residues of $L \in \mathcal{L}$.

4.7 The ABL planarity criterion

MacLane's well-known planarity criterion [59] characterises planar graphs in terms of the cycle space. MacLane observed that, in (finite) plane graphs, the set of facial walks is a double cover that generates the cycle space. Then he proved that, conversely, any double cover of closed walks with this property can be realised as a set of facial walks and is therefore a certificate for planarity.

The planarity criterion of Archdeacon, Bonnington and Little [6] works in a similar way with the difference that they list the essential properties of the left-right tours. These properties are rather more elaborate and necessitate a number of definitions, which we will give below. In this section it is our aim to show that the ABL criterion remains true in locally finite graphs.

Consider a locally finite graph G, and let $\mathcal{W}$ be a double cover of tours in $|G|$, i.e. every edge is traversed twice by $\mathcal{W}$. For any l, let $\mathcal{H}$ be a cyclic sequence $e = f_1, W_1, \dots, f_l, W_l, f_{l+1} = e$ where the W_i are distinct members of $\mathcal{W}$ and the f_j are distinct edges of G, so that W_i contains both of f_i and f_{i+1}. We call such a sequence $\mathcal{H}$ a *ladder (with respect to $\mathcal{W}$)*, and we say that the f_i are the *rungs* of $\mathcal{H}$.

For each i, let W_i' be one of the two orientations of W_i, and denote by P_i the topological subpath in W_i' between f_i and f_{i+1}, and by P_i' the one between f_{i+1} and f_i; i.e. traversing f_i, then following P_i, traversing f_{i+1} and finally running along P_i' describes the same tour in $|G|$ as W_i'. An edge that is traversed both times in the same direction by the W_i' (either by one W_i', in which it appears twice, or by two distinct tours), is said to be *consistent*; otherwise it is *inconsistent*. We call the family $(P_i)_{i=1,\dots,l}$ together with the set of inconsistent rungs (with respect to the W_i') a *side of $\mathcal{H}$*. Furthermore, if the side is denoted by S, then we write $\triangledown S$ for $\sum_{i=1}^{l} \triangledown P_i + \sum_{j\in J} f_j$ where $J = \{j : 1 \leq j \leq l \text{ and } f_j \text{ is inconsistent}\}$.

Finally, a double cover $\mathcal{D}$ of tours of G is called a *diagonal* if both $\triangledown D$ and $\triangledown S$ are cuts, for every $D \in \mathcal{D}$ and every side S of any ladder in $\mathcal{D}$.

We can now state the ABL criterion:

Theorem 4.20 (Archdeacon, Bonnington and Little [6])**.** *A finite graph is planar if and only if it has a diagonal. In particular, the set of LRTs of a finite plane graph is a diagonal.*

A simple proof of the ABL criterion can be found in Keir and Richter [55]. Theorem 4.20 extends to locally finite graphs:

Theorem 4.21. [21] *A locally finite graph is planar if and only if it has a diagonal.*

Proof. Let G be a locally finite graph. First, assume G to be planar. From Theorem 4.4 we know that $|G|$ has an embedding in the sphere, and thus Lemma 4.18 yields (with, for instance, $B = \emptyset$) a set $\mathcal{L}$ of LRTs so that each LRS of G lies in exactly one element of $\mathcal{L}$. Hence, $\mathcal{L}$ is a double cover of G (Lemma 4.13 (iii)). Furthermore, Lemma 4.16 implies that $\triangledown L$ is a cut for each $L \in \mathcal{L}$.

For $\mathcal{L}$ to be a diagonal, it remains to show that for any side S of any ladder $\mathcal{H}$ (with respect to $\mathcal{L}$), $\triangledown S$ is a cut as well. We show that $\triangledown S$ meets every finite cycle C in an even number of edges, thereby proving $\triangledown S$ to be a cut (Lemma 4.3).

If R is the set of rungs of $\mathcal{H}$, then we define H to be the plane subgraph of G consisting of C and all the edges in R together with their incident vertices. We apply Lemma 4.15 to H and the LRTs in $\mathcal{H}$, which yields a finite plane supergraph H' and a set $\mathcal{H}'$ of LRTs of H'. It is straightforward to see that $\mathcal{H}'$ is a ladder in H' with a side S' for which it holds that $\triangledown S' \cap E(H) = \triangledown S \cap E(H)$. Since $\triangledown S'$ is a cut, by Theorem 4.20, the intersection $\triangledown S' \cap E(C) = \triangledown S \cap E(C)$ is even. This proves $\mathcal{L}$ to be a diagonal.

For the converse direction, let us now suppose that G has a diagonal $\mathcal{D}$ but also contains a subdivision X of $K_{3,3}$ or of K_5. Denote by H the (finite) induced subgraph of G on $V(X)$, and set $F := E(H, G - H)$, which is a finite cut. One by one, we delete the edges of F from G. We claim that after each edge deletion, the graph G still has a diagonal. For finite graphs, this is proved in Archdeacon, Bonnington and Little [6]. As their arguments remain still valid in locally finite graphs, we will not repeat them.

Once we have deleted all of F, the diagonal will split into two parts: into the set $\mathcal{D}'$ of those that are completely contained in H, and into those tours that are disjoint from H. Clearly, $\mathcal{D}'$ is then a diagonal of the finite non-planar graph H, which is impossible by Theorem 4.20. □

For pedestrian graphs, i.e. those graphs G for which $\mathcal{B}(G) = \{\emptyset\}$, Read and Rosenstiehl [74] gave a slightly simpler planarity criterion. Let a tour W traverse an edge $e = uv$ twice. If e is consistent, and traversed from u to v, say, then W decomposes into four topological subpaths uv, H_1, uv and H_2. We call each of H_1 and H_2 a *half of* W (with respect to e). If e is inconsistent, then W is equally comprised of four topological subpaths: namely of uv, H_1', vu and of H_2'. In this case we call the topological subpaths uvH_1' and vuH_2' *halves of* W.

We note two facts: first, if e is inconsistent in W then it is contained in each half of W; and second, if e, W, e is seen as a ladder then a half is simply a side of this ladder (more precisely, they have the same residues).

We say that a tour D in $|G|$ is an *algebraic diagonal of* G if D is a double cover and if for every edge e, every half of D is a cut.

Theorem 4.22 (Read and Rosenstiehl [74])**.** *A finite connected pedestrian graph is planar if and only if it has an algebraic diagonal.*

Theorem 4.23. **[21]** *A locally finite connected pedestrian graph is planar if and only if it has an algebraic diagonal.*

Proof. Let G be a locally finite connected pedestrian graph. If G is planar, then $|G|$ can be embedded in the sphere (Theorem 4.4) and there is a family $\mathcal{L}$ of LRTs of G that forms a double cover (by Lemma 4.18). We already know (from the proof of Theorem 4.21) that $\mathcal{L}$ is a diagonal. If $\mathcal{L}$ has only a single member D, then D is an algebraic diagonal of G: since every half H of D is the side of a ladder, it follows that $\triangledown H$ is a cut.

So, assume that $\mathcal{L}$ has two members, and denote one of them by L. Since G is pedestrian, Lemma 4.16 implies $\triangledown L = \emptyset$. As G is connected there is therefore a vertex v which is met by L but also incident with edges not lying in L. Consider an edge e incident with v that lies in L. Let η be the half of e with $v \notin \partial\eta$, and let σ be a side of e. Since L traverses e twice (as $\bigtriangledown L = \emptyset$), L (or, more precisely, the LRS lying in L) contains one corner of each of $\{(e, \eta, \sigma), (e, \overline{\eta}, \overline{\sigma})\}$ and $\{(e, \eta, \overline{\sigma}), (e, \overline{\eta}, \sigma)\}$. Let (e_1, η_1, σ_1) be the corner that is matched with $(e, \overline{\eta}, \overline{\sigma})$, and let (e_2, η_2, σ_2) be the one matched with $(e, \overline{\eta}, \sigma)$. By definition, if L contains (e, η, σ) then it also contains (e_1, η_1, σ_1). If, on the other hand, $(e, \overline{\eta}, \overline{\sigma})$ lies in L, then $(e_1, \overline{\eta}_1, \overline{\sigma}_1)$ is a corner of L. In any case, e_1 is traversed by L. As, in a similar way, we see that e_2 lies in L as well, it follows that the predecessor and the successor of e in the local rotation at v both lie in L, and thus that all of $E(v)$ is covered by L, a contradiction to our assumption.

If, conversely, G has an algebraic diagonal D, then the set $\{D\}$ is a diagonal. Theorem 4.21 shows that G is planar. □

Chapter 5

MacLane's planarity criterion

5.1 MacLane's planarity criterion for locally finite graphs

MacLane's well-known planarity criterion [59, 31] characterises the finite planar graphs in terms of their cycle space. For a graph G, call a family $\mathcal{F}$ of sets $F \subseteq E(G)$ *sparse* if every edge of G lies in at most two members of $\mathcal{F}$.

MacLane's planarity criterion can then be stated as follows:

MacLane's Theorem. *A finite graph is planar if and only if its cycle space is generated by some sparse family of (edge sets of) cycles.*

The question how MacLane's criterion could be extended to encompass infinite graphs was already raised by Wagner [91]. Thomassen [78] gave a partial answer by characterising those infinite graphs that satisfy MacLane's original condition. For this, we say that a *vertex accumulation point*, abbreviated VAP, of a plane graph Γ is a point p of the plane such that every neighbourhood of p contains an infinite number of vertices of Γ. Moreover, as in Chapter 4 we denote by $\mathcal{C}_{\mathit{fin}}(G)$ the set of all finite sums of finite circuits.

Theorem 5.1 (Thomassen [78])**.** *Let G be an infinite 2-connected graph. Then G has a VAP-free embedding in the plane if and only if $\mathcal{C}_{\mathrm{fin}}(G)$ has a sparse generating set consisting of finite circuits.*

An extension in another direction is due to Bonnington and Richter [13], who proved a necessary and sufficient condition for a graph to have a drawing with k VAPs.

Neither of these two results fully extends MacLane's planarity criterion to infinite graphs. This only becomes possible, at least for locally finite graphs, within the framework of the topological cycle space:

Theorem 5.2. **[23]** *A countable locally finite graph is planar if and only if its topological cycle space has a sparse generating set.*

In this chapter, which is based on [17], we generalise MacLane's theorem to embeddability criteria for arbitrary closed surfaces. While we only treat finite graphs, we conjecture that our criteria extend to locally finite graphs.

5.2 MacLane's theorem for arbitrary surfaces

Our approach is motivated by simplicial homology, as follows. Let a finite connected graph G be embedded in a closed surface S of minimum Euler genus $\varepsilon := 2 - \chi(S)$. Then S can be viewed as the underlying space of a 2-dimensional CW-complex C with 1-skeleton G. Its first homology group $Z_1(C;\mathbb{Z}_2)/B_1(C;\mathbb{Z}_2)$ is $\mathbb{Z}_2^\varepsilon$, the direct product of ε copies of $\mathbb{Z}_2$.

In graph theoretic language this means that the subspace $\mathcal{B}$ $(= B_1(C;\mathbb{Z}_2))$ spanned in $\mathcal{C}(G)$ $(= Z_1(C;\mathbb{Z}_2))$ by the set of face boundaries of G in S has codimension ε in $\mathcal{C}(G)$. Now the set of face boundaries is a sparse set of cycles. Thus, if G embeds in a surface of small Euler genus, at most ε, then G has a sparse set of cycles spanning a large subspace in $\mathcal{C}(G)$, one of codimension at most ε.

MacLane's theorem says that, for $\varepsilon = 0$, the converse implication holds too: if G has a sparse set of cycles whose span in $\mathcal{C}(G)$ has codimension at most $\varepsilon = 0$, then G embeds in the (unique) surface of Euler genus at most $\varepsilon = 0$, the sphere. Our initial aim, then, would be to prove this converse implication for arbitrary ε.

This naive extension soon runs into difficulties, and indeed is not true. In Section 5.4 we discuss the obstructions encountered as they arise, and modify our naive conjecture accordingly. The result will be a collection of theorems, presented in Section 5.5, which each characterise embeddability in a given surface, or in a surface of given Euler genus, by a condition akin to MacLane's planarity criterion that is both necessary and sufficient. All proofs are given in Section 5.6.

Some previous work in this direction can be found in the literature. Lefschetz [57] characterises the graphs that are embeddable in a given surface so that every face is bounded by a cycle. His theorem for orientable surfaces will follow from Theorem 5.7 (i). Lefschetz's theorem for non-orientable surfaces, stated in [57] without a formal proof, is incorrect; our Theorem 5.7 (ii) corrects and strengthens his result. Mohar [60] starts out from the necessary condition discussed earlier for embeddability in a surface of Euler genus at most ε, namely, that the graph must have a sparse set of cycles whose span in its cycle space has codimension at most ε. Unlike our plan here, Mohar does not strengthen this condition to one that is also sufficient, but establishes how much it implies as it is; the (best possible) result is that it implies embeddability in a surface of Euler genus at most 2ε. Širáň and Škoviera [89, 90] investigate when a given family of closed walks in a graph G can appear as face boundaries in an embedding of G in some surface, not necessarily of small genus (as will be our aim). Their work extends our discussions in Section 5.4 and provides an interesting backdrop for our proofs in Section 5.6, some of which use techniques they developed. We shall also use techniques of Edmonds [37], who studies embeddability in arbitrary surfaces in terms of duality.

5.3 General definitions and background

All graphs we consider in this chapter are finite, unless otherwise noted. In the statements of some of our results we do not allow loops, but only to avoid unnecessary complication in our terminlogy: those theorems can be applied to graphs with loops by subdividing (and thereby eliminating) these.

The set of edges of a graph $G = (V, E)$ incident with a given vertex v is denoted by $E(v)$. When W is a walk in G, we denote the subgraph of G that consists of the edges on W and their incident vertices by $G[W]$; note that this need not be an induced subgraph of G. The (*unoriented*) *edge space* $\mathcal{E}(G)$ of G is the $\mathbb{Z}_2$ vector space of all functions $E \to \mathbb{Z}_2$ under pointwise addition. We usually write these as subsets of E, so vector addition becomes symmetric difference of edge sets. As before, the (*unoriented*) *cycle space* $\mathcal{C}(G)$ of G is the subspace of $\mathcal{E}(G)$ generated by *circuits*, the edge sets of cycles.

A triple (e, u, v) consisting of an edge $e = uv$ together with its ends listed in a specific order is an *oriented edge*. The two oriented edges corresponding to e are its two *orientations*, denoted by $\vec{e}$ and $\overleftarrow{e}$. Thus, $\{\vec{e}, \overleftarrow{e}\} = \{(e, u, v), (e, v, u)\}$, but we cannot generally say which is which. Given a set E of edges, we write $\vec{E}$ for the set of their orientations, two for every edge in E.

The *oriented edge space* $\vec{\mathcal{E}}(G)$ of $G = (V, E)$ is the real vector space of all functions $\phi\colon \vec{E} \to \mathbb{R}$ satisfying $\phi(\overleftarrow{e}) = -\phi(\vec{e})$ for all $\vec{e} \in \vec{E}$. When $v_0 \ldots v_{k-1} v_0$ is a cycle and $e_i := v_i v_{i+1}$ (with $v_k := v_0$), the function mapping the oriented edges (e_i, v_i, v_{i+1}) to 1, their inverses (e_i, v_{i+1}, v_i) to -1, and every other oriented edge to 0, is an *oriented circuit*. The *oriented cycle space* $\vec{\mathcal{C}}(G)$ is the subspace of $\vec{\mathcal{E}}(G)$ generated by the oriented circuits.

If G is connected and has n vertices and m edges, its oriented and its unoriented cycle space both have dimension

$$\dim \mathcal{C}(G) = \dim \vec{\mathcal{C}}(G) = m - n + 1. \tag{5.1}$$

A (closed) *surface* is a compact connected 2-manifold without boundary. It is *orientable* if it admits a triangulation whose 2-simplices (triangles) can be compatibly oriented. Equivalent conditions are that every triangulation has this property, and that the surface does not contain a Möbius strip [7].

An *n-dimensional CW-complex*, or *n-complex*, is a finite set C of open balls $B_j^i \subseteq \mathbb{R}^i$ with $i \leq n$, called *i-cells*, that have disjoint closures and whose union is made into a topological space $|C|$ as follows. The union C^0 of all 0-cells (which are singletons, so C^0 is just a set of points) carries the discrete topology. Assume now that the union of all i-cells with $i \leq k < n$, the *k-skeleton* C^k of C, has been given a topology, and denote this space by $|C^k|$. For every $(k+1)$-cell $B_j^{k+1} \in C$ choose a continuous *attachment map* $f_j : \partial B_j^{k+1} \to |C^k|$ from its boundary ∂B_j^{k+1} in $\mathbb{R}^{k+1}$ to $|C^k|$. Then give $|C^{k+1}|$ the quotient topology of the (disjoint) union of $|C^k|$ with all the closures of the B_j^{k+1} obtained by identifying every $x \in \partial B_j^{k+1}$ with $f_j(x)$.

Every graph G is a 1-complex, with vertices as 0-cells and edges as 1-cells. A topological embedding of G in another space S is a *2-cell-embedding* if G is the 1-skeleton of a 2-complex C such that the embedding of G in S extends to a homeomorphism $\varphi \colon |C| \to S$. The images under φ of the 2-cells of C are the *faces* of G in S. If S is a surface, their attachment maps define closed walks in G. These walks are unique up to cyclic shifts and orientation, a difference we shall often ignore. We thus have one such walk (with two orientations) assigned to each face, and call this family the (unique) *family of facial walks*. If W is the facial walk of some face f, then φ maps the subgraph $G[W]$ onto the frontier of f in S, and we call $G[W]$ the *boundary* of the face f.

Given a surface S, consider any 2-cell-embedding of any graph in S. Let n be its number of vertices, m its number of edges, and ℓ its number of faces in S. Euler's theorem tells us that $n - m + \ell$ is equal to a constant $\chi(S)$ depending only on S (not on the graph), the Euler characteristic of S. The *Euler genus* $\varepsilon(S)$ of S is defined as the number $2 - \chi(S)$. Euler's theorem then takes the following form, which we refer to as *Euler's formula*:

$$\varepsilon(S) = m - n - \ell + 2. \tag{5.2}$$

Given a graph G, let $\varepsilon = \varepsilon(G)$ be minimum such that G has a topological embedding φ in a surface of Euler genus at most ε. This ε is the *Euler genus of* G, and any such φ is a *genus-embedding* of G. Every connected graph has a genus-embedding that is a 2-cell-embedding [61, p. 95]. If G has components $G_1, \dots, G_n$, then $\varepsilon(G) = \varepsilon(G_1)+\dots+\varepsilon(G_n)$, a fact referred to as *genus additivity* [61]. (The same is true for blocks rather than components, but we do not need this.)

We say that a family $\mathcal{W}$ of walks *covers* a subgraph H of G (often given in terms of its edge set) if every edge of H lies on some walk of $\mathcal{W}$. It covers an edge e *k times* if $k = \sum_{W\in\mathcal{W}} k_W(e)$, where $k_W(e)$ is the number of occurrences of e on W (irrespective of the direction in which W traverses e). $\mathcal{W}$ is a *double cover* of G if it covers every edge of G exactly twice. A walk is *non-trivial* if it contains an edge.

Given a walk W in G, we write $c(W)\colon E(G) \to \mathbb{Z}_2$ for the function that assigns to every edge e the number of times that W traverses e (in either direction), taken mod 2. Informally, we think of $c(W)$ as its support, the set of edges that appear an odd number of times in W. The *dimension* of a family $\mathcal{W}$ of walks, $\dim \mathcal{W}$, is the dimension of the subspace spanned in $\mathcal{E}(G)$ by the functions (or sets) $c(W)$ with $W \in \mathcal{W}$. If the walks are closed, their $c(W)$ lie in $\mathcal{C}(G)$; then the *codimension* of $\mathcal{W}$ in $\mathcal{C}(G)$ is the number $\dim \mathcal{C}(G) - \dim \mathcal{W}$.

Taking the natural orientation of W into account, we write $\vec{c}(W)$ for the function that assigns to every $\vec{e} \in \vec{E}$ the number of times that W traverses e in the direction of $\vec{e}$ minus the number of times that W traverses e in the direction of $\overleftarrow{e}$, and assigns 0 to any $\vec{e}$ with e not on W. The *oriented dimension* of a family $\mathcal{W}$ of walks, $\overrightarrow{\dim}\, \mathcal{W}$, is the dimension of the subspace of $\vec{\mathcal{E}}(G)$ spanned by the functions $\vec{c}(W)$ with $W \in \mathcal{W}$. If the walks are closed, their $\vec{c}(W)$ lie in $\vec{\mathcal{C}}(G)$; then the *codimension* of $\mathcal{W}$ in $\vec{\mathcal{C}}(G)$ is the number $\dim \vec{\mathcal{C}}(G) - \overrightarrow{\dim}\, \mathcal{W}$.

5.4 Reconstructing a surface

MacLane's theorem offers a necessary and sufficient condition for embeddability in a fixed surface, the sphere. Our aim is to find a similar condition characterising embeddability in an arbitrary but fixed surface S.

To illustrate what we mean by 'similar', let us think of MacLane's theorem as listing some properties of the facial cycles of a plane graph—sparseness and generating the entire cycle space—which, together, imply the following: that whenever we have *any* collection of cycles with these properties and attach a 2-cell to each of them, the 2-complex obtained is homeomorphic to the sphere. (This, indeed, is the outline of the standard topological proof of MacLane's theorem.)

For an arbitrary surface S, we are thus looking for a similar list of properties shared by the facial cycles of all graphs suitably embedded in S (with a genus-embedding, say) such that, given any family of cycles with these properties in a graph G, attaching a 2-cell along every cycle in this family turns G into a copy of S. One of those properties should be sparseness: if more than two 2-cells meet in an edge, the complex obtained will not be a surface. Following the homological approach outlined in the introduction, we might complement this by requiring that our cycles span a large enough subspace of the cycle space of G:

Naive Conjecture. *A graph G embeds in a surface S if and only if G has a sparse set of cycles whose span in $\mathcal{C}(G)$ has codimension at most $\varepsilon(S)$ in $\mathcal{C}(G)$.*

Notice that this conjecture can be true only if embeddability in a surface S depends only on $\varepsilon(S)$. For $\varepsilon = 0$ this is not an issue, since the sphere is the only surface with $\varepsilon = 0$. For even $\varepsilon > 0$, however, there are two surfaces of Euler genus ε—one orientable and one non-orientable—and the corresponding classes of graphs embeddable in them do not coincide. (Indeed, large projective-planar grids have unbounded orientable genus [8], while K_7 can be embedded in the torus but not in the Klein bottle [44].) Our first aim, therefore, will be to characterise embeddability not in a given surface S, but in 'some' surface of given Euler genus—in other words, to characterise the graphs of given Euler genus.

Another flaw in the Naive Conjecture is its reference to cycles: for surfaces other than the sphere, even genus-embeddings of 2-connected graphs can have facial walks that are not cycles. (For example, we can embed the graph G of Figure 5.2 in the torus by running the edge $e = uv$ along a handle added to the sphere to join two triangular faces containing u and v, respectively. Then e lies on the boundary of only one face, whose facial walk contains it twice and therefore is not a cycle. Zha [96] constructed for every surface S other than the sphere and the projective plane a 2-connected graph that has a genus-embedding in S but no embedding whose facial walks are all cycles.)

With these two modifications, our conjecture might become the following:

Revised Conjecture. *For every integer $\varepsilon \geq 0$, a graph G embeds in a surface of Euler genus at most ε if and only if it has a family of closed walks that covers every edge at most twice and whose codimension in $\mathcal{C}(G)$ is at most ε.*

However, as noticed by various authors [57, 60, 90], this is still not true: our list of properties of facial cycles—so far, sparseness and large dimension—needs a further addition.

To illustrate this, consider the plane graph A_1 shown in Figure 5.1. Let G be obtained from A_1 by identifying the vertices u and v. This graph G is one of the 35 forbidden minors that characterise embeddability in the projective plane (Archdeacon [5]), so $\varepsilon(G) \geq 2$.

Let $\mathcal{W}$ denote the family of facial walks of A_1. The subspace it spans in $\mathcal{C}(G)$ is the cycle space of A_1. By (5.1), and since G has one vertex less than A_1 but the same number of edges, we deduce that

$$\dim \mathcal{W} = \dim \mathcal{C}(A_1) = \dim \mathcal{C}(G) - 1.$$

By the Revised Conjecture for $\varepsilon = 1$, this implies that G can be embedded in the projective plane—which it cannot.

To see what went wrong, let us form the 2-complex obtained by pasting a 2-cell on every walk in $\mathcal{W}$: the complex that 'should' be the projective plane but is not. The solution to the paradox is that this complex is not a surface at all: it is the pseudosurface obtained from a sphere by identifying two points.

To rule out this type of counterexample we could require that, for every vertex v, no proper subset of those of our given walks that pass through v can combine to a flat neighbourhood of v when we attach 2-cells to these walks. Since the facial walks in any 2-cell embedding of a graph have this property, it would certainly be an acceptable addition to our list. (In MacLane's theorem no such requirement is needed, because it follows; we shall prove this after stating Theorem 5.3 below.) If $\mathcal{W}$ is a double cover of G as in the example, or at least a cover, this flatness condition is sufficient and, indeed, the additional requirement we shall impose in Section 5.5 will then reduce to this. However if $\mathcal{W}$ does not cover G we need to be yet more careful. Our next example shows why.

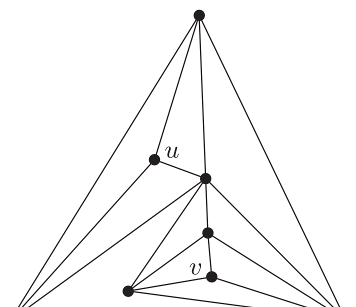

Figure 5.1: Identify u and v to obtain a graph G with $\varepsilon(G) \geq 2$

Consider the plane graph A_2 shown in solid lines in Figure 5.2. Let G be obtained from A_2 by adding the edge uv. This graph G is another of Archdeacon's 35 forbidden minors for the projective plane, so again $\varepsilon(G) \geq 2$.

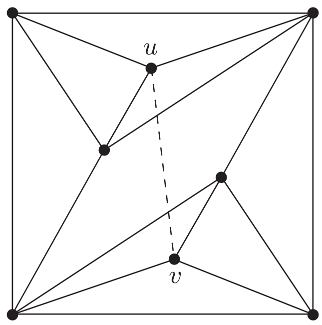

Figure 5.2: Add the edge uv to obtain a graph G with $\varepsilon(G) \geq 2$

As before, the subspace $\mathcal{W}$ spanned in $\mathcal{C}(G)$ by the facial walks of A_2 is the cycle space of A_2. By (5.1), and since G has one more edge than A_2 but the same number of vertices, we deduce that

$$\dim \mathcal{W} = \dim \mathcal{C}(A_2) = \dim \mathcal{C}(G) - 1.$$

By the Revised Conjecture for $\varepsilon = 1$, this implies that G can be embedded in the projective plane—which it cannot.

Again let us see what goes wrong in the 2-complex C formed by attaching a 2-cell to each walk in $\mathcal{W}$. While we obtain from the walks in $\mathcal{W}$ a flat neighbourhood around each vertex, it is impossible to extend $\mathcal{W}$ so as to include the edge uv without producing a non-flat neighbourhood at u and at v: the only way to do this is to add a walk just along uv and back, and pasting a disc onto this walk will add a second sphere, touching the sphere of $|C|$ in u and v.

To rule out counterexamples such as this, we could simply require that our family $\mathcal{W}$ should cover all the edges of G. This would be an acceptable addition to our list of requirements on $\mathcal{W}$ in that every collection of facial walks of an embedded graph satisfies it. However it would be against the spirit of MacLane's theorem—that marrying a purely graph-theoretical sparseness condition on $\mathcal{W}$ to a purely algebraic richness condition can yield a characterisation of planarity. Requiring that $\mathcal{W}$ cover the edges of G would spoil this dichotomy by adding a graph-theoretical richness condition.

Our solution to this dilemma will be to strengthen the requirement of 'sparseness at vertices' on $\mathcal{W}$, discussed after our first example, as follows: we shall require that no subfamily $\mathcal{U}$ of $\mathcal{W}$ (plus discs) shall form a flat neighbourhood of a vertex v unless it covers all the edges at v, regardless of whether they lie on walks of $\mathcal{W} \setminus \mathcal{U}$ or not. This condition will rule out both the above examples, while the graph-theoretical richness requirement it entails—that $\mathcal{U}$ must cover $E(v)$—will at least be kept local.

There is one more twist. Although, as we shall prove, our sparseness condition is now strong enough to ensure that graphs with a sparse family of closed walks of codimension at most ε embed in a surface of Euler genus at most ε, it is not true that pasting a disc on each of those walks will yield such as surface. For example, consider in a graph drawn on the sphere two vertices that lie on a common face boundary W. Identifying these two vertices into a new vertex v turns the sphere into a pseudosurface S on which the old facial walks still bound discs, so attaching discs to the walks after identification yields this pseudosurface. But those facial walks still form a sparse family: any non-empty subfamily summing to zero at v must contain W, but then it contains edges from both of the 'two disjoint disc neighbourhoods' of v on S and hence contains all the facial walks through v and thus covers $E(v)$.

Fortunately, there is a way to dissolve such singularities: rather than pasting a disc on W along its original orientation (which will result in a pseudosurface), we change the orientation of half of W, reversing it as shown in Figure 5.3 on one of its two segments between its two visits to v. (Recall that, after identification, W passes through v twice.) As the reader may verify, this alteration dissolves the singularity of our pseudosurface, turning it into a projective plane. In general, we shall prove that all singularities that can arise from pasting discs on a sparse family of closed walks can be dissolved in this way.

The ideas discussed so far will enable us to characterise, for any given ε, the graphs embeddable in either the orientable or the non-orientable surface of Euler genus ε (Theorem 5.3). To distinguish between the two, we shall have to refine our sparseness condition at vertices once more, and make use of the oriented cycle space. The key observation is that, given a 2-cell embedding of a graph G in a surface S, the facial walks—suitably oriented—will sum to zero in $\vec{\mathcal{C}}(G)$ when S is orientable, but will never sum to zero if S is non-orientable. With this observation suitably implemented, we shall finally be able to derive our desired MacLane-type characterisation of the graphs embeddable in a given surface, Theorem 5.7.

5.5 Statement of results

Recall that a family $\mathcal{F}$ of subsets of $E(G)$ is *sparse* if every edge of G lies in at most two members of $\mathcal{F}$. Similarly, we shall call a family of walks *sparse at* an edge e if it covers e at most twice. In view of our discussion in Section 5.4, we now wish to supplement this by a sparseness requirement at vertices.

Given a family $\mathcal{W}$ of walks and a vertex v, let us call a non-empty subfamily $\mathcal{U}$ of the walks in $\mathcal{W}$ through v a *cluster at* v if $\sum_{W \in \mathcal{U}} c(W) \cap E(v) = \emptyset$ but $\mathcal{U}$ fails to cover $E(v)$. We say that $\mathcal{W}$ is *sparse* if it is sparse at all edges and does not have a cluster at any vertex. For families of edge sets rather than walks we retain our earlier notion of sparseness, meaning sparseness at edges.

We can now state our first extension of MacLane's theorem. It can be read as a characterisation of the graphs of given Euler genus:

Theorem 5.3. **[17]** *For every integer $\varepsilon \geq 0$, a graph G can be embedded in some surface of Euler genus at most ε if and only if there is a sparse family of closed walks in G whose codimension in $\mathcal{C}(G)$ is at most ε.*

For $\varepsilon = 0$, Theorem 5.3 implies MacLane's theorem. This is not immediately obvious: one has to show that a sparse family $\mathcal{B}$ of edge sets of cycles generating $\mathcal{C}(G)$ (as in MacLane's theorem) must be sparse also as a family of walks, i.e., that it does not have any clusters. We may assume that G is 2-connected. Suppose that $\mathcal{B}$ has a cluster at a vertex v. Thus, there is a non-empty subfamily $\mathcal{F}$ of $\mathcal{B}$ whose edges at v sum to zero but which fails to cover some other edge vw at v. Choose $\mathcal{F}$ minimal, and pick an edge uv from a cycle in $\mathcal{F}$. As G is 2-connected, $G - v$ contains a u–w path P; then $C = uPwvu$ is a cycle. We claim that no set $\mathcal{B}' \subseteq \mathcal{B}$ can sum to C, contradicting the choice of $\mathcal{B}$. Indeed, since $\mathcal{B}$ is sparse and $\mathcal{F}$ sums to zero at v, every edge in $D := E(v) \cap \bigcup \mathcal{F}$ lies on exactly two cycles in $\mathcal{F}$ but not on any cycle in $\mathcal{B} \setminus \mathcal{F}$. The set of cycles in $\mathcal{B}'$ with an edge in D, therefore, is precisely $\mathcal{B}' \cap \mathcal{F}$. In particular, if $uv \in \sum \mathcal{B}'$ then $uv \in E' := \sum(\mathcal{B}' \cap \mathcal{F})$. Since every cycle in $\mathcal{B}' \cap \mathcal{F}$ has two edges in D, we know that $|E' \cap D|$ is even. Hence if $uv \in \sum \mathcal{B}'$, there must be another edge $e \neq uv$ in $E' \cap D = (\sum \mathcal{B}') \cap D$. This edge cannot be $vw \notin D$, so it does no lie on C. Thus, $\sum \mathcal{B}'$ differs from C either in uv or in e, i.e. $\sum \mathcal{B}' \neq C$ as claimed.

The forward implication of Theorem 5.3 is well known, and its proof will not be hard. In our proof of the backward implication we shall take a detour via 'locally sparse' families of walks, which we define next. (We shall also need this concept again to state and prove our second main result, Theorem 5.7 below.) In order to keep our terminology simple we shall now ban loops; this will be easy to undo when we later prove Theorem 5.3.

Let $W = v_1 e_1 \dots v_n e_n v_1$ be a closed walk in a loopless graph G, where the v_i are vertices and the e_i are edges. For a vertex v we call a subsequence $e_{j-1} v_j e_j$ of W with $v_j = v$ (where $e_0 := e_n$) a *pass* of W *through* the vertex v.

Extending our earlier notation for walks, we write $c(e_{j-1}v_je_j) := \{e_{j-1}, e_j\}$ if $e_{j-1} \neq e_j$, and $c(e_{j-1}v_je_j) := \emptyset$ if $e_{j-1} = e_j$. In order keep track of how often a given walk passes through a given vertex, we shall consider the *family of all passes of W through v*, the family $(e_{j-1}v_je_j)_{j\in J}$ where $J = \{j : v_j = v, 1 \leq j \leq n\}$. Similarly, if $\mathcal{W} = (W_i)_{i\in I}$ is a family of walks then the *family of all passes of $\mathcal{W}$ through v* is the family $\mathcal{A}(\mathcal{W}, v) := (p_{ij})_{i\in I, j\in J_i}$ where, for each i, $(p_{ij})_{j\in J_i}$ is the family of all passes of W_i through v. Let us call a non-empty subfamily $\mathcal{F} \subseteq \mathcal{A}(\mathcal{W}, v)$ a *local cluster at v* if $\sum_{p\in\mathcal{F}} c(p) = \emptyset$ but $\mathcal{F}$ fails to cover $E(v)$. We say that $\mathcal{W}$ is *locally sparse* if $\mathcal{W}$ is sparse at all edges and has no local cluster at any vertex. Note that any locally sparse family of closed walks in G is sparse, since for every vertex v and every closed walk W we have $c(W) \cap E(v) = \sum_{p\in\mathcal{A}((W),v)} c(p)$.

The following equivalence, whose implication (ii)→(i) will be a lemma in our proof of the backward implication of Theorem 5.3, is weaker than that implication in that it requires local sparseness rather than just sparseness in (ii). But it is also stronger, in that it allows us to make our *given* walks into face boundaries.

Lemma 5.4. [17] *Let $G = (V, E)$ be a loopless connected graph, $\mathcal{W}$ a family of closed walks in G, and $\varepsilon \geq 0$ an integer. Then the following two statements are equivalent:*

(i) *There is a surface S of Euler genus at most ε in which G can be 2-cell-embedded so that $\mathcal{W}$ is a subfamily of the family of facial walks.*

(ii) *There is a locally sparse family of closed walks in G that has codimension at most ε in $\mathcal{C}(G)$ and includes $\mathcal{W}$.*

In order to make Lemma 5.4 usable for the proof of Theorem 5.3, we next have to address the task of turning a sparse family $\mathcal{W}$ of closed walks into a locally sparse family $\mathcal{W}'$ without changing its codimension in $\mathcal{C}(G)$. In fact, we shall be able to do much more: we shall obtain $\mathcal{W}'$ from $\mathcal{W}$ by merely changing the order in which a walk traverses its edges. This is not unremarkable: it means, for example, that by merely changing the order in which the offending boundary walk W in the example discussed at the end of Section 5.4 traverses its edges we can turn the resulting pseudosurface into a surface.

To do this formally, consider any family $\mathcal{W}$ of closed walks in G. Call a family $\mathcal{W}' = (\, W' : W \in \mathcal{W}\,)$ of closed walks *similar* to $\mathcal{W}$ if, for every $e \in E(G)$ and every $W \in \mathcal{W}$, the edge e occurs on W' as often as it does on W. Thus if $\mathcal{W}'$ is similar to $\mathcal{W}$ then $G[W'] = G[W]$ and $c(W') = c(W)$ for every $W \in \mathcal{W}$, and in particular $\dim \mathcal{W}' = \dim \mathcal{W}$. Note that although a family similar to a locally sparse family need not itself be locally sparse (which indeed is our reason for defining similarity), a family similar to a sparse family will always be sparse.

Our next step, then, will be to prove the following lemma:

Lemma 5.5. [**17**] *For every sparse family $\mathcal{W}$ of closed walks in a connected loopless graph G there exists a locally sparse family $\mathcal{W}'$ similar to $\mathcal{W}$.*

Using Lemmas 5.4 and 5.5 it will be easy to prove the following equivalence, a more explicit version of Theorem 5.3:

Theorem 5.6. [**17**] *Let G be a connected graph, $\mathcal{W}$ a family of closed walks in G, and $\varepsilon \geq 0$ an integer. Then the following statements are equivalent:*

(i) *There is a surface of Euler genus at most ε in which G can be 2-cell-embedded so that the family of facial walks has a subfamily similar to $\mathcal{W}$.*

(ii) *There is a sparse family of closed walks in G that has codimension at most ε in $\mathcal{C}(G)$ and includes $\mathcal{W}$.*

While Theorem 5.3 characterises the graphs of given Euler genus, our initial aim was to characterise the graphs embeddable in a given surface S. This will be achieved by the following theorem, which is our main result.

Theorem 5.7. [**17**] *Let S be any surface, and let ε denote its Euler genus. Let G be any loopless graph, and let k denote the number of its components.*

(i) *If S is orientable, then G can be embedded in S if and only if G has a double cover by a locally sparse family $\mathcal{W}$ of closed walks whose oriented dimension is at most $|\mathcal{W}| - k$ and which has codimension at most ε in $\vec{\mathcal{C}}(G)$.*

(ii) *If G is connected and S is not orientable, then G can be embedded in S if and only if there is a sparse family $\mathcal{W}$ of closed walks in G whose codimension in $\vec{\mathcal{C}}(G)$ is at most $\varepsilon - 1$.*

We conjecture that 'locally sparse' cannot be replaced by 'sparse' in (i). And we remark that the connectivity requirement in (ii) cannot be dropped. Indeed, consider a graph G consisting of k disjoint copies of a graph that can be embedded in the projective plane but not in the sphere. By (ii), G can be covered by a sparse family of closed walks that has codimension 0 in $\vec{\mathcal{C}}(G)$. However, G cannot be embedded in any surface of Euler genus less than k.

5.6 The proofs

Let $\mathcal{W}$ be a family of closed walks in a loopless graph G that is sparse at edges. Recall that, for each vertex $v \in G$, we denoted by $\mathcal{A}(\mathcal{W}, v)$ the family of all passes of $\mathcal{W}$ through v. As a tool for our proofs, let us define for every vertex v an auxiliary graph $H = H(\mathcal{W}, v)$ with vertex set $\mathcal{A}(\mathcal{W}, v)$. Its edge set will be a subset of $E(G)$, with incidences defined as follows. Whenever two distinct vertices p, q of H (i.e., passes that are distinct as family members—they may be equal as triples) share an edge $e \in G$, we let e be an edge of H joining p and q. If $\mathcal{W}$ contains a pass $p = eve$, we let e be a loop at p. Clearly, H has maximum degree at most 2, since a pass evf can be incident only with the edges e and f. (For example, if there are three edges e, f, g at v in G, and $\mathcal{W}$ contains the passes evf, fvg, gve, then these three passes and the three edges e, f, g form a triangle in H. As another example, if $\mathcal{W}$ has two passes consisting of the triple evf, or one pass evf and another pass fve, then these two passes are joined by the pair $\{e, f\}$ of double edges in H and have no other incident edge.) If $\mathcal{W}$ is a double cover of G, then every $H(\mathcal{W}, v)$ is 2-regular.

Note that if $\mathcal{W}$ covers $E(v)$, then $\mathcal{W}$ has a local cluster at v if and only if $H = H(\mathcal{W}, v)$ contains a non-spanning cycle. Thus, $\mathcal{W}$ is locally sparse if and only if (it is sparse at edges and) each of the graphs $H(\mathcal{W}, v)$ is either a forest—possibly empty—or, if $\mathcal{W}$ covers $E(v)$, a single cycle.

We begin with a lemma which says that sparse double covers by closed walks[1] are nearly independent: that $\dim \mathcal{W} = |\mathcal{W}| - 1$. We shall need this for the family of face boundaries in the proof of (i)→(ii) of Theorem 5.6, and again for arbitrary sparse families in the proof of Theorem 5.7.

Lemma 5.8.[17] *Let $G = (V, E)$ be a connected graph, and let $\mathcal{W}$ be a sparse family of non-trivial walks.*

(i) *For every non-empty subfamily $\mathcal{U}$ of $\mathcal{W}$ that is not a double cover of G, the family $(\, c(U) : U \in \mathcal{U}\,)$ is linearly independent in $\mathcal{C}(G)$.*

(ii) *If $\mathcal{W}$ is a double cover then* $\dim \mathcal{W} = |\mathcal{W}| - 1$.

Proof. It suffices to prove (i), since this implies that $\dim \mathcal{W} \geq |\mathcal{W}| - 1$: then (ii) follows, since $\mathcal{W}$ covers every edge twice and hence $\sum_{W \in \mathcal{W}} c(W) = \emptyset$.

For a proof of (i), let $\mathcal{U}$ be given as stated. Suppose the assertion fails; then $\mathcal{U}$ has a non-empty subfamily $\mathcal{U}' \subseteq \mathcal{U}$ such that $\sum_{U \in \mathcal{U}'} c(U) = \emptyset$. Then any edge covered by $\mathcal{U}'$ is covered by it twice, so as $\mathcal{U}$ is not a double cover

[1] Indeed by any edge sets without clusters: our proof of Lemma 5.8 will not use the fact that $\mathcal{W}$ is a family of walks.

there exists an edge not covered by $\mathcal{U}'$. On the other hand, since $\mathcal{U}'$ is non-empty and its walks are non-trivial, $\mathcal{U}'$ covers some edge of G. Since G is connected, it therefore has a vertex v that is incident both with an edge that is covered by $\mathcal{U}'$ and an edge that is not. Denote by $\mathcal{U}'(v)$ the non-empty family of all walks in $\mathcal{U}'$ containing v. As

$$\sum_{U \in \mathcal{U}'(v)} c(U) \cap E(v) \subseteq \sum_{U \in \mathcal{U}'} c(U) = \emptyset,$$

and as $\mathcal{U}'(v)$ does not cover $E(v)$, $\mathcal{U}'(v)$ is a cluster at v, contradicting that $\mathcal{W}$ is sparse. □

Next, we show that locally sparse families extend to double covers. It is possible to deduce this from results of Širáň and Škoviera [90], but for simplicity we sketch a direct proof.

Lemma 5.9. **[17]** *Let G be a loopless graph and $\mathcal{W}$ a locally sparse family of closed walks in G. Then $\mathcal{W}$ can be extended to a locally sparse double cover $\mathcal{W}'$ of G by closed walks.*

Proof. Let $\mathcal{W}' \supseteq \mathcal{W}$ be a maximal family of closed walks that is locally sparse. We show that $\mathcal{W}'$ is a double cover.

Suppose not. Let F be the set of edges in G not covered twice. Our aim is to find a closed walk W in (V, F) such that $\mathcal{W}'' := \mathcal{W}' \cup \{W\}$ is again locally sparse; this will contradict our maximal choice of $\mathcal{W}'$.

For every vertex v incident with an edge in F, consider the auxiliary graph $H(v) := H(\mathcal{W}', v)$ defined at the start of this section. Let us show that $H(v)$ is a (possibly empty) forest. Suppose not, and let U be the vertex set of a cycle in $H(v)$. Then $\sum_{u \in U} c(u) = 0$. By assumption v is incident with an edge $f \in F$, which thus lies in at most one pass of $\mathcal{W}'$ through v. As this pass has degree at most 1 in $H(v)$ it cannot be in U, which implies that U, as a family of passes, does not cover $E(v)$. Then, however, U is a local cluster at v—a contradiction to our assumption that $\mathcal{W}'$ is locally sparse.

The components of $H(v)$, therefore, are paths. The edges of these paths are precisely the edges at v which $\mathcal{W}'$ covers twice, those in $E(v) \setminus F$. For every such path P put $\partial P := \sum_{p \in V(P)} c(p)$; this is a set of two edges in $F \cap E(v)$, and all these 2-sets are disjoint. Let $C(v)$ be a cycle on $F \cap E(v)$ as its vertex set such that $E(C(v)) \supseteq \{\, \partial P : P \text{ is a component of } H(v) \,\}$. Call the edges in this last set *red*, and the other edges of $C(v)$ *green*. (We allow $C(v)$ to be a loop or to consist of two parallel edges.) Call the number of green edges of $C(v)$ incident with a given vertex f of $C(v)$ the *green degree* of f in $C(v)$.

The green degree in $C(v)$ of an edge $f \in F \cap E(v)$ equals $2-k$, where $k \in \{0, 1\}$ is the number of times that $\mathcal{W}'$ covers f. (5.3)

To construct our additional walk W in (V, F), we start by picking a vertex v_0 of G that is incident with an edge $f_0 \in F$. Then $H(v_0)$ and $C(v_0)$ are defined. Let us construct a maximal walk $W = v_0 f_0 v_1 f_1 \dots f_{n-1} v_n$ in (V, F) such that $f_{i-1} f_i$ is a green edge of $C(v_i)$ and these green edges are distinct for different i. To ensure that we do not use a green edge again, let us delete the green edges as we construct W inductively, $f_{i-1} f_i$ at the time we add $f_{i-1} v_i f_i$ to W. Note, for $i = 1, \dots, n-1$ inductively, that assertion (5.3) still holds for f_{i-1} and f_i at v_i with $W_i := v_0 f_0 \dots f_i v_{i+1}$ added to $\mathcal{W}'$ and the green edges $f_{j-1} f_j$ deleted for all j with $1 \le j \le i$. This implies that when W gets to v_i via f_{i-1}, there is still a green edge $f_{i-1} f$ in $C(v_i)$ at f_{i-1} at that time, so W can continue and leave v_i via $f =: f_i$—unless $v_i = v_0$ and $f = f_0$, for which the extended assertion (5.3) does not hold (and was not proved above). Hence when our construction of W terminates we have $v_n = v_0$, and f_{n-1} is joined to f_0 by a green edge of $C(v_0)$. Thus, W is indeed a closed walk, and $\mathcal{W}'' := \mathcal{W} \cup \{W\}$ is again sparse at edges.

It remains to show that $\mathcal{W}''$ has no local clusters at vertices. The passes of W through a vertex v are all triples evf such that ef is a green edge of $C(v)$. Adding these passes as new vertices to $H(v)$, with adjacencies as defined before, turns $H(v)$ into a graph $H'(v)$ that is either a single cycle containing all of $E(v)$ (if W 'traverses' every green edge of $C(v)$) or a disconnected graph whose components are still paths: $H'(v)$ cannot contain cycles other than a Hamilton cycle, because $C(v)$ is a single cycle. Therefore, as any family $\mathcal{F}$ of passes of $\mathcal{W}''$ through v with $\sum_{p \in \mathcal{F}} c(p) = \emptyset$ induces a cycle in $H'(v)$, this can happen only when $\mathcal{F}$ covers $E(v)$. Thus, $\mathcal{W}''$ is again locally sparse, contradicting the maximal choice of $\mathcal{W}'$. $\square$

We remark that Lemma 5.9 remains true if we replace 'locally sparse' with 'sparse', but we will not need this.

Proof of Lemma 5.4. (i)→(ii) Extend $\mathcal{W}$ to the family $\mathcal{W}'$ of all the facial walks of G in S. Since S is locally homeomorphic to the plane, $\mathcal{W}'$ covers every edge of G twice, and elementary topological arguments show that $\mathcal{W}'$ cannot have a local cluster at any vertex. Hence $\dim \mathcal{W}' = |\mathcal{W}'| - 1$ by Lemma 5.8 (ii). Using (5.1) and Euler's formula (5.2), we deduce that

$$\dim \mathcal{C}(G) - \varepsilon = |E(G)| - |V(G)| + 1 - \varepsilon \le |\mathcal{W}'| - 1 = \dim \mathcal{W}'$$

as desired.

(ii)→(i) Replacing $\mathcal{W}$ with the extension of $\mathcal{W}$ whose existence is asserted in (ii), we may assume that $\mathcal{W}$ itself is locally sparse and has codimension at most ε in $\mathcal{C}(G)$. Extending $\mathcal{W}$ by Lemma 5.9 if necessary, we may further assume that $\mathcal{W}$ is a double cover of G.

Let C be the 2-dimensional CW-complex obtained as follow. We start with G as its 1-skeleton. As the 2-cells we take disjoint open discs $D_W \subseteq \mathbb{R}^2$, one for each walk $W \in \mathcal{W}$, divide the boundary of D_W into as many segments as W is long, and map consecutive segments homeomorphically to consecutive edges in W.

In order for $S := |C|$ to be a surface, we have to check that every point has an open neighbourhood that is homeomorphic to $\mathbb{R}^2$. For points in the interior of 2-cells or edges, this is clear; recall that $\mathcal{W}$ is a double cover. Now consider a vertex v of G. Define $H(v)$ as earlier. Since $\mathcal{W}$ is a double cover, $H(v)$ is now 2-regular, and since $\mathcal{W}$ has no local cluster at v it contains no cycle properly. Hence, $H(v)$ is a cycle. For each pass $p = evf \in V(H(v))$ we let $D(p)$ be a closed disc whose interior lies inside a disc D_W such that p is a pass of W, choosing each $D(p)$ so that its boundary contains v and intersects W in one segment contained in $e \cup f$ and meeting both e and f. These discs $D(p)$ can clearly be chosen with disjoint interiors for different p. Using the elementary fact that the union of two closed discs intersecting in a common segment of their boundaries is again a disc, one easily shows inductively that the interior of the union of all the discs $D(p)$ is an open disc, and hence homeomorphic to $\mathbb{R}^2$. This completes the proof that S is a surface.

Since C is finite, S is compact. Since G is connected, so is S. Finally, Euler's formula (5.2) applied to C, together with (5.1), the trivial inequality of Lemma 5.8 (ii), and our assumption that $\mathcal{W}$ has codimension at most ε in $\mathcal{C}(G)$, yields

$$\begin{aligned}\varepsilon(S) &= 2 - (|V(G)| - |E(G) + |\mathcal{W}|) \\ &= (|E(G)| - |V(G)] + 1) - (|\mathcal{W}| - 1) \\ &= \dim \mathcal{C}(G) - \dim \mathcal{W} \\ &\leq \varepsilon\,.\end{aligned}$$

Thus, (i) is proved. □

We need an easy technical lemma relating $\overrightarrow{\dim}\,\mathcal{W}$ to $\dim \mathcal{W}$.

Lemma 5.10. [17] *Let $G = (V, E)$ be a connected graph, and let $\mathcal{W} = (W_1, \ldots, W_n)$ be a sparse family of non-trivial walks.*

(i) $\overrightarrow{\dim}\,\mathcal{W} \geq \dim \mathcal{W}$.[2]

(ii) *If $\overrightarrow{\dim}\,\mathcal{W} < |\mathcal{W}|$ then there exist $\mu_i \in \{1, -1\}$ such that $\sum_{i=1}^{n} \mu_i \vec{c}(W_i) = 0$.*

[2]This is true regardless of whether $\mathcal{W}$ is sparse. But the special case proved here is all we need.

Proof. Assertion (i) will follow at once from the following claim:

$$\text{If there exist } \lambda_1, \dots, \lambda_n \in \mathbb{R} \setminus \{0\} \text{ such that } \sum_{i=1}^n \lambda_i \vec{c}(W_i) = 0 \text{ in } \vec{\mathcal{E}}(G), \text{ then there are also } \mu_1, \dots, \mu_n \in \{-1, 1\} \text{ such that } \sum_{i=1}^n \mu_i \vec{c}(W_i) = 0. \tag{5.4}$$

Indeed, whenever two walks W_i, W_j share an edge e, we have $|\lambda_i| = |\lambda_j|$ because $\mathcal{W}$ is sparse at e. Let H be the graph on $\{1, \dots, n\}$ in which ij is an edge whenever W_i and W_j share an edge. Then the values of $|\lambda_i|$ coincide for all i in a common component C of H, and letting $\mu_j := \lambda_j/\lambda_i$ for some fixed i and all j in C satisfies (5.4).

Let us now prove (ii). If $\vec{\dim}\, \mathcal{W} < |\mathcal{W}|$ there are $\lambda_1, \dots, \lambda_n \in \mathbb{R}$ not all zero such that $\sum_{i=1}^n \lambda_i \vec{c}(W_i) = 0$. By (5.4), we may assume the λ_i to be in $\{1, 0, -1\}$. Applying Lemma 5.8 (i) to the subfamily $\mathcal{U}$ of the W_i with $\lambda_i \neq 0$ we see that the λ_i are in fact all non-zero, as desired. $\square$

Next, let us prove Lemma 5.5, our tool for turning a sparse family of walks into a locally sparse one without changing the edge sets of its walks. The proof employs a trick from surface surgery to dissolve singularities, which we learnt from Edmonds [37]. In fact, we prove a slightly stronger statement:

Lemma 5.11. **[17]** *For every sparse family $\mathcal{W}$ of closed walks in a connected loopless graph G there exists a locally sparse family $\mathcal{W}'$ similar to $\mathcal{W}$. If $\mathcal{W}$ is not locally sparse, then $\mathcal{W}'$ can be chosen so that $\vec{\dim}\, \mathcal{W}' = |\mathcal{W}'|$.*

Proof. For families $\mathcal{W}'$ of closed walks, define $\gamma(\mathcal{W}') := \sum_{v \in V(G)} \gamma_{\mathcal{W}'}(v)$ where $\gamma_{\mathcal{W}'}(v)$ denotes the number of components of $H(\mathcal{W}', v)$. Assuming that $\mathcal{W}$ is not locally sparse, we will construct a family $\mathcal{W}'$ similar to $\mathcal{W}$ such that $\gamma(\mathcal{W}') < \gamma(\mathcal{W})$; we will further ensure that $\vec{\dim}\, \mathcal{W}' = |\mathcal{W}'|$. Since $\gamma(\mathcal{W})$ is bounded below by 0, this will prove the lemma.

Let us construct $\mathcal{W}'$. As $\mathcal{W}$ is not locally sparse, there must exist a local cluster at some vertex v. Seen in $H := H(\mathcal{W}, v)$ this local cluster forms a cycle. Since $\mathcal{W}$ is sparse, one of the vertices of C must be a pass $p = eve'$ of a walk $W \in \mathcal{W}$ which also contains a pass $q = fvf'$ that is a vertex in another component $D \neq C$ of H. Choose these passes so that W has a subwalk $ve' \dots fv$ not containing e or f'. Let W' be the closed walk obtained from W by reversing this subwalk (Figure 5.3), and let $\mathcal{W}'$ be obtained from $\mathcal{W}$ by replacing W with W'. Clearly, W' is again a closed walk, and $\mathcal{W}'$ is similar to $\mathcal{W}$.

Let us show that $\gamma(\mathcal{W}') < \gamma(\mathcal{W})$. For vertices $u \neq v$ of G we have $H(\mathcal{W}', u) = H(\mathcal{W}, u)$, so $\gamma_{\mathcal{W}'}(u) = \gamma_{\mathcal{W}}(u)$. At v, however, we have $\gamma_{\mathcal{W}'}(v) < \gamma_{\mathcal{W}}(v)$, so $\gamma(\mathcal{W}') < \gamma(\mathcal{W})$. Indeed, $H' := H(\mathcal{W}', v)$ arises from H by the replacement of $p = eve' \in V(C)$ and $q = fvf' \in V(D)$ with two new

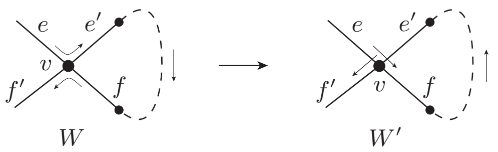

Figure 5.3: Turning W into W' by reversing the segment $ve' \dots fv$

vertices, $p' := evf$ and $q' := e'vf'$, and redefining the incidences for the edges $e, f, e', f' \in E(H) = E(H')$ accordingly. As one easily checks (see Figure 5.4), this has the effect of merging the components C and D of H into one new component, leaving the other components of H intact. Thus, the components of H' are those of H other than C and D, plus one new component arising from $(C - p) \cup (D - q)$ by adding the new vertex p' incident with e and f and the new vertex q' incident with e' and f' (leaving the other incidences of e, e', f, f' in H' as they were in H).

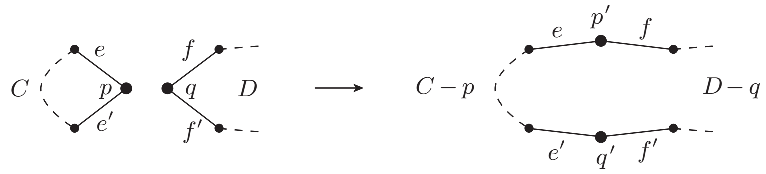

Figure 5.4: Merging the components C and D of H to form H'

It remains to show that $\overrightarrow{\dim}\, \mathcal{W}' = |\mathcal{W}'|$. First note that, if $C = e_1 \dots e_m$ where $e = e_1$ and $e' = e_m$ then $fe_1 \dots e_m f'$ is a subpath of C', the new component that arose from merging C and D.

Suppose now that $\overrightarrow{\dim}\, \mathcal{W}' < |\mathcal{W}'|$. Then for all $U \in \mathcal{W}'$ there are $\mu_U \in \{1, -1\}$ such that $\sum_{U \in \mathcal{W}'} \mu_U \vec{c}(U) = 0$ (Lemma 5.10 (ii)), and we may assume that $\mu_{W'} = 1$. Reversing the orientation of each $U \in \mathcal{W}'$ with $\mu_U = -1$ we obtain $\sum_{U \in \mathcal{W}'} \vec{c}(U) = 0$. Since the orientation of W' has not changed, $p' = e_1 v f$ and $q' = e_m v f'$ are still subwalks of W'. The orientations of the walks in $\mathcal{W}'$ induce orientations on the passes at v; therefore $\sum_{U \in \mathcal{W}'} \vec{c}(U) = 0$ implies that $\sum_{r \in V(C')} \vec{c}(r) = 0$, the passes r being interpreted as subwalks. Hence as $p' = e_1 v f \in V(C')$, each of the passes $e_{i+1} v e_i$ is traversed by some walk in $\mathcal{W}'$ in this order: e_{i+1} towards v, and e_i away from v ($i = 1, \dots, m-1$). In particular, e_m is traversed towards v in the pass $e_m v e_{m-1} \neq q'$. However, this is also the case in q'. As $\mathcal{W}'$ is sparse at edges, this implies $\sum_{r \in V(C')} \vec{c}(r) \neq 0$, a contradiction. $\square$

Proof of Theorem 5.6. Denote by $\dot{G}$ the loopless graph obtained from G by subdividing every loop once. Note that there is an obvious isomorphism

$\mathcal{C}(G) \doteq \mathcal{C}(\dot{G})$, and in particular, the two spaces have the same dimension.

To prove the implication (i)→(ii), consider an embedding of G as in (i). The embedding of G immediately induces an embedding of $\dot{G}$, so that there is a 1-1 correspondence between the facial walks $\dot{\mathcal{U}}$ of the embedding of $\dot{G}$ and the facial walks $\mathcal{U}$ of the embedding of G. Applying Lemma 5.4 to $\dot{\mathcal{U}}$, which is a double cover, we see that $\dot{\mathcal{U}}$ is locally sparse and of codimension $\leq \varepsilon$ in $\mathcal{C}(\dot{G})$. Then the same holds for $\mathcal{U}$ with respect to $\mathcal{C}(G)$. Replacing in $\mathcal{U}$ the subfamily of $\mathcal{U}$ similar to $\mathcal{W}$ with $\mathcal{W}$ preserves both the sparseness of $\mathcal{U}$ and its dimension, so (ii) follows.

For a proof of the implication (ii)→(i), let $\mathcal{W}' \supseteq \mathcal{W}$ be the sparse family of codimension $\leq \varepsilon$ in $\mathcal{C}(G)$ provided by (ii). Then the subdivided walks $\dot{\mathcal{W}}'$ in $\dot{G}$ are still sparse and have codimension $\leq \varepsilon$ in $\mathcal{C}(\dot{G})$. We use Lemma 5.11 to turn $\dot{\mathcal{W}}'$ into a locally sparse family $\dot{\mathcal{W}}''$ similar to $\mathcal{W}'$, which, by Lemma 5.4, is a subfamily of the family $\dot{\mathcal{U}}$ of facial walks of an embedding of $\dot{G}$ in a surface of Euler genus at most ε. If each walk W in $\dot{\mathcal{U}}$ is a subdivision of a walk in G then the embedding of $\dot{G}$ induces one of G in which $\mathcal{W}$ is similar to a subfamily of the facial walks, since $\dot{\mathcal{U}}$ contains $\dot{\mathcal{W}}'' \sim \dot{\mathcal{W}}'$. This can fail only if W contains a pass eve through a subdividing vertex v. If it does, let f be the other edge of $\dot{G}$ at v. Then the subfamily $\mathcal{F} = \{eve\}$ of $\dot{\mathcal{U}}$ satisfies $\sum_{p \in \mathcal{F}} c(p) = 0$, but fails to cover f. Thus the local cluster $\mathcal{F}$ at v contradicts that $\dot{\mathcal{U}}$ is locally sparse. $\square$

Theorem 5.6 immediately implies Theorem 5.3 for connected graphs. To complete the proof of Theorem 5.3, it remains to reduce the disconnected to the connected case.

Proof of Theorem 5.3. For the forward direction, let G and ε be such that G embeds in a surface of Euler genus at most ε. Our aim is to find a certain family of closed walks of codimension at most ε, so there is no loss of generality in choosing ε minimum, i.e., in assuming that $\varepsilon = \varepsilon(G)$. Let $G_1, \ldots, G_n$ be the components of G. For each $i = 1, \ldots, n$ choose a genus-embedding $G_i \hookrightarrow S_i$. These can be chosen to be 2-cell-embeddings, and by genus additivity we have $\varepsilon_1 + \cdots + \varepsilon_n = \varepsilon$ for $\varepsilon_i := \varepsilon(S_i) = \varepsilon(G_i)$. For each i let $\mathcal{W}_i$ be the family of facial walks of G_i in S_i. By Theorem 5.6, the $\mathcal{W}_i$ are sparse and have codimension at most ε_i in $\mathcal{C}(G_i)$: as $\mathcal{W}_i$ already covers every edge of G_i twice, it cannot be extended to a larger sparse family. Since the G_i are vertex-disjoint, $\mathcal{W} := \mathcal{W}_1 \cup \cdots \cup \mathcal{W}_n$ is again sparse, and it has codimension at most $\varepsilon_1 + \cdots + \varepsilon_n = \varepsilon$ in $\mathcal{C}(G)$, since $\mathcal{C}(G)$ is the direct sum of the spaces $\mathcal{C}(G_i)$.

For a proof of the backward direction, let $\mathcal{W}$ be a sparse family of closed walks in G that has codimension at most ε in $\mathcal{C}(G)$. If G has components

$G_1, \ldots, G_k$, say, write $\mathcal{W}_i$ for the subfamily of walks contained in G_i, and ε_i for the codimension of $\mathcal{W}_i$ in $\mathcal{C}(G_i)$. Then $\varepsilon(G_i) \leq \varepsilon_i$, by (ii)→(i) of Theorem 5.6. Moreover, $\sum_{i=1}^{k} \varepsilon_i \leq \varepsilon$, since $\mathcal{C}(G)$ is the direct sum of the spaces $\mathcal{C}(G_i)$. Hence, by genus additivity,

$$\varepsilon(G) = \sum_{i=1}^{k} \varepsilon(G_i) \leq \sum_{i=1}^{k} \varepsilon_i \leq \varepsilon \, .$$

Thus, G can be embedded in a surface of Euler genus at most ε. $\square$

We finally come to the proof of Theorem 5.7. We need another easy lemma.

Lemma 5.12. [**17**] *Let G be a loopless and connected graph. If $\mathcal{W}$ is the family of facial walks of an embedding of G in a surface S, then S is orientable if and only if $\overrightarrow{\dim}\, \mathcal{W} < |\mathcal{W}|$.*

Proof. If $\mathcal{W}$ is the family of facial walks of an embedding of G in S, insert a new vertex in every face and join it to all the vertices on the boundary of that face. This yields a triangulation of S. If S is orientable, we can orient the 2-simplices of this complex C (i.e., the newly created triangles) compatibly, so that every edge receives opposite orientations from the orientations of the two 2-simplices containing it. Then the 2-simplices triangulating a given face induce orientations on the edges of its boundary walk $W \in \mathcal{W}$ that either all coincide with their orientations induced by W or are all opposite to them. Let $\lambda_W := 1$ or $\lambda_W := -1$ accordingly. Then $\sum_{W \in \mathcal{W}} \lambda_W \vec{c}(W) = 0$, showing that $\overrightarrow{\dim}\, \mathcal{W} < |\mathcal{W}|$.

Conversely, if $\overrightarrow{\dim}\, \mathcal{W} < |\mathcal{W}|$ then, by Lemma 5.10 (ii), there are $\mu_W \in \{1, -1\}$, $W \in \mathcal{W}$, so that $\sum_{W \in \mathcal{W}} \mu_W \vec{c}(W) = 0$. Reversing the orientation of every W with $\mu_W = -1$ yields $\sum_{W \in \mathcal{W}} \vec{c}(W) = 0$. These new orientations of the boundary walks W therefore extend to compatible orientations of the 2-simplices of C, showing that S is orientable. $\square$

Proof of Theorem 5.7. (i) We assume that G is connected; the general case then follows as in the proof of Theorem 5.3.[3] Suppose first that G can be embedded in S. Replacing S with a surface of smaller oriented genus if necessary, we may assume that this is a 2-cell embedding. (Any such replacement reduces ε, so this assumption entails no loss of generality.) By Lemma 5.4, the family $\mathcal{W}$ of facial walks is locally sparse and has codimension at most ε in $\mathcal{C}(G)$. Its codimension in $\vec{\mathcal{C}}(G)$ is no greater, since $\overrightarrow{\dim}\, \mathcal{W} \geq$

[3]Use the additivity of oriented genus rather than of Euler genus.

$\dim \mathcal{W}$ by Lemma 5.10 (i), and $\dim \vec{\mathcal{C}}(G) = \dim \mathcal{C}(G)$ by (5.1). It remains to show that $\overrightarrow{\dim}\, \mathcal{W} \leq |\mathcal{W}| - 1$, which follows from Lemma 5.12.

For the converse implication of (i), Lemmas 5.8 (ii) and 5.10 (i) and our assumption about $\overrightarrow{\dim}\, \mathcal{W}$ give

$$\dim \mathcal{W} \leq \overrightarrow{\dim}\, \mathcal{W} \leq |\mathcal{W}| - 1 = \dim \mathcal{W},$$

with equality. By (5.1), then, also the codimension of $\mathcal{W}$ is the same in $\mathcal{C}(G)$ as in $\vec{\mathcal{C}}(G)$, at most ε. By (ii)$\to$(i) of Lemma 5.4, there exists a surface S' with $\varepsilon' := \varepsilon(S') \leq \varepsilon$ in which G has a 2-cell-embedding with $\mathcal{W} =: (W_1, \dots, W_n)$ as the family of facial walks. By Lemma 5.12, S' is orientable. Adding $(\varepsilon - \varepsilon')/2$ handles turns S' into a copy of S with G embedded in it, as desired.

(ii) For the forward implication let $\mathcal{W}$ be the family of facial walks of the given embedding. By Lemma 5.4, $\mathcal{W}$ is sparse. By Lemma 5.12, $\overrightarrow{\dim}\, \mathcal{W} = |\mathcal{W}|$. By (5.1) and (5.2), the codimension of $\mathcal{W}$ in $\vec{\mathcal{C}}(G)$ is $\varepsilon - 1$.

For the backward implication in (ii), let us assume first that the (unoriented) codimension of $\mathcal{W}$ in $\mathcal{C}(G)$ is also at most $\varepsilon - 1$. By Theorem 5.3, we can embed G in a surface S' of Euler genus $\varepsilon' \leq \varepsilon - 1$. The addition of $\varepsilon - \varepsilon' \geq 1$ crosscaps turns S' into a copy of S with G embedded in it.

We may therefore assume that $\mathcal{W}$ has codimension at least ε in $\mathcal{C}(G)$. Let us show that $\mathcal{W}$ is a double cover of G. If not, then Lemmas 5.10 (i) and 5.8 (i) imply

$$|\mathcal{W}| \geq \overrightarrow{\dim}\, \mathcal{W} \geq \dim \mathcal{W} = |\mathcal{W}|$$

with equality, so $\overrightarrow{\dim}\, \mathcal{W} = \dim \mathcal{W}$. By (5.1), this contradicts our assumption that the codimensions of $\mathcal{W}$ in $\mathcal{C}(G)$ and $\vec{\mathcal{C}}(G)$ differ. Moreover, by assumption and Lemma 5.8 we have

$$\dim \mathcal{C}(G) - \varepsilon \geq \dim \mathcal{W} \geq |\mathcal{W}| - 1 \geq \overrightarrow{\dim}\, \mathcal{W} - 1 \geq \dim \vec{\mathcal{C}}(G) - \varepsilon.$$

By (5.1), we have equality throughout. In particular, $\mathcal{W}$ has codimension exactly ε in $\mathcal{C}(G)$, and $\overrightarrow{\dim}\, \mathcal{W} = |\mathcal{W}|$. By Lemma 5.11 there is a locally sparse family $\mathcal{W}'$ similar to $\mathcal{W}$ such that $\overrightarrow{\dim}\, \mathcal{W}' = |\mathcal{W}'|$. Since $\mathcal{W}'$, like $\mathcal{W}$, is a double cover, $\mathcal{W}'$ is by Lemma 5.4 the family of facial walks of an embedding of G in a surface S' of Euler genus $\varepsilon' \leq \varepsilon$. By Lemma 5.12, S' is not orientable. Adding $\varepsilon - \varepsilon'$ crosscaps we turn S' into a copy of S with G embedded in it. $\square$

Chapter 6

Bases and closures under thin sums

6.1 Abstract thin sums

One essential property of the topological cycle space is that, in it, it is necessary to use well-defined infinite sums, the thin sums. In this chapter, which is based on [20], we study two problems about the topological cycle space in which thin sums play an important role. These problems were previously solved by ad-hoc methods appealing to the structure of the underlying graph. We show that the graph theoretic formulation is unnecessary: the problems can be rephrased, and solved, in a purely algebraic setting. Our solutions yield more general results of independent interest. Moreover, our algebraic approach sheds some light on a technique that has appeared in several proofs about the topological cycle space; indeed, in Section 6.4 we obtain a corollary that can be used to simplify those proofs.

Let R be a ring and let M be an arbitrary set. If $\mathcal{N} \subseteq R^M$ is an infinite family for which in each coordinate almost all entries are zero, then there is an obvious way to define its sum, namely by pointwise addition. More formally, we call $\mathcal{N}$ *thin* if for all $m \in M$ the number of members N of $\mathcal{N}$ with $N(m) \neq 0$ is finite. For a thin family $\mathcal{N}$ we define the sum $\sum_{N\in\mathcal{N}} N = \sum \mathcal{N}$ to be the element S of R^M with $S(m) = \sum_{N\in\mathcal{N},\, N(m)\neq 0} N(m)$ for all $m \in M$. We remark that thin families, also called *summable* families, occur in the context of slender modules, see Göbel and Trlifaj [48, Chapter 1].

For a (not necessarily thin) family $\mathcal{N} \subseteq R^M$ denote by $\langle\mathcal{N}\rangle$ the space consisting of all sums of thin subfamilies of $\mathcal{N}$. Our first problem, discussed in Section 6.2, concerns the existence of *bases*: is there a subfamily $\mathcal{B}$ of $\mathcal{N}$ with $\langle\mathcal{N}\rangle = \langle\mathcal{B}\rangle$ such that each element of $\langle\mathcal{N}\rangle$ has a unique representation in $\mathcal{B}$? We will see that even if R is a field there are generating sets that do not contain a basis (Theorem 6.2), although we can always find a basis if the underlying set M is countable (Theorem 6.1). Bases were used as a tool in [23] in order to characterise planar locally finite graphs. There, bases in certain generating sets were needed, and their existence was proved under very restrictive additional requirements. Theorem 6.1 yields these bases immediately.

The second problem we consider is whether the space $\langle\mathcal{N}\rangle$ is closed under taking thin sums, i.e. whether $\langle\mathcal{N}\rangle = \langle\langle\mathcal{N}\rangle\rangle$. While this is false in general, we will show in Section 6.3 that it is true if $\mathcal{N}$ is thin—provided R is a field or a finite ring (Theorem 6.7). We will see that, in a sense, this is best possible. Closedness is relevant to the study of the topological cycle space: it is important to know that the cycle space as well as the space of all cuts is closed under taking thin sums. Both of these facts follow immediately from Theorem 6.7.

The question whether a space $\mathcal{N}$ is closed under taking thin sums is related to the question whether $\mathcal{N}$ is topologically closed as a subspace of the product space R^M. We investigate this connection in the last section of this chapter.

A line of research very much in the direction of this chapter, albeit more oriented towards graph theory, has been pursued by Casteels and Richter [27].

6.2 Bases

Before we start we need some definitions. Let M be a set, R a ring, and let $\mathcal{N}$ be a family with its members in R^M. Often we will multiply families with coefficients before adding them: if $a = (a(N))_{N\in\mathcal{N}}$ is a family of coefficients in R, one for each $N \in \mathcal{N}$, we use the shorthand $a\mathcal{N}$ for the family $(a(N)N)_{N\in\mathcal{N}}$. For a $K \in R^M$, we call a family a of coefficients a *representation of K in $\mathcal{N}$* if $a\mathcal{N}$ is thin and if $K = \sum a\mathcal{N}$—that is, $K(m) = \sum_{N\in\mathcal{N}} a(N)N(m)$ for every $m \in M$. Denote by $\langle\mathcal{N}\rangle$ the set of elements of R^M that have a representation in $\mathcal{N}$. Intuitively, $\mathcal{N}$ is a generating set, and $\langle\mathcal{N}\rangle$ is the space it generates.

For a family $\mathcal{N} \subseteq R^M$, we call a subfamily $\mathcal{B}$ of $\mathcal{N}$ a *basis of* $\langle\mathcal{N}\rangle$, if $\langle\mathcal{B}\rangle = \langle\mathcal{N}\rangle$ and $\mathbf{0} \in R^M$ has a unique representation in $\mathcal{B}$. Note that $\mathbf{0}$ has a unique representation in some family $\mathcal{N}$ if and only if every element in $\langle\mathcal{N}\rangle$ has a unique representation in $\mathcal{N}$.

It is well known that a generating set in a module does not need to contain a basis (in the classical sense), and, clearly, this is also the case in our setting: Take for example $R = \mathbb{Z}$, $M = \{0\}$, and $\mathcal{N} = \{a_2, a_3\}$, where a_2, a_3 are defined by $a_2(0) = 2$ and $a_3(0) = 3$. Then $\mathcal{N}$ does not contain a basis of $\langle\mathcal{N}\rangle$. When R is a field, however, we can say more about the existence of a basis:

Theorem 6.1. **[20]** *Let M be a countable set, F be a field and let $\mathcal{N}$ be a family with its members in F^M. Then $\mathcal{N}$ contains a basis of $\langle\mathcal{N}\rangle$.*

In linear algebra the analogous assertion is usually proved with Zorn's lemma as follows. Given a chain $(\mathcal{B}_\lambda)_\lambda$ of linearly independent subsets of the generating set, it is observed that $\bigcup_\lambda \mathcal{B}_\lambda$ is still linearly independent since any violation of linear independence is witnessed by finitely many elements, and these would already lie in one of the $\mathcal{B}_\lambda$. Thus, each chain has an upper bound, which implies, by Zorn's lemma, that there is a maximal linearly independent set, a basis. This approach, however, fails in our context, as dependence does not need to be witnessed by only finitely many elements, thus we cannot get the contradiction that already one of the $\mathcal{B}_\lambda$ was not independent.

As an illustration, put $F = \mathbb{Z}_2$, $M := \mathbb{Z}$ and $\mathcal{B}_i := \{\{j, j{+}1\} : -i \leq j < i\}$ for $i = 1, 2, \ldots$. (Here, and later we shall freely identify elements of $\mathbb{Z}_2^M$ with subsets of M.) Now, while no nonempty finite subset of $\mathcal{B}_\infty := \bigcup_{i=1}^\infty \mathcal{B}_i$ is dependent, the whole set is: $\sum_{B\in\mathcal{B}_\infty} B = \emptyset$.

The standard proof outlined above pursues a bottom-up approach. There is also a, perhaps more pedestrian, top-down proof, where successively those elements of $\mathcal{N}$ are weeded out that can be replaced by others. Slightly more precisely, a (possibly transfinite) enumeration $N_1, N_2, \ldots$ of $\mathcal{N}$ is processed step-by-step, and in each step it is checked whether N_λ can be expressed as a sum in earlier N_μ (i.e. $\mu < \lambda$) that are still left. If yes, N_λ is not needed to generate $\langle\mathcal{N}\rangle$, and therefore deleted. If no, N_λ is kept. It is not hard to check that this process yields a basis.

Our proof of Theorem 6.1 relies on an adaptation of this argument. Clearly, we cannot expect it to work as it is, since it fails to take infinite sums into account. However, if we restrict ourselves to those elements of $\mathcal{N}$ that share a given $m \in M$ then all sums are finite, and we can employ the argument. So, we will partition $\mathcal{N}$ into sets $\mathcal{N}_1, \mathcal{N}_2, \ldots$ so that all the elements in each of the $\mathcal{N}_j$ share an $m \in M$. Then we will use an (adapted) top-down argument on each of the $\mathcal{N}_i$.

Proof of Theorem 6.1. Let $m_1, m_2, \ldots$ be a (possibly finite) enumeration of M, and for $i = 1, 2, \ldots$ define $\mathcal{N}_i$ to be the set of those elements $N \in \mathcal{N} \setminus \bigcup_{j<i} \mathcal{N}_j$ for which $N(m_i) \neq 0$. Clearly, $\{\mathcal{N}_i : i \in \mathbb{N}\}$ is a partition of $\mathcal{N}$. For every $i \in \mathbb{N}$, let $N_{i1}, N_{i2}, \ldots, N_{i\lambda}, \ldots$ be a (possibly transfinite) enumeration of $\mathcal{N}_i$.

Now, for each $i = 1, 2, \dots$ we perform a transfinite induction as follows. Start by setting $\mathcal{B}_{i0} = \emptyset$, and then for every ordinal $\lambda > 0$ define the set $\mathcal{B}_{i\lambda} \subseteq \mathcal{N}_i$ as follows ($\mathcal{B}_{i\lambda}$ is the set of those elements among the first λ $N \in \mathcal{N}_i$ that we will put in our basis): If $L := N_{i\lambda}$ has a representation a_L in $\mathcal{N}$ such that

$$a_L(N) = 0 \textit{ for } N \notin \bigcup_{\mu<\lambda} \mathcal{B}_{i\mu} \cup \bigcup_{k>i} \mathcal{N}_k, \tag{6.1}$$

then let $\mathcal{B}_{i\lambda} = \bigcup_{\mu<\lambda} \mathcal{B}_{i\mu}$. Otherwise, set $\mathcal{B}_{i\lambda} = \bigcup_{\mu<\lambda} \mathcal{B}_{i\mu} \cup \{N_{i\lambda}\}$. Having defined all $\mathcal{B}_{i\lambda}$, we put $\mathcal{B}_i := \bigcup_\lambda \mathcal{B}_{i\lambda}$.

For later use we note that

$$\textit{if } L \in \mathcal{N}_i \setminus \mathcal{B}_i,\ N \in \mathcal{N}_j \setminus \mathcal{B}_i \textit{ and } j \leq i \textit{ then } a_L(N) = 0. \tag{6.2}$$

Indeed, if $a_L(N) \neq 0$ then by (6.1) we obtain $N \in \bigcup_{\mu<\lambda} \mathcal{B}_{i\mu} \cup \bigcup_{k>i} \mathcal{N}_k$ for some λ, and as $j \leq i$ we have $N \in \bigcup_{\mu<\lambda} \mathcal{B}_{i\mu}$. But this means that $N \in \mathcal{B}_i = \bigcup_\mu \mathcal{B}_{i\mu}$, a contradiction.

We claim that $\mathcal{B} := \bigcup_i \mathcal{B}_i$ is a basis of $\langle \mathcal{N} \rangle$. To show that $\mathbf{0} \in R^M$ has a unique representation, suppose there are coefficients $b : \mathcal{B} \to F$, not all of which are zero, such that $b\mathcal{B}$ is thin and $\sum b\mathcal{B} = 0$. Let $i \in \mathbb{N}$ be minimal so that there is an ordinal μ with $b(N_{i\mu}) \neq 0$, and observe that since for all the elements B in $\mathcal{B}_i$ we have $B(m_i) \neq 0$, there is a maximal ordinal λ such that $b(N_{i\lambda}) \neq 0$ (because $b\mathcal{B}$ is thin). Then $N_{i\lambda} = \sum_{N \in \mathcal{B} \setminus \{N_{i\lambda}\}} b^{-1}(N_{i\lambda}) b(N) N$ is (or, more precisely, can be extended to) a representation of $N_{i\lambda}$ that satisfies (6.1), a contradiction to that $N_{i\lambda} \in \mathcal{B}_i$.

Next, consider a $K \in \langle \mathcal{N} \rangle$. We will show that K has a representation in $\mathcal{B}$. Starting with any representation b^0 of K in $\mathcal{N}$, we inductively define for $i = 1, 2, \dots$ representations $b^i : \mathcal{N} \to F$ as follows. (Intuitively, b^i is a representation of K using only elements of $\mathcal{N}$ that are left after step i of the construction of $\mathcal{B}$, that is, after we have finished deleting elements of $\mathcal{N}_i$.) Set $\mathcal{E}_i := \{N \in \mathcal{N}_i \setminus \mathcal{B}_i : b^{i-1}(N) \neq 0\}$. Since $b^{i-1}\mathcal{N}$ is thin and since $N(m_i) \neq 0$ for all $N \in \mathcal{E}_i \subseteq \mathcal{N}_i$, it follows that $\mathcal{E}_i$ is a finite set. Put

$$\begin{aligned} b^i(N) &= 0 \textit{ for } N \in \mathcal{E}_i, \textit{ and} \\ b^i(N) &= b^{i-1}(N) + \sum_{L \in \mathcal{E}_i} b^{i-1}(L) a_L(N) \textit{ for } N \notin \mathcal{E}_i. \end{aligned} \tag{6.3}$$

(Note that a_L is defined for every $L \in \mathcal{E}_i$, since $\mathcal{E}_i \subseteq \mathcal{N}_i \setminus \mathcal{B}_i$.)

We claim that this definition yields a representation of K that uses only those elements of $\mathcal{N}_1, \dots, \mathcal{N}_i$ that lie in $\mathcal{B}$, in other words, we claim that for every $i \in \mathbb{N}$ it is true that

$$b^i(N) = 0 \textit{ if } N \in \bigcup_{j=1}^{i} \mathcal{N}_j \setminus \mathcal{B} \tag{6.4}$$

and

$$K = \sum b^i \mathcal{N} \text{ (in particular, } b^i\mathcal{N} \text{ is thin).} \tag{6.5}$$

To prove the two claims we proceed by induction. For (6.4), consider an $N \in \mathcal{N}_j \setminus \mathcal{B}$ where $j \leq i$. If $N \in \mathcal{E}_i$ then $b^i(N) = 0$ by definition, so consider the case when $N \notin \mathcal{E}_i$. If $j = i$ this implies that $b^{i-1}(N) = 0$; if $j < i$ then we get $b^{i-1}(N) = 0$ too, this time using induction. Since, by (6.2), $a_L(N) = 0$ for every $L \in \mathcal{E}_i$, (6.3) implies $b^i(N) = 0$, as desired.

For (6.5), first note that $b^i\mathcal{N}$ is indeed thin as $\mathcal{E}_i$ is finite and both of $a_L\mathcal{N}$ and $b^{i-1}\mathcal{N}$ thin (the latter by induction). Furthermore:

$$\begin{aligned}
&\sum b^i\mathcal{N} - \sum b^{i-1}\mathcal{N} = \\
&= \sum_{N\in\mathcal{N}\setminus\mathcal{E}_i} (b^i(N) - b^{i-1}(N))N + \sum_{N\in\mathcal{E}_i} (b^i(N) - b^{i-1}(N))N \\
&= \sum_{N\in\mathcal{N}\setminus\mathcal{E}_i} (\sum_{L\in\mathcal{E}_i} b^{i-1}(L)a_L(N)N) + \sum_{N\in\mathcal{E}_i} (\mathbf{0} - b^{i-1}(N)N) \\
&= \sum_{L\in\mathcal{E}_i} b^{i-1}(L) \sum_{N\in\mathcal{N}\setminus\mathcal{E}_i} a_L(N)N - \sum_{N\in\mathcal{E}_i} (\mathbf{b}^{i-1}(N)N).
\end{aligned}$$

As $a_L(N) = 0$ if $N, L \in \mathcal{E}_i$ by (6.2), we obtain for the first sum in the last line

$$\sum_{L\in\mathcal{E}_i} b^{i-1}(L) \sum_{N\in\mathcal{N}\setminus\mathcal{E}_i} a_L(N)N = \sum_{L\in\mathcal{E}_i} b^{i-1}(L) \sum_{N\in\mathcal{N}} a_L(N)N = \sum_{L\in\mathcal{E}_i} b^{i-1}(L)L.$$

Together with the previous equation this yields $\sum b^i\mathcal{N} - \sum b^{i-1}\mathcal{N} = \mathbf{0}$, which proves (6.5).

For every $N \in \mathcal{N}$, define $b^\infty(N) := b^j(N)$ if $N \in \mathcal{N}_j$, and note that

$$b^i(N) = b^\infty(N) \text{ for } N \in \mathcal{N}_j \text{ and } i \geq j. \tag{6.6}$$

Indeed, consider $i > j$ and observe that, by (6.2), $a_L(N) = 0$ for all $L \in \mathcal{E}_i \subseteq \mathcal{N}_i \setminus \mathcal{B}_i$. So, from (6.3) it follows that $b^i(N) = b^{i-1}(N)$.

We immediately get from (6.4) that

$$\text{if } N \in \mathcal{N} \setminus \mathcal{B} \text{ then } b^\infty(N) = 0. \tag{6.7}$$

Thus, the $b^\infty(N)$ can be seen as coefficients on $\mathcal{B}$, and our next claim states that b^∞ is what we are looking for, namely a representation of K in $\mathcal{B}$:

$$K = \sum_{B\in\mathcal{B}} b^\infty(B)B \text{ (in particular, } b^\infty\mathcal{N} \text{ is thin).} \tag{6.8}$$

Consider an $m_i \in M$. By definition of the $\mathcal{N}_j$, every N with $N(m_i) \neq 0$ lies in $\bigcup_{j=1}^{i} \mathcal{N}_j$. By (6.6), b^i and b^∞ are identical on $\bigcup_{j=1}^{i} \mathcal{N}_j$. Since $b^i\mathcal{N}$ is thin, there are therefore only finitely many $N \in \mathcal{N}$ so that $b^\infty(N)N(m_i) \neq \infty$. Thus $b^\infty\mathcal{N}$ is thin. Furthermore, we obtain with (6.5) and (6.6):

$$\sum_{N\in\mathcal{N}} b^\infty(N)N(m_i) = \sum_{N\in\bigcup_{j=1}^{i}\mathcal{N}_j} b^\infty(N)N(m_i) = \sum_{N\in\mathcal{N}} b^i(N)N(m_i) = K(m_i).$$

Claim (6.8) now follows from (6.7). This completes the proof. $\square$

Observe that contrary to conventional linear algebra, two bases do not need to have the same cardinality—even over a field. Indeed, putting $F = \mathbb{Z}_2$ and $M = \{m_0, m_1, \ldots\}$ we see that $\mathcal{B} := \{\{m_i\} : i \geq 0\}$ is a countable basis of F^M. On the other hand, $\mathcal{N} := \{\{m_0\} \cup N : N \subseteq M\}$ clearly generates F^M too, and contains, by Theorem 6.1, a basis $\mathcal{B}'$. Since all thin subsets of $\mathcal{N}$ are finite, $\mathcal{B}'$ needs to be uncountable to generate the uncountable set F^M. Thus $\mathcal{B}$ and $\mathcal{B}'$ are two bases of F^M that do not have the same cardinality.

We have formulated Theorem 6.1 only for countable sets M. The following result shows that this is indeed best possible.

Theorem 6.2. **[20]** *For an uncountable set M there exists a family $\mathcal{N}$ of elements of $\mathbb{Z}_2^M$, so that $\mathcal{N}$ does not contain a basis of $\langle\mathcal{N}\rangle$.*

Proof of Theorem 6.2. Let A, B be two disjoint sets with cardinalities $|A| = \aleph_0$ and $|B| = \aleph_1$. Define G to be the graph with vertex set $M := A \cup B$ and edge set $\mathcal{N} := A \times B$. As $\mathcal{N} \subseteq \mathcal{P}(M)$, we may ask whether $\mathcal{N}$ contains a basis of $\langle\mathcal{N}\rangle$. We claim that it does not.

Let us show that each countable subset N of M is an element of $\langle\mathcal{N}\rangle$. Indeed, let $n_1, n_2, \ldots$ be a (possibly finite) enumeration of N, and choose for $i = 1, 2, \ldots$ a ray R_i that starts at n_i, and does not meet the first $i-1$ vertices of each of $R_1, \ldots, R_{i-1}$ except, possibly, at n_i. Then, the set $\bigcup_{i\in\mathbb{N}} E(R_i)$ of edges of these rays is thin, and its sum equals N since $\sum_{e\in E(R_i)} e = \{n_i\}$.

Suppose that $\langle\mathcal{N}\rangle$ has a basis $\mathcal{B} \subseteq \mathcal{N}$, and let H be the graph with vertex set M and edge set $\mathcal{B}$. Since $\mathcal{B}$ must contain for each element in B at least one edge incident with it, $\mathcal{B}$ is uncountable. Therefore, one of the vertices in the countable set A, say v, is incident with infinitely many edges in $\mathcal{B}$. Delete from H the vertex v (and its incident edges) and denote by $\mathcal{C}$ the set of components of the resulting graph that (in H) are adjacent to v.

Observe that for each $C \in \mathcal{C}$ there is exactly one edge in H between v and some vertex, u_C say, in C; indeed, if there were two edges between v and vertices u, u' in C, then the union of these edges with an u–u' path in C

would be a cycle in H, contradicting that $\emptyset$ has a unique representation in $\mathcal{B}$ since the sum of the edges of a cycle equals $\emptyset$.

Next, suppose there are distinct $C, D \in \mathcal{C}$ each containing a ray; then, C (respectively D) also contains a ray R (resp. S) starting at u_C (resp. at u_D). Then $R \cup S$ together with the two edges between v and $\{u_C, u_D\}$ yields a set of edges which sums to $\emptyset$, again a contradiction.

Pick a countably infinite number of $C \in \mathcal{C}$ none of which contains a ray, and denote the set of these by $\mathcal{C}'$. As $N := \{u_C : C \in \mathcal{C}'\}$ is countable it lies in $\langle \mathcal{N} \rangle$, thus there is a $\mathcal{B}_N \subseteq \mathcal{B}$ such that $\sum_{e \in \mathcal{B}_N} e = N$.

Suppose there is a $C \in \mathcal{C}'$ such that an edge $e \in \mathcal{B}_N$ incident with u_C lies in C. As C does not contain any cycle or any ray, we can run from e along edges in $E(C) \cap \mathcal{B}_N$ to a vertex $w \neq u_C$ that is only incident with one edge in $\mathcal{B}_N$. This implies that $w \in \sum_{e \in \mathcal{B}_N} e = N$, a contradiction since $w \notin N$. However, $u_C \in N$ must be incident with an edge from $\mathcal{B}_N$. Consequently, for each $C \in \mathcal{C}'$ the edge between v and u_C lies in $\mathcal{B}_N$, contradicting that $\mathcal{B}_N$ is thin. $\square$

6.3 Closedness under taking thin sums

In this section we investigate the following question:

Question 6.3. *Let M be a set, R be a ring and $\mathcal{N} \subseteq R^M$. When is $\langle \mathcal{N} \rangle$ closed under taking thin sums, i.e. when is $\langle \mathcal{N} \rangle = \langle \langle \mathcal{N} \rangle \rangle$?*

In conventional algebra, the answer is easy: always. Once we allow infinite sums, however, the answer is not that straightforward. Consider, for instance, the case when $M = \mathbb{N}$, $R = \mathbb{Z}_2$ and $\mathcal{N} := \{\{1, i\} : i \in \mathbb{N}\}$. Clearly, we have $\{i\} \in \langle \mathcal{N} \rangle$ for all $i \in \mathbb{N}$ and thus $\mathbb{N} = \sum_{i \in \mathbb{N}} \{i\} \in \langle \langle \mathcal{N} \rangle \rangle$. On the other hand, $\mathbb{N} \notin \langle \mathcal{N} \rangle$ as all thin sums of elements in $\mathbb{N}$ are necessarily finite. Thus, $\langle \mathcal{N} \rangle$ is indeed a proper subset of $\langle \langle \mathcal{N} \rangle \rangle$, and therefore not closed under taking thin sums.

The example seems to indicate that $\mathcal{N}$ needs to be thin. Indeed, if we require $\mathcal{N}$ to be thin then we will see that $\langle \mathcal{N} \rangle$ is closed under taking thin sums—provided that R is a field or a finite ring (Theorem 6.7). At the end of the section we will give an example showing that this is a fairly complete answer to Question 6.3: If R is neither a field nor finite, then we cannot guarantee that $\langle \mathcal{N} \rangle = \langle \langle \mathcal{N} \rangle \rangle$.

We remark that there is another way to overcome the counterexample above. Vella [86] shows that a family $\mathcal{N}$ of elements in $\mathbb{Z}_2^{\mathbb{N}}$ is closed under taking thin sums if, instead of being thin, $\mathcal{N}$ has the property that every sum of finitely many members of $\mathcal{N}$ is the disjoint union of members of $\mathcal{N}$.

It turns out that Question 6.3 is closely related to the topological closure in the product space $R^M = \prod_{m \in M} R_m$ where each R_m is a copy of R endowed with the discrete topology. In what follows we denote by $\overline{\mathcal{N}}$ the topological closure of a subset $\mathcal{N}$ of R^M in the space $\prod_{m \in M} R_m$. Given a $K \in R^M$, note that $K \in \overline{\mathcal{N}}$ if and only if for every finite subset M' of M there is a $N \in \mathcal{N}$ with $K(m) = N(m)$ for all $m \in M'$. In the next two lemmas we will see that for a thin family $\mathcal{T}$ of elements of R^M, $\langle \mathcal{T} \rangle$ is topologically closed if R is a finite ring or a field. Moreover, we will conclude in Lemma 6.6 that $\langle \langle \mathcal{T} \rangle \rangle$ lies in $\overline{\langle \mathcal{T} \rangle}$, and combining these results yields an answer to Question 6.3. We will pursue the relationship between topological closedness and closedness under taking thin sums further in the next section.

The first proof uses a typical compactness argument.

Lemma 6.4. *Let M be a set, R be a finite ring, and let $\mathcal{T}$ be a thin family of elements of R^M. Then $\overline{\langle \mathcal{T} \rangle} = \langle \mathcal{T} \rangle$.*

Proof. Consider an element $K \in \overline{\langle \mathcal{T} \rangle}$. By Tychonoff's theorem, the product space $X := \prod_{T \in \mathcal{T}} R$ where R bears the discrete topology is compact. For any finite subset $M' \subseteq M$, we consider the set $A_{M'}$ of families of coefficients a, such that $a\mathcal{T}$ agrees with K on M'; formally, define

$$A_{M'} := \{a \in X : \sum_{T \in \mathcal{T}} a(T)T(m) = K(m) \textit{ for every } m \in M'\}.$$

(Note that we view the elements of X as coefficients for the family $\mathcal{T}$.) We claim that these sets are closed in X, and that their collection has the finite intersection property.

To show that each $A_{M'}$ is closed, let $\mathcal{S}_{M'}$ be the subfamily of all those $T \in \mathcal{T}$ for which there is a $m \in M'$ with $T(m) \neq 0$. As $\mathcal{T}$ is thin, $\mathcal{S}_{M'}$ is finite. Since R is finite as well, there are only finitely many $b : \mathcal{S}_{M'} \to R$ such that $\sum_{S \in \mathcal{S}_{M'}} b(S)S(m) = K(m)$ for all $m \in M'$. For each such b, $B_b := \{a \in X : a(S) = b(S)$ for every $S \in \mathcal{S}_{M'}\}$ is closed in X, and since $A_{M'}$ is the union of these finitely many sets B_b it is closed too.

Next, if we have finite sets $M_1, \ldots, M_l \subseteq M$ then, clearly, $\bigcap_{i=1}^{l} A_{M_i} = A_{M''}$ for $M'' := \bigcup_{i=1}^{l} M_i$. As $K \in \overline{\langle \mathcal{T} \rangle}$, there is an element L of $\langle \mathcal{T} \rangle$ that agrees with K on M''. But any representation of L in $\mathcal{T}$ is an element of $A_{M''}$, thus the $A_{M'}$ have the finite intersection property as claimed.

Now, the assertion of the theorem follows easily: X is compact, therefore, the intersection of all $A_{M'}$ is non-empty. For an element a of that intersection, we have $\sum_{T \in \mathcal{T}} a(T)T(m) = K(m)$ for all $m \in M$, i.e. it is a representation of K in $\mathcal{T}$. $\square$

We need a completely different approach to prove a similar result in the case when R is a (possibly infinite) field:

Lemma 6.5. [20] *Let M be a set, F be a field, and let $\mathcal{T}$ be a thin family of elements of F^M. Then $\overline{\langle \mathcal{T} \rangle} = \langle \mathcal{T} \rangle$.*

Proof. Consider a $K \in \overline{\langle \mathcal{T} \rangle}$. We will reduce the problem of finding a representation of K in $\mathcal{T}$ to the solution of an infinite system of equations. To do this, we associate a variable x^T with every member T of $\mathcal{T}$, and for each $m \in M$ we define e_m to be the linear equation

$$\sum_{T \in \mathcal{T} : T(m) \neq 0} x^T T(m) = K(m)$$

in the variables x^T. Note that as $\mathcal{T}$ is thin, each e_m contains only finitely many variables. Let $E = \{e_m : m \in M\}$. By construction, if there is an assignment $a : \mathcal{T} \to F$ such that setting $x^T = a(T)$ for every $T \in \mathcal{T}$ yields a solution to every equation in E, then a is a representation of K in $\mathcal{T}$. So in what follows, our task is to find such a solution.

For every $T \in \mathcal{T}$, define d_T to be the linear equation $x^T = 1$, put $D = \{d_T : T \in \mathcal{T}\}$, and denote by $\mathcal{E}$ the set of all those sets E' with $E \subseteq E' \subseteq E \cup D$ such that every finite subset $E'' \subseteq E'$ has a solution. We claim that $\mathcal{E}$ contains a $\subseteq$-maximal element E^*. First, note that $\mathcal{E} \neq \emptyset$ as $E \in \mathcal{E}$; indeed, for a finite subset $E'' \subseteq E$, the set M' of all $m \in M$ for which $e_m \in E''$ is finite. As $K \in \overline{\langle \mathcal{T} \rangle}$, there is an element L of $\langle \mathcal{T} \rangle$ that agrees with K on M', and any representation a of L in $\mathcal{T}$ yields a solution of E''. Second, if $(E_i)_{i \in I}$ is a chain in $\mathcal{E}$ then, clearly, the union $\bigcup_{i \in I} E_i$ lies in $\mathcal{E}$, too. Thus, Zorn's lemma ensures the existence of E^*.

Next, we show that for every $T \in \mathcal{T}$ there is a finite $E_T \subseteq E^*$ and an element f_T of the field F such that $x^T = f_T$ in every solution of E_T. Suppose not, and observe that then, clearly, $d_T \notin E^*$. Consider a finite subset E' of E^*, and note that, by assumption, there are two solutions in which x^T takes two distinct values, g_1 and g_2, say. Then, for every $f \in F$ there is a solution of E' where $x^T = f g_1 + (1 - f) g_2$. Setting $f = (1 - g_2)(g_1 - g_2)^{-1}$ yields a solution of E' with $x^T = 1$, which means that $E' \cup \{d_T\}$ has a solution. As E' was an arbitrary finite subset of E^*, it follows that every finite subset of $E^* \cup \{d_T\}$ has a solution, contradicting the maximality of E^*. Thus, E_T and f_T exist, as we have claimed.

Finally, define coefficients $a(T) := f_T$ for $T \in \mathcal{T}$. To see that a is a solution of E, consider an arbitrary $m \in M$. As $\mathcal{T}$ is thin, the subfamily $\mathcal{T}_m$ of those members T of $\mathcal{T}$ with $T(m) \neq 0$ is finite. Thus, $E' := \{e_m\} \cup \bigcup_{T \in \mathcal{T}_m} E_T$ has, as a finite subset of E^*, a solution $b : \mathcal{T} \to F$. Since for every $T \in \mathcal{T}_m$, we have $E_T \subseteq E'$ it follows that $b(T) = f_T = a(T)$. As b solves e_m we see that a solves e_m, too. Thus a is a solution of E, and hence a representation of K in $\mathcal{T}$. $\square$

We need one more simple lemma:

Lemma 6.6. **[20]** *Let M be a set, R be a ring, and let $\mathcal{N}$ be a family of elements of R^M. Then $\langle\langle\mathcal{N}\rangle\rangle \subseteq \overline{\langle\mathcal{N}\rangle}$.*

Proof. Let $K \in \langle\langle\mathcal{N}\rangle\rangle$, and consider an arbitrary finite subset M' of M. Denote the canonical projection of R^M to $R^{M'}$ by π, and observe that $\pi(K) \in \langle\langle\pi(\mathcal{N})\rangle\rangle$. Since rings are closed under addition, and since all thin families in $R^{M'}$ are finite, it holds that $\langle\langle\pi(\mathcal{N})\rangle\rangle = \langle\pi(\mathcal{N})\rangle$. Choosing a representation of $\pi(K)$ in $\langle\pi(\mathcal{N})\rangle$, and replacing each element of $\pi(\mathcal{N})$ in it by one of its preimages with respect to π yields an $N \in \langle\mathcal{N}\rangle$ so that $K(m) = N(m)$ for all $m \in M'$. This proves that $K \in \overline{\langle\mathcal{N}\rangle}$. $\square$

Lemma 6.6 combined with Lemmas 6.5 and 6.4 immediately implies the following:

Theorem 6.7. **[20]** *Let M be a set, R be a ring, and let $\mathcal{T}$ be a thin family of elements of R^M. Then $\langle\mathcal{T}\rangle = \langle\langle\mathcal{T}\rangle\rangle$ if R is a field or a finite ring.*

As mentioned in Section 6.1, an immediate consequence of Theorem 6.7 is that the topological cycle space as well as the cut space of a locally finite graph is closed under taking thin sums. Both these spaces are generated by thin sets: the former by the fundamental circuits of a topological spanning tree (see Theorem 1.5), and the latter by the cuts separating a single vertex from the rest of the graph.

Let us now argue that Theorem 6.7 gives, in a sense, a comprehensive answer to Question 6.3. In fact, we shall construct an example where R is neither a field nor finite, and where there exists a thin family $\mathcal{T} \subseteq R^M$ such that $\langle\mathcal{T}\rangle$ is not closed under taking thin sums.

For this, set $R := \mathbb{Z}$ and $M := \mathbb{N}$. Define $N \in \mathbb{Z}^{\mathbb{N}}$ by $N(i) = 1$ for every $i \in \mathbb{N}$. For $j = 1, 2, \ldots$, define $N_j \in \mathbb{Z}^{\mathbb{N}}$ by $N_j(j) = p_j$ and $N_j(i) = 0$ for every $i \neq j$, where p_j is the jth prime number. Let $\mathcal{T} = \{N, N_1, N_2, \ldots\}$, and note that $\mathcal{T}$ is thin. We will show that the function $K \in \mathbb{Z}^{\mathbb{N}}$ defined by $K(i) = i$ is in $\langle\langle\mathcal{T}\rangle\rangle$ but not in $\langle\mathcal{T}\rangle$.

Let us first prove that $K \notin \langle\mathcal{T}\rangle$. Suppose for contradiction, there is a representation $a : \mathcal{T} \to \mathbb{Z}$ of K in $\langle\mathcal{T}\rangle$. We distinguish two cases, depending on whether $n := a(N)$ is non-negative or not. If $n \geq 0$, then $n + 1 = K(n+1) = \sum_{L \in \mathcal{T}} a(L)L(n+1) = a(N)N(n+1) + a(N_{n+1})N_{n+1}(n+1) = n + a(N_{n+1})p_{n+1}$, which implies that $1 = a(N_{n+1})p_{n+1}$, a contradiction. If, on the other hand, $n < 0$ then we have for $n' = -n$ that $n' = K(n') = n + a(N_{n'})p_{n'}$, i.e. $2n' = a(N_{n'})p_{n'}$. With $p_{n'} > n'$ we obtain $p_{n'} = 2$ and thus $n = -1$. This again leads to a contradiction, as it implies that $K(3) = 3 = -1 + a(N_3) \cdot 5$.

Having shown $K \notin \langle \mathcal{T} \rangle$, we now prove $K \in \langle\langle \mathcal{T} \rangle\rangle$. For this, we construct for every $i \in \mathbb{N}$ an $S_i \in \langle \mathcal{T} \rangle$ so that $S_i(i) = 1$ and $S_i(j) = 0$ for every $j < i$. Once we have done that we can represent K with these S_i: put $d(S_1) = 1$ and, inductively, set $d(S_i) = i - \sum_{j=1}^{i} d(S_j)S_j(i)$. Then $K = \sum_{i \in \mathbb{N}} d(S_i)S_i$.

Let us now find coefficients $a_0, \ldots, a_i \in \mathbb{Z}$ so that $S_i := a_0 N + \sum_{j=1}^{i} a_j N_j$ is as desired. It follows from the Chinese remainder theorem that the system of congruences

$$\begin{aligned} x &\equiv 0 \pmod{p_1} \\ &\vdots \\ x &\equiv 0 \pmod{p_{i-1}} \\ x &\equiv 1 \pmod{p_i} \end{aligned}$$

has a solution $a_0 \in \mathbb{Z}$. This allows us to choose $a_j \in \mathbb{Z}$ so that $a_0 + a_j p_j = 0$, for every $1 \le j < i$, and $a_i \in \mathbb{Z}$ so that $a_0 + a_i p_i = 1$.

6.4 Thin sums and topological closure

In [86], Vella introduced for a family $\mathcal{N}$ of elements of R^M the following notation and spaces:

- the *weak span* $\mathcal{W}(\mathcal{N})$ is the set of all finite sums of elements of $\mathcal{N}$ with coefficients in R, i.e. $\mathcal{W}(\mathcal{N})$ is the R-module generated by $\mathcal{N}$;
- the *algebraic span* $\mathcal{A}(\mathcal{N})$ is the set of all thin sums of elements of $\mathcal{N}$ with coefficients in R, i.e. what we have called (and will continue to call) $\langle \mathcal{N} \rangle$; and
- the *strong span* $\mathcal{S}(\mathcal{N})$ is the intersection of all sets $\mathcal{M} \supseteq \mathcal{N}$ that are closed under taking thin sums, i.e. the smallest set $\mathcal{S}$ containing $\mathcal{N}$ with $\langle \mathcal{S} \rangle = \mathcal{S}$.

Let us add to this list a fourth space, namely $\overline{\mathcal{W}(\mathcal{N})}$, the topological closure of $\mathcal{W}(\mathcal{N})$ in the product space R^M. It is not hard to see that it contains the strong span of $\mathcal{N}$:

Lemma 6.8. **[20]** *For any ring R, any set M and any family $\mathcal{N}$ of elements of R^M, $\overline{\mathcal{W}(\mathcal{N})}$ is closed under taking thin sums, i.e.* $\left\langle \overline{\mathcal{W}(\mathcal{N})} \right\rangle = \overline{\mathcal{W}(\mathcal{N})}$.

Proof. Consider a thin family $\mathcal{T}$ of elements of $\overline{\mathcal{W}(\mathcal{N})}$. We need to show that $S := \sum_{T \in \mathcal{T}} T \in \overline{\mathcal{W}(\mathcal{N})}$. For this, let M' be an arbitrary finite subset of M. As $\mathcal{T}$ is thin, the subfamily $\mathcal{T}'$ of those members T of $\mathcal{T}$ for which $T(m) \neq 0$ for some $m \in M'$ is finite. For each $T \in \mathcal{T}'$ there exists a finite subfamily $\mathcal{N}_T$ of $\mathcal{N}$ and coefficients $r_N^T \in R$ for $N \in \mathcal{N}_T$ so that $\sum_{N \in \mathcal{N}_T} r_N^T N(m) = T(m)$ for all $m \in M'$ since T lies in $\overline{\mathcal{W}(\mathcal{N})}$. Then, $S' := \sum_{T \in \mathcal{T}'} \sum_{N \in \mathcal{N}_T} r_N^T N$ is an element of $\mathcal{W}(\mathcal{N})$ and $S'(m) = S(m)$ for all $m \in M'$. As M' was arbitrary, this means that $S \in \overline{\mathcal{W}(\mathcal{N})}$. □

We thus have the following inclusions:

$$\mathcal{W}(\mathcal{N}) \subseteq \langle \mathcal{N} \rangle \subseteq \mathcal{S}(\mathcal{N}) \subseteq \overline{\mathcal{W}(\mathcal{N})}.$$

Clearly, the first two inclusions can be proper. But can the third one also be proper?

Question 6.9. *Is $\mathcal{S}(\mathcal{N}) = \overline{\mathcal{W}(\mathcal{N})}$ for every family $\mathcal{N}$ of elements of R^M?*

In the special case when $\mathcal{N}$ is thin and R a field or a finite ring we obtain from the results of the previous section (Lemmas 6.4 and 6.5) that $\langle \mathcal{N} \rangle = \overline{\langle \mathcal{N} \rangle}$, which clearly implies $\mathcal{S}(\mathcal{N}) = \overline{\mathcal{W}(\mathcal{N})}$. This answers a question of Manfred Droste (personal communication).

For countable M, Question 6.9 has an affirmative answer, too. This can easily be seen using a telescoping sum argument:

Proposition 6.10.[20] *If R is any ring and M a countable set then $\overline{\mathcal{W}(\mathcal{N})} = \mathcal{S}(\mathcal{N})$ for any family $\mathcal{N}$ of elements of R^M.*

Proof. We only need to prove that $\overline{\mathcal{W}(\mathcal{N})} \subseteq \mathcal{S}(\mathcal{N})$. Consider an arbitrary $K \in \overline{\mathcal{W}(\mathcal{N})}$, and let $m_1, m_2 \ldots$ be an enumeration of M. As K lies in $\overline{\mathcal{W}(\mathcal{N})}$ there is for every $i \in \mathbb{N}$ an $N_i \in \mathcal{W}(\mathcal{N})$ so that $N_i(m_j) = K(m_j)$ for all $j \leq i$. Set $N_0 = \mathbf{0}$ and define $L_i = N_i - N_{i-1} \in \mathcal{W}(\mathcal{N})$ for every i. Then, $L_i(m_j) = 0$ for $j < i$ and, consequently, the L_i form a thin family. Furthermore:

$$\left(\sum_{j=1}^{\infty} L_j \right)(m_i) = \left(\sum_{j=1}^{i} L_j \right)(m_i) = N_i(m_i) = K(m_i).$$

As $L_j \in \mathcal{W}(\mathcal{N}) \subseteq \mathcal{S}(\mathcal{N})$ for every j and $\mathcal{S}(\mathcal{N})$ is closed under taking thin sums, we obtain that $K = \sum_{j=1}^{\infty} L_j$ lies in $\mathcal{S}(\mathcal{N})$. □

In the case of the topological cycle space $\mathcal{C}$ of a graph, the set M, the set of edges of the graph, is usually countable and so Proposition 6.10 is applicable (with $\mathcal{N} = \mathcal{C}$). Moreover, $\mathcal{C}$ is generated by a thin set (see Theorem 1.5), thus we obtain with Theorem 6.7 that $\mathcal{C} = \overline{\mathcal{C}}$. A technique that appears in a number of proofs, see e.g. [15, 16, 19, 45, 46, 47], makes implicit use of this fact. In those proofs, an infinite cycle or element of $\mathcal{C}$ with certain properties is sought. The standard way to construct the desired object is to approximate it by a sequence of finite cycles or elements of $\mathcal{C}$ and to consider the limit of this sequence. That this limit lies indeed in $\mathcal{C}$ is usually proved explicitly, but follows directly from our corollary that $\mathcal{C} = \overline{\mathcal{C}}$.

For uncountable M, Question 6.9 is not as easy to answer and indeed Proposition 6.10 becomes false. In the rest of this chapter we will present a family $\mathcal{N}$ of elements of $\mathbb{Z}_2^{[0,1]}$ for which the inclusion $\mathcal{S}(\mathcal{N}) \subseteq \overline{\mathcal{W}(\mathcal{N})}$ is proper.

Again, we will view elements of $\mathbb{Z}_2^{[0,1]}$ as subsets of $[0,1]$. In particular, the elements of $\mathcal{N}$ will consist of disjoint unions of intervals of $[0,1]$. These intervals will be chosen from an ever finer subdivision of $[0,1]$. More precisely, we will construct $\mathcal{N}$ in ω steps each of which corresponds to a certain level of coarseness: we start in step 0 with the whole interval, so the first elements of $\mathcal{N}$ will be $[0,1]$ and the empty set. In the next step, we cut $[0,1]$ in half, and in the following step we subdivide each of these halves again into two halves, and so on. Thus, in step n, we have the intervals $[0,1/2^n],\ldots,[(2^n-1)/2^n,1]$ at our disposal, and an element N constructed in this step will be the union of some of these intervals. However, not every such union will be put in $\mathcal{N}$.

Before we start with the formal definition, let us make one amendment. When we add (perhaps many) intervals of the form $[\frac{i-1}{2^n},\frac{i}{2^n}]$ it is not so easy to keep track of what happens with the points on the boundary of the intervals. While this is not a serious problem, it complicates the matter. To circumvent this, we will delete from our ground set all those points that can ever arise as a boundary of some interval. These are precisely the points $J := \{\frac{i}{2^n} : 0 \le i \le 2^n,\, i,n \in \mathbb{N}\}$, and the subsets of $\mathcal{N}$ will therefore be subsets of $[0,1] \setminus J$.

We begin with the definition of the "intervals": for $n \in \mathbb{N}$ and every $i \in \{1,2,\ldots 2^n\}$ let $I_n^i = [\frac{i-1}{2^n},\frac{i}{2^n}] \setminus J$, and define $\mathcal{I}_n := \{I_n^i : i \in \{1,2,\ldots 2^n\}\}$. In step 0, we set $S_0 := [0,1] \setminus J$ and $\mathcal{N}_0 := \{\emptyset, S_0\}$. Then, in step $n+1$, assuming that we have already defined nested sets $\mathcal{N}_0 \subseteq \ldots \subseteq \mathcal{N}_n$ in the previous steps, we construct a new "seed" element by taking every second interval in $\mathcal{I}_{n+1}$:

$$S_{n+1} := \bigcup_{i=1}^{2^n} I_{n+1}^{2i}$$

By adding this seed to the existing elements we define the new ones:

$$\mathcal{N}_{n+1} := \mathcal{N}_n \cup \{N + S_{n+1} : N \in \mathcal{N}_n\}.$$

Once all these $\mathcal{N}_n$ are constructed, we put $\mathcal{N} := \bigcup_{n=0}^{\infty} \mathcal{N}_n$. See Figure 6.1 for the first few elements of $\mathcal{N}$.

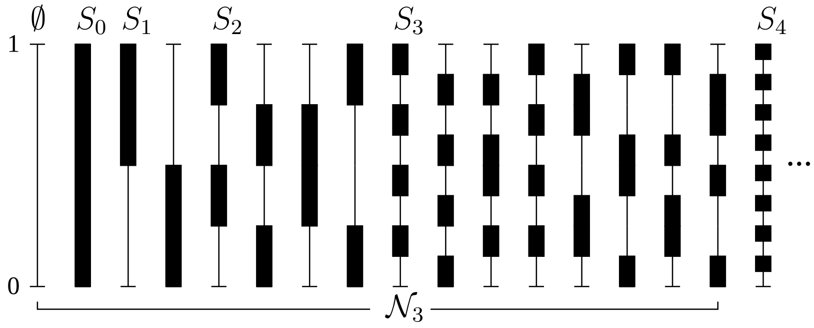

Figure 6.1: A schematic drawing of the first few elements of $\mathcal{N}$

We will accomplish our aim, to show that $\mathcal{S}(\mathcal{N}) \neq \overline{\mathcal{W}(\mathcal{N})}$, in three steps. In each of these we prove one of the following assertions:

(i) $\mathcal{N}$ *is closed under taking finite sums, i.e.* $\mathcal{N} = \mathcal{W}(\mathcal{N})$;

(ii) $\mathcal{W}(\mathcal{N}) = \mathcal{S}(\mathcal{N})$; and

(iii) $\mathcal{N} \neq \overline{\mathcal{N}}$.

Combining (i), (ii) and (iii) we immediately obtain $\mathcal{S}(\mathcal{N}) \neq \overline{\mathcal{W}(\mathcal{N})}$.

In order to establish (i), we will show inductively that each $\mathcal{N}_n$ is already closed under taking finite sums, i.e. that $\mathcal{W}(\mathcal{N}_n) = \mathcal{N}_n$. For this, consider $N, L \in \mathcal{N}_n$. If $N, L \in \mathcal{N}_{n-1}$ then we are done by induction. So we may assume that $N \in \mathcal{N}_n \setminus \mathcal{N}_{n-1}$, i.e. that there is an $N' \in \mathcal{N}_{n-1}$ with $N = S_n + N'$. Now if $L \in \mathcal{N}_{n-1}$ then $N' + L \in \mathcal{N}_{n-1}$ (by induction) and thus $N + L = S_n + (N' + L) \in \mathcal{N}_n$. If, on the other hand, $L = S_n + L'$ for some $L' \in \mathcal{N}_{n-1}$, then $N + L = S_n + N' + S_n + L' = N' + L' \in \mathcal{N}_{n-1}$, which completes the proof of (i).

For the proof of (ii), we will need some intermediate assertions. The first one states that

$$\textit{if } N \in \mathcal{N}_n \textit{ and } I \in \mathcal{I}_n \textit{ then either } I \subseteq N \textit{ or } I \cap N = \emptyset. \tag{6.9}$$

To prove this, we perform induction on n. Note that I is contained in an $I' \in \mathcal{I}_{n-1}$. Thus, if $N \in \mathcal{N}_{n-1}$ then the assertion holds by the induction hypothesis. If, however, $N \in \mathcal{N}_n \setminus \mathcal{N}_{n-1}$, then $N = S_n + N'$ for some $N' \in \mathcal{N}_{n-1}$. The assertion holds for N' (by induction) and for S_{n+1} (by construction), and therefore it is also true for $N = S_{n+1} + N'$.

Next, we prove that

$$\textit{every } N \in \mathcal{N}_{n+1} \setminus \mathcal{N}_n \textit{ meets every } I \in \mathcal{I}_n. \tag{6.10}$$

Indeed, it is easy to see that this holds for $n = 0$. Now, suppose that $N = S_{n+1} + N'$ where $N' \in \mathcal{N}_n$. By construction, $I \cap S_{n+1}$ is a non-empty proper subset of I. Since, by (6.9), either $I \subseteq N'$ or $I \cap N' = \emptyset$, it follows that $S_{n+1} + N'$ meets I.

Before we deduce (ii), we need one final assertion:

$$\begin{array}{l}\textit{if } \mathcal{L} \textit{ is an infinite subset of } \mathcal{N} \textit{ and if } I \in \mathcal{I}_n \textit{ for some } n \in \mathbb{N} \\ \textit{then there exists an } r \in I \textit{ lying in infinitely many } L \in \mathcal{L}.\end{array} \tag{6.11}$$

To prove this, let $N_1, N_2, \ldots$ be distinct elements of $\mathcal{L}$, and define $n_k \in \mathbb{N}$ by $N_k \in \mathcal{N}_{n_k} \setminus \mathcal{N}_{n_k - 1}$. We may assume that the N_k are ordered so that $n_1 \leq n_2 \leq \ldots$. By going to a subsequence we may even assume that $n_{k+1} - n_k \geq 2$ for all k, and that $n_1 \geq n+1$ (note that $|\mathcal{N}_n| < \infty$ for all n). Thus there is, by (6.9) and (6.10), an $I_1 \in \mathcal{I}_{n_1}$ with $I_1 \subseteq N_1 \cap I$. Now $I_1 = I^j_{n_1}$ for some j, and I_1 is the union of two elements of $\mathcal{I}_{n_1+1}$: the "left half" $I_l := I^{2j-1}_{n_1+1}$ and the "right half" $I_l := I^{2j}_{n_1+1}$. Since $n_2 - n_1 \geq 2$ it follows from (6.10) that N_2 meets both of I_l and I_r, and therefore both of $I_l \cap N_2$ and $I_r \cap N_2$ contain, by (6.9), an element of $\mathcal{I}_{n_2}$ as a subset. We choose an $I_2 \in \mathcal{I}_{n_2}$ with $I_2 \subseteq I_l \cap N_2$. Continuing in this manner, we find nested sets $I \supseteq I_1 \supseteq I_2 \supseteq \ldots$ with $I_k \subseteq N_k$. Since $\mathbb{R}$ is complete and since the lengths of the intervals $\overline{I_k}$ converge to zero, there is precisely one point $r \in \mathbb{R}$ lying in all $\overline{I_k}$ (where $\overline{I_k}$ is the closure of I_k in the usual topology of $\mathbb{R}$). By choosing I_k to lie in the right half of I_{k-1} for odd k and in the left half for even k, we ensure that $r \notin J$. Thus, r lies in $\overline{I_k} \setminus J = I_k \subseteq N_k$ for every k. Since $I_1 \subseteq I$ this proves (6.11).

Now (ii) follows directly from (6.11), as the latter implies that no thin sum can have infinitely many non-zero summands.

Let us deduce an easy corollary of (6.11) that we will use in order to prove (iii):

$$\begin{array}{l}\textit{if } \mathcal{L} \textit{ is an infinite subset of } \mathcal{N} \textit{ and if } I \in \mathcal{I}_n \textit{ for some } n \in \mathbb{N} \\ \textit{then there exists an } s \in I \textit{ that is missed by infinitely many} \\ L \in \mathcal{L}.\end{array} \tag{6.12}$$

Indeed, this follows immediately if we apply (6.11) to $\mathcal{L}' := \{([0,1] \setminus J) \setminus L : L \in \mathcal{L}\}$. (Note that $\mathcal{L}' \subseteq \mathcal{N}$ as $\mathcal{L}' = \{L + S_0 : L \in \mathcal{L}\}$.)

To prove (iii), take any infinite subset $\mathcal{L}_0$ of $\mathcal{N}$, and choose an $I_1 \in \mathcal{I}_1$. By subsequent application of (6.11) and (6.12) we find $r_1, s_1 \in I_1$ and an infinite subset $\mathcal{L}_1$ of $\mathcal{L}$ so that $r_1 \in L$ but $s_1 \notin L$ for all $L \in \mathcal{L}_1$. Pick any element of $\mathcal{L}_1$ and denote it by L_1. Next, we choose some $I_2 \in \mathcal{I}_2$ and find, again, $r_2, s_2 \in I_2$ and an infinite subset $\mathcal{L}_2$ of $\mathcal{L}_1$ so that r_2 lies in all the $L \in \mathcal{L}_2$ and s_2 in none. Pick any $L_2 \in \mathcal{L}_2$ and continue in this manner.

This process yields a sequence $L_1, L_2, \ldots$ of elements of $\mathcal{N}$. By Tychonoff's theorem, the space $\mathbb{Z}_2^{[0,1]}$ is compact. Hence, the sequence $L_1, L_2, \ldots$ has an accumulation point $X \in \mathbb{Z}_2^{[0,1]}$, and clearly $X \in \overline{\mathcal{N}}$.

We claim that $X \notin \mathcal{N}$. To see this, first note that since this is the case for almost all L_k, $r_i \in X$ but $s_i \notin X$ for any i. Now, suppose that $X \in \mathcal{N}_n$ for some n. As $I_n \in \mathcal{I}_n$, it follows from (6.9) that either $I_n \subseteq X$ or $I_n \cap X = \emptyset$. However, the former contradicts $s_n \notin X$ and the latter contradicts $r_n \in X$. This completes the proof of (iii).

Chapter 7

Degree constrained orientations in countable graphs

7.1 Degree constrained orientations

Orientations of finite graphs are well-studied. An early result is the theorem of Robbins [73] on the existence of a strongly connected orientation. This result has been widely generalised by Nash-Williams [62] in 1960, who described orientations satisfying global or (symmetric) local edge-connectivity requirements. Ford and Fulkerson [40] investigated when a partial orientation can be completed to a di-eulerian one. As a last example, let us cite Frank [41] who characterised the graphs that can be oriented in such a way that there are k directed paths between a specified vertex and every other vertex.

In contrast, not much is known about orientations of infinite graphs. An exception is an old result of Egyed [39] that extends Robbins' theorem on strongly connected orientations. We mention also Thomassen [80] who raised some related conjectures.

In this chapter, which is based on [12], we will focus on degree constrained orientations in infinite (but countable) graphs. These are orientations where the in-degree function, i.e. the function counting the number of ingoing edges at each vertex, satisfies given lower and upper bounds. Degree constrained orientations have a close relationship to Hall's marriage theorem, and are also used by Berg and Jordán [10] in the context of graph rigidity. For finite graphs, Frank and Gyárfás [43] gave a necessary and sufficient condition for the existence of a degree constrained orientation. In infinite graphs, however, their condition is no longer sufficient. By strengthening the Frank-Gyárfás condition we will recover sufficiency while maintaining necessity.

Let us briefly recall some standard notation. For subsets U, W of the vertex set of a graph $G = (V, E)$ denote by $i_G(U)$ the number of edges in G having both endvertices in U and by $d_G(U, W)$ the number of edges in G with one endvertex in $U \setminus W$ and the other in $W \setminus U$. For a directed graph $\vec{G}$ and $X \subseteq V(\vec{G})$ let $\rho_{\vec{G}}(X)$ (resp. $\delta_{\vec{G}}(X)$) denote the number of edges entering (resp. leaving) the set X. If x is a vertex, we write $\rho_{\vec{G}}(x)$ instead of $\rho_{\vec{G}}(\{x\})$, and if no confusion can arise we will omit the subscripts G and $\vec{G}$. For a function $m : V \to \mathbb{R}$ and $X \subseteq V$ we will use the notation $m(X)$ to mean $\sum_{x \in X} m(x)$. Unfortunately, this notation is slightly inconsistent, in so far as $\rho_{\vec{G}}(X)$ is, in general, not the same as $\sum_{x \in X} \rho_{\vec{G}}(x)$.

Theorem 7.1 (Frank and Gyárfás [43])**.** *Let $G = (V, E)$ be a finite graph, and let $l, u : V(G) \to \mathbb{Z}$ be such that $l(v) \leq u(v)$ for all $v \in V$. Then*

(i) *there exists an orientation $\vec{G}$ of G such that $l(v) \leq \rho_{\vec{G}}(v) \leq u(v)$ for each vertex v if and only if*

(ii) *$l(X) \leq i(X) + d(X, V \setminus X)$ and $u(X) \geq i(X)$ for all $X \subseteq V(G)$.*

For a proof see also Frank [42].

The result carries over to locally finite graphs by an easy compactness argument, i.e. using Kőnig's Lemma 4.6. For non-locally finite graphs, however, the condition (ii) is too weak for the lower bound, as can be seen by considering an infinite star and setting $l \equiv 1$. There is no orientation satisfying the lower bounds while (ii) clearly holds.

Before we look at this example in more depth, let us rephrase Theorem 7.1. If we define the *surplus* to be $s(X) = i(X) + d(X, V \setminus X) - l(X)$ for a graph $G = (V, E)$ and a set $X \subseteq V$, then the theorem states that there is an orientation satisfying the lower bounds if and only if there is no set of negative surplus. Our aim is to find a condition in this vein.

Compare the infinite star with a finite star with the same lower bound of 1 everywhere. The whole finite star has negative surplus of -1, showing that there is no orientation satisfying the lower bound. Instead of computing this surplus directly let us do it in two steps. First, we observe that the set L of all leaves has surplus $s(L) = 0$. Now, if we add the centre c to L we do not gain any new edges since every edge is already incident with a leaf but since $l(c) > 0$ the demand for ingoing edges increases. Hence, $L \cup \{c\}$ has negative surplus.

Let us try to do the same for the infinite star. We immediately encounter the problem that the set L of all leaves is incident with infinitely many edges but has infinite demand for ingoing edges, i.e. $l(L) = \infty$. This results in $s(L) = \infty - \infty$, for which it is not clear which value this should be. So, let us compute the surplus of L in a similar stepwise fashion as above. Indeed, enumerate the leaves of the infinite star and denote by L_n the set of the first n leaves, which then has surplus 0. As L is the limit of the sets L_n it seems justified to define the surplus of L as the limit of $s(L_n)$, which therefore yields 0. Now, adding the centre c to L we see as for the finite star that the set $L \cup \{c\}$ has negative surplus. Consequently, the set $L \cup \{c\}$ is a witness for the non-existence of an orientation respecting the lower bounds.

We will now turn the ad hoc reasoning in the preceding paragraph into a formal condition. Fix a graph $G = (V, E)$, and for an ordinal number θ call a family $\mathcal{U}_\theta := (U_\mu)_{\mu \le \theta}$ of subsets of V a *queue in* G if

- $U_0 = \emptyset$;
- $U_\mu \subseteq U_\lambda$ for all $\mu \le \lambda \le \theta$;
- $U_\lambda = \bigcup_{\mu < \lambda} U_\mu$ for each limit ordinal $\lambda \le \theta$.

We write $\mathcal{U}_\lambda$ for the initial segment up to λ of $\mathcal{U}_\theta$, i.e. $\mathcal{U}_\lambda = (U_\mu)_{\mu \le \lambda}$.

Let $l : V \to \mathbb{Z}$ be a non-negative function, and let $\mathcal{U}_\theta = (U_\lambda)_{\lambda \le \theta}$ be a queue in G. Putting $\eta(\mathcal{U}_0, l) = 0$, we define by transfinite induction a function η such that

$$\eta(\mathcal{U}_{\lambda+1}, l) = \eta(\mathcal{U}_\lambda, l) + i(U_{\lambda+1} \setminus U_\lambda) + d(U_{\lambda+1} \setminus U_\lambda, V \setminus U_{\lambda+1}) - l(U_{\lambda+1} \setminus U_\lambda)$$

and such that $\eta(\mathcal{U}_\lambda, l) = \liminf_{\mu < \lambda} \eta(\mathcal{U}_\mu, l)$ for limit ordinals λ. In the computation of η we might need to calculate with ∞; we use the convention that $\infty - \infty = \infty$. Sometimes, if confusion can arise, we will write η_G to specify the underlying graph. We remark that for a finite vertex set the η-function provides merely an overly complicated way of computing its surplus. For infinite sets, however, η can be seen as a refinement of the surplus.

A set $U \subseteq V$ will be called *l-deficient* (or simply *deficient* if l is clear from the context) if there exists a queue $\mathcal{U}_\theta = (U_\lambda)_{\lambda \le \theta}$ with $U = U_\theta$ and $\eta(\mathcal{U}_\theta, l) < 0$. Deficient sets will play the same role as sets of negative surplus in the finite case.

We can now state the main result of this chapter, which we will prove in the next section:

Theorem 7.2. **[12]** *Let $G = (V, E)$ be a countable graph, and let $l, u : V \to \mathbb{Z} \cup \{\infty\}$ be non-negative functions with $l \le u$. Then the following statements are equivalent:*

(i) *there exists an orientation $\vec{G}$ of G such that $l(v) \le \rho_{\vec{G}}(v) \le u(v)$ for each vertex v; and*

(ii) *there are no l-deficient sets and $u(X) \geq i(X)$ for all finite $X \subseteq V(G)$.*

We mention that the theorem is very much in spirit of [64], in which Nash-Williams extends Hall's marriage theorem to countable graphs. This is perhaps not at all surprising since for finite graphs Theorem 7.1 can be reduced to the marriage theorem and vice versa. For infinite graphs, there are several versions of Hall's theorem. From the one in [64] one can indeed obtain our main result. However, as our proof is not a simple translation of Nash-Williams' arguments and as the reduction of Theorem 7.2 to Nash-Williams' theorem is not at all immediate (it takes about two pages), we see merit in providing a direct proof.

Nash-Williams' idea to refine a finite condition by using transfinite sequences is also used in Wojciechowski [93], who investigates when an infinite family of matroids on the same ground set has a system of disjoint bases. We will see this approach again in Section 8.9.

7.2 Proof of main result

In this section G will always denote a countable graph with vertex set V and edge set E, and $l, u : V \to \mathbb{Z} \cup \{\infty\}$ will always be non-negative functions such that $l \leq u$.

We shall prove Theorem 7.2 in the course of this section. Let us start with the observation that the function η satisfies a submodularity-type inequality (a set function $b : 2^V \to \mathbb{R}$ is called submodular if $b(X) + b(Y) \geq b(X \cap Y) + b(X \cup Y)$ for any $X, Y \subseteq V$). More precisely, it is easy to see that for the surplus function it holds that

$$s(U) + s(W) = s(U \cup W) + s(U \cap W) + d(U \setminus W, W \setminus U),$$

where U, W are vertex sets. (Clearly, this implies that s is submodular.) The lemma below states that a similar relation is true for η.

For a set $X \subseteq V$, we will denote its complement $V \setminus X$ by $\overline{X}$ if the base set V is clear from the context. To ease notation further, we will say that $\mathcal{U} = (U_\lambda)_{\lambda \leq \theta}$ *is a queue for* a set U if $U = U_\theta$. For a successor ordinal λ, we write (slightly abusing notation) $\lambda - 1$ for the ordinal μ for which $\lambda = \mu + 1$. We also introduce the notation $U'_\lambda := U_\lambda \setminus U_{\lambda-1}$.

Lemma 7.3. **[12]** *Let $\mathcal{U} = (U_\lambda)_{\lambda \leq \theta}$ be a queue for U and $\mathcal{W} = (W_\lambda)_{\lambda \leq \kappa}$ be a queue for W. Define queues $\mathcal{X} = (X_\lambda)_{\lambda \leq \kappa}$ with $X_\lambda = U \cap W_\lambda$ for every $\lambda \leq \kappa$ and $\mathcal{Y} = (Y_\lambda)_{\lambda \leq \theta+\kappa}$ with $Y_\lambda = U_\lambda$ for $\lambda \leq \theta$ and $Y_{\theta+\lambda} = U \cup W_\lambda$ for $\lambda \leq \kappa$. Then*

$$\eta(\mathcal{U}, l) + \eta(\mathcal{W}, l) \geq \eta(\mathcal{X}, l) + \eta(\mathcal{Y}, l) + d(W \setminus U, U \setminus W).$$

Proof. We shall show that

$$\eta(\mathcal{U}, l) + \eta(\mathcal{W}_\lambda, l) \geq \eta(\mathcal{X}_\lambda, l) + \eta(\mathcal{Y}_{\theta+\lambda}, l) + d(W_\lambda \setminus U, U \setminus W_\lambda) \tag{7.1}$$

for all $\lambda \leq \kappa$, which will give the statement with $\lambda = \kappa$.

We have $\eta(\mathcal{W}_0, l) = \eta(\mathcal{X}_0, l) = 0$, $\eta(\mathcal{Y}_\theta, l) = \eta(\mathcal{U}, l)$ and $d(W_0 \setminus U, U \setminus W_0) = 0$ since $W_0 = \emptyset$. Therefore, (7.1) holds with equality for $\lambda = 0$. We proceed by transfinite induction. Let λ be the smallest ordinal for which (7.1) is not yet shown.

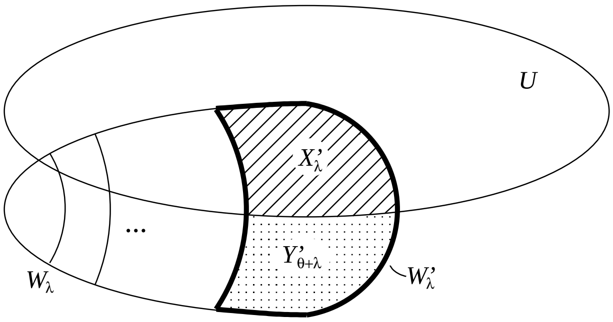

Figure 7.1: Relevant sets in Lemma 7.3

First, assume λ to be a successor ordinal. Observe that

$$d_x := d(X'_\lambda, \overline{X_\lambda}) \quad = \quad d(X'_\lambda, W_\lambda \setminus U) + d(X'_\lambda, \overline{W_\lambda}).$$

and

$$d_y := d(Y'_{\theta+\lambda}, \overline{Y_{\theta+\lambda}}) \quad = \quad d(W'_\lambda \setminus U, \overline{U \cup W_\lambda}).$$

We use these two relations in what follows:

$$\begin{aligned} & d(X'_\lambda, W_\lambda \setminus U) + d(W'_\lambda, \overline{W_\lambda}) \\ & \qquad = \quad d(X'_\lambda, W_\lambda \setminus U) + d(Y'_{\theta+\lambda}, \overline{Y_{\theta+\lambda}}) \\ & \qquad \qquad + d(Y'_{\theta+\lambda}, U \setminus W_\lambda) + d(X'_\lambda, \overline{W_\lambda}) \\ & \qquad = \quad d_x + d_y + d(Y'_{\theta+\lambda}, U \setminus W_\lambda) \end{aligned}$$

Noting that

$$i(W'_\lambda) = i(X'_\lambda) + i(Y'_{\theta+\lambda}) + d(X'_\lambda, Y'_{\theta+\lambda}),$$

and that

$$d(X'_\lambda, W_\lambda \setminus U) = d(X'_\lambda, Y'_{\theta+\lambda}) + d(X'_\lambda, W_{\lambda-1} \setminus U),$$

we obtain

$$\begin{aligned} & d(X'_\lambda, W_{\lambda-1} \setminus U) + d(W'_\lambda, \overline{W_\lambda}) + i(W'_\lambda) \\ & \quad = \quad d_x + i(X'_\lambda) + d_y + i(Y'_{\theta+\lambda}) + d(Y'_{\theta+\lambda}, U \setminus W_\lambda) \end{aligned} \tag{7.2}$$

Using this and the induction hypothesis for $\lambda - 1$ we get

$$\begin{aligned} \eta(\mathcal{U}, l) + \eta(\mathcal{W}_\lambda, l) &= \eta(\mathcal{U}, l) + \eta(\mathcal{W}_{\lambda-1}, l) + d(W'_\lambda, \overline{W_\lambda}) + i(W'_\lambda) - l(W'_\lambda) \\ &\overset{(7.1)}{\geq} \eta(\mathcal{X}_{\lambda-1}, l) + \eta(\mathcal{Y}_{\theta+(\lambda-1)}, l) + d(W_{\lambda-1} \setminus U, U \setminus W_{\lambda-1}) \\ &\quad + d(W'_\lambda, \overline{W_\lambda}) + i(W'_\lambda) - l(W'_\lambda) \\ &= \eta(\mathcal{X}_{\lambda-1}, l) + \eta(\mathcal{Y}_{\theta+(\lambda-1)}, l) + d(W_{\lambda-1} \setminus U, U \setminus W_\lambda) \\ &\quad + d(W_{\lambda-1} \setminus U, X'_\lambda) + d(W'_\lambda, \overline{W_\lambda}) + i(W'_\lambda) - l(W'_\lambda) \\ &\overset{(7.2)}{=} \eta(\mathcal{X}_{\lambda-1}, l) + d_x + i(X'_\lambda) - l(X'_\lambda) \\ &\quad + \eta(\mathcal{Y}_{\theta+(\lambda-1)}, l) + d_y + i(Y'_{\theta+\lambda}) - l(Y'_{\theta+\lambda}) \\ &\quad + d(W_{\lambda-1} \setminus U, U \setminus W_\lambda) + d(Y'_{\theta+\lambda}, U \setminus W_\lambda) \\ &= \eta(\mathcal{X}_\lambda, l) + \eta(\mathcal{Y}_{\theta+\lambda}, l) + d(W_\lambda \setminus U, U \setminus W_\lambda). \end{aligned}$$

This proves the induction step when λ is a successor ordinal.

Second, assume that λ is a limit ordinal. Then

$$\begin{aligned} \eta(\mathcal{U}, l) + \eta(\mathcal{W}_\lambda, l) &= \eta(\mathcal{U}, l) + \liminf_{\mu<\lambda} \eta(\mathcal{W}_\mu, l) \\ &\geq \liminf_{\mu<\lambda} (\eta(\mathcal{X}_\mu, l) + \eta(\mathcal{Y}_{\theta+\mu}, l) + d(W_\mu \setminus U, U \setminus W_\mu)) \\ &\geq \eta(\mathcal{X}_\lambda, l) + \eta(\mathcal{Y}_{\theta+\lambda}, l) + \liminf_{\mu<\lambda} d(W_\mu \setminus U, U \setminus W_\mu). \end{aligned}$$

Furthermore, for any $\mu < \lambda$ we get

$$\begin{aligned} d(W_\mu \setminus U, U \setminus W_\mu) &= d(W_\mu \setminus U, U \setminus W_\lambda) + d(W_\mu \setminus U, U \cap (W_\lambda \setminus W_\mu)) \\ &\geq d(W_\mu \setminus U, U \setminus W_\lambda). \end{aligned}$$

It is easy to see that $\liminf_{\mu<\lambda} d(W_\mu \setminus U, U \setminus W_\lambda) = d(W_\lambda \setminus U, U \setminus W_\lambda)$ since $W_\lambda = \bigcup_{\mu<\lambda} W_\mu$. Putting all this together we obtain (7.1). □

We call a vertex set U *l-tight* if (it is not l-deficient and) there exists a queue $(U_\lambda)_{\lambda \leq \theta}$ for U with $\eta(\mathcal{U}_\theta, l) = 0$. If it is clear in regard to which function l a set is tight, we will suppress the l. Tight sets are the most critical sets, and it can be seen that in an orientation respecting the lower bound l there can be no edge leaving a tight set.

Lemma 7.3 immediately implies that the intersection and the union of two tight sets is tight, too. We will need a little bit more, namely that this also holds for the union of countably many tight sets:

Lemma 7.4. [12] *Assume that there are no deficient sets in G, and let $U_1, U_2, \ldots$ be countably many tight sets. Then also their union is tight.*

Proof. Let $\mathcal{U}^i = (U^i_\lambda)_{\lambda \leq \theta_i}$ be queues witnessing the tightness of U_i for each i, i.e. $\eta(\mathcal{U}^i, l) = 0$ and $U^i_{\theta_i} = U_i$. For any $n \in \mathbb{N}$, set $\kappa_n = \sum_{i=1}^n \theta_i$ and $\kappa = \sum_{i=1}^\infty \theta_i$. Then we can define the queue $\mathcal{Y} = (Y_\lambda)_{\lambda \leq \kappa}$ with $Y_{\kappa_{n-1}+\lambda} = Y_{\kappa_{n-1}} \cup U^n_\lambda$ if $\lambda \leq \theta_n$ (where $\kappa_0 = 0$ and $Y_0 = \emptyset$) and $Y_\kappa = \bigcup_{n=1}^\infty Y_{\kappa_n}$. By Lemma 7.3 and induction we get $\eta(\mathcal{Y}_{\kappa_n}, l) = 0$ for all $n \geq 0$:

$$\eta(\mathcal{Y}_{\kappa_n}, l) \leq \eta(\mathcal{Y}_{\kappa_{n-1}}, l) + \eta(\mathcal{U}^n, l) = 0 + 0.$$

(Note, that there are no deficient sets.) From this it follows that $\eta(\mathcal{Y}, l) = \liminf_{\lambda < \kappa} \eta(\mathcal{Y}_\lambda, l) \leq 0$. Again, as there are no deficient sets, this implies that $Y_\kappa = \bigcup_{i=1}^\infty U_i$ is tight. □

As for the lower bound we will define deficiency and tightness of sets with respect to the upper bound, too. We call a finite vertex set X *u-faulty*, if $u(X) - i(X) < 0$, and we call it *u-taut* if $u(X) - i(X) = 0$. Again, if u is clear from the context, we will omit it.

In the last lemma we saw that the union of tight sets is tight. In contrast, for taut sets we will need that their intersection is taut:

Lemma 7.5. [12] *If there are no faulty sets in G then the following is true:*

(i) *if X and Y are two taut sets then $X \cap Y$ is taut and there is no edge between $X \setminus Y$ and $Y \setminus X$; and*

(ii) *the intersection of arbitrarily many taut sets is taut.*

Proof. (i) On the one hand, we get

$$i(X) + i(Y) = u(X) + u(Y) = u(X \cup Y) + u(X \cap Y) \geq i(X \cup Y) + i(X \cap Y)$$

and on the other hand, i is supermodular, i.e. it holds that:

$$i(X) + i(Y) \leq i(X \cup Y) + i(X \cap Y).$$

Thus, we have equality everywhere. In particular, if there was an edge between $X \setminus Y$ and $Y \setminus X$ then $i(X) + i(Y) < i(X \cup Y) + i(X \cap Y)$, which is not the case.

(ii) Let X_i, $i \in I$ be taut sets. Since by definition each of the X_i is finite, their intersection is also finite. Hence, there are already finitely many X_j, $j \in J \subseteq I$ with $\bigcap_{i \in I} X_i = \bigcap_{j \in J} X_j$. Therefore, we only need to check that the intersection of two taut sets is taut, which is true by (i). □

In Theorem 7.2 (ii) the conditions regarding the lower and the upper bound are independent of each other. The following lemma provides a link between tight and taut sets.

Lemma 7.6. **[12]** *Let there be neither deficient sets nor faulty sets in* G*, and let* U *be a taut set and* L *be a tight set. Then* $U \setminus L$ *is taut and* $L \setminus U$ *is tight.*

Proof. Let $\mathcal{L} = (L_\lambda)_{\lambda \leq \theta}$ be a queue with $\eta(\mathcal{L}, l) = 0$ and $L_\theta = L$, and define $\mathcal{M} = (L_\lambda \setminus U)_{\lambda \leq \theta}$. By transfinite induction, we show that for any ordinal $\lambda \leq \theta$ it holds that

$$\eta(\mathcal{L}_\lambda, l) \geq \eta(\mathcal{M}_\lambda, l) + i(L_\lambda \cap U) - l(L_\lambda \cap U) + d(L_\lambda \cap U, \overline{L_\lambda}). \tag{7.3}$$

This is trivially true for $\lambda = 0$. Let λ be such that the induction hypothesis holds for all $\mu < \lambda$. First, assume that λ is a successor ordinal. We use the induction hypothesis for $\lambda - 1$ in what follows:

$$\begin{aligned}
\eta(\mathcal{L}_\lambda, l) &= \eta(\mathcal{L}_{\lambda-1}, l) + i(L'_\lambda) + d(L'_\lambda, \overline{L_\lambda}) - l(L'_\lambda) \\
&\overset{(7.3)}{\geq} \eta(\mathcal{M}_{\lambda-1}, l) + i(L_{\lambda-1} \cap U) - l(L_{\lambda-1} \cap U) \\
&\quad + d(L_{\lambda-1} \cap U, \overline{L_{\lambda-1}}) + i(L'_\lambda) + d(L'_\lambda, \overline{L_\lambda}) - l(L'_\lambda) \\
&= \eta(\mathcal{M}_{\lambda-1}, l) + i(L_{\lambda-1} \cap U) + d(L_{\lambda-1} \cap U, \overline{L_{\lambda-1}} \setminus M'_\lambda) \\
&\quad d(L_{\lambda-1} \cap U, M'_\lambda) + i(L'_\lambda) + d(L'_\lambda, \overline{L_\lambda}) \\
&\quad -l(L_\lambda \cap U) - l(M'_\lambda)
\end{aligned}$$

With

$$\begin{aligned}
&d(L_{\lambda-1} \cap U, M'_\lambda) + i(L'_\lambda) + d(L'_\lambda, \overline{L_\lambda}) \\
&\quad = d(L_{\lambda-1} \cap U, M'_\lambda) + i(L'_\lambda \cap U) + d(L'_\lambda \cap U, M'_\lambda) \\
&\qquad + i(M'_\lambda) + d(L'_\lambda \cap U, \overline{L_\lambda}) + d(M'_\lambda, \overline{L_\lambda}) \\
&\quad = i(M'_\lambda) + d(M'_\lambda, \overline{M_\lambda}) + i(L'_\lambda \cap U) + d(L'_\lambda \cap U, \overline{L_\lambda})
\end{aligned} \tag{7.4}$$

we get

$$\begin{aligned}
\eta(\mathcal{L}_\lambda, l) &\overset{(7.4)}{\geq} \eta(\mathcal{M}_{\lambda-1}, l) + i(L_{\lambda-1} \cap U) + d(L_{\lambda-1} \cap U, \overline{L_{\lambda-1}} \setminus M'_\lambda) \\
&\quad i(M'_\lambda) + d(M'_\lambda, \overline{M_\lambda}) + i(L'_\lambda \cap U) + d(L'_\lambda \cap U, \overline{L_\lambda}) \\
&\quad -l(L_\lambda \cap U) - l(M'_\lambda) \\
&= \eta(\mathcal{M}_\lambda, l) + i(L_{\lambda-1} \cap U) + d(L_{\lambda-1} \cap U, \overline{L_{\lambda-1}} \setminus M'_\lambda) \\
&\quad +i(L'_\lambda \cap U) + d(L'_\lambda \cap U, \overline{L_\lambda}) - l(L_\lambda \cap U) \\
&= \eta(\mathcal{M}_\lambda, l) + i(L_{\lambda-1} \cap U) + d(L_{\lambda-1} \cap U, \overline{L_\lambda}) \\
&\quad +d(L_{\lambda-1} \cap U, L'_\lambda \cap U) + i(L'_\lambda \cap U) + d(L'_\lambda \cap U, \overline{L_\lambda}) - l(L_\lambda \cap U) \\
&= \eta(\mathcal{M}_\lambda, l) + i(L_\lambda \cap U) + d(L_{\lambda-1} \cap U, \overline{L_\lambda}) \\
&\quad +d(L'_\lambda \cap U, \overline{L_\lambda}) - l(L_\lambda \cap U) \\
&= \eta(\mathcal{M}_\lambda, l) + i(L_\lambda \cap U) - l(L_\lambda \cap U) + d(L_\lambda \cap U, \overline{L_\lambda})
\end{aligned}$$

So, let λ be a limit ordinal. Then observe that $\liminf_{\mu \leq \lambda} d(L_\mu \cap U, \overline{L_\mu}) = d(L_\lambda \cap U, \overline{L_\lambda})$ as U is finite. Furthermore, $l(L_\mu \cap U)$ is bounded for the same reason. Thus

$$\begin{aligned}
\eta(\mathcal{L}_\lambda, l) &\geq \liminf_{\mu \leq \lambda} \big(\eta(\mathcal{M}_\mu, l) + i(L_\mu \cap U) - l(L_\mu \cap U) + d(L_\mu \cap U, \overline{L_\mu})\big) \\
&\geq \eta(\mathcal{M}_\lambda, l) + i(L_\lambda \cap U) - l(L_\lambda \cap U) + d(L_\lambda \cap U, \overline{L_\lambda}).
\end{aligned}$$

Now, for $\lambda = \theta$ this yields

$$\begin{aligned}
0 &= \eta(\mathcal{L}, l) + u(U) - i(U) \\
&\geq \eta(\mathcal{M}, l) + i(L \cap U) - l(L \cap U) + u(U) - i(U) + d(L \cap U, \overline{L}) \\
&\geq \eta(\mathcal{M}, l) + u(U \setminus L) - i(U \setminus L) + (u - l)(L \cap U)
\end{aligned}$$

Since $\eta(\mathcal{M}, l) \geq 0$, $u \geq l$ and since $u(U \setminus L) \geq i(U \setminus L)$ it follows that $U \setminus L$ is taut. This then also implies that $\eta(\mathcal{M}, l) = 0$, and hence $L \setminus U$ is tight. □

Lemma 7.7. **[12]** *Let there be no u-faulty sets. Assume that for an edge e with endvertices v, w there is no u-taut set U with $v \in U$ but $w \notin U$. Then, setting $u'(x) = u(x)$ for all vertices $x \neq v$ and $u'(v) = u(v) - 1$, there are no u'-faulty sets in $G - e$.*

Proof. If $u(v) = \infty$ then for every set $X \subseteq V$ we get $u(X) = \infty$, and thus there cannot be any u'-faulty set in $G - e$. So, let $u(v) < \infty$, and suppose U is u'-faulty in $G - e$. Clearly, $v \in U$ but $w \notin U$ since there are no u-faulty sets in G. But then U is u-taut in G, a contradiction. □

We can finally prove our main result:

Proof of Theorem 7.2. (ii)⇒(i) Let $v_1, v_2, \ldots$ be a sequence of the vertices of G such that every vertex v appears exactly $l(v)$ times in it. Putting $l_0 = l$ and $u_0 = u$ we recursively

(a) set $l_n(v) = l_{n-1}(v)$ if $v \neq v_n$ and $l_n(v_n) = l_{n-1}(v_n) - 1$;

(b) set $u_n(v) = u_{n-1}(v)$ if $v \neq v_n$ and $u_n(v_n) = u_{n-1}(v_n) - 1$; and

(c) find distinct edges $e_1, e_2, \ldots$ such that $G_n := G - \{e_1, \ldots, e_n\}$ has no l_n-deficient and no u_n-faulty sets and such that e_n is incident with v_n.

Assume that this has been achieved for $i < n$. It is not difficult to check directly that picking any loop at v_n for e_n we satisfy (a)–(c). However, if we agree that v_n is a neighbour of itself if there is a loop at v_n then what follows covers also loops.

For each neighbour w of v_n in G_{n-1} for which this is possible pick an l_{n-1}-tight set X with $w \in X$ but $v_n \notin X$, and consider the union L of these sets. By Lemma 7.4, L is still l_{n-1}-tight. In a similar way, consider a minimal u_{n-1}-taut set U that contains v_n (where we set $U = \emptyset$ if there is no such set). From Lemma 7.5 (ii) it follows that for a neighbour w of v_n in G_{n-1} for which there is an u_{n-1}-taut set Y with $v_n \in Y$ but $w \notin Y$ it holds that $w \notin U$. By Lemma 7.6, $U \setminus L$ is u_{n-1}-taut, too. As U is minimal this implies $U = U \setminus L$ and therefore that U and L are disjoint.

Next, if $U \neq \emptyset$ then there is a neighbour w_n of v_n in G_{n-1} with $w_n \in U$. For if that was not the case, then, recalling that $u(v_n) > 0$ by definition of v_n, we would have

$$i(U \setminus \{v_n\}) = i(U) = u(U) > u(U \setminus \{v_n\}),$$

which is a contradiction, as there are no u_{n-1}-faulty sets. Note that $w_n \notin L$ since U and L are disjoint. If, on the other hand, $U = \emptyset$ then there is a neighbour $w_n \notin L$ of v_n in G_{n-1}. Indeed, suppose not. Let $\mathcal{L} := (L_\lambda)_{\lambda \leq \theta}$ be a queue with $\eta(\mathcal{L}, l_{n-1}) = 0$ and $L_\theta = L$ (recall, that L is l_{n-1}-tight). Put $L_{\theta+1} = L \cup \{v_n\}$, and observe that $i(L_{\theta+1} \setminus L_\theta) = 0$ (since there is no loop at v_n) and $d_{G_{n-1}}(L_{\theta+1} \setminus L_\theta, V(G_{n-1}) \setminus L_{\theta+1}) = 0$. Thus, $\eta(\mathcal{L}_{\theta+1}, l_{n-1}) = -l_{n-1}(v_n) < 0$, (from the definition of our sequence $v_1, v_2, \ldots$ it follows that $l_{n-1}(v_n) > 0$). Consequently, $L_{\theta+1}$ is l_{n-1}-deficient, contrary to our induction hypothesis. In any case, let e_n be any edge between v_n and w_n and observe that, by Lemma 7.7, there are no u_n-faulty sets in G_n. In addition, by construction of L and because of $w_n \notin L$ we get

$$\textit{there is no } l_{n-1}\textit{-tight set } X \textit{ in } G_{n-1} \textit{ with } w_n \in X \textit{ but } v_n \notin X \tag{7.5}$$

Let us check that there are also no l_n-deficient sets in G_n. So, suppose there is a l_n-deficient set M in G_n, and let $\mathcal{M}_\theta = (U_\lambda)_{\lambda\leq\theta}$ be a queue in G_n with $M = M_\theta$ and $\eta_{G_n}(\mathcal{M}_\theta, l_n) < 0$. Since G_n differs from G_{n-1} only in the edge e_n we get, if neither $v_n \in M$ nor $w_n \in M$, that $\eta_{G_{n-1}}(\mathcal{M}_\theta, l_{n-1}) = \eta_{G_n}(\mathcal{M}_\theta, l_n)$, which is impossible since M is not l_{n-1}-deficient in G_{n-1}. In a similar way, if $l_{n-1}(v_n) < \infty$ we can exclude the case when $v_n \in M$ as we lose an edge but also have less demand of ingoing edges. If, on the other hand, $l_{n-1}(v_n) = \infty$ we can get rid of this case, too: Denote by λ the smallest ordinal for which $v_n \in M_\lambda$, which is a successor ordinal, by definition of a queue. Then as $\eta_{G_{n-1}}(\mathcal{M}_\lambda, l_{n-1}) \geq 0$ and as

$$\eta_{G_{n-1}}(\mathcal{M}_\lambda, l_{n-1}) = \eta_{G_{n-1}}(\mathcal{M}_{\lambda-1}, l_{n-1}) + i(M'_\lambda) + d(M'_\lambda, \overline{M_\lambda}) - \infty$$

it follows that $\eta_{G_{n-1}}(\mathcal{M}_\lambda, l_{n-1}) = \infty$, and hence $\eta_{G_{n-1}}(\mathcal{M}_\theta, l_{n-1}) = \infty$. Because $\eta_{G_{n-1}}$ and η_{G_n} can differ by at most one, we obtain $\eta_{G_n}(\mathcal{M}_\theta, l_n) = \infty$, a contradiction.

Therefore, we may assume that $w_n \in M$ but $v_n \notin M$ (independent of the value of $l_{n-1}(v_n)$). Now, let λ be the smallest ordinal for which $w_n \in M_\lambda$, which is a successor ordinal. Then,

$$d_{G_{n-1}}(M_\lambda \setminus M_{\lambda-1}, \overline{M_\lambda}) = d_{G_n}(M_\lambda \setminus M_{\lambda-1}, \overline{M_\lambda}) + 1,$$

and thus $\eta_{G_{n-1}}(\mathcal{M}_\lambda, l_{n-1}) = \eta_{G_n}(\mathcal{M}_\lambda, l_n) + 1$ (since $v_n \notin M_\lambda$ implies that $l_{n-1}(M_\lambda \setminus M_{\lambda-1}) = l_n(M_\lambda \setminus M_{\lambda-1})$). Hence, $\eta_{G_{n-1}}(\mathcal{M}_\theta, l_{n-1}) = \eta_{G_n}(\mathcal{M}_\theta, l_n) + 1$. Now, since $\eta_{G_n}(\mathcal{M}_\theta, l_n) < 0$ but $\eta_{G_{n-1}}(\mathcal{M}_\theta, l_{n-1}) \geq 0$ we obtain that $\eta_{G_{n-1}}(\mathcal{M}_\theta, l_{n-1}) = 0$. Therefore, M is an l_{n-1}-tight set with $v_n \notin M$ but $w_n \in M$, contradicting (7.5). Thus, there are no l_n-deficient sets in G_n, as required.

Having terminated the transfinite induction, we put $G^0 = G - \{e_1, e_2, \ldots\}$. We think of each edge e_n as already directed towards v_n. In this way, each vertex v has an indegree of exactly $l(v)$ (by definition of the vertex enumeration). So, what remains is to direct the edges in G^0 in such a way, that the reduced upper bound $u^0 := u - l$ is respected.

First, let us show that there are no u^0-faulty sets in G^0. Indeed, consider a finite vertex set U in G^0. Then there is an N such that $u_N(U) = u^0(U)$ and $i_{G_N}(U) = i_{G^0}(U)$, and thus $u^0(U) \geq i_{G^0}(U)$ since $u_N(U) \geq i_{G_N}(U)$. As a u^0-faulty set is by definition finite, there is therefore no such set.

Second, let $f_1, f_2, \ldots$ be an enumeration of the edges of G^0. Denote the endvertices of f_1 by x and y, and observe that if there is a u^0-taut set X with $x \in X$ but $y \notin X$ then there is no u^0-taut set Y with $x \notin Y$ but $y \in Y$, by Lemma 7.5 (i).

Now, if there is such a set X, then direct f_1 towards y, and define $u^1(v) = u^0(v)$ for $v \neq y$ and $u^1(y) = u^0(y) - 1$. If not, direct f_1 in the other way, and define u^1 accordingly. Lemma 7.7 ensures that $G^1 = G^0 - f_1$ has no u^1-faulty sets. Continuing in this way, we obtain the desired orientation. Indeed, suppose a vertex v receives more ingoing edges than $u(v)$. Then there is an N such that $u^N(v) < 0$, which implies that $\{v\}$ is u^N-faulty, a contradiction.

(i)⇒(ii) Let $\vec{G}$ be an orientation as in (i). Then trivially $u(X) \geq \sum_{v\in X} \rho_{\vec{G}}(v) \geq i(X)$ holds for any finite set $X \subseteq V$. In order to prove that there is no l-deficient set, pick any queue $\mathcal{U}_\theta := (U_\lambda)_{\lambda\leq\theta}$. We will show by transfinite induction that $\eta(\mathcal{U}_\lambda, l) \geq \delta_{\vec{G}}(U_\lambda)$ for every $\lambda \leq \theta$. (Recall that $\delta_{\vec{G}}(U)$ denotes the number of edges leaving U.) This is true for $\lambda = 0$. Let λ be the smallest ordinal for which this is not yet shown.

First, let λ be a successor ordinal, and assume that $\sum_{v\in U'_\lambda} \rho_{\vec{G}}(v) < \infty$. Then

$$\begin{aligned}
\eta(\mathcal{U}_\lambda, l) &= \eta(\mathcal{U}_{\lambda-1}, l) + i(U'_\lambda) + d(U'_\lambda, \overline{U_\lambda}) - l(U'_\lambda) \\
&\geq \delta_{\vec{G}}(U_{\lambda-1}) + i(U'_\lambda) + d(U'_\lambda, \overline{U_\lambda}) - \sum_{v\in U'_\lambda} \rho_{\vec{G}}(v) \\
&= \delta_{\vec{G}}(U_{\lambda-1}) + d(U'_\lambda, \overline{U_\lambda}) - \rho_{\vec{G}}(U'_\lambda) = \delta_{\vec{G}}(U_\lambda).
\end{aligned}$$

If, on the other hand, $\sum_{v\in U'_\lambda} \rho_{\vec{G}}(v) = \infty$ then either there are infinitely many edges directed from $U_{\lambda-1}$ to U'_λ, in which case $\eta(\mathcal{U}_{\lambda-1}, l) \geq \delta_{\vec{G}}(U_{\lambda-1}) = \infty$, or $i(U'_\lambda) = \infty$, or there are infinitely many edges directed from $\overline{U_\lambda}$ towards U'_λ, which implies $d(U'_\lambda, \overline{U_\lambda}) = \infty$. In all of these cases we obtain

$$\eta(\mathcal{U}_\lambda, l) = \eta(\mathcal{U}_{\lambda-1}, l) + i(U'_\lambda) + d(U'_\lambda, \overline{U_\lambda}) - l(U'_\lambda) \geq \infty - \infty = \infty.$$

Next, let λ be a limit ordinal. Denoting by $A(X, Y)$ the edges directed from $X \subseteq V$ to $Y \subseteq V$ we obtain

$$\begin{aligned}
\eta(\mathcal{U}_\lambda, l) &= \liminf_{\mu<\lambda} \eta(\mathcal{U}_\mu, l) \geq \liminf_{\mu<\lambda} (\delta_{\vec{G}}(U_\mu)) \\
&\geq \liminf_{\mu<\lambda} |A(U_\mu, \overline{U_\lambda})| \quad = \quad \delta_{\vec{G}}(U_\lambda).
\end{aligned}$$

Finally, with $\lambda = \theta$ we get $\eta(\mathcal{U}_\theta, l) \geq \delta_{\vec{G}}(U_\theta) \geq 0$, as desired. □

7.3 Open questions

Let us formulate two directions for future research. First, Theorem 7.2 treats only countable graphs, and indeed our proof does not seem to be adaptable to higher cardinalities. On the other hand, we do not have any example showing that our condition fails in uncountable graphs.

Problem 7.8. *Can Theorem 7.2 be extended to uncountable graphs?*

Second, in finite graphs, Theorem 7.1 allows to impose lower bounds on the in-degree and the out-degree at the same time. Indeed, $d(v) - u(v)$ gives a lower bound on the out-degree of a vertex v. In contrast, for a vertex v of infinite degree we can only demand all or nothing. Setting $u(v)$ to a finite value in Theorem 7.2 is the same as requiring infinitely many outgoing edges at v, whereas putting $u(v) = \infty$ will not impose any restrictions on the out-degree at all. To regain a finer control, we propose the following conjecture:

Conjecture 7.9. [12] *Let G be a countable graph, and let $l, r : V(G) \to \mathbb{N} \cup \{\infty\}$ be two non-negative functions with $l(v) + r(v) \leq d(v)$ for all vertices v. Then*

(i) *there exists an orientation $\vec{G}$ of G such that $\rho_{\vec{G}}(v) \geq l(v)$ and $\delta_{\vec{G}}(v) \geq r(v)$ for each vertex v if and only if*

(ii) *there are no l-deficient sets and no r-deficient sets.*

Chapter 8

Connectivity in infinite matroids

8.1 Infinite matroids

The cycle matroid of a finite graph is the matroid in which those edge sets are independent that do not contain any circuit of the graph. In Chapters 2–4 we have seen that the introduction of infinite circuits allows more natural and smoother extensions of finite results. So, it appears conceivable that the infinite circuits together with the finite ones yield a matroid associated to an infinite graph. But how should then a matroid on an infinite set be defined?

The current entry on matroids in wikipedia [1] states "that there are many reasonable and useful definitions [of infinite matroids], none of which captures all the important aspects of finite matroid theory". Fortunately, the entry only demonstrates that while it is often a highly useful source of knowledge, wikipedia still needs to be taken with a grain, and sometimes with a spoonful, of salt. Indeed, there is such a notion of a matroid that features all the important properties of a finite matroid, namely the B-matroids of Higgs [52, 53, 54]. Among the properties of B-matroids are that every independent set in a B-matroid is contained in a basis; every dependent set contains a circuit; the dual of a B-matroid is well defined; and there exist contraction and deletion operations such that one is the dual operation of the other. We discuss the basic properties of B-matroids in more detail in Sections 8.2 and 8.4. See also Oxley [67] for a general discussion of B-matroids. A more recent result involving graphs and B-matroids is due to Christian, Richter and Rooney [28].

The main part of this chapter focuses on connectivity in B-matroids. A finite matroid is connected if and only if every two elements are contained in a common circuit. This definition can be extended verbatim to infinite matroids. Higher connectivity is less straightforward. A finite matroid M is k-connected unless there exists an ℓ-separation for some $\ell < k$, i.e. a partition (X, Y) of the ground set so that $r(X) + r(Y) - r(M) \leq \ell - 1$ and $|X|, |Y| \geq \ell$ where r is the rank function of the matroid. In an infinite matroid this definition is obviously useless as all of the involved ranks will usually be infinite. In Section 8.6 we will give an equivalent definition of k-connectivity that does generalise to infinite matroids. We will examine the circuit definition of connectivity in Section 8.5 and treat k-connectivity in more detail in Section 8.7.

As an application of our definition of higher connectivity we prove Tutte's linking theorem for an important class of B-matroids. Tutte's linking theorem is often said to be the analogue of Menger's theorem in matroids and concerns how well connected two subsets of the ground set are. In order to formulate this more precisely, we define for disjoint sets $X, Y \subseteq E(M)$ in a finite matroid M the *connectivity function* by $\kappa_M(X, Y) = \min\{r(U) + r(E - U) - r(E) : X \subseteq U \subseteq E - Y\}$. Tutte's linking theorem is then as follows:

Theorem 8.1 (Tutte [84])**.** *Let M be a finite matroid, and let X and Y be two disjoint subsets of $E(M)$. Then there exists a partition (C, D) of $E(M) - (X \cup Y)$ such that $\kappa_{M/C-D}(X, Y) = \kappa_M(X, Y)$.*

We conjecture that Theorem 8.1 holds also for arbitrary B-matroids.

Conjecture 8.2.[18] *Let M be a B-matroid, and let X and Y be two disjoint subsets of $E(M)$. Then there exists a partition (C, D) of $E(M) - (X \cup Y)$ such that $\kappa_{M/C-D}(X, Y) = \kappa_M(X, Y)$.*

In Section 8.8 we verify the conjecture for finitary matroids, i.e. B-matroids where every circuit is of finite size.

After a short interlude in Section 8.9 on matchability of B-matroids we will, in Sections 8.10 and 8.11, investigate the cycle matroid associated with a graph. We will conclude this thesis by calculating the connectivity function of the cycle matroid. This will allow us to revisit a result from Chapter 2, namely the invariance of Tutte-connectivity under taking duals.

8.2 Definition of B-matroids

Higgs defined B-matroids by giving a set of axioms for the closure operator. Oxley [65, 67] provided equivalent independence axioms that are far more accessible.

Let E be some set, and let $\mathcal{I}$ be a set of subsets of E. We call $M = (E, \mathcal{I})$ a *B-matroid* if the following conditions are satisfied:

(I1) $\{\emptyset\} \in \mathcal{I}$;

(I2) if $I' \subseteq I$ and $I \in \mathcal{I}$ then $I' \in \mathcal{I}$;

(IB1) for all $I \in \mathcal{I}$, if $I \subseteq X \subseteq E$ then there exists a maximal subset B of X with $B \in \mathcal{I}$ that contains I; and

(IB2) if B_1 and B_2 are maximal elements of $\mathcal{I}$ and $x \in B_1 - B_2$ then there exists a $y \in B_2 - B_1$ so that $(B_1 - x) + y$ is maximal in $\mathcal{I}$.

As usual, any set in $\mathcal{I}$ is called *independent* and any subset of E not in $\mathcal{I}$ is *dependent*.

The main feature of B-matroids is that they have bases and circuits while maintaining duality at the same time. The existence of bases, i.e. maximal independent sets, is guaranteed by (IB1). The *dual matroid* $M^* = (E, \mathcal{I}^*)$ of a matroid $M = (E, \mathcal{I})$ is defined by requiring that $J \subseteq E$ lies in $\mathcal{I}^*$ if and only if there is a basis B of M with $J \subseteq E - B$. Higgs [54] proved that M^* is a B-matroid if M is one. Finally, a minimal dependent set is a *circuit*, and a circuit of the dual matroid a *co-circuit*. Every dependent set contains a circuit, see Oxley [65].

B-matroids and finite matroids share more properties, among them the most natural ones, such as that the minor of a B-matroid is a B-matroid. Moreover, Oxley [67] demonstrated that B-matroids are the largest class that satisfy a small number of properties that a matroid should have. For these reason, we contend that B-matroids are the right generalisation of finite matroids to infinite ground sets. To emphasise that we shall from now on simply speak of *matroids* instead of B-matroids.

An important subset of matroids are the finitary matroids. A matroid is *finitary* if every of its circuits is finite, and a matroid is called *co-finitary* if its dual matroid is finitary.

Finitary matroids, who can alternatively be defined in terms of simple independence axioms, predate (B-)matroids. A natural and rich source of finitary matroids are graphs. Let G be a graph, and define $\mathcal{I}_{\mathrm{ST}}(G)$ to be the set of all edge sets I so that I is contained in a subgraph of G that is a spanning tree on each component of G. Let $M_{\mathrm{ST}}(G) = (E(G), \mathcal{I}_{\mathrm{ST}}(G))$.

Proposition 8.3 (Higgs[1] [54])**.** *$M_{ST}(G)$ is a finitary matroid for every graph G.*

[1] The attribution to Higgs is perhaps a bit of a stretch. It was known before, and in fact is quite simple to see, that $M_{\mathrm{ST}}(G)$ is finitary if the older definition via independence axioms is used. Higgs' contribution lies in recognising that finitary matroids are B-matroids.

Graphs give also rise to co-finitary matroids. With infinite (as well as finite) circuits we can define a co-finitary matroid M_{TST} in a similar way as M_{ST} above. We declare every edge set that does not contain a (finite or infinite) circuit to be independent. We discuss M_{TST} and its relationship to $M_{\mathrm{ST}}(G)$ in more depth in Section 8.10.

Let us consider an example. The double ladder, as shown in Figure 8.1, has two ends, one on the right and one on the left. Adding these ends produces new infinite circuits. For instance the edge set of the double ray R_1eR_2 becomes a circuit, as well as S_1eS_2. Moreover, the edge set of the union of the two disjoint double rays S_1R_1 and S_2R_2 is a circuit too. Consequently, a spanning tree as $S_1R_1 \cup S_2R_2 + e$, which is a basis in M_{ST}, is no longer a basis in M_{TST}. On the other hand, the disconnected set $(S_1R_1 - f) \cup S_2R_2$ can be seen to be maximally independent.

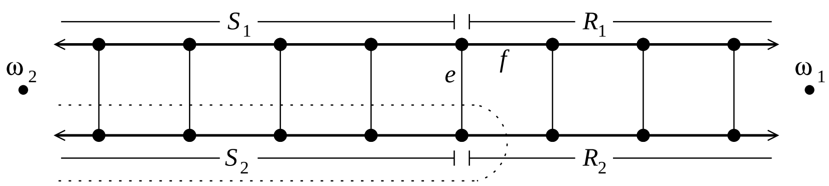

Figure 8.1: A co-finitary matroid on the double ladder

By only adding some of the ends of a graph we can obtain a matroid that is neither finitary nor co-finitary. If we delete the left end ω_2 in Figure 8.1, then the only infinite circuits that remain are edge sets of double rays with two disjoint subrays going towards ω_1, such as R_1eR_2. However, the double ray S_1eS_2 then becomes independent. An example for an infinite co-circuit is the set of edges with one endvertex on S_2 and the other outside (indicated by the dashed line). In this way we obtain a matroid that is neither finitary nor co-finitary.

8.3 A matroid on the ω-regular tree

To illustrate the concept of infinite matroids a bit further, let us give another example. We shall define a matroid MT_∞ on the edge set of the ω-regular tree T_∞. As the double ladder with only one end in the previous section MT_∞ will neither be finitary nor co-finitary. In fact, we shall see that MT_∞ has only infinite circuits and co-circuits.

Declare every subset of $E(T_\infty)$ to be independent if it does not contain (the edge set of) any double ray. Let us first verify that this constitutes a matroid. For this purpose we call an edge set F an *edge-double ray* if it is the edge set of a double ray.

Clearly, the axioms $(I0)$ and $(I1)$ are trivially satisfied for MT_∞. To verify $(IB1)$, consider an independent set I, i.e. a set of edges not containing any edge-double ray, and an arbitrary set $X \subseteq E(T_\infty)$ with $I \subseteq X$. Let $e_1, e_2, \ldots$ be an enumeration of the countably many edges in X.

First, we pick an $\subseteq$-maximal set $\mathcal{R}$ of pairwise disjoint rays so that $E(R) \subseteq X$ for every $R \in \mathcal{R}$, and so that $I^* := I \cup \bigcup_{R \in \mathcal{R}} E(R)$ does not contain any edge-double ray. This can indeed be accomplished. For instance, we can greedily choose a ray R_i with $E(R_i) \subseteq X$ through the edge with lowest index so that R_i is disjoint from $R_1, \ldots, R_{i-1}$ and so that $I \cup \bigcup_{j=1}^{i} E(R_j)$ does not contain any edge-double ray. (If there is no such R_i, we stop.) Now, suppose there is a double ray D with $E(D) \subseteq I^*$. Then D has to meet two rays of the R_i, let us say R_k and R_l. In D there is a finite path P between some vertex in R_k and some vertex in R_l. Therefore the edge-double ray in $R_k \cup P \cup R_l$ is already contained in $I \cup \bigcup_{j=1}^{n} E(R_j)$ for any large enough n.

Next, we go through the edges in $X - I^*$ in order of their index and add them to a set J if their inclusion does not produce an edge-double ray with I^* and the previously added edges. Now suppose there exists a double ray D with $E(D) \subseteq I^* \cup J$. Observe that D cannot meet two disjoint rays R, S whose edge set lies in I^*. Indeed, if D meets two such rays then there is a finite path in D between R and S, and the edge set of this path is contained in $I^* \cup (J \cap \{e_1, \ldots, e_n\})$ for some large enough n, yielding an edge-double ray together with $E(R \cup S)$, in contradiction to the definition of J. But then D contains a tail T that is disjoint from any ray in I^*, and hence $\mathcal{R} \cup \{T\}$ contradicts the maximality of $\mathcal{R}$. To conclude, $I^* \cup J \subseteq X$ is maximally independent in X, which shows that $(IB1)$ holds.

In order to check $(IB2)$ consider two bases B_1 and B_2 of MT_∞, and let $x \in B_1 - B_2$ be given. One of the two (graph-theoretical) components of $B_1 - x$ that contain an endvertex of x has to be rayless. Denote this component by K, and note that as there is an edge-double ray in $B_2 + x$, the basis B_2 must contain an edge y in the cut $E(K, T_\infty - K)$. We claim that $(B_1 - x) + y$ is maximally independent. As y has an endvertex in K and as K is rayless, it is clear that $(B_1 - x) + y$ is independent. On the other hand, let z be an edge outside $(B_1 - x) + y$. If $z = x$ then there is an edge-double ray in $B_1 + y = (B_1 - x) + y + z$. So, assume $z \neq x$ and denote by D the edge-double ray in $B_1 + z$ through z. If D misses x, we are done, so let $x \in D$. Next, denote by D' the edge-double ray in $B_1 + y$ through y. Edge-double rays in a tree satisfy the circuit exchange property $(C2)$, which applied to D and D' yields an edge-double ray in $(D \cup D') - x$. This shows that $(B_1 - x) + y$ is a basis, as desired.

By definition it is clear that the circuits of MT_∞ are precisely the edge-double rays. Moreover, the co-circuits are all infinite as well. This can easily be deduced from the fact that no circuit and co-circuit may meet in precisely one element. (For a proof of this, see Lemma 8.7.) Indeed, for any finite set of edges we can find an edge-double ray in T_∞ containing exactly one of its edges.

8.4 Basic properties of matroids

Throughout this chapter we shall need a number of standard properties that are taken for granted in finite matroids. Fortunately, most properties that have a rank-free formulation still hold true in infinite matroids.

Higgs [52] showed that every two bases have the same cardinality if the generalised continuum hypothesis is assumed. We shall need a weaker statement:

Lemma 8.4. [18] *If B, B' are bases of a matroid with $|B_1 - B_2| < \infty$ then $|B_1 - B_2| = |B_2 - B_1|$.*

Proof. Assume the claim is false and let B_1 and B_2 be a counterexample chosen to minimise $|B_1 - B_2|$. Thus, it follows that $|B_1 - B_2| < |B_2 - B_1|$. Let $x \in B_2 - B_1$ and consider the independent set $B_2 - x$. By (IB1), there exists a base B' such that $B_2 - x \subseteq B' \subseteq B_1 \cup (B_2 - x)$. Since B' must contain at least one element $y \in B_1 - B_2$, it follows that $|B_1 - B'| < |B_1 - B_2|$. By the choice of B_1, B_2, the statement of the lemma holds true for B_1, B', i.e. we have $|B_1 - B'| = |B' - B_1|$. As a consequence, we deduce from $B' - B_1 = (B_2 - B_1) - x$ and $B_1 - B' \subseteq (B_1 - B_2) - y$ that $|B_2 - B_1| \leq |B_1 - B_2|$, which contradicts our choice of B_1 and B_2. □

Let $M = (E, \mathcal{I})$ be a matroid, and let $X \subseteq E$. We define the restriction of M to X, denoted by $M|X$, as follows: $I \subseteq X$ is independent in $M|X$ if and only if it is independent in M. We write $M - X$ for $M|(E - X)$. A set $I \subseteq X$ is independent in the *contraction of M to X*, denoted by $M.X$, if and only if there exists a $I' \subseteq E - X$ so that $I \cup I'$ is a basis of M. We abbreviate $M.(E - X)$ by M/X. It is easy to check that $M|X$ is a matroid. For the fact that $M.X$ is a matroid too, we refer the reader to Oxley [67].

Lemma 8.5 (Oxley [67])**.** *For all $X \subseteq E$ it holds that $(M|X)^* = M^*.X$.*

Lemma 8.6 (Oxley [67])**.** *Let $X \subseteq E$, and let $B_X \subseteq X$. Then the following are equivalent*

(i) *B_X is a basis of $M.X$;*

(ii) *there exists a basis B of $M - X$ so that $B_X \cup B$ is a basis of M; and*

(iii) *for all bases B of $M - X$ it holds that $B_X \cup B$ is a basis of M.*

The circuits of a matroid satisfy the following three properties.

($C0$) $\emptyset$ is not a circuit;

($C1$) two distinct circuits are incomparable; and

($C2$) if C_1 and C_2 are two distinct circuits that contain a common element x then for all $z \in C_1 - C_2$ there exists a circuit $C \subseteq (C_1 \cup C_2) - x$ such that $z \in C$.

Indeed, these three properties characterise the circuits of a finite matroid. In infinite matroids this is no longer the case, and we will find it useful to have a stronger version of the circuit exchange axiom ($C2$).

($C2\omega$) Let C_ω be a circuit and $\{C_i : i \in I\}$ be a collection of circuits indexed by some index set I. For every $i \in I$, let $x_i \in C_i \cap C_\omega$ such that $x_i \notin C_j$ for all $j \neq i$. Then for every $z \in C_\omega - \left(\bigcup_{i \in I} C_i\right)$, there exists a circuit C contained in $\left(C_\omega \cup \bigcup_{i \in I} C_i\right) - \{x_i : i \in I\}$ such that $z \in C$.

We will prove ($C2\omega$) in Lemma 8.9. If B is a basis of a matroid M then for any element x outside B there is exactly one circuit contained in $B \cup x$, the *fundamental circuit of* x, see Oxley [65]. A *fundamental co-circuit* is a fundamental circuit of the dual matroid.

Lemma 8.7. **[18]** *Let M be a matroid and $X \subseteq E(M)$ with $X \neq \emptyset$. If for every co-circuit D of M such that $D \cap X \neq \emptyset$, it holds that $|D \cap X| \geq 2$, then X is dependent. If C is a circuit of M, then it holds that $|D \cap C| \geq 2$ for every co-circuit D such that $D \cap C \neq \emptyset$.*

Proof. First, assume that $|D \cap X| \geq 2$ for every co-circuit D such that $D \cap X \neq \emptyset$, but contrary to the claim, that X is independent. Then there exists a co-basis $\overline{B}$ contained in $E(M) - X$. If we fix an element $x \in X$ and consider the fundamental co-circuit contained in $x \cup \overline{B}$, we find a co-circuit intersecting X in exactly one element, namely x, a contradiction. Thus we conclude that X is dependent.

Now consider a circuit C and assume, to reach a contradiction, that there exists a co-circuit D such that $D \cap C = \{x\}$ for some element $x \in C$. The set $D - x$ is co-independent, so there exists a basis B of M disjoint from $D - x$. Let $Y \subseteq B$ be a set such that $(C - x) \cup Y$ is a basis of M. Then $D \cap ((C - x) \cup Y) = \emptyset$, a contradiction to the fact that D is co-dependent. $\square$

Lemma 8.8. [18] *Let M be a matroid, and let X be any subset of $E(M)$. Then for every circuit C of M/X, there exists a subset $X' \subseteq X$ such that $X' \cup C$ is a circuit of M.*

Proof. Let B_X be a base of $M|X$. The independent sets of M/X are all sets I such that $I \cup B_X$ is independent in M. Since C is a circuit of M/X, it follows that $C \cup B_X$ is dependent and so it contains a circuit C'. By the minimality of circuits, it follows that $C' - X = C$. The set $X' := C' \cap X$ is as desired by the claim. $\square$

We now prove directly that the circuits of a matroid do in fact satisfy $(C2\omega)$.

Lemma 8.9. [18] *The circuits of a matroid M satisfy $(C2\omega)$.*

Proof. Let C_ω, the circuits $\{C_i : i \in I\}$ and z and $\{x_i : i \in I\}$ be given as in the statement of $(C2\omega)$. Let M' be the matroid M restricted to the ground set $C_\omega \cup \left(\bigcup_{i \in I} C_i\right)$. Assume the claim is false and that $\{z\}$ is in no circuit of $M' - \{x_i : i \in I\}$. It follows that every basis of $M' - \{x_i : i \in I\}$ contains z, implying that $\{z\}$ is a co-circuit of $M' - \{x_i : i \in I\}$. Thus, $\{z\}$ is a circuit of $(M' - \{x_i : i \in I\})^*$, which means that $\{z\}$ is a circuit of $(M')^*/\{x_i : i \in I\}$. By Lemma 8.8, there exists a subset $I' \subseteq I$ such that $z \cup \{x_i : i \in I'\}$ is a co-circuit of M'. Furthermore, the set I' cannot be empty as then C_ω and $\{z\}$ would be a pair of circuit and co-circuit of M' intersecting in precisely one element, in contradiction to Lemma 8.7. But if we consider the circuit C_i for any $i \in I'$, C_i intersects this co-circuit exactly in the element x_i, contradicting Lemma 8.7. Thus z is contained in a circuit of $M' - \{x_i : i \in I\}$ as desired. $\square$

We conclude this section by demonstrating that the circuit axioms $(C0)$, $(C1)$ and $(C2\omega)$ are not sufficient to describe the circuits of an infinite matroid. For this, we construct a non-matroid in which $(C0)$, $(C1)$ and $(C2\omega)$ hold.

Let A and B be two disjoint copies of $\mathbb{Z}$, set $E := A \cup B$, and define $\mathcal{I}$ to be the set of all subsets I of E so that $|I \cap A| \leq |B - I|$ and $|I \cap A| < \infty$.

Observe that $M := (E, \mathcal{I})$ is a not a matroid. Indeed, in $M|A$ all finite sets are independent but A is dependent in $M|A$ as it is dependent in M. Hence, $M|A$ does not contain a maximal independent set and is therefore not a matroid.

While M is not a matroid, the minimal sets not contained in $\mathcal{I}$, i.e. the *circuits* of M, satisfy $(C0)$, $(C1)$, and $(C2\omega)$. To see this, first note that a set $C \subseteq E$ is a circuit if and only if $|C \cap A| = |B - C| + 1 < \infty$. Thus, $(C0)$ and $(C1)$ do trivially hold.

Let us next show that $(C2)$ is satisfied. Let C_1 and C_2 be two circuits that intersect in x, and let $z \in C_1 - C_2$. Then $(C_2 - x) + z$ is easily seen to be a circuit. Next, consider circuits C_ω and C_i, $i \in I$, as in $(C2\omega)$, i.e. for each $i \in I$ there is a $x_i \in C_\omega \cap C_i$ but $x_i \notin C_j$ for $j \neq i$. Then I is a finite index set. Indeed, since C_ω contains all the distinct x_i it follows from $|C_\omega \cap A| < \infty$ that at most finitely many of the x_i lie in A. On the other hand, each C_i is disjoint from all x_j, $j \neq i$, but contains all of B except for a finite number of elements. Thus, B contains only finitely many x_j. This proves that $|I| < \infty$. Then, however, $(C2\omega)$ reduces to finitely many repetitions of $(C2)$.

In conclusion, we have shown that the circuit axioms $(C0)$, $(C1)$ and $(C2\omega)$ do not suffice to characterise a matroid.

8.5 Connectivity

A finite matroid is connected if and only if every two elements are contained in a common circuit. Clearly, this definition can be extended verbatim to infinite matroids. It is, however, not clear anymore that this definition makes much sense in infinite matroids. Notably, the fact that being in a common circuit is an equivalence relation needs proof. To provide that proof is the main aim of this section.

Let $M = (E, \mathcal{I})$ be a fixed matroid in this section. Define a relation $\sim$ on E by: $x \sim y$ if and only if there is a circuit in M that contains x and y. As for finite matroids, we say that M is *connected* if $x \sim y$ for all $x, y \in E$.

Lemma 8.10. [18] *$\sim$ is an equivalence relation.*

The proof will require two simple facts that we note here.

Lemma 8.11. [18] *If C is a circuit and $X \subsetneq C$, then $C - X$ is a circuit in M/X.*

Proof. If $C - X$ is not a circuit then there exists a set $C' \subsetneq C - X$ such that C' is a circuit of M/X. Now, Lemma 8.8 yields a set $X' \subseteq X$ such that $C' \cup X'$ is a circuit of M, and this will be a proper subset of C, a contradiction. □

Lemma 8.12. [18] *Let $e \in E$ be contained in a circuit of M, and consider $X \subseteq E - e$. Then e is contained in a circuit of M/X.*

Proof. Let e be contained in a circuit C of M, and suppose that e does not lie in any circuit of M/X. Then $\{e\}$ is a a co-circuit of M/X, and thus also a co-circuit of M. This, however, contradicts Lemma 8.7 since the circuit C intersects the co-circuit $\{e\}$ in exactly one element. □

Proof of Lemma 8.10. Symmetry and reflexivity are immediate. To see transitivity, let e, f, and g in E be given such that e, f lie in a common circuit C_1, and f, g are contained in a circuit C_2. We will find a subset X of the ground set such that M/X contains a circuit containing both e and g. By Lemma 8.8, this will suffice to prove the claim.

First, we claim that without loss of generality we may assume that

$$E(M) = C_1 \cup C_2 \text{ and } C_1 \cap C_2 = \{f\}. \tag{8.1}$$

Indeed, as any circuit in any restriction of M is still a circuit of M, we may delete any element outside $C_1 \cup C_2$. Moreover, we may contract $(C_1 \cap C_2) - f$. Then $C_1 - C_2 \cup \{f\}$ is a circuit containing e and f, and similarly, $C_2 - C_1 \cup \{f\}$ is a circuit containing both f and g by Lemma 8.11. Any circuit C with $e, g \in C$ in $M/((C_1 \cap C_2) - f)$ will extend to a circuit in M, by Lemma 8.8. Hence, we may assume (8.1).

Next, we attempt to contract the set $C_2 - \{f, g\}$. If C_1 is a circuit of $M/(C_2 - \{f, g\})$, then we can find a circuit containing both e and g by applying the circuit exchange axiom $(C2)$ to the circuit C_1 and the circuit $\{f, g\}$. Thus we may assume that C_1 is not a circuit, but by Lemma 8.12, it contains a circuit C_3 containing the element e. If the circuit C_3 also contains the element f, then again by the circuit exchange axiom, we can find a circuit containing both e and g. Therefore, we instead assume that C_3 does not contain the element f. Consequently, there exists a non-empty set $A \subseteq C_2 - \{f, g\}$ such that $C_3 \cup A$ is a circuit of M.

Contract the set $C_2 - (\{f, g\} \cup A)$. We claim that the set $C_3 \cup A$ is a circuit of the contraction. If not, there exist sets $D \subseteq C_3$ and $B \subseteq A$ such that $D \cup B$ is a circuit of $M/(C_2 - (\{f, g\} \cup A))$. Furthermore, $D \cup B \cup X$ is a circuit of M for some set $X \subseteq C_2 - (\{f, g\} \cup A)$. This implies that $D = C_3$, since D contains a circuit of $M/(C_2 - \{f, g\})$. If $A \neq B$, we apply the circuit exchange axiom to the two circuits $C_3 \cup A$ and $C_3 \cup B \cup X$ to find a circuit contained in their union that does not contain the element e. However, the existence of such a circuit is a contradiction. Either it would be contained as a strict subset of C_2, or upon contracting $C_2 - \{f, g\}$ we would have a circuit contained as a strict subset of C_3. This final contradiction shows that $C_3 \cup A$ is a circuit of $M/\left(C_2 - (A \cup \{f, g\})\right)$.

We now consider two circuits in $M/\left(C_2 - (A \cup \{f, g\})\right)$. The first is $C_1' := C_3 \cup A$, which contains e. The second is $C_2' := \{f, g\} \cup A$, the remainder of C_2 after contracting $C_2 - (A \cup \{f, g\})$ (note Lemma 8.11). We have shown that in attempting to find a circuit containing e and g utilising two circuits C_1 containing e and C_2 containing g, we can restrict our attention to the case when $C_2 - C_1$ consists of exactly two elements. The argument was symmetric, so in fact we may assume that $C_1 - C_2$ also consists of only two elements. In (8.1) we observed that we may assume that C_1 and C_2 intersect in exactly one element. Thus we have reduced to a matroid on five elements, in which it is easy to find a circuit containing both e and g. $\square$

As an application of Lemma 8.10 we shall show that every matroid is the direct sum of its connected components. With a little extra effort this will allow us to re-prove a characterisation matroids that are both finitary and co-finitary, that had been noted by Las Vergnas [88], and by Bean [9] before.

Let $M_i = (E_i, \mathcal{I}_i)$ be a collection of matroids indexed by a set I. We define the *direct sum* of the M_i, written $\bigoplus_{i \in I} M_i$, to have ground set consisting of $E := \bigcup_{i \in I} E_i$ and independent sets $\mathcal{I} = \left\{ \bigcup_{i \in I} J_i : J_i \in \mathcal{I}_i \right\}$.

As noted by Oxley [67] for finitary matroids, it is easy to check that:

Lemma 8.13. *The direct sum of matroids M_i for $i \in I$ is in fact a matroid.*

Lemma 8.14. **[18]** *Every matroid is the direct sum of the restrictions to its connected components.*

Proof. Let $M = (E, \mathcal{I})$ be a matroid. As $\sim$ is an equivalence relation, the ground set E partitions into connected components E_i, for some index set I. Setting $M_i := M|E_i$, we claim that $\bigoplus_{i \in I} M_i$ and M have the same independent sets.

Clearly, if I is independent in M, then $I \cap E_i$ is independent in M_i for every $i \in I$, which implies that I is independent in $\bigoplus_{i \in I} M_i$. Conversely, consider a set $X \subseteq E$ that is dependent in M. Then, X contains a circuit C, which, in turn, lies in E_j for some $j \in I$. Therefore, $X \cap E_j$ is dependent, implying that X is dependent in $\bigoplus_{i \in I} M_i$ as well. $\square$

We now give the characterisation of matroids that are both finitary and co-finitary.

Theorem 8.15. *A matroid M is both finitary and co-finitary if and only if there exists an index set I and finite matroids M_i for $i \in I$ such that $M = \bigoplus_{i \in I} M_i$.*

Theorem 8.15 is a direct consequence of the following lemma, which has previously been proved by Bean [9]. The theorem was first proved by Las Vergnas [88]. Our proof is different from the proofs of Las Vergnas and of Bean.

Lemma 8.16. *An infinite, connected matroid contains either an infinite circuit or an infinite co-circuit.*

Proof. Assume, to reach a contradiction, that M is a connected matroid with $|E(M)| = \infty$ such that every circuit and every co-circuit of M is finite. Fix an element $e \in E(M)$ and let $C_1, C_2, C_3, \dots$ be an infinite sequence of distinct circuits each containing e. Let $M' = M| \left(\bigcup_{i=1}^{\infty} C_i\right)$ be the restriction of M to the union of all the circuits C_i. Note that M' contains a countable number of elements by our assumption that every circuit is finite. Let $e_1, e_2, \dots$ be an enumeration of $E(M')$ such that $e_1 = e$. We now recursively define an infinite set $\mathcal{C}_i$ of circuits and a finite set X_i for $i \geq 1$. Let $\mathcal{C}_1 = \{C_i : i \geq 1\}$ and $X_1 = \{e_1\}$. Assuming $\mathcal{C}_i$ and X_i are defined for $i = 1, 2, \dots, k$, we define $\mathcal{C}_{k+1}$ as follows. If infinitely many circuits in $\mathcal{C}_k$ contain e_{k+1}, we let $\mathcal{C}_{k+1} = \{C \in \mathcal{C}_k : e_{k+1} \in C\}$, and $X_{k+1} = X_k \cup \{e_{k+1}\}$. Otherwise we set $\mathcal{C}_{k+1} = \{C \in \mathcal{C}_k : e_{k+1} \notin C\}$ and $X_{k+1} = X_k$. Let $X = \bigcup_1^{\infty} X_i$. Note that $\mathcal{C}_k$ is always an infinite set, and for all i, j, $i < j$, if $e_i \in X_i$, then $e_i \in C$ for all circuits $C \in \mathcal{C}_j$.

We claim that the set X is dependent in M' . By Lemma 8.7, if X is independent then there is a co-circuit D of M' that meets X in exactly one element. As D is finite, we may pick an integer k such that $D \subseteq \{e_1, \dots, e_k\}$. Choose any $C \in \mathcal{C}_k$. Since $C \cap \{e_1, e_2, e_3, \dots, e_k\} = X_k$, we see that C also intersects D in exactly one element, a contradiction to Lemma 8.7. Thus, X is dependent and therefore contains a circuit C'. As M is finitary, C' contains a finite number of elements, and so $C' \subseteq X_\ell$ for some integer ℓ. However, $\mathcal{C}_\ell$ contains an infinite number of circuits, each containing the set X_ℓ. It follows that some circuit strictly contains C', a contradiction. $\square$

Proof: Theorem 8.15. If we let $\mathcal{C}(M)$ be the set of circuits of a matroid M, an immediate consequence of the definition of the direct sum is that $\mathcal{C}\left(\bigoplus_{i\in I} M_i\right) = \bigcup_{i\in I} \mathcal{C}(M_i)$. Moreover, the dual version of Lemma 8.7 shows that every co-circuit is completely contained in some M_i. It follows that if $M = \bigoplus_{i\in I} M_i$ where M_i is finite for all $i \in I$, then M is both finitary and co-finitary.

To prove the other direction of the claim, let M be a matroid that is both finitary and co-finitary. Let M_i for $i \in I$ be the restriction of M to the connected components. For every $i \in I$, the matroid M_i is connected, so by our assumptions on M and by Lemma 8.16, M_i must be a finite matroid. Lemma 8.14 implies that $M = \bigoplus_{i\in I} M_i$, and the theorem is proved.

$\square$

8.6 Higher connectivity

Let us recapitulate the definition of k-connectivity in finite matroids and see what we can keep of that in infinite matroids. So, in a finite matroid M on a ground set E the *connectivity function* κ is defined as

$$\kappa_M(X) := r_M(X) + r_M(E - X) - r_M(E) \text{ for } X \subseteq E. \tag{8.2}$$

(We note that some authors define a connectivity function λ by $\lambda(X) = \kappa(X) + 1$. In dropping the $+1$ we follow Oxley [68].) We call a partition (X, Y) of E a *k-separation* if $\kappa_M(X) \le k - 1$ and $|X|, |Y| \ge k$. The matroid M is *k-connected* if there exists no ℓ-separation with $\ell < k$.

Of these notions only the connectivity function is obviously useless in an infinite matroid, as the involved ranks will usually be infinite. We shall therefore only redefine κ and leave the other definitions unchanged. For this we have two aims. First, the new κ should coincide with the ordinary connectivity function if the matroid is finite. Second, κ should be consistent with connectivity as defined in the previous section.

Our goal is to find a rank-free formulation of (8.2). Observe that (8.2) can be interpreted as the number of elements we need to delete from the union of a basis of $M|X$ and a basis of $M|Y$ in order to obtain a basis of the whole matroid. To show that this number does not depend on the choice of bases is the main purpose of the next lemma.

Let $M = (E, \mathcal{I})$ be a matroid, and and let I, J be two independent sets. We define

$$\mathrm{del}_M(I, J) := \min\{|F| : F \subseteq I \cup J,\ (I \cup J) - F \in \mathcal{I}\}.$$

Thus, $\mathrm{del}_M(I, J)$ is either a non-negative integer or infinity. If there is no chance of confusion, we will simply write $\mathrm{del}(I, J)$ rather than $\mathrm{del}_M(I, J)$

Lemma 8.17.[18] *Let $M = (E, \mathcal{I})$ be a matroid, let (X, Y) be a partition of E, and let B_X be a basis of $M|X$ and B_Y a basis of $M|Y$. Then*

(i) $\mathrm{del}(B_X, B_Y) = |F|$ *for any $F \subseteq B_X \cup B_Y$ so that $(B_X \cup B_Y) - F$ is a basis of M;*

(ii) $\mathrm{del}(B_X, B_Y) = |F|$ *for any $F \subseteq B_X$ so that $(B_X - F) \cup B_Y$ is a basis of M;*

(iii) $\mathrm{del}(B_X, B_Y) = \mathrm{del}(B'_X, B'_Y)$ *for every basis B'_X of $M|X$ and basis B'_Y of $M|Y$.*

Proof. Let us first prove that

$$\text{if for } F_1, F_2 \subseteq B_X \cup B_Y \text{ it holds that } B_i := (B_X \cup B_Y) - F_i \text{ is a basis of } M \ (i = 1, 2), \text{ then } |F_1| = |F_2|. \tag{8.3}$$

We may assume that one of $|F_1|$ and $|F_2|$ is finite, say $|F_2|$. Then as $|B_1 - B_2| = |F_2 - F_1| < \infty$, it follows from Lemma 8.4 that $|F_1 - F_2| = |F_2 - F_1|$, and hence $|F_1| = |F_2|$.

(i) If $F' \subseteq B_X \cup B_Y$ has minimal cardinality so that $(B_X \cup B_Y) - F'$ is independent, then $(B_X \cup B_Y) - F'$ is maximally independent in $B_X \cup B_Y$, and hence a basis of M. Moreover, if F is as in (i) then $|F| = |F'|$ by (8.3).

(ii) Follows directly from (8.3) and (i).

(iii) Let $F \subseteq B_X$ as in (ii), i.e. $|F| = \mathrm{del}(B_X, B_Y)$. Because of the equivalence of (ii) and (iii) in Lemma 8.6, we obtain that $(B_X - F) \cup B'_Y$ is a basis of M as well, and it follows that $\mathrm{del}(B_X, B_Y) = |F| = \mathrm{del}(B_X, B'_Y)$. By exchanging the roles of X and Y we get then that $\mathrm{del}(B_X, B'_Y) = \mathrm{del}(B'_X, B'_Y)$, which finishes the proof. □

We now give a rank-free definition of the *connectivity function*. Let X be a subset of $E(M)$ for some matroid M. We pick an arbitrary basis B of $M|X$, and a basis B' of $M - X$ and define $\kappa_M(X) := \mathrm{del}_M(B, B')$. Lemma 8.17 (iii) ensures that κ is well defined, i.e. that the value of $\kappa_M(X)$ only depends on X (and M) and not on the choice of the bases. The next two propositions demonstrate that κ extends the connectivity function of a finite matroid and, furthermore, is consistent with connectivity defined in terms of circuits.

Proposition 8.18. [18] *If M is a finite matroid on gound set E, and if $X \subseteq E$ then*

$$r(X) + r(E - X) - r(E) = \kappa(X).$$

Proof. Let B be a basis of $M|X$, B' a basis of $M - X$, and choose $F \subseteq B \cup B'$ so that $(B \cup B') - F$ is a basis of M. Then

$$\begin{aligned} \kappa(X) &= \mathrm{del}(B, B') = |F| = |B| + |B'| - |(B \cup B') - F| \\ &= r(X) + r(E - X) - r(E). \end{aligned}$$

□

Proposition 8.19. [18] *A matroid is 2-connected if and only if it is connected.*

Proof. Let $M = (E, \mathcal{I})$ be a matroid. First, assume that there is a 1-separation (X, Y) of M. We need to show that M cannot be connected. Pick $x \in X$ and $y \in Y$ and suppose there is a circuit C containing both, x and y. Then $C \cap X$ as well as $C \cap Y$ is independent, and so there are bases $B_X \supseteq C \cap X$ of $M|X$ and $B_Y \supseteq C \cap Y$ of $M|Y$, by (IB1). As (X, Y) is a 1-separation, $B_X \cup B_Y$ is a basis. On the other hand, we have $C \subseteq B_X \cup B_Y$, a contradiction.

Conversely, assume M to be 2-connected and pick a $x \in E$. Define X to be the set of all x' so that x' lies in a common circuit with x. If $X = E$ then we are done, with Lemma 8.10. So suppose that $Y := E - X$ is not empty, and choose a basis B_X of $M|X$ and a basis B_Y of $M|Y$. Since there are no 1-separations of M, $B_X \cup B_Y$ is dependent and thus contains a circuit C. But then C must meet X as well as Y, yielding together with Lemma 8.10 a contradiction to the definition of X. $\square$

To illustrate the definition of κ and since it is relevant for the open problem stated below let us compute the connectivity of the matroid MT_∞ described in Section 8.3. Since every two edges in the ω-regular tree T_∞ are contained in a common double ray, we see that MT_∞ is 2-connected. On the other hand, MT_∞ contains a 2-separation: deleting an edge e splits the graph T_∞ into two components K_1, K_2. Put $X := E(K_1) \cup \{e\}$ and $Y := E(K_2)$, and pick a basis B_X of $MT_\infty|X$ and a basis B_Y of $MT_\infty|Y$. Clearly, neither B_X nor B_Y contains a double ray, while every double ray in $B_X \cup B_Y$ has to use e. Thus, $(B_X \cup B_Y) - e$ is a basis of MT_∞, and (X, Y) therefore a 2-separation.

It is easy to construct matroids of connectivity k for arbitrary positive integers k. Moreover, there are matroids that have infinite connectivity, namely the uniform matroid $U_{r,k}$ where $k \simeq r/2$. However, it can be argued that these matroids are simply too small for their high connectivity, and therefore more a fluke of the definition than a true example of an infinitely connected matroid. Such a matroid should certainly have an infinite ground set.

Problem 8.20. *Find an infinite infinitely connected matroid.*

As for finite matroids the minimal size of a circuit or co-circuit is an upper bound on the connectivity. So, an infinitely connected infinite matroid cannot have finite circuits or co-circuits. The matroid mT_∞ based on T_∞ has that property but, as we have seen, fails to be 3-connected.

8.7 Properties of the connectivity function

The main goal of this section is to list a number of lemmas that will be necessary in extending Tutte's linking theorem to finitary matroids. As a by-product we will see that the choice of connectivity function does in fact satisfy many of the standard properties of the connectivity function in a finite matroid, specifically Proposition 8.21 and Lemmas 8.22 and 8.24.

Let us start by showing that connectivity is invariant under duality.

Proposition 8.21. [18] *Let M be a matroid, and let $X \subseteq E(M)$. It holds that $\kappa_M(X) = \kappa_{M^*}(X)$.*

Proof. Set $Y := E(M) - X$, let B_X be a basis of $M|X$, and B_Y a basis of $M|Y$. Pick $F_X \subseteq B_X$ and $F_Y \subseteq B_Y$ so that $(B_X - F_X) \cup B_Y$ and $B_X \cup (B_Y - F_Y)$ are bases of M.

Then $B_X^* := (X - B_X) \cup F_X$ and $B_Y^* := (Y - B_Y) \cup F_Y$ are bases of $M^*|X$ and $M^*|Y$, respectively. Indeed, $B_X - F_X$ is a basis of $M.X$, by Lemma 8.6, which implies that $X - (B_X - F_X) = B_X^*$ is a basis of $(M.X)^* = M^*|X$ (Lemma 8.5). For B_Y^* we reason in a similar way.

Moreover, since $B_X \cup (B_Y - F_Y)$ is a basis of M we see from

$$(B_X^* - F_X) \cup B_Y^* = (X - B_X) \cup (Y - B_Y) \cup F_Y = E(M) - (B_X \cup (B_Y - F_Y))$$

that $(B_X^* - F_X) \cup B_Y^*$ is a basis of M^*. Therefore

$$\mathrm{del}_M(B_X, B_Y) = |F_X| = \mathrm{del}_{M^*}(B_X^*, B_Y^*),$$

and thus $\kappa_M(X) = \kappa_{M^*}(X)$, as desired. $\square$

Lemma 8.22. [18] *The connectivity function κ is submodular, i.e. for all $X, Y \subseteq E(M)$ for a matroid M it holds that*

$$\kappa(X) + \kappa(Y) \geq \kappa(X \cup Y) + \kappa(X \cap Y).$$

Proof. Denote the ground set of M by E. Choose a basis $B_\cap$ of $M|(X \cap Y)$, and a basis $B_{\overline{\cap}}$ of $M - (X \cup Y)$. Pick $F \subseteq B_\cap \cup B_{\overline{\cap}}$ so that $I := (B_\cap \cup B_{\overline{\cap}}) - F$ is a basis of $M|(X \cap Y) \cup (E - (X \cup Y))$. Next, we extend I into $(X - Y) \cup (Y - X)$: let $I_{X-Y} \subseteq X - Y$ and $I_{Y-X} \subseteq Y - X$ so that $I \cup I_{X-Y} \cup I_{Y-X}$ is a basis of M.

We claim that $I_\cup := B_\cap \cup I_{X-Y} \cup I_{Y-X}$ (and by symmetry also $I_{\overline{\cup}} := B_{\overline{\cap}} \cup I_{X-Y} \cup I_{Y-X}$), is independent. Suppose that $I_\cup$ contains a circuit C_ω. For each $x \in F \cap C_\omega$, denote by C_x the (fundamental) circuit in $I \cup \{x\}$. As C_ω meets $I_{X-Y} \cup I_{Y-X}$, we have $C_\omega \not\subseteq \bigcup_{x \in F \cap C} C_x$. Thus, $(C2\omega)$ is applicable and yields a circuit $C \subseteq (C_\omega \cup \bigcup_{x \in F \cap C} C_x) - F$. As therefore C is a subset of the independent set $I \cup I_{X-Y} \cup I_{Y-X}$, we obtain a contradiction.

Since $I_\cup$ is independent and $B_\cap \subseteq I_\cup$ a basis of $M|(X \cap Y)$, we can pick $F_\cup^X \subseteq X - (Y \cup I_\cup)$ and $F_\cup^Y \subseteq Y - (X \cup I_\cup)$ so that $I_\cup \cup F_\cup^X \cup F_\cup^Y$ is a basis of $M|(X \cup Y)$. In a symmetric way, we pick $F_{\overline{\cup}}^X \subseteq X - (Y \cup I_{\overline{\cup}})$ and $F_{\overline{\cup}}^Y \subseteq Y - (X \cup I_{\overline{\cup}})$ so that $I_{\overline{\cup}} \cup F_{\overline{\cup}}^X \cup F_{\overline{\cup}}^Y$ is a basis of $M - (X \cap Y)$.

Let us compute a lower bound for $\kappa(X)$. Both sets $I_X := B_\cap \cup I_{X-Y} \cup F_\cup^X$ and $I_{\overline{X}} := B_{\overline{\cap}} \cup I_{Y-X} \cup F_{\overline{\cup}}^Y$ are independent. As furthermore $I_X \subseteq X$ and $I_{\overline{X}} \subseteq E - X$, we obtain that $\kappa(X) \geq \mathrm{del}(I_X, I_{\overline{X}})$. Since each $x \in F \cup F_\cup^X \cup F_{\overline{\cup}}^Y$

gives rise to a circuit in $I \cup I_{X-Y} \cup I_{Y-X} \cup \{x\}$, we get that $\mathrm{del}(I_X, I_{\overline{X}}) \geq |F| + |F^X_{\cup}| + |F^Y_{\overline{\cup}}|$. In a similar way we obtain a lower bound for $\kappa(Y)$. Together these result in

$$\kappa(X) + \kappa(Y) \geq 2|F| + |F^X_{\cup}| + |F^Y_{\cup}| + |F^X_{\overline{\cup}}| + |F^Y_{\overline{\cup}}|.$$

To conclude the proof we compute upper bounds for $\kappa(X \cap Y)$ and $\kappa(X \cup Y)$. Since $B_\cap$ is a basis of $M|(X \cap Y)$ and $B_{\overline{\cup}} := I_{\overline{\cup}} \cup F^X_{\overline{\cup}} \cup F^Y_{\overline{\cup}}$ is one of $M - (X \cap Y)$, it holds that $\kappa(X \cap Y) = \mathrm{del}(B_\cap, B_{\overline{\cup}})$. Since $I \cup I_{X-Y} \cup I_{Y-X}$ is independent, we get that $\mathrm{del}(B_\cap, B_{\overline{\cup}}) \leq |F| + |F^X_{\overline{\cup}}| + |F^Y_{\overline{\cup}}|$. For $\kappa(X \cup Y)$ the computation is similar, so that we obtain

$$\kappa(X \cap Y) + \kappa(X \cup Y) \leq 2|F| + |F^X_{\cup}| + |F^Y_{\cup}| + |F^X_{\overline{\cup}}| + |F^Y_{\overline{\cup}}|,$$

as desired. □

Lemma 8.23. [18] *In a matroid M let $(X_i)_{i \in I}$ be a family of nested subsets of $E(M)$, i.e. $X_i \subseteq X_j$ if $i \geq j$, and set $X := \bigcap_{i \in I} X_i$. If $\kappa(X_i) \leq k$ for all $i \in I$ then $\kappa(X) \leq k$.*

Proof. Set $Y_i := E(M) - X_i$ for $i \in I$, $Y := \bigcap_{i \in I} Y_i = Y_1$ and $Z := E(M) - (X \cup Y)$. Pick bases B_X of $M|X$ and B_Y of $M|Y$. Choose $I_Z \subseteq Z$ so that $B_Y \cup I_Z$ is a basis of $M|(Y \cup Z)$. Moreover, there exists a finite set (of size $\leq k$) $F \subseteq B_Y$ so that $B_X \cup (B_Y - F)$ is a basis of $M|(X \cup Y)$, and a (possibly infinite) set $F_Z \subseteq I_Z$ so that $B_X \cup (B_Y - F) \cup (I_Z - F_Z)$ is a basis of M.

Suppose that $k + 1 \leq \kappa(X) = |F| + |F_Z|$. Then choose $j \in I$ large enough so that $|F| + |F_Z \cap Y_j| \geq k + 1$. Extend the independent subset $B_X \cup (I_Z \cap X_j) - F_Z$ of X_j to a basis B of $M|X_j$. The set $B_Y \cup (I_Z \cap Y_j)$ is independent too, and we may extend it to a basis B' of $M|Y_j$. As $B_X \cup B_Y \cup (I_Z - (F_Z \cap X_j)) \subseteq B \cup B'$ we obtain with

$$\kappa(X_j) = \mathrm{del}(B, B') \geq |F| + |F_Z \cap Y_j| \geq k + 1$$

a contradiction. □

For disjoint sets $X, Y \subseteq E(M)$ define

$$\kappa_M(X, Y) := \min\{\kappa_M(U) : X \subseteq U \subseteq E(M) - Y\}.$$

Again, we may drop the subscript M if no confusion is likely.

Lemma 8.24. [18] *Let $X, Y \subseteq E(M)$ be disjoint, and let N be a minor of a matroid M so that $X \cup Y \subseteq E(N)$. Then $\kappa_N(X, Y) \leq \kappa_M(X, Y)$.*

Proof. Let $U \subseteq E(M)$ be such that $X \subseteq U \subseteq E(M) - Y$ and $\kappa_M(U) = \kappa_M(X,Y)$. First suppose that $N = M - D$ for $D \subseteq E(M) - (X \cup Y)$. Pick a basis B'_U of $M|(U - D)$ and extend it to a basis B_U of $M|U$. Define B'_W and B_W analogously for $W := E(M) - U$. Let $F \subseteq B_U \cup B_W$ be such that $(B_U \cup B_W) - F$ is a basis of M. Since B'_U and B'_W are bases of $M|(U-D) = N|(U - D)$ resp. of $N|W - D$, and since clearly $(B'_U \cup B'_W) - (F - D)$ is independent it follows that $\kappa_N(X,Y) \leq \kappa_N(U - D) \leq |F| = \kappa_M(X,Y)$.

Next, assume that $N = M/C$ for some $C \subseteq E(M)$. Then, using Lemma 8.5 and Proposition 8.21 we obtain

$$\kappa_N(X,Y) = \kappa_{(M^*-C)^*}(X,Y) = \kappa_{M^*-C}(X,Y) \leq \kappa_{M^*}(X,Y) = \kappa_M(X,Y).$$

As $N = M/C - D$ for some C, D the lemma follows. ☐

Lemma 8.25.[18] *In a matroid M let X, Y be two disjoint subsets of $E(M)$, and let $X' \subseteq X$ and $Y' \subseteq Y$ be such that $\kappa(X', Y') = k-1$. Then $\kappa(X,Y) \geq k$ if and only if there exist $x \in X$ and $y \in Y$ so that $\kappa(X' + x, Y' + y) = k$.*

Proof. Necessity is trivial. To prove sufficiency, assume that there exist no x, y as in the statement. For $x \in X$ denote by $\mathcal{U}_x$ the sets U with $X' + x \subseteq U \subseteq E(M) - Y'$ and $\kappa(U) = k - 1$. By our assumption, $\mathcal{U}_x \neq \emptyset$. By Zorn's Lemma and Lemma 8.23 there exist an $\subseteq$-minimal element U_x in $\mathcal{U}_x$.

Suppose there is a $y \in Y \cap U_x$. Again by the assumption, we can find a set Z with $X' + x \subseteq Z \subseteq E(M) - (Y' + y)$ and $\kappa(Z) = k - 1$. From Lemma 8.22 it follows that $\kappa(U_x \cap Z) = k - 1$, and thus an element of $\mathcal{U}_x$. As $y \notin U_x$ it is strictly smaller than U_x and therefore a contradiction to the minimality of U_x. Hence, U_x is disjoint from Y.

Next, let $\mathcal{W}$ be the set of sets W with $Y \subseteq W \subseteq E(M) - X'$ and $k(W) = k - 1$. As $E(M) - U_x \in \mathcal{W}$ for every $x \in X$, $\mathcal{W}$ is non-empty and we can apply Zorn's Lemma and Lemma 8.23 in order to find an $\subseteq$-minimal element W' in $\mathcal{W}$. Suppose that there is a $x \in X \cap W'$. But then Lemma 8.22 shows that $W' \cap (E(M) - U_x) \in \mathcal{W}$, a contradiction to the minimality of W'.

In conclusion, we have found that $Y \subseteq W \subseteq E(M) - X$ and $k(W) = k-1$, which implies $\kappa(X,Y) \leq k - 1$. This contradiction proves the claim. ☐

8.8 The linking theorem

We prove our main theorem in this chapter, Tutte's linking theorem for finitary (and co-finitary) matroids:

Theorem 8.26.[18] *Let M be a finitary or co-finitary B-matroid, and let X and Y be two disjoint subsets of $E(M)$. Then there exists a partition (C, D) of $E(M) - (X \cup Y)$ such that $\kappa_{M/C-D}(X,Y) = \kappa_M(X,Y)$.*

A fact that is related to Tutte's linking theorem, but quite a bit simpler to prove, is that for every element e of a finite 2-connected matroid M, one of M/e or $M-e$ is still 2-connected. This fact extends to infinite matroids in a straightforward manner. Yet, in an infinite matroid it is seldom necessary to only delete or contract a single element or even a finite set. Rather, to be useful we would need that

for any set $F \subseteq E(M)$ of a 2-connected matroid $M = (E, \mathcal{I})$ there always exists a partition (A, B) of F so that $M/A - B$ is still 2-connected. (8.4)

Unfortunately, such a partition of F does not need to exist. Indeed, consider the matroid M_{ST} obtained from the double ladder (see Figure 8.1), i.e. the matroid on the edge set of the double in which an edge set is independent if and only if it does not contain a finite (graph-)circuit. If F is the set of rungs then we cannot contract any element in F without destroying 2-connectivity, but if we delete all rungs we are left with two disjoint double rays.

In view of the failure of (8.4) in infinite matroids, even in finitary matroids like the example, it appears somewhat striking that Tutte's linking theorem does extend to, at least, finitary matroids.

Before we can finally prove Theorem 8.26 we need one more definition and one lemma that will be essential when $\kappa(X, Y) < \infty$. Let M' be a minor of a matroid M. We say a k-separation (X', Y') of M' *extends* to a k-separation of M if there exists a k-separation (X, Y) of M such that $X' \subseteq X$, $Y' \subseteq Y$. The k-separation (X', Y') is *exact* if $\kappa(X', Y') = k - 1$.

Lemma 8.27. [18] *Let M be a matroid and let $X \cup Y \subseteq E(M)$ be disjoint subsets of $E(M)$. Let (X, Y) be an exact k-separation of $M|(X \cup Y)$ that does not extend to a k-separation of M. Then there exist circuits C_1 and C_2 of M such that (X, Y) does not extend to a k-separation of $M|(X \cup Y \cup C_1 \cup C_2)$*

Proof. We define

$$Comp_X := \{D : \ D \text{ a component of } M/X \text{ such that } D \cap Y = \emptyset\}$$

to be the set of connected components of M/X that do not contain an element of Y. Symmetrically, we define $Comp_Y$ to be the components of M/Y that do not contain an element of X.

We claim that

if $A \in Comp_X$ and $B \in Comp_Y$ such that $A \cap B \neq \emptyset$, then $A = B$. (8.5)

Assume the claim to be false and let A and B be a counterexample. Without loss of generality, we may assume there exists an element $x \in A \cap B$ and an element $y \in B - A$. By definition (and Lemma 8.8), there exists a circuit C_Y of M such that $C_Y - Y$ is a circuit of M/Y containing x and y. Now consider the circuit C_Y in the matroid M/X. By Lemma 8.12, the dependent set $C_Y - X$ (in fact, C_Y is disjoint from X but we will not need this) contains a circuit of M/X that contains x but not y as y and x lie in distinct components of M/X. It follows that there exists a circuit C_X in M such that $x \in C_X - X \subseteq C_Y - y$. By the circuit exchange axiom, there exists a circuit $C \subseteq C_X \cup C_Y$ of M containing y but not x. Hence, there is a circuit $D \subseteq C - Y$ in M/Y with $y \in D$ and $x \notin D$. Since $y \in B$, D cannot meet X, which implies $D \subseteq C - (X \cup Y) \subseteq C_Y - Y$. As $x \notin D$, the circuit D in M/Y is a strict subset of the circuit $C_Y - Y$, a contradiction. This proves the claim.

Note that it is certainly possible that a set A lies in both $Comp_X$ and $Comp_Y$. In a slight abuse of notation, we let $E(Comp_X) = \bigcup_{A \in Comp_X} A$, and similarly define $E(Comp_Y)$.

Next, let us prove that

If $E(Comp_X) \cup E(Comp_Y) \cup X \cup Y = E(M)$, then the separation (X, Y) extends to a k-separation of M. (8.6)

Indeed, consider the partition (L, R) for $L := X \cup E(Comp_X)$ and $R := E(M) - L \supseteq Y$. We claim that (L, R) is a k-separation of M. Pick bases B_X and B_Y of $M|X$ resp. of $M|Y$, extend B_X to a basis B_L of $M|L$, and let B_R be a basis of $M|R$ containing B_Y.

Consider a circuit $C \subseteq B_L \cup B_R$ in M, and suppose C to contain an element $x \in B_L - X$. The set $C - X$ contains a circuit C' in M/X containing x. Since $(C \cap B_L) - X$ is independent in M/X, the circuit C' must contain an element $y \in B_R$. This implies that x and y are in the same component of M/X, and consequently, $y \in E(Comp_X)$. This contradicts the definition of the partition, implying that no such circuit C and element x exist. A similar argument implies that $B_L \cup B_R$ does not contain any circuit containing an element of $B_R - Y$ by considering M/Y. We conclude that every circuit contained in $B_L \cup B_R$ must lie in $B_X \cup B_Y$. As $\kappa(X, Y) = k - 1$, there exists a set of $k-1$ elements intersecting every circuit contained in $B_X \cup B_Y$, and thus in $B_L \cup B_R$, which implies that (L, R) forms a k-separation. This completes the proof of (8.6).

Before finishing the proof of the lemma, we need one further claim.

Let C be a circuit of M such that $C - (X \cup Y)$ is a circuit of $M/(X \cup Y)$. Then the only k-separations of $M|(X \cup Y \cup C)$ that extend (X, Y) are $(X \cup (C - Y), Y)$ and $(X, Y \cup (C - X))$. (8.7)

Assume that (X', Y') is a k-separation that extends (X, Y) in the matroid $M|(X \cup Y \cup C)$. Let $C' := C - (X \cup Y)$. Assume that (X', Y') induces a proper partition of C', i.e. that $C' \cap X' \neq \emptyset$ and $C' \cap Y' \neq \emptyset$. Picking bases B_X of $M|X$ and B_Y of $M|Y$ we observe that $B_X \cup (C' - Y')$ and $B_Y \cup (C' - X')$ form bases of $M|X'$ and $M|Y'$ respectively. Assume there exists a set F of $k-1$ elements intersecting every circuit contained in $B_X \cup B_Y \cup C$. By our assumption that (X, Y) is an exact k-separation, we see that $F \subseteq B_X \cup B_Y$. However, C' is a circuit of $M/(X \cup Y)$, or, equivalently, C' is a circuit of $M/((B_X \cup B_Y) - F)$. It follows that there exists a circuit contained in $C' \cup B_X \cup B_Y$ avoiding the set F, a contradiction. This completes the proof of (8.7).

Since, by assumption, (X, Y) does not extend to a k-separation of M it follows from (8.6) that there is an $e \notin E(Comp_X) \cup E(Comp_Y)$. Then there exists a circuit C_1 of M containing e with $C_1 \cap Y \neq \emptyset$ such that the following hold:

- $C_1 - X$ is a circuit of M/X, and
- $C_1 - (X \cup Y)$ is a circuit of $M/(X \cup Y)$.

To see that such a circuit C_1 exists, recall first that $e \notin E(Comp_X)$ implies that there is a circuit C_X in M/X containing e so that $C_X \cap Y \neq \emptyset$. Then $C_X - Y$ contains a circuit C' in $M/(X \cup Y)$ with $e \in C'$ (see Lemma 8.12). For suitable $A_X \subseteq X$ and $A_Y \subseteq Y$ it therefore holds, by Lemma 8.8, that $C' \cup A_X \cup A_Y$ is a circuit of M. If $A_Y = \emptyset$ then C' would be a dependent set of M/X strictly contained in the circuit C_X. Thus, $A_Y \neq \emptyset$ and $C_1 := A_X \cup A_Y \cup C'$ has the desired properties. Symmetrically, there exists a circuit C_2 containing e and intersecting X in at least one element such that $C_2 - Y$ is a circuit of M/Y and $C_2 - (X \cup Y)$ is a circuit of $M/(X \cup Y)$.

Let us now see that the circuits C_1 and C_2 are as required in the statement of the lemma. To reach a contradiction, suppose that (X, Y) extends to a k-separation (X', Y') of $M|(X \cup Y \cup C_1 \cup C_2)$. By symmetry, we may assume that $e \in X'$. By (8.7), we see that $C_1 - (X \cup Y) \subseteq X'$ and $C_2 - (X \cup Y) \subseteq X'$ as well. Pick a basis B_X of $M|X$, and a basis B_Y of $M|Y$. We extend the set $B_X \cup (C_1 - Y)$, which is independent because $C_1 - Y$ is independent in M/X, to a basis $B_{X'}$ of $M|X'$. The set B_Y forms a basis of $M|Y'$ as $Y' = Y$. Choose $F \subseteq B_{X'} \cup B_Y$ so that $(B_{X'} \cup B_Y) - F$ is independent. As (X, Y) is an exact k-separation and (X', Y) thus too, it follows that $|F| = k - 1$. As $\kappa_{M|X \cup Y}(X, Y) = k-1$, we see $F \subseteq B_X \cup B_Y$. However, the set $C_1 - (X \cup Y)$ is dependent in $M/(X \cup Y)$ and thus in $M/((B_X \cup B_Y) - F)$. Hence, there is a set $S \subseteq (B_X \cup B_Y) - F$ so that $C_1 - (X \cup Y) \cup S \subseteq B_{X'} \cup B_Y - F$ is dependent in M, contradicting our choice of F. This contradiction implies that the separation (X, Y) does not extend to a k-separation of $M|(X \cup Y \cup C_1 \cup C_2)$, which concludes the proof of the lemma. $\square$

We now proceed with the proof of the linking theorem for finitary matroids.

Proof of Theorem 8.26. Since for any set $F \subseteq E(M)$ it holds, by Proposition 8.21, that $\kappa_M(F) = \kappa_{M^*}(F)$, we may assume M to be finitary. We will consider the case when $\kappa_M(X, Y) = \infty$ and $\kappa_M(X, Y) < \infty$ separately.

Assume that $\kappa_M(X, Y) = \infty$. We inductively define a series of disjoint circuits $C_1, C_2, \ldots$ in different minors of M as follows. Starting with C_1 to be chosen as a circuit in M intersecting both X and Y, assume $C_1, \ldots, C_t$ to be defined for $t \geq 1$. Note that as $Z := \bigcup_{i=1}^{t} C_i$ is finite, we still have $\kappa_{M/Z}(X - Z, Y - Z) = \infty$. Thus, there exists a circuit C_{t+1} in M/Z that meets both both $X - Z$ and $Y - Z$. Having finished this construction, we let $C_X = (\bigcup_{i=1}^{\infty} C_i) \cap X$, $C_Y = (\bigcup_{i=1}^{\infty} C_i) \cap Y$, and $C = (\bigcup_{i=1}^{\infty} C_i) - C_X - C_Y$. Set $D = E(M) - (X \cup Y \cup C)$.

In order to show $\kappa_{M/C-D}(X, Y) = \infty$ observe first that C_X (and symmetrically, C_Y) is an independent set in M/C. If not then there exists a circuit $A \subseteq C_X \cup C$. Given that M is finitary and A thus finite, there exists a minimal index t such that $A \subseteq \bigcup_{i=1}^{t} C_i$. It follows that $A - (\bigcup_{i=1}^{t-1} C_i)$ is dependent in $M/(\bigcup_{i=1}^{t-1} C_i)$. Since A is disjoint from Y but $C_t \cap Y \neq \emptyset$, the dependent set $A - (\bigcup_{i=1}^{t-1} C_i)$ is also a strict subset of the circuit C_t, a contradiction.

Let B_X be a basis of X containing C_X, and let B_Y be a basis of Y containing C_Y in $M/C - D$. Assume there exists a finite set F such that $(B_X \cup B_Y) - F$ is a basis of $M/C - D$. Then for all $f \in F$, there exists a (fundamental) circuit $A_f \subseteq B_X \cup B_Y$ of $M/C - D$ with $A_f \cap F = \{f\}$. Since the circuits A_f are finite and the C_i pairwise disjoint, we may choose t large enough so that $C_t \cap A_f = \emptyset$ for all $f \in F$. Lemma 8.12 ensures the existence of a circuit $K \subseteq C_X \cup C_Y$ in $M/C - D$ with $K - \bigcup_{f \in F} A_f \neq \emptyset$. By the finite circuit exchange axiom $(C2)$, there exists a circuit contained in $(K \cup \bigcup_{f \in F} A_f) - F \subseteq (B_X \cup B_Y) - F$, a contradiction. It follows that $\kappa_{M/C-D}(X, Y) = \infty$, as claimed.

We now consider the case when $\kappa_M(X, Y) = k < \infty$. By repeatedly applying Lemma 8.25, there exists a set $X' \subseteq X$ and $Y' \subseteq Y$ such that $\kappa_M(X', Y') = k$ and $|X'| = |Y'| = k$. (Observe that $\kappa(X', Y') = k$ implies $|X'|, |Y'| \geq k$.) We shall find a partition (C', D') of $E(M) - (X' \cup Y')$ such that $\kappa_{M/C'-D'}(X', Y') = k$. Then, Lemma 8.24 implies that setting $C = C' - (X \cup Y)$ and $D = D' - (X \cup Y)$ results in $\kappa_{M/C-D}(X, Y) = k$ as desired.

In order to find such C' and D' we will inductively define for $t \leq k$ finite sets $Z_t \subseteq E(M)$ with $Z_{t-1} \subseteq Z_t$ such that in the restriction $M|Z_t$ it holds that $\kappa_{M|Z_t}(X', Y') \geq t$. For $t = 1$, pick a circuit A intersecting both X' and Y', and let $Z_1 = X' \cup Y' \cup A$. As M is finitary Z_1 is finite, and its choice ensures $\kappa_{M|Z_1}(X', Y') \geq 1$.

Assume that for $t < k$ we have defined Z_{t-1}, and and observe that as Z_{t-1} is finite, there are only finitely many t-separations in $M|Z_{t-1}$ separating X' and Y', all of which are exact. By applying Lemma 8.27 to each of those, we conclude that there exists a finite set of circuits $A_1, A_2, \ldots, A_l$ such that for $Z_t := Z_{t-1} \cup A_1 \cup \cdots \cup A_l$ we get $\kappa_{M|Z_t}(X', Y') \geq t$.

To conclude, note that the matroid $M|Z_k$ is finite and that $\kappa_{M|Z_k}(X', Y') = k$. By Theorem 8.1, there exists a partition $(C', \tilde{D})$ of $Z_k - (X' \cup Y')$ such that $\kappa_{(M|Z_k)/C'-\tilde{D}}(X', Y') = k$. Consequently, we obtain $\kappa_{M/C'-D'}(X', Y') = k$ for $D' := \tilde{D} \cup (E(M) - Z_k)$, which completes the proof of the theorem. $\square$

8.9 μ-admissibility is not sufficient for matchability

Let $\mathcal{A}$ and $\mathcal{B}$ be two matroids on the same ground set E. A *matching of* $(\mathcal{A}, \mathcal{B})$ is a subset $S \subseteq E$ that is spanning in $\mathcal{A}$ and independent in $\mathcal{B}$. This definition, due to Aharoni and Ziv [4], draws its motivation from matchings in bipartite graphs. Indeed, let G be a bipartite graph with partition classes (A, B) and edge set E, and consider $\mathcal{A}$ to be the *partition matroid* on A, i.e. the matroid in which a set $I \subseteq E$ is independent if and only if no two edges in I have a common endvertex in A, and let $\mathcal{B}$ be the partition matroid on B. Then it can be easily checked that a matching of the matroid pair $(\mathcal{A}, \mathcal{B})$ is, or rather contains, a matching of A into B.

Hall's marriage theorem characterises when the partition class A can be matched into B if the bipartite graph is finite. Nash-Williams [64] among others extended Hall's theorem to countable graphs. Inspired by Nash-Williams' approach, Wojciechowski developed a Hall-like condition first for infinite graphs [94], and then for matchings of matroids [95]. This condition, called μ-admissibility, turns out to be necessary in general, and in the case of bipartite graphs it is also sufficient. The purpose of this section, which is based on [25], is to demonstrate that μ-admissibility fails to be sufficient for matchings of matroids.

Let $\mathcal{A}$ and $\mathcal{B}$ be two matroids on the same ground set E. We call an injective function f a *string (in E)* if its domain is an ordinal and $\operatorname{rge} f \subseteq E$.

For an ordinal θ we write f_θ for the restriction of f to ordinals $\gamma < \theta$. We call $e = f(\theta)$ *f-positive* if e is not spanned by rge f_θ in $\mathcal{B}$, and e is *f-negative* if e is not spanned by $E -$ rge $f_{\theta+1}$ in $\mathcal{A}$. For each string we define by transfinite induction a number in $\mathbb{Z} \cup \{-\infty, \infty\}$ as follows. Set $\mu(f_0) = 0$, let

$$\mu(f_{\theta+1}) = \begin{cases} \mu(f_\theta) + 1 & \text{if } f(\theta) \text{ is } f\text{-positive but not } f\text{-negative,} \\ \mu(f_\theta) - 1 & \text{if } f(\theta) \text{ is } f\text{-negative but not } f\text{-positive, and} \\ \mu(f_\theta) & \text{otherwise,} \end{cases}$$

and $\mu(f_\lambda) = \liminf_{\theta<\lambda} \mu(f_\theta)$ if λ is a limit ordinal. The pair of matroids $(\mathcal{A}, \mathcal{B})$ is said to be *μ-admissible* if $\mu(f) \geq 0$ for all strings f in E.

We remark that for finite matroids $\mathcal{A}$ and $\mathcal{B}$ μ-admissibility reduces to the following condition

$$r_{\mathcal{B}}(W) \geq r_{\mathcal{A}}(E) - r_{\mathcal{A}}(E - W) \text{ for all } W \subseteq E. \tag{8.8}$$

It follows from Edmonds' [38] matroid intersection theorem that this condition is necessary and sufficient for $(\mathcal{A}, \mathcal{B})$ to have a matching (provided $\mathcal{A}$ and $\mathcal{B}$ are finite).

Let us say that a matroid M *is SCF* if it is the countable sum of a family of matroids of finite rank.

Theorem 8.28 (Wojciechowski [95])**.** *Let $\mathcal{A}$ and $\mathcal{B}$ be matroids on the same ground set. If $(\mathcal{A}, \mathcal{B})$ has a matching then $(\mathcal{A}, \mathcal{B})$ is μ-admissible. Conversely, if $\mathcal{A}$ is SCF and $\mathcal{B}$ finitary, and if $(\mathcal{A}, \mathcal{B})$ is μ-admissible then $(\mathcal{A}, \mathcal{B})$ has a matching.*

We shall prove that μ-admissibility is not sufficient for the existence of a matching in two steps. In the first step we introduce a necessary condition for the existence of k disjoint bases in a matroid and show that this condition, called π_k-admissibility, fails to be sufficient. In the second step we relate π_k-admissibility of a matroid M to μ-admissibility in a specific pair of matroids.

The following theorem is the matroidal variant of the well-known Nash-Williams/Tutte tree packing theorem (Theorem 1.6). The classical theorem states that a graph has k edge-disjoint spanning trees if and only if for every partition $\mathcal{P}$ of the vertex set the number of crossing edges is at least $k(|\mathcal{P}|-1)$.

Theorem 8.29. *A finite matroid M on ground set E has k disjoint bases if and only if*

$$|W| \geq k \cdot r_{M.W}(W) \text{ for all } W \subseteq E.$$

In a similar way as μ-admissibility tries to generalise (8.8) we define a measure with the aim to extend the condition in the theorem. Let M be a matroid on a ground set E, let k be a positive integer, and let g be a string in E. Starting with $\pi_k(g_0) = 0$ we set by transfinite induction

$$\pi_k(g_{\theta+1}) = \pi_k(g_\theta) + 1 - k \cdot r_{M.\mathrm{rge}\ g_{\theta+1}}(g(\theta))$$

and $\pi_k(g_\lambda) = \liminf_{\theta<\lambda} \pi_k(g_\theta)$ if λ is a limit ordinal. The matroid M is then π_k-*admissible* if $\pi_k(g) \geq 0$ for all strings g in E.

Now, π_k-admissibility seems a very plausible necessary and sufficient condition for the existence of k disjoint bases in a (countable, say) matroid. While one can check that π_k-admissibility is indeed necessary, there are unfortunately matroids without k disjoint bases that are π_k-admissible.

More precisely, we shall show:

Lemma 8.30. **[25]** *If G is a $2k$-edge-connected graph then $M_{ST}(G)$ is π_k-admissible.*

As Aharoni and Thomassen [3] constructed a $2k$-edge-connected graph without even two edge-disjoint spanning trees we obtain the following corollary.

Corollary 8.31. **[25]** *There is a (finitary) π_k-admissible matroid in which every k bases meet.*

For subsets U, W of the vertex set of a graph G, we denote by $d(U, W)$ the number (which may be infinite) of edges with one endvertex in U and the other in W. We write $d(U) := d(U, V(G) - U)$.

Proof of Lemma 8.30. Let g be a string in $E(G)$ with domain κ. Define $\mathcal{P}(g)$ to be the partition of $V(G)$ given by the components of $G - \mathrm{rge}\ g$. Define $\varepsilon(g)$ to be the number of edges that have both their endvertices in the same element of $P(g)$ (i.e. $\varepsilon(g)$ counts the edges that do not go across the partition classes). We will prove by transfinite induction that for each $\theta < \kappa$ it holds that

$$\pi_k(g_\theta) \geq \varepsilon(g_\theta) + \sum_{U\in\mathcal{P}(g_\theta)} \left(\frac{1}{2}d(U) - k\right). \tag{8.9}$$

Since G is $2k$-edge-connected the right side of the equation will always be non-negative, which then proves the lemma.

In order to see that (8.9) holds true, let us first consider a successor ordinal, $\theta + 1$ say, and let us assume that (8.9) is true for θ. We distinguish two cases.

First, assume that $\mathcal{P}(g_{\theta+1}) = \mathcal{P}(g_\theta)$. Then, the edge $g(\theta)$ must have both its endvertices in the same partition class, which implies $\varepsilon(g_{\theta+1}) = \varepsilon(g_\theta) + 1$.

Moreover, $g(\theta)$ is a loop in $M_{\mathrm{ST}}(G).\mathrm{rge}\ g_{\theta+1}$, and thus we obtain $\pi_k(g_{\theta+1}) = \pi_k(g_\theta) + 1$. As the sum over the partition in (8.9) did not change, we get that (8.9) holds for $\theta + 1$ as well.

Second, consider the case when $\mathcal{P}(g_{\theta+1}) \neq \mathcal{P}(g_\theta)$. Let $W \in \mathcal{P}(g_\theta)$ be such that $\mathcal{P}(g_\theta) - \{W\} = \mathcal{P}(g_{\theta+1}) - \{W_1, W_2\}$, i.e. deleting $g(\theta)$ splits W into W_1 and W_2. If $\pi_k(g_\theta) = \infty$ then clearly $\pi_k(g_{\theta+1}) = \infty$, and (8.9) holds for $\theta + 1$, too. So, let $\pi_k(g_\theta) < \infty$. Now, as $g(\theta)$ has one endvertex in W_1 and the other in W_2, it follows that $\pi_k(g_{\theta+1}) = \pi_k(g_\theta) + 1 - k$. On the other hand, we have $\varepsilon(g_\theta) = \varepsilon(g_{\theta+1}) + d(W_1, W_2) - 1$ and $d(W_1) + d(W_2) = d(W) + 2d(W_1, W_2)$. As $\pi_k(g_\theta) < \infty$ implies $d(W_1, W_2) < \infty$ by (8.9), we get

$$\begin{aligned}
\pi_k(g_{\theta+1}) &= \pi_k(g_\theta) + 1 - k \\
&\geq \varepsilon(g_\theta) + \sum_{U \in \mathcal{P}(g_\theta)} \left(\frac{1}{2} d(U) - k\right) + 1 - k \\
&= \varepsilon(g_{\theta+1}) + d(W_1, W_2) + \sum_{U \in \mathcal{P}(g_\theta)} \left(\frac{1}{2} d(U) - k\right) - k \\
&= \varepsilon(g_{\theta+1}) + d(W_1, W_2) + \sum_{U \in \mathcal{P}(g_\theta) - \{W\}} \left(\frac{1}{2} d(U) - k\right) \\
&\quad + \left(\frac{1}{2} d(W_1) + \frac{1}{2} d(W_2) - d(W_1, W_2) - k\right) - k \\
&= \varepsilon(g_{\theta+1}) + \sum_{U \in \mathcal{P}(g_{\theta+1})} \left(\frac{1}{2} d(U) - k\right).
\end{aligned}$$

This proves that (8.9) holds for successor ordinals.

Now, let λ be a limit ordinal. Set $\rho(g_\theta) := \varepsilon(g_\theta) + \sum_{U \in \mathcal{P}(g_\theta)} \left(\frac{1}{2} d(U) - k\right)$. In order show that $\pi_k(g_\lambda) \geq \rho(g_\lambda)$ it suffices to prove that $\liminf_{\theta < \lambda} \rho(g_\theta) \geq \rho(g_\lambda)$. For this purpose pick any $K \in \mathbb{Z}$ so that $\rho(g_\lambda) \geq K$. Then there exists a finite set $F_1 \subseteq \mathrm{rge}\ g_\lambda$ of edges with both their endvertices in the same element of $\mathcal{P}(g_\lambda)$, i.e. these are non-crossing edges, and a finite set $F_2 \subseteq \mathrm{rge}\ g_\lambda$ of (crossing) edges with their endvertices in different partition classes so that $|F_1| + \sum_{U \in \mathcal{P}'} \left(\frac{1}{2} |E(U, E - U) \cap F_2| - k\right) \geq K$, where $\mathcal{P}'$ is the set of those $U \in \mathcal{P}(g_\lambda)$ for which $E(U, E - U) \cap F_2 \neq \emptyset$. Then it is easy to check that for any ordinal $\alpha < \lambda$ with $F_1 \cup F_2 \subseteq \mathrm{rge}\ \alpha$ it holds that $\rho(g_\alpha) \geq K$, which proves that $\liminf_{\theta < \lambda} \rho(g_\theta) \geq K$. $\square$

Next, starting with a matroid M on a ground set E and $k \in \mathbb{N}$ we construct a pair of matroids $(\mathcal{A}, \mathcal{B})$ so that a matching in $(\mathcal{A}, \mathcal{B})$ yields k disjoint bases in M. For this, set $S := E \times \{1, \ldots, k\}$, and define M_i, $i =$

$1, \ldots, k$, to be a copy of M on $E \times \{i\}$. Denote by $\mathcal{A}(M,k)$ the disjoint union of $M_1, \ldots, M_k$ and by $\mathcal{B}(M,k)$ the partition matroid on S with partition classes $\{(e,i) : 1 \leq i \leq k\}$. We remark that $\mathcal{A}(M,k)$ is finitary (resp. co-finitary) if M is finitary (resp. co-finitary), and that $\mathcal{B}(M,k)$ is always SCF (assuming that E is countable).

Lemma 8.32. **[25]** *Let M be a matroid, and let $k \in \mathbb{N}$. If M is π_k-admissible then $(\mathcal{A}(M,k), \mathcal{B}(M,k))$ is μ-admissible.*

Since a matching in $(\mathcal{A}(M,k), \mathcal{B}(M,k))$ yields k disjoint spanning sets in M, the lemma applied to Corollary 8.31 implies the following result:

Proposition 8.33. **[25]** *There is a finitary matroid $\mathcal{A}$, and a SCF matroid $\mathcal{B}$ so that $(\mathcal{A}, \mathcal{B})$ is μ-admissible but not matchable.*

Proof of Lemma 8.32. Let f be a string in $S = E \times \{1, \ldots, k\}$. There is, for every $i \in \{1, \ldots, k\}$, a function γ^i mapping ordinals to ordinals, which counts the number of $f(\theta)$ in $E \times \{i\}$, and a string $\bar{g}^i$ in E so that $f(\theta) = \bar{g}^i(\gamma^i(\theta))$ if $f(\theta) \in E \times \{i\}$. Abusing notation we treat $g^i := \bar{g}^i \circ \gamma^i$ as string in E and pretend that expressions like $\pi_k(g^i_\theta)$ are defined, when in fact we mean $\pi_k(\bar{g}^i_{\gamma^i(\theta)})$. For any ordinal θ denote by $p(f_\theta)$ the number (which may be infinite) of $(e,i) \notin \text{rge } f_\theta$ for which at least one of $(e,1), \ldots, (e,k)$ lies in rge f_θ.

We claim that for all ordinals θ in the domain of f it holds that

$$\mu(f_\theta) \geq \frac{1}{k}\left(p(f_\theta) + \sum_{i=1}^{k} \pi_k(g^i_\theta)\right). \tag{8.10}$$

If (8.10) is true then the lemma is proved, as M is π_k-admissible and as $p(f_\theta)$ is never negative.

We prove (8.10) by transfinite induction. Consider first a successor ordinal $\theta + 1$ and assume (8.10) to hold for θ. Let $(e,i) = f(\theta)$ for $e \in E$ and $i \in \{1, \ldots, l\}$. Observe first that (e,i) is f-positive if and only if $\{(e,1), \ldots, (e,k)\} \cap \text{rge } f_{\theta+1} = \{(e,i)\}$, i.e. if θ is the first ordinal for which f chooses a copy of e. (Recall that $f(\theta)$ is f-positive precisely when it is not spanned by rge f_θ in $\mathcal{B}$.) So, it follows that

$$\begin{array}{l} \textit{if } (e,i) \textit{ is } f\textit{-positive then } p(f_{\theta+1}) = p(f_\theta) + k - 1; \\ \textit{if } (e,i) \textit{ is not } f\textit{-positive then } p(f_{\theta+1}) = p(f_\theta) - 1. \end{array} \tag{8.11}$$

Next, (e,i) is defined to be f-negative when (e,i) is not spanned by $S - \text{rge } f_{\theta+1}$ in $\mathcal{A}$. Since $\mathcal{A}$ is the union of the disjoint copies $M_1, \ldots, M_k$ of M it follows that (e,i) is f-negative if and only e is not spanned by

$E - \text{rge } g^i_{\theta+1}$ in M, i.e. if e is not a loop in $M.\text{rge } g^i_{\theta+1}$. Thus, (e,i) is f-negative if and only if $r_{M.\text{rge } g^i_{\theta+1}}(e) = 1$, which leads to

$$\begin{array}{l} \textit{if } (e,i) \textit{ is } f\textit{-negative then } \pi_k(g^i_{\theta+1}) = \pi_k(g^i_\theta) + 1 - k; \\ \textit{if } (e,i) \textit{ is not } f\textit{-negative then } \pi_k(g^i_{\theta+1}) = \pi_k(g^i_\theta) + 1. \end{array} \tag{8.12}$$

Moreover, whether (e,i) is f-negative or not we have $\pi_k(g^j_{\theta+1}) = \pi_k(g^j_\theta)$ for $j \neq i$ as $\text{rge } g^j_{\theta+1} = \text{rge } g^j_\theta$.

Assume that (e,i) is f-positive but not f-negative. Then $\mu(f_{\theta+1}) = \mu(f_\theta) + 1$, and (8.11) and (8.12) together with (8.10) for θ yield

$$\begin{aligned} \frac{1}{k}\left(p(f_{\theta+1}) + \sum_{j=1}^k \pi_k(g^j_{\theta+1})\right) &= \frac{1}{k}\left(p(f_\theta) + k - 1 + \sum_{j=1}^k \pi_k(g^j_\theta) + 1\right) \\ &= \frac{1}{k}\left(p(f_\theta) + \sum_{j=1}^k \pi_k(g^j_\theta)\right) + 1 \\ &\leq \mu(f_\theta) + 1 \quad = \quad \mu(f_{\theta+1}). \end{aligned}$$

The other cases follow in a similar way.

So, let λ be a limit ordinal and assume (8.10) to hold for all $\theta < \lambda$. As $\liminf_{\theta<\lambda} p(f_\theta) \geq p(f_\lambda)$, inequality (8.10) follows easily. □

8.10 Graph and matroid duality

As briefly touched upon in Section 8.2 it is possible to associate a finitary matroid, the matroid $M_{\text{ST}}(G)$, as well as a co-finitary matroid, called $M_{\text{TST}}(G)$, with a graph G. While $M_{\text{ST}}(G)$ is based on the spanning trees of G (provided G is connected), the bases of $M_{\text{TST}}(G)$ are the (edge sets of) topological spanning trees. More precisely, if G satisfies (1.1) define $\mathcal{I}_{\text{TST}}(G)$ to be the set of all edge sets I so that I is contained in a subspace of $\tilde{G}$ that is a TST on each component of G. Let $M_{\text{TST}}(G) = (E, \mathcal{I}_{\text{TST}}(G))$. We note that the circuits of the matroid $M_{\text{TST}}(G)$ are precisely the edge sets of circles in $\tilde{G}$.

Let us define two more matroids for G. Define $M_{\text{CUT}}(G)$ to be the matroid on $E(G)$ in which a subset of $E(G)$ is independent if and only if it does not contain any (edge-)cut. A related matroid is $M_{\text{FCUT}}(G)$: the matroid on $E(G)$ in which an edge set is precisely then independent if it does not contain any *finite* cut. It is easy to see that $M_{\text{CUT}}(G)$ and $M_{\text{ST}}(G)$ are dual matroids. A similar relationship exists between $M_{\text{TST}}(G)$ and the finitary matroid $M_{\text{FCUT}}(G)$.

Proposition 8.34. [18] *Let G be a graph.*

(i) *It holds that* $(M_{ST}(G))^* = M_{CUT}(G)$.

(ii) *If* G *satisfies* (1.1) *then* $(M_{TST}(G))^* = M_{FCUT}(G)$.

Proof. We may assume that $G = (V, E)$ is connected.

(ii) Let I be an independent set in $(M_{\mathrm{TST}}(G))^*$, which means there exists a TST with edge set T so that I and T are disjoint. As a TST is connected in the space $\tilde{G}$ it needs to meet every finite cut (this can be seen as in Lemma 8.5.5 in [31]). Hence, I cannot contain any finite cut and is thus independent in $M_{\mathrm{FCUT}}(G)$. Conversely, consider an independent set J of $M_{\mathrm{FCUT}}(G)$. Then the topological space $\tilde{G} - J$ (here, we only delete the interior points of the edges in J) is still topologically connected, since J does not contain any finite cut; see Lemma 8.5.12 in [31]. On the other hand, Lemma 8.5.13 in [31] asserts that in a connected subspace such as $\tilde{G} - J$ that contains every vertex we can always find a TST. Clearly, none of the edges of this TST lie in J, showing that J is independent in $(M_{\mathrm{ST}}(G))^*$.

The proof of (i) is similar. $\square$

Let G be a graph satisfying (1.1). Recall from Chapter 2 that a graph G^* is called the *dual of* G if there is a bijection $* : E(G) \to E(G^*)$ so that every set $F \subseteq E(G)$ is a bond of G, i.e. a non-empty minimal cut, if and only F^* is the edge set of a circle (in $\tilde{G}^*$).

As stated in Chapter 2, the duals defined in this way have similar properties to their well-known finite counterparts. In particular, a graph has a dual if and only if the graph is planar; a graph is the dual of its dual; and the dual of a 3-connected graph is unique (or non-existent).

The motivation of this definition of graph duality is of purely graph-theoretical nature. It appears to be the only natural way to force all the three facts mentioned above to become true in infinite graphs. Somewhat surprisingly, graph duality is compatible with matroid duality:

Corollary 8.35. **[18]** *Let* G *and* G^* *be a pair of dual graphs (each satisfying* (1.1)*) defined on the same edge set* E*. Then*

$$M_{ST}(G) = (M_{CUT}(G))^* = M_{FCUT}(G^*) = (M_{TST}(G^*))^*.$$

Proof. The first and the last equality are due to Proposition 8.34. The remaining equality follows from $M_{\mathrm{CUT}}(G) = M_{\mathrm{TST}}(G^*)$, which is a direct consequence of the definition of a dual graph. $\square$

8.11 Matroid and graph connectivity

In this final section of this chapter, and indeed of this thesis, we will determine the connectivity function of the two matroids, $M_{\mathrm{ST}}(G)$ and $M_{\mathrm{TST}}(G)$, associated with a graph G. In particular, we shall find in Theorem 8.36 that the connectivity function κ of $M_{\mathrm{ST}}(G)$ behaves exactly as in finite graphs, a result which yields further evidence that κ was defined in the right way. Moreover, we will see in Theorem 8.41 that $M_{\mathrm{ST}}(G)$ and $M_{\mathrm{TST}}(G)$ cannot be distinguished by the connectivity function. Finally, returning to the beginning of this thesis we will relate k-connectivity in the matroid $M_{\mathrm{ST}}(G)$ to Tutte-connectivity. This will allow us to give a matroidal proof of Theorem 2.16, ie of the invariance of Tutte-connectivity under taking duals. This section is based on [14].

In a graph G, denote for $X \subseteq E(G)$ by $V[X]$ the set of vertices that are incident with an edge in X. Let $c(X)$ be the number of components of the subgraph $(V[X], X)$ of G.

Our first aim is the following theorem:

Theorem 8.36. **[14]** *Let G be a 2-connected graph satisfying* (1.1)*, and let $X \subseteq E(G)$, and $Y := E(G) - X$. Then the following statements hold:*

(i) $\kappa_{M_{ST}(G)}(X) = \infty$ *if and only* $|V[X] \cap V[Y]| = \infty$*; and*

(ii) *if* $\kappa_{M_{ST}(G)}(X) < \infty$ *then*

$$\kappa_{M_{ST}(G)}(X) = |V[X] \cap V[Y]| - c(X) - c(Y) + 1.$$

Statement (ii) is exactly as for finite graphs when the traditional connectivity function is used, see Tutte [85]. We shall need two lemmas for the proof of Theorem 8.36.

Lemma 8.37. **[14]** *Let G be a graph satisfying* (1.1)*, and let $\mathcal{D}$ be an infinite set of edge-disjoint finite cycles. Then there exists an infinite subset $\mathcal{D}'$ of $\mathcal{D}$ and a vertex v of G so that any two distinct cycles in $\mathcal{D}'$ are disjoint outside v.*

Proof. Let $C_1, C_2, \ldots$ be an enumeration of (countably many of) the cycles in $\mathcal{D}$. Inductively we will delete certain cycles from $\mathcal{D}$ while ensuring in each step that we keep infinitely many cycles. In step i, assuming C_i has not been deleted, we go through the finitely many vertices of C_i, one by one. Then for a vertex w of C_i, unless w lies in all but finitely many of the remaining C_j, we delete from $\mathcal{D}$ all those C_j that contain w. If w lies in all but finitely many of the remaining C_j we skip to the next vertex of C_i without deleting any cycles. Denote the resulting infinite subset of $\mathcal{D}'$ by $\mathcal{D}$.

Now, if the cycles in $\mathcal{D}'$ are pairwise disjoint, choose any vertex of G for v and observe that this choice of $\mathcal{D}'$ and v is as desired. So, assume that there is a vertex shared by two cycles in $\mathcal{D}'$. Pick the smallest index i for which there is a $j \neq i$ so that C_i and C_j have a vertex, v say, in common, and so that $C_i, C_j \in \mathcal{D}'$. Note that v, as well as any other vertex that lies in two cycles of $\mathcal{D}'$, is contained in infinitely many cycles in $\mathcal{D}'$; otherwise we would have deleted all but one of those cycles incident with v.

Suppose there exists a second vertex w contained in two cycles of $\mathcal{D}'$. If k is the lowest index with $w \in V(C_k)$ and $C_k \in \mathcal{D}'$ then why have we not deleted all those cycles C_l containing w with $l > k$ from $\mathcal{D}'$ in step k? Precisely because all but finitely many of the cycles in $\mathcal{D}'$ contain w. In particular, infinitely many of those cycles in $\mathcal{D}'$ that contain v must also contain w. By picking a v–w path in each of those cycles we obtain infinitely many edge-disjoint v–w paths in contradiction of (1.1). $\square$

Lemma 8.38. **[14]** *Let G be a 2-connected graph satisfying* (1.1)*, and let $X \subseteq E(G)$, and $Y := E(G) - X$. If $|V[X] \cap V[Y]| = \infty$ then $\kappa_{M_{ST}(G)}(X) = \kappa_{M_{TST}(G)}(X) = \infty$.*

Proof. For each vertex in $V[X] \cap V[Y]$ pick one incident edge in X and one in Y; denote the set of these edges by $X' \subseteq X$ and $Y' \subseteq Y$, respectively. Lemma 2.9 yields an end $\omega \in \overline{X'}$ (where the closure is taken in $|G|$). It is easy to check that then also $\omega \in \overline{Y'}$ (again with respect to $|G|$).

By Lemma 2.10, there exists an infinite set $\mathcal{D}$ of edge-disjoint finite cycles each of which meets both X' and Y'. Applying Lemma 8.37 yields a vertex v and an infinite subset $\mathcal{D}'$ of $\mathcal{D}$ so that any two cycles either meet only in v or not at all. As every cycle in $\mathcal{D}'$ contains an edge in X as well as in Y it follows that neither $I_X := X \cap \bigcup_{C \in \mathcal{D}'} E(C)$ nor $I_Y := Y \cap \bigcup_{C \in \mathcal{D}'} E(C)$ contains a (finite or infinite) circuit. To see that neither set contains an infinite circuit observe that as all the $C \in \mathcal{D}'$ are finite neither $(V[I_X], I_X)$ nor $(V[I_Y], I_Y)$ contain a ray.

Thus I_X and I_Y are independent in both matroids $M_{\mathrm{ST}}(G)$ and $M_{\mathrm{TST}}(G)$. Let T_X be a basis of $M|X$ containing I_X, and let $T_Y \supseteq I_Y$ be a basis of $M|Y$, where M is either $M_{\mathrm{ST}}(G)$ or $M_{\mathrm{TST}}(G)$. Choose $F \subseteq T_X \cup T_Y$ so that $(T_X \cup T_Y) - F$ is a basis of M. Since $I_X \cup I_Y$ contains the (edge-)disjoint circuits $E(C)$, $C \in \mathcal{D}'$, F must contain at least one edge from each of those infinitely many circuits. Hence $\kappa_M(X) = |F| = \infty$. $\square$

Proof of Theorem 8.36. (i) By Lemma 8.38 we only need to consider the case when $|V[X] \cap V[Y]| < \infty$. Pick a basis T_X of $M_{\mathrm{ST}}(G)|X$, and let T_Y

be a basis of $M_{\mathrm{ST}}(G)|Y$. For every of the finitely many pairs of vertices $u, v \in V[X] \cap V[Y]$ there is by (1.1) a finite set of edges separating u from v in $(V[X], X)$. Denote by F the union of all those edges, and observe that F is a finite edge set. By the choice of F the set $(T_X \cup T_Y) - F$ cannot contain any finite circuit, and is thus independent in $M_{\mathrm{ST}}(G)$. As $|F| < \infty$ is therefore an upper bound for $\kappa_{M_{\mathrm{ST}}(G)}(X)$ the result follows.

(ii) Pick a spanning tree on every component of $(V[X], X)$ and denote the union of their edge sets by T_X. We define T_Y for $(V[Y], Y)$ in a similar way.

We claim that

$$\textit{if } c(X) = c(Y) = 1 \textit{ then } \kappa_{M_{ST(G)}}(X, Y) = |V[X] \cap V[Y]| - 1. \tag{8.13}$$

Let us prove the claim. Each vertex of $U := V[X] \cap V[Y]$ must lie in a distinct component of $(V[T_X], T_X) - F$ since otherwise there exists a U-path in $(V[T_X], T_X)$ that misses F. This path can be extended with edges in T_Y to a cycle that still misses F, a contradiction. As $(V[T_X], T_X)$ is connected and as each deletion of a single edge increases the number of components by at exactly one, we obtain $|F| \geq |U| - 1$. Suppose, on the other hand, that $|F| > |U| - 1$. Then there exists a component of $(V[T_X], T_X) - F$ that contains no vertex of U. Pick an edge $e \in F$ with one of its endvertices in this component. Setting $T := (T_X - F) \cup T_Y$, we observe that $\{e\}$ is a cut of $(V[T], T + e)$. However, as T is (the edge set of) a spanning tree of G, there has to be a circuit in $T + e$ containing e, a contradiction. This proves the claim.

We now proceed by induction on $c(X) + c(Y)$, which is indeed a finite number as $|V[X] \cap V[Y]|$ is an upper bound for both $c(X)$ and $c(Y)$. Since the induction start is established by (8.13), we may assume that $(V[X], X)$ has two components K and K'. Insert a new edge f between K and K', and set $G' := G + f$ and $X' := X \cup \{f\}$. Clearly, (X', Y) is a partition of $E(G')$. Since $c(X') = c(X) - 1$, the induction yields

$$\kappa_{M_{\mathrm{ST}}(G')}(X', Y) = |V[X] \cap V[Y]| - (c(X) - 1) - c(Y) + 1.$$

We shall now show that $\kappa_{M_{\mathrm{ST}}(G')}(X', Y) = \kappa_{M_{\mathrm{ST}}(G)}(X, Y) + 1$. Observe that then $T_X + f$ is (the edge set of) a maximal spanning forest of $(V[X'], X') \subseteq G'$. Moreover, $(T_X - F) \cup T_Y = ((T_X + f) - (F \cup \{f\})) \cup T_Y$ is a spanning tree of G', too. Thus

$$\kappa_{M_{\mathrm{ST}}(G')}(X, Y') = |F \cup \{f\}| = |F| + 1 = \kappa_{M_{\mathrm{ST}}(G)}(X, Y) + 1,$$

which finishes the proof. $\square$

Next, let us show that the connectivity functions of $M_{\mathrm{ST}}(G)$ and $M_{\mathrm{TST}}(G)$ coincide. For this, we should be able to modify the proof of Theorem 8.36 in order to make it work for $M_{\mathrm{TST}}(G)$, too. Rather then repeating the argument we will pursue a different approach, for which we will need a small lemma and a result from [16].

Lemma 8.39. [14] *Let G be a graph satisfying* (1.1)*, and let H be an induced subgraph of G so that $N(G-H)$ is a finite set. Then if $Z \subseteq E(H)$ lies in $\mathcal{C}(\tilde{G})$ then also $Z \in \mathcal{C}(\tilde{H})$.*

Proof. Consider $Z \subseteq E(H)$ so that $Z \notin \mathcal{C}(\tilde{H})$. By Theorem 1.9 there is a finite cut F of H so that $Z \cap F$ is an odd set. The cut F partitions $N(G-H)$ into two sets A and B (one of them possibly empty). Since every two vertices in G can separated by finitely many edges there is a finite subset of $E(G)-E(H)$ that separates A from B in $G[(G-H) \cup N(G-H)]$. Choosing a minimal such set F' ensures that $F \cup F'$ is a cut of G. Then $|Z \cap (F \cup F')| = |Z \cap F|$ is odd, implying with Theorem 1.9 that $Z \notin \mathcal{C}(\tilde{G})$. $\square$

Theorem 8.40. [16] *Every connected graph G satisfying* (1.1) *has a spanning tree that does not contain any (infinite) circuit (with respect to $\mathcal{C}(\tilde{G})$).*

Theorem 8.41. [14] *Let G be a 2-connected graph satisfying* (1.1)*. Then $\kappa_{M_{ST}(G)}(X) = \kappa_{M_{TST}(G)}(X)$ for all $X \subseteq E(G)$.*

Proof. Consider a set $X \subseteq E(G)$ and put $Y := E(G) - X$. If $V[X] \cap V[Y]$ is an infinite set then $\kappa_{M_{\mathrm{ST}}(G)}(X) = \kappa_{M_{\mathrm{TST}}(G)}(X)$ by Lemma 8.38.

So, assume $V[X] \cap V[Y]$ to be finite. By Theorem 8.40 there is for each component K of $(V[X], X)$ a spanning tree not containing any circuit of $\mathcal{C}(\tilde{K})$. Lemma 8.39 ensures that this spanning tree is also acirclic (ie, without any circuits with respect to $\mathcal{C}(\tilde{G})$). Consequently, the union T_X of the edge sets of those spanning trees is a basis of $M_{\mathrm{ST}}(G)|X$ as well as of $M_{\mathrm{TST}}(G)|X$. We define T_Y analogously for $(V[Y], Y)$.

Next, pick $F \subseteq T_X \cup T_Y$ so that $(T_X \cup T_Y) - F$ is a basis of $M_{\mathrm{TST}}(G)$. Clearly, the set $(T_X \cup T_Y) - F$ is independent in $M_{\mathrm{ST}}(G)$, too. If it is even a basis in $M_{\mathrm{ST}}(G)$ then we have $\kappa_{M_{\mathrm{ST}}(G)}(X) = |F| = \kappa_{M_{\mathrm{TST}}(G)}(X)$ as desired. So, suppose $T := (T_X \cup T_Y) - F$ fails to be a basis, which implies that $(V[T], T)$ is not (graph-theoretically) connected. Since T is the edge set of a TST in $\tilde{G}$ there must therefore be an arc with infinitely edges between two vertices, so that the arc is completely contained in $\overline{T}$. Since $|V[X] \cap V[Y]| < \infty$ there is then also an arc of infinite length between two vertices in $\overline{T_X}$ or $\overline{T_Y}$, let us say in $\overline{T_X}$. However, as any two vertices in $(V[T_X], T_X)$ are connected by a (finite) path as well this yields a circuit contained in T_X, contradicting the definition of T_X. Thus, $(V[T], T)$ is connected and hence T a basis of $M_{\mathrm{ST}}(G)$. $\square$

Let us recall the definition of Tutte-connectivity. A *k-Tutte-separation* of a graph G is a partition (X, Y) of $E(G)$ so that $|X|, |Y| \geq k$ and so that $|V[X] \cap V[Y]| \leq k$. We say that a graph G is *k-Tutte-connected* if G has no ℓ-Tutte-separation for any $\ell < k$.

Theorem 8.42. **[14]** *Let G be a graph satisfying* (1.1). *Then for integers $k \geq 2$ the following statements are equivalent:*

(i) *G is k-Tutte-connected;*

(ii) *$M_{ST}(G)$ is k-connected; and*

(iii) *$M_{TST}(G)$ is k-connected.*

Proof. Observe that we may assume G to be 2-connected and that G is an infinite graph. (For finite graphs, see Tutte [85]—note that $M_{\mathrm{ST}}(G)$ and $M_{\mathrm{TST}}(G)$ coincide in this case.) In light of Theorem 8.41 we only need to prove that G has a k-Tutte-separation with $k \leq m$ if and only if $M_{\mathrm{ST}}(G)$ has an ℓ-separation with $\ell \leq m$.

First, let (X, Y) be a k-Tutte-separation (X, Y) of G, which implies $|V[X] \cap V[Y]| \leq k$. Since $c(X), c(Y) \geq 1$ this yields with Theorem 8.36 that $\kappa_{M_{\mathrm{ST}}(G)} \leq k - 1$. Consequently, (X, Y) is a k-separation of $M_{\mathrm{ST}}(G)$.

Conversely, let there be an ℓ-separation in $M_{\mathrm{ST}}(G)$, and choose an ℓ-separation (X, Y) of $M_{\mathrm{ST}}(G)$ so that $c(X) + c(Y)$ is minimal among all ℓ-separations of $M_{\mathrm{ST}}(G)$. Since G is infinite, we may assume that Y is an infinite set.

First, we claim that

$$(V[Y], Y) \text{ is connected.} \tag{8.14}$$

If $(V[Y], Y)$ is not connected then there is a component K of $(V[Y], Y)$ so that $Y' := Y - E(K)$ is an infinite set. With $X' := X \cup E(K)$ we see that both X' and Y' have at least ℓ elements. Moreover, it holds that $|V[X] \cap V[Y]| = |V[X'] \cap V[Y']| + |V[X] \cap V[K]|$ and $c(Y) = c(Y') + 1$. The set of components of $(V[X'], X')$ is comprised of components of $(V[X], X)$ and of the union of those components of $(V[X], X)$ that have a vertex with K in common together with K. Since there are at most $|V[X] \cap V[K]|$ components of the latter kind, we obtain $c(X) \leq c(X') + |V[X] \cap V[K]| - 1$. It follows with Theorem 8.36 that

$$\begin{aligned}
\kappa_{M_{\mathrm{ST}}(G)}(X', Y') &= |V[X'] \cap V[Y']| - c(X') - c(Y') + 1 \\
&\leq |V[X] \cap V[Y]| - |V[X] \cap V[K]| - c(X) \\
&\quad + |V[X] \cap V[K]| - 1 - c(Y) + 1 + 1 \\
&= |V[X] \cap V[Y]| - c(X) - c(Y) + 1 \leq \ell - 1.
\end{aligned}$$

Thus, (X', Y') is an ℓ-separation with $c(X') + c(Y') \leq c(X) + c(Y) + 1$, contradicting the choice of (X, Y).

Second, we show that

$$\textit{for every component } K \textit{ of } (V[X], X) \textit{ holds that } |V[K] \cap V[Y]| \leq \ell. \quad (8.15)$$

Suppose there exists a component M of $(V[X], X)$ with $|V[M)] \cap V[Y]| \geq \ell + 1$. Denoting by $\mathcal{K}$ the components of $(V[X], X)$ we get

$$\begin{aligned} \ell - 1 &\geq |V[X] \cap V[Y]| - c(X) - c(Y) + 1 \\ &\geq \sum_{K \in \mathcal{K} - \{M\}} |V[K] \cap V[Y]| + (\ell + 1) - c(X) - c(Y) + 1. \end{aligned}$$

That G is connected implies $|V[K] \cap V[Y]| \geq 1$ for every $K \in \mathcal{K}$. Hence

$$\ell - 1 \geq (c(X) - 1) + (\ell + 1) - c(X) - c(Y) + 1 = \ell + 1 - c(Y).$$

This yields $c(Y) \geq 2$, which is impossible by (8.14). Therefore, (8.15) is proved.

Next, we see that

$$\textit{there is a component } M \textit{ of } (V[X], X) \textit{ with } |E(M)| \geq |V[M] \cap V[Y]|. \quad (8.16)$$

If (8.16) is false then we have $|V[K] \cap V[Y]| \geq |E(K)| + 1$ for all $K \in \mathcal{K}$. This, however, implies with $c(Y) = 1$ that

$$\begin{aligned} \ell - 1 &\geq |V[X] \cap V[Y]| - c(X) - c(Y) + 1 \\ &= \sum_{K \in \mathcal{K}} |V[K] \cap V[Y]| - c(X) \\ &\geq \sum_{K \in \mathcal{K}} (|E(K)| + 1) - c(X) \; = \; |X|. \end{aligned}$$

As (X, Y) is an ℓ-separation, X is required to have at least ℓ elements, which shows that (8.16) holds.

Finally, with the component M from (8.16) we set $\bar{X} := E(M)$ and $\bar{Y} := E(G) - E(M)$. Then $k := |V[\bar{X}] \cap V[\bar{Y}]| = |V[M] \cap V[Y]| \leq \ell$, by (8.15). As $|\bar{X}| \geq k$ and $|\bar{Y}| = \infty$ it follows that $(\bar{X}, \bar{Y})$ is a k-Tutte-separation with $k \leq \ell$, as desired. □

We remark that the arguments in the proof are not new. Indeed, (8.14) is inspired by Tutte [85] and steps (8.15), (8.16) are quite similar to the proof of Lemma 2.14.

The relation between the connectivity of M_{ST} and Tutte-connectivity now yields a matroidal proof for the invariance of Tutte-connectivity under taking duals that we have already discussed in Chapter 2. In fact, if G and G^* is a pair of dual graphs then, by Theorem 8.42, G is k-Tutte-connected if and only if $M_{\mathrm{ST}}(G)$ is k-connected. Since $M_{\mathrm{ST}}(G) = (M_{\mathrm{TST}}(G^*))^*$ by Corollary 8.35 and since matroid connectivity is invariant under taking duals (Proposition 8.21) this is precisely the case when $M_{\mathrm{TST}}(G^*)$ is k-connected. Finally, Theorem 8.42 again shows that $M_{\mathrm{TST}}(G^*)$ is k-connected if and only if G^* is k-Tutte-connected. Therefore we have reproved:

Theorem 2.16. **[22]** *Let G and G^* be a pair of dual graphs, and let $k \geq 2$. Then G is k-Tutte-connected if and only if G^* is k-Tutte-connected.*

Bibliography

[1] April 2009, `http://en.wikipedia.org/wiki/Matroid`.

[2] R. Aharoni and E. Berger, *Menger's theorem for infinite graphs*, to appear in *Inventiones mathematicae*.

[3] R. Aharoni and C. Thomassen, *Infinite, highly connected digraphs with no two arc-disjoint spanning trees*, J. Graph Theory **13** (1989), 71–74.

[4] R. Aharoni and R. Ziv, *The intersection of two infinite matroids*, J. London Math. Soc. **58** (1998), 513–525.

[5] D. Archdeacon, *A Kuratowski theorem for the projective plane*, J. Graph Theory **5** (1981), 243–246.

[6] D. Archdeacon, P. Bonnington, and C. Little, *An algebraic characterization of planar graphs*, J. Graph Theory **19** (1995), 237–250.

[7] M.A. Armstrong, *Basic topology*, Springer-Verlag, 1983.

[8] L. Auslander, T.A. Brown, and J.W.T. Youngs, *The imbedding of graphs in manifolds*, J. Math. and Mech. **12** (1963), 629–634.

[9] D.T. Bean, *A connected finitary and co-finitary matroid is finite*, Proc. 7th southeast. Conf. Comb., Graph Theory, Comput. (Baton Rouge), 1976, pp. 115–119.

[10] A.R. Berg and T. Jordán, *Algorithms for graph rigidity and scene analysis*, Preprint 2003.

[11] E. Berger and H. Bruhn, *Eulerian edge sets in locally finite graphs*, Preprint 2008.

[12] A. Bernáth and H. Bruhn, *Degree constrained orientations in countable graphs*, Electronic J. Comb. **15** (2008), #R122.

[13] C.P. Bonnington and R.B. Richter, *Graphs embedded in the plane with finitely many accumulation points*, J. Graph Theory **44** (2003), 132–147.

[14] H. Bruhn, unpublished.

[15] ______, *The cycle space of a* 3*-connected locally finite graph is generated by its finite and infinite peripheral circuits*, J. Combin. Theory (Series B) **92** (2004), 235–256.

[16] H. Bruhn and R. Diestel, *Duality in infinite graphs*, Comb., Probab. Comput. **15** (2006), 75–90.

[17] ______, *MacLane's theorem for arbitrary surfaces*, (2008), to appear in *J. Combin. Theory (Series B).*

[18] H. Bruhn, R. Diestel, M. Kriesell, and P. Wollan, *A new look at infinite matroids*, Preprint 2009.

[19] H. Bruhn, R. Diestel, and M. Stein, *Cycle-cocycle partitions and faithful cycle covers for locally finite graphs*, J. Graph Theory **50** (2005), 150–161.

[20] H. Bruhn and A. Georgakopoulos, *Bases and closures under infinite sums*, preprint 2007.

[21] H. Bruhn, S. Kosuch, and M. Win Myint, *Bicycles and left-right tours in locally finite graphs*, Europ. J. Combin. **30** (2009), 356–371.

[22] H. Bruhn and M. Stein, *Duality of ends*, Preprint 2008.

[23] ______, *MacLane's planarity criterion for locally finite graphs*, J. Combin. Theory (Series B) **96** (2006), 225–239.

[24] ______, *On end degrees and infinite circuits in locally finite graphs*, Combinatorica **27** (2007), 269–291.

[25] H. Bruhn and J. Wojciechowski, unpublished.

[26] H. Bruhn and X. Yu, *Hamilton circles in planar locally finite graphs*, SIAM J. Discrete Math. **22** (2008), 1381–1392.

[27] K. Casteels and B. Richter, *The bond and cycle spaces of an infinite graph*, J. Graph Theory **59** (2008), 162–176.

[28] R. Christian, B. Richter, and B. Rooney, *The planarity theorems of MacLane and Whitney for graph–like spaces*, Preprint 2009.

[29] Q. Cui, J. Wang, and X. Yu, *Hamilton circles in infinite planar graphs*, Preprint 2007.

[30] R. Diestel, *The cycle space of an infinite graph*, Comb., Probab. Comput. **14** (2005), 59–79.

[31] ______, *Graph theory* (3rd edition), Springer-Verlag, 2005.

[32] ______, *End spaces and spanning trees*, J. Combin. Theory (Series B) **96** (2006), 846–854.

[33] R. Diestel and D. Kühn, *On infinite cycles I*, Combinatorica **24** (2004), 69–89.

[34] ______, *On infinite cycles II*, Combinatorica **24** (2004), 91–116.

[35] ______, *Topological paths, cycles and spanning trees in infinite graphs*, Europ. J. Combinatorics **25** (2004), 835–862.

[36] R. Diestel and P. Sprüssel, *The homology of locally finite graphs with ends*, Preprint 2007.

[37] J. Edmonds, *On the surface duality of linear graphs*, J. Res. Natl. Bur. Stand., Sect. B **69** (1965), 121–123.

[38] ______, *Submodular functions, matroids, and certain polyhedra.*, Combinat. Struct. Appl., Proc. Calgary internat. Conf. combinat. Struct. Appl., Calgary 1969, 1970, pp. 69–87.

[39] L. Egyed, *Ueber die wohlgerichteten unendlichen Graphen*, Math. phys. Lapok **48** (1941), 505–509, in Hungarian with German summary.

[40] L.R. Ford and D.R. Fulkerson, *Flows in networks*, Princeton Univ. Press, 1962.

[41] A. Frank, *On disjoint trees and arborescences*, Algebraic Methods in Graph Theory (Szeged 1978), vol. 25, Soc. J. Bolyai, 1978, pp. 402–412.

[42] ______, *Applications of submodular functions*, Surveys in Combinatorics (K. Walker, ed.), Cambridge Univ. Press, 1993.

[43] A. Frank and A. Gyárfás, *How to orient the edges of a graph*, Coll. Math. Soc. J. Bolyai **18** (1976), 353–364.

[44] P. Franklin, *A six color problem*, J. Math. Massachusetts **13** (1934), 363–369.

[45] A. Georgakopoulos, *Infinite Hamilton cycles in squares of locally finite graphs*, to appear in Adv. Math.

[46] A. Georgakopoulos, *Topological circles and Euler tours in locally finite graphs*, Electronic J. Comb. **16** (2009), #R40.

[47] A. Georgakopoulos and P. Sprüssel, *Geodesic topological cycles in locally finite graphs*, Preprint 2006.

[48] R. Göbel and J. Trlifaj, *Approximations and endomorphism algebras*, Walter de Gruyter, 2006.

[49] C. Godsil and G. Royle, *Algebraic graph theory*, Springer-Verlag, 2001.

[50] G. Hahn, F. Laviolette, and J. Širáň, *Edge-ends in countable graphs*, J. Combin. Theory (Series B) **70** (1997), 225–244.

[51] R. Halin, *Über die Maximalzahl fremder unendlicher Wege in Graphen*, Math. Nachr. **30** (1965), 63–85.

[52] D.A. Higgs, *Equicardinality of bases in B-matroids*, Can. Math. Bull. **12** (1969), 861–862.

[53] ______, *Infinite graphs and matroids*, Recent Prog. Comb., Proc. 3rd Waterloo Conf., 1969, pp. 245–253.

[54] ______, *Matroids and duality*, Colloq. Math. **20** (1969), 215–220.

[55] J. Keir and R.B. Richter, *Walks through every edge exactly twice II*, J. Graph Theory **21** (1996), 301–309.

[56] A.K. Kelmans, *The concept of a vertex in a matroid, the non-separating cycles and a new criterion for graph planarity*, Algebraic Methods in Graph Theory (Conf. Szeged 1978), vol. 1, Colloq. Math. Soc. János Bolyai 25, 1981, pp. 345–388.

[57] S. Lefschetz, *Applications of algebraic topology*, Springer-Verlag, 1975.

[58] S. Lins, R.B. Richter, and H. Shank, *The Gauss code problem off the plane*, Aequationes Math. **33** (1987), 81–95.

[59] S. MacLane, *A combinatorial condition for planar graphs*, Fund. Math. **28** (1937), 22–32.

[60] B. Mohar, *Generalization of MacLane's characterization of planar graphs*, Congr. Numerantium **56** (1987), 231–236.

[61] B. Mohar and C. Thomassen, *Graphs on surfaces*, Johns Hopkins, 2001.

[62] C.St.J.A. Nash-Williams, *On orientations, connectivity and odd pairings in finite graphs*, Canad. J. Math. **12** (1960), 555–567.

[63] ______, *Edge-disjoint spanning trees of finite graphs*, J. London Math. Soc. **36** (1961), 445–450.

[64] ______, *Another criterion for marriage in denumerable societies*, Ann. Disc. Math. **3** (1978), 165–179.

[65] J.G. Oxley, *Infinite matroids*, Proc. London Math. Soc. **37** (1978), 259–272.

[66] ______, *On a packing problem for infinite graphs and independence*, J. Combin. Theory (Series B) **26** (1979), 123–130.

[67] ______, *Infinite matroids*, Matroid applications (N. White, ed.), Encycl. Math. Appl., vol. 40, Cambridge University Press, 1992, pp. 73–90.

[68] ______, *Matroid theory*, Oxford University Press, New York, 1992.

[69] N. Polat, *A Mengerian theorem for infinite graphs with ideal points*, J. Combin. Theory (Series B) **51** (1991), 248–255.

[70] B. Richter and B. Rooney, *Relating different cycle spaces of the same infinite graph*, Preprint 2008.

[71] R.B. Richter and H. Shank, *Left-right paths off the plane*, Congr. Numerantium **45** (1984), 115–128.

[72] R.B. Richter and C. Thomassen, 3-*connected planar spaces uniquely embed in the sphere*, Trans. Am. Math. Soc. **354** (2002), 4585–4595.

[73] H.E. Robbins, *A theorem on graphs with an application to a problem on traffic control*, American Math. Monthly **46** (1939), 281–283.

[74] R. Rosenstiehl and R. Read, *On the principal edge tripartition of a graph*, Ann. Disc. Math. **3** (1978), 195–226.

[75] H. Shank, *The theory of left-right paths*, Combinatorial Mathematics III (A. Penfold Street and W.D. Wallis, eds.), Lecture notes in Math., vol. 452, Springer-Verlag, 1975, pp. 42–54.

[76] M. Stein, *Arboricity and tree-packing in locally finite graphs*, J. Combin. Theory (Series B) **96** (2006), 302–312.

[77] ———, *Forcing highly connected subgraphs*, J. Graph Theory **54** (2007), 331–349.

[78] C. Thomassen, *Planarity and duality of finite and infinite graphs*, J. Combin. Theory (Series B) **29** (1980), 244–271.

[79] ———, *Duality of infinite graphs*, J. Combin. Theory (Series B) **33** (1982), 137–160.

[80] ———, *Configurations in graphs of large minimum degree, connectivity or chromatic number*, Combinatorial mathematics (Proc. 3rd Int. Conf. 1985), vol. 555, Ann. N.Y. Acad. Sci., 1989, pp. 402–412.

[81] C. Thomassen and A. Vella, *Graph-like continua, augmenting arcs, and Menger's theorem*, Combinatorica **28** (2008), 595–623.

[82] W.T. Tutte, *On the problem of decomposing a graph into n connected factors*, J. London Math. Soc. **36** (1961), 221–230.

[83] ———, *How to draw a graph*, Proc. London Math. Soc. **13** (1963), 743–768.

[84] ———, *Menger's theorem for matroids*, J. Res. Natl. Bur. Stand., Sect. B **69** (1965), 49–53.

[85] ———, *Connectivity in matroids*, Can. J. Math. **18** (1966), 1301–1324.

[86] A. Vella, *A fundamentally topological perspective on graph theory*, Ph.D. thesis, University of Waterloo, 2005.

[87] A. Vella and R.B. Richter, *Cycle spaces of topological spaces*, Preprint.

[88] M. Las Vergnas, *Sur la dualité en théorie des matroides*, C. R. Acad. Sci. Paris, Sér. A **270** (1971), 67–85.

[89] J. Širáň and M. Škoviera, *Oriented relative embeddings of graphs*, Applicationes Math. **XIX** (1987), 589–597.

[90] ———, *Relative embeddings of graphs on closed surfaces*, Math. Nachr. **136** (1988), 275–284.

[91] K. Wagner, *Graphentheorie*, Bibliogr. Inst., Mannheim, 1970.

[92] H. Whitney, *Non-separable and planar graphs*, Trans. Am. Math. Soc. **34** (1932), 339–362.

[93] J. Wojciechowski, *A necessary condition for the existence of disjoint bases of a family of infinite matroids*, Congr. Numerantium **105** (1994), 97–115.

[94] ______, *A criterion for the existence of transversals of set systems*, J. London Math. Soc. **56** (1997), 491–503.

[95] ______, *Infinite matroidal version of hall's matching theorem*, J. London Math. Soc. **71** (2005), 563–578.

[96] X. Zha, *On minimum-genus embeddings*, Disc. Math. **149** (1996), 261–278.