

Infinite circuits in locally finite graphs

Dissertation
zur Erlangung des Doktorgrades
des Fachbereichs Mathematik
der Universität Hamburg

vorgelegt von
Henning Bruhn

Hamburg 2005

Contents

1	Introduction	1
1.1	Cycles in finite graphs	1
1.2	Infinite cycles	2
1.3	A topological definition of circles	3
1.4	The topological cycle space	6
1.5	The identification topology	7
1.6	Overview	10
2	Gallai's theorem and faithful cycle covers	11
2.1	Introduction	11
2.2	Proof of Gallai's theorem	13
2.3	Faithful cycle covers	15
3	MacLane's and Kelmans' planarity criteria	19
3.1	Introduction	19
3.2	Discussion of MacLane's criterion	20
3.3	Simple generating sets	21
3.4	The backward implication	27
3.5	The forward implication	29
3.6	Kelmans' planarity criterion	31
4	Duality	33
4.1	Introduction	33
4.2	Duality in infinite graphs	35
4.3	Proof of the duality theorem	38
4.4	Locally finite duals	42
4.5	Duality in terms of spanning trees	43
4.6	Alternative proof of Theorem 4.20	46
4.7	Colouring-flow duality and circuit covers	47
4.8	MacLane's criterion revisited	48

5	End degrees and infinite circuits	51
5.1	Introduction	51
5.2	Defining end degrees	52
5.3	A cut criterion for the end degree	56
5.4	Proof of Theorem 5.4	61
5.5	Properties of the end degree	65
5.6	Weakly even ends	70
6	Hamilton cycles and long cycles	73
6.1	Introduction	73
6.2	Hamilton circuits and covers	74
6.3	An infinite Herschel-type graph	75
6.4	Proof of Theorem 6.3	76
6.5	Infinite circuits generate the cycle space	83
7	Linear algebra with thin sums	87
7.1	Introduction	87
7.2	Bases	88
7.3	Closedness	91
8	Menger's theorem in infinite graphs with ends	93
8.1	Introduction	93
8.2	Definitions and statement of results	94
8.3	The reduction lemma	96
8.4	Countable separators	106
8.5	Disjoint closures	108
	Summary and conclusion	115
	Bibliography	117

Acknowledgement

I am very grateful to Reinhard Diestel for supervising this thesis, and for a lot of advice and support. I would also like to thank the Studienstiftung des deutschen Volkes for supporting me throughout all these years. Last but not least, I would like to express my gratitude towards Princeton University, Univerzita Karlova and also, even though it was only a very brief stay, towards GeorgiaTech for accommodating me and for their hospitality. In all these places I felt right at home.

Chapter 1

Introduction

Many of the properties of the cycle space of finite graphs break down in infinite graphs, when the cycle space is extended to infinite graphs in a naive way. This is mostly attributable to the lack of infinite cycles. Diestel and Kühn [33, 34] realised this and introduced infinite cycles that are defined in a topological manner. That their definition is extremely and almost surprisingly successful has since then been demonstrated in a series of papers, some of which make up the main part of this thesis. The objective of this thesis is to show that the properties of the cycle space in a finite graph carry over to locally finite graphs, if (and only if) infinite cycles are allowed.

In this chapter, we will recall the main properties of the cycle space of finite graphs, introduce the main definitions and review prior work. Finally, we will give an outlook over the thesis. In our notation, we follow Diestel [31].

1.1 Cycles in finite graphs

Let $G = (V, E)$ be a finite graph. A *cycle* C in G is a connected 2-regular subgraph, its edge set is called a *circuit*. Together with the symmetric difference as addition, the set of sums of circuits becomes a \mathbb{Z}_2 -vector space. The cycle space has a number of well-known basic properties. In particular, its elements are precisely those edge sets $Z \subseteq E$ for which

- Z meets every cut in an even number of edges;
- Z is the disjoint union of circuits;
- Z is the sum of fundamental circuits of any spanning tree of G ; and
- every vertex of (V, Z) has even degree.

(The fundamental circuits of a spanning tree T are the circuits C_e , $e \in E \setminus E(T)$, that arise from adding e to the edge set of $E(T)$.)

There are also some more advanced theorems that involve the cycle space:

- Tutte's generating theorem;
- MacLane's planarity criterion;
- Kelmans' planarity theorem;
- Duality/Whitney's planarity criterion; and
- Gallai's theorem.

In the course of this thesis we will come back to each of these theorems except the first one. Therefore, we shall only state Tutte's generating theorem here. Say that an induced circuit is *peripheral* if deleting its incident vertices does not separate the graph. Then:

Theorem 1.1 (Tutte [63]). *Every element of the cycle space of a finite 3-connected graph is a sum of peripheral circuits.*

1.2 Infinite cycles

The most immediate way to extend the cycle space to infinite graphs is what we will call the *finite-cycle space* $\mathcal{C}_{\text{fin}}(G)$: Its members are precisely the symmetric differences of finitely many (finite) circuits. Thus, the elements of $\mathcal{C}_{\text{fin}}(G)$ are always finite edge sets. While all of the basic properties that we have listed in the previous section remain true in an infinite graph (if we restrict ourselves to finite edge sets), almost all of the more advanced theorems either are considerably weakened or fail completely.

In the case of Tutte's generating theorem this was already noticed by Halin [41], who provided the counterexample in Figure 1.1: Consider the




Figure 1.1: The circuit C is not the finite sum of peripheral circuits

cartesian product of a *double ray* (an infinite 2-way path) with a pentagon. The peripheral circuits of this graph are exactly its 4-circuits. We see that

the deletion of all vertices incident with C separates the graph but C is not the sum of any peripheral circuits, so Tutte's theorem fails for this graph.

If we want to extend Theorem 1.1 then this example seems to imply that we need to allow certain infinite sums. In fact, summing up all the infinitely many peripheral circuits to the right, say, of C we obtain C . However, this alone is not enough. Figure 1.2 shows a graph in which an edge, denoted by e ,




Figure 1.2: There is no (finite) peripheral circuit containing the edge e

lies in no (finite) peripheral circuit at all. Therefore, no circuit containing e can be a sum of peripheral circuits, even allowing infinite sums. This example is from [14]; see also [15].

In a finite planar graph, the peripheral circuits are precisely the (edge sets of) face boundaries (assuming a 2-connected graph). Both of the face boundaries in Figure 1.2 that are incident with e are double rays, and therefore infinite. This indicates that, in order to make Theorem 1.1 valid in infinite graphs, we need infinite circuits.

Diestel and Kühn introduced a cycle space, which provides infinite circuits and allows (well-defined) infinite sums. In this space, the example in Figure 1.2 ceases to be a counterexample, and more generally Tutte's generating theorem (along with all the other advanced theorems) becomes true. Infinite circuits and the corresponding cycle space are defined in the next two sections.

1.3 A topological definition of circles

Diestel and Kühn [33, 34] define their infinite circuits (more precisely, any circuit—finite or infinite) based on a topological space whose point set consists of the graph together with its ends.

So, let us first recall what the ends of a graph are. A 1-way infinite path is called a *ray*, a 2-way infinite path is a *double ray*, and the subrays of a ray or double ray are its *tails*. Let $G = (V, E)$ be any graph. Two rays in G are *equivalent* if no finite set of vertices separates them; the corresponding

equivalence classes of rays are the *ends* of G . We denote the set of these ends by $\Omega = \Omega(G)$. Ends for graphs were introduced by Halin [40]; see also Diestel and Kühn [32] for their relationship to ends in topological spaces.

Let us define a topology on G together with its ends. We shall call this topology VTOP ; if G is locally finite, then this topology is usually called its *Freudenthal compactification*. We begin by viewing G itself (without ends) as the point set of a 1-complex. Then every edge is a copy of the real interval $[0, 1]$, and we give it the corresponding metric and topology. For every vertex v we take as a basis of open neighbourhoods the open stars of radius $1/n$ around v . (That is to say, for every integer $n \geq 1$ we declare as open the set of all points on edges at v that have distance less than $1/n$ from v , in the metric of that edge.)¹ In order to extend this topology to Ω , we take as a basis of open neighbourhoods of a given end $\omega \in \Omega$ the sets of the form

$$\hat{C}(S, \omega) := C(S, \omega) \cup \Omega(S, \omega) \cup \mathring{E}(S, \omega),$$

where $S \subseteq V$ is a finite set of vertices, $C(S, \omega)$ is the unique component of $G - S$ in which every ray in ω has a tail, $\Omega(S, \omega)$ is the set of all ends $\omega' \in \Omega$ whose rays have a tail in $C(S, \omega)$, and $\mathring{E}(S, \omega)$ is the set of all inner points of edges between S and $C(S, \omega)$.² Note that $\hat{C}(S, \omega)$ is an open neighbourhood of ω . Let $|G|$ denote the topological space on the point set $V \cup \Omega \cup \bigcup E$ thus defined. (When A is a set, we write $\bigcup A$ for the union of all its elements.) We shall freely view G and its subgraphs either as abstract graphs or as subspaces of $|G|$. Note that in $|G|$ every ray converges to the end of which it is an element. Furthermore, if G is locally finite then $|G|$ is compact and Hausdorff. If G is not locally finite then it might not be Hausdorff; for a characterisation of when $|G|$ is compact, see Diestel [27].

For any subset $X \subseteq |G|$, put $V(X) := X \cap V$, and let $E(X)$ be the set of edges e with $e \subseteq X$. We write \overline{X} for the closure of a set $X \subseteq |G|$ in $|G|$. For example, the set $\overline{C}(S, \omega)$ defined above is the closure in $|G|$ of the set $C(S, \omega)$. As a convenience and by slight abuse of notation, we also write \overline{Z} for the closure of the point set $\bigcup Z \subseteq |G|$ of an edge set Z . A subset of $|G|$ which is homeomorphic to the unit interval $[0, 1]$ is called an *arc*. The

¹If G is locally finite, this is the usual identification topology of the 1-complex. Vertices of infinite degree, however, have a countable neighbourhood basis in VTOP , which they do not have in the 1-complex.

²In the early papers on this topic, such as Diestel and Kühn [33, 34, 35], some more basic open sets were allowed: in the place of $\mathring{E}(S, \omega)$ we could take an arbitrary union of open half-edges from C towards S , one from every S - C edge. When G is locally finite, this yields the same topology. When G has vertices of infinite degree, our topology is slightly sparser but still yields the same topological cycle space; see the end of this section for more discussion.

images of 0 and of 1 are its *endpoints*. A set $C \subseteq |G|$ is a *circle* if it is a homeomorphic image of the unit circle. The next lemma will be of help when dealing with arcs and circles.

Lemma 1.2 (Diestel and Kühn [34]). *For every arc A and every circle C in $|G|$ the following is true:*

- (i) *The sets $A \cap G$ and $C \cap G$ are dense in A and C , respectively;*
- (ii) *every arc in $|G|$ whose endpoints are vertices or ends, and every circle C in $|G|$, includes every edge of G of which it contains an inner point; and*
- (iii) *if x is a vertex in $\overset{\circ}{A}$ (respectively in C), then A (respectively C) contains precisely two edges or half-edges of G at x .*

Therefore, $\overline{C \cap G} = C$ and $\overline{E(C)} = C$ for a circle $C \subseteq |G|$. We call the subgraph $C \cap G$ of G a *cycle* and the edge set $E(C)$ a *circuit*. Thus, a circle is uniquely determined by a cycle (resp. circuit) and vice versa. During the thesis we will switch back and forth between circles, cycles and circuits, depending on which is more appropriate for the concrete discussion.

Clearly, the definition of cycles includes traditional finite cycles but also allows infinite cycles. Such an infinite cycle is the disjoint union of double rays whose ends fit together nicely. A simple example of an infinite cycle consisting of two double rays is seen in Figure 1.3. There are, however, rather more complicated cycles. For instance, we will encounter in Chapter 6 a circle that is made up of countable infinitely many double rays and uncountably many ends.




Figure 1.3: A simple infinite cycle in the double ladder

We remark that, although the topology for $|G|$ considered in Diestel and Kühn [33, 34, 35] is slightly larger than ours (see the earlier footnote), the above lemma, as well as Theorems 1.4 and 1.5 below, is nevertheless applicable in our context. This is because the circles and arcs in $|G|$ coincide for these topologies: as one readily checks, the identity on $|G|$ between the two spaces is bicontinuous when restricted to a circle in either space.

1.4 The topological cycle space

In Section 1.2, we have seen that we need to allow certain infinite sums, if we want to extend Tutte's generating theorem. Let us make this more precise. For a graph G , call a family $(D_i)_{i \in I}$ of subsets of $E(G)$ *thin* if no vertex of G is incident with an edge in D_i for infinitely many i . (Thus in particular, no edge lies in more than finitely many D_i .) Let the *sum* $\sum_{i \in I} D_i$ of this family be the set of all edges that lie in D_i for an odd number of indices i , and let the *topological cycle space* $\mathcal{C}(G)$ of G be the set of all sums of (thin families of) circuits, finite or infinite. Symmetric difference as addition makes $\mathcal{C}(G)$ into a \mathbb{Z}_2 -vector space, which coincides with the usual cycle space of G when G is finite. We remark that $\mathcal{C}(G)$ is closed also under taking infinite thin sums (Diestel and Kühn [33, 34]), which is not obvious from the definitions.

In the topological cycle space the first three basic properties listed in Section 1.1 become true (with some adaptations), and also Tutte's theorem extends to locally finite graphs.

Theorem 1.3. [15] *Let G be a locally finite 3-connected graph. Then the peripheral circuits generate the topological cycle space.*

We state here two of the basic properties, which will serve as useful tools during the course of this thesis. The first is the cut criterion. A set $F \subseteq E(G)$ is a *cut* of a graph G if there is a partition (A, B) of $V(G)$ such that F is the set of all the edges of G with one vertex in A and the other in B . We shall also denote this set by $E_G(A, B)$, or by $E(A, B)$ if there is no confusion possible. A cut F is said to be *finitely covered* if there is a finite set of vertices covering F .

Theorem 1.4 (Diestel and Kühn [34]). *Let G be a graph. Then the following statements are equivalent for every $Z \subseteq E(G)$:*

- (i) $Z \in \mathcal{C}(G)$; and
- (ii) $|F \cap Z|$ is even for every finitely covered cut F of G .

In particular, if G is locally finite then (ii) becomes

- (ii') $|F \cap Z|$ is even for every finite cut F of G .

The second tool allows the decomposition of cycle space elements into circuits:

Theorem 1.5 (Diestel and Kühn [34]). *Every element of the topological cycle space of a graph is a disjoint union of circuits.*

From a combinatorial viewpoint the definition of the topological cycle space $\mathcal{C}(G)$ might at first seem to be rather complicated and perhaps even slightly intimidating, as it relies on topological concepts. And it is true that its elements can become quite complex; see Chapter 4 or 6 for an example. However, apart from being a very successful notion, an insight to which we hope to contribute in this thesis, there are two other reasons to embrace the topological cycle space. First, there is with Theorem 1.4 a simple combinatorial description after all. Second, $\mathcal{C}(G)$ is, in a certain sense, the smallest cycle space to which one can hope to extend the facts of Section 1.1; see Diestel [29].

1.5 The identification topology

While with the topological cycle space (almost) all of the theorems in Section 1.1 carry over to locally finite graphs, many fail in non-locally finite graphs. One of these that fail is, again, Tutte's generating theorem. In Figure 1.4, we see that every finite cycle containing the edge e is separating. Earlier, in the example in Figure 1.2, we were able to overcome a similar problem by introducing infinite cycles. Here, however, every infinite cycle containing e has a chord and is, therefore, not peripheral.




Figure 1.4: There is no peripheral circuit containing the edge e

Let us see what happens if we try to construct an infinite peripheral cycle. We start off in v_0 , go through e and then might run along $R := v_1 v_2 \dots$ into the end of R . Then we are stuck: In order to leave the end, which is necessary to ‘close up’ the cycle, we would have to traverse infinitely many neighbours of v_0 . But perhaps we do not need to leave the end after all, since it is already very ‘close’ to v_0 . Indeed, although R does not converge to v_0 , it nearly does: VTOP cannot separate its end from v_0 by two disjoint open sets. If we adjust our topology so that R does converge against v_0 , by identifying v_0 with the

end containing R , the ray v_0v_1R becomes a cycle, which then is peripheral, as desired.

Unfortunately, this is not enough to make Tutte's generating theorem become true. Consider the separating triangle $v_0v_1u_1v_0$ and suppose that (its edge set) is a thin sum of peripheral circuits. We may assume that the triangle $v_0u_1u_2v_0$ is one of the summands. Since each of the triangles $v_0u_nu_{n+1}v_0$ contains v_0 there is a minimal $n \geq 2$ such that $v_0u_nu_{n+1}v_0$ is not a summand. But then the edge v_0u_n lies in the sum, a contradiction.

It turns out that the problems that arise here are due to our too restrictive definition of a thin sum. In fact, if we allow infinite sums for which a vertex is incident with infinitely many summands, as long as no edge does, we easily obtain $v_0u_1u_2v_0$ as the sum of the (edges sets of the) peripheral cycles v_0v_1R , $v_0u_nu_{n+1}v_0$ and $u_nv_nv_{n+1}u_{n+1}u_n$ for $n \in \mathbb{N}$. Of course, there was a good reason for not allowing these kinds of thin sums when we defined the topological cycle space $\mathcal{C}(G)$; we will come back to this below.

Let us make all that more precise. We say that a vertex v *dominates* an end ω in a graph G if there is ray $R \in \omega$ and an infinite set of v - R paths that meet pairwise only in v . Assuming that

$$\text{every end of } G \text{ is dominated by at most one vertex,} \quad (1.1)$$

we now identify each vertex with all the ends it dominates, to obtain from $|G|$ the space \tilde{G} whose (quotient) topology we denote by ITOP. Note that, by (1.1), the vertices of G remain distinct in this identification. The identification space \tilde{G} is Hausdorff (unlike $|G|$, when G has a dominated end), and compact if G is 2-connected and satisfies condition (1.2) below; see Diestel and Kühn [35] and Diestel [27].

In the current setting we might experience a curious phenomenon: we might find that (the edge set of) a path is the sum of circuits, which should therefore be in our cycle space. Indeed, consider two vertices x and y that are linked by infinitely many independent paths. Then, since we now allow sums in which a vertex is incident with infinitely many of the summands, we can generate each of these paths P as a sum of cycles. To avoid this, we require the following:

$$\text{No two vertices of } G \text{ are joined by infinitely many independent paths.} \quad (1.2)$$

Note that (1.2) implies (1.1). As before, we define *circles* to be homeomorphic images of the unit circle in \tilde{G} , *circuits* to be their edge sets and *cycles* to be those subgraphs of G whose closure in \tilde{G} is a circle. The *topological cycle space* $\mathcal{C}(\tilde{G})$ of \tilde{G} is defined as the span of all sums of circuits such that no

edge appears in infinitely many of the summands. In the context of \tilde{G} we will always tacitly assume that all circles, cycles and circuits are defined with respect to \tilde{G} . It is worth stressing that for locally finite graphs not only the spaces $|G|$ and \tilde{G} coincide but also the cycle spaces $\mathcal{C}(G)$ and $\mathcal{C}(\tilde{G})$. The topology ITOP and $\mathcal{C}(\tilde{G})$ have been introduced by Diestel and Kühn in [35], where more details can be found.³

With this definition, several more results can be extended to graphs with infinite degrees. This usually necessitates the usage of the following two tools, that are analogous to Theorems 1.4 and 1.5.

Theorem 1.6 (Diestel and Kühn [35]). *Let G be a graph satisfying (1.2). Then $\mathcal{C}(\tilde{G})$ consists of precisely those sets of edges that meet every finite cut in an even number of edges.*

To decompose elements of $\mathcal{C}(\tilde{G})$ into circuits we need the following strengthening of (1.2):

$$\text{No two vertices are joined by infinitely many edge-disjoint paths.} \quad (1.3)$$

This is indeed stronger, see Diestel and Kühn [35].

Theorem 1.7 (Diestel and Kühn [35]). *Let G be a graph satisfying (1.3). Then every element of $\mathcal{C}(\tilde{G})$ is a disjoint union of circuits.*

Let us make a final useful observation. A connected locally finite graph is countable, which simplifies a lot of arguments. While there are uncountable connected graphs satisfying (1.3), there are none that are 2-connected. Indeed, as every uncountable connected graph has a vertex of uncountable degree, it is easy to show that an uncountable 2-connected graph contains two vertices joined by uncountably many independent paths. (See eg. Thomassen [60].) Thus:

Lemma 1.8. *A 2-connected graph satisfying (1.3) is countable.*

As the space $\mathcal{C}(\tilde{G})$ is the direct product of the cycle space of the blocks of G , we can usually assume G to be 2-connected.

Returning to the example in Figure 1.4, we remark that while it ceases to be a counterexample to Theorem 1.1 when $\mathcal{C}(\tilde{G})$ is used, it is still not known whether Tutte's generating theorem can be extended to graphs satisfying (1.3). On the other hand, no counterexamples are known and I suspect that there are none. In Chapters 2 and 4, we will see examples for theorems that fail in $|G|$ but become true in \tilde{G} .

³Here, we obtain ITOP from VTOP, which is slightly sparser than the topology TOP from which ITOP is derived in Diestel and Kühn [35]. However, in a similar way as detailed at the end of Section 1.3, we see that both topologies yield the same cycle space. In particular, Theorems 1.6 and 1.7 are still applicable.

1.6 Overview

We will start off in Chapter 2 by extending Gallai's theorem to graphs satisfying (1.3). In addition, we will show that if Seymour's faithful cycle cover conjecture, which generalises the famous cycle double cover conjecture, holds for finite graphs then it also does so for graphs satisfying (1.3). Both of these theorems will be seen to fail if only finite cycles are admitted.

In Chapter 3, we will extend MacLane's and Kelmans' planarity criteria to locally finite graphs. While there was already a partial extension of MacLane's criterion to infinite graphs by Thomassen, ours will be a full characterisation of planar graphs. Both criteria will be seen to be false if infinite cycles are disregarded.

In Chapter 4, we will deal with duality in graphs satisfying (1.3). We will build on work by Thomassen, who provided a concept of duality based on only finite cycles and cuts. Our notion of duality, which will draw on all cycles and cuts, finite or infinite, will fix all of the problems that arise from Thomassen's definition. Furthermore, we will be able to express duality in terms of trees, which was previously not possible.

In Chapter 5, we will pick up one of the problems of Diestel and Kühn, and try to characterise the elements of the cycle space in terms of vertex and 'end' degrees. For this, we will introduce a notion of an end degree that is based on counting edge-disjoint rays in an end. While a full characterisation is still out of reach, we will provide an important special case. Moreover, we will be able to describe cycles in terms of vertex and end degrees.

In Chapter 6, we will be concerned with Hamilton cycles. Tutte proved that a finite planar 4-connected graph is always hamiltonian. We will conjecture that this extends to locally finite graphs and provide a partial result to support this. Furthermore, we will prove that infinite circuits generate the cycle space.

In Chapter 7, we will try to understand the consequences of allowing infinite sums a bit more thoroughly. In particular, we will consider minimal generating sets, and prove their existence, assuming a countable ground set. If the ground set is uncountable, we will see that the existence cannot be guaranteed.

In the final chapter, we will turn to another problem that is closely related to the space $|G|$, namely the Erdős-Menger conjecture with ends. Building on work from Diestel, we will see how the topological space $|G|$ can be used to prove certain cases of the end version of the Erdős-Menger conjecture.

Chapter 2

Gallai's theorem and faithful cycle covers

2.1 Introduction

Let us start the main part of this thesis with two results that, while not trivial, are also not overly difficult. The main purpose of this chapter is to give the reader the opportunity to get used to the definitions and to introduce some of the methods.

By a result of Gallai (see Lovász [47]), every finite graph has a ‘cycle-cocycle’ partition of its edge set induced by a bipartition of its vertex set:

Theorem 2.1 (Gallai). *Every finite graph G admits a vertex partition into (possibly empty) sets V_1, V_2 such that both $E(G[V_1])$ and $E(G[V_2])$ are elements of the cycle space of G .*

As stated above, Gallai's theorem has no obvious extension to infinite graphs, if only finite circuits are admitted. Indeed, when G is infinite, the elements of its finite-cycle space $\mathcal{C}_{\text{fin}}(G)$ are still finite sets of edges, so a partition as in Theorem 2.1 does not exist, for instance, when G is an infinite disjoint union of triangles.

One way to deal with the problem is to look for an equivalent reformulation of Theorem 2.1 and extend that. For example:

Theorem 2.2. [18] *Every locally finite graph G admits a vertex partition into (possibly empty) sets V_1, V_2 such that in both $G[V_1]$ and $G[V_2]$ all vertex degrees are even.*

(The proof of Theorem 2.2 is an easy exercise in compactness. It is also an immediate corollary of Theorem 2.4 below.)

However, the requirement that all degrees of a subgraph H of a finite graph G should be even is only one equivalent reformulation among many of saying that $E(H)$ lies in the cycle space of G . Another is that H should be an edge-disjoint union of cycles (and isolated vertices). This would be just as meaningful for infinite H , and for locally finite H it implies the even-degree condition but not conversely. (Consider a double ray, which is 2-regular but not a union of cycles.) But with this latter reformulation, Theorem 2.1 no longer extends to infinite graphs, when only finite cycles are allowed:




Figure 2.1: A graph with no bipartition into edge-disjoint unions of cycles

Example 2.3. [18] *The graph G shown in Figure 2.1 has a unique vertex partition into two induced even-degree subgraphs. One of these is edgeless, the other a double ray.*

Proof. Consider any partition (V_1, V_2) of $V(G)$. Note that if two vertices x, y (such as u_1 and w_1) have a common neighbour z (such as v_1) not adjacent to any other vertex, then x and y must lie in the same partition class: otherwise, z would have degree 1 in its partition class. Thus if $u_1 \in V_1$, say, we deduce inductively that $w_1, w_3, w_5, \dots \in V_1$ and hence also $u_2, u_3, u_4, \dots \in V_1$. But u_2, u_3, u_4, \dots must not have degree 3 in $G[V_1]$, so $v_2, v_3, v_4, \dots \in V_2$. Finally, v_1 lies in V_1 because u_2 does, so inductively $w_2, w_4, \dots \in V_1$.

Thus, V_2 is the independent set $\{v_2, v_3, \dots\}$, while V_1 consists of the remaining vertices, which span a double ray. \square

Our aim in this chapter is to show that, despite Example 2.3, Theorem 2.2 is not the strongest possible extension of Theorem 2.1. Indeed, we can say more of the double ray $G[V_1]$ in Figure 2.1 than that its degrees are even: the double ray forms an infinite cycle in the topological cycle space $\mathcal{C}(G)$. (It does so, because its tails converge to the same end of G , which thus ‘closes it up’.) So for the space $\mathcal{C}(G)$, the graph of Figure 2.1 is no longer a counterexample to Theorem 2.1. And indeed, we have the following extension of Theorem 2.1 to infinite graphs, which implies Theorem 2.2 but is quite a bit stronger:

Theorem 2.4. [18] *For every locally finite graph G there is a partition of $V(G)$ into two (possibly empty) sets V_1, V_2 such that $E(G[V_i]) \in \mathcal{C}(G)$ for both $i = 1, 2$.*

The chapter, which is based on [18], is organised as follows. We shall prove Theorem 2.4 in the next section. In Section 2.3 we use similar techniques to extend the *cycle double cover conjecture* and Seymour's *faithful cycle cover conjecture* to locally finite graphs: if these conjectures are true for finite graphs, they also hold for locally finite graphs with our notion of an infinite topological cycle space. (The latter conjecture fails unless infinite cycles are admitted; for the former it is not known whether infinite cycles are really needed.)

2.2 Proof of Gallai's theorem

Before starting with its proof let us remark in this section [18] that Theorem 2.4 does not extend to arbitrary graphs with vertices of infinite degree. For example, consider the graph G obtained by joining a vertex v_0 to every vertex of a ray $R := v_1v_2v_3 \dots$. Suppose there is a partition as in Theorem 2.4, and assume that $v_0 \in V_1$. By the definition of thin sums, no element of $\mathcal{C}(G)$ can have infinitely many edges incident with v_0 . So there is a maximal $n \geq 0$ with $v_n \in V_1$. But then v_{n+1} has degree 1 in $G[V_2]$, a contradiction.

The problem here is that no element of the topological cycle space $\mathcal{C}(G)$ is allowed to have a vertex of infinite degree. However, in $\mathcal{C}(\tilde{G})$, in which only those infinite sums are forbidden where some *edge* lies in infinitely many of the summands (i.e. making no restrictions on vertices), our counterexample ceases to be one: for $V_1 := \{v_3, v_6, v_9, \dots\}$, the set $V_2 := V \setminus V_1$ induces an element of the cycle space. Of course, there was a good reason for forbidding these sums: summing up the triangles $v_0v_1v_2v_0, v_0v_2v_3v_0, v_0v_3v_4v_0, \dots$ yields the ray v_0v_1R , which should then also be a member of the cycle space. While this is not the case in the topological cycle space $\mathcal{C}(G)$, v_0v_1R is an element of $\mathcal{C}(\tilde{G})$, since v_0 dominates the end of R . Consequently, we will work within \tilde{G} and prove the following somewhat stronger version of Theorem 2.4:

Theorem 2.5. [18] *Let G be a graph satisfying (1.2). Then there is a partition of $V(G)$ into two (possibly empty) sets V_1, V_2 such that $E(G[V_i]) \in \mathcal{C}(\tilde{G})$ for both $i = 1, 2$.*

Our proof of Theorem 2.5 will be a compactness proof. Recall that while Theorem 2.2 has a straightforward compactness proof, the naive extension of

Theorem 2.1 to locally finite graphs does not (and is in fact false). The reason is, roughly speaking, that having all degrees even is a ‘local’ property of finite subsets $S \subseteq V(G)$ (one that S will satisfy in every large enough induced subgraph or in none), while inducing part of an element of the finite-cycle space $\mathcal{C}_{\text{fin}}(G)$, which is based on only finite cycles, is not: the sequence of finite cycles $C_n = P_n + e_n$, for example, where the $P_n = v_{-n}v_{-(n-1)} \dots v_{n-1}v_n$ are nested paths and e_n is the edge $v_{-n}v_n$, ‘tends’ for $n \rightarrow \infty$ to the double ray $D = \dots v_{-1}v_0v_1 \dots$ whose edge set does not lie in the finite-cycle space of $\bigcup_{n \in \mathbb{N}} C_n$. However, D is an infinite cycle in $\bigcup_{n \in \mathbb{N}} C_n$, and more generally it turns out that all such ‘limits’ of finite cycles in a graph G are elements of $\mathcal{C}(G)$ (though not necessarily single infinite cycles).

We shall cast our compactness proof in terms of König’s infinity lemma (see Diestel [31]), which we restate:

Lemma 2.6. *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 $v_1v_2 \dots$ with $v_n \in W_n$ for all n .*

Proof of Theorem 2.5. We may assume G to be 2-connected, because $\mathcal{C}(\tilde{G})$ is the direct product of the topological cycle spaces of its blocks. (Recall that vertices are allowed to lie in infinitely many summands as long as no edge does.) Then G is countable, by Lemma 1.8. Let v_1, v_2, \dots be an enumeration of $V(G)$. For $n \in \mathbb{N}$ set $S_n := \{v_1, \dots, v_n\}$, and define W_n as the set of all tuples (V_1, V_2) such that

- (i) (V_1, V_2) is a partition of S_n into two (possibly empty) sets; and
- (ii) for $i = 1, 2$, there is a $Z \in \mathcal{C}(\tilde{G})$ such that $Z \cap E(G[S_n]) = E(G[V_i])$.

Each set W_n is clearly finite. It is non-empty by Theorem 2.1 applied to $G[S_n]$.

Let us define a graph H on $\bigcup_{n=1}^{\infty} W_n$. For $n \geq 2$, let $(V_1, V_2) \in W_n$ be adjacent to $(V'_1, V'_2) \in W_{n-1}$ if and only if, for both $i = 1, 2$, $V'_i \subseteq V_i$. Observe that for $n \geq 2$ every vertex in W_n has a neighbour in W_{n-1} .

By the infinity lemma (2.6), there is a ray $v_1v_2 \dots$ in H with $(V_1^n, V_2^n) := v_n \in W_n$ for all n . Clearly, $V_1 := \bigcup_{n=1}^{\infty} V_1^n$ and $V_2 := \bigcup_{n=1}^{\infty} V_2^n$ form a partition of $V(G)$.

We shall use Theorem 1.6 to show that $E(G[V_1]) \in \mathcal{C}(\tilde{G})$, and in a similar way that $E(G[V_2]) \in \mathcal{C}(\tilde{G})$. Write $Z_n := E(G[V_1^n])$ for each n . Consider a finite cut F of G . Choose n large enough that $F \subseteq E(G[S_n])$. By (ii), there

is a $Z \in \mathcal{C}(\tilde{G})$ with $Z \cap E(G[S_n]) = Z_n$. Then

$$\begin{aligned} F \cap E(G[V_1]) &= F \cap E(G[S_n] \cap G[V_1]) = F \cap Z_n \\ &= F \cap Z \cap E(G[S_n]) = F \cap Z. \end{aligned}$$

Since $Z \in \mathcal{C}(\tilde{G})$, the last intersection is even. Hence $E(G[V_1]) \in \mathcal{C}(\tilde{G})$ by Theorem 1.6, as desired. \square

2.3 Faithful cycle covers

Another problem concerning cycles is the well-known cycle double cover conjecture, which states that every bridgeless finite graph has a cycle double cover. (A *cycle double cover* of a graph G is a family of cycles such that each edge of G lies on exactly two of those cycles.) We will discuss cycle covers, and a more general kind of cover, namely faithful covers, in this section [18]. Using the same techniques as in the proof of Theorem 2.5 one can show that if the cycle double cover conjecture is true for finite graphs then it also holds for locally finite graphs, possibly with infinite cycles. However, it is unknown whether there exists an example where infinite cycles are really needed; I suspect there is none.

The situation is different for the following related conjecture of Seymour, which extends with infinite cycles but fails with finite cycles only. For a graph G and a map $p : E(G) \rightarrow \mathbb{N}$ ($\ni 0$) a *faithful cycle cover* of (G, p) is a family of cycles such that every edge $e \in G$ lies on exactly $p(e)$ of those cycles. Such a map p is *admissible* if $p(F) = \sum_{f \in F} p(f)$ is even and $p(e) \leq p(F)/2$ for every finite cut F and every edge $e \in F$. We call p *even* if all its values $p(e)$ are even numbers. If (G, p) is to have a faithful cycle cover, then obviously p has to be admissible, and we shall see below that for some G it has to be even. Since the constant map with value 2 is admissible for bridgeless graphs, the following *faithful cycle cover conjecture* extends the cycle double cover conjecture:

Conjecture 2.7 (Seymour [56]). *Let G be a finite graph, and p an even admissible map. Then (G, p) has a faithful cycle cover.*

Unlike the cycle double cover conjecture, we know that Conjecture 2.7 fails for locally finite graphs unless we allow infinite cycles. Here is a simple example. Let G be the double (= two-way infinite) ladder, and let p assign 0 to every rung and 2 to all the other edges. By our current definition of admissibility (which requires $p(e) \leq p(F)/2$ only for finite cuts F), the function p is admissible. But G contains no finite cycle that avoids all rungs,

so (G, p) has no faithful cover consisting of finite cycles. (It does, however, have a faithful cover consisting of two copies of the infinite cycle spanned by the edges for which $p = 2$.)

The above example is no longer a counterexample to the infinite analogue of Conjecture 2.7 if we require of an admissible map p that it satisfies $p(e) \leq p(F)/2$ also for infinite cuts F (and edges $e \in F$): if e is any edge with $p(e) = 2$ and R is a maximal ray in the subgraph of $G - e$ spanned by all its remaining edges with $p = 2$, then e and the edges with $p = 0$ incident with R form an infinite cut F such that $p(e) = p(F)$. Thus, p is no longer admissible, and we no longer have a contradiction.

Our next example, however, shows that strengthening the definition of ‘admissible’ as above is not enough to make Conjecture 2.7 true for locally finite graphs—if only finite cycles are admitted. Consider the ladder G shown in Figure 2.2 and the admissible map $p : E(G) \rightarrow \mathbb{N}$ defined by $p(e_i) = p(e'_i) = 2i$ and $p(f_i) = 2$ for all i . (Since $p(e) > 0$ for all e , we trivially have $p(e) \leq p(F)/2$ also for infinite cuts F .) Suppose there is a faithful cycle cover which contains a finite cycle D . Obviously, D contains exactly two rungs f_m, f_n , with $m < n$, say. Let \mathcal{C} be the subfamily of the cover consisting of those cycles which pass through the edge e_n . Each but at most one (which might go through f_n) of the cycles in \mathcal{C} must use the edge e_{n-1} . Thus, at least $|\mathcal{C}| - 1 = 2n - 1$ cycles of the cover meet the edge e_{n-1} , contradicting $p(e_{n-1}) = 2n - 2$. Therefore, the only faithful cycle cover that (G, p) can have (and which is easily seen to exist) must be one consisting of infinite cycles.




Figure 2.2: The unique faithful cycle cover consists of infinite cycles only

As soon as we allow infinite cycles, however, Conjecture 2.7 does extend to locally finite graphs, and more generally to graphs that satisfy (1.3):

Theorem 2.8. [18] *Let G be a graph satisfying (1.3), and let $p : E(G) \rightarrow \mathbb{N}$ be an even admissible map. If Conjecture 2.7 is true then (\tilde{G}, p) has a faithful cycle cover.*

First, we need a lemma, that will prove to be useful in Chapter 4 too.

Lemma 2.9.[18] *Let G be a 2-connected multigraph satisfying (1.3), and let U be a finite set of vertices in G . Then we can contract edges of G , deleting loops but keeping any multiple edges that arise, so that no two vertices from U are identified, the multigraph H obtained has only finitely many edges and vertices, and every cut of H is also a cut of G .*

Proof. First, note that the set \mathcal{K} of components of $G - U$ is finite. Indeed, as G is 2-connected, every component C of $G - U$ has distinct neighbours u, v in U . If \mathcal{K} is infinite, then infinitely many $C \in \mathcal{K}$ are joined to the same two vertices u, v (because U is finite), so these are linked by infinitely many independent paths. This contradicts (1.3).

Next, consider a component $C \in \mathcal{K}$. For every two vertices $u, v \in U$ that both send infinitely many edges to C there is a finite cut $F_{u,v} \subseteq E(C)$ separating $N(u) \cap V(C)$ from $N(v) \cap V(C)$ in C , because of (1.3). Let F_C be the union of all such cuts $F_{u,v}$. Note that F_C is finite, as there are only finitely many pairs u, v . Then the set \mathcal{K}_C of components of $C - F_C$ is also finite, and so is $\mathcal{K}' := \bigcup_{C \in \mathcal{K}} \mathcal{K}_C$. Each $D \in \mathcal{K}'$ sends only finitely many edges to $G - U - D$, and at most one vertex in U sends infinitely many edges to D . If such a vertex exists, we denote it by u_D .

In G , contract every $D \in \mathcal{K}'$ to a vertex v_D , keeping parallel edges but deleting loops. If two vertices of the resulting multigraph are joined by infinitely many edges, then these are u_D and v_D for some $D \in \mathcal{K}'$. In a second step, we now contract all these edges $u_D v_D$, again keeping parallel edges. We obtain a finite multigraph H in which no two vertices from U are identified. (Note in particular that the edge set of H is finite, despite the parallel edges that arose in the contraction.) Since we did not delete any edges except loops, every cut of H is also a cut of G . \square

Proof of Theorem 2.8. Consider a block B of G . Every cut of B is a cut of G , so the restriction of p to B is an even admissible map on B . As $\mathcal{C}(\tilde{G})$ is the direct product of the topological cycle spaces of the blocks of G , we may therefore assume G to be 2-connected. (Note that p assigns zero to bridges, so we need not cover these.) By Lemma 1.8, G is countable.

Consider an enumeration v_1, v_2, \dots of $V(G)$, and set $G_n := G[\{v_1, \dots, v_n\}]$. Define W_n as the set of all families \mathcal{E} of sets $E \subseteq E(G_n)$ such that

- (i) every edge $e \in E(G_n)$ lies in exactly $p(e)$ members of \mathcal{E} ; and
- (ii) for every $E \in \mathcal{E}$ there is a finite cycle $C \subseteq G$ with $E(C \cap G_n) = E$.

Let us show that the sets W_n are not empty. Apply Lemma 2.9 with $U = \{v_1, \dots, v_n\}$, and denote the multigraph H obtained by G'_n . Since every cut of G'_n is also one of G , the map p induces an admissible even map p'_n on G'_n .

By subdividing edges we obtain from G'_n a simple graph G''_n with admissible even map p''_n (induced by p'_n). Then by assumption there is a faithful cycle cover of (G''_n, p''_n) . Every cycle in that cover can be extended to a finite cycle in G . The family of these cycles then satisfies (i) and (ii), thus proving $W_n \neq \emptyset$.

The rest of the proof is analogous to that of Theorem 2.5: applying the infinity lemma to an auxiliary graph H , we obtain a family of elements of $\mathcal{C}(\tilde{G})$ such that every edge e lies on exactly $p(e)$ members of this family. By Theorem 1.7, we can modify this into a faithful cover consisting of single cycles. Therefore, if the faithful cycle cover conjecture holds for finite graphs, it is also true for graphs satisfying (1.3). \square

Conjecture 2.7 requires p to be even, and indeed if p is allowed to assume odd values the conjecture becomes false: take the Petersen graph, and give p the value 2 on a perfect matching and 1 on all other edges.

Take any subgraph of an infinite graph G , and contract some—possibly infinitely many—of its edges; the resulting graph will be called a *minor* of G . Then the following result, whose finite version is a theorem of Alspach, Goddyn and Zhang [7], can be proved like Theorem 2.8.

Theorem 2.10.[18] *Let G be a graph that satisfies (1.3) and does not contain the Petersen graph as a minor, and let $p : E(G) \rightarrow \mathbb{N}$ be any admissible map. Then (G, p) has a faithful cycle cover.*

Chapter 3

MacLane's and Kelmans' planarity criteria

3.1 Introduction

A set (or family) \mathcal{E} of edge sets $E \subseteq E(G)$ of a graph G is called *simple*, if every edge of G lies in at most two elements of \mathcal{E} . MacLane's planarity criterion states:

Theorem 3.1 (MacLane [48]). *A finite graph G is planar if and only if its cycle space has a simple generating set.*

Wagner [64] raised the question if MacLane's result could be extended so that it characterises planar graphs which are infinite. Rather than modifying the planarity criterion, Thomassen [59] describes all infinite graphs that satisfy MacLane's condition. For this, recall 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 .

Theorem 3.2 (Thomassen [59]). *Let G be an infinite 2-connected graph. Then G has a VAP-free embedding in the plane if and only if $\mathcal{C}_{\text{fin}}(G)$ has a simple generating set consisting of finite circuits.*

(Recall that $\mathcal{C}_{\text{fin}}(G)$ consists of all finite symmetric differences of finite circuits; see Section 1.2.) Bonnington and Richter [12] also provide a generalisation of MacLane's theorem using the *even cycle space* $\mathcal{Z}(G)$, defined as the set of all subgraphs of G with all vertex degrees even. With this space they investigate which graphs have an embedding with k VAPs.

Our main result in this chapter, which is based on [22], is a verbatim generalisation of MacLane's theorem to locally finite graphs, using the topological cycle space $\mathcal{C}(G)$:

Theorem 3.3. [22] *Let G be a countable locally finite graph. Then, G is planar if and only if $\mathcal{C}(G)$ has a simple generating set.*

We discuss Theorem 3.3 in the next section, and prove it in the course of Sections 3.4 and 3.5. In Section 3.3, we investigate some properties of simple generating sets. In Section 3.6, we extend Kelmans' planarity criterion to locally finite graphs.

3.2 Discussion of MacLane's criterion

In this section we shall discuss Theorem 3.3 and have a closer look at generating sets [22]. First let us make the notion of a generating set more precise. A *generating set* of the topological cycle space will be a set $\mathcal{F} \subseteq \mathcal{C}(G)$ such that every element of $\mathcal{C}(G)$ can be written as a thin sum of elements of \mathcal{F} . Thus, in contrast to a generating set in the vector space sense we allow (thin) infinite sums. There are two reasons for this. First, thin sums are integral to the topological cycle space of an infinite graph, so it seems unnatural to forbid them. Second, MacLane's criterion is false if we insist that every $Z \in \mathcal{C}(G)$ is a finite sum of elements of a simple subset of $\mathcal{C}(G)$, as we shall see in Proposition 3.6.

To show that, in a certain sense, Theorem 3.3 is as strong as possible, we need the following theorem, which is of interest on its own. It will be proved in the next section. Recall that a circuit C is called *peripheral* if it is induced and if deleting all incident vertices does not separate the graph.

Theorem 3.4. [22] *Let G be a 3-connected graph, and let \mathcal{F} be a simple generating set of $\mathcal{C}(G)$ consisting of circuits. Then every element of \mathcal{F} is a peripheral circuit.*

First note that because of Theorem 1.5, if the topological cycle space has a simple generating set then it also has a simple generating set consisting of circuits.

Theorem 3.3 is formulated for locally finite graphs, and indeed it is false for arbitrary infinite graphs. Indeed, consider the 3-connected graph G in Figure 3.1, which is not locally finite; the example is from [15]. By Theorem 3.4 and the remark following it, we may assume that a simple generating set \mathcal{F} of $\mathcal{C}(G)$ consists of peripheral circuits (finite or infinite). In particular, no circuit which contains the edge e is in \mathcal{F} . But then such a circuit cannot be generated by any sum of circuits of \mathcal{F} . Thus, there is no simple generating set of $\mathcal{C}(G)$, but G is clearly planar. We remark that, nevertheless, it is possible to extend MacLane's criterion to certain non-locally finite graphs, namely those satisfying (1.3); we will discuss this in Section 4.8.




Figure 3.1: A planar graph whose topological cycle space has no simple generating set

Infinite circuits might seem, although topologically natural, combinatorially unwieldy. They are, however, inevitable in a certain sense: there is not always in a planar graph a simple generating set comprised of only finite circuits. In Figure 1.2, we presented an example for a graph in which the finite peripheral circuits fail to generate the cycle space. As that graph is 3-connected, we know from Theorem 3.4 that there is then also no simple generating set consisting only of finite circuits. Another indication that we cannot do without infinite circuits is provided by Georgakopoulos [38], who constructs highly connected plane graphs in which the edge set of every face boundary is an infinite circuit. Later, in Lemma 3.21, we will see that (the edge sets of) the face boundaries are precisely the peripheral circuits. Thus, there are 3-connected planar graphs with no finite peripheral circuits at all.

3.3 Simple generating sets

We prove Theorem 3.4 in this section [22]. As a tool, we introduce the notion of a 2-basis. For this, let $\mathcal{B} \subseteq \mathcal{C}(G)$ be a simple generating set of the topological cycle space of G . We call \mathcal{B} a 2-basis of $\mathcal{C}(G)$ if for every element $Z \in \mathcal{C}(G)$ there is a unique (thin) subset of \mathcal{B} , henceforth denoted by \mathcal{B}_Z , with $Z = \sum_{B \in \mathcal{B}_Z} B$. Observe that in a finite graph the 2-bases are exactly the simple bases of $\mathcal{C}(G)$, and thus conform with the traditional definition of a 2-basis in a finite graph.

Since we have left linear algebra with our definition of a 2-basis (allowing thin infinite sums), it is not clear if the properties usually expected of a basis are still retained. One of these, which we shall need later on, is that a generating set always contains a basis. The next lemma asserts that at least for simple sets this is true. We will investigate bases in more generality than is necessary here in Chapter 7.

Lemma 3.5.[22] *Let G be a 2-connected graph, and let \mathcal{F} be a simple generating set of $\mathcal{C}(G)$. If \mathcal{F} is not a 2-basis, then for any $Z \in \mathcal{F}$ the set $\mathcal{F} \setminus \{Z\}$*

is a 2-basis of $\mathcal{C}(G)$.

Proof. Observe first that it suffices to check the uniqueness required in the definition of a 2-basis for the empty set: a simple generating subset \mathcal{B} of $\mathcal{C}(G)$ is a 2-basis if and only if for every $\mathcal{B}' \subseteq \mathcal{B}$ with $\sum_{B \in \mathcal{B}'} B = \emptyset$ it follows that $\mathcal{B}' = \emptyset$.

Let us assume there is a non-empty set $\mathcal{D} \subsetneq \mathcal{F}$ with $\sum_{B \in \mathcal{D}} B = \emptyset$. Since G is 2-connected every edge of G appears in a finite circuit, and thus in at least one element of \mathcal{F} . But as \mathcal{F} is simple and $\sum_{B \in \mathcal{D}} B = \emptyset$ no edge of G can lie in an element of \mathcal{D} and at the same time in an element of $\mathcal{F} \setminus \mathcal{D}$.

So, $E_1 := \bigcup \mathcal{D}$ and $E_2 := \bigcup (\mathcal{F} \setminus \mathcal{D})$ define a partition of $E(G)$ (note that both sets are non-empty). Because G is 2-connected there is, by Menger's theorem, for any two edges a finite circuit through both of them. Therefore, there is a circuit D which shares an edge e_1 with E_1 and another edge e_2 with E_2 . Let $\mathcal{D}' \subseteq \mathcal{F}$ be such that $D = \sum_{B \in \mathcal{D}'} B$. Then $D' := \sum_{B \in \mathcal{D} \cap \mathcal{D}'} B \subseteq D$, since for any edge $e \in D' \setminus D$ both $\mathcal{D}' \setminus \mathcal{D}$ and $\mathcal{D} \cap \mathcal{D}'$ have an element which contains e ; thus $e \in E_1 \cap E_2$, which is impossible. Therefore, D' is a subset of the circuit D , and thus either $D' = \emptyset$ or $D' = D$. Since $e_1 \in D'$ the former case is impossible; the latter, however, is so too, as $D' \subseteq E_1$ cannot contain $e_2 \in E_2$, a contradiction.

We thus have shown:

$$\sum_{B \in \mathcal{D}} B = \emptyset \text{ for } \mathcal{D} \subseteq \mathcal{F} \text{ implies } \mathcal{D} = \emptyset \text{ or } \mathcal{D} = \mathcal{F}.$$

So, if \mathcal{F} is not a 2-basis, then none of its subsets but itself generates the empty set. In particular, \mathcal{F} is thin. For any $Z \in \mathcal{F}$,

$$Z = \sum_{B \in \mathcal{F} \setminus \{Z\}} B,$$

thus, the thin simple set $\mathcal{F} \setminus \{Z\}$ certainly generates the topological cycle space. It also is a 2-basis, as none of its non-empty subsets generates the empty set. \square

With our definition of a generating set, which allows infinite sums, we shall show that MacLane's criterion holds for locally finite graphs. Since, in a vector space context, one usually allows only finite sums for a generating set, there is one obvious question: Does Theorem 3.3 remain true if we consider simple generating sets in the vector space sense? The answer is a strikingly clear no:

Proposition 3.6. [22] *There is no locally finite 2-connected infinite graph in which the topological cycle space has a simple generating set in the vector space sense (i.e. allowing only finite sums).*

Proof. Suppose there is such a graph G so that $\mathcal{C}(G)$ has a simple set $\mathcal{A} \subseteq \mathcal{C}(G)$ which generates every $Z \in \mathcal{C}(G)$ through a finite sum. We determine the cardinality of $\mathcal{C}(G)$ in two ways.

First, since \mathcal{A} is simple, every of the countably many edges of G lies in at most two elements of \mathcal{A} . Therefore, \mathcal{A} is a countable set, and thus, $\mathcal{C}(G)$ also.

Second, there is, by Lemma 3.5, a 2-basis $\mathcal{B} \subseteq \mathcal{A}$. As $\mathcal{C}(G)$ is an infinite set (since G is infinite and 2-connected), so is \mathcal{B} . Hence, there are distinct $B_1, B_2, \dots \in \mathcal{B}$. Also, as G is locally finite and \mathcal{B} simple, all subsets of \mathcal{B} are thin. Therefore, all the sums

$$\sum_{i \in I} B_i \text{ for } I \subseteq \mathbb{N}$$

are distinct elements of $\mathcal{C}(G)$. Since the power set of \mathbb{N} has uncountable cardinality, it follows that $\mathcal{C}(G)$ is uncountable, a contradiction. \square

The rest of this section is devoted to the proof of Theorem 3.4, which we restate:

Theorem 3.4. [22] *Let G be a 3-connected graph, and let \mathcal{F} be a simple generating set of $\mathcal{C}(G)$ consisting of circuits. Then every element of \mathcal{F} is a peripheral circuit.*

A basic tool when dealing with finite circuits are bridges, see for instance Bondy and Murty [11]. As our circuits may well be infinite, we need an adaption of the notion of a bridge, which we introduce together with a number of related results before proving the theorem.

Definition 3.7. [15] *Let $C \subseteq |G|$ be a circle in a graph G . We call the closure B of a topological component of $|G| \setminus C$ a bridge of C . The points in $B \cap C$ are called the attachments of B in C .*

There is a close relationship between bridges and peripheral circuits. Indeed, in a 3-connected graph a circuit D is peripheral if and only if the circle \overline{D} has a single bridge [15].

For the subgraph $H := C \cap G$, the following can be shown: a set $B \subseteq |G|$ is a bridge of C if and only if it is induced by a chord of H or if there is a component K of $G - H$ such that B is the closure of K plus the edges between K and H together with the incident vertices. Thus, our definition coincides with the traditional definition of a bridge in a finite graph.

Lemma 3.8. [15] *Let $C \subseteq |G|$ be a circle in a graph G , and let B be a bridge of C . Let x be an attachment of B . Then:*

- (i) x is a vertex or an end;
- (ii) if x is an end then every neighbourhood of x contains attachments of B that are vertices;
- (iii) every edge of which B contains an inner point lies entirely in B ; and
- (iv) either B is induced by a chord of C or the subgraph $(B \cap G) - V(C)$ is non-empty and connected.

We define a *residual arc* of the bridge B in the circle C to be the closure of a topological component of $C \setminus B$. Note that if B has at least two attachments every residual arc is indeed an arc (if not then the circle C itself is a residual arc, and it is the only one).

Lemma 3.9. [15] *Let G be a 2-connected graph, and let $C \subseteq |G|$ be a circle with a bridge B . Then:*

- (i) *the endpoints of a residual arc L of B in C are attachments of B ; and*
- (ii) *for a point $x \in C \setminus B$ there is exactly one residual arc L of B in C containing x .*

Let us introduce some more terminology that are standard in a finite setting; see Bondy and Murty [11]. We say a bridge B of C *avoids* another bridge B' of C if there is a residual arc of B that contains all attachments of B' . Otherwise, they *overlap*. Note that overlapping is a symmetric relation. Two bridges B and B' of C are called *skew* if C contains four (distinct) points v, v', w, w' in that cyclic order such that v, w are attachments of B and v', w' attachments of B' . Clearly, if two bridges B and B' are skew, they overlap. On the other hand, in a 3-connected graph, overlapping bridges are either skew or *3-equivalent*, i.e. they both have only three attachments which are the same:

Lemma 3.10. [22] *Let G be a 3-connected graph. Let $C \subseteq |G|$ be a circle, and let B and B' be two overlapping bridges of C . Then B and B' are either skew or 3-equivalent.*

Proof. First, if either B or B' is induced by a chord, it is easy to see, that they are skew because they overlap. Thus, by Lemma 3.8 (iv), we may assume that each of the bridges has three attachments. Next, assume that $B \cap C = B' \cap C$. If $|B \cap C| = 3$ then B and B' are 3-equivalent, otherwise they are clearly skew.

So, suppose there is an attachment u of B with $u \notin B'$. The attachment u is contained in a residual arc L of B' . Its endpoints u', v' are attachments

of B' . Since B and B' are overlapping, not all attachments of B may lie in L . Thus, there is an attachment $v \in C \setminus L$ of B . Then, the sequence u, u', v, v' shows that B and B' are skew. \square

For a set $X \subseteq |G|$, an X -path is a path that starts in X , ends in X and is otherwise disjoint from X .

Lemma 3.11. [22] *Let B and B' be two skew bridges of a circle $C \subseteq |G|$ in a graph G . Then there are two disjoint C -paths $P = u \dots v$ and $P' = u' \dots v'$ such that u, u', v, v' appear in that order on C .*

Proof. Since B and B' are skew there are points x, x', y, y' appearing in that cyclic order on C such that x, y are attachments of B and x', y' are attachments of B' . If x is a vertex put $u := x$. If not, then there is a whole arc $A \subseteq C$ around x disjoint from any of the other points. In A we find, by Lemma 3.8 (ii), an attachment u of B that is a vertex. Doing the same for x', y and y' , if necessary, we end up with vertices u, u', v, v' appearing in that cyclic order on C such that $u, v \in B$ and $u', v' \in B'$. As $(B \cap G) - V(C)$ is connected, by Lemma 3.8 (iv), we find an u - v path P through B , and analogously an u' - v' path P' through B' . Since bridges meet only in attachments, P and P' are disjoint. \square

We need that in a 3-connected graph, for any circle, there are always two overlapping bridges (if there is more than one bridge at all). For this, we define for a circle C in the graph G the *overlap graph* of C in G as the graph on the bridges of C such that two bridges are adjacent if and only if they overlap. The next lemma ensures that there are always overlapping bridges.

Lemma 3.12. [15] *For every circle C in a 3-connected graph G the overlap graph of C in G is connected.*

The next simple lemma will be used repeatedly in the proof of Theorem 3.4. Denote by $M \triangle N$ the symmetric difference of two sets M and N .

Lemma 3.13. [22] *Let G be a 3-connected graph, and let \mathcal{B} be a 2-basis of $\mathcal{C}(G)$ consisting of circuits. Let C and D be circuits in G such that $\overline{C} \cap \overline{D}$ is an arc. Suppose that $\mathcal{B}_C \cap \mathcal{B}_D \neq \emptyset$. Then, either $\mathcal{B}_C \subseteq \mathcal{B}_D$ or $\mathcal{B}_D \subseteq \mathcal{B}_C$.*

Proof. Put $K := \sum_{B \in \mathcal{B}_C \cap \mathcal{B}_D} B$ and consider an edge $e \notin C \cup D$. Then both \mathcal{B}_C and \mathcal{B}_D contain either both or none of the at most two circuits $B \in \mathcal{B}$ with $e \in B$. Thus, both or none of them is in $\mathcal{B}_C \cap \mathcal{B}_D$, and hence $e \notin K$. Therefore, K is an element of the topological cycle space contained in $C \cup D$.

These are precisely \emptyset , C , D and $C + D$ (since $\overline{C} \cap \overline{D}$ is an arc). Note that $K \neq \emptyset$ as $\mathcal{B}_C \cap \mathcal{B}_D \neq \emptyset$. Also, $K \neq C + D$, since otherwise

$$\mathcal{B}_C \cap \mathcal{B}_D = \mathcal{B}_K = \mathcal{B}_{C+D} = \mathcal{B}_C \Delta \mathcal{B}_D,$$

which is impossible. Consequently, we obtain either $K = C$ and thus, $\mathcal{B}_C \subseteq \mathcal{B}_D$, or $K = D$ and $\mathcal{B}_D \subseteq \mathcal{B}_C$. \square

Proof of Theorem 3.4. Note that it suffices to prove the theorem for a 2-basis \mathcal{B} . Indeed, if \mathcal{F} is not a 2-basis, consider two distinct elements Z_1 and Z_2 of \mathcal{F} . By Lemma 3.5, both $\mathcal{F} \setminus \{Z_1\}$ and $\mathcal{F} \setminus \{Z_2\}$ are a 2-basis of $\mathcal{C}(G)$, and, if Theorem 3.4 holds for these, it clearly also holds for \mathcal{F} .

Consider a non-peripheral circuit C . Then, the circle \overline{C} has more than one bridge [15]. Two of these, B and B' say, are, by Lemma 3.12, overlapping. By Lemma 3.10, they are either skew or 3-equivalent. We show that $C \notin \mathcal{B}$ for each of the two cases.

(i) Suppose that B and B' are skew. By Lemma 3.11, there are two disjoint \overline{C} -paths $P = u \dots v$ and $P' = u' \dots v'$ such that u, u', v, v' appear in this order on \overline{C} . Denote by $L_{uu'}, L_{u'v}, L_{vv'}, L_{v'u}$ the closures of the topological components of $\overline{C} \setminus \{u, u', v, v'\}$ such that x, y are the endpoints of L_{xy} . Define the circuits

$$\begin{aligned} C_1 &:= E(L_{uu'} \cup L_{u'v} \cup P), & C_2 &:= E(L_{vv'} \cup L_{v'u} \cup P), \\ D_1 &:= E(L_{u'v} \cup L_{vv'} \cup P') & \text{and} & \quad D_2 := E(L_{v'u} \cup L_{uu'} \cup P'). \end{aligned}$$

Observe that $C_1 + C_2 = C = D_1 + D_2$, and additionally, that $\overline{C}_i \cap \overline{D}_j$ is an arc for any $i, j \in \{1, 2\}$.

Suppose $C \in \mathcal{B}$. Since

$$\mathcal{B}_{C_1} \Delta \mathcal{B}_{C_2} = \mathcal{B}_{C_1+C_2} = \mathcal{B}_C = \{C\},$$

not both of \mathcal{B}_{C_1} and \mathcal{B}_{C_2} may contain C . As the same holds for D_1 and D_2 we may assume that

$$C \notin \mathcal{B}_{C_1} \text{ and } C \notin \mathcal{B}_{D_1}. \quad (3.1)$$

Consider an edge $e \in C_1 \cap D_1 \subseteq C$. Both of \mathcal{B}_{C_1} and \mathcal{B}_{D_1} must contain a circuit which contains e . By (3.1), this cannot be C . Therefore, and since \mathcal{B} is simple, \mathcal{B}_{C_1} and \mathcal{B}_{D_1} contain the same circuit K with $e \in K$. Consequently, $\mathcal{B}_{C_1} \cap \mathcal{B}_{D_1} \neq \emptyset$, and applying Lemma 3.13 we may assume that

$$\mathcal{B}_{C_1} \subseteq \mathcal{B}_{D_1}. \quad (3.2)$$

Now, consider an edge $e' \subseteq L_{uu'}$, hence $e' \in C_1 \cap D_2$. There is a circuit $K' \in \mathcal{B}_{C_1}$ with $e' \in K' \neq C$. By (3.2), $K' \in \mathcal{B}_{D_1}$, but since e' lies in $L_{uu'}$ we

have $e' \notin D_1$. Thus, \mathcal{B}_{D_1} also contains the other circuit in \mathcal{B} that contains e' , which is C , a contradiction to (3.1). Therefore, $C \notin \mathcal{B}$.

(ii) Suppose that B and B' are 3-equivalent. Let v_1, v_2, v_3 be their attachments, which then are vertices (by Lemma 3.8 (ii)). Then there is a vertex $x \in V(B \setminus C)$ and three $x\text{-}\overline{C}$ paths $P_i = x \dots v_i \subseteq B$, $i = 1, 2, 3$ whose interiors are pairwise disjoint. Let $Q_i = y \dots v_i$ be analogous paths in B' . The closures of the topological component of $\overline{C} \setminus \{v_1, v_2, v_3\}$ are three arcs; denote by $L_{i,i+1}$ the one that has v_i and v_{i+1} as endpoints (where indices are taken mod 3). For $i = 1, 2, 3$, define the circuits

$$C_i := E(L_{i,i+1} \cup P_i \cup P_{i+1}) \text{ and } D_i := E(L_{i,i+1} \cup Q_i \cup Q_{i+1}).$$

Note that $C_1 + C_2 + C_3 = C = D_1 + D_2 + D_3$.

Now suppose $C \in \mathcal{B}$. As

$$\mathcal{B}_{C_1} \Delta \mathcal{B}_{C_2} \Delta \mathcal{B}_{C_3} = \mathcal{B}_{C_1+C_2+C_3} = \mathcal{B}_C = \{C\},$$

either C lies in all of the \mathcal{B}_{C_i} or in only one of them, in \mathcal{B}_{C_3} , say. In both cases, we have $C \notin \mathcal{B}_{C_1+C_2}$. We obtain the same result for the D_i : either C lies in all of the \mathcal{B}_{D_i} or in only one of them. In any case, we can define D as either D_1 or $D_2 + D_3$ such that $C \notin \mathcal{B}_D$. Put $D' := C + D$, and note that $\mathcal{B}_{D'} = \mathcal{B}_D \cup \{C\}$.

Then, since $C_1 + C_2$ shares an edge in C with D , and neither $\mathcal{B}_{C_1+C_2}$ nor \mathcal{B}_D contains C , we have $\mathcal{B}_{C_1+C_2} \cap \mathcal{B}_D \neq \emptyset$. Applying Lemma 3.13, we obtain that one of the two sets $\mathcal{B}_{C_1+C_2}, \mathcal{B}_D$ is contained in the other.

First assume that $\mathcal{B}_{C_1+C_2} \subseteq \mathcal{B}_D$, and consider an edge $e \in C$ that lies in both $C_1 + C_2$ and D' . Such an edge exists since $D' = D_1$ or $D' = D_2 + D_3$. Since $C \notin \mathcal{B}_{C_1+C_2}$, e lies in a circuit $K \neq C$ in $\mathcal{B}_{C_1+C_2}$, and thus also $K \in \mathcal{B}_D$. On the other hand, $e \in C \in \mathcal{B}_{D'}$ contradicts $e \in D'$.

So, we may assume that $\mathcal{B}_D \subseteq \mathcal{B}_{C_1+C_2}$. Because $\mathcal{B}_{D'} = \mathcal{B}_D \cup \{C\}$ we even have $\mathcal{B}_D \subseteq \mathcal{B}_{C_1+C_2} \cap \mathcal{B}_{D'}$. Thus, by Lemma 3.13, either $\mathcal{B}_{C_1+C_2} \subseteq \mathcal{B}_{D'}$ or $\mathcal{B}_{D'} \subseteq \mathcal{B}_{C_1+C_2}$. The latter is impossible as $C \notin \mathcal{B}_{C_1+C_2}$. Therefore, we obtain

$$\mathcal{B}_D \subseteq \mathcal{B}_{C_1+C_2} \subseteq \mathcal{B}_{D'} = \mathcal{B}_D \cup \{C\}.$$

Now, from $C \notin \mathcal{B}_{C_1+C_2}$ follows that $\mathcal{B}_{C_1+C_2} = \mathcal{B}_D$, contradicting $C_1+C_2 \neq D$. Thus, $C \notin \mathcal{B}$. \square

3.4 The backward implication

In this section, we show the backward implication of Theorem 3.3 namely that if the topological cycle space has a simple generating set then G is planar [22]. But first, let us remark that it is sufficient to show Theorem 3.3 for

2-connected graphs. Indeed, the Kuratowski planarity criterion for countable graphs below asserts that a countable graph is planar if and only if its blocks are planar.

Theorem 3.14 (Dirac and Schuster [37]). *Let G be a countable graph. Then, G is planar if and only if G contains neither a subdivision of K_5 nor a subdivision of $K_{3,3}$.*

The backward direction will follow from the next lemma.

Lemma 3.15. [22] *Let G be a 2-connected graph such that $\mathcal{C}(G)$ has a 2-basis, and let $H \subseteq G$ be a finite 2-connected subgraph. Then $\mathcal{C}(H)$ has a 2-basis.*

Proof. Let \mathcal{B} be the 2-basis of $\mathcal{C}(G)$. Since H is finite, there are $Z \in \mathcal{C}(H)$ with a non-empty generating set $\mathcal{B}_Z \subseteq \mathcal{B}$ which is \subseteq -minimal among all B_Z with $Z \in \mathcal{C}(H)$. Let us denote these by Z_1, \dots, Z_k .

Consider a $D \in \mathcal{C}(H)$ with $\mathcal{B}_D \cap \mathcal{B}_{Z_i} \neq \emptyset$ for some i . We claim that $\mathcal{B}_{Z_i} \subseteq \mathcal{B}_D$. First, note that

$$C := \sum_{B \in \mathcal{B}_D \cap \mathcal{B}_{Z_i}} B \subseteq E(H).$$

Indeed, consider an edge $e \notin E(H)$. Since $Z_i, D \subseteq E(H)$, and since \mathcal{B} is simple, e either lies on exactly two or on none of the elements of \mathcal{B}_{Z_i} , and the same holds for \mathcal{B}_D . Furthermore, if e lies on two elements of \mathcal{B}_{Z_i} and on two of \mathcal{B}_D , these must be the same. So, $e \notin C$.

Therefore, $C \subseteq E(H)$, and thus $C \in \mathcal{C}(H)$. As $\mathcal{B}_C \subseteq \mathcal{B}_{Z_i}$ we obtain, by the minimality of \mathcal{B}_{Z_i} , that $C = Z_i$. Consequently, $\mathcal{B}_{Z_i} = \mathcal{B}_C \subseteq \mathcal{B}_D$, as claimed.

This result also implies $\mathcal{B}_{Z_i} \cap \mathcal{B}_{Z_j} = \emptyset$ for all $1 \leq i < j \leq k$. Thus, every edge of H appears in at most two of the Z_i . Furthermore, we claim that $\{Z_1, \dots, Z_k\}$ is a generating set for $\mathcal{C}(H)$. Then, $\{Z_1, \dots, Z_k\}$ contains a 2-basis of $\mathcal{C}(H)$, and we are done.

So consider a $D \in \mathcal{C}(H)$, and let I denote the set of those indices i with $\mathcal{B}_{Z_i} \cap \mathcal{B}_D \neq \emptyset$. We may assume $I = \{1, \dots, k'\}$ for a $k' \leq k$. Then, by $\mathcal{B}_{Z_i} \subseteq \mathcal{B}_D$ and $\mathcal{B}_{Z_i} \cap \mathcal{B}_{Z_j} = \emptyset$ for $i, j \in I$, it follows that \mathcal{B}_D is the disjoint union of the sets $\mathcal{B}_{Z_1}, \mathcal{B}_{Z_2}, \dots, \mathcal{B}_{Z_{k'}}$ and

$$\mathcal{B}' := \mathcal{B}_D \setminus \bigcup_{i=1}^{k'} \mathcal{B}_{Z_i}.$$

Consequently,

$$\sum_{B \in \mathcal{B}'} B = \sum_{B \in \mathcal{B}_D} B + \sum_{B \in \mathcal{B}_{Z_1}} B + \dots + \sum_{B \in \mathcal{B}_{Z_{k'}}} B = D + Z_1 + \dots + Z_{k'} \subseteq E(H)$$

as all the summands lie in H . Now, if $\mathcal{B}' \neq \emptyset$ then there is a $Z \in \mathcal{C}(H)$ with a non-empty and minimal $\mathcal{B}_Z \subseteq \mathcal{B}'$ which then must be one of the Z_i , a contradiction. Thus, \mathcal{B}' is empty and we have $D = \sum_{i=1}^{k'} Z_i$. \square

For the backward implication of Theorem 3.3, we use the well-known fact that the cycle space of every subdivision of K_5 or of $K_{3,3}$ fails to have a 2-basis (see, for instance, Diestel [31]).

Lemma 3.16. [22] *Let G be a locally finite 2-connected graph such that $\mathcal{C}(G)$ has a simple generating set. Then G is planar.*

Proof. Suppose not. Then G contains, by Theorem 3.14, a subdivision H of K_5 or of $K_{3,3}$ as subgraph. By Lemma 3.5, $\mathcal{C}(G)$ has a 2-basis. Then, by Lemma 3.15, $\mathcal{C}(H)$ also has a 2-basis, which is impossible. \square

3.5 The forward implication

For the forward implication of Theorem 3.3, we need to show that the topological cycle space of a planar graph has a simple generating set. By proceeding mainly as for the finite case, we will establish the forward implication in this section [22]. We first embed our graph G in the sphere and then show that the set of the face boundaries' edge sets is a simple generating set. So, our first priority is to ensure that every face is indeed bounded by a circle of $|G|$. As for the backward direction we may assume that G is 2-connected.

This, however, is certainly not the case when a VAP of the embedded graph coincides with a vertex or an inner point of an edge. To avoid this problem we consider topological embeddings of the space $|G|$ in the sphere (rather than graph embeddings of G), which, in our context, is no restriction:

Theorem 3.17 (Richter and Thomassen [55]). *Let G be a locally finite 2-connected planar graph. Then $|G|$ embeds in the sphere.*

We call a topological space *2-connected* if it is connected and remains so after the deletion of any point. Thus, any embedding of the (standard) compactification $|G|$ of a 2-connected graph G in the sphere clearly is 2-connected. Note that also, any such embedding is compact if G is locally finite and connected. A *face* of a compact subset K of the sphere is a component of the complement of K . A face boundary $\partial f \subseteq K$ of a face f is

simply the boundary of f . If K is the image of $|G|$ under an embedding, then it can be shown in a similar way as for finite plane graphs (see for instance Diestel [31]) that if an inner point of an edge lies in a face boundary then the whole edge lies in it.

Theorem 3.18 (Richter and Thomassen [55]). *Every face of a compact 2-connected locally connected subset of the sphere is bounded by a simple closed curve.*

Another result of Richter and Thomassen [55] states that $|G|$ is locally connected if G is locally finite and connected¹. As a simple closed curve by definition is homeomorphic to the unit circle, we obtain:

Corollary 3.19. [22] *Let G be a locally finite 2-connected graph with an embedding $\varphi : |G| \rightarrow S^2$. Then the face boundaries of $\varphi(|G|)$ are circles of $|G|$.*

Showing the forward implication, we now complete the proof of Theorem 3.3:

Lemma 3.20. [22] *Let G be a locally finite 2-connected planar graph. Then, $\mathcal{C}(G)$ has a simple generating set.*

Proof. By Theorem 3.17, $|G|$ has an embedding $\varphi : |G| \rightarrow S^2$ in the sphere. Put $\Gamma := \varphi(|G|)$. We show that the set \mathcal{F} which we define to consist of the edge sets of the face boundaries of Γ , is a simple generating set of $\mathcal{C}(G)$. Certainly, \mathcal{F} is simple, and, by Corollary 3.19, a subset of $\mathcal{C}(G)$. So, we only have to prove that every element of the topological cycle space is the sum of certain elements of \mathcal{F} .

Fix a face f^* of Γ . First, consider a circuit C in G . Then for the circle \overline{C} , $\varphi(\overline{C})$ is homeomorphic to the unit circle and, thus, bounds two faces (by the Jordan-curve theorem). Let f_C be the face not containing f^* . As G is 2-connected, every edge e lies on a finite circuit, and therefore on the boundaries of exactly two faces of Γ , which we denote by f_e and f'_e . Hence, the set

$$\mathcal{B}^C := \{E(\partial f) : f \subseteq f_C \text{ is a face of } \Gamma\}$$

is thin. Moreover, as we have $f_e, f'_e \subseteq f_C$ or $f_e, f'_e \not\subseteq f_C$ if and only if $e \notin C$, it follows that

$$\sum_{B \in \mathcal{B}^C} B = C. \tag{3.3}$$

¹They show this to be true for all *pointed* compactifications of G , which are those obtained from the standard compactification by identifying some ends.

Now, consider an arbitrary element Z of the topological cycle space. By definition, there is a thin family \mathcal{D} of circuits with $Z = \sum_{C \in \mathcal{D}} C$. If none of the elements of \mathcal{F} appears in \mathcal{B}^C for infinitely many $C \in \mathcal{D}$, then the family \mathcal{B} , which we define to be the (disjoint) union of all \mathcal{B}^C with $C \in \mathcal{D}$, is thin (since every edge lies on exactly two face boundaries). Then, $Z = \sum_{B \in \mathcal{B}} B$, and we are done. Therefore, if $F(\Gamma)$ is the set of faces of Γ , it suffices to show that the set

$$F := \{f \in F(\Gamma) : f \subseteq f_C \text{ for infinitely many } C \in \mathcal{D}\}$$

is empty.

So suppose $F \neq \emptyset$. By definition of f_C , we have $f^* \not\subseteq f_C$ for all $C \in \mathcal{D}$, and thus also $F \neq F(\Gamma)$. Hence, there is an edge e such that one of its adjacent faces, say f_e , lies in F and the other, f'_e , in $F(\Gamma) \setminus F$. Then, $E(\partial f_e)$ appears in infinitely many \mathcal{B}^C while e lies on only finitely many $C \in \mathcal{D}$. Thus, also $E(\partial f'_e)$ lies in infinitely many \mathcal{B}^C , which implies $f'_e \in F$, a contradiction. \square

3.6 Kelmans' planarity criterion

For finite 3-connected graphs there is another well-known planarity criterion, namely Kelmans' criterion² (see Kelmans [44]). It follows from MacLane's criterion together with Tutte's generating theorem, both of which are known to be true for locally finite graphs. We extend Kelmans' criterion to locally finite graphs in this section [22].

Corollary 3.19 implies that the face boundaries of a locally finite 2-connected planar graph are circles. When G is 3-connected then, as for finite graphs (see Diestel [31]), the Jordan curve theorem implies that these circles are precisely the closures in $|G|$ of the peripheral circuits of G :

Lemma 3.21. [22] *Let G be a locally finite 3-connected graph with an embedding $\varphi : |G| \rightarrow S^2$ in the sphere. Then, the face boundaries are precisely the closures in $\varphi(|G|)$ of the peripheral circuits of G . \square*

Now, we easily obtain a verbatim generalisation of Kelmans' criterion for locally finite graphs.

Theorem 3.22. [22] *Let G be a locally finite 3-connected graph. If G is planar then every edge appears in exactly two peripheral circuits. Conversely, if every edge appears in at most two peripheral circuits then G is planar.*

²This criterion is sometimes also known as Tutte's planarity criterion, as it easily follows from Tutte's generating theorem. However, this seems to have gone unnoticed until Kelmans published a direct proof of the criterion neither using MacLane's nor Tutte's theorem.

Proof. If G is planar then there is, by Theorem 3.17, also an embedding of $|G|$, in which, by Lemma 3.21, the closure of every peripheral circuit is a face boundary. Since G is 2-connected every edge lies in exactly two face boundaries, hence in exactly two peripheral circuits of G .

For the backward implication let \mathcal{F} be the set of all peripheral circuits of G , which then is simple. Thus, \mathcal{F} is, by Theorem 1.3, a simple generating set, and hence G planar, by Theorem 3.3. \square




Figure 3.2: Infinite circuits are necessary for Kelmans' criterion

As MacLane's planarity criterion, Kelmans', too, fails when infinite circuits are disallowed. Indeed, there are 3-connected non-planar graphs in which every edge lies on at most two finite peripheral circuits. The graph G shown in Figure 3.2 is such an example [22]. It consists of a $K_{3,3}$ (bold) to which three disjoint infinite 3-ladders are added. First observe that any finite peripheral circuit that contains edges of $G - \{u, w\}$ cannot contain any edge incident with either one of u, w , as otherwise it also contains (the edges of) a finite $\{x, y, z\} - \{x, y, z\}$ path in $G - \{u, w\}$, and thus is separating. Therefore, every finite peripheral circuit of G has either none or all of its edges incident with $\{u, w\}$; in the latter case, it is a circuit of $G[u, w, x, y, z]$.

Now, assume that there is an edge of G that appears in three finite peripheral circuits. All of these circuits then lie either in $G - \{u, w\}$ or in $G[u, w, x, y, z]$, where they are also peripheral. Now, it is easy to check that none of the edges of the finite graph $G[u, w, x, y, z]$ lies on three peripheral circuits, and by Theorem 3.22 this is also impossible for any edge of the planar 3-connected graph $G - \{u, w\}$. This shows that Kelmans' criterion fails if only finite circuits are admitted.

Chapter 4

Duality

4.1 Introduction

Let G and G^* be multigraphs, and let there exist a bijection $E(G) \rightarrow E(G^*)$ that maps the circuits of G precisely to the minimal non-empty cuts (or *bonds*) of G^* . Then G^* is called a *dual* of G . Whitney used duality in order to characterise planar graphs:

Theorem 4.1 (Whitney [65]). *A finite multigraph G has a dual if and only if it is planar.*

In this chapter we shall extend the duality of finite graphs in what appears to be a complete and best-possible way. However, all our results build on previous work of Thomassen [59, 60], who extended finite duality to infinite multigraphs as far as it will go without considering infinite circuits. Lacking infinite circuits, Thomassen was confronted with a disparity between finite circuits and possibly infinite cuts. His approach was to disregard the latter, and thus, in his terms, G^* qualifies as a dual of G as soon as the edge bijection between the two multigraphs maps all the finite circuits of G to the finite bonds of G^* .

This weaker notion of duality already permits a very satisfactory extension of Whitney's theorem to 2-connected multigraphs. Our stronger notion strengthens this (in one direction), and in addition re-establishes two aspects of finite duality that cannot be achieved for infinite multigraphs when infinite cuts (and circuits) are disregarded: the uniqueness of duals for 3-connected multigraphs, and symmetry, the fact that a multigraph is always a dual of its duals.

Concerning symmetry, note that taking duals may force us out of the class of locally finite multigraphs: while the half-grid, shown in unbroken




Figure 4.1: A locally finite graph with a non-locally finite dual

lines in Figure 4.1, is locally finite, its (geometric) dual, in dotted lines, is not. But graphs with vertices of infinite degree do not, in general, have duals. Indeed, Thomassen showed that any infinite multigraph with a dual (even in his weaker sense) must satisfy condition (1.3), which we restate here for convenience:

$$\textit{No two vertices are joined by infinitely many edge-disjoint paths.} \quad (1.3)$$

To achieve symmetry, we thus need that the class of multigraphs satisfying (1.3) is closed under taking duals. While this is not the case for Thomassen's notion, it will be true for ours.

Before we go on with the discussion let us briefly make a technical remark. In this chapter we will, in contrast to the rest of the thesis, primarily consider multigraphs. Multigraphs are the natural context when considering duality as the dual of a (simple) graph may well have parallel edges or loops.

In Chapter 1, we have defined the topological cycle space for graphs rather than for multigraphs. However, this was for the sake of simplicity, and all the concepts apply to multigraphs too. In particular, our main tools, Theorems 1.4 and 1.5, resp. Theorems 1.6 and 1.7, still hold.¹

This chapter, which is based on [16], is organised as follows. We continue the above discussion in more precise terms in Section 4.2, which leads up to the statement of our basic duality theorem, Theorem 4.6. We prove this theorem in Section 4.3. In Section 4.4, we characterise the graphs that have locally finite duals. In Section 4.5, we treat duality in terms of spanning trees. In Section 4.7, we apply our results to colouring-flow duality. Finally, in the last section, we use duality to extend MacLane's planarity criterion from Chapter 3 to graphs satisfying (1.3).

¹Observe that if we subdivide every edge in a multigraph G then the topological cycle space of the subdivision induces a space in G that is identical to $\mathcal{C}(G)$, which we obtain by directly applying the definitions to G . The theorems are valid in the subdivision, which is simple, and can thus be pulled back to G .

4.2 Duality in infinite graphs

As discussed in the previous section, Thomassen pursued an approach to duality in infinite graphs that is based solely on finite circuits and cuts. While being very successful in some respects, such as conditions for the existence of duals and extensions of Whitney’s theorem, this approach leads to unavoidable problems in others, such as symmetry and the uniqueness of duals. Our aim in this section is to discuss these problems, to indicate why considering infinite circuits and working in ITOP is both, in essence, necessary and sufficient to cure them, and to state the main result of this chapter [16].

Consider multigraphs G and G^* , possibly infinite, and assume that there is a bijection $*$: $E(G) \rightarrow E(G^*)$. Given a set $F \subseteq E(G)$, put $F^* := \{e^* \mid e \in F\}$, and vice versa. (That is, given a subset of $E(G^*)$ denoted by F^* , we write F for the subset $\{e \mid e^* \in F^*\}$ of $E(G)$.) A cut is called a *bond* if it is minimal among the non-empty cuts. Call G^* a *finitary dual* of G if, for every finite set $F \subseteq E(G)$, the set F is a circuit in G if and only if F^* is a bond in G^* . Following Thomassen [59, 60], we require 2-connected multigraphs to be loopless, and 3-connected multigraphs to have no parallel edges.² That is, a 3-connected multigraph is in fact simple.

Expressed in these terms, Thomassen obtained the following extension of Whitney’s theorem:

Theorem 4.2 (Thomassen [60]). *A 2-connected³ multigraph G has a finitary dual if and only if G is planar and satisfies (1.3).*

Going back to Figure 4.1, we see that the dotted graph G^* is a finitary dual of the half-grid G . However, splitting the vertex v into two vertices u and w , and making each of these adjacent to infinitely many neighbours of v in such a way that every neighbour of v is adjacent to exactly one of u and w , we obtain another finitary dual H of G . This violates the intended uniqueness of duals for 3-connected graphs such as G . (Recall that duals of 3-connected finite graphs are unique.)

Moreover, admitting H as a dual violates symmetry, since G is not a finitary dual of H . In fact, H has no finitary dual at all, and it might not even be planar, depending on how we join u and w to the neighbours of v .

²This is motivated by matroid theory. Disallowing loops is also necessary for uniqueness: a multigraph with loops never has exactly one dual (unless it is itself a loop).

³We expect that Thomassen’s theorem extends to multigraphs of smaller connectivity. There is no mention of this in [60], however, and we note that the canonical proof for the forward implication fails: when G^* is a finitary dual of G , then the duality map $*$ need not map the blocks of G to blocks of G^* , so it is not obvious that the blocks of G have finitary duals too. Compare Lemma 4.13 below.

Thomassen realised these problems, as is witnessed by the following two theorems.

Theorem 4.3 (Thomassen [59]). *Let G be a 2-connected multigraph having a finitary dual. Then G has a finitary dual G^* satisfying (1.3), and every such finitary dual G^* has G as its finitary dual.*

We say that a multigraph H is a *finitary predual* of G if G is a finitary dual of H .

Theorem 4.4 (Thomassen [60]). *If a multigraph G has a 3-connected finitary predual then this is its only predual, up to isomorphism.*

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 $|G| = \tilde{G}$ achieves this.

To restore symmetry, we have to allow vertices of infinite degree. (Note that G^* has one, and we want G to be its dual.) We thus have to decide now whether to work in $|G|$ or in \tilde{G} . That is to say, should we take the circles that define our infinite circuits in the topology ITOP specifically designed for graphs satisfying (1.3), or in the simpler VTOP?

To answer this question, let us consider the graph G shown in unbroken lines in Figure 4.2, and let G^* be a hypothetical dual of G (by the definition we are seeking). We want G to be a dual of G^* , and in particular a finitary dual. Thus, G^* will be a finitary predual of G . Now G is certainly a finitary dual of the dotted graph H shown in Figure 4.2, so H is also a finitary predual of G . Since H is 3-connected, Theorem 4.4 implies that $H = G^*$.




Figure 4.2: The self-dual graph G

Since the dotted edges at v form a bond of G^* , we thus have to make the edge set F of the double ray D a circuit of G . Now in $|G|$ the set F is not

a circuit, because G has two ends and D has a tail in each. In \tilde{G} , however, both ends of G are identified with the vertex v , so the double ray D and the vertex v together do form a circle, making F into a circuit as desired.

We therefore propose the following stronger notion of duality for infinite graphs, in which the duality condition is required of all sets of edges, finite or infinite, and circuits are defined as in \tilde{G} under ITOP.

Definition 4.5. [16] *Let G be a multigraph satisfying (1.3). Let G^* be another multigraph, with a bijection $*$: $E(G) \rightarrow E(G^*)$. 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^* .*

Note that every dual in this sense is also a finitary dual, but not conversely.

Figure 4.3 shows that infinite circuits can get pretty wild, even in locally finite graphs. The following theorem, which is our main result, can thus deviate more from the corresponding finite situation than it might at first appear.




Figure 4.3: The bold edges form an infinite circuit in G , the broken edges indicate the corresponding cut in G^*

Theorem 4.6. [16] *Let G be a countable⁴ multigraph satisfying (1.3).*

⁴By Lemma 1.8 below, the countability assumption is redundant for 2-connected multigraphs, and therefore inessential. If we agree to call a multigraph ‘planar’ as soon as it has neither a K_5 nor a $K_{3,3}$ minor, then Theorem 4.6 becomes true also for uncountable multigraphs.

- (i) G has a dual if and only if G is planar.
- (ii) If G^* is a dual of G , then G^* satisfies (1.3), G is a dual of G^* , and this is witnessed by the inverse bijection of $*$.
- (iii) If G^* is a dual of G and $F \subseteq E(G)$, then $F \in \mathcal{C}(\tilde{G})$ if and only if F^* is a cut in G^* .

We shall prove Theorem 4.6 in the next section.

Since all finitary duals of a 3-connected graph are again 3-connected (see Thomassen [60]), Theorems 4.4 and 4.6 (ii) together imply at once that the dual of a 3-connected graph is unique:

Corollary 4.7. [16] *A 3-connected graph has at most one dual, up to isomorphism.*

4.3 Proof of the duality theorem

Recall that, by Theorem 4.2, any multigraph G with a finitary dual G^* satisfies (1.3). Our first aim is to show that if G^* is a dual of G , then G^* too satisfies (1.3). We need two lemmas. The first is from Chapter 2, which we restate here.

Lemma 2.9. [18] *Let G be a 2-connected multigraph satisfying (1.3), and let U be a finite set of vertices in G . Then we can contract edges of G , deleting loops but keeping any multiple edges that arise, so that no two vertices from U are identified, the multigraph H obtained has only finitely many edges and vertices, and every cut of H is also a cut of G .*

Lemma 4.8. [16] *Let G be a 2-connected multigraph satisfying (1.3), and let C be an infinite circuit in \tilde{G} . Let X be a finite set of edges meeting C in exactly one edge e . Then there is a finite circuit in G that meets X precisely in e .*

Proof. Apply Lemma 2.9 to G , taking as U the set of endvertices of the edges in X . Consider the finite multigraph H returned by the lemma. Applying Theorem 1.6 twice, we deduce from $C \in \mathcal{C}(\tilde{G})$ and the separation property of H that $C \cap E(H) \in \mathcal{C}(H)$. Let $C' \subseteq C \cap E(H)$ be a circuit containing e . As C meets X only in e , so does C' . Since no two vertices from U were identified when G was contracted to H , the branch sets of the contraction induce no edge from X in G . We can therefore expand C' to a finite circuit in G that still meets X only in e . \square

We need the following strong version of Menger's theorem for countable graphs.

Theorem 4.9 (Aharoni [3]). *For any countable graph G and two sets A, B of vertices in G there exist a set \mathcal{P} of disjoint A - B paths and an A - B separator X in G such that X consists of a choice of one vertex from each of the paths in \mathcal{P} .*

We can now show that, unlike with finitary duals, the class of multigraphs satisfying (1.3) is closed under taking duals.

Lemma 4.10.[16] *Any dual G^* of a 2-connected multigraph G satisfies (1.3).*

Proof. The existence of G^* implies, by definition of a dual, that G satisfies (1.3). By Lemma 1.8, then, G is countable.

Suppose G^* violates (1.3). Then there are two vertices in G^* , x and y say, that cannot be separated by finitely many edges. By Theorem 4.9 applied to the line graph of G^* (which is countable because G is), we can find in G^* an infinite set \mathcal{P} of edge-disjoint x - y paths and a set C^* of edges separating x from y , such that C^* consists of a choice of one edge from each path in \mathcal{P} . Then C^* is a bond in G^* , and $C = \{e \mid e^* \in C^*\}$ is an infinite circuit in \tilde{G} .

Pick an edge $e \in C$. We claim the following:

There is an infinite sequence of finite circuits C_1, C_2, \dots in G , each containing e , and such that $C_i \setminus C$ and $C_j \setminus C$ are non-empty and disjoint for all $i \neq j$. (4.1)

To prove (4.1), assume inductively that C_1, \dots, C_{i-1} have been constructed, and put

$$X := \{e\} \cup \bigcup_{j < i} C_j \setminus C.$$

Since C meets X only in e , Lemma 4.8 gives us a finite circuit C_i that contains e and does not meet $C_1 \cup \dots \cup C_{i-1}$ outside C . As both C and C_i are circuits in \tilde{G} , neither contains the other properly, so $C_i \setminus C \neq \emptyset$. This proves (4.1).

Let u, v be the endvertices of e^* in G^* . Each of the sets C_i^* is a cut in G^* that contains e^* , and hence separates u from v in G^* . Denote by P the path in \mathcal{P} that contains e^* . Since $E(P)$ meets C^* only in e^* , no edge of P other than e^* lies in more than one of the sets C_i^* (by (4.1)). Therefore only finitely many of the sets C_i^* meet $E(P - e^*)$; let C_n^* be one that does not. Since C_n^* is finite, there is a path $Q \in \mathcal{P}$ that has no edge in C_n^* . But then $(P - e^*) \cup Q$ is a connected subgraph of G^* that avoids C_n^* but contains both u and v , a contradiction. \square

As pointed out before, every dual of a multigraph is also its finitary dual, and we have just seen that it satisfies (1.3). Our next aim is to show that, conversely, every finitary dual satisfying (1.3) is even a dual. We need the following lemma:

Lemma 4.11.[16] *A set $F \subseteq E(G)$ in a multigraph G is a cut if and only if it meets every finite circuit in an even number of edges.*

Proof. Clearly, a cut meets every finite circuit in an even number of edges, so let us prove the other direction.

Let G' be the multigraph obtained from G by contracting every edge not in F . Then G' is bipartite (in particular, loopless), since any odd circuit would give rise to a finite circuit in G meeting F in an odd number of edges. The bipartition of G' induces a partition (A, B) of the vertex set of G such that every edge in F has one vertex in A and the other in B and such that no edge outside F has that property. Thus, $F = E(A, B)$ is a cut. \square

Lemma 4.12.[16] *Let G be a 2-connected multigraph, and let G^* be a finitary dual of G that satisfies (1.3). Then the following assertions hold:*

- (i) G^* is a dual of G (witnessed by the same map $*$).
- (ii) A set $F \subseteq E(G)$ lies in $\mathcal{C}(\tilde{G})$ if and only if F^* is a cut in G^* .

Proof. We first prove (ii). Let $F \in \mathcal{C}(\tilde{G})$ be given. To show that F^* is a cut in G^* , it suffices by Lemma 4.11 to show that F^* meets every finite circuit Z^* in G^* in an even number of edges. Since G is a finitary dual of G^* (Theorem 4.3), Z is a finite cut in G . Hence by Theorem 1.6, $|F \cap Z| = |F^* \cap Z^*|$ is even.

Similarly, consider a cut F^* in G^* . To show that $F \in \mathcal{C}(\tilde{G})$, it suffices by Theorem 1.6 to show that F meets every finite cut D of G evenly. But $D^* \in \mathcal{C}(\tilde{G}^*)$ since G is a finitary dual of G^* (Theorem 4.3). So $|F \cap D| = |F^* \cap D^*|$ is even by the trivial direction of Lemma 4.11 and the fact that D^* is a disjoint union of circuits.

To prove (i), let now C be a circuit in \tilde{G} ; we have to show that C^* is a bond in G^* . We have already shown that C^* is a cut in G^* , and that any cut $F^* \subseteq C^*$ corresponds to a set $F \in \mathcal{C}(\tilde{G})$. Since F cannot be a proper subset of C unless it is empty, we deduce that C^* is a bond.

Conversely, let F^* be a bond in G^* . Then F is a minimal non-empty element of $\mathcal{C}(\tilde{G})$. By Theorem 1.7, F must be a circuit. \square

For a proof of Theorem 4.6, it remains to combine Lemmas 4.10 and 4.12 with Thomassen's results on finitary duals, and to extend the result from 2-connected to arbitrary multigraphs.

The latter is standard for finite graphs, but we have to be more careful here. (Indeed, Lemma 4.13 below fails for finitary duals.) If G has a finitary dual G^* and B is a block of G , let B^* denote the submultigraph of G^* formed by the edges e^* with $e \in B$ and their incident vertices.

Lemma 4.13. [16] *Let a multigraph G have a finitary dual G^* that satisfies (1.3). If B is a block of G , then B^* is a block of G^* and a finitary dual of B . If G^* is even a dual of G then B^* is a dual of B .*

Proof. It is not hard to check that two edges $e, f \in G$ lie in a common block of G if and only if they lie in a common finite circuit of \tilde{G} . For a proof that B^* is a block of G^* , it therefore suffices to show that the edges e^* and f^* lie in a common block of G^* if and only if they lie in a common finite bond of G^* .

If e^* and f^* lie in a common bond F^* of G^* , then the edges in F^* that lie in the same block as e^* suffice to separate the endvertices of e^* in G^* . By the minimality of F^* , these are all its edges, including f^* .

Now suppose that e^* and f^* lie in a common block B^* of G^* . Then B^* has a finite circuit containing both e^* and f^* . Deleting e^* and f^* from this circuit, we obtain the edge sets of two paths, P and Q . Suppose that every $X \subseteq E(G^*)$ separating P and Q is infinite. Then there are also two vertices, one in P and the other in Q , that cannot be separated in G^* by finitely many edges. Consequently, we find infinitely many edge-disjoint paths connecting these two vertices, a contradiction to (1.3). Therefore, there is a finite set $F^* \subseteq E(G^*)$ separating P from Q in G^* , which clearly contains e^* and f^* . If we choose F^* to be minimal then it is a bond.

It remains to show that B^* is a finitary dual (resp. dual) of B . But since B^* is a block of G^* , a set $F^* \subseteq E(B^*)$ is a bond of B^* if and only if it is a bond of G^* . Similarly, a set $F \subseteq B$ is a circuit in \tilde{B} if and only if it is a circuit in \tilde{G} . The assertion therefore follows from the assumption that G^* is a finitary dual (resp. dual) of G . \square

Lemma 4.14. [16] *If G and G^* are two multigraphs and $*$: $E(G) \rightarrow E(G^*)$ maps the blocks B of G to the blocks of G^* so that B^* is a dual of B , then G^* is a dual of G .*

Proof. It is easily checked that a subset of $E(G)$ is a circuit in \tilde{G} if and only if it is a circuit in \tilde{B} for some block B of G . Similarly, a subset of $E(G^*)$ is a bond in G^* if and only if it is a bond in some block of G^* . \square

Proof of Theorem 4.6. (i) By Lemmas 4.13 and 4.14, G has a dual if and only if its blocks do. (To obtain a dual of G from duals of its blocks, take their disjoint union.) Similarly, a countable multigraph is planar if and only

if its blocks are (Dirac and Schuster [37]). We may therefore assume that G is 2-connected.

If G is planar then, by Theorems 4.2 and 4.3, G has a finitary dual G^* that satisfies (1.3). By Lemma 4.12, G^* is even a dual of G . Conversely, if G has a dual G^* , then G is planar by Theorem 4.2.

(ii) Suppose that G^* is a dual of G . By Lemma 4.13, the subgraphs B^* of G^* , where B ranges over the blocks of G , are the blocks of G^* , and each B^* is a dual of B . We show that, conversely, B is a dual of B^* . Then, by Lemma 4.14, G is a dual of G^* .

By Lemma 4.10, every B^* satisfies (1.3). (Hence so does G^* .) By Theorem 4.3, B is a finitary dual of B^* . Now B satisfies (1.3), because G does so by assumption. Hence by Lemma 4.12, B is a dual of B^* .

(iii) For 2-connected graphs, this is Lemma 4.12(ii). The general case reduces easily to this with the help of Lemma 4.11 and Theorem 1.6. \square

4.4 Locally finite duals

We started out by observing that a dual of a locally finite graph may have vertices of infinite degree. This raises the question under what circumstances the dual is locally finite. For 3-connected graphs, Thomassen gave the following characterisation in terms of peripheral circuits. Recall that a peripheral circuit C is a circuit whose incident vertices do not separate the graph and do not span any edges not in C .

Theorem 4.15 (Thomassen [59]). *Let G be a locally finite 3-connected graph. Then G has a locally finite finitary dual if and only if G is planar and every edge lies in exactly two finite peripheral circuits.*

Since locally finite graphs trivially satisfy condition (1.3), Lemma 4.12 implies that Theorem 4.15 still holds if the word ‘finitary’ is dropped.

To obtain another characterisation, we need the extension of Kelmans’ planarity criterion from Chapter 3, which we restate:

Theorem 3.22. [22] *Let G be a locally finite 3-connected graph. If G is planar then every edge appears in exactly two peripheral circuits. Conversely, if every edge appears in at most two peripheral circuits then G is planar.*

Theorem 4.16. [16] *A locally finite 3-connected graph has a locally finite dual if and only if it is planar and all its peripheral circuits are finite.*

Proof. Let G be a locally finite 3-connected graph. If G has a locally finite dual then, by Theorem 4.15, G is planar and every edge lies in exactly two

finite peripheral circuits. By Theorem 3.22, its edges cannot lie in any other peripheral circuits, so all peripheral circuits are finite.

Conversely, if G is planar and all its peripheral circuits are finite then, by Theorems 3.22 and 4.15, G has a locally finite finitary dual. By Lemma 4.12, this is in fact a dual. \square

4.5 Duality in terms of spanning trees

In this section we show that our notion of duality permits the extension of another well-known duality theorem for finite multigraphs: that the complement of the edge set of any spanning tree of G defines a spanning tree in any dual of G , and conversely that any two multigraphs whose edge sets are in bijective correspondence so that their spanning trees complement each other as above form a pair of duals [16].

It is not difficult to see that the verbatim analogue of this fails for infinite multigraphs. Indeed, the edge set of an ordinary spanning tree of G might contain an infinite circuit C (such as the edges of the double ray D in Figure 4.1), in which case C^* would be a cut in G^* , and $G^* - C^*$ could not contain a spanning tree of G^* . However, the following adjustment to the notion of a spanning tree makes an extension possible.

Let us call a spanning tree T of G *acirclic* (under ITOP) if its closure in \tilde{G} contains no circle – or equivalently, if its edges contain no circuit of \tilde{G} . (We remark that if G is locally finite then its acirclic spanning trees are precisely its end-faithful spanning trees; see Diestel and Kühn [35].)

Theorem 4.17. [16] *Let $G = (V, E)$ and $G^* = (V^*, E^*)$ be connected⁵ multigraphs satisfying (1.3), and let $*$: $E \rightarrow E^*$ 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 multigraph (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 ITOP).*

Before we prove Theorem 4.17, let us show that those acirclic spanning trees always exist. We need the following standard lemma; see Diestel [28] or Diestel and Kühn [32] for a proof.

⁵This assumption is for convenience only. For disconnected multigraphs, one has to replace ‘acirclic spanning tree’ with ‘subgraph inducing an acirclic spanning tree in every component’.

Lemma 4.18. *Let U be an infinite set of vertices in a connected graph T . Then T contains either a ray R and infinitely many disjoint U - R paths (possibly trivial) or a subdivided star with infinitely many leaves in U .*

In order to apply a lemma of Hahn, Laviolette and Širáň [39], we need a few more definitions. Call a ray R in a graph G *edge-dominated* if there exists a vertex $v \in G$ such that G contains infinitely many edge-disjoint v - R paths whose last vertices are distinct. Given any spanning tree T of G , let $\eta: \Omega(T) \rightarrow \Omega(G)$ map every end of T to the unique end of G that contains it as a subset (of rays). Given $\Omega' \subseteq \Omega(G)$, let us say that T *represents only Ω'* if η is injective and its image is (only) Ω' .

Lemma 4.19 (Hahn, Laviolette & Širáň [39]). *Let G be a countable connected graph in which every edge-dominated ray is dominated. Then G has a spanning tree T that represents only the set Ω' of undominated ends of G .*

We can now prove the existence of acirclic spanning trees, which settles Problem 7.9 of Diestel and Kühn [35].

Theorem 4.20.[16] *Every connected graph G satisfying (1.3) has a spanning tree whose closure in \tilde{G} contains no circle.*

Proof. Let G be a connected graph satisfying (1.3). We may assume that G is 2-connected (and hence countable, by Lemma 1.8), since any union of acirclic spanning trees of its blocks is clearly an acirclic spanning tree of G .

In order to apply Lemma 4.19, let us show that every edge-dominated ray R in G is dominated. Let \mathcal{P} be an infinite set of edge-disjoint v - R paths in G , for some vertex v . If v does not dominate R , then v is separated from R in G by a finite set of vertices distinct from v . One of these, w say, lies on infinitely many paths from \mathcal{P} . Then G contains infinitely many edge-disjoint v - w paths, contradicting (1.3).

Let T be a spanning tree as provided by Lemma 4.19. We show that T is acirclic. Suppose not, let $C \subseteq E(T)$ be a circuit in \tilde{G} , and let $e = xy$ be an edge in C . Let T_x and T_y be the components of $T - e$ containing x and y , respectively, and let F be the set of all T_x - T_y edges in G . Then C meets F exactly in e ; to obtain a contradiction (to Theorem 1.6), it thus suffices to show that F is finite.

Suppose not. Then we may assume that F has infinitely many incident vertices in T_x since, by (1.3), G has no multiple edges of infinite multiplicity. Applying Lemma 4.18 in T_x with the set U_1 of these vertices as U , we obtain an infinite subset U_2 of U_1 linked disjointly to a ray in T_x or to some common vertex $v \in T_x$. Applying Lemma 4.18 again, in the graph T'_y obtained from

T_y by adding the edges from F and their endvertices in T_x , and with U_2 as U , we obtain a similar set of paths in T'_y . Together, the two path systems either violate condition (1.3) or our assumption that T represents only Ω' . (If both applications of Lemma 4.18 yield a ray, then these two rays belong to distinct ends of T but to the same end of G , ie. the map η was not injective.) \square

Proof of Theorem 4.17. (i) \Rightarrow (ii): Let $T = (V, F)$ be an acirclic spanning tree of G in ITOP. Then $E^* \setminus F^*$ contains no circuit C^* of \tilde{G}^* , since C would then be a cut of G missed by T . Similarly $(V^*, E^* \setminus F^*)$ must be connected: if not, then F^* contains a bond of G^* , and F contains the corresponding circuit of \tilde{G} . The converse implication follows by symmetry, since G is a dual of G^* .

(ii) \Rightarrow (i): We show that the map $*$ makes G^* a finitary dual of G . Then Lemma 4.13 implies that for each block B of G the block B^* of G^* is a finitary dual. Then B^* is even a dual of B (with the same map $*$, by Lemma 4.12), and by Theorem 4.6 the inverse of $*$ makes B a dual of B^* . With Lemma 4.14 we obtain (i).

So consider a finite circuit C of G . We first show that C^* contains a cut of G^* . By subdividing each edge of G , we may apply Theorem 4.20 in order to obtain an acirclic spanning tree S^* of G^* . If C^* contains no cut, we can join up the components of $S^* - C^*$ by finitely many edges from $E^* \setminus C^*$ to form another spanning tree T^* of G^* . Then T^* , too, is acirclic: any circle in its closure contains an arc A^* that contains infinitely many edges but avoids the (finitely many) new edges and hence lies in the closure of S^* , so the union of A^* with a suitable path from S^* contains a circle in the closure of S^* . Now use (ii) to find an acirclic spanning tree T of G corresponding to T^* . Since T^* contains no edge from C^* , the edges of T include C , a contradiction.

To show that C^* is even a minimal cut in G^* , we show that for every $e \in C$ and $A := C \setminus \{e\}$ the multigraph $G^* - A^*$ is connected. To do so, it suffices to find a spanning tree T^* of G^* with no edge in A^* , and hence by (ii) to find an acirclic spanning tree of G whose edges include A . Let S be any acirclic spanning tree of G . Since A is finite but contains no circuit, we can obtain another spanning tree T from S by adding all the edges from A and deleting some (finitely many) edges not in A . As before, T is acirclic in \tilde{G} because S was, and hence is as desired.

It remains to show that if B^* is a finite bond in G^* then B is a circuit in G . As before, we first show that B contains a circuit. If not, we can modify an acirclic spanning tree of G into one whose edges include B , which by (ii) corresponds to a spanning tree of G^* that has no edge in B^* (contradiction). On the other hand, given any proper subset D^* of B^* , we can modify an acirclic spanning tree of G^* into one missing D^* , because D^* contains no cut

of G^* . Then this tree corresponds by (ii) to a spanning tree of G whose edges include D , so D is not a circuit in G . \square

4.6 Alternative proof of Theorem 4.20

Lemma 4.19, which is used in the proof of Theorem 4.20, is not trivial. Let us therefore give a (in our context) direct proof of Theorem 4.20 that is not much more complicated than the one above.

We need the following easy fact. Its proof is the same as for finite graphs, eg. as in Diestel [31, Lemma 1.9.4].

Lemma 4.21. *Every cut in a graph is a disjoint union of bonds.*

Alternative proof of Theorem 4.20. [13] We may assume that $G = (V, E)$ is 2-connected. Thus, there is, by Lemma 1.8, an enumeration e_1, e_2, \dots of E . Put $S_0 = T_0 = E$, and inductively define $S_n, T_n \subseteq E$ in the following way. Denote by i the minimal index such that there is a circuit $C \subseteq S_{n-1}$ with $e_i, e_n \in C$; if there is no such circuit set $i = \infty$. Analogously, choose j minimal such that there is a bond $B \subseteq T_{n-1}$ with $e_i, e_n \in B$, and put $j = \infty$ if there is no such bond. If $i < j$ put $S_n := S_{n-1} - e_n$ and $T_n := T_{n-1}$; if $i > j$ put $S_n := S_{n-1}$ and $T_n := T_{n-1} - e_n$; and if $i = j$ choose arbitrarily whether to delete e_n from S_{n-1} or from T_{n-1} .

Define $S := \bigcap_{n=1}^{\infty} S_n$ and $T := \bigcap_{n=1}^{\infty} T_n$, and observe that $S \cup T = E$. Moreover, note that

$$e_n \in S_l \text{ if and only if } e_n \notin T_l \text{ for all } l \geq n \quad (4.2)$$

If we can show that neither S contains a circuit nor T a bond, then (V, S) is an acirclic spanning tree, and we are done. Indeed, if (V, S) is not connected then there is a bond B in G which is disjoint from S . But then $B \subseteq E \setminus S = T$, a contradiction.

So, assume that S contains a circuit or T a bond, and choose i minimal such that there is a set $C \subseteq E$ with $e_i \in C$, and such that $C \subseteq S$ is a circuit in \tilde{G} or such that $C \subseteq T$ is a bond of G . Let us first suppose that $C \subseteq S$, i.e. that C is a circuit.

If C is finite then there is a maximal k so that $e_k \in C$. From $e_k \in S$ it follows that $e_k \in S_k$ but $e_k \notin T_k$. Thus, there is a bond $D \subseteq T_{k-1}$ with $e_k \in D$. As C is finite, and as D is a cut in G , we obtain from $e_k \in C \cap D$ that there is a second shared edge e_j , say. Because of the choice of k , we have $j < k$. But now $e_j \in C \subseteq S_{k-1}$ and $e_j \in D \subseteq T_{k-1}$, a contradiction to (4.2).

Therefore, C is infinite. Let $e_i, e_{k_1}, e_{k_2}, \dots$ be distinct edges in C , and observe that, by the minimality of i , it holds that $i < k_l$ for each l . Since $e_{k_l} \in S_{k_l}$, there is a bond $D_l \subseteq T_{k_l-1}$ containing e_{k_l} and an edge e_{m_l} with $m_l \leq i$; otherwise we would have deleted e_{k_l} from S_{k_l-1} in order to obtain S_{k_l} . Because of $e_i \notin T_{k_l-1}$ as $e_i \in S_{k_l-1}$ it follows that $m_l < i$. As there are only finitely many edges with index smaller than i , infinitely many of the D_l share an edge e_m with $m < i$. Denote the set of these D_l by \mathcal{D}_0 .

Put $E_n := \{e_1, \dots, e_n\}$ for each n , and define inductively infinite sets $\mathcal{D}_0 \supseteq \mathcal{D}_1 \supseteq \mathcal{D}_2 \supseteq \dots$ such that $D \cap E_n = D' \cap E_n$ for all $D, D' \in \mathcal{D}_n$ and all n . Indeed, if infinitely many of the bonds in \mathcal{D}_{n-1} share the edge e_n then let the set of these be \mathcal{D}_n . Otherwise define \mathcal{D}_n to be those of the bonds in \mathcal{D}_{n-1} which do not contain e_n . Finally, choose for each n a $K_n \in \mathcal{D}_n$, so that all the K_n are distinct.

Put $D := \bigcup_{n=1}^{\infty} (K_n \cap E_n)$. We claim that D contains a bond B of G with $e_m \in B$. Indeed, let F be a finite circuit, and choose N large enough so that $F \subseteq E_N$. Then,

$$F \cap D = F \cap D \cap E_N = F \cap K_N,$$

where the last intersection is even, as K_N is a bond (Lemma 4.11). Thus, by the other direction of Lemma 4.11, D is a cut and therefore a disjoint union of bonds, by Lemma 4.21. One of these contains $e_m \in D$.

We claim that $D \subseteq T$. Indeed, consider $e_j \in D$. Then, as e_j lies in all but finitely many of the K_l there is one $K_l = D_s$ such that $k_s > j$. As $e_j \in D_s \subseteq T_{k_s-1}$, we obtain $e_j \in T$. Therefore, $B \subseteq D$ is a bond contained in T but $e_m \in B$ with $m < i$ is a contradiction to the minimal choice of i .

Now, if $C \subseteq T$ is a bond then the proof is analogous with the roles of S and T , and of circuits and bonds interchanged. Instead of Lemmas 4.11 and 4.21 we use Theorems 1.6 and 1.7. \square

4.7 Colouring-flow duality and circuit covers

As an application of Theorem 4.6 and our results from Section 4.3, we now show that the edge set of every bridgeless locally finite planar graph can be covered by two elements of its cycle space [16]. For finite graphs, this is a well-known reformulation of the four colour theorem. For infinite graphs, of course, it must fail as long as the cycle space contains only finite sets of edges.

In our setting, however, Theorem 4.6 enables us to imitate the finite result (and its proof from the four colour theorem), because 4-colourability extends by compactness; see de Bruijn and Erdős [25]. Rather than assuming that G is locally finite, we work slightly more generally in \tilde{G} .

Theorem 4.22. [16] *Let G be a bridgeless planar graph satisfying (1.3). Then there are $Z_1, Z_2 \in \mathcal{C}(\tilde{G})$ such that $E(G) = Z_1 \cup Z_2$.*

Proof. Assume that we find for every block B of G elements Z_1^B, Z_2^B of the cycle space of B such that $E(B) = Z_1^B \cup Z_2^B$. From Theorem 1.6 follows that for $i = 1, 2$, $Z_i^B \in \mathcal{C}(G)$ and then also $Z_i := \sum_B Z_i^B \in \mathcal{C}(G)$, where the sum ranges over the blocks of G . Clearly, we get $E(G) = Z_1 \cup Z_2$. As G is bridgeless, we may therefore assume that G is 2-connected.

By Theorem 4.6, G has a dual G^* and is itself a dual of G^* , which therefore is planar too. By the four colour theorem (Appel and Haken [9, 10]) and compactness [25], G^* has chromatic number at most 4. Choose a 4-colouring $c: V(G^*) \rightarrow \mathbb{Z}_2 \times \mathbb{Z}_2$ of G^* . For $i = 1, 2$, let $c_i: V(G^*) \rightarrow \mathbb{Z}_2$ be c followed by the projection to the i th coordinate, define $f_i: E(G) \rightarrow \mathbb{Z}_2$ by $f_i(e) := c_i(v) + c_i(w)$ where v and w are the endvertices of e^* , and put $Z_i := f_i^{-1}(1)$.

Let us show that every edge e of G lies in Z_1 or Z_2 . If not, then $f_1(e) = f_2(e) = 0$, and hence $c(v) = c(w)$ for $e^* =: vw$. But this contradicts our assumption that c is a proper colouring of G^* .

Next we show that $Z_i \in \mathcal{C}(\tilde{G})$, for both $i = 1, 2$. By Theorem 1.6, it suffices to show that Z_i meets every finite cut F of G in an even number of edges, ie. that

$$f_i(F) := \sum_{e \in F} f_i(e) = 0.$$

As every cut is a disjoint union of bonds (Lemma 4.21) we may assume that F is a bond. Then F^* is a circuit in G^* . Hence,

$$f_i(F) = \sum_{e^* = vw \in F^*} (c_i(v) + c_i(w)) = 2 \sum_{u \in U} c_i(u) = 0,$$

where U is the vertex set of the cycle in G^* whose edge set is F^* . □

4.8 MacLane's criterion revisited

Let us return to MacLane's planarity criterion from Chapter 3. We have seen in Section 3.2 that MacLane's planarity criterion fails for non-locally finite graphs. Nevertheless, duality provides a tool with which MacLane's criterion can be extended to graphs satisfying (1.3), and that in a rather simple way.

Theorem 4.23. [22] *Let G be a countable graph satisfying (1.3). Then G is planar if and only if $\mathcal{C}(\tilde{G})$ has a simple generating system.*

The backward direction of Theorem 4.23 can be shown in exactly the way as detailed in Section 3.4. For the forward direction, we proceed as in

the finite case. Since G is planar, there is, by Theorem 4.6, a dual G^* of G . We claim that $\mathcal{B} := \{F : F^* = E(v) \text{ for some } v \in V(G^*)\}$ is a simple generating set of $\mathcal{C}(\tilde{G})$. Certainly, \mathcal{B} is simple, and, by Theorem 4.6, every of its elements is a member of the cycle space. In addition, \mathcal{B} generates every $Z \in \mathcal{C}(\tilde{G})$. Indeed, Theorem 4.6 implies that Z^* is a cut in G^* . Let (A, B) be the corresponding partition of $V(G^*)$, i.e. $Z^* = E_{G^*}(A, B)$. Then,

$$Z^* = \sum_{v \in A} E(v), \text{ and therefore } Z = \sum_{F \in \mathcal{B}'} F,$$

where $\mathcal{B}' := \{F : F^* = E(v), v \in A\} \subseteq \mathcal{B}$. This completes the forward direction of Theorem 4.23.

Chapter 5

End degrees and infinite circuits

5.1 Introduction

The main objective of this thesis is to extend all, or as many as possible, of the properties of the cycle space of a finite graph to locally finite graphs. So far, we have been quite successful in attaining this objective. In the course of Chapters 2–4 we have shown that all of the theorems we listed in Section 1.1 become true in locally finite graphs once infinite circuits are allowed. That the more basic properties of the cycle space generalise had already been demonstrated by Diestel and Kühn—only one property has been conspicuously absent from the discussion so far: the characterisation of the cycle space elements by degrees.

Theorem 5.1. *Let H be a subgraph of a finite graph G . Then $E(H)$ is an element of the cycle space of G if and only if every vertex of G has even degree in H .*

Simple examples show that for infinite graphs it is not sufficient to consider vertex degrees. Consider, for instance, the double ray D . Since $|D|$ is homeomorphic to the unit interval, it does not contain any circles, hence, it follows that $\mathcal{C}(D) = \{\emptyset\}$. Thus $E(D) \notin \mathcal{C}(D)$, even though every vertex of D has degree 2 in D . The problem here seems to arise from the ends rather than the vertices of the considered graph. Diestel and Kühn [34] raised the following problem:

Problem 5.2. *Characterise the circles and the elements of the cycle space of an infinite graph in purely combinatorial terms, such as vertex degrees and ‘degrees of ends’.*

We introduce a concept of end degrees that is a natural extension of vertex degrees; instead of incident edges we count edge-disjoint rays (for degrees in subgraphs it will be necessary to substitute ‘rays’ with ‘arcs’). This notion allows us to solve the first part of Problem 5.2; we prove a straightforward adaptation of the well-known fact that the cycles in a finite graph are exactly its 2-regular connected subgraphs.

Theorem 5.3. [23] *Let C be a subgraph of a locally finite graph G . Then \overline{C} is a circle if and only if \overline{C} is topologically connected and every vertex or end x of G with $x \in \overline{C}$ has degree two in C .*

Depending on its degree, an end can be assigned a parity, i.e. the label ‘even’ or ‘odd’—as long as it has finite degree. Inspired by Laviolette [45], who introduced a concept to measure the parity of vertices of infinite degree, we assign a parity also to ends of infinite degree. A classification of ends into even and odd ends has already been achieved by Nash-Williams [50] for the case of eulerian graphs with only finitely many ends. Our definition coincides with Nash-Williams’ in these graphs but covers all locally finite graphs. Moreover, with our definition the following important special case of Problem 5.2 becomes true, which is the main result of this chapter.

Theorem 5.4. [23] *Let G be a locally finite graph. Then $E(G) \in \mathcal{C}(G)$ if and only if every vertex and every end of G has even degree.*

An extension of this characterisation to arbitrary subgraphs of G would solve Problem 5.2 completely. We shall offer a conjecture in that respect (see Section 5.3).

This chapter, which is based on [23], is organised as follows. We introduce and discuss our end degree concept in Sections 5.2 and 5.3. Theorem 5.4 will be proved in Section 5.4. In Section 5.5, we show Theorem 5.3 and other results, and in the last section, we briefly discuss an alternative concept of an end degree.

5.2 Defining end degrees

As ends are equivalence classes of rays, the degree of an end should in some way be related to its rays. Also, the rays may be seen as somewhat analogous to the incident edges of a vertex, whose number is the degree of the vertex. We will argue in this section that counting rays, in the right fashion, leads to a suitable notion of an end degree [23].

As a first try, we might count the maximal number of disjoint rays in an end (sometimes called the *multiplicity* of the end). For this notion of

an end degree, however, Theorem 5.4 fails as the graph G in Figure 5.1 demonstrates. Since G contains an odd cut, its edge set is not an element of $\mathcal{C}(G)$ by Theorem 1.4. But each vertex degree is even, and the maximal number of disjoint rays in each end is two.




Figure 5.1: The multiplicity of each end is even, but $E(G) \notin \mathcal{C}(G)$

Looking more closely we see that although the maximal number of disjoint rays in each end is even, the maximal number of *edge-disjoint* rays is odd, namely three. Thus with this measure instead we would have correctly decided that $E(G) \notin \mathcal{C}(G)$. Let us make that more precise. We call a ray R in an end ω an ω -ray, and define the *degree* of an end ω in a graph G as

$$d(\omega) := \sup\{|\mathcal{R}| : \mathcal{R} \text{ is a set of edge-disjoint } \omega\text{-rays}\} \in \mathbb{N} \cup \{\infty\}.$$

Although we will not use this, we shall prove later on (Lemma 5.13) that the supremum is attained, i.e. if $d(\omega) = \infty$, then there exists an infinite set of edge-disjoint ω -rays. Andreae [8] proves a similar result.

Our degree concept clearly divides the ends of finite degree into even and odd ends, but how are we to deal with ends of infinite degree? We may not simply treat them as odd ends, since the edge set of the infinite grid obviously is an element of its cycle space but the single end of the grid has infinite degree; see Figure 5.2 on the left.

On the other hand, classifying all ends of infinite degree as even is not any better: consider the graph G on the right in Figure 5.2. All vertex degrees are even and both ends have infinite degree, but G has an odd cut (which together with Theorem 1.4 implies that $E(G) \notin \mathcal{C}(G)$).

Consequently, the degree, if infinite, is not sufficiently fine enough to determine the parity of an end. For an adequate refinement we will use the following characterisation of ends with even finite degree.

Lemma 5.5. [23] *In a locally finite graph G let $\omega \in \Omega(G)$ have finite degree k . Then the following statements are equivalent:*

- (i) k is even;
- (ii) there is a finite $S \subseteq V(G)$ such that for every finite set $S' \supseteq S$ of vertices the maximal number of edge-disjoint ω -rays starting in S' is even.




Figure 5.2: Even and odd ends of infinite degree

Proof. Consider a set \mathcal{R} of edge-disjoint ω -rays of maximal cardinality $|\mathcal{R}| = k$, and let U be the set of starting vertices of \mathcal{R} . Then, for every finite set $S' \supseteq U$, \mathcal{R} has maximal cardinality among all sets of edge-disjoint ω -rays starting in S' . Thus, putting $S := U$, we deduce that (i) implies (ii). Also, (ii) implies (i), which we see by choosing $S' = S \cup U$. \square

Let us call an induced subgraph C of G a *region*, if the cut $E(C, G - C)$ consists of only finitely many edges. Observe that as every finite set $S \subseteq V(G)$ gives (essentially) rise to a neighbourhood $\hat{C}(S, \omega)$ of ω , condition (ii) in Lemma 5.5 can be alternatively formulated using these neighbourhoods, or using regions whose closures contain ω :

- (ii') There is a region A of G with $\omega \in \overline{A}$ such that for every region $B \subseteq A$ of G with $\omega \in \overline{B}$ the maximal number of edge-disjoint rays of ω starting outside B is even.

This motivates the following definition of the parity of an end: an end ω of a locally finite graph is said to be *even* if ω satisfies (ii) of Lemma 5.5. Otherwise ω is *odd*. Thus, ω is odd if and only if for all finite $S \subseteq V(G)$ there is a finite set $S' \supseteq S$ such that the maximal number of edge-disjoint ω -rays starting in S' is odd. By Lemma 5.5, an end ω of finite degree is even if and only if $d(\omega)$ is even.

However, it is not possible to extend this notion literally to subgraphs H of G . First of all, we cannot simply measure the degrees of the ends of H (as opposed to those of G). This is not surprising as H is embedded in the space $|G|$. If H is a double ray, for instance, then (viewed as a graph on its own and not as a subgraph) it has two ends, each of which has degree 1. On the other hand, the tails of H may lie in the same end of G , in which case \overline{H} is a circle in $|G|$. Thus any end contained in \overline{H} should have degree 2 in H , not 1. Therefore, we only consider ends of G (and not of H).

Even taking that into account, the literal extension to subgraphs fails: Consider the bold subgraph of the graph in Figure 5.3, and let ω be the end of G “to the right”. Then, if we count the edge-disjoint ω -rays that lie in H , we find that apart from tail-equivalence there is only one ω -ray. But as \overline{H} is a circle, we would expect the end to have degree 2. In contrast, if we count *arcs*




Figure 5.3: In subgraphs, counting edge-disjoint rays is not enough.

with endpoint ω rather than ω -rays we obtain the desired degree 2. Counting arcs will indeed turn out to be successful, and the following proposition, which we shall prove in the next section, shows that in G it makes actually no difference whether we count rays or arcs. In analogy to ω -rays we call an arc of which one of its endpoints is an end ω an ω -arc.

Proposition 5.6. [23] *Let G be a locally finite graph, and let $\omega \in \Omega(G)$. Then for every finite $S \subseteq V(G)$ the maximal number of edge-disjoint ω -rays starting in S equals the maximal number of edge-disjoint ω -arcs starting in S .*

Hence, for a subgraph H of a locally finite graph G , and $\omega \in \Omega(G)$, we define, analogously to the definition of $d(\omega)$ given above, the degree of ω in H as

$$d_H(\omega) := \sup\{|\mathcal{R}| : \mathcal{R} \text{ is a set of edge-disjoint } \omega\text{-arcs in } \overline{H}\} \in \mathbb{N} \cup \{\infty\}.$$

We note that the supremum is attained (see Lemma 5.13). Furthermore, observe that $d(\omega) = d_G(\omega)$. Indeed, suppose otherwise, ie. $d(\omega) < d_G(\omega)$. So, in particular, $d(\omega)$ is finite. For a set of $d(\omega) + 1$ edge-disjoint ω -arcs, let $S \subseteq V(G)$ be a choice of exactly one vertex from each of the arcs. Then, by Proposition 5.6, there are also $d(\omega) + 1$ edge-disjoint ω -rays starting in S , a contradiction.

The parity of an end in H is defined as follows:

Definition 5.7. [23] *An end ω of G is even in H if there is a finite $S \subseteq V(G)$ such that for every finite $S' \subseteq V(G)$ with $S' \supseteq S$ the maximal number of edge-disjoint ω -arcs in \overline{H} starting in S' is even. Otherwise, ω is odd in H .*

Note that by Proposition 5.6, the definition of parity is consistent with the one given previously. Furthermore, it can be seen in a similar way as in the proof of Lemma 5.5 that for an end ω with finite degree in H , ω has even degree in H if and only if $d_H(\omega)$ is even.

A complete solution of Problem 5.2 requires an analogon of Theorem 5.4 for subgraphs H of G . The forward direction of such an analogon can be proved easily with the same methods as used for Theorem 5.4. Moreover, if G has only countably many ends the problem is not overly difficult (Proposition 5.22). In view of this, and in view of Theorems 5.3 and 5.4, and two more results in Section 5.5, which demonstrate that the end degrees behave in many aspects similar to vertex degrees in finite graphs, we offer the following conjecture:

Conjecture 5.8. [23] *Let H be a subgraph of a locally finite graph G . Then $E(H) \in \mathcal{C}(G)$ if and only if every vertex and every end has even degree in H .*

5.3 A cut criterion for the end degree

In this section we prove Proposition 5.6. The other result of this section is Corollary 5.15, which yields a criterion for the parity of an end in terms of cut cardinalities. Let us start with a simple lemma that shows how we can construct a topological path by piecing together infinitely many arcs.

For convenience, we introduce the following notation: if C is a subgraph of a graph G , write $\partial_G C$ or, where no confusion is possible, ∂C , for the cut $E_G(C, G - C)$.

Lemma 5.9. [23] *Let G be a locally finite graph, and let for $n \in \mathbb{N}$, $\phi_n : [0, 1] \rightarrow |G|$ be a homeomorphism such that if $A_n := \phi_n([0, 1])$ it holds that:*

- (i) $A_n \cap A_m \subseteq V(G) \cup \Omega(G)$ for $n \neq m$; and
- (ii) $\phi_n(1) = \phi_{n+1}(0)$ for all n .

Then there is an $x \in |G|$ such that $\bigcup_{n=1}^{\infty} A_n \cup \{x\}$ is a topological path from $\phi_1(0)$ to x .

Proof. Instead of the ϕ_n let us consider compositions with suitable homeomorphisms $\phi'_n : [1 - 2^{-(n-1)}, 1 - 2^{-n}] \rightarrow A_n$. Together the ϕ'_n define, by (ii), a continuous function $\phi' : [0, 1] \rightarrow |G|$. As $|G|$ is compact, the sequence $\phi_1(0) = \phi'(1/2), \phi_2(0) = \phi'(3/4), \dots$ has an accumulation point x . We claim that $\phi : [0, 1] \rightarrow |G|$ defined by $\phi(s) := \phi'(s)$ for $s \in [0, 1)$ and by $\phi(1) := x$ is continuous.

Let a neighbourhood V of x be given, and note that because of (i) and (ii), none of the $\phi_n(0)$ is an inner point of an edge, and thus x is an end. Then there is a basic open neighbourhood $\hat{C}(S, x) \subseteq V$ that contains all but finitely many of the $\phi_n(0)$. By (i), only finitely many of the A_n meet the finite cut $\partial C(S, x)$. So, there is an N such that $A_n \subseteq \hat{C}(S, x)$ for $n \geq N$. Consequently, $\phi^{-1}(V)$ contains the open set $(1 - 2^{-N}, 1]$, and thus is a neighbourhood of 1 in $[0, 1]$. \square

Menger's theorem applied to the line graph implies that between any two finite edge sets E_1, E_2 in a graph there are as many edge-disjoint E_1 - E_2 paths as the minimal number of edges needed in order to separate E_1 and E_2 . The following lemma generalises this result to arcs. Let us say that an arc A is an E_1 - E_2 arc if it has exactly one edge in E_1 , exactly one in E_2 , and these are incident with an endpoint of A .

Lemma 5.10. [23] *Let H be a subgraph of a locally finite graph G . Let $E_1, E_2 \subseteq E(H)$ be finite. Then the maximal number of edge-disjoint E_1 - E_2 arcs in $\bar{H} \subseteq |G|$ equals the minimum k such that there is a finite set $X \subseteq E(G)$ separating E_1 from E_2 in G with $k = |X \cap E(H)|$.*

We need König's infinity lemma, which we restate:

Lemma 2.6 (König'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, such that for $n \geq 2$ 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 Lemma 5.10. Let S be a finite vertex set such that $E_1 \cup E_2 \subseteq E(G[S])$, let v_1, v_2, \dots be an enumeration of $V(G)$, and put $G_n := G[S \cup \{v_1, \dots, v_n\}]$ for $n \in \mathbb{N}$. Let \mathcal{L}_n be the set of all sets M satisfying

- (i) M is a set of pairwise edge-disjoint subgraphs of H ;
- (ii) for each $L \in M$ there is an E_1 - E_2 path P with $P \cap G_n = L$; and
- (iii) $|M| \geq k$.

Let us show that \mathcal{L}_n is non-empty for each n . Contract each component of $G - G_n$ to a vertex (keeping parallel edges but deleting loops), and denote the resulting finite graph by \tilde{G}_n . Let \tilde{H}_n be the subgraph of \tilde{G}_n that consists of the edges in $E(H) \cap E(\tilde{G}_n)$ together with the incident vertices. Now, let \tilde{M} be a set of edge-disjoint E_1 - E_2 paths in \tilde{H}_n of maximal cardinality. By Menger's theorem applied to the line graph, there is a finite set $\tilde{X} \subseteq E(\tilde{H})$ of cardinality $|\tilde{M}|$ that separates E_1 from E_2 in \tilde{H}_n . Then $X := \tilde{X} \cup (E(\tilde{G}_n) \setminus$

$E(\tilde{H}_n)$ separates E_1 from E_2 in \tilde{G}_n , and $|X \cap E(H)| = |\tilde{M}|$. Next, observe that X is also an E_1 - E_2 separator in G , implying $|M| = |X \cap E(H)| \geq k$. Put $M := \{\tilde{L} \cap G_n : \tilde{L} \in \tilde{M}\}$, and note that this choice satisfies (i) and (iii). Furthermore, by replacing each vertex of $\tilde{G}_n - G$ that \tilde{L} meets with a path through the respective component of $G - G_n$, we easily find for the corresponding $L \in M$ a path such that (ii) is satisfied. Thus, $\mathcal{L}_n \neq \emptyset$.

Define a graph on the vertex set $\bigcup_{n=1}^{\infty} \mathcal{L}_n$ with edges $M_n M_{n+1}$ for $M_n \in \mathcal{L}_n$ and $M_{n+1} \in \mathcal{L}_{n+1}$ if $M_n = \{L \cap G_n : L \in M_{n+1}\}$. As every $M \in \mathcal{L}_{n+1}$ has a neighbour in \mathcal{L}_n we may apply Lemma 2.6, which yields a ray $M_1 M_2 \dots$ with $M_n \in \mathcal{L}_n$ for all $n \geq 1$. For each $L_1 \in M_1$ there is a sequence L_1, L_2, \dots with $L_n \in M_n$ and $L_n = L_{n+1} \cap G_n$. More precisely, there are, by (iii), k such sequences so that their respective unions $L^1 = \bigcup L_n^1 \dots, L^k = \bigcup L_n^k$ are mutually edge-disjoint. Furthermore, $L^i \subseteq \overline{H}$, and $L^i \cap G_n = L_n^i \in M_n$ for $i = 1, \dots, k$.

We claim that each of the L^i contains an E_1 - E_2 arc. Indeed, consider an $i \in \{1, \dots, k\}$, and let v and w be the endvertices of the path P for which $P \cap G_1 = L_1^i$. Introduce a new vertex z to G , and link it to v and w . Denote by Q the resulting v - w path vzw . We want to show that $E(L^i \cup Q)$ is an element of the cycle space of $G' := G \cup Q$. To this end, let F be a finite cut of G' , and choose n large enough so that $F \subseteq E(G_n \cup Q)$. Denote by P_n the v - w path with $P_n \cap G_n = L_n^i$, which exists by (ii). Then

$$\begin{aligned} F \cap E(L^i \cup Q) &= F \cap (E(G_n \cap L^i) \cup Q) \\ &= F \cap (E(G_n \cap L_n^i) \cup Q) = F \cap E(P_n \cup Q). \end{aligned}$$

The last intersection is an even set as $E(P_n \cup Q)$ is a circuit in G' , and hence, $E(L^i \cup Q)$ an element of the cycle space of G' , by Theorem 1.4. Thus, by Theorem 1.5 there is a circle D with $Q \subseteq D \subseteq L^i \cup Q$. Hence, $A^i := D \setminus (Q \setminus \{v, w\})$ is an arc from v to w . Therefore, there are k edge-disjoint E_1 - E_2 arcs A^1, \dots, A^k in \overline{H} .

Since there is a finite set $X \subseteq E(G)$ separating E_1 from E_2 in G such that $k = |X \cap E(H)|$, there cannot be more than $k + 1$ arcs in \overline{H} connecting E_1 and E_2 , as each of them meets X . \square

Corollary 5.11. [23] *Let G be a locally finite graph, let H be a subgraph, and let $C_1 \supseteq C_2$ be regions of G . Then the maximal number of edge-disjoint $\partial C_1 \cap E(H)$ - $\partial C_2 \cap E(H)$ arcs in \overline{H} equals the minimum k such that there is a region D with $C_1 \supseteq D \supseteq C_2$ and $|\partial D \cap E(H)| = k$.*

Proof. By Lemma 5.10, the maximal number of edge-disjoint $\partial C_1 \cap E(H)$ - $\partial C_2 \cap E(H)$ arcs in \overline{H} equals the minimal $k' \in \mathbb{N}$ such that there is a finite $X \subseteq E(G)$ with $|X \cap E(H)| = k'$ which separates $\partial C_1 \cap E(H)$ from $\partial C_2 \cap$

$E(H)$ in G . Any such X with $|X|$ minimal gives rise to a region D as above, hence $k = k'$. \square

Let us say that a finite cut $F \subseteq E(G)$ of G separates a set $S \subseteq V(G)$ from an end $\omega \in \Omega$ if every ray of ω that starts in S uses an edge of F . This is equivalent to that the (unique) component C of $G - F$ with $\omega \in \overline{C}$ is disjoint from S . Similarly, F separates two ends ω and ω' , if the closure of each component of $G - F$ contains at most one of ω, ω' .

We obtain a Mengerian criterion:

Lemma 5.12. [23] *Let G be a locally finite graph, let H be a subgraph, let $\omega \in \Omega(G)$, and let $S \subseteq V(G)$ be finite. Then the maximal number of edge-disjoint ω -arcs in \overline{H} starting in S equals the minimum $|F \cap E(H)|$ over all cuts F separating S from ω .*

Proof. First, choose a region $C_1 \subseteq C_0 := C(S, \omega)$ with $\omega \in \overline{C_1}$ such that $|\partial C_1 \cap E(H)|$ is minimal among all regions $C \subseteq C_0$ with $\omega \in \overline{C}$. To prove the assertion it suffices to find a set of edge-disjoint ω -arcs in \overline{H} starting in S , where each of the arcs uses exactly one edge from $\partial C_1 \cap E(H)$.

Next, observe that the closure of one of the components C of $C_1 - N(G - C_1)$ contains ω . Choose C_2 such that $|\partial C_2 \cap E(H)|$ is minimal among all regions $C \subseteq C_1$ satisfying $C \cup N(C) \subseteq C_1$ and $\omega \in \overline{C}$. Continuing in this manner, we obtain regions C_i with $\omega \in \overline{C_i}$ that satisfy for $i \geq 1$:

$$C_{i+1} \cup N(C_{i+1}) \subseteq C_i; \text{ and} \quad (5.1)$$

$$|\partial C_i \cap E(H)| \leq |\partial D \cap E(H)| \text{ for every region } D \text{ with } C_i \supseteq D \supseteq C_{i+1} \quad (5.2)$$

We claim that

there exists a set \mathcal{A} of edge-disjoint ω -arcs in \overline{H} starting outside C_1 such that $|\partial C_1 \cap E(H)| = |\mathcal{A}|$. (5.3)

Indeed, because of (5.2), Corollary 5.11 yields for every $i \geq 1$ a set \mathcal{P}_i of edge-disjoint $(\partial C_i \cap E(H)) - (\partial C_{i+1} \cap E(H))$ arcs in \overline{H} , with $|\mathcal{P}_i| = |\partial C_i \cap E(H)|$. Hence, for any edge $e \in \partial C_1 \cap E(H)$ there is an $A_1 \in \mathcal{P}_1$ that starts in e . The arc A_1 ends in an edge $e' \in \partial C_2 \cap E(H)$, and thus there exists an arc $A_2 \in \mathcal{P}_2$ such that $A_1 \cap A_2 = e'$. In this manner we find a sequence A_1, A_2, \dots so that A_i and A_{i+1} overlap in exactly one edge of $\partial C_{i+1} \cap E(H)$. As each $A_i \in \mathcal{P}_i$ is an $(\partial C_i \cap E(H)) - (\partial C_{i+1} \cap E(H))$ arc, A_i and A_j are disjoint for $|i - j| > 1$. Thus, by deleting the last vertex and all inner points of the last edge in each A_i , we obtain a sequence of edge-disjoint arcs to which we may apply Lemma 5.9. This yields an $x \in |G|$ together with an arc

$A_e \subseteq \bigcup_{i=1}^{\infty} A_i \cup \{x\}$ in \overline{H} that starts in e and ends in x , which uses exactly one edge in each $\partial C_i \cap E(H)$.

Suppose $x \neq \omega$. Then, there exists a finite set $T \subseteq V(G)$ that separates ω and x . By (5.1), there is an $N \in \mathbb{N}$ such that $T \cap V(C_N) = \emptyset$. Thus, C_N is completely contained in one component of $G - T$, and only one of ω, x lies in $\overline{C_N}$, a contradiction. So, we obtain for each e an ω -arc A_e , and all these arcs are edge-disjoint since for each i the arcs in \mathcal{P}_i are.

Finally, using Corollary 5.11 we lengthen the A_e in order to obtain a set of edge-disjoint ω -arcs which start in S (this is possible by the minimal choice of C_1). Note that each of these arcs indeed uses exactly one edge from $\partial C_1 \cap E(H)$. \square

Similarly we can show that for a given subgraph H of G and an end $\omega \in \Omega(G)$ there indeed exists a maximal set of edge-disjoint ω -arcs in \overline{H} .

Lemma 5.13. [23] *Let G be a locally finite graph, let H be a subgraph, and let $\omega \in \Omega(G)$ such that $d_H(\omega) = \infty$. Then there is an infinite set of edge-disjoint ω -arcs in \overline{H} .*

Proof. As in the proof of Lemma 5.12 we define a sequence of regions C_i satisfying (5.1) and (5.2). Again, Corollary 5.11 yields for each $i \in \mathbb{N}$ a set of edge-disjoint $(\partial C_{i-1} \cap E(H)) - (\partial C_i \cap E(H))$ arcs, which we piece together with the help of Lemma 5.9 to obtain a set of edge-disjoint ω -arcs in \overline{H} (whose union contains all edges of all of the $\partial C_i \cap E(H)$). Note that this set is indeed infinite, since $d_H(\omega) = \infty$, and hence by Lemma 5.12, we may assume $\partial C_{i-1} \cap E(H) < \partial C_i \cap E(H)$ for all $i \in \mathbb{N}$. \square

We finally prove Proposition 5.6, which we restate:

Proposition 5.6. [23] *Let G be a locally finite graph, and let ω be an end of G . Then for every finite set $S \subseteq V(G)$ the maximal number of edge-disjoint ω -rays starting in S equals the maximal number of edge-disjoint ω -arcs starting in S .*

Proof. By Lemma 5.12, the maximal number of edge-disjoint ω -arcs starting in S equals the minimal cardinality of a finite cut that separates S from ω . This minimal cardinality, on the other hand, equals the maximal number of edge-disjoint ω -rays starting in S : there clearly cannot be more ω -rays, and conversely, using Menger's theorem applied to the line graph, we can piece together the rays we need along minimal separating cuts, in a similar fashion as in Lemma 5.12. \square

Two further results that follow immediately from Lemma 5.12 are characterisations of end degrees in terms of cut cardinalities:

Corollary 5.14. [23] *Let G be a locally finite graph, let H be a subgraph, and let $\omega \in \Omega(G)$. Then $d_H(\omega) = k \in \mathbb{N}$ if and only if k is the smallest integer such that every finite $S \subseteq V(G)$ can be separated from ω with a finite cut that shares exactly k edges with $E(H)$.*

Corollary 5.15. [23] *Let G be a locally finite graph, let H be a subgraph, and let $\omega \in \Omega(G)$. Then ω has even degree in H if and only if there is a finite $S \subseteq V(G)$ such that for every finite $S' \subseteq V(G)$ with $S' \supseteq S$ it holds: if $F \subseteq E(G)$ is a finite cut separating S' and ω with $|F \cap E(H)|$ minimal, then $|F \cap E(H)|$ is even.*

5.4 Proof of Theorem 5.4

The forward direction follows from Theorem 1.4, which ensures that every finite cut of G is even, and thus together with Corollary 5.15 implies the assertion.

The backward direction, which will occupy us in this section, takes some more effort [23]. Suppose that $E(G) \notin \mathcal{C}(G)$. Observe that we may assume G to be connected, which means in particular that G is countable. We shall find a sequence $C_1 \supseteq C_2 \supseteq \dots$ of regions of G that satisfy

- (i) $\partial_G C_n$ is an odd cut, for $n \geq 1$;
- (ii) $C_n \cup N(C_n) \subseteq C_{n-1}$, for $n \geq 2$; and
- (iii) if D is a region of G with $C_{n-1} \supseteq D \supseteq C_n$ then $|\partial_G D| \geq |\partial_G C_{n-1}|$, for $n \geq 2$,

Then G has an odd end, contradicting the assumption, as desired. Indeed, by piecing together paths in the C_n , we see that there is a ray R which has a tail in every C_n . Let ω be the end with $R \in \omega$, and consider any finite $S \subseteq V(G)$. Choose I large enough such that $C_I \subseteq C(S, \omega)$, which is possible by (ii). So, every cut that separates $S' := N(C_I)$ from ω has cardinality at least $|\partial_G C_I|$, by (iii). Thus by (i) and Corollary 5.15, ω has odd degree.

For our construction, we need a further condition for $n \geq 1$. Let us call a region C of a graph H a k -region if $|\partial_H C| = k$, and let us say that C is *even* (resp. *odd*) if k is even (resp. odd).

- (iv) for every k -region $D \subseteq C_n$ of G with $k < |\partial_G C_n|$ there is an $\ell \in \mathbb{N}$ and even regions $K_1, \dots, K_\ell \subseteq C_n$ such that $|\partial_G K_i| \leq k$ for all i , and $V(D) \subseteq \bigcup_{i=1}^{\ell} V(K_i)$.

This condition, of course, is trivially satisfied if k is even.

As $E(G) \notin \mathcal{C}(G)$, Theorem 1.4 ensures the existence of odd regions in G . Choose any odd region C_1 such that $\partial_G C_1$ has minimal cardinality. This choice satisfies (i) and (iv), and that is all we required for $n = 1$.

Now, suppose the C_i to be defined for $i \leq n$. In order to find a suitable C_{n+1} , we shall contract certain even k -regions D of G (contained in C_n) for which $k < |\partial_G C_n|$. In the resulting minor, which has only big cuts, we will choose a small odd cut, which in G induces the desired region C_{n+1} .

We will construct this minor in several steps. More precisely, for each even integer $m < |\partial_G C_n|$ we define a minor G^m of $G =: G^0$, which will have the properties:

- (a) G^m is obtained from G^{m-2} by contracting disjoint infinite m -regions K of G^{m-2} with $E(K) \subseteq E(C_n)$ for $m \geq 2$; and
- (b) $|E(D) \cap E(G^m)| < \infty$ for every k -region D of G with $D \subseteq C_n$ and $k \leq m$.

Observe that, by (a) and as all vertices of G are even, all vertices of G^m have even degree too. We claim that (b) together with (iv) implies for $m < |\partial_G C_n| - 2$:

- (c) every k -region D of G^m with $E(D) \subseteq E(C_n)$ and $k \leq m + 1$ is finite.

Indeed, consider a k -region D of G^m with $E(D) \subseteq E(C_n)$ and $k \leq m + 1$. By uncontracting, we obtain from D a region D' of G with $\partial_G D' = \partial_{G^m} D$ and $E(D) \subseteq E(D')$. Since by assumption $m < |\partial_G C_n| - 2$, we get that $k < |\partial_G C_n|$. Then (iv) implies that there is a finite set \mathcal{K} of regions $K \subseteq C_n$ such that their union contains all vertices of D' (and thus also all but finitely many edges of D'). Each $K \in \mathcal{K}$ is an ℓ -region with even $\ell \leq k = m + 1$. As $m + 1$ is odd we get $\ell \leq m$, and hence by (b), that $E(K) \cap E(G^m)$ is finite. Thus $|E(D') \cap E(G^m)| < \infty$, and hence D is finite. This establishes (c).

As G is connected, $G^0 = G$ obviously satisfies (b), which is all we required for $m = 0$. So, assume $m \geq 2$, and G^i to be constructed for all even $i < m$. We define a sequence $(L_j)_{j \in \mathbb{N}}$ of (not necessarily induced) subgraphs of G^{m-2} ; by contracting the components of their union L we obtain G^m .

Consider an enumeration R_1, R_2, \dots of all infinite m -regions of G^{m-2} with $E(R_i) \subseteq E(C_n)$ (such an enumeration is possible since $E(G^{m-2}) \subseteq E(G)$ is countable). Put $L_1 := R_1$, and let for $j > 1$,

$$L_j := L_{j-1} \cup R_j \text{ if } \partial_{G^{m-2}} R_j \cap E(L_{j-1}) = \emptyset \quad (5.4)$$

and $L_j := L_{j-1}$ otherwise. Note that in the former case each component of L_{j-1} is either contained in R_j or disjoint from R_j . Thus, by induction on j , every component K of L_j is an infinite m -region of G^{m-2} .

Put $L := \bigcup_{j \in \mathbb{N}} L_j$, and consider a component K of L . Certainly, K is an infinite induced subgraph in G^{m-2} with $E(K) \subseteq E(C_n)$. We claim that $k := |\partial_{G^{m-2}} K| = m$. Clearly, $k \leq m$ as otherwise there would already be a component $K' \subseteq K$ of some L_j with $|\partial_{G^{m-2}} K'| > m$, which is impossible. On the other hand, $k \geq m$, by (c) for $m-2$; thus $k = m$, as desired. We now obtain G^m from G^{m-2} by contracting the components of L to one vertex each (keeping multiple edges but deleting loops). Obviously, G^m satisfies (a).

Before we show the validity of (b), let us prove for all even $i < m < |\partial_G C_n|$ that it holds that:

- (*) for every region $D \subseteq C_n$ of G with $\partial_G D \subseteq E(G^i)$ there is a (possibly empty) induced subgraph $D' \subseteq C_n$ that satisfies
 - (I) there are finitely many regions K_1, \dots, K_ℓ of G each of which contracts to a vertex of degree $\leq i+2$ in G^{i+2} such that $V(D) \setminus V(D') \subseteq \bigcup_{j=1}^\ell V(K_j)$;
 - (II) if there is a region C of G with $D \subseteq C$ and $\partial_G C \subseteq E(G^{i+2})$, then also $D' \subseteq C$;
 - (III) $|\partial_G D'| \leq |\partial_G D|$; and
 - (IV) $\partial_G D' \subseteq E(G^{i+2})$.

Observe that D' has at most $|\partial_G D'|$ components (since G is connected). Each of these is a region of G with properties (II)–(IV).

Let us now show (*). Given D as above, choose an induced subgraph $\tilde{D} \subseteq C_n$ such that $|\partial_G \tilde{D} \setminus E(G^{i+2})|$ is minimal among all induced subgraphs that satisfy (I), (II), (III) and $\partial_G \tilde{D} \subseteq E(G^i)$ (which is possible as D itself has these four properties). If $|\partial_G \tilde{D} \setminus E(G^{i+2})| = 0$, we may put $D' := \tilde{D}$, so suppose otherwise. Then, by (a), there is an $(i+2)$ -region K of G^i with $E(K) \subseteq E(C_n)$, which is contracted to a vertex in G^{i+2} and for which holds that $\partial_G \tilde{D} \cap E(K) \neq \emptyset$. Denote by \tilde{D}^i the image of \tilde{D} in G^i , ie. the induced subgraph of G^i with $\partial_{G^i} \tilde{D}^i = \partial_G \tilde{D}$ and $E(\tilde{D}) \cap E(G^i) = E(\tilde{D}^i)$.

Suppose that one of $|E_{G^i}(K \cap \tilde{D}^i, \tilde{D}^i \setminus K)|$, $|E_{G^i}(K \setminus \tilde{D}^i, G^i - (\tilde{D}^i \cup K))|$ is smaller than or equal to $|E_{G^i}(K \cap \tilde{D}^i, K \setminus \tilde{D}^i)|$. Then, putting either $\hat{D}^i = \tilde{D}^i \setminus K$ or $\hat{D}^i = G^i[\tilde{D}^i \cup K]$ we get

$$|\partial_{G^i} \hat{D}^i| \leq |\partial_{G^i} \tilde{D}^i| = |\partial_G \tilde{D}|. \quad (5.5)$$

Observe that $\partial_{G^i} \hat{D}^i \cap E(K) = \emptyset$, and denote by \hat{D} the induced subgraph of G that we obtain from \hat{D}^i by uncontracting. Then $\partial_G \hat{D} = \partial_{G^i} \hat{D}^i$ has fewer edges outside $E(G^{i+2})$ than $\partial_G \tilde{D}$. We claim that this contradicts the minimal choice of \tilde{D} . Indeed, $\partial_G \hat{D} \subseteq E(G^i)$, and also (I) and (III) hold for \hat{D} : the

latter by (5.5), and for the former observe that each K_i either still contracts to a vertex of degree $\leq i + 2$ in G^{i+2} or is contained in a region which does so. Adding K to these regions, we obtain the K_i as desired for (I).

To see that \hat{D} satisfies (II), consider a region $C \supseteq D$ with $\partial_G C \subseteq E(G^{i+2})$. Observe that $\tilde{D} \subseteq C$ because \tilde{D} satisfies (II). Now, since $\partial_G C \cap E(K) = \emptyset$, either $E(K) \subseteq E(C)$ or $E(K) \subseteq E(G - C)$ because K is connected. The latter case is impossible, as $\partial_G \tilde{D} \cap E(K) \neq \emptyset$. Hence, $E(K) \subseteq E(C)$, and thus, as $\hat{D}^i \subseteq G^i[\tilde{D}^i \cup K]$, we get $\hat{D} \subseteq C$, as desired. Note also that $\hat{D} \subseteq C_n$, as $\partial_G C_n \subseteq E(G^{i+2})$ by (a).

We may therefore assume that

$$|E_{G^i}(K \cap \tilde{D}^i, K \setminus \tilde{D}^i)| < |E_{G^i}(K \cap \tilde{D}^i, \tilde{D}^i \setminus K)|, |E_{G^i}(K \setminus \tilde{D}^i, G^i - (\tilde{D}^i \cup K))|,$$

and thus $|\partial_{G^i}(K \cap \tilde{D}^i)|, |\partial_{G^i}(K \setminus \tilde{D}^i)| < |\partial_{G^i} K|$. As K is infinite and as $K \cap \tilde{D}^i$ and $K \setminus \tilde{D}^i$ have only finitely many components (since G is connected), one of these components, say K' , is infinite. Now, K' is a region of G^i with $\partial_{G^i} K' \subseteq \partial_{G^i}(K \cap \tilde{D}^i)$ or $\partial_{G^i} K' \subseteq \partial_{G^i}(K \setminus \tilde{D}^i)$. In both cases, $|\partial_{G^i} K'| < |\partial_{G^i} K| = i + 2$. Because $E(K') \subseteq E(K) \subseteq E(C_n)$ and $i \leq m - 2 < |\partial_G C_n| - 2$, this contradicts (c). We have thus shown (*).

Let us prove that G^m also satisfies (b). For this, consider a region $D \subseteq C_n$ of G with $|\partial_G D| \leq m$, and suppose that $E(D) \cap E(G^m)$ is infinite. Assume D to be chosen among all such regions such that i is maximal with $\partial_G D \subseteq E(G^i)$. Now, if $i < m$, then (*) yields a subgraph D' . By (I), all but finitely many of the edges in $E(D) \cap E(G^m)$ lie in $E(D')$. Since D' has only finitely many components, there is one, C say, such that $E(C) \cap E(G^m)$ is infinite. Since, by (III), $|\partial_G C| \leq |\partial_G D|$, and since, by (IV), $\partial_G C \subseteq E(G^{i+2})$, we obtain a contradiction to the choice of D . Thus, we may assume that $i = m$, ie. $\partial_G D \subseteq E(G^m)$.

Therefore, by performing the according contractions we obtain from D an infinite region \tilde{D} of G^{m-2} such that $\partial_{G^{m-2}} \tilde{D} = \partial_G D$ and $E(\tilde{D}) \subseteq E(C_n)$. Because of (c) for $m - 2$ and because of $|\partial_G D| \leq m$, we get $|\partial_{G^{m-2}} \tilde{D}| = m$. Hence, the region \tilde{D} appears in the enumeration R_1, R_2, \dots used in the construction of G^m , ie. there is a j with $\tilde{D} = R_j$. Since $\partial_{G^{m-2}} R_j \subseteq E(G^m)$ it follows from (5.4) that $E(R_j) \subseteq E(L_j) \subseteq E(L)$. Thus, in G^m all edges of $\tilde{D} = R_j$ are contracted, and consequently, $E(D) \cap E(G^m) = E(R_j) \cap E(G^m) = \emptyset$, a contradiction to $|E(D) \cap E(G^m)| = \infty$.

Having constructed G^m for all $m \leq M := |\partial_G C_n| - 1$, we finally find the region C_{n+1} . Observe that, by (a), $\partial_G C_n$ is a cut of G^M , and that the cut F of G^M that consists of those edges in $E(G^M) \cap E(C_n)$ that in G^M are adjacent to $\partial_G C_n$ has odd cardinality (because by (a), all vertices of G^M are even). Thus, since F is also a cut of G , there exists a region C of G with $\partial_G C \subseteq E(G^M)$ that satisfies (i) and (ii) for $n + 1$.

We claim that

any region C of G that for $n+1$ satisfies, (i), (ii) and $\partial_G C \subseteq E(G^M)$ (5.6) also satisfies (iii).

Indeed, consider a k -region D of G with $C_n \supseteq D \supseteq C$. Then $E(C) \cap E(G^M)$ is infinite, by (i). (For this, contract all edges in $E(G^M) \setminus (E(C) \cup \partial_G C)$, and recall that a finite graph always has an even number of odd vertices.) From (b) for $m = M$ it follows that $k \geq M + 1 = |\partial_G C_n|$, as desired for (iii).

Now, choose C_{n+1} such that $\partial_G C_{n+1}$ has minimal (odd) cardinality among all regions satisfying (i), (ii) and (iii) for $n + 1$.

To see (iv) for $n + 1$, consider a k -region $D \subseteq C_{n+1}$ with $k < |\partial_G C_{n+1}|$. If $k \leq M$, then we can apply (iv) for n , so suppose $k > M$. Furthermore, we may assume that k is odd, as otherwise we can choose $\ell := 1$ and $K_1 := D$. Then D satisfies (i) and (ii) for $n + 1$, and we should have chosen D as C_{n+1} , if not (iii) and thus, by (5.6), also $\partial_G D \subseteq E(G^M)$ fails for D . Repeated use of (*), where we apply (*) in each step to every component of the subgraph D' obtained in the previous step, yields a subgraph D^* of C_n such that each of its components K_1, K_2, \dots, K_ℓ has properties (II)–(IV) for $m = M$. In particular, (II) implies for $1 \leq i \leq \ell$ that $K_i \subseteq C_{n+1}$. Let $\{K_{\ell+1}, \dots, K_L\}$ be the set of all regions that arose as one of the K_i in one of the applications of (*). Then $V(D) \subseteq \bigcup_{i=1}^L V(K_i)$, by (I).

By (III), $|\partial_G K_i| \leq k$ for $i = 1, \dots, \ell$. Now, if there is an $j \in \{1, \dots, \ell\}$ such that $|\partial_G K_j|$ is odd, then $K_j \subseteq C_{n+1}$ satisfies (i), (ii), and, by (IV) and (5.6), also (iii) for $n + 1$, contradicting the choice of C_{n+1} . So, $|\partial_G K_i|$ is even and $\leq k$ for $i = 1, \dots, L$ (for $i > \ell$ this follows from (I) and $k > M$). As $\partial_G C_{n+1} \subseteq E(G^M)$, and thus $\partial_G C_{n+1} \cap \bigcup_{i=\ell+1}^L E(K_i) = \emptyset$, and as $K_1, \dots, K_\ell \subseteq C_{n+1}$, each of the K_1, \dots, K_L either lies completely in C_{n+1} or is disjoint from it. Together with $D \subseteq C_{n+1}$ this implies that $V(D) \subseteq \bigcup_{K \in \mathcal{K}} V(K)$ for $\mathcal{K} := \{K_i : K_i \subseteq C_{n+1} \text{ and } 1 \leq i \leq L\}$, which proves (iv) for $n + 1$. This completes the proof of the theorem.

5.5 Properties of the end degree

As an indication that the end degree indeed behaves as expected of a degree, we extend, in this section, three basic properties of the vertex degree in finite graphs to end degrees in locally finite graphs [23]. At the end of this section, however, we present two examples where end degrees differ from vertex degrees.

The number of odd vertices in a finite graph is always even. We prove the following easy analogon.

Proposition 5.16. [23] *Let G be a locally finite graph. Then the number of odd vertices and ends in G is even or infinite.*

Proof. Suppose that the set \mathcal{O} of odd vertices and ends has odd cardinality. Observe that there is a finite set $S \subseteq V(G)$ that contains all vertices of \mathcal{O} and separates the ends in \mathcal{O} pairwise. By Corollary 5.15, there is for each end $\omega \in \mathcal{O}$ an odd region $A_\omega \subseteq G - S$ with $\omega \in \overline{A_\omega}$. Observe that the A_ω are pairwise disjoint. So, contracting each A_ω to a vertex a_ω we arrive at a graph G' that has an odd number of odd vertices, and in which all end degrees are even. Now, consider two copies of G' and add all edges vv' , where v is an odd vertex of G' and v' its copy. The resulting graph has an odd cut, but no odd vertices or ends, a contradiction to Theorems 5.4 and 1.4. \square

Dirac [36] proved that if a finite graph has minimum vertex degree $k \geq 2$ then it contains a circuit of length $k + 1$. This becomes false for infinite graphs: an easy counterexample is the k -regular infinite tree. But the tree ceases to be a counterexample if a minimum degree is also imposed on the ends, and indeed, then Dirac's result extends to locally finite graphs:

Theorem 5.17. [23] *Let G be a locally finite graph, and let $H \subseteq G$ be a subgraph so that every vertex and every end $x \in \overline{H}$ has degree at least $k \geq 2$ in H . Then there is a circuit $C \subseteq E(H)$ of G with length $\geq k + 1$.*

First, note that the theorem is best possible, even for infinite graphs. Indeed, consider disjoint copies G_1, G_2, \dots of K^{k+1} . Identify a vertex in G_1 with a vertex in G_2 . Then identify a different vertex in G_2 with a vertex of G_3 and so on. In the resulting graph the minimum vertex degree is k , and it is easy to see that the single end has degree k too, but there is no circuit of length greater than $k + 1$. See Figure 5.4 for an example with $k = 2$.




Figure 5.4: Theorem 5.17 is best possible for $k = 2$

Next, let us remark that the long circuit provided by the theorem may be infinite, and indeed the result becomes false if we require finite circuits. To see this, consider a k -regular tree H with root r . Let G be the graph obtained by adding an edge between any two vertices which the same distance to r . Then G has a single end, which has infinite degree in H , but H does not contain any finite circuits.

Before we can prove Theorem 5.17, we need a standard lemma, which can be found in Hall and Spencer [42, p. 208].

Lemma 5.18. *Every topological path with distinct endpoints x, y in a Hausdorff space X contains an arc between x and y .*

Proof of Theorem 5.17. First, observe that if H has an infinite block then H contains two disjoint rays that are equivalent in H (and thus also in G). By linking these by a path in H we obtain a double ray whose edge set is an infinite circuit of G .

Therefore, we may assume that every block of H is finite. Next, suppose that there is a block B of H that contains at most one vertex v with $d_B(v) < k$. Pick a longest path in B . One of the endvertices has at least k neighbours on that path, and hence there is a finite circuit of length $\geq k + 1$ in B .

So, every block B of H is finite and contains at least two vertices of H with degree $< k$ in B , which then are cutvertices of H . Now, replace every block B of H by a tree $T \subseteq B$ whose leaves are exactly the cutvertices of H incident with B . Then every vertex of the resulting forest $H' \subseteq H$ has degree ≥ 2 as every block contains two cutvertices.

Assume that $E(H')$ does not contain infinite circuits, and let v_1, v_2, \dots be an enumeration of $V(H')$. We will inductively construct for $n \in \mathbb{N}$ homeomorphisms $\phi_n : [0, 1] \rightarrow \overline{H'} \subseteq |G|$. Choosing b_0 as any vertex in H' and putting $A_n := \phi_n([0, 1])$, we require that for $n \geq 1$ both $a_n := \phi_n(0)$ and $b_n := \phi_n(1)$ are vertices, and satisfy:

- (i) $a_n = b_{n-1}$ for $n \geq 2$;
- (ii) $A_m \cap A_n = \emptyset$ for $1 \leq m \leq n - 2$ and $A_{n-1} \cap A_n = \{b_{n-1}\}$;
- (iii) there is a cutvertex v incident with two blocks B, B' of H such that $d_B(v) < k$ and such that A_n contains two edges incident with v , one in $E(B)$ and the other in $E(B')$ (let us call any arc with that property *deficient*); and
- (iv) if there is a topological path in $\overline{H'}$ from b_{n-1} to v_n that is edge-disjoint from $B_{n-1} := \bigcup_{i=1}^{n-1} A_i$, then $v_n \in A_n$.

Note that for $n \geq 1$, B_n is a topological path.

In order to construct ϕ_n , assume $\phi_1, \dots, \phi_{n-1}$ to be defined already. First, suppose there is a topological path as required by (iv). By Lemma 5.18, either b_{n-1} and v_n are the endpoints of an arc A that is edge-disjoint from B_{n-1} , or $b_{n-1} = v_n$, in which case we put $A := \{v_n\}$. We claim that $A \cap B_{n-1} = \{b_{n-1}\}$. Indeed, otherwise let v be the vertex with $b_{n-1}v \subseteq A$. Then $A \cup B_{n-1}$ contains a topological path from v to b_{n-1} that avoids all inner points of $b_{n-1}v$, and hence, by Lemma 5.18, also a $b_{n-1}v$ arc A' . Thus, $A' \cup b_{n-1}v \subseteq A \cup B_{n-1} \subseteq \overline{H'}$ is a circle, contradicting our assumption.

We now lengthen A so that it also satisfies (iii). Because every vertex has degree ≥ 2 in H' , and because $\overline{H'}$ does not contain any circles, v_n has a neighbour in $H' \setminus A \cup B_{n-1}$. Continuing in this way, we obtain a v_n - B path in H' that meets $A \cup B_{n-1}$ only in v_n , where B is a block of H which is adjacent to the block that contains v_n . As $B \cap \overline{H'}$ is connected and as $\overline{H'}$ does not contain any circles, B is disjoint from $A \cup B_{n-1}$. So, since B has a cutvertex b with $d_B(b) < k$, there is a deficient path $P \subseteq H'$ that starts in v_n and is otherwise disjoint from $A \cup B_{n-1}$. Thus, we easily find a homeomorphism $\phi_n : [0, 1] \rightarrow A \cup P$ which satisfies (i)–(iv).

So suppose there is no topological path as in (iv). Again we find a deficient path $P \subseteq H'$ starting in b_{n-1} which is disjoint from $B_{n-1} \setminus \{b_{n-1}\}$, and the respective homeomorphism $\phi_n : [0, 1] \rightarrow P$ has properties (i)–(iv).

This process yields a set of arcs A_n , to which we apply Lemma 5.9. We obtain an $x \in |G|$, which is necessarily an end, such that $A^* := \bigcup_{n=1}^{\infty} A_n \cup \{x\}$ is a topological path from b_0 to x .

The end x has degree k in H , and hence there are k edge-disjoint arcs $R_1, \dots, R_k \subseteq \overline{H}$ that start in x . Each of the R_i meets $A^* \setminus \{x\}$ in every neighbourhood of x . Indeed, suppose there is a neighbourhood U of x and an index j such that $R_j \cap U$ is disjoint from $A^* \setminus \{x\}$. Since R_j is continuous, there is a subarc of R_j which starts in x and is completely contained in U . Pick a vertex v_m on this subarc, and denote by R the subarc of R_j between x and v_m . Then $\bigcup_{n=m-1}^{\infty} A_n \cup R$ clearly is a topological path from b_{m-1} to v_m which is edge-disjoint from B_{m-1} , a contradiction to (iv) as $v_m \notin A^* \supseteq A_m$.

Let $\phi : [0, 1] \rightarrow A^*$ be a continuous function with range A^* and $\phi(1) = x$. Choose an $s \in [0, 1)$ such that each of the R_i hits A^* in a $\phi(r_i)$ with $r_i < s$. Because of (iii), we may assume that $v := \phi(s)$ is a cutvertex incident with two blocks B, B' of H such that $d_B(v) < k$ and such that A^* contains two edges incident with v , one in $E(B)$ and the other in $E(B')$. Not all of the k arcs R_i can go through the cut $F := E_H(v, B - v)$ of H , which has cardinality $d_B(v) < k$; so assume R_j does not contain any edge of F . Let uw be the (unique) edge in $E(A) \cap F$, and assume $\phi^{-1}(u) \leq \phi^{-1}(w)$. Then $(A \cup R_j) \setminus uw \cup \{u, w\}$ contains a topological path from w to u (simply run from w to x along A , then from x to $\phi(r_j)$ along R_j and finally from $\phi(r_j)$ to u along A). Therefore, there is also an arc $R \subseteq (A \cup R_j) \setminus uw \cup \{u, w\}$ with endpoints u and w , by Lemma 5.18. Consequently, $R \cup uw \subseteq \overline{H}$ is a circle. Since $E(R)$ is disjoint from F and every B – $(B' - v)$ path in H has to go through F , $|E(R)|$ is infinite. Thus, $E(R \cup uw) \subseteq E(H)$ is an infinite circuit, as desired. \square

In a finite graph the cycles are exactly the connected 2-regular subgraphs. We extend this characterisation to locally finite graphs.

Theorem 5.3. [23] *Let C be a subgraph of a locally finite graph G . Then \overline{C} is a circle if and only if \overline{C} is topologically connected and every vertex or end x of G with $x \in \overline{C}$ has degree two in C .*

We need a lemma:

Theorem 5.19 (Diestel and Kühn [35]). *When a graph G is locally finite, every closed connected subset of $|G|$ is path-connected.*

Proof of Theorem 5.3. If \overline{C} is a circle, then it is clearly topologically connected and every vertex and every end $x \in \overline{C}$ has degree two in C .

For the converse direction, Theorem 5.17 implies that there is a circle $D \subseteq \overline{C}$. Suppose there exists a point $z \in \overline{C} \setminus D$. Theorem 5.19 yields an arc $A \subseteq \overline{C}$ that starts at z and ends in D . As both A and D are closed, A has a first point in D , ie. a point x such that the subarc A' of A between z and x meets D only in x . Thus, there are three edge-disjoint arcs in C with common endpoint x , two in D and the arc A' . So, x is either a vertex or an end and has degree at least 3 in C , a contradiction. Thus, $\overline{C} = D$. \square

Let us now turn to two areas in which end degrees differ in their behaviour from vertex degrees in finite graphs.

For a subgraph H of a graph G , deleting $E(H)$ reduces the degree of a vertex $v \in V(G)$ by its degree in H , ie. $d_G(v) = d_H(v) + d_{G-E(H)}(v)$. Although for an end ω it clearly holds that $d_G(\omega) \geq d_H(\omega) + d_{G-E(H)}(\omega)$, equality is in general not ensured. Consider the $4 \times \infty$ -grid, which has a single end. As depicted in Figure 5.5, the removal of (the edge set of) a ray R leads to a decrease of the end degree from 4 to any of 3, 2, 1 or 0, depending on how R is chosen [23]. Similarly, deleting a circuit can lead to an odd decrease in the degree.




Figure 5.5: Removal of a ray lets the end degree decrease by 1 or more

The second area where the end degree differs in its behaviour concerns extremal results. A classical theorem by Mader [49], for instance, states

that high average degree forces a finite graph to contain a large complete minor. This, however, fails for locally finite graphs even if every end has high degree. Figure 5.6 indicates how for every $k \geq 5$ a planar k -regular




Figure 5.6: High degree in all vertices and in the single end but planar

graph with a single end of infinite degree can be constructed. Being planar, such a graph can never contain even a K^5 as a minor.

In finite graphs high degree guarantees not only the existence of a large clique minor but also that of a highly connected subgraph. In contrast, Stein [58] provides an example showing that high degree in all vertices and ends is not sufficient for a locally finite graph to have a subgraph of high connectivity. On the other hand, if we only require high *edge*-connectivity then we are assured that there is such a subgraph.

5.6 Weakly even ends

Finally, let us briefly discuss an alternative degree concept [23]. It arises from the observation that (ii) of Lemma 5.5 is equivalent to:

- (iii) for every finite $S \subseteq V(G)$ there is a finite set $S' \supseteq S$ of vertices such that the maximal number of edge-disjoint ω -rays starting in S' is even.

Lemma 5.5 (ii) was our main motivation for our definition of an even end. In the same vein, (iii) leads to the following alternative definition of parity, which differs from our original definition only in that the quantifiers are exchanged:

Definition 5.20.[23] *Let H be a subgraph of a locally finite graph G . Call ω weakly even in H if for every finite $S \subseteq V(G)$ there is a finite set $S' \supseteq S$ of vertices such that the maximal number of edge-disjoint ω -arcs in \overline{H} starting in S' is even. Otherwise, ω is strongly odd in H .*

Observe that an even end is weakly even, and that a strongly odd end is odd. For ends of finite degree the two parity concepts are equivalent; this can be seen in a similar way as the equivalence of (ii) and (iii). For ends of infinite degree, however, this need not be true: consider a ray $v_1v_2\dots$, and replace each edge v_iv_{i+1} by i (subdivided) parallel edges. The obtained graph has a single end, which is both odd and weakly even.

This construction only works because there are odd vertices present. But could an odd end exist in a graph that has all vertices even and all ends weakly even? Or, on the contrary:

Problem 5.21. [23] *Does Theorem 5.4 remain true if we substitute “even ends” by “weakly even ends”?*

We have been unable to settle the problem. However, we can answer both this question and Conjecture 5.8 positively for locally finite graphs with only countably many ends:

Proposition 5.22. [23] *Let G be a locally finite graph with only countably many ends, and let H a subgraph. Then $E(H) \in \mathcal{C}(G)$ if and only if every vertex has even degree in H and if every end has weakly even degree in H .*

Proof. The forward direction follows immediately from Theorem 1.4 and Corollary 5.15. For the backward direction, suppose $E(H) \notin \mathcal{C}(G)$, which by Theorem 1.4 means that G has a finite cut F with $|F \cap E(H)|$ odd. Let $\omega_1, \omega_2, \dots$ be an enumeration of $\Omega(G)$. We successively define a sequence $\mathcal{A}_0 \subseteq \mathcal{A}_1 \subseteq \dots$ of finite sets of disjoint regions $A \subseteq G - F$ of G with $|\partial A \cap E(H)|$ even and such that for each ω_i with $i \leq n$ there is an $A \in \mathcal{A}_n$ with $\omega_i \in \overline{A}$. Put $\mathcal{A}_0 := \emptyset$. In order to define the set \mathcal{A}_n first check whether there is an $A \in \mathcal{A}_{n-1}$ such that $\omega_n \in \overline{A}$, in which case we put $\mathcal{A}_n := \mathcal{A}_{n-1}$. Otherwise consider the (finite) set S of all neighbours of each $A \in \mathcal{A}_{n-1}$ and of the endvertices of the edges in F . As ω_n is weakly even, Lemma 5.12 yields a region $B \subseteq G - F$ with $\omega_n \in \overline{B}$ and $A \cap B = \emptyset$ for all $A \in \mathcal{A}_{n-1}$. Put $\mathcal{A}_n := \mathcal{A}_{n-1} \cup \{B\}$. Finally, contracting all the disjoint regions $A \in \bigcup_{n=1}^{\infty} \mathcal{A}_n$ to a vertex each yields a finite graph with all vertex degrees even in H that has a cut F with $|F \cap E(H)|$ odd, a contradiction. \square

Chapter 6

Hamilton cycles and long cycles

6.1 Introduction

While Hamilton cycles have been investigated intensively in finite graphs, less attention has been paid to Hamilton cycles in infinite graphs. One reason for this is that it was not entirely clear what the infinite analogon of a cycle should be.

Adomaitis [1] avoided this question by defining a graph to be hamiltonian if for every finite subset of the vertex set there is a spanning cycle. In contrast, Nash-Williams [52] addressed the problem and proposed spanning double rays as infinite analogons of Hamilton cycles. He noticed that for a spanning double ray to exist the graph needs to be 3-indivisible. (A graph is k -indivisible if the deletion of finitely many vertices leaves at most $k - 1$ infinite components.) Generalising the following classical result by Tutte, Nash-Williams conjectured that a 3-indivisible 4-connected planar graph contains a spanning double ray.

Theorem 6.1 (Tutte [62]). *Every finite 4-connected planar graph has a Hamilton cycle.*

Recently, Yu [66] announced a proof of Nash-Williams' conjecture.

The restriction to 3-indivisible graphs is a quite serious one that at first appears unavoidable. Yet, while double rays are the obvious first choice for an infinite analogon of cycles we have seen in the previous chapters that infinite cycles as we have defined them are far more suited. In particular, our infinite cycles overcome the limitation to 3-indivisible graphs. Indeed, in Figure 6.1 we see a *Hamilton circle*, ie. a spanning circle, in a graph that is not k -indivisible for any k . The example is due to Diestel and Kühn [33]. I conjecture that, in this sense, Theorem 6.1 extends to locally finite graphs; see Diestel [29].




Figure 6.1: A Hamilton circle (drawn bold)

Conjecture 6.2. *Let G be a locally finite 4-connected planar graph. Then G has a Hamilton circle.*

In this chapter, which is mainly based on [24], we shall present a partial result in this direction:

Theorem 6.3. [24] *Let G be a locally finite 6-connected planar graph that is k -indivisible for some finite $k \in \mathbb{N}$. Then G has a Hamilton circle.*

A Hamilton cycle is the longest cycle a finite graph can have. Of interest are also those cycles that while not Hamilton are still, in a certain sense, long. We will turn to this subject in the last section of this chapter and show that the longest cycles an infinite (locally finite) graph can have, namely the infinite cycles, generate the topological cycle space.

In the next section, we will see that Hamilton circles behave in some respect as their finite counterparts. In Section 6.3, we will show that Conjecture 6.2 fails for 3-connected infinite graphs, and in Section 6.4 we shall prove Theorem 6.3.

6.2 Hamilton circuits and covers

Let us briefly give an indication that our concept of a Hamilton circle, i.e. a circle that contains every vertex (and then also every end) of a graph, behaves in similar ways as its finite analogon.

In Section 4.7, we showed that every locally finite graph that is planar can be covered by two elements of the cycle space. If instead of being planar the

graph has a *Hamilton circuit*, i.e. a circuit that is the edge set of a Hamilton circle, then we can still find such a cover. This is, of course, well-known for finite graphs.

Proposition 6.4. [13] *Let G be a locally finite graph with a Hamilton circuit. Then there are $Z_1, Z_2 \in \mathcal{C}(G)$ with $E(G) = Z_1 \cup Z_2$.*

Proof. Let H be the Hamilton circuit and fix a vertex u of G . We say that an edge set C has *distance* n to u if the shortest path that starts in u and ends in a vertex that is incident with C has length n . Since every edge $e \notin H$ is a chord of H there are two circuits $\subseteq H \cup \{e\}$ containing e . Denote by C_e the one which has greater distance to u (or choose arbitrarily if both have the same distance). We claim that

$$\{C_e : e \notin H\} \text{ is a thin set.} \tag{6.1}$$

If (6.1) is true, we are done. Putting $Z_1 := H$ and $Z_2 := \sum_{e \notin H} C_e$, we see that $E(G) = Z_1 \cup Z_2$.

So suppose (6.1) to be false. Then there is an N such that for infinitely many edges $e \notin H$ the circuit C_e is incident with $S_N := \{v \in V(G) : v \text{ has at most distance } N \text{ to } u\}$. As $|G|$ is compact there is a sequence e_1, e_2, \dots among these edges converging against an end ω . Let A be the unique topological component of $\overline{C}(S_N, \omega) \cap \overline{H}$ that contains ω . Then A is an arc. All but finitely many of the e_n have both their endvertices in A . Indeed, $\overline{C}(S_N, \omega) \cap \overline{H}$ has only finitely many topological components, and any of these components which is incident with infinitely many of the e_n must contain ω . Since ω is contained in exactly one topological component there is, thus, a $e_n \in E(C(S_N, \omega))$ with both its endvertices in A , and we easily find a circuit in $E(A) \cup \{e_n\} \subseteq H \cup \{e_n\}$ containing e_n . This circuit then avoids S_N , a contradiction to the choice of C_{e_n} . \square

6.3 An infinite Herschel-type graph

The Herschel graph (see Figure 6.2) is a well-known example for a 3-connected planar graph without a Hamilton cycle, which shows that in a sense Theorem 6.1 is best possible. In this section, let us briefly demonstrate that Conjecture 6.2 is also false for infinite graphs if we only assume 3-connectivity [24].

Consider Figure 6.3. There we have arranged copies of the Herschel graph (greyed) in a hexagonal grid. The copies are glued together in such a way that each copy of the vertex v in Figure 6.2 does not receive an extra edge and thus still has degree 3. Now assume that the resulting graph has a Hamilton




Figure 6.2: The Herschel graph

circle C , and consider a copy of the Herschel graph H in that graph. If the Hamilton circle enters H in u and leaves H in either x or y then C induces a Hamilton circle of H , which is impossible. Thus, C enters H in x and leaves it in y , which implies that H has a Hamilton path, i.e. a spanning path, between x and y . However, this is impossible since H is an odd bipartite graph, where x and y are in the same partition class.




Figure 6.3: An infinite 3-connected planar graph without a Hamilton circle

6.4 Proof of Theorem 6.3

Before we start proving Theorem 6.3, let us reformulate the theorem slightly. The notion of ends is central to the definition of infinite cycles, and we will therefore express the theorem in terms of ends. It is straightforward to see that a locally finite graph is k -indivisible if and only if it has at most $k - 1$

ends. Thus, we obtain the following alternative version of Theorem 6.3:

Theorem 6.5. [24] *Let G be a locally finite 6-connected planar graph with at most finitely many ends. Then G has a Hamilton circle.*

In the context of Hamilton cycles two concepts have been quite successful, namely that of a bridge and a Tutte path; see for instance Thomassen [61]. Let us define these. For a subgraph H of a (finite) graph G , an H -bridge is either a chord $e \notin E(H)$ together with its endvertices both of which lie in $V(H)$ or a component K of $G - H$ together with all edges between K and H , denoted by $E(K, H)$, and their incident vertices. Note that this definition of a bridge differs from the one given in Chapter 3. However, we will only apply it to finite graphs, in which case the two definitions coincide. We say that an H -bridge B is *chordal* if $E(B)$ consists of a single edge. All the vertices of an H -bridge B in H are *attachments* of B . For a subgraph F of G , we call a path (resp. cycle) P a F -Tutte path (resp. cycle) in G if every P -bridge of G has at most three attachments, and if every bridge containing an edge of F has at most two attachments.

Our main tool in the proof will be the following result of Thomassen, which itself implies Tutte's theorem:

Theorem 6.6 (Thomassen [61]). *Let G be a finite 2-connected plane graph with a face boundary C . Assume that $u \in V(C)$, $e \in E(C)$ and $v \in V(G) \setminus \{u\}$. Then G contains a C -Tutte path P from u to v and through e .*

Roughly, our strategy to prove Theorem 6.5 is as follows: we apply Theorem 6.6 in a finite subgraph H (plus some extra vertices) of our graph G and then extend the resulting Tutte cycle to (finite parts of) the components $G - H$. For this to work, we need that all these components, which will be infinite, are 3-connected. This is the task of our next lemma, and, in particular, of its consequence, Lemma 6.9.

Lemma 6.7. [24] *Let G be a 4-connected locally finite planar graph with minimum degree at least 6, let $U \subseteq V(G)$ be a finite vertex set, and let ω be an end of G . Then, there is a finite vertex set S such that for $C_\omega := C(S, \omega)$ holds*

- (i) C_ω is disjoint from U ; and
- (ii) if $X \subseteq V(C_\omega)$ with $|X| \leq 2$, then every component of $C_\omega - X$ is infinite.

Proof. Choose a finite vertex set S such that $C_\omega := C(S, \omega)$ is disjoint from U , and such that $|E(C_\omega, G - C_\omega)|$ is minimal with that property.

Suppose there is a set $X \subseteq V(C_\omega)$ with $|X| \leq 2$ such that there is a finite component K of $C_\omega - X$. Denote by K' the subgraph obtained by adding the vertices in X and the edges between X and K to K . Put $r := |E(K, G - C_\omega)|$, and $s := |E(X, K)|$. We shall show that $s < r$. Then $S' := S \cup V(K)$ leads to a smaller cut between $C(S', \omega)$ and the rest of the graph, a contradiction.

First, assume that K' is not a triangulation. Thus, putting $n := |V(K)|$, Euler's formula implies $|E(K')| < 3(n + |X|) - 6 \leq 3(n + 2) - 6 = 3n$. On the other hand,

$$2|E(K')| = \sum_{v \in V(K')} d_{K'}(v) = |E(X, K)| + \sum_{v \in V(K)} d_K(v) \geq s + 6n - r.$$

Hence, $3n + (s - r)/2 \leq |E(K')| < 3n$ and thus $s < r$, as desired.

Second, let K' be a triangulation. We may assume that some vertices in S lie in the outer face of K' . Since G is 4-connected and since the face boundary of the outer face of K' is a triangle, T , say, no vertices of G can be contained in the interior face of T . Thus, $T = K'$, and consequently there are exactly two edges between X and K , ie. $s = 2$. Since the minimum degree is at least 6 in G , and as the at most two vertices in K can have at most one edge between them, it follows that $r \geq 3$. Again, we get $s < r$, as desired. \square

Lemma 6.8. [24] *Let G be a k -connected graph, and let $V(G) = A \cup B$ be a partition such that $G[B]$ is l -connected. Consider $X \subseteq A$ with $|X| \leq k - l$. Then, for every component K of $G[A] - X$ the graph $G[K \cup B]$ is still l -connected.*

Proof. Put $B' := V(K) \cup B$, and suppose there is a separator Y of $G[B']$ with $|Y| < l$, and consider two distinct components C and D of $G[B'] - Y$. If both C and D contain vertices of B then $Y \cap B$ is a separator of $G[B]$, contradicting that $G[B]$ is l -connected. So we may assume that $C \subseteq K$. Then, $X \cup Y$ separates C from B in G , but $|X \cup Y| < (k - l) + l = k$, a contradiction. \square

Lemma 6.9 will be used in each of the induction steps of the proof of Theorem 6.5.

Lemma 6.9. [24] *Let G be a locally finite planar 6-connected graph, let U be a finite vertex set and ω an end of G . Let $S \subseteq V(G)$ be a finite vertex set such that $C := C(S, \omega)$ is 2-connected. Then there is a finite vertex set $T \supseteq U$ such that $C(T, \omega) \subseteq C$ is 3-connected and $C - C(T, \omega)$ 2-connected.*

Proof. Add to the union of $U \cap V(C)$ and the neighbours of S in C finitely many vertices such that for the resulting set $U' \subseteq V(C)$ the subgraph $G[U']$ is 2-connected. Then Lemma 6.7 applied to $U' \cup S$ yields a finite vertex set T' so that $D := C(T', \omega) \subseteq C$ is 3-connected and disjoint from U' (and thus from U). Since $G[U'] \subseteq C$ is 2-connected there is a block B of $C - D$ containing U' . Observe that, as G is 6-connected, every component of $C - D - B$ has a neighbour in the finite set $S \cup T'$. Thus there are only finitely many of them. Consider one such component, K say, and note that, as K is disjoint from U' all its neighbours except possibly one, which lies in B , are contained in D . Thus, by Lemma 6.8, $G[D \cup K]$ is still 3-connected. Therefore, the neighbours of $C - B$ together with U give rise to a set T as desired. \square

We will construct our Hamilton circle in a piecewise manner. Slightly more precise, we shall construct finite nested subgraphs G_i in which the application of Theorem 6.6 will yield subgraphs C_i that are finite approximations of the desired infinite Hamilton circle. The following definition and lemma make sure that these subgraphs C_i indeed tend to a cycle, ie. that $\bigcup_{i=1}^{\infty} C_i$ is a cycle.

Call a sequence $(G_1, C_1), \dots, (G_k, C_k)$ of finite induced subgraphs $G_i \subseteq G$ and subgraphs $C_i \subseteq G$ *good*, if for $i = 1, \dots, k$ holds

- (i) $C_{i-1} \subseteq C_i$ and $G_{i-1} \cup N(G_{i-1}) \subseteq G_i$ for $i \geq 2$;
- (ii) if K is an infinite component of $G - G_i$ then $|E(K, G - K) \cap E(C_i)| = 2$;
and
- (iii) there is a cycle Z such that $Z \cap G_i = C_i \cap G_i$.

We say that a finite vertex set S *separates* an end ω from a vertex set U if every ray $R \in \omega$ that starts in a vertex of U meets S . In a similar manner, S separates two ends ω and ω' if every double ray with one tail in ω and the other in ω' goes through S .

Lemma 6.10. [24] *Let G be a locally finite graph with at most finitely many ends, and let every finite initial segment of $(G_1, C_1), (G_2, C_2), \dots$ be good. Then the closure of $C := \bigcup_{i=1}^{\infty} C_i$ is a circle of $|G|$.*

Proof. From (iii) and (i) it follows that C is either a finite cycle or a disjoint union of double rays. Assume the latter, and let ω be an end of one of the double rays. Because of (i) there is an n such that the component K of $G - G_n$ to which ω belongs is separated from all the other ends by G_n . Condition (ii) implies that $C \cap K$ consists of exactly two rays in ω , which

shows that \overline{C} is the disjoint union of circles. It is easy to see that (iii) implies that \overline{C} is topologically connected. Therefore, \overline{C} is a circle. \square

Let us introduce one useful definition before we finally start with the proof. We call a vertex set U in a graph G *externally k -connected*, if $|U| \geq k$ and if for every set $X \subseteq V(G)$ with $|X| < k$ there is a path between any two vertices of $U \setminus X$ in $G - X$.

Proof of Theorem 6.5. By Theorem 3.17, $|G|$ can be embedded in the sphere. We shall identify $|G|$ with that embedding, and thus view G as a plane graph. Note that, by Corollary 3.19, all face boundaries are circles.

Inductively, we shall construct connected finite induced subgraphs $G_1 \subseteq G_2 \subseteq \dots \subseteq G$ and subgraphs $H_1 \subseteq H_2 \subseteq \dots \subseteq G$ such that for every $i \geq 1$ it holds that

- (iv) $(G_1, H_1), \dots, (G_i, H_i)$ is good;
- (v) $V(G_i) \subseteq V(H_i)$;
- (vi) if \mathcal{C}_i is the set of components of $G - G_i$ then each $C \in \mathcal{C}_i$ is 3-connected and belongs to exactly one end; and
- (vii) for each $C \in \mathcal{C}_i$ it follows that $|V(C) \cap V(H_i)| = 2$.

Assume this can be achieved. Then, Lemma 6.10 shows that $H := \overline{\bigcup_{i=1}^{\infty} H_i}$ is a circle. In addition, H contains every vertex of G , by (v) and (i). Therefore, H is a Hamilton circle.

As the construction of the base case, i.e. when $i = 1$, is quite similar to the general case, i.e. when $i \geq 2$, we will treat both at once. However, for some of the steps we shall need to make case distinctions. Put $G_0 := \emptyset$, $H_0 := \emptyset$ and $\mathcal{C}_0 = \{G\}$, and assume (G_{i-1}, H_{i-1}) to be constructed for some $i \geq 1$. Consider a component $C \in \mathcal{C}_{i-1}$.

First, let $i = 1$, and note that as $C = G$ has only finitely many ends, we can choose a finite vertex set W such that none of the infinite components of $G - W$ belong to more than one end. Denote by F_C a face boundary of $|G|$, and pick an edge $x_C y_C \in E(F_C)$, a vertex $u_C \in V(F_C)$ with $u_C \notin \{x_C, y_C\}$, and a vertex $v_C \notin \{u_C, x_C, y_C\}$ that is adjacent to u_C . Put $U_C := W \cup \{u_C, v_C, x_C, y_C\}$.

For $i \geq 2$, observe that, by (vii), H_{i-1} contains exactly two vertices of C , u_C and v_C say. In the by $|G|$ induced embedding of \overline{C} all the neighbours of the connected graph G_{i-1} that lie in C are incident with the same face boundary F_C , say. Pick an edge $x_C y_C \in E(F_C)$ such that $x_C y_C \neq u_C v_C$, and

let U_C be the union of $\{u_C, v_C, x_C, y_C\}$ together with all neighbours of G_{i-1} in C . Thus, $U_C \subseteq V(F_C)$ separates C from the rest of G .

In any case, we obtain

$$u_C \neq v_C \text{ and } |\{u_C, v_C, x_C, y_C\}| \geq 3 \text{ (resp. } = 4 \text{ for } i = 1). \quad (6.2)$$

Since C is 3-connected (resp. 4-connected for $i = 1$) there are three (resp. four) internally disjoint paths in C between any two vertices in U_C . Denote by T_C the vertex set of the union of three (resp. four) such paths for each pair of vertices in U_C . Observe that U_C is externally 3-connected (resp. externally 4-connected for $i = 1$) in $C[T_C]$.

Let $\omega_1, \dots, \omega_n$ be an enumeration of the ends belonging to C . Put $C_1 := C$, and apply Lemma 6.9 to C_1, ω_1, T_C in order to find a finite vertex set S_1 for which $C(S_1, \omega_1) \subseteq C_1$ is 3-connected and disjoint from T_C and for which $C_2 := C_1 - C(S_1, \omega_1)$ is 2-connected. Continuing in this manner, we see that

$$G'_C := C - \bigcup_{i=1}^n C(S_i, \omega_i) \text{ is a finite 2-connected graph, and } U_C \text{ is} \quad (6.3) \\ \text{externally 3-connected (resp. 4-connected for } i = 1) \text{ in } G'_C.$$

Since $U_C \subseteq V(G'_C)$ contains all neighbours of G_{i-1} in C it follows that

$$G'_i := G[G_{i-1} \cup \bigcup_{B \in \mathcal{C}_{i-1}} G'_B] \text{ is connected, and every component of} \quad (6.4) \\ G - G'_i \text{ is 3-connected and belongs to exactly one end of } G.$$

G'_i will serve as a precursor to G_i .

Next, let \mathcal{D}_C be the components D of $G - G'_i$ with $D \subseteq C$, and consider $D \in \mathcal{D}_C$. (Note that for $i > 1$ it holds that $\mathcal{D} = \{D\}$). Let M_D be obtained from $E(D, G - D)$ by deleting for each vertex in D all but one of the incident edges in $E(D, G - D)$. Thus,

$$\text{every neighbour of } G'_C \text{ in } D \text{ is incident with exactly one edge of } M_D. \quad (6.5)$$

Starting from G'_C define a 2-connected finite plane graph \tilde{G}_C as follows: for every $D \in \mathcal{D}_C$ put a vertex z_D into the face of G'_i that contains D and link z_D to the vertices in G'_C that are incident with an edge in M_D . We shall identify these linking edges with the edges in M_D . Note that the resulting graph, which is a minor of G'_C , may have parallel edges. Then, since, for $i = 1$, F_C is a face boundary of $|G|$ such that $u_C, x_C, y_C \in V(F_C)$ and since, for $i \geq 2$, $U_C \subseteq V(F_C)$ this also holds for a face boundary \tilde{F}_C of \tilde{G}_C , i.e.

$$\text{for } i = 1 : u_C, x_C, y_C \in V(\tilde{F}_C) \text{ and } U_C \text{ separates the ends of } G \quad (6.6) \\ \text{for } i \geq 2 : U_C \subseteq V(\tilde{F}_C) \text{ and } U_C \text{ separates } C \text{ from } G - C$$

We apply Theorem 6.6 to \tilde{G}_C , and obtain a \tilde{F}_C -Tutte path (resp. circle for $i = 1$) \tilde{H}_C from u_C to v_C and through $x_C y_C$. (More precisely, if \tilde{G}_C has

parallel edges we first subdivide these before using Theorem 6.6; the obtained Tutte-path then induces one in \tilde{G}_C .) From (6.2) it follows that

$$|V(\tilde{H}_C) \cap U_C| \geq 3 \text{ (resp. } \geq 4 \text{ for } i = 1). \quad (6.7)$$




Figure 6.4: Illustration of the proof of Theorem 6.5

Next, we show that

for every nonchordal \tilde{H}_C -bridge K in \tilde{G}_C , $(K - \tilde{H}_C) \cap G$ is disjoint from U_C , and all its neighbours in G lie in $G'_C \cup D$ for some $D \in \mathcal{D}_C$. (6.8)

Suppose that $K - \tilde{H}_C$ meets U_C . If $i \geq 2$ then K contains an edge of \tilde{F}_C , as $U_C \subseteq V(\tilde{F}_C)$, by (6.6). Thus, K has at most two (resp. three for $i = 1$) attachments as \tilde{H}_C is a \tilde{F}_C -Tutte path (resp. cycle). Since U_C is externally 3-connected (resp. externally 4-connected) in G'_C by (6.3), this implies $U_C \subseteq V(K)$. Thus, by (6.7), $K - \tilde{H}_C$ contains a vertex in $V(\tilde{H}_C) \cap U_C$, a contradiction. Therefore, $K - \tilde{H}_C$ is disjoint from U_C . The second assertion follows from the first together with (6.6) and the fact that $|\mathcal{D}_C| = 1$ if $i \geq 2$.

It holds that

$$z_D \in V(\tilde{H}_C) \text{ for every } D \in \mathcal{D}_C. \quad (6.9)$$

Indeed, suppose there is a $D \in \mathcal{D}_C$ with $z_D \notin V(\tilde{H}_C)$. Thus, $z_D \in V(K - \tilde{H}_C)$ for some \tilde{H}_C -bridge K . Denote by f_D the face boundary of G'_i that contains D . Then there exists a vertex $v \notin V(K)$ in the face boundary of f_D . Indeed, otherwise the at most three attachments of K separate D from $U_C \setminus V(K) \neq \emptyset$ in G , contradicting that G is 6-connected. Note that v is not incident with any edge in M_D as $z_D \in V(K - \tilde{H}_C)$. Let $f \subseteq f_D$ be the (unique) face of $G'_i \cup D \cup M_D$ whose face boundary F contains v . Because G'_i and D are connected there are exactly two vertices in $V(F \cap D)$, a and b say, that are incident with an edge in M_D .

We claim that the at most three attachments of K together with a, b separate $D - \{a, b\}$ from v in G , which then contradicts that G is 6-connected,

thus establishing (6.9). So suppose there is a path from D to v that avoids the attachments of K and a, b , and let xy be its last edge in $E(D, G - D) = E(D, G'_i)$. Assume that $x \in V(D)$ and $y \in V(G'_i)$. If xy is disjoint from f then y is easily seen to be separated in G'_i from v by the attachments of K , which is impossible. Therefore, the interior of xy lies in f , and thus x in its boundary F . By (6.5), x is incident with an edge in M_D . Since a, b are the only vertices in $V(F \cap D)$ incident with an edge in M_D , we obtain with $x \in \{a, b\}$ a contradiction. This establishes (6.9).

Consider $E(\tilde{H}_C)$ as a subset of $E(G)$, and let H_C be the subgraph of G consisting of the edges $E(\tilde{H}_C)$ and the incident vertices. Put $H_i := H_{i-1} \cup \bigcup_{C \in \mathcal{C}_{i-1}} H_C$, and observe that the pair (G'_i, H_i) already satisfies almost all of the desired properties. In particular, it satisfies (i)–(iii) of the definition of a good sequence: (i) holds because of $U_C \subseteq V(G'_C)$ for every $C \in \mathcal{C}_{i-1}$ (for $i \geq 2$), and (ii) because of (6.9). To see (iii), add a path through D between the two vertices of H_i in D for every $D \in \mathcal{D}_C$, and denote the resulting subgraph by Z . Clearly, Z is connected, and every $w \in V(Z)$ has degree two: if $w \in V(G_{i-1})$ then because of (iii) for $i - 1$, if $w \in \{u_C, v_C\}$ then because $u_C \neq v_C$, if $w \in V(H_i - H_{i-1})$ because H_i is a disjoint union of paths for every C , and finally if $w \in V(Z - H_i)$ because $Z - H_i$ is a disjoint union of paths. Thus, Z is a cycle. Moreover, of (iv)–(vii) only (v) is not satisfied: (vi) is (6.4), and (vii) holds because of (6.9) and the definition of M_D .

To fix (v), consider a nonchordal H_i -bridge K in G'_i . By (v) for $i - 1$, we deduce that $K - H_i$ is disjoint from G_{i-1} . Observe that for each $C \in \mathcal{C}_{i-1}$, $K \cap G'_C$ is either empty, a chord or a union of H_C -bridges. Thus, $K - H_i$ is disjoint from U_C , by (6.8), which implies with Lemma 6.8 that $G'_i - (K - H_i)$ still satisfies (i) and (vi), and then also (iv) and (vii). Thus, putting $G_i := G'_i - (G'_i - H_i)$ we see that for the pair (G_i, H_i) conditions (iv)–(vii) hold. \square

6.5 Infinite circuits generate the cycle space

A triangle is certainly a *short* cycle, while a Hamilton cycle should clearly be considered as *long*. Let us turn in this section to cycles that although possibly far shorter than Hamilton cycles are still considerably longer than triangles [21].

In Locke [46], it is conjectured that there is a constant $0 < m < 1$ so that the cycle space of any finite graph in which every two vertices can be joined by a path of length at least k is generated by the cycles of length $\geq mk$. To motivate that conjecture Locke shows that if between any two vertices there is a path of length $\geq (k - 1)^2 + 1$ then the circuits of length $\geq k$ generate the cycle space.

We prove the following theorem, which is a natural extension of Locke's observation to locally finite graphs. We say that an arc has *infinite length* if it contains infinitely many edges.

Theorem 6.11.[21] *Let G be a locally finite graph in which every two vertices are linked by an arc of infinite length. Then $\mathcal{C}(G)$ is generated by the circuits of infinite length.*

Note that every arc linking two vertices of the same block lies entirely in that block. As, in addition, the topological cycle space of a locally finite graph is the direct sum of the topological cycle spaces of its blocks, it is sufficient to prove Theorem 6.11 for 2-connected graphs. Therefore, Theorem 6.11 is an immediate corollary of the following theorem. Call a circuit which is the edge set of a double ray, a *double ray-circuit*.

Theorem 6.12.[21] *Let G be an infinite 2-connected locally finite graph. Then there is a thin set of double ray-circuits that generates $\mathcal{C}(G)$.*

Diestel and Kühn [33] show that there is always a thin generating set consisting of finite circuits. Thus, Theorem 6.12 may be understood as a converse of that fact. A spanning tree T is *end-faithful* if for an arbitrary root r , T has exactly one ray starting in r in every end of G .

Theorem 6.13 (Diestel and Kühn [33]). *Let G be a locally finite connected graph, and let T be any spanning tree of G . Then every element of $\mathcal{C}(G)$ is generated by fundamental circuits of T if and only if T is end-faithful. Also, if T is end-faithful, then the set of fundamental circuits of T is a thin set.*

Since every countable graph has a normal spanning tree (Jung [43]), and since every normal spanning tree is end-faithful, there is always a thin set of finite circuits that generates the topological cycle space in a locally finite graph.

Lemma 6.14.[21] *Let G be an infinite 2-connected locally finite graph. Then there is for every edge vw of G a double ray D containing vw such that $E(D)$ is a double ray-circuit.*

Proof. Choose an end ω of G . Put $S_1 := \{v, w\}$ and for $i = 2, 3, \dots$ choose S_i with minimal cardinality subject to $S_i \cup C(S_i, \omega) \subseteq C(S_{i-1}, \omega)$. (Thus, S_i separates S_{i-1} from the end ω .) Then S_i is a minimal S_i - S_{i+1} separator (which for $i = 1$ is ensured by the 2-connectivity of G and for $i = 2, 3, \dots$ by the minimal choice of the S_i). Thus for each i Menger's Theorem yields $|S_i|$ disjoint S_i - S_{i+1} paths, the union of which contains two disjoint rays in ω one starting in v and the other in w . Using these we obtain the desired double ray D . \square

Proof of Theorem 6.12. Let e_1, e_2, \dots be an enumeration of the edge set of G . By Lemma 6.14, there is for each $e \in E(G)$ a double ray-circuit D_e containing e ; let us assume D_e to be chosen such that the minimal edge index in D_e , i.e. $\min\{i : e_i \in D_e\}$, is maximal. We claim that

$$\text{the set } \mathcal{D} := \{E(D_e) : e \in E(G)\} \text{ is thin.} \quad (6.10)$$

Assume (6.10) to be true. By Theorem 6.13 there is a thin set \mathcal{C} of finite circuits that generates $\mathcal{C}(G)$. For each $C \in \mathcal{C}$ pick an edge $e^C \in C$. Since C and D_{e^C} share at least e^C we easily find two double ray-circuits $D_1^C, D_2^C \subseteq C \cup D_{e^C}$ with $D_1^C + D_2^C = C$. As both \mathcal{C} and \mathcal{D} are thin, the set $\{D_1^C, D_2^C : C \in \mathcal{C}\}$ is a thin set that generates the topological cycle space of G (since \mathcal{C} is a generating set).

So, suppose that (6.10) does not hold. Then, there is an edge e_n which is met by infinitely many of the D_e . Putting $E_n := \{e_1, \dots, e_n\}$, we note that $G - E_n$ has a component K that contains an infinite number of edges e with $e_n \in D_e$. We claim that

$$\text{the number of blocks of } K \text{ is finite.} \quad (6.11)$$

Assume that K has an infinite number of blocks. Denote by S the finite set of vertices in K that are incident with an edge in E_n . Since G is 2-connected, there is, by Menger's theorem, for every block B of K an S - S path P_B in K containing an edge of B . Then there are two vertices in S that are the endvertices of infinitely many P_B . We fix B as one of these blocks, and then choose B' as another which is disjoint from the finite path P_B . This yields two blocks B, B' of K such that P_B and $P_{B'}$ have the same endvertices and $P_B \cap E(B') = \emptyset$.

Let x be the last vertex in $P_{B'}$ that also lies in P_B before $P_{B'}$ enters B' , and let y be the first vertex in $P_{B'}$ that also lies in P_B after $P_{B'}$ has left B' . Then, $E(xP_{B'}y \cup xP_By)$ is a circuit of K that meets both $E(B')$ and $E(K) \setminus E(B')$, and thus contains a cutvertex of B' . This yields a contradiction, as a cutvertex may not lie on a circuit.

Having proved (6.11), we find an infinite block B^∞ of K that contains an edge e with $e_n \in D_e$. Lemma 6.14 yields a double ray-circuit containing e in B^∞ , which, avoiding E_n , contradicts the choice of D_e as D_e meets E_n in e_n . \square

We remark that, in a similar way, Theorem 6.11 can be extended to the space $\mathcal{C}(\tilde{G})$ for graphs G satisfying (1.3). As in the case of Gallai's theorem in Chapter 2 or of dual graphs in Chapter 4, it is here, too, necessary to work in \tilde{G} rather than in $|G|$. Indeed, consider a ray R each of whose vertices

are joined to a (dominating) vertex v . Clearly, there is for any two vertices an arc of infinite length linking them. However, in $|G|$ there are no infinite circuits at all, and thus Theorem 6.11 fails if $\mathcal{C}(G)$ is used. In contrast, it is straightforward to see that the infinite circuits in \tilde{G} generate $\mathcal{C}(\tilde{G})$.

Finally, let us observe that in a non-locally finite graph also the finite circuits may not be sufficient to generate the topological cycle space (although this is still true in any countable graph). Let v_1, v_2, \dots be some distinguished vertices, and let there be a double ray $D^r = \dots w_{-1}^r w_0^r w_1^r \dots$ for every $r \in \mathbb{R}$. Join v_n to w_{-n}^r and to w_n^r for all n and r . Then, the resulting graph G has a single end, and the edge set of each D^r is a circuit. Suppose $Z := \sum_{r \in \mathbb{R}} E(D^r) = \bigcup_{r \in \mathbb{R}} E(D^r) \in \mathcal{C}(G)$ is the sum of finite circuits. Since Z is uncountable the sum contains uncountably many distinct finite circuits, each of which is incident with one of the v_n . Thus, there is a v_n which is incident with infinitely many of the summands, contradicting the definition of a thin sum.

Chapter 7

Linear algebra with thin sums

7.1 Introduction

The topological cycle space differs in two points from the traditional definition, $\mathcal{C}_{\text{fin}}(G)$, of a cycle space in infinite graphs, namely in that it allows infinite cycles and infinite sums. In this chapter, we will focus on that second point and study the ramifications of allowing infinite sums.

The effects of admitting infinite sums are most prominently seen in generating sets. For instance, Theorem 6.13 asserts that in a 2-connected infinite locally finite graph G , the space $\mathcal{C}(G)$ is spanned by the fundamental circuits of an end-faithful spanning tree. Thus, this is an example for a space of uncountable cardinality which contains infinite objects that is nevertheless generated by only countably many finite objects.

Generating sets of the cycle space have been of interest in previous chapters: Tutte's generating theorem (Theorem 1.3), MacLane's planarity criterion (Theorem 3.3), as well as Theorem 6.11 state that a certain set of circuits is sufficient in order to generate the whole space. A standard approach in linear algebra is to look at minimal generating sets, i.e. bases. We have successfully employed a similar approach in Chapter 3 when we defined 2-bases. There, in Chapter 3, it was not clear from the outset whether such a 2-basis always existed in a simple generating set, and we used some ad-hoc methods to prove their existence. We will pursue the question of the existence of bases in this chapter in more generality. The second problem we will be interested in, is, under which circumstances a family of edge sets is closed under taking infinite sums; we know that this is the case for the topological cycle space.

As the context of graphs neither simplifies nor complicates these problems we drop it and instead consider the problems in terms of abstract set systems.

7.2 Bases

For the purpose of this chapter we call a family \mathcal{T} of subsets of a set M a *thin* family if each $m \in M$ lies in at most finitely many of the members of \mathcal{T} . The sum $\sum_{T \in \mathcal{T}} T$ of a thin family is the set of all elements $m \in M$ that appear in exactly an odd number of members of \mathcal{T} . This definition is consistent with the one for locally finite graphs that we have used throughout this thesis; our ground set M was usually the edge set of a graph, and \mathcal{T} a family of circuits. For a family \mathcal{N} of subsets of M , we denote by $\langle \mathcal{N} \rangle$ the set of sums of all thin subfamilies of \mathcal{N} . Calling a minimal subfamily \mathcal{B} of \mathcal{N} for which $\langle \mathcal{B} \rangle = \langle \mathcal{N} \rangle$ a *basis* of $\langle \mathcal{N} \rangle$, we shall prove the following theorem:

Theorem 7.1.[20] *Let M be a countable set, and let \mathcal{N} be a family of subsets of M . Then \mathcal{N} contains a basis of $\langle \mathcal{N} \rangle$.*

In linear algebra the analogous theorem is usually proved with Zorn's lemma as follows. Starting from a chain $(\mathcal{B}_\lambda)_\lambda$ of linear independent subfamilies of \mathcal{N} it is observed that $\bigcup_\lambda \mathcal{B}_\lambda$ is still linear independent since any violation of linear independence is witnessed by finitely many elements, and these would already lie in one of the \mathcal{B}_λ . Thus, each chain has an upper bound, which implies, by Zorn's lemma, that there is a maximal linear independent set, a basis. This approach fails in our context. Indeed, 'dependence' need not be witnessed by only finitely many elements, that is, we might observe that while no nonempty finite sum in $\bigcup_\lambda \mathcal{B}_\lambda$ adds up to \emptyset some infinite sum does. Thus, we cannot get the contradiction that already one of the \mathcal{B}_λ was not 'independent'.

Before proving the theorem let us give an alternative definition of a basis, which is also standard in linear algebra. We call a subfamily \mathcal{L} of \mathcal{N} a *representation of N in \mathcal{N}* if $N = \sum_{L \in \mathcal{L}} L$. Then it is easy to see that a subfamily \mathcal{B} of \mathcal{N} is basis of \mathcal{N} if and only if $\langle \mathcal{B} \rangle = \langle \mathcal{N} \rangle$ and if \emptyset has a unique representation in \mathcal{B} (namely the empty set).

Proof of Theorem 7.1. If M is a finite set the result is immediate, so we assume M to be infinite. By deleting members that appear more than once, we may assume \mathcal{N} to be a set. Let m_1, m_2, \dots be an enumeration of M , and define \mathcal{N}_i to be the set of those elements $N \in \mathcal{N} \setminus \bigcup_{j < i} \mathcal{N}_j$ for which $m_i \in N$. Clearly, $\{\mathcal{N}_i : i \in \mathbb{N}\}$ is a partition of \mathcal{N} .

For every $i \in \mathbb{N}$, let $N_{i1}, N_{i2}, \dots, N_{i\lambda}, \dots$ be a (possibly transfinite) enumeration of \mathcal{N}_i , and perform the following transfinite recursion. Start by setting $\mathcal{N}_{i0} = \mathcal{N}_i$, and then for every ordinal $\lambda > 0$ define sets $\mathcal{N}_{i\lambda} \subseteq \mathcal{N}_i$ and

$\mathcal{X}_{i\lambda}^i$ as follows: If there is a

$$\mathcal{X}_{i\lambda}^i \subseteq \left(\bigcap_{\mu < \lambda} \mathcal{N}_{i\mu} \cap \{N_{i\mu} : \mu < \lambda\} \right) \cup \bigcup_{k=i+1}^{\infty} \mathcal{N}_k$$

with $\sum_{X \in \mathcal{X}_{i\lambda}^i} X = N_{i\lambda}$ then set $\mathcal{N}_{i\lambda} := \bigcap_{\mu < \lambda} \mathcal{N}_{i\mu} \setminus \{N_{i\lambda}\}$. Otherwise, set $\mathcal{N}_{i\lambda} := \bigcap_{\mu < \lambda} \mathcal{N}_{i\mu}$ and $\mathcal{X}_{i\lambda}^i := \emptyset$. Having defined all $\mathcal{N}_{i\lambda}$, we put $\mathcal{B}_i := \bigcap_{\lambda} \mathcal{N}_{i\lambda}$.

We claim that \emptyset has a unique representation in $\mathcal{B} := \bigcup_i \mathcal{B}_i$. Indeed, suppose there is a nonempty thin set $\mathcal{C} \subseteq \mathcal{B}$ whose elements sum up to \emptyset . Let $i \in \mathbb{N}$ be minimal so that \mathcal{C} contains an element of \mathcal{B}_i , and observe that as all the elements in \mathcal{B}_i share m_i there is a maximal ordinal λ such that $N_{i\lambda} \in \mathcal{C}$. Then $N_{i\lambda} = \sum_{C \in \mathcal{C} \setminus \{N_{i\lambda}\}} C$ and

$$\mathcal{C} \setminus \{N_{i\lambda}\} \subseteq \left(\bigcap_{\mu < \lambda} \mathcal{N}_{i\mu} \cap \{N_{i\mu} : \mu < \lambda\} \right) \cup \bigcup_{k=i+1}^{\infty} \mathcal{N}_k,$$

a contradiction to that $N_{i\lambda} \in \mathcal{B}_i \subseteq \mathcal{N}_{i\lambda}$.

Next, consider a $N_{i\lambda} \notin \mathcal{B}_i$ for some i, λ . We will show that it has a representation in \mathcal{B} . Inductively, define for integers $k > i$ a thin set $\mathcal{X}_{i\lambda}^k$ as follows. Let \mathcal{E} be the set of ordinals η for which $N_{k\eta} \in \mathcal{X}_{i\lambda}^{k-1}$ but $N_{k\eta} \notin \mathcal{B}_k$. Since $\mathcal{X}_{i\lambda}^{k-1}$ is thin and since all these $N_{k\eta} \in \mathcal{N}_k$ share m_k it follows that \mathcal{E} is a finite set, and so

$$\mathcal{X}_{i\lambda}^k := (\mathcal{X}_{i\lambda}^{k-1} \setminus \{N_{k\eta} : \eta \in \mathcal{E}\}) \Delta_{\eta \in \mathcal{E}} \mathcal{X}_{k\eta}^k.$$

is thin, where Δ denotes the symmetric difference. By induction we easily obtain that

$$\mathcal{X}_{i\lambda}^k \text{ is a thin set and } N_{i\lambda} = \sum_{X \in \mathcal{X}_{i\lambda}^k} X. \quad (7.1)$$

Indeed

$$N_{i\lambda} = \sum_{X \in \mathcal{X}_{i\lambda}^{k-1}} X = \sum_{X \in \mathcal{X}_{i\lambda}^{k-1} \setminus \{N_{k\eta} : \eta \in \mathcal{E}\}} X + \sum_{\eta \in \mathcal{E}} \sum_{X \in \mathcal{X}_{k\eta}^k} X = \sum_{X \in \mathcal{X}_{i\lambda}^k} X.$$

By construction, and as $\mathcal{X}_{i\lambda}^{k+1}$ is disjoint from $\bigcup_{s=1}^k \mathcal{N}_s$ it follows that

$$\mathcal{X}_{i\lambda}^k \cap \bigcup_{l=1}^k \mathcal{N}_l = \mathcal{X}_{i\lambda}^{k+1} \cap \bigcup_{l=1}^k \mathcal{N}_l \subseteq \bigcup_{l=i}^k \mathcal{B}_l. \quad (7.2)$$

Putting $\mathcal{X}_{i\lambda}^* := \bigcup_{k=i}^{\infty} (\mathcal{X}_{i\lambda}^k \cap \mathcal{B}_k)$, we obtain from (7.2) that

$$\mathcal{X}_{i\lambda}^* \subseteq \bigcup_{l=i}^{\infty} \mathcal{B}_l. \quad (7.3)$$

We claim that

$$\mathcal{X}_{i\lambda}^* \text{ is a thin set and } N_{i\lambda} = \sum_{X \in \mathcal{X}_{i\lambda}^*} X. \quad (7.4)$$

Indeed, consider an element $m_s \in M$. Since no element in \mathcal{N}_t for $t > s$ contains m_s , it holds that of the elements in $\mathcal{X}_{i\lambda}^*$ only those in $\bigcup_{t=1}^s \mathcal{N}_t$ may contain m_s . Since $\mathcal{X}_{i\lambda}^* \cap \bigcup_{t=1}^s \mathcal{N}_t = \mathcal{X}_{i\lambda}^s \cap \bigcup_{t=1}^s \mathcal{N}_t$ by (7.2) and since $\mathcal{X}_{i\lambda}^s$ is thin there may be only finitely many elements in $\mathcal{X}_{i\lambda}^*$ containing m_s . Consequently, $\mathcal{X}_{i\lambda}^*$ is thin, as claimed. Furthermore, $\mathcal{X}_{i\lambda}^* \cap \bigcup_{t=1}^s \mathcal{N}_t = \mathcal{X}_{i\lambda}^s \cap \bigcup_{t=1}^s \mathcal{N}_t$ implies that $m_s \in \sum_{X \in \mathcal{X}_{i\lambda}^*} X$ if and only if $m_s \in \sum_{X \in \mathcal{X}_{i\lambda}^s} X$. Together with (7.1) this establishes the claim.

Finally, consider a $N \in \langle \mathcal{N} \rangle$, and let $\mathcal{L} \subseteq \mathcal{N}$ be a representation of N in \mathcal{N} . We will show that N has a representation in \mathcal{B} . Let I and J_i for $i \in I$ be the sets of indices such that $\mathcal{L} \setminus \mathcal{B} = \{N_{i\lambda} : i \in I, \lambda \in J_i\}$, and define the family $\mathcal{X}' := (X)_{X \in \mathcal{X}_{i\lambda}^*, i \in I, \lambda \in J_i}$. Then, \mathcal{X}' is a thin family. Indeed, suppose there is a m_s that lies in infinitely many members of \mathcal{X}' . The fact that for $t > s$ no element in \mathcal{N}_t contains m_s together with (7.3) implies that there are infinitely many triples (X, λ, i) with $m_s \in X \in \mathcal{X}_{i\lambda}^*$, $\lambda \in J_i$ and $i \in \{1, \dots, s\}$. Since \mathcal{L} is thin and since for fixed i all the $N_{i\lambda}$ contain m_i it follows that J_i is a finite set. Thus, there are $i' \in I$ and $\lambda' \in J_{i'}$ such that there are infinitely many $X \in \mathcal{X}_{i'\lambda'}^*$ with $m_s \in X$, contradicting that, by (7.4), $\mathcal{X}_{i'\lambda'}^*$ is thin. Therefore, \mathcal{X}' is thin and we get with (7.4)

$$\begin{aligned} N &= \sum_{L \in \mathcal{L} \cap \mathcal{B}} L + \sum_{L \in \mathcal{L} \setminus \mathcal{B}} L = \sum_{L \in \mathcal{L} \cap \mathcal{B}} L + \sum_{i \in I} \sum_{\lambda \in J_i} N_{i\lambda} \\ &= \sum_{L \in \mathcal{L} \cap \mathcal{B}} L + \sum_{i \in I} \sum_{\lambda \in J_i} \sum_{X \in \mathcal{X}_{i\lambda}^*} X = \sum_{L \in \mathcal{L} \cap \mathcal{B}} L + \sum_{X \in \mathcal{X}'} X. \end{aligned}$$

As $\mathcal{X}' \subseteq \mathcal{B}$, by (7.3), we have thus found a representation of N in \mathcal{B} . Therefore, $\langle \mathcal{B} \rangle = \langle \mathcal{N} \rangle$. \square

We have formulated Theorem 7.1 only for countable sets M . The following example shows that this is indeed best possible.

Example 7.2. [20] *If M is the vertex set of the bipartite graph K_{\aleph_0, \aleph_1} , and if \mathcal{N} is the set of all pairs of adjacent vertices then \mathcal{N} contains no basis of $\langle \mathcal{N} \rangle$.*

Proof. Treating the edges as pairs of vertices we may view \mathcal{N} as the edges of K_{\aleph_0, \aleph_1} . First note, that each $N \subseteq M$ of countable cardinality is contained in $\langle \mathcal{N} \rangle$. Indeed, let n_1, n_2, \dots be a (possibly finite) enumeration of N , and choose for each n_i in the enumeration a ray R_i that starts in n_i but avoids

the first $i - 1$ vertices of each of the rays R_1, \dots, R_{i-1} . Then, the family $(e)_{e \in E(R_i), i \in \mathbb{N}}$ is thin and its sum equals to N since $\sum_{e \in E(R_i)} e = n_i$.

Suppose that $\mathcal{B} \subseteq \mathcal{N}$ is a basis of $\langle \mathcal{N} \rangle$. Since each of the uncountably many vertices of K_{\aleph_0, \aleph_1} must lie in one of the edges in \mathcal{B} , there is one vertex v with (uncountably) infinite degree in $H := (M, \mathcal{B})$. Next, note that H cannot contain any cycles or double rays because the sum of the corresponding edges would be \emptyset , contradicting that \emptyset has a unique representation in \mathcal{B} . Thus, for the component C of H with which v is incident, infinitely many of the components of $C - v$ are rayless trees. Pick a leaf not adjacent with v from countable infinitely many of these, and denote their set by N . Since N is countable there is a thin set $\mathcal{B}_N \subseteq \mathcal{B}$ whose sum equals to N . Consider a leaf $l \in N$, let C_l be the component of $C - v$ that is incident with l , and let $u_l v$ be the edge in \mathcal{B} between v and C_l . Then, \mathcal{B} induces a graph in C_l that has even degree in all vertices apart from l and possibly u_l . Since l has odd degree, u_l is odd too, which implies that $u_l v \in \mathcal{B}_N$. Thus, v lies in infinitely many $u_l v$, which contradicts that \mathcal{B}_N is a thin set. \square

The example is in stark contrast to linear algebra, where we always find a basis. Another point in which our bases differ is that two bases do not need to have the same cardinality. Indeed, putting $M := \{m_0, m_1, \dots\}$ we see that $\mathcal{B} := \{\{m_i\} : i \geq 0\}$ is a countable basis of the power set $\mathcal{P}(M)$ of M . On the other hand, $\mathcal{N} := \{\{m_0\} \cup N : N \subseteq M\}$, for which $\langle \mathcal{N} \rangle = \mathcal{P}(M)$, contains, by Theorem 7.1, a basis \mathcal{B}' . Since all thin subsets of \mathcal{N} are finite, \mathcal{B}' needs to be uncountable to generate the uncountable set $\mathcal{P}(M)$.

7.3 Closedness

In linear algebra, closedness is not an issue. Clearly, if only finite sums are allowed then the span of some set is always closed under taking (finite) sums. For graphs G it follows from Theorem 1.5 that $\mathcal{C}(G)$, which is the span of all circuits, is closed under taking thin sums. In general, however, this need not be the case. For instance, consider the set M consisting of the distinct elements m_0, m_1, \dots and set $\mathcal{N} := \{\{m_0, m_i\} : i \geq 0\}$. Clearly, $\langle \mathcal{N} \rangle$ is countable. On the other hand, the set $\mathcal{L} := \{\{m_i\} : i \geq 0\}$ for which $\langle \mathcal{L} \rangle = \mathcal{P}(M)$ is contained in $\langle \mathcal{N} \rangle$. Thus, $\langle \mathcal{N} \rangle$ cannot be closed under taking thin sums since $\mathcal{P}(M)$ is uncountable.

Proposition 7.3. [20] *Let M be a countable set and let \mathcal{T} be a thin family of subsets of M . Then $\langle \mathcal{T} \rangle$ is closed under taking (thin) sums.*

Proof. By Theorem 7.1, we may assume that \emptyset has a unique representation in \mathcal{T} , and by deleting duplicate members we may view \mathcal{T} as a set. Consider

a thin family $(C_i)_{i \in I}$ of elements $C_i \in \langle \mathcal{T} \rangle$. Then for each C_i there is a $\mathcal{T}_i \subseteq \mathcal{T}$ with $C_i = \sum_{T \in \mathcal{T}_i} T$. Now, if $\mathcal{S} := (T)_{i \in I, T \in \mathcal{T}_i}$ is a thin family, then $\sum_{i \in I} C_i = \sum_{S \in \mathcal{S}} S$, and $\langle \mathcal{T} \rangle$ is closed. Therefore, we suppose that there is a $m \in M$ that lies in infinitely many members of \mathcal{S} . Since \mathcal{S} consists of elements of the thin set \mathcal{T} it follows that there is a $T \in \mathcal{T}$ and an infinite index set $I_0 \subseteq I$ for which $T \in \mathcal{T}_i$ for $i \in I_0$. Put $\mathcal{C}_0 := \{T\}$, and $K_0 := T$. Let m_1, m_2, \dots be an enumeration of M . Inductively, we will define infinite sets $I_0 \supseteq I_1 \supseteq \dots$, finite sets $\mathcal{C}_0 \subseteq \mathcal{C}_1 \subseteq \dots \subseteq \mathcal{T}$, and sets $K_i \subseteq M$ such that

- (i) $\mathcal{C}_i \subseteq \mathcal{T}_j$ for all $j \in I_i$; and
- (ii) if $k \in \mathbb{N}$ is the smallest index for which $m_k \in K_{i-1} := \sum_{C \in \mathcal{C}_{i-1}} C$ then $m_k \notin K_i$.

Assume I_0, \dots, I_{i-1} and $\mathcal{C}_0, \dots, \mathcal{C}_{i-1}$ to be defined. Let $k \in \mathbb{N}$ be the smallest index such that $m_k \in K_{i-1}$. As $(C_j)_{j \in I}$ is thin, there is an infinite set $I'_i \subseteq I_{i-1}$ such that $m_k \notin C_j = \sum_{T \in \mathcal{T}_j} T$ for $j \in I'_i$. Thus, there is for each $j \in I'_i$ a $C \in \mathcal{T}_j \setminus \mathcal{C}_{i-1}$ with $m_k \in C$. Since, m_k is contained in only finitely many different $C \in \mathcal{T}$, there is an infinite set $I_i \subseteq I'_i$ and a $D \in \mathcal{T}$ such that $m_k \in D \in \mathcal{T}_j \setminus \mathcal{C}_{i-1}$ for all $j \in I_i$. Setting $\mathcal{C}_i := \mathcal{C}_{i-1} \cup \{D\}$ we see that both (i) and (ii) are satisfied.

We distinguish two cases, both of which will lead to contradiction. First, suppose there is a $k \in \mathbb{N}$ and an infinite set $N \subseteq \mathbb{N}$ so that $m_k \in K_n$ for each $n \in N$. Indeed, assume k to be chosen minimal. Then there is a $n' \in \mathbb{N}$ with $m_l \notin K_n$ for all $n \geq n'$ and $1 \leq l < k$. For each $n \in N$, we see by (ii) that the number of members of \mathcal{C}_n that contain m_k is increased since k is also the smallest index with respect to (ii). Since N is infinite this implies that there are infinitely many members of \mathcal{T} that contain m_k , contradicting that \mathcal{T} is thin.

Next, suppose there is for each k a $n_k \in \mathbb{N}$ such that $m_k \notin K_n$ for all $n \geq n_k$. Putting $\mathcal{C}_* := \bigcup_{i=1}^{\infty} \mathcal{C}_i$, we consequently get that $\sum_{D \in \mathcal{C}_*} D = \emptyset$, which, as $\mathcal{C}_* \supseteq \mathcal{C}_0 = \{T\} \neq \emptyset$, contradicts that \emptyset has a unique representation in \mathcal{T} . \square

For a locally finite graph G , Proposition 7.3 provides an alternative way to show that $\mathcal{C}(G)$ is closed. Indeed, we only need to apply Proposition 7.3 to the set of fundamental circuits of an end-faithful spanning tree, which is, by Theorem 6.13, a thin generating set of $\mathcal{C}(G)$.

In a similar way, we may conclude from Proposition 7.3 that the cuts are closed under thin sums. However, this is more easily done using the simple Lemma 4.11.

Chapter 8

Menger's theorem in infinite graphs with ends

8.1 Introduction

The topological space $|G|$, for a graph G , is essential in defining the topological cycle space $\mathcal{C}(G)$. In this last chapter, we will encounter another problem that is intimately related to the space $|G|$.

Erdős conjectured (see Nash-Williams [51]) that Menger's theorem should extend to infinite graphs as follows:

Erdős-Menger Conjecture. *For every graph $G = (V, E)$ and any two sets $A, B \subseteq V$ there is a set \mathcal{P} of disjoint A - B paths in G and an A - B separator X consisting of a choice of one vertex from each of the paths in \mathcal{P} .*

There are several partial results; see Aharoni [4]. In particular, Aharoni [3] proved the conjecture for countable graphs; we have used this result already in Chapter 4. In Aharoni and Diestel [6], this was extended as follows:

Theorem 8.1 (Aharoni and Diestel [6]). *The Erdős-Menger conjecture holds for all graphs $G = (V, E)$ and sets $A, B \subseteq V$ that are separated in G by a countable set of vertices.*

In particular, the conjecture holds whenever A is countable, regardless of the cardinality of G .

Another approach was pursued in Diestel [30]. Again, an assumption is made that A and B are easy to separate. But this time, the notion of separation used is topological:

Theorem 8.2 (Diestel [30]). *The Erdős-Menger conjecture holds for all graphs $G = (V, E)$ and sets $A, B \subseteq V$ whose closures in the topological space $|G|$ are disjoint.*

Although Theorem 8.2 refers implicitly to the ends of G by its closure condition, the conclusion is the original one from Erdős's conjecture, which makes no reference to ends. However, there is also a natural extension of the conjecture that does refer to ends. Here, the sets A and B may contain ends as well as vertices. The A - B paths in \mathcal{P} can be either finite paths linking two vertices, or rays linking a vertex to an end, or double rays linking two ends. Similarly, the separator X may contain ends (that lie in A or B), thus blocking any ray belonging (= converging) to that end. This extension was proposed in Diestel [26], and found to be true for countable graphs under certain necessary restrictions for A and B .

In this chapter, which is based on [19], we extend both Theorems 8.1 and 8.2 to ends, in the spirit of Diestel [26]. Our extension of Theorem 8.2 will build on our extension of Theorem 8.1.

We remark that recently Aharoni and Berger [5] have announced a proof of the Erdős-Menger conjecture. If their proof holds up to scrutiny then our results imply that the Erdős-Menger conjecture is true in the end version too:

Theorem 8.3. [19] *Let G be a graph and let $A, B \subseteq V(G) \cup \Omega(G)$ be such that $A \cap \overline{B} = \emptyset = \overline{A} \cap B$, the closures being taken in $|G|$. Then G satisfies the Erdős-Menger conjecture for A and B .*

(For a formal definition of when a graph satisfies the Erdős-Menger conjecture if A, B are allowed to contain ends see next section.)

We note that the condition $A \cap \overline{B} = \emptyset = \overline{A} \cap B$ is really necessary; see Diestel [26]. The condition means that any ray whose end lies in A can be separated from B by a finite set of vertices, and vice versa with A and B interchanged. Note that this does not imply the much stronger condition that A and B can be finitely separated, in which case the proof is immediate by standard alternating path techniques; see Diestel [26]. A more typical example for the disjointness condition is to take as A and B distinct levels of vertices in a tree: if the tree is \aleph_0 -regular, for example, it contains infinitely many disjoint paths between these levels, so A and B have disjoint closures (in fact, are closed and disjoint) but cannot be finitely separated.

8.2 Definitions and statement of results

Fix a graph $G = (V, E)$. Recall that the set of ends of G is denoted by $\Omega = \Omega(G)$, and we write $G = (V, E, \Omega)$ to refer to G together with its set of ends.

Following [19], we provide some more required terminology. In contrast to previous chapters, we define paths here in a more general way: *Paths* in G can be finite paths (which contain at least one vertex), rays, double rays, or singleton sets $\{\omega\}$, where ω is an end of G . The closure of an infinite path P in $|G|$ contains one or two ends of G . We will often consider such an end as the first or last point of P , and when we say that two paths are *disjoint* then these points too shall be distinct. (The first and last point of a path $P = \{\omega\}$, of course, is ω .) For $A, B \subseteq V \cup \Omega$, a path is an A - B *path* if its first but no other point lies in A and its last but no other point lies in B .

The union of a ray R and infinitely many disjoint paths starting on R but otherwise disjoint from R is a *comb* with *spine* R . The last points (vertices or ends) of those paths are the *teeth* of the comb. We will frequently use the following simple lemma:

Lemma 8.4. *In the graph $G = (V, E, \Omega)$ let R be a ray of an end ω , and let $X \subseteq V \cup \Omega$ such that $\omega \notin X$. Then $\omega \in \overline{X}$ if and only if G contains a comb with spine R and teeth in X . \square*

A set $X \subseteq V \cup \Omega$ is an A - B *separator* in a subspace $T \subseteq |G|$ if every path P in T with its first point in A and its last point in B satisfies $\overline{P} \cap X \neq \emptyset$. (We express this informally by saying that “ P meets X ”, though strictly speaking we shall mean \overline{P} rather than just P .) We say that a set $Y \subseteq V \cup \Omega$ lies *on* a set \mathcal{P} of disjoint A - B paths if Y consists of a choice of exactly one vertex or end from every path in \mathcal{P} . We say that G *satisfies the Erdős-Menger conjecture for A and B* , or that *the Erdős-Menger conjecture holds for G, A, B* , if $|G|$ contains a set \mathcal{P} of disjoint A - B paths and an A - B separator on \mathcal{P} . (Thus, officially, we always refer to the ends version of the conjecture. But this is compatible with the traditional terminology: if neither A nor B contains an end then neither can any A - B path, so the conjecture with ends automatically defaults to the original conjecture in this case.)

We can now state the two main results of this chapter. First, our extension of Theorem 8.1:

Theorem 8.5. [17] *Let $G = (V, E, \Omega)$ be a graph, let $A, B \subseteq V \cup \Omega$ satisfy $A \cap \overline{B} = \emptyset = \overline{A} \cap B$, and suppose that there exists a countable A - B separator $X \subseteq V \cup \Omega$ in G . Then G satisfies the Erdős-Menger conjecture for A and B .*

In particular, the ends version of the Erdős-Menger conjecture is true whenever A is countable.

Our second main result, whose proof builds on Theorem 8.5, extends Theorem 8.2 fully to its natural topological setting:

Theorem 8.6. [17] *Every graph $G = (V, E, \Omega)$ satisfies the Erdős-Menger conjecture for all sets $A, B \subseteq V \cup \Omega$ that have disjoint closures in $|G|$.*

Let us complete this section with an outline of the proofs to come, and of how this chapter is organised.

The proof of Theorem 8.5 will occupy us for the next two sections. It runs roughly as follows. Most of the proof – all of Section 8.3 – will be spent on transferring our problem to an equivalent problem in which G is replaced with a suitable minor and A and B consist of vertices only. In this new situation, the countable A – B separator X – which likewise may be assumed to consist of vertices only – divides G into two parts: one between X and A (including both) and the rest (which includes B). We now apply Theorem 8.1 to obtain an A – X separator Y on a system of disjoint A – X paths in the first part. Note that Y is again countable, and it separates A from B in G . Repeating the same procedure for the part of G between Y and B yields a system of Y – B paths with a separator Z on it, which again separates A from B . These paths can be concatenated with the A – Y segments of the first, to give a system of disjoint A – B paths (with the A – B separator Z on it). It remains to transfer this solution back to the original sets A, B containing ends, in the original graph G .

In Section 8.5 we prove Theorem 8.6. Employing techniques developed in Section 8.3 and in Diestel [30], we will eliminate all ends in the closures of A and B . Then the remaining ends can be discarded as well. In this way, the problem is reduced to a rayless graph, for which the Erdős-Menger conjecture is known to hold:

Theorem 8.7 (Aharoni [2], Polat [53]). *The Erdős-Menger conjecture holds for rayless graphs.*

8.3 The reduction lemma

In this section we develop further some techniques from Diestel [26] designed to reduce the ends versions of the Erdős-Menger conjecture to the related vertex versions. Observe that in the finite Menger theorem we can ignore all the vertices in $A \cap B$ and work with the graph $G - (A \cap B)$ instead. In an infinite graph, however, we have to take care that no end in $A \cup B$ is destroyed or split when the vertices of $A \cap B$ are deleted from G .

Lemma 8.8.[19] *Let $G = (V, E, \Omega)$ be a graph, and let $A, B \subseteq V \cup \Omega$ satisfy*

$$A \cap (\overline{B} \setminus B) = \emptyset = (\overline{A} \setminus A) \cap B.$$

Then for the graph $G' := G - (A \cap B \cap V)$ there are sets $A', B' \subseteq V(G') \cup \Omega(G')$ satisfying

- (i) $|A'| \leq |A|$;
- (ii) if $A \subseteq V$ then $A' \subseteq A$, and if $B \subseteq V$ then $B' \subseteq B$;
- (iii) $A' \cap \overline{B'} = \emptyset = \overline{A'} \cap B'$;
- (iv) if G' satisfies the Erdős-Menger conjecture for A' and B' , then G satisfies it for A and B .

Proof. Put $A' := A \setminus B$ and $B' := B \setminus A$, both of which are subsets of $|G|$. Consider a ray R of an end α in A' or B' , say in A' . Then R has a tail in G' . Indeed, if not then there are vertices of $A \cap B \cap V \subseteq B$ in every neighbourhood of $\alpha \in A \setminus B$. Consequently, $\alpha \in A \cap (\overline{B} \setminus B)$, which is a contradiction. Similarly, two rays R_1, R_2 in G' of which R_1 is a ray of an end $\omega \in A' \cup B'$ are equivalent in G' if and only if they are equivalent in G . Indeed, if R_1 and R_2 are equivalent in G then there is a ray $R_3 \in \omega$ that meets both of R_1 and R_2 infinitely often. Now R_3 has a tail in G' , showing that R_1 and R_2 are also equivalent in G' .

Thus, mapping every end of G in $A' \cup B'$ to the unique end of G' that contains tails of its rays defines a bijection between the ends in $A' \cup B'$ and certain ends in G' . Using this bijection (and a slight abuse of notation) we may view A' and B' also as subsets of $V(G') \cup \Omega(G')$. Clearly, $A' \cap B'$ is empty and hence (iii) is satisfied. Also, (i) and (ii) are trivial.

For (iv), let X' be an A' - B' separator on a set of disjoint A' - B' paths \mathcal{P}' in G' . Adding to \mathcal{P}' the trivial paths $\{x\}$ for all $x \in A \cap B$ yields a set \mathcal{P} of disjoint A - B paths on the A - B separator $X := X' \cup (A \cap B)$. \square

Later on, in Lemma 8.12, we shall need a family of disjoint subgraphs of G (with certain properties) such that every end of A lies in the closure of one of these subgraphs. Such a family cannot always be found. But our next lemma finds instead a family of subgraphs such that the ends of A not contained in their closures form a set I that can be ignored: those ends will automatically be separated from B by any $(A \setminus I)$ - B separator on a set of disjoint A - B paths.

Lemma 8.9. [19] *Let $G = (V, E, \Omega)$ be a graph, and let $A, B \subseteq V \cup \Omega$ be such that $A \cap \overline{B} = \emptyset = \overline{A} \cap B$. Then for every set $A_\Omega \subseteq A \cap \Omega$ there exist a set $I \subseteq A_\Omega$, an ordinal μ^* , and families $(G_\mu)_{\mu < \mu^*}$ and $(S_\mu)_{\mu < \mu^*}$ such that, for every $\mu < \mu^*$, the graph $G_\mu - S_\mu$ is a component of $G - S_\mu$ with S_μ as its finite set of neighbours, and*

- (i) $\overline{G_\mu - S_\mu} \cap B = \emptyset$;
- (ii) if $G_\mu \neq \emptyset$ then $\overline{G_\mu} \cap A_\Omega \neq \emptyset$;

(iii) $V(G_\nu \cap G_\mu) \subseteq S_\nu \cap S_\mu$ for all $\nu < \mu$.

Moreover,

(iv) for every end $\alpha \in A_\Omega \setminus I$ there is a $\mu < \mu^*$ with $\alpha \in \overline{G_\mu}$;

(v) every $(A \setminus I)$ - B separator on a set of disjoint $(A \setminus I)$ - B paths is also an A - B separator.

Proof. We construct the families $(G_\mu)_{\mu < \mu^*}$ and $(S_\mu)_{\mu < \mu^*}$ and a transfinite sequence $I_0 \subseteq I_1 \subseteq \dots \subseteq A_\Omega$ recursively. The sets I_μ ($\mu < \mu^*$) will serve as precursors to I . To simplify notation, we write $C_\mu := G_\mu - S_\mu$ for every μ . For the construction, we will in addition to (i)–(iii) require for every μ that

(vi) $I_\mu \cap \overline{G_\nu} = \emptyset$ for all $\nu \leq \mu$.

We start by setting $I_0, G_0, S_0 := \emptyset$. Consider the least ordinal $\mu > 0$ such that the above sets are already defined for all $\lambda < \mu$. If μ is a limit, we set

$$I_\mu := \bigcup_{\lambda < \mu} I_\lambda$$

and $G_\mu, S_\mu := \emptyset$. This choice clearly satisfies (i)–(iii) and (vi).

Suppose now that μ is a successor, $\mu = \lambda + 1$ say. If every end in $A_\Omega \setminus I_\lambda$ lies in some $\overline{G_\nu}$ with $\nu < \mu$, we set $\mu^* := \mu$ and terminate the recursion. So suppose there is an end $\alpha \in A_\Omega \setminus I_\lambda$ that lies in no earlier $\overline{G_\nu}$. Then, if possible, choose a finite vertex set S such that $C(S, \alpha)$ avoids all G_ν with $\nu < \mu$.

Such a choice of S is impossible if and only if

$$\text{for every finite } S \subseteq V \text{ there is a } \nu < \mu \text{ with } C(S, \alpha) \cap G_\nu \neq \emptyset. \quad (8.1)$$

In this case we choose to ignore α , i.e. set $I_\mu := I_\lambda \cup \{\alpha\}$ and $G_\mu, S_\mu := \emptyset$. Again the requirements (i)–(iii) are clearly met, while (vi) holds by the choice of α .

Now suppose we can find S as desired. As $A \cap \overline{B} = \emptyset$, we can also find a basic open neighbourhood $\hat{C}(S', \alpha)$ of α in $|G|$ that is disjoint from B . We now define S_μ as the set of neighbours of $C(S \cup S', \alpha)$ and $G_\mu := G[S_\mu \cup C(S_\mu, \alpha)]$. Then (i) holds since $S_\mu \supseteq S'$, while (ii) holds as $\alpha \in \overline{G_\mu}$. To see (iii), first note that

$$G_\nu \cap C_\mu = \emptyset \text{ for all } \nu < \mu$$

by the choice of S . So, all we have to show is that $G_\nu \cap S_\mu \subseteq S_\nu$. Consider a vertex $v \in G_\nu \cap S_\mu$. Since S_μ is the set of neighbours of C_μ , there is a

vertex $w \in C_\mu$ adjacent to v . As noted above, $w \notin G_\nu$. So v is a vertex in $G_\nu = C_\nu \cup N(C_\nu)$ with a neighbour outside G_ν , implying $v \notin C_\nu$ and hence $v \in S_\nu$, as desired.

Let us finally set $I_\mu := I_\lambda$ and verify (vi). We only need to show that $I_\mu \cap \overline{G_\mu} = \emptyset$. Suppose that intersection contains an end α' . Let $\mu' < \mu$ be minimal such that $\alpha' \in I_{\mu'}$. Then (8.1) should have been satisfied for μ' and α' , but fails with $S := S_{\mu'}$ as $C(S_{\mu'}, \alpha') = C_\mu$, a contradiction.

Having defined I_μ , G_μ and S_μ for all $\mu < \mu^*$ so that (i)–(iii) and (vi) are satisfied, we put

$$I := \bigcup_{\mu < \mu^*} I_\mu.$$

Together with the definition of μ^* this implies (iv). Observe that from (vi) we obtain $I \cap \overline{G_\mu} = \emptyset$ for all $\mu < \mu^*$.

To establish (v) let \mathcal{P} be a system of disjoint $(A \setminus I)$ – B paths and X an $(A \setminus I)$ – B separator on \mathcal{P} . Now suppose that X is not an A – B separator in $|G|$, i.e. there is a path Q from A to B that avoids X . By turning Q into a path \tilde{Q} from $A \setminus I$ to B that avoids X , we will obtain a contradiction.

We may assume that Q starts at an end $\alpha \in I$. Let μ be the step at which α was added to I , i.e. let μ be minimal with $\alpha \in I_\mu$. Choose a finite vertex set S such that $\overline{C}(S, \alpha)$ is disjoint from B (this is possible, as $A \cap \overline{B} = \emptyset$). Then any path of \mathcal{P} that meets $C(S, \alpha)$ must pass through S . Hence only finitely many paths of \mathcal{P} can meet $C(S, \alpha)$, and so $X_\alpha := X \cap \overline{C}(S, \alpha)$ is also finite. Conditions (iii) and (iv) ensure that every end in X_α lies in exactly one $\overline{C_\lambda}$; let $\{\lambda_1, \dots, \lambda_m\}$ be the set of these λ . Then for

$$S' := S \cup (X_\alpha \cap V) \cup \bigcup_{i=0}^m S_{\lambda_i}$$

we have

$$\overline{C}(S', \alpha) \cap X = \emptyset.$$

Now, all we need is a point of $A \setminus I$ that lies in $\overline{C}(S', \alpha)$ (and thus can be used to change Q into the desired path). Indeed, if there is an ordinal $\lambda < \mu$ such that $G_\lambda \neq \emptyset$ and

$$C_\lambda \subseteq C(S', \alpha), \tag{8.2}$$

we can complete the proof as follows. By (ii) for λ there will be an end $\alpha' \in A$ in $\overline{C_\lambda} \subseteq \overline{C}(S', \alpha)$. Since $I \cap \overline{G_\lambda} = \emptyset$, we have $\alpha' \in A \setminus I$. Take an α' – Q path P in $\overline{C}(S', \alpha)$ with last vertex x , say. Then P avoids X , and hence so does the path $\tilde{Q} := PxQ$. Thus, \tilde{Q} is as desired.

So suppose there is no ordinal $\lambda < \mu$ satisfying (8.2). Then for all $\lambda < \mu$ we have either $C_\lambda \cap C(S', \alpha) = \emptyset$ or $C_\lambda \cap S' \neq \emptyset$. As all the C_λ are

disjoint by (iii), only finitely many of them meet S' ; let $\lambda_{m+1}, \dots, \lambda_n$ be the corresponding ordinals. Then

$$S'' := S' \cup \bigcup_{i=m+1}^n S_{\lambda_i}$$

satisfies $C(S'', \alpha) \cap C_\lambda = \emptyset$ for all $\lambda < \mu$.

However, $G_\lambda \cap C(S'', \alpha)$ cannot be empty for all $\lambda < \mu$, as this would contradict (8.1) for step μ with $S := S''$. So there exists an ordinal $\lambda < \mu$ with $S_\lambda \cap C(S'', \alpha) \neq \emptyset$. A vertex v in this intersection must have a neighbour in C_λ , which then also lies in $S' \cup C(S', \alpha)$ because $C(S'', \alpha) \subseteq C(S', \alpha)$. Thus,

$$(S' \cup C(S', \alpha)) \cap C_\lambda \neq \emptyset.$$

Since $C_\lambda \not\subseteq C(S', \alpha)$ by assumption, this implies that C_λ meets S' . But then $\lambda \in \{\lambda_{m+1}, \dots, \lambda_n\}$ and hence $S_\lambda \subseteq S''$, contradicting the fact that v lies in both S_λ and $C(S'', \alpha)$. \square

For our end-to-vertex reduction we need two more lemmas.

Lemma 8.10 (Diestel [26]). *Let H be a subgraph of a graph G , let $S \subseteq V(H)$ be finite, and let $T \subseteq V(H) \cup \Omega(G)$ be such that $T \subseteq \overline{H}$. Then \overline{H} contains a set \mathcal{P} of disjoint S - T -paths and an S - T -separator (in \overline{H}) on \mathcal{P} .*

For a set T of vertices in a graph H , a T -path is a path that meets T only in its first and last vertex. A set of paths will be called *disjoint outside* a given subgraph $Q \subseteq H$ if distinct paths meet only in Q .

Lemma 8.11 (Stein [57]; see also Diestel [26]). *Let H be a graph, $T \subseteq V(H)$ finite, and $k \in \mathbb{N}$. Then H has a subgraph H' containing T such that for every T -path $Q = s \dots t$ in H meeting $H - H'$ there are k distinct T -paths from s to t in H' that are disjoint outside Q .*

Our next lemma allows us to replace the set $A \subseteq V \cup \Omega$ in Theorem 8.5 with a set A' consisting only of vertices.

Lemma 8.12. [19] *Let $G = (V, E, \Omega)$ be a graph, and let $A, B \subseteq V \cup \Omega$ be such that $A \cap \overline{B} = \emptyset = \overline{A} \cap B$. Then there are a minor $G' = (V', E', \Omega')$ of G and sets $A' \subseteq V', B' \subseteq V' \cup \Omega'$ satisfying*

- (i) $|A'| \leq |A|$;
- (ii) if $B \subseteq V$ then $B' \subseteq B$;
- (iii) $A' \cap \overline{B'} = \emptyset = \overline{A'} \cap B'$;

(iv) G satisfies the Erdős-Menger-conjecture for A and B if G' satisfies it for A' and B' .

Proof. Applying Lemma 8.9 with $A_\Omega := A \cap \Omega$ we obtain an ordinal μ^* , subgraphs G_μ , finite vertex sets S_μ and a set of ends $I \subseteq A$. Our aim is to change G into G' by deleting and contracting certain connected subgraphs of our graphs $G_\mu - S_\mu$. By Lemma 8.9 (iii) we shall be able to do this independently for the various G_μ : for each $\mu < \mu^*$ separately, we shall find in $G_\mu - S_\mu$ a set $\mathcal{D}_1(\mu)$ of connected subgraphs to be deleted, and another set $\mathcal{D}_2(\mu)$ of connected subgraphs that will be contracted.

Fix $\mu < \mu^*$. If G_μ is empty we let $\mathcal{D}_1(\mu) = \mathcal{D}_2(\mu) = \emptyset$. Assume now that $G_\mu \neq \emptyset$. Put $A_\mu := A \cap \overline{G_\mu}$. Applying Lemma 8.10 to $H = G_\mu$ we find in $\overline{G_\mu}$ a finite set \mathcal{P} of disjoint S_μ - A_μ paths and an S_μ - A_μ separator X_μ on \mathcal{P} . We write $X_\mu = U_\mu \cup O_\mu$, where $U_\mu = X_\mu \cap V$ and $O_\mu = X_\mu \cap \Omega$, both of which are finite since $|X_\mu| \leq |\mathcal{P}| \leq |S_\mu|$. Moreover,

$$U_\mu \text{ separates } S_\mu \text{ from } A_\mu \setminus O_\mu \text{ in } G. \quad (8.3)$$

Indeed, every S_μ - $(A_\mu \setminus O_\mu)$ path in G lies in $\overline{G_\mu}$ and hence meets X_μ , and since it cannot meet O_μ unless it ends there, it meets X_μ in U_μ .

We define $\mathcal{D}_1(\mu)$ as the set of all the components D of $G - U_\mu$ whose closure \overline{D} meets $A_\mu \setminus O_\mu$. By (8.3), these components satisfy $D \subseteq G_\mu - S_\mu$, and their neighbourhood $N(D) \subseteq U_\mu$ in G is finite. In addition,

$$\overline{D} \cap O_\mu = \emptyset \text{ for all } D \in \mathcal{D}_1(\mu). \quad (8.4)$$

For if $\alpha \in \overline{D} \cap O_\mu$, say, and P is the S_μ - A_μ path in \mathcal{P} that ends in α , then P has a tail in D . Since P does not meet $U_\mu \supseteq N(D)$, this implies $P \subseteq \overline{D}$. Consequently, $S_\mu \cap D$ is not empty as it contains at least the first vertex of P . This contradicts $D \subseteq G_\mu - S_\mu$.

Put

$$H_\mu := G_\mu - \bigcup \mathcal{D}_1(\mu).$$

Note that, as every $v \in U_\mu$ lies on a path in \mathcal{P} ,

$$G_\mu \text{ contains a set of disjoint } H_\mu\text{-}A_\mu \text{ paths whose set of first points is } U_\mu. \quad (8.5)$$

By (8.3) and the definition of H_μ , we have $\overline{H_\mu} \cap A \subseteq U_\mu \cup O_\mu = X_\mu$. Since O_μ is finite, we can extend $U_\mu \cup S_\mu$ to a finite set $T_\mu \subseteq V(H_\mu)$ that separates the ends in O_μ pairwise in G . Let H'_μ be the finite subgraph of H_μ containing T_μ which Lemma 8.11 provides for $k := |S_\mu| + 1$, and for each $\alpha \in O_\mu$ let D_α be the component of $G - H'_\mu$ to which α belongs. Finally, we conclude our definitions for μ by setting $\mathcal{D}_2(\mu) := \{D_\alpha \mid \alpha \in O_\mu\}$.

Define for $i = 1, 2$

$$\mathcal{D}_i := \bigcup_{\mu < \mu^*} \mathcal{D}_i(\mu).$$

Observe that, by Lemma 8.9 (iii) and since their neighbourhoods in G are finite, the elements of $\mathcal{D}_1 \cup \mathcal{D}_2$ have pairwise disjoint closures.

Before we can define G' , we first have to introduce a graph $G^\sharp = (V^\sharp, E^\sharp, \Omega^\sharp)$ from which we will obtain G' by deleting certain vertices. Let G^\sharp be obtained from $G - \bigcup \mathcal{D}_1$ by contracting every $D_\alpha \in \mathcal{D}_2$ to a single vertex a_α , and put

$$A^* := \{a_\alpha \mid D_\alpha \in \mathcal{D}_2\}.$$

Then for $Z := \bigcup \mathcal{D}_1 \cup \bigcup \mathcal{D}_2$ we have

$$G - Z = G \cap G^\sharp = G^\sharp - A^*.$$

By Lemma 8.9 (iii) and by (8.3), the union of the sets of paths in (8.5) for all $\mu < \mu^*$ is a set of disjoint paths. Thus, for $U := \bigcup_{\mu < \mu^*} U_\mu$

there is a set of disjoint U - A paths whose set of first points is U , and whose paths meet G^\sharp only in U . (8.6)

An important property of G^\sharp is that the ends of G in $B \cap \Omega$ correspond closely to ends of G^\sharp . To establish this correspondence formally, we begin with the following observation:

Every ray of an end $\beta \in B$ has a tail in $G - Z$. (8.7)

To see this, recall that all the $D \in \mathcal{D}_1 \cup \mathcal{D}_2$ have pairwise disjoint closures, and that each of them is a connected subgraph of G whose closure contains an end or a vertex of A . Hence, a ray R of β meets only finitely many $D \in \mathcal{D}_1 \cup \mathcal{D}_2$, as we could otherwise find infinitely many disjoint R - A paths, giving $\overline{A} \cap B \neq \emptyset$ by Lemma 8.4 – a contradiction. Also, R meets every $D \in \mathcal{D}_1 \cup \mathcal{D}_2$ only finitely often. Indeed, D lies in G_μ for some $\mu < \mu^*$ and is thus, by Lemma 8.9 (i), separated from β by its finite set of neighbours $N(D)$. This establishes (8.7).

Let R_1, R_2 be two rays in $G \cap G^\sharp$, and assume that the end of R_1 lies in B . Then R_1 and R_2 are equivalent in G if and only if they are equivalent in G^\sharp . (8.8)

To prove (8.8), suppose first that R_1, R_2 are equivalent in G , i.e. belong to the same end $\beta \in B$. Then there is a ray R_3 that meets both R_1 and R_2 infinitely

often, and hence ends in β . By (8.7), R_3 has a tail in $G - Z = G^\# - A^*$, showing that R_1 and R_2 are equivalent also in $G^\#$.

Conversely, if R_1 and R_2 are joined in $G^\#$ by infinitely many disjoint paths, we can replace any vertices $a_\alpha \in V^\# \setminus V = A^*$ on these paths by finite paths in D_α to obtain infinitely many disjoint R_1 - R_2 paths in G . This completes the proof of (8.8).

We can now define our correspondence between the ends in B and certain ends of $G^\#$. For every end $\beta \in B$ there is by (8.7) an end $\beta' \in \Omega^\#$ such that $\beta \cap \beta' \neq \emptyset$. By (8.8), this end β' is unique and the map $\beta \mapsto \beta'$ is injective. Moreover,

$$B^\# := (B \cap V) \cup \{\beta' \mid \beta \in B \cap \Omega\} \subseteq V^\# \cup \Omega^\#$$

by Lemma 8.9 (i). For each $\mu < \mu^*$, let

$$A_\mu^\# := U_\mu \cup \{a_\alpha \mid \alpha \in O_\mu\},$$

if $G_\mu \neq \emptyset$; if $G_\mu = \emptyset$, put $A_\mu, A_\mu^\# := \emptyset$. Then let

$$A^\# := \left(A \setminus \left(\bigcup_{\mu < \mu^*} A_\mu \cup I \right) \right) \cup \bigcup_{\mu < \mu^*} A_\mu^\#,$$

which is a subset of $V^\#$ by Lemma 8.9 (iii),(iv). Finally, let

$$G' := G^\# - (A^\# \cap B^\#).$$

To show the assertions (i)–(iv), we will apply Lemma 8.8 to the graph $G^\#$ and the sets $A^\#$ and $B^\#$.

So, let us show that

$$\overline{(A^\# \setminus A^\#)} \cap B^\# = \emptyset = A^\# \cap \overline{(B^\# \setminus B^\#)}$$

(with closures taken in $|G^\#|$). We trivially have $A^\# \cap \overline{(B^\# \setminus B^\#)} = \emptyset$ because $A^\# \subseteq V^\#$. To prove that $\overline{(A^\# \setminus A^\#)} \cap B^\# = \emptyset$, consider an end $\beta' \in B^\#$. The corresponding end $\beta \in B$ has a neighbourhood $C := \hat{C}(S, \beta)$ in $|G|$ that avoids A . By (8.6), and since S is finite, the intersection $C \cap U =: U_C$ is finite. Also, as in the proof of (8.7), C may meet only finitely many $D_\alpha \in \mathcal{D}_2$. Denote by O_C the set of the corresponding $a_\alpha \in G^\#$. Adding to $S \setminus Z$ the sets U_C and O_C then yields a finite set $S' \subseteq V^\#$ such that the neighbourhood $\hat{C}(S', \beta')$ in $|G^\#|$ even avoids $A^\#$.

Thus, Lemma 8.8 is applicable and yields sets $A' \subseteq V'$ and $B' \subseteq V' \cup \Omega'$ satisfying (iii). For (i) use Lemma 8.8 (i), and observe that $|A^\#| \leq |A|$. (Indeed, $A^\#$ is comprised of two sets, one of which is contained in A . The

other set, $\bigcup_{\mu < \mu^*} A_\mu^\sharp$, has cardinality at most $|A|$ by (8.6).) Assertion (ii) follows from the definition of B^\sharp and Lemma 8.8 (ii).

We now prove assertion (iv) of the lemma. Suppose G' satisfies the Erdős-Menger conjecture for A' and B' . Then, by Lemma 8.8, there is also in G^\sharp a set \mathcal{P}^\sharp of disjoint A^\sharp - B^\sharp paths and an A^\sharp - B^\sharp separator X^\sharp on \mathcal{P}^\sharp . In order to turn \mathcal{P}^\sharp into a set $\mathcal{P} := \{P \mid P^\sharp \in \mathcal{P}^\sharp\}$ of disjoint A - B paths in G , consider any $P^\sharp \in \mathcal{P}^\sharp$. If the first point a of P^\sharp lies in A we leave P^\sharp unchanged, i.e. set $P := P^\sharp$. If $a \in A^\sharp \setminus (A \cup A^*)$, then $a \in U_\mu$ for some $\mu < \mu^*$, and we let P be the union of P^\sharp with an A_μ - U_μ path in G_μ that ends in a ; this can be done disjointly for different $P^\sharp \in \mathcal{P}^\sharp$ if we use the paths from (8.6). Moreover, the A_μ - H_μ path concatenated with P^\sharp in this way has only its last vertex in G^\sharp , so it will not meet any other vertices on \mathcal{P}^\sharp . Finally if $a = a_\alpha \in A^*$, we let P be obtained from P^\sharp by replacing a with a path in D_α that starts at the end α and ends at the vertex of D_α incident with the first edge of P^\sharp (the edge incident with a). In all these cases we have $P \subseteq G$, because P^\sharp has no vertex in A^* other than possibly a . And no vertex of P other than possibly its last vertex lies in B , because $B \cap V = B^\sharp \cap V^\sharp$ and any new initial segment of P lies in a subgraph $G_\lambda - S_\lambda$ of G which avoids B by Lemma 8.9 (i).

It remains to check that the paths P just defined have distinct last points in B even when the last points of the corresponding paths P^\sharp are ends. However if P^\sharp ends in $\beta' \in B^\sharp$ then its tail $P^\sharp - a \subseteq P \subseteq G$ is equivalent in G^\sharp to some ray in $\beta' \cap \beta$, by definition of β' . By (8.8) this implies $P^\sharp - a \in \beta$, so the last point of P is $\beta \in B$. And since the map $\beta \mapsto \beta'$ is well defined, these last points differ for distinct P , because the corresponding paths P^\sharp have different endpoints β' by assumption.

We still need an A - B separator on \mathcal{P} . The only vertices $x \in X^\sharp$ that do not lie on the path P obtained from the path P^\sharp containing x are points in A^* . So let X be obtained from X^\sharp by replacing every end $\beta' \in X^\sharp \cap B^\sharp$ with the corresponding end $\beta \in B$ and replacing every $a_\alpha \in X^\sharp \cap A^*$ with the end $\alpha \in A$. Since $P \in \mathcal{P}$ starts in α if P^\sharp starts in a_α (and P ends in β if P^\sharp ends in β'), this set X consists of a choice of one point from every path in \mathcal{P} .

Let us then show that

$$X \text{ is an } A\text{-}B \text{ separator in } G. \quad (8.9)$$

Suppose there exists a path $Q \subseteq G - X$ that starts in A and ends in B . Lemma 8.9 (v) enables us to choose Q as a path starting in $A \setminus I$. Our aim is to turn Q into an A^\sharp - B^\sharp path Q' in G^\sharp that avoids X^\sharp , which contradicts the choice of X^\sharp .

If Q meets $\bigcup \mathcal{D}_1$, it has a last vertex there by (8.7), in $D \in \mathcal{D}_1(\lambda)$, say. Its next vertex a lies in U_λ , by the definition of D . We then define (for the

time being) Q' as the final segment aQ of Q starting at a . If Q has no vertex in $\bigcup \mathcal{D}_1$, then either the first point of Q is a vertex $a \in A \cap A^\sharp$ (in which case we put $Q' := Q$), or Q starts at an end $\alpha \in A \setminus I$. By Lemma 8.9 (iv), there exists a $\lambda < \mu^*$ such that $\alpha \in \overline{G}_\lambda$, which implies $\alpha \in O_\lambda$. We make $a := a_\alpha$ the starting vertex of Q' and continue Q' along Q , beginning with the last D_α - G^\sharp edge on Q . Our assumption of $\alpha \notin X$ implies that $a_\alpha \notin X^\sharp$, by the definition of X . Thus in the first two cases, Q' is now a path in $G - \bigcup \mathcal{D}_1$; in the third, Q' is a path in $(G - \bigcup \mathcal{D}_1)/D_\alpha$, which starts at the vertex $a \in A^\sharp$ and avoids X^\sharp .

However, Q' may still meet \mathcal{D}_2 . And although we know from (8.7) that Q' has a last vertex in $\bigcup \mathcal{D}_2$, say in $D_{\alpha'}$, we cannot simply shorten Q' to a path $a_{\alpha'}Q'$ in G^\sharp , because it may happen that $a_{\alpha'} \in X^\sharp$. Instead, we will use Lemma 8.11 to replace any segments of Q' that meet some $D_\alpha \in \mathcal{D}_2$ (with $a_\alpha \neq a$) by paths through the corresponding G_μ that avoid X^\sharp . As we only have to deal with a finite initial segment of Q' and the D_α are all disjoint, we are able to modify Q' step by step. Eventually, we will obtain a (walk that can be pruned to a) path Q' in G^\sharp that avoids X^\sharp , yielding the desired contradiction.

So consider a segment of Q' that meets some $D_\alpha \in \mathcal{D}_2$. By definition of D_α we may assume that segment to be a T_μ -path $sQ't$ in H_μ , where μ is such that $D_\alpha \subseteq G_\mu$. By definition of H'_μ (which is a subgraph of G^\sharp by Lemma 8.9 (iii), i.e. no parts of H'_μ were deleted or contracted when we defined G^\sharp), there are $|S_\mu| + 1$ paths from s to t in H'_μ that are disjoint outside $sQ't$. But H'_μ contains at most $|S_\mu|$ vertices from X^\sharp : since these lie on disjoint paths ending in B^\sharp and S_μ separates $H'_\mu \subseteq G_\mu$ from B in G and hence from B^\sharp in G^\sharp , all of these paths must meet S_μ . So one of our $|S_\mu| + 1$ s - t paths in H'_μ avoids X^\sharp , and we can use this path to replace $sQ't$ on Q' . This completes the proof of (8.9). \square

We can now repeat the reduction for the ends of B .

Lemma 8.13. [17] *Let $G = (V, E, \Omega)$ be a graph, and let $A, B \subseteq V \cup \Omega$ be such that $A \cap \overline{B} = \emptyset = \overline{A} \cap B$. Then there are a minor G' of G and sets $A', B' \subseteq V(G')$ satisfying*

- (i) $|A'| \leq |A|$,
- (ii) G satisfies the Erdős-Menger conjecture for A and B if G' satisfies it for A' and B' .

Proof. Apply Lemma 8.12 twice, once for A and once for B . \square

One immediate consequence of Lemma 8.13 is the validity of the Erdős-Menger conjecture for ends, i.e. of Theorem 8.3, assuming the announced proof of Aharoni and Berger [5] to be correct. As another consequence we obtain the main part of Theorem 8.5:

Proposition 8.14. [17] *Let $G = (V, E, \Omega)$ be a graph, let $A, B \subseteq V \cup \Omega$ be such that $(\overline{A} \setminus A) \cap B = \emptyset = A \cap (\overline{B} \setminus B)$, and let A be countable. Then G satisfies the Erdős-Menger conjecture for A and B .*

Proof. Use Lemmas 8.8 and 8.13, and Theorem 8.1. □

The next section will be spent on strengthening Proposition 8.14 to the full generality of Theorem 8.5.

8.4 Countable separators

In this section we prove Theorem 8.5. But first, let us establish a weaker version in which the separator X also satisfies $A \cap (\overline{X} \setminus X) = \emptyset$:

Lemma 8.15. [17] *Let $G = (V, E, \Omega)$ be a graph, let $A, B \subseteq V \cup \Omega$ satisfy $A \cap \overline{B} = \emptyset = \overline{A} \cap B$, and suppose there exists a countable A - B separator $X \subseteq V \cup \Omega$ in G with $A \cap (\overline{X} \setminus X) = \emptyset$. Then G satisfies the Erdős-Menger conjecture for A and B .*

Proof. First, we find a countable A - B separator Y and a set \mathcal{P}_Y of disjoint Y - A paths satisfying

- (a) in every $y \in Y$ there starts a path of \mathcal{P}_Y ;
- (b) $Y \cap \Omega \subseteq A \cup B$;
- (c) $(\overline{Y} \setminus Y) \cap B = \emptyset = Y \cap (\overline{B} \setminus B)$.

We may clearly assume that $X \cap \Omega \subseteq A \cup B$. As $\overline{A} \cap B = \emptyset$, this implies that $X \cap (\overline{A} \setminus A) = \emptyset$. As also $A \cap (\overline{X} \setminus X) = \emptyset$ and X is countable, we can use Proposition 8.14 to obtain a set \mathcal{P}_1 of disjoint X - A paths and an X - A separator Y on \mathcal{P}_1 in G .

We claim that Y together with the set

$$\mathcal{P}_Y := \{yP \mid P \in \mathcal{P}_1 \text{ and } y \in Y \cap P\}$$

of disjoint Y - A paths is as desired.

Indeed, Y is countable because \mathcal{P}_1 is. Further, Y is an A - B separator in G : any path starting in B and ending in A meets X and thus has a

subpath P starting in X and ending in A . But then P also meets Y . The conditions (a) and (b) are easily checked; for the latter recall that we assumed $X \cap \Omega \subseteq A \cup B$.

Next, we show (c). By (b), any end $\alpha \in Y \cap (\overline{B} \setminus B)$ lies in $A \cap \overline{B}$, which is empty by assumption. Now consider an end $\beta \in (\overline{Y} \setminus Y) \cap B$. Every neighbourhood of β contains a point of Y ; it even contains infinitely many points of Y , as otherwise we could find a neighbourhood containing no point of Y . Choose a neighbourhood $\hat{C}(S, \beta)$ of β that contains no point of A (which is possible, as $\overline{A} \cap B = \emptyset$). The infinitely many points of Y lying in $\hat{C}(S, \beta)$ are linked to A by disjoint paths in \mathcal{P}_Y . All of these infinitely many paths must meet the finite set S , a contradiction. Therefore, (c) is proved.

Having found Y and \mathcal{P}_Y , we now apply Proposition 8.14 again, this time for the sets Y and B . Thus, we get a system \mathcal{P}_2 of disjoint B - Y paths and a B - Y separator Z on \mathcal{P}_2 in G .

Finally, let \mathcal{P} be the system of disjoint B - A paths obtained by concatenating every B - Y path $P \in \mathcal{P}_2$ with the unique Y - A path $P' \in \mathcal{P}_Y$ from (a) that starts at the endpoint of P . Note that these are indeed paths: if P terminates in an end $\omega \in Y$ then, by (b), $\omega \in A$ or $\omega \in B$. In the former case P' is trivial, in the latter case P . Clearly, Z lies on \mathcal{P} . All that is left to show is that Z is also an A - B separator. So consider a path P that starts in A and ends in B . By definition of Y , P meets Y and thus has a subpath starting in Y and ending in B . This subpath cannot avoid the Y - B separator Z . \square

We can now complete the proof of Theorem 8.5, which we restate:

Theorem 8.5. [17] *Let $G = (V, E, \Omega)$ be a graph, let $A, B \subseteq V \cup \Omega$ satisfy $A \cap \overline{B} = \emptyset = \overline{A} \cap B$, and suppose that there exists a countable A - B separator $X \subseteq V \cup \Omega$ in G . Then G satisfies the Erdős-Menger conjecture for A and B .*

Proof. We will start by constructing a set $A' \subseteq A$ and a countable A' - B separator X' with

$$(\overline{X'} \setminus X') \cap A' = \emptyset, \quad (8.10)$$

to which we will then apply Lemma 8.15.

Using Lemma 8.9 with

$$A_\Omega := (\overline{X} \setminus X) \cap A,$$

we find I, μ^* and families $(G_\mu)_{\mu < \mu^*}$ and $(S_\mu)_{\mu < \mu^*}$ satisfying the assertions (i)–(v) of Lemma 8.9. As before, write $C_\mu := G_\mu - S_\mu$. Setting $A' := A \setminus I$ and

$$X' := \bigcup_{\mu < \mu^*} S_\mu \cup \left(X \setminus \bigcup_{\mu < \mu^*} \overline{C_\mu} \right)$$

we claim that X' is a countable A' - B separator satisfying (8.10).

To see that X' is countable, recall that the sets \overline{C}_μ are disjoint for different μ (Lemma 8.9 (iii)), and that each \overline{C}_μ with $S_\mu \neq \emptyset$ contains an end $\alpha \in A_\Omega$ (Lemma 8.9 (ii)). Since $\alpha \in \overline{X}$ (by definition of A_Ω), we have $\overline{C}_\mu \cap X \neq \emptyset$, so by the countability of X there are only countably many such μ .

Let us now show that X' is an A' - B separator. Consider a path Q from A' to B . Since X is an A - B separator, Q must meet X . If Q meets X outside X' , it meets X in \overline{C}_μ for some $\mu < \mu^*$. By Lemma 8.9 (i), however, Q cannot be contained in \overline{C}_μ , so Q meets $S_\mu \subseteq X'$.

To prove (8.10), suppose there is an end $\omega \in (\overline{X'} \setminus X') \cap A'$. Let us show first that

$$\omega \notin \overline{X}. \quad (8.11)$$

Suppose otherwise; then we have either $\omega \in (\overline{X} \setminus X) \cap A = A_\Omega$ or $\omega \in X \setminus X'$. In both cases (in the first by Lemma 8.9 (iv), observe that $\omega \notin I$ as $\omega \in A'$; in the second by construction of X') there is a $\mu < \mu^*$ with $\omega \in \overline{G}_\mu$. Then by Lemma 8.9 (iii), the sets of the form $\hat{C}(S_\mu, \omega)$ are neighbourhoods of ω that avoid X' , contradicting $\omega \in \overline{X'} \setminus X'$.

By (8.11), there exists a finite set S such that $C(S, \omega) \cap X = \emptyset$. By our choice of ω , however, $C(S, \omega)$ contains infinitely many vertices from X' , and hence from $X' \setminus X \subseteq \bigcup_{\mu < \mu^*} S_\mu$. Since each S_μ is finite, we can thus find infinitely many $\mu < \mu^*$ and corresponding vertices $s_\mu \in S_\mu \cap C(S, \omega)$ that are distinct for different μ . By Lemma 8.9 (ii) and $S_\mu = N(C_\mu)$, each s_μ sends a path $P_\mu \subseteq \overline{G[C_\mu \cup \{s_\mu\}]}$ to an end in $A_\Omega \subseteq \overline{X}$. These P_μ are disjoint by Lemma 8.9 (iii), so only finitely many of them meet S . Every other P_μ lies entirely in $\overline{C}(S, \omega)$, so $\overline{C}(S, \omega) \cap \overline{X} \neq \emptyset$. But then also $C(S, \omega) \cap X \neq \emptyset$, contradicting our choice of S . This establishes (8.10).

Applying Lemma 8.15 to A', B and the separator X' (note that $\overline{A'} \cap B = \emptyset = A' \cap \overline{B}$, as $A' \subseteq A$), we obtain a set \mathcal{P} of disjoint A' - B paths and an A' - B separator on \mathcal{P} , which by Lemma 8.9 (v) is also an A - B separator. \square

8.5 Disjoint closures

We restate the second main result of this chapter, which we shall prove in this section.

Theorem 8.6. [17] *Every graph $G = (V, E, \Omega)$ satisfies the Erdős-Menger conjecture for all sets $A, B \subseteq V \cup \Omega$ that have disjoint closures in $|G|$.*

Our proof follows that of Theorem 8.2 as given in Diestel [30], and in addition we will draw on techniques from the proof of Lemma 8.12. Our

aim is to reduce our problem to rayless graphs, and then apply Theorem 8.7. We thus need to dispose of the ends in G , which will be achieved in three steps. First, we delete all ends in \overline{A} . More precisely, we reduce the problem to a minor G' of G and to sets A', B' so that the closure of A' contains no ends. In the next step we repeat this procedure for B' . To preserve what we have gained in the first step, we have to be careful that no new ends are introduced into $\overline{A'}$. All this amounts to the following lemma:

Lemma 8.16. [17] *Let $G = (V, E, \Omega)$ be a graph, and let $A, B \subseteq V \cup \Omega$ be such that $\overline{A} \cap \overline{B} = \emptyset$. Then there exists a minor $G' = (V', E', \Omega')$ of G and sets $A', B' \subseteq V' \cup \Omega'$ that satisfy the following conditions:*

- (a) $\Omega' \cap \overline{A'} = \emptyset$ (in particular $A' \subseteq V'$);
- (b) if $\Omega \cap \overline{B} = \emptyset$ then $\Omega' \cap \overline{B'} = \emptyset$ (and in particular $B' \subseteq V'$);
- (c) $\overline{A'} \cap \overline{B'} = \emptyset$; and
- (d) the Erdős-Menger conjecture holds for A and B in G if it holds for A' and B' in G' .

Two applications of Lemma 8.16, one for A and another for B , reduces our problem to the case that G has no ends in $\overline{A} \cup \overline{B}$. We then eliminate the remaining ends by the following lemma from Diestel [30]:

Lemma 8.17 (Diestel [30]). *Let $G = (V, E, \Omega)$ be a graph, and let $A, B \subseteq V$ be such that $\Omega \cap (\overline{A} \cup \overline{B}) = \emptyset$. Then G has a rayless subgraph $G' \subseteq G$ containing $A \cup B$ such that the Erdős-Menger conjecture for A and B holds in G if it does in G' .*

Finally, we apply Theorem 8.7 to complete the proof. It thus remains to establish Lemma 8.16. We shall need the following easy lemma from Diestel and Kühn [32]:

Lemma 8.18 (Diestel and Kühn [32]). *Let G be a connected graph, and let $U \subseteq V(G)$ be an infinite set of vertices. Then G contains either a comb with $|U|$ teeth in U or a subdivided star with $|U|$ leaves in U .*

Proof of Lemma 8.16. First, we construct a subgraph $M \subseteq G$ whose closure does not contain ends of \overline{A} . More formally, our aim is that in $|G|$

$$\Omega \cap \overline{A} \cap \overline{M} = \emptyset. \tag{8.12}$$

Our desired graph G' will then be obtained from a supergraph of M .

We define M by transfinite ordinal recursion, as a limit $M = \bigcap_{\mu \leq \mu^*} M_\mu$ of a well-ordered descending family of subgraphs M_μ indexed by ordinals. Put $M_0 := G$, and for a limit ordinal $\mu \neq 0$ let $M_\mu := \bigcap_{\mu' < \mu} M_{\mu'}$. Now, consider a successor ordinal $\mu + 1$. If $\Omega \cap \overline{A} \cap \overline{M}_\mu = \emptyset$ put $\mu^* := \mu$ and $M := M_\mu$, and terminate the recursion. Otherwise, there is an $\alpha_\mu \in \Omega \cap \overline{A} \cap \overline{M}_\mu$. Since $\overline{A} \cap \overline{B} = \emptyset$, we can choose a finite vertex set L_μ such that the open neighbourhood $\hat{C}(L_\mu, \alpha_\mu)$ is disjoint from B . Put $C_\mu := C(L_\mu, \alpha_\mu)$ and $M_{\mu+1} := M_\mu - C_\mu$. Observe that $C_\mu \cap M_\mu$ is never empty as $\alpha_\mu \in \overline{M}_\mu$. Thus, the recursion terminates.

Let \mathcal{C} be the set of components of $G - M$. For every $C \in \mathcal{C}$ put $S_C := N_G(C)$, $G_C := G[V(C) \cup S_C]$ and $A_C := A \cap \overline{G_C}$. We shall now proceed in a similar way as in the proof of Lemma 8.12: in order to obtain G' we will delete and contract certain connected subgraphs, each of which will lie in a $C \in \mathcal{C}$. Thus, the G_C here play a similar role as the G_μ in Lemma 8.12. Now, in contrast to the set of neighbours S_μ we deal with there, the sets S_C here are not necessarily finite. Consequently, it may happen that two $C \in \mathcal{C}$ do not have disjoint closures. However, for our purposes it is sufficient to know that no end in A lies in the closure of two elements of \mathcal{C} :

$$A_C \cap A_D \subseteq S_C \cap S_D \text{ for two different } C, D \in \mathcal{C} \quad (8.13)$$

Indeed, suppose there is an end $\alpha \in A_C \cap A_D = A \cap \overline{G_C} \cap \overline{G_D}$. Then every neighbourhood $\hat{C}(T, \alpha)$ of α contains vertices of both G_C and G_D . Being connected, $\hat{C}(T, \alpha)$ also contains a G_C - G_D path, and therefore a vertex of $S_C \subseteq V(M)$. This implies $\alpha \in \overline{M}$, a contradiction to (8.12). As no vertices other than those of $S_C \cap S_D$ lie in $A_C \cap A_D$, (8.13) is proved.

Now, consider a given $C \in \mathcal{C}$. Although S_C may be infinite it can be separated from A_C by finitely many vertices and ends. Indeed, we claim that

$$\overline{G_C} \text{ contains a finite set } \mathcal{P}_C \text{ of disjoint } S_C\text{-}A_C \text{ paths and a } S_C\text{-}A_C \text{ separator } X_C \text{ on } \mathcal{P}_C. \quad (8.14)$$

First, consider a non-trivial S_C - A_C path P . P is either completely contained in G_C or in $G - C$. Suppose the latter. Then P terminates in an end $\alpha \in A_C$. By (8.13), P can meet every $D \in \mathcal{C}$ only finitely often, hence it meets M infinitely often, a contradiction to (8.12). Thus, every S_C - A_C path lies completely in G_C .

Next, let us show that each set of disjoint S_C - A_C paths in G_C is finite. So suppose there is an infinite set P_1, P_2, \dots of such paths. Let $S \subseteq S_C$ be the set of first vertices of the P_i . We claim there is a comb in G with teeth in S . If not, applying Lemma 8.18 to $H := C \cup \bigcup_{i \in \mathbb{N}} P_i$ we obtain a subdivided infinite star with leaves in S . Since each vertex $s \in S$ has degree 1 in H , the

centre v of the star must lie in C . Let μ be the step when v was deleted from G , i.e. $\mu := \min\{\mu' \mid v \in V(C_{\mu'})\}$. Then, the finite set L_μ separates v from S , which is impossible. Thus, there is a comb with teeth in S . Let $\omega \in \Omega$ be the end of its spine. Then every neighbourhood of ω contains infinitely many vertices of S and then also infinitely many of the P_i . Consequently, infinitely many elements of A lie in the neighbourhood. Thus, $\omega \in \overline{S_C} \cap \overline{A} \subseteq \overline{M} \cap \overline{A}$, a contradiction to (8.12).

Taking an inclusion-maximal (and hence finite) set of disjoint S_C - A_C paths in G we see that S_C is separated from A_C in G by a countable separator (namely by the union of the vertex sets of the paths together with the set of last points). Furthermore, (8.12) implies $\overline{S_C} \cap A_C \cap \Omega = \emptyset$. Thus, by Lemma 8.8 and Theorem 8.5, there is a set \mathcal{P}_C of disjoint S_C - A_C paths in G and a separator X_C on \mathcal{P}_C . By the preceding arguments, \mathcal{P}_C is a finite set of paths in G_C , which establishes (8.14).

For $C \in \mathcal{C}$, put $U_C := X_C \cap V$ and $O_C := X_C \cap \Omega$. As in the proof of Lemma 8.12 we will delete certain subgraphs of G_C that lie “behind” U_C and contract others that contain ends of O_C . First of all, we see that U_C separates S_C from $A_C \setminus O_C$. This is exactly the same as (8.3) in the proof of Lemma 8.12.

We then define $\mathcal{D}_1(C)$ as the set of all the components D of $G - U_C$ whose closure \overline{D} meets $A_C \setminus O_C$. Because U_C is an S_C - $(A_C \setminus O_C)$ separator, these components satisfy $D \subseteq G_C - S_C$, and their neighbourhood $N(D) \subseteq U_C$ in G is finite. Also, $\overline{D} \cap O_C = \emptyset$ for all $D \in \mathcal{D}_1(C)$, which can be proved as (8.4) of Lemma 8.12. We put $H_C := G_C - \bigcup \mathcal{D}_1(C)$ and see that G_C contains a set of disjoint H_C - A_C paths whose set of first points is U_C .

To find the subgraphs which will be contracted, we extend in the proof of Lemma 8.12 the finite set $U_\mu \cup S_\mu$ to a finite set T_μ that separates the ends in O_μ pairwise. This enables us to apply Lemma 8.11. Here, S_C may be an infinite set, so to find a suitable finite set T_C we cannot use a superset of S_C . Instead, consider for each $\alpha \in O_C$ the ordinal μ_α for which $\alpha \in \overline{C_{\mu_\alpha}}$. Then, the finite set L_{μ_α} separates α from $S_C \subseteq V(M)$. The union of these finitely many finite sets together with U_C can be taken as the finite set T_C we need for Lemma 8.11.

Define H'_C as the finite subgraph of H_C containing T_C which Lemma 8.11 provides for $k := |T_C| + 1$, and for each $\alpha \in O_C$ let D_α be the component of $G - H'_C$ to which α belongs. Finally, we set $\mathcal{D}_2(C) := \{D_\alpha \mid \alpha \in O_C\}$ and $\mathcal{D}_i := \bigcup_{C \in \mathcal{C}} \mathcal{D}_i(C)$, $i = 1, 2$. By deleting all subgraphs in \mathcal{D}_1 and by contracting each subgraph $D_\alpha \in \mathcal{D}_2$ to a vertex a_α we obtain the graph

$G^\sharp = (V^\sharp, E^\sharp, \Omega^\sharp)$. We put

$$A^\sharp := (A \setminus \bigcup_{C \in \mathcal{C}} A_C) \cup \bigcup_{C \in \mathcal{C}} U_C \cup \{a_\alpha \mid D_\alpha \in \mathcal{D}_2\}.$$

Observe that $A^\sharp \subseteq V^\sharp$ because of (8.12).

Let us show that (a) holds for A^\sharp in G^\sharp . Indeed, suppose not. Then, there is by Lemma 8.4 a comb $K^\sharp \subseteq G^\sharp$ with teeth in A^\sharp and spine R , say. Consider a R - A^\sharp path P of the comb. If P ends in a vertex of U_C for some $C \in \mathcal{C}$ then we extend P using the corresponding path in \mathcal{P}_C to a R - A path. If P ends in one of the contracted vertices a_α we substitute its last edge by a path through D_α ending in α so that again we obtain a R - A path. Because of (8.13), and since the D_α and all the paths in $\bigcup_{C \in \mathcal{C}} \mathcal{P}_C$ are disjoint, all these changed paths are still disjoint. Thus we have found a comb $K \subseteq G$ with spine R and teeth in A . Only finitely many of the R - A paths of K may meet M as otherwise the end of R lies in the closure of both A and M , a contradiction to (8.12). Thus, we may assume that K is contained in a $C \in \mathcal{C}$. In particular, $R \subseteq C \cap G^\sharp$. But the finite set X_C separates $C \cap G^\sharp$ from A_C , contradicting that K contains infinitely many disjoint R - A_C paths. This proves (a) for A^\sharp .

Making use of (8.13), we see in a similar way as in the proof of Lemma 8.12, that every ray R of an end $\beta \in B$ has a tail in $G \cap G^\sharp$. There thus is a $\beta' \in \Omega^\sharp$ with $\beta \cap \beta' \neq \emptyset$ and it can be shown as in the proof of Lemma 8.12 that this mapping is injective. Put

$$B^\sharp := (B \cap V) \cup \{\beta' \mid \beta \in B \cap \Omega\} \subseteq V^\sharp \cup \Omega^\sharp.$$

The graph G^\sharp together with the sets A^\sharp and B^\sharp is almost what we want. Indeed, (a) is satisfied and below we shall see that (b) and (d) hold as well. Only (c) fails since A^\sharp and B^\sharp may share some vertices. However, $\overline{A^\sharp}$ does not contain ends by (a). Thus

$$\emptyset = \overline{A^\sharp} \cap \overline{B^\sharp} \cap \Omega^\sharp = (\overline{A^\sharp} \setminus A^\sharp) \cap (\overline{B^\sharp} \setminus B^\sharp).$$

Now, in a similar way as in Lemma 8.8 we find sets A' and B' in the graph $G' := G^\sharp - (A^\sharp \cap B^\sharp)$ such that (c) holds and properties (a), (b) and (d) are preserved. Therefore, to finish the proof it suffices to show (b) and (d) for A^\sharp and B^\sharp in G^\sharp .

For (b), assume that $\Omega^\sharp \cap \overline{B^\sharp} \neq \emptyset$. Then either B^\sharp contains ends, or there is an end $\omega \in \Omega^\sharp$ such that every neighbourhood $\hat{C}(S, \omega)$ in G^\sharp contains vertices of B^\sharp . In the latter case, by (a), we may suppose that $\hat{C}(S, \omega)$ does not contain vertices of A^\sharp , and thus view it as a neighbourhood in G containing vertices of B . Thus, in both cases, $\Omega \cap \overline{B} \neq \emptyset$, which proves (b).

The proof of (d) is completely analogous to the proof of (iv) in Lemma 8.12. Note that when we change the set \mathcal{P}^\sharp of disjoint A^\sharp - B^\sharp paths into A - B paths, the paths still have, by (8.13), distinct first points. \square

Summary and conclusion

Diestel and Kühn were motivated by a question of Richter [54], who asked how the fact that the fundamental circuits of every spanning tree generate the cycle space might best be generalised to locally finite graphs. When they developed the topological cycle space $\mathcal{C}(G)$ they answered this question and showed that two other basic properties extend as well, which resulted in Theorems 1.4 and 1.5. Only afterwards it was realised that infinite cycles and the space $\mathcal{C}(G)$ are much more powerful concepts.

The main objective of this thesis was to contribute to the effort of showing the usefulness of $\mathcal{C}(G)$ in extending finite theorems to locally finite graphs. We have presented a number of results in support of this. In particular, we have extended the following theorems to either locally finite graphs or, slightly more generally, to graphs satisfying (1.3):

- Gallai's theorem (Theorem 2.4);
- MacLane's planarity criterion (Theorem 3.3);
- Kelmans' planarity criterion (Theorem 3.22);
- Whitney's planarity criterion (Theorem 4.6); and
- duality in terms of spanning trees (Theorem 4.17).

We have seen that each of these results either fails or is weakened considerably if infinite cycles are disregarded.

Diestel and Kühn did not deal with the characterisation of cycle space elements in terms of degrees but indicated that this would necessitate a notion of an end degree. We have introduced such a notion, which enabled us to prove the following important special case:

- for a locally finite graph G , $E(G) \in \mathcal{C}(G)$ if and only if every vertex and every end has even degree (Theorem 5.4).

Moreover, we demonstrated that spanning infinite cycles are a also worthy contender for being the right infinite analogons of Hamilton cycles. Indeed, we showed that:

- every locally finite planar 6-connected graph with at most finitely many ends has a Hamilton circle (Theorem 6.5).

This supports our conjecture that Tutte's theorem about Hamilton cycles extends to locally finite graphs.

The finite versions of these theorems together with Tutte's generating theorem make up a good part of what is known about the cycle space of a finite graph. That all these extend to $\mathcal{C}(G)$ seems to be a very good indication of the usefulness of $\mathcal{C}(G)$, and make it a quite successful and fruitful notion.

We also studied two related problems. First, we asked whether and when minimal generating systems exist if infinite sums are allowed. Second, we used the topological space $|G|$ to prove certain special cases of the end version of the Erdős-Menger conjecture.

Bibliography

- [1] R. Adomaitis, *Infinite hamiltonian graphs*, preprint 1995.
- [2] R. Aharoni, *Menger's theorem for graphs containing no infinite paths*, Europ. J. Combinatorics **4** (1983), 201–204.
- [3] ———, *Menger's theorem for countable graphs*, J. Combin. Theory (Series B) **43** (1987), 303–313.
- [4] ———, *A few remarks on a conjecture of Erdős on the infinite version of Menger's theorem*, The Mathematics of Paul Erdős II (R.L. Graham and J. Nešetřil, eds.), Springer-Verlag, Heidelberg 1997.
- [5] R. Aharoni and E. Berger, *Menger's theorem for infinite graphs*, preprint 2005.
- [6] R. Aharoni and R. Diestel, *Menger's theorem for a countable source set*, Comb., Probab. Comput. **3** (1994), 145–156.
- [7] B. Alspach, L. Goddyn, and C-Q. Zhang, *Graphs with the circuit cover property*, Trans. Am. Math. Soc. **344** (1994), 131–154.
- [8] Th. Andreae, *Über maximale Systeme von kantendisjunkten unendlichen Wegen in Graphen*, Math. Nachr. **101** (1981), 219–228.
- [9] K. Appel and W. Haken, *Every planar map is four colorable. I: Discharging*, Ill. J. Math. **21** (1977), 429–490.
- [10] K. Appel, W. Haken, and J. Koch, *Every planar map is four colorable. II: Reducibility*, Ill. J. Math. **21** (1977), 491–567.
- [11] J.A. Bondy and U.S.R. Murty, *Graph theory with applications*, Macmillan Press, London, 1976.
- [12] C.P. Bonnington and R.B. Richter, *Graphs embedded in the plane with finitely many accumulation points*, J. Graph Theory **44** (2003), 132–147.

- [13] H. Bruhn, unpublished.
- [14] ———, *Generating the cycle space by induced non-separating cycles in locally finite graphs and in graphs with at most one end*, 2001, Diploma thesis.
- [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*, to appear in *Comb., Probab. Comput.*
- [17] H. Bruhn, R. Diestel, and M. Stein, unpublished.
- [18] ———, *Cycle-cocycle partitions and faithful cycle covers for locally finite graphs*, to appear in *J. Graph Theory*.
- [19] ———, *Menger's theorem for infinite graphs with ends*, to appear in *J. Graph Theory*.
- [20] H. Bruhn and A. Georgakopoulos, *Bases for thin sums*, unpublished.
- [21] H. Bruhn and M. Stein, unpublished.
- [22] ———, *MacLane's planarity criterion for locally finite graphs*, submitted.
- [23] ———, *On end degrees and infinite circuits in locally finite graphs*, submitted.
- [24] H. Bruhn and X. Yu, *Hamilton circles in planar locally finite graphs*, preprint 2005.
- [25] N.G. de Bruijn and P. Erdős, *A colour problem for infinite graphs and a problem in the theory of relations*, *Indag. Math.* **13** (1951), 371–373.
- [26] R. Diestel, *The countable Erdős-Menger conjecture with ends*, to appear in *J. Combin. Theory (Series B)*.
- [27] ———, *End spaces and spanning trees*, preprint 2004.
- [28] ———, *Spanning trees and k -connectedness*, *J. Combin. Theory (Series B)* **56** (1992), 263–277.

- [29] ———, *The cycle space of an infinite graph*, *Comb., Probab. Comput.* **14** (2005), 59–79.
- [30] ———, *The Erdős-Menger conjecture for source/sink sets with disjoint closures*, *J. Combin. Theory (Series B)* **93** (2005), 107–114.
- [31] ———, *Graph theory* (3rd ed., in preparation), Springer-Verlag, 2005.
- [32] R. Diestel and D. Kühn, *Graph-theoretical versus topological ends of graphs*, *J. Combin. Theory (Series B)* **87** (2003), 197–206.
- [33] ———, *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] G.A. Dirac, *Some theorems on abstract graphs*, *Proc. London Math. Soc.* **2** (1952), 69–81.
- [37] G.A. Dirac and S. Schuster, *A theorem of Kuratowski*, *Nederl. Akad. Wetensch. Proc. Ser. A* **57** (1954), 343–348.
- [38] A. Georgakopoulos, *Infinite highly connected planar graphs of large girth*, In Preparation.
- [39] G. Hahn, F. Laviolette, and J. Širáň, *Edge-ends in countable graphs*, *J. Combin. Theory (Series B)* **70** (1997), 225–244.
- [40] R. Halin, *Über unendliche Wege in Graphen*, *Math. Annalen* **157** (1964), 125–137.
- [41] ———, *Miscellaneous problems on infinite graphs*, *J. Graph Theory* **35** (2000), 128–151.
- [42] D.W. Hall and G.L. Spencer, *Elementary topology*, John Wiley, New York 1955.
- [43] H.A. Jung, *Wurzelbäume und unendliche Wege in Graphen*, *Math. Nachr.* **41** (1969), 1–22.
- [44] 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.

- [45] F. Laviolette, *Decompositions of infinite graphs: Part II - Circuit decompositions*, preprint 2003.
- [46] S.C. Locke, *Long paths and the cycle space of a graph*, *Ars Combin.* **33** (1992), 77–85.
- [47] L. Lovász, *Combinatorial problems and exercises* (2nd edition), North-Holland, Amsterdam, 1993.
- [48] S. MacLane, *A combinatorial condition for planar graphs*, *Fund. Math.* **28** (1937), 22–32.
- [49] W. Mader, *Homomorphieeigenschaften und mittlere Kantendichte von Graphen*, *Math. Ann.* **174** (1967), 265–268.
- [50] C.St.J.A. Nash-Williams, *Decomposition of graphs into two-way infinite paths*, *Can. J. Math.* **15** (1963), 479–485.
- [51] ———, *Infinite graphs—a survey*, *J. Combin. Theory* **3** (1967), 286–301.
- [52] ———, *Hamiltonian lines in infinite graphs with few vertices of small valency*, *Aequationes Math.* **7** (1971), 59–81.
- [53] N. Polat, *Sur une extension du théorème de menger aux graphes infinis*, *European J. Combin.* **4** (1983), 359–365.
- [54] B. Richter, *Communication at the 1999 oberwolfach conference on Graph Theory*.
- [55] R.B. Richter and C. Thomassen, *3-connected planar spaces uniquely embed in the sphere*, *Trans. Am. Math. Soc.* **354** (2002), 4585–4595.
- [56] P.D. Seymour, *Sums of circuits*, *Graph Theory and Related Topics* (J.A. Bondy and U.S.R. Murty, eds.), Academic Press, 1979, pp. 341–355.
- [57] M. Stein, *Die Erdős-Menger'sche Vermutung für abzählbare Graphen*, Diplomarbeit Universität Hamburg 2002.
- [58] ———, *Forcing highly connected subgraphs in locally finite graphs*, preprint 2004.
- [59] C. Thomassen, *Planarity and duality of finite and infinite graphs*, *J. Combin. Theory (Series B)* **29** (1980), 244–271.

- [60] ———, *Duality of infinite graphs*, J. Combin. Theory (Series B) **33** (1982), 137–160.
- [61] ———, *A theorem on paths in planar graphs*, J. Graph Theory **7** (1983), 169–176.
- [62] W.T. Tutte, *A theorem on planar graphs*, Trans. Am. Math. Soc. **82** (1956), 99–116.
- [63] ———, *How to draw a graph*, Proc. London Math. Soc. **13** (1963), 743–768.
- [64] K. Wagner, *Graphentheorie*, Bibliogr. Inst., Mannheim, 1970.
- [65] H. Whitney, *Non-separable and planar graphs*, Trans. Am. Math. Soc. **34** (1932), 339–362.
- [66] X. Yu, *Infinite paths in planar graphs I–V*.