Bicycles and Left-Right Tours in Locally Finite Graphs

von
Melanie Win Myint, M. S.

Dem Department Mathematik
der Universität Hamburg
zur Erlangung des akademischen Grades

Doktor der Naturwissenschaften
– Dr. rer. nat. –

vorgelegte Dissertation

Hamburg
2009

Als Dissertation angenommen vom Department Mathematik der Universität Hamburg

aufgrund der Gutachten

von ________________________

und ________________________

Hamburg, den ______________

Prof. Dr. ________________________
Sprecher des Departments Mathematik

Acknowledgments

I would like to thank Professor Diestel and Professor Andreae for advising and evaluating this thesis. I especially thank my adviser Reinhard Diestel for the opportunities he has given me and the room for mathematical and personal growth working on this thesis has offered.

I owe many thanks to Dr. Henning Bruhn with whom I enjoyed working and who gave me valuable feedback on my thesis. I am very appreciative of the support from Evan Morgan and Moritz Schmitt regarding non-mathematical matters of this work.

This thesis is dedicated to my family and friends. I am very grateful for their constant support and love. With respect to this work, I especially thank my parents Angela Terfloth and Donald Win Myint, Beate Terfloth, Nikolas Win Myint, Moritz Schmitt, Frank Ciesinski, Evan Morgan, Esther Hartwich, Dirk Schlatter, and Markus Suing. Last but not least, I would like to thank Benny and Lindy for coming into my life.

Contents

Acknowledgments iii

1 Introduction 1

2 Definitions and Preliminaries 7

3 The Tripartition Theorem 15

4 Principal Cuts 23

5 Left-Right Tours 29

6 LRTs Generate the Bicycle Space 43

7 The ABL Planarity Criterion 53

8 Pedestrian Graphs 69

Bibliography 85

Chapter 1

Introduction

The aim of this thesis is to portray some interesting aspects of graph theory and, more precisely, infinite graph theory. In particular, we will focus on objects called *bicycles* and some other concepts they relate to, such as *left-right tours* and *pedestrian graphs.*

The origins of graph theory are usually dated back to Leonhard Euler and his paper from 1736 about the problem of the Seven Bridges of Königsberg. While in the beginning finite and infinite graphs were studied with equal interest, in recent decades there has been far more emphasis only on finite structures. Yet infinite graphs remain fascinating and important, despite the fact that many standard theorems for finite graphs are very difficult to extend to infinite graphs. These extensions often require new methods or viewpoints in order to prove theorems analogous to the ones for finite graphs. One of the challenges in infinite graph theory is that we often do not know how certain structures will behave at the point of infinity. The points at infinity that occur in an infinite graph are called *ends*, and they were first introduced in the 1940s.

The importance of cycles in a graph is illustrated by the fact that they were already the objects of study in Euler's paper. Hence the question of how cycles behave in infinite graphs is fundamental, yet the definition of the cycle space from finite graph theory has proven to be insufficient in infinite graph theory. In 2004, Reinhard Diestel and Daniela Kühn introduced the notion of *infinite cycles* and provided a definition for the cycle space of an infinite graph which included these together with the finite cycles [12], [13]. This definition has proven to be the most natural. Many famous theorems about the cycles in finite graphs (which do not directly extend to infinite graphs when considering only finite cycles) have now been generalized to infinite graphs: MacLane's planarity criterion [6], Tutte's generating theorem [2], Whitney's theorem that a finite graph has a dual if and only if it is planar [3], Gallai's theorem that every finite graph has a vertex partition into two parts each inducing an element of its cycle space [4] and, last but not least, Euler's theorem from 1736 [12], [7].

This thesis is yet another illustration of the importance of infinite cycles. Since a bicycle is, in particular, also an element of the cycle space, there is a sensible way to define infinite bicycles. There are a number of important results involving bicycles that hold for finite graphs, but which fail when we consider only finite cycles in infinite graphs. In this thesis we extend four such results to infinite graphs using the cycle space as introduced by Diestel and Kühn. This thesis evolved from joint work with Henning Bruhn and Stefanie Kosuch which resulted in [5]. Therefore, some main results of this thesis, such as Theorems 3.2, 5.5, 6.4, and 7.2, have been published in [5].

So, let us get a better idea of these bicycles. The set of edge (sub)sets in a graph G together with symmetric difference as addition forms a $\mathbb{Z}_2$-vector space, which we call its *edge space* $\mathcal{E}(G)$. Two important subspaces of $\mathcal{E}(G)$ are the *cycle space* $\mathcal{C}(G)$, which is the set of all sums of (edge sets of) cycles, and the *cut space* $\mathcal{C}^*(G)$, which is the set of all (edge) cuts. Although cuts (or co-cycles) and cycle space elements are orthogonal to each other, it is possible for an edge set to be an element of both $\mathcal{C}(G)$ and $\mathcal{C}^*(G)$. Such an edge set is then called a *bicycle* and the space $\mathcal{B}(G){:=}\mathcal{C}(G) \cap \mathcal{C}^*(G)$ is the *bicycle space*. The set of bold edges in the graph in Figure 1.1 (a) is an example of a finite bicycle.

Bicycles in finite graphs have been widely studied, and a number of fundamental results involving bicycles are known. We will extend four of these to an important class of infinite graphs, namely to *locally finite* graphs, ie. to graphs in which every vertex has finite degree. Although this may seem very

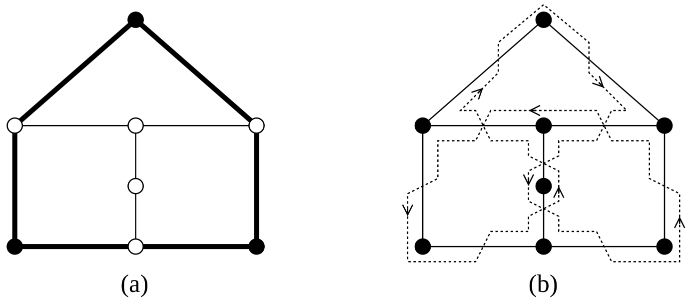

Figure 1.1: (a) a bicycle; (b) a left-right tour

restrictive, locally finite graphs are a very natural class of infinite graphs since many important properties of (and intuitions about) graphs fail completely if we also allow vertices with infinite degree.

The first theorem we will extend is Rosenstiehl and Read's tripartition theorem:

Theorem (Rosenstiehl & Read [24]). *Let e be an edge in a finite graph G. Then exactly one of the following holds:*

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

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

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

One way of regarding this theorem is to consider the statement in terms of the cycle space and the cut space. Then given any edge e in a finite graph G, either $e \in B$ for some $B \in \mathcal{C}(G) \cap \mathcal{C}^*(G)$ (which is what (i) says), or $\{e\} \in \mathcal{C}(G) + \mathcal{C}^*(G)$, but not both. In the latter case, we can then consider the two symmetric subcases that either e lies in some cycle space element, and when we delete e from that, it becomes a cut (this is the statement in (ii)), or else e lies in some cut, and when we delete e from it, this edge set becomes an element of the cycle space (which corresponds to (iii)).

If we now naively use the finite version of the cycle space on an infinite, locally finite graph, every element of the cycle space will necessarily be finite, and we will see that this theorem fails for locally finite graphs. One counterexample is the double ladder, shown in Figure 1.2. Here, no finite bicycle contains the edge e since every element of the cycle space will be a sum (modulo 2) of the induced 4-cycles and hence cannot form a cut at the same time. Similarly, some consideration yields that on the other hand (for the same reason), there is neither a finite Y nor a finite Z as in (ii) or (iii) of the theorem.

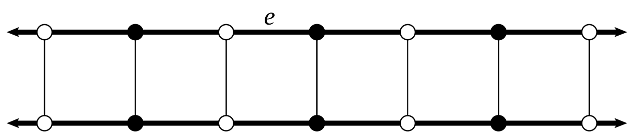

Figure 1.2: There is no finite B, Y or Z as in the theorem for e.

Fortunately, as mentioned earlier, we have a broader definition of the cycle space which allows for infinite structures. Diestel and Kühn [12, 13] provided a definition of cycles that introduces infinite cycles but encompasses the usual finite cycles as well. They defined a *circle* to be the homeomorphic image of the unit circle in the graph compactified by its *ends*. (Where *ends* are equivalence classes of rays. These definitions will be introduced formally in the next chapter.) This definition has proven to be very fruitful, insofar as almost all of the properties of the cycle space in a finite graph remain valid in locally finite graphs. There has also been studied a more general approach by Vella and Richter [26], that covers other compactifications of infinite graphs as well.

In the double ladder, the two double rays (the bold edges in Figure 1.2) together with their two ends are homeomorphic to the unit circle, and hence they form such an infinite cycle. If we now reconsider the tripartition theorem, we see that the set of bold edges now forms an infinite cycle and also constitutes a cut, and hence it forms an infinite bicycle which contains our edge e. The proof of the tripartition theorem for locally finite graphs and lemmas relating to it will be the topic of Chapters 3 and 4.

There is a very interesting connection between bicycles and *left-right tours* in planar graphs, which we will investigate next. If G is a finite plane graph, we can easily obtain a left-right tour of G: start at an arbitrary edge uv, traverse it in one direction, say from u to v, then 'turn left' when we reach v (meaning choose the leftmost edge at v in our embedding), follow it along, then 'turn right' at the next vertex and continue traversing edges and alternately turn left and right until we reach the edge uv again. There we stop, provided we are about to traverse uv again from u to v and provided our turn at v would, again, be a left turn. In this way, we produce a closed walk which is our left-right tour. Depending on the graph, the left-right tour might then contain a certain edge e of G once, twice or not at all. The set of edges that are traversed exactly once is called the *residue* of our left-right tour.

Shank [25] showed that in finite graphs, this residue is always a bicycle. This gives us a particularly easy way to find bicycles in planar graphs.

Theorem (Shank [25])**.** *If G is a finite plane graph, then the residue of a left-right tour is a bicycle.*

We will extend his theorem and prove that this also holds for locally finite graphs in Chapter 5. In order to be able to do this, we will first need to generalize the concept of left-right tours to infinite, locally finite graphs. Figure 1.1 (b) shows a left-right tour in the graph. Looking at the residue of that tour we can see the same bicycle that we observed earlier in Figure 1.1 (a).

Not only do the left-right tours of a graph offer an easy way to find bicycles; they in fact determine *all* bicycles in the graph:

Theorem (Shank [25])**.** *In a finite plane graph the residues of the left-right tours generate the bicycle space.*

In Chapter 6 we will show that this is also true for locally finite graphs. See also Richter and Shank [22] and Lins, Richter and Shank [19].

In finite graphs, Archdeacon, Bonnington and Little [1] used *ladders* (which are certain substructures involving left-right tours and bicycles in an unusual way) to give a criterion for planarity. Planarity criteria are very important tools, and this one in particular is purely algebraical and requires no geometrical knowledge about the graph. In Chapter 7 we will extend this criterion to locally finite graphs.

This planarity criterion becomes particularly intuitive when considering *pedestrian* graphs, which are graphs that do not contain any nonempty bicycles. When considering only finite graphs, we have the following simple characterization.

Theorem (Chen [9]). *A finite graph G is pedestrian if and only if the number of spanning forests of G is odd.*

Since the number of spanning forests in a locally finite graph is usually infinite, this theorem certainly cannot hold for infinite graphs. Unfortunately, neither can we use our usual methods or infinite cycles to simply generalize it in some way. In this last chapter we will show that if a locally finite graph contains only a finite number of spanning forests, then the theorem holds. As yet there is no adequate extension to graphs with an infinite number of spanning forests.

We will see examples of locally finite pedestrian graphs with any given number of ends and any given degree in each end (where the *degree* of an end is the maximum number of disjoint rays it contains). We will even present a plane graph that has a thick end, ie. an end that contains an infinite number of disjoint rays; furthermore, surprisingly, this graph contains no bicycles. Using what we show in Chapter 5, this means that the residue of any left-right tour in this graph must be empty, and hence, the entire edge set is covered twice by only a single left-right tour. At the end of the chapter, we will briefly discuss other possible approaches for finding a characterization of infinite pedestrian graphs.

Chapter 2

Definitions and Preliminaries

All of our graphs are undirected and simple, unless otherwise noted. In general, we follow the notation of [11], which also provides a good introduction to the topological cycle space. A more thorough introduction gives the expository paper [10].

Sometimes we will write $\bigcup A$ to denote the union of all elements of a set A. A graph of the form $K_{1,n}$ is called a *star*. A *normal spanning tree* T of a graph G is a rooted spanning tree of G such that any two vertices that are adjacent in G are comparable in the tree-order of T. We will make use of certain double covers of the edge set of a graph G, thus, let us call a set of walks $\mathcal{W}$ a *double cover* of G if every edge $e \in E(G)$ is traversed exactly twice by walks in $\mathcal{W}$ (ie. either once in two walks or twice in one walk).

In this thesis we will only consider *locally finite* graphs, which are graphs whose vertex degrees are finite. Let G be an infinite, locally finite graph. A one-way infinite path is called a *ray*, and a two-way infinite path is a *double ray*. A subray of a ray is called a *tail* of its ray. We will introduce an equivalence relation on the rays of a graph, where two rays R and S are *equivalent* if there are infinitely many disjoint R–S paths. The equivalence classes of rays are called the *ends* of G. The set of ends of G is denoted by $\Omega(G)$. As an example, the double ladder in Figure 5.2 has two ends, one to the left and one to the right, and the upper and lower double ray both converge to one of these ends on each side. By contrast, the 3-regular tree has uncountably many ends. The *degree of an end* is the maximum number of disjoint rays it contains. An end is said to be *thick* if it contains infinitely many disjoint rays. A *comb* is a graph that is the union of a ray with infinitely many disjoint finite paths each of whose first vertex lies on this ray. The ray is then called the *spine* of the comb, and the last vertices of the paths are its *teeth*.

In order to introduce infinite cycles, we need to view G as a topological space. So let us introduce a topology on G. We will view G as a 1-complex, meaning that every vertex and every end of G is represented by a distinct point, and for every edge e of G we add a set of continuum many points to our topological space. These sets shall be disjoint from each other and from the set $V \cup \Omega$. Every edge shall be homeomorphic to the real interval $(0, 1)$, and we extend this bijection to one from the edge including its endpoints to the interval $[0, 1]$. We denote the topological space thus constructed by $|G|$.

Now, to actually define a topology on $|G|$, let us define a basis of open sets. For every edge, let the open neighborhoods of inner points of this edge correspond to the open neighborhoods around points in $(0, 1)$. For a vertex v and every $\epsilon > 0$, declare as open all open stars around v of radius ϵ, ie. for every edge at v all points on this edge with distance less than ϵ from v (in the metric of the respective edge). Given an end ω of G, let S be a finite set of vertices. Let $C(S, \omega)$ be the component of $G - S$ that contains a ray in ω, then $C(S, \omega)$ contains a subray for every ray in ω. By $\Omega(G, \omega)$ we denote those ends of G which have a ray in $C(S, \omega)$. We write $\hat{C}(S, \omega)$ for the union of $C(S, \omega)$ with $\Omega(G, \omega)$ and with all interior points of edges between $C(S, \omega)$ and S. Then for every such S, the set $\hat{C}(S, \omega)$ shall be an open neighborhood of ω. With the basis thus defined, the open sets of $|G|$ shall be all unions of these sets.

The *closure* of a set X in $|G|$ is denoted by $\overline{X}$. Also note that if H is a subgraph of G, then the ends of H need not correspond to the ends of G. For instance, consider the double ladder from our example in Figure 1.2. The double ladder itself has two ends (one to the left and one to the right), but if we consider the subgraph H of G consisting only of the two (horizontal) double rays, then H has four ends, two for each double ray. Therefore, $\overline{H} \setminus H$ (where, again, $\overline{H}$ is the closure of H *in* G) has two ends, which are the ends of G, not of H.

The space $|G|$ is sometimes also referred to as the *Freudenthal compactification* of G. We note that the topological space $|G|$ is Hausdorff. When G is connected and locally finite, then $|G|$ is compact (cf. [11]).

In the topological space thus obtained, every ray converges to the end it belongs to. Since there is a correspondence between G and $|G|$, we will now translate some graph theoretic concepts to the topological space: The image of a continuous mapping from the unit interval $[0, 1]$ to $|G|$ is called a *topological path*, and the images of 0 and 1 are called its *endpoints*. The homeomorphic image of $[0, 1]$ in $|G|$ is called an *arc*. A *circle* of $|G|$ is a homeomorphic image C in $|G|$ of the unit circle. We call the subgraph $C \cap G$ a *cycle*, and its edge set a *circuit*. Hence, a cycle may be finite or infinite; in the latter case it is a disjoint union of double rays. Since $C \cap G$ is dense in C (cf. [12]), every circle is the closure in $|G|$ of its cycle. Thus, there is a unique correspondence between a circle and its cycle.

The closure $\overline{T}$ in $|G|$ of a subgraph T of G is a *topological spanning tree* of G if it is path-connected and contains all the vertices and ends of G, but no circles of $|G|$. In finite graph theory, for a connected graph G and a spanning tree T of G, we know that for every edge of $E(G) \setminus E(T)$ there exists a unique cycle C_e in $T + e$ which we call a *fundamental cycle*. Similarly, consider a connected, locally finite graph G with a topological spanning tree $\overline{T}$ of G. Then for an edge $e \in E(G) \setminus E(T)$ there exists a unique circle in $\overline{T} \cup e$. Its edge set is called the *fundamental circuit* C_e of e with respect to $\overline{T}$.

As mentioned in Chapter 1, the collection of all subsets of the edge set $E(G)$ of a graph G is the *edge space* of G, which we denote by $\mathcal{E}(G)$. Together with symmetric difference as addition, it forms a vector space over $\mathbb{Z}_2$. In order to introduce the topological cycle space as defined by Diestel and Kühn, we need to allow certain infinite sums as well. So let us call a family $\mathcal{F}$ of subsets of the edge set of a graph *thin* if no edge appears in infinitely many members of $\mathcal{F}$. The *sum* $\sum_{F \in \mathcal{F}} F$ is the set of edges that appear in precisely an odd number of members of $\mathcal{F}$ and hence is well-defined. Whenever we take the sum over an infinite family we will assume this family to be thin (and sometimes refer to the sum as being *thin*) without mentioning it explicitly.

We are now ready to define the *topological cycle space*. For an infinite, locally finite graph G, let its cycle space $\mathcal{C}(G)$ be the set of all thin sums of circuits — where, again, a circuit is the edge set of the subgraph $C \cap G$, and C is a homeomorphic image of S^1 in $|G|$. If G is a finite graph, then this space coincides with the usual cycle space as traditionally defined. We remark that the topological cycle space is closed under the taking of infinite thin sums (cf. [12, 13]).

A *tour* T in $|G|$ is a continuous map $T : S^1 \to |G|$ that is locally injective at every $x \in S^1$ for which $T(x)$ is an interior point of an edge. Note that, therefore, every edge with an interior point in the image of T, denoted by rge T, is completely contained in rge T. We denote the set of all edges that lie in rge T by $E(T)$. The *residue* $\triangle T$ of a tour T is the set of those edges that are traversed exactly once by T.

A *cut* (or *cocycle*) in G is a set of edges $F \subseteq E(G)$ such that either $F = \emptyset$, or there is a set $U \subseteq V(G)$ such that an edge is in F if and only if it has precisely one endvertex in U and one outside of U. In case that $U = \{v\}$, we denote this cut by $E(v)$. Just as the cycle space, the space consisting of all cuts in G forms a subspace of $\mathcal{E}(G)$, which is called the *cut space* and denoted by $\mathcal{C}^*(G)$. The cuts of the form $E(v)$ with $v \in V(G)$ generate $\mathcal{C}^*(G)$. A minimal non-empty cut in G is called a *bond*.

A graph is *plane* if it is drawn in the plane in such a way that the vertices are distinct points, and the intersection of any two edges is precisely their common endpoints. A graph is said to be *planar* if it *can* be drawn in such a way. Let $G = (V, E)$ and $G^* = (V^*, E^*)$ be two plane multigraphs, and let F resp. F^* be the faces of G, resp. G^*. Then we call G and G^* *plane duals* if there exist bijections

$$\begin{array}{lll} F \to V^* & E \to E^* & V \to F^* \\ f \mapsto v^*(f) & e \mapsto e^* & v \mapsto f^*(v) \end{array}$$

such that

(i) $v^*(f) \in f$ for all $f \in F$;

(ii) $|e^* \cap G| = |\mathring{e}^* \cap \mathring{e}| = |G^* \cap e| = 1$ for all $e \in E$; and

(iii) $v \in f^*(v)$ for all $v \in V$.

If G and G^* are not planar, we have the following generalization. We call G^* an *abstract dual* of G if there exists a bijection $E(G^*) \to E(G)$ that maps the circuits of G precisely to the bonds of G^*.

There exists a certain orthogonality between the cycle space and the cut space of a graph. To make this more precise, let us recall that there is a scalar product $*$ defined on $\mathcal{E}(G)$ for a multigraph G as follows: for $X, Y \subseteq E(G)$, we let $X * Y = 0$ if $|X \cap Y|$ is even, and we set $X * Y = 1$ otherwise. With this product, for a set of edge sets $\mathcal{X}$, we can define the orthogonal space $\mathcal{X}^{\perp}{:=}\{Y \subseteq E(G) : Y * X = 0 \text{ for all } X \in \mathcal{X}\}$. For a finite (multi-) graph G, it holds that $\mathcal{C}(G) = \mathcal{C}^{*\perp}(G)$ and $\mathcal{C}^*(G) = \mathcal{C}^{\perp}(G)$ (cf. [11]). At the end of this chapter we will see how this property generalizes to the infinite case.

The intersection of these two spaces is another subspace of $\mathcal{E}(G)$ and is called the *bicycle space* of G. It is denoted by $\mathcal{B}(G){:=}\mathcal{C}(G) \cap \mathcal{C}^*(G)$, and an element of $\mathcal{B}(G)$ is a *bicycle*.[1] A graph that contains no non-empty bicycles is said to be *pedestrian*. When dealing with infinite graphs, we may sometimes wish to refer to only the finite sets in a subspace of $\mathcal{E}(G)$. Therefore, let us denote by $\mathcal{C}_{\text{fin}}(G)$ (and $\mathcal{C}^*_{\text{fin}}(G)$ or $\mathcal{B}_{\text{fin}}(G)$, resp.) the set of all finite edge sets in $\mathcal{C}(G)$ (in $\mathcal{C}^*(G)$ or in $\mathcal{B}(G)$, resp.).

As already mentioned in Chapter 1, the topological cycle space has proven to be the most natural and has allowed for a number of fundamental results involving cycles to be extended to infinite graphs. We mention a few such results now, for use later in this thesis.

A very valuable tool in extending results from finite graphs to infinite

[1] There is a certain inconsistency here. Following Diestel [11], we use "cycle" to denote a *subgraph* stemming from a homeomorphic image of S^1. In particular, a finite cycle is a connected subgraph. On the other hand, a finite bicycle, which is an *edge set*, does not need to span a connected graph.

graphs is the following lemma. We will make use of it in the proof of Theorem 3.2.

Lemma 2.1 (Kőnig's Infinity Lemma).
Let $W_1, W_2, \ldots$ 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 neighbor in W_{n-1}. Then H contains a ray $v_1 v_2 \ldots$ with $v_n \in W_n$ for all n.

For a proof we refer the reader to [11].

The aforementioned orthogonality between the cycle space and the cut space of a finite graph generalizes to infinite graphs as follows:

Theorem 2.2 (Diestel & Kühn [12]).
Let F be a set of edges in a locally finite graph G. Then F is an element of the cycle space of G if and only if it meets every finite cut in an even number of edges.

In the cut space $\mathcal{C}^*(G)$, a result analogous to Theorem 2.2 holds; a proof can be found in [3].

Lemma 2.3 (Bruhn & Diestel [3]).
Let F be a set of edges in a graph G. Then F is a cut in G if and only if it meets every finite circuit in an even number of edges.

We will need another property of the topological cycle space:

Theorem 2.4 (Diestel & Kühn [12]).
Every element of the cycle space of a locally finite graph is the (edge-) disjoint union of circuits.

The following lemma tells us something about the structure in infinite graphs:

Lemma 2.5 (Star-Comb Lemma).
Let U be an infinite set of vertices in a connected graph G. Then G contains either a comb with all teeth in U or a subdivision of an infinite star with all leaves in U.

In Chapters 5 and 6 we will be concerned with plane graphs.

Theorem 2.6 (Kuratowski 1930; Wagner 1937)**.**
A graph G is planar if and only if it contains neither K^5 nor $K_{3,3}$ as a minor.

This also holds, more generally, for countable graphs (cf. [15]).

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

Theorem 2.7 (Richter & Thomassen [23])**.**
Let G be a locally finite 2-connected planar graph. Then $|G|$ embeds in the sphere.

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

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

Chapter 3

The Tripartition Theorem

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

Theorem 3.1 (Rosenstiehl & Read [24]).
Let e be an edge in a finite graph G. Then exactly one of the following holds:

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

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

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

As mentioned already in Chapter 1, one way of regarding this theorem is to consider the statement in terms of the cycle space and the cut space. Then given any edge e in a finite graph G, the edge e either lies in some element from the intersection of the cycle and the cut space, or e is the symmetric difference of an element from the cycle space and a cut, but not both. The latter case then splits into the two symmetric subcases that e either lies in the cycle space element, or in the cut. When we extend this theorem to the infinite case in Theorem 3.2, we will assume this viewpoint again. Now, let us first give the proof for the finite case.

Proof. Assume that there is no bicycle containing e. Recall that for any $X, Y \subseteq E(G)$ we have $X * Y = 0$ if $|X \cap Y|$ is even, and $X * Y = 1$ otherwise.

Therefore, since $\mathcal{C}(G)^\perp = \mathcal{C}^*(G)$, we know that

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

On the other hand, if (i) holds, then $\{e\} \in \mathcal{B}(G)$ and hence, with the same reasoning, $\{e\} \notin \mathcal{B}(G)^\perp = \mathcal{C}^*(G) + \mathcal{C}(G)$ and therefore neither (ii) nor (iii) holds.

Finally, assume that both (ii) and (iii) hold and that there exist such Y and Z. Then $e = Y + Y + e$ and $e = Z + Z + e$. Thus, $Y + Y + e = Z + Z + e$ which gives $Y + Z = Y + e + Z + e$. Since $Y + Z \in \mathcal{C}(G)$ and $Y + e + Z + e = Y + Z \in \mathcal{C}^*(G)$ it follows that $Y + Z \in \mathcal{B}(G)$. But since $e \in Y + Z$ and since by assumption there is no bicycle that contains e, it follows that $Y + Z = \emptyset$ and hence $Y = Z$. But this is a contradiction, since $e \in Y$ and $e \notin Z$.

□

We have seen in Chapter 1 that Theorem 3.1 fails for locally finite graphs when we consider only the definition of the cycle space stemming from finite graph theory. Hence, we use the cycle space as defined by Diestel and Kühn which is applicable to both finite and infinite graphs.

In trying to apply this same proof to infinite graphs we encounter a problem with the definition of the scalar product. In a finite graph G, for any finite subsets $X, Y \subseteq E(G)$ we have $X * Y = 0$ if $|X \cap Y|$ is even, and $X * Y = 1$ otherwise. In a locally finite graph, however, X and Y need not be finite, so what should the value of $X * Y$ be if the edge sets X and Y have an infinite intersection? Fortunately, we will be able to circumvent this issue by only using the scalar product for those $X, Y \in \mathcal{E}(G)$ with $|X \cap Y| < \infty$. A proper concept for orthogonal spaces appears to be more difficult, as however defined they seem to lose a number of their usual properties. For this reason, we will make do without them here. We remark that, these problems notwithstanding, Casteels and Richter [8] introduce orthogonal spaces in infinite graphs that still retain many of the usual properties.

Remembering that $\mathcal{C}_{\text{fin}}(G)$ (and $\mathcal{C}^*_{\text{fin}}(G)$ or $\mathcal{B}_{\text{fin}}(G)$, resp.) denotes the set of all finite edge sets in $\mathcal{C}(G)$ (in $\mathcal{C}^*(G)$ or in $\mathcal{B}(G)$, resp.) let us now state the tripartition theorem for locally finite graphs.

Theorem 3.2. [5]
Let e be an edge of a locally finite graph G. Then either

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

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

but not both.

Proof. For the proof we will use Kőnig's Infinity Lemma (Lemma 2.1 from Chapter 2). To do so, we need to define suitable sets W_i. The idea is to have sets W_i such that a bicycle in W_{n+1} induces a bicycle in the smaller set W_n. In order to achieve this, we will construct two different sequences of graphs. For all n, the graphs G_n will simulate the cuts in G, and the graphs $\widetilde{G}_n$ will on the other hand simulate the cycle space elements of G. For an edge set to be an element of W_n then, it will be necessary to be in the corresponding (cut, resp. cycle) spaces of *both* graphs.

Now, let us start with our construction. We may assume G to be connected and therefore countable. For each $n \in \mathbb{N}$, denote by S_n the set of the first $n+1$ vertices in some fixed enumeration of the vertices of G that starts with the endvertices of our given edge e. We define G_n to be the graph $G[S_n]$ induced by these first $n+1$ vertices, together with the edges in $E(S_n, V(G) \setminus S_n)$ and their incident vertices. Let $\widetilde{G}_n$ be the minor of G obtained by contracting the components of $G - S_n$ (where we keep parallel edges but delete loops). Note that $E(G_n) = E(\widetilde{G}_n)$, and that by our choice of S_n, we have that $e \in E(G_n) = E(\widetilde{G}_n)$ for every n. By our construction, it holds that $\mathcal{C}^*(G_n) \subseteq \mathcal{C}^*(G_{n+1})$ and $\mathcal{C}(\widetilde{G}_n) \subseteq \mathcal{C}(\widetilde{G}_{n+1})$. Also observe that $\mathcal{C}(G_n) \subseteq \mathcal{C}(\widetilde{G}_n)$ and $\mathcal{C}^*(\widetilde{G}_n) \subseteq \mathcal{C}^*(G_n)$ hold for all n. Now we let $W_n := \{B \in \mathcal{C}^*(G_n) \cap \mathcal{C}(\widetilde{G}_n) : e \in B\}$.

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

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

and hence (ii) holds.

Second, assume $W_n \neq \emptyset$ for all n. We will first show that for each $K \in \mathcal{C}^*(G_{n+1})$ it holds that $K \cap E(G_n) \in \mathcal{C}^*(G_n)$. Assume that this is not the case. Let $K' := K \cap E(G_n) \notin \mathcal{C}^*(G_n)$. Since $W_n \neq \emptyset$ for all n we know that there exists some finite circuit $Z \in \mathcal{C}(\widetilde{G}_n)$ (recall that, by Theorem 2.4, every element of the cycle space is the disjoint union of circuits). Since $K' \notin \mathcal{C}^*(G_n)$ we know by our observation above that also $K' \notin \mathcal{C}^*(\widetilde{G}_n)$. By Lemma 2.3 it then follows that $|Z \cap K'|$ is odd. Since $Z \subseteq E(G_n) = E(\widetilde{G}_n)$, we have $Z \cap K' = Z \cap K \cap E(G_n) = Z \cap K$ and therefore, $|Z \cap K|$ is also odd. Since $Z \in \mathcal{C}(\widetilde{G}_n)$, again applying Lemma 2.3 yields $K \notin C^*(\widetilde{G}_n) \subseteq C^*(G_n) \subseteq C^*(G_{n+1})$, which is a contradiction.

Similarly, let $Z \in \mathcal{C}(\widetilde{G}_{n+1})$. We will show that the restriction $Z' := Z \cap E(\widetilde{G}_n)$ lies in $\mathcal{C}(\widetilde{G}_n)$. Assume that this is not the case. Since for all n we know that $W_n \neq \emptyset$, there exists some finite cut $K \in \mathcal{C}^*(G_n)$. Again, since $Z' \notin \mathcal{C}(\widetilde{G}_n)$ we use our observation above to see that therefore, $Z' \notin \mathcal{C}(G_n)$. Applying Theorem 2.2 we obtain that $|Z' \cap K|$ is odd. Because $K \subseteq E(G_n) = E(\widetilde{G}_n)$, we have that $Z' \cap K = Z \cap E(\widetilde{G}_n) \cap K = Z \cap K$ and hence $|Z \cap K|$ is also odd. Since $K \in \mathcal{C}^*(G_n)$ applying Theorem 2.2 once more gives $Z \notin \mathcal{C}(G_n) \subseteq \mathcal{C}(\widetilde{G}_n) \subseteq C(\widetilde{G}_{n+1})$, which is a contradiction.

Thus we have shown that the restriction of an element from the cut space of G_{n+1} to $E(G_n)$ lies again in the cut space of G_n, and similarly the restriction of an element from the cycle space of $\widetilde{G}_{n+1}$ to $E(G_n)$ lies again in the cycle space of $\widetilde{G}_n$. Hence, it follows that for every $B \in W_{n+1}$ we have $B \cap E(G_n) \in W_n$. Now, let us define a graph on $\bigcup_{n=1}^{\infty} W_n$ such that $B \in W_{n+1}$ is adjacent to $B' \in W_n$ if and only if $B \cap E(G_n) = B'$. Thus, the conditions for Lemma 2.1 are satisfied, and we obtain for each $n \in \mathbb{N}$ a $B_n \in W_n$ such that $B_{n+1} \cap E(G_n) = B_n$ for all n. Then (by definition of the W_n) the union $B := \bigcup_{n\in\mathbb{N}} B_n$ contains e.

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

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

where the last equality follows since $B_N \in \mathcal{C}(\widetilde{G}_N)$. Since F was chosen arbitrarily, Theorem 2.2 implies that $B \in \mathcal{C}(G)$. In a similar way, let Z be a finite circuit of G. Now let $N \in \mathbb{N}$ be large enough such that $Z \subseteq E(G_N)$. Then it follows that Z is also a circuit in G_N. We conclude

$$B * Z = B * (Z \cap E(G_N)) = (B \cap E(G_N)) * Z = B_N * Z = 0,$$

where the last equality follows from $B_N \in \mathcal{C}^*(G_N)$. Lemma 2.3 then implies that $B \in \mathcal{C}^*(G)$. Hence, we have seen that $B \in \mathcal{C}(G) \cap \mathcal{C}^*(G)$ and therefore, $B \in \mathcal{B}(G)$ and (i) holds.

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

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

which is a contradiction. □

The reader will have noticed that the theorem only divides the edges into two classes rather than three. We will address this issue at the end of this chapter.

Casteels and Richter [8] independently proved a complementary result:

Theorem 3.3 (Casteels & Richter [8]).
Let e be an edge of a locally finite graph G. Then either

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

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

but not both.

It should be noted that Casteels and Richter in fact prove a more general result of which Theorem 3.3 is but a consequence.

Theorems 3.2 and 3.3 look interestingly similar, the difference being whether infinite bicycles and only a finite sum of the cycle- and the cut space are allowed in the theorem, or vice versa. The next lemma gives us a better understanding of their relation.

Lemma 3.4. [5]
Let G be a locally finite graph. If for an edge e of G any two of the following conditions hold, then the remaining condition holds as well:

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

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

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

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

The lemma is reminiscent of a theorem by Richter and Shank [22] about finite surface duals. In fact, our proof uses similar arguments. We mention, moreover, that all of (i)–(iii) can hold for an edge. In Figure 1.2 we have already seen that e lies in an infinite bicycle, while in Figure 3.1 we witness the other two cases.

Proof of Lemma 3.4. First, assume (iii) and one of either (i) or (ii) to hold. Then, there exists a $B \in \mathcal{B}(G)$ with $e \in B$ and an $X \in \mathcal{C}(G)$ so that $X + e \in \mathcal{C}^*(G)$. Since $X + B \in \mathcal{C}(G)$, we know that if X contains e (and hence satisfies (i)), then $X + B$ satisfies (ii); and if, on the other hand, X does not contain e (and satisfies (ii)), then $X + B$ satisfies (i).

Secondly, assume that (i) and (ii) both hold, and let Y be as in (i) and Z be as in (ii). Then $B{:=}Y + Z \in \mathcal{C}(G)$ since $Y, Z \in \mathcal{C}(G)$. From the fact that

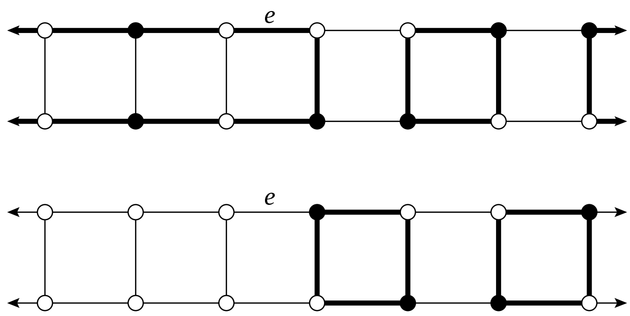

Figure 3.1: (i) and (ii) in Lemma 3.4 both hold for e

$B = (Y + e) + (Z + e)$ it follows that B is also a cut. Finally, since $e \in Y$ but $e \notin Z$, we know that $e \in B$.

For the second part of the lemma, assume that (i)–(iii) all hold for e, and let $e \in B \in \mathcal{B}(G)$. By (the trivial part of) Theorem 3.3, it follows that B cannot be finite. On the other hand, Y and Z as in (i) resp. (ii) need to be infinite sets, too, since otherwise this would yield a contradiction to Theorem 3.2. □

Rosenstiehl and Read's theorem partitions the edges of a finite graph into three classes. So far, our theorem yields only two classes. So, let us refine Theorem 3.2. For this, we say that an edge e in a locally finite graph G is of *cut-type* if there is a finite cut K containing e so that $K \setminus \{e\} \in \mathcal{C}(G)$. We say that e is of *cycle-type* if there is a finite element Z of the cycle space with $e \in Z$ and $Z \setminus \{e\} \in \mathcal{C}^*(G)$. (Note that although we use the term *cycle-type* and Z is an element of the cycle space, it need not be a circuit itself.) Now the following immediate corollary of Lemma 3.4 turns Theorem 3.2 into a true tripartition theorem:

Corollary 3.5. [5]
No edge in a locally finite graph can be of cut-type and of cycle-type at the same time.

We should point out that it is possibly a bit misleading to denote the set of all cuts in G by $\mathcal{C}^*(G)$, since it might give the impression that it is the dual space of $\mathcal{C}(G)$. That, however, is not the case. Rather, Theorem 2.2

shows that, at least in some sense, $\mathcal{C}(G)$ and $\mathcal{C}^*_{\mathrm{fin}}(G)$ are dual to each other. On the other hand, the dual space of $\mathcal{C}^*(G)$ is $\mathcal{C}_{\mathrm{inf}}(G)$, see for instance [3].

In this respect, our bicycle space $\mathcal{B}(G)$ is situated between these two dualities. The graph in Figure 1.2, among other examples, indicates that this is nevertheless justified since in order to make the tripartition theorem work for infinite graphs, whether it is in the form of Theorem 3.2 or in the form of Theorem 3.3, we need both spaces, $\mathcal{C}(G)$ and $\mathcal{C}^*(G)$.

Chapter 4

Principal Cuts

At the end of the last chapter, we saw that if in a locally finite graph G there is no bicycle that contains an edge e, then e is either of cycle- or of cut-type; in which case there exists by definition a $Z \in \mathcal{C}_{\text{fin}}(G)$ so that $Z + e \in \mathcal{C}^*(G)$. We call Z a *principal cycle* of e and $Z + e$ a *principal cut* of e. In this chapter, we demonstrate how the properties of principal cuts translate from finite graphs to locally finite graphs.

If G is a pedestrian graph, ie. a graph for which $\mathcal{B}(G) = \{\emptyset\}$, then every edge must be either of cycle- or of cut-type. In this case, we have the noteworthy property that the principal cuts in G are unique. To see this, let K and K' be two principal cuts for e. Hence, K and $K' \in \mathcal{C}^*(G)$ with $K + e$ and $K' + e \in \mathcal{C}(G)$. Thus, $K + K' = K + K' + e + e = (K + e) + (K' + e) \in \mathcal{C}(G) \cap \mathcal{C}^*(G) = \mathcal{B}(G)$, which implies that $K = K'$, since $\mathcal{B}(G) = \{\emptyset\}$. Therefore, given a pedestrian graph, we denote the principal cut of an edge e by K_e and the principal cycle by Z_e.

We need the following lemma, whose statement and proof are direct extensions of the finite case which appears in Rosenstiehl and Read [24].

Lemma 4.1. [5]
Let e and f be edges in a locally finite pedestrian graph G. Then:

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

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

Proof. To prove (i) consider

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

This holds using the fact that $\mathcal{C}_{\text{fin}}(G) = \mathcal{C}_{\text{fin}}^{*\perp}(G)$ and $\mathcal{C}^*_{\text{fin}}(G) = \mathcal{C}^{\perp}_{\text{fin}}(G)$. Assertion (ii) is shown analogously:

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

Note that all these scalar products are well-defined since the Z_e and K_e are finite sets.

□

We note that Lemma 4.1 remains true in non-pedestrian graphs; Z_e (resp. Z_f) is then simply *any* principal cycle through e (resp. f), since there is no longer a unique one. And similarly, this also holds for K_e and K_f.

Proposition 4.2. [5]
In a locally finite pedestrian graph G both of the families $(Z_e)_{e\in E(G)}$ and $(K_e)_{e\in E(G)}$ are thin.

Proof. Recall that a family of subsets of, for example, $E(G)$ is *thin* when no edge appears in infinitely many of its members. So suppose that this is not the case, ie. there is an edge e lying in infinitely many Z_f. Since G is a pedestrian graph, e must be of cycle- or of cut-type and Z_e is therefore defined. Thus, Lemma 4.1 implies that also $f \in Z_e$ for all these infinitely many f, contradicting the fact that Z_e is finite. It follows that $(Z_e)_{e\in E(G)}$ is thin.

Similarly, assume that the family of principal cuts is not thin. Then again, we have an edge e that is part of infinitely many K_f. Since e must be either of cycle- or of cut-type, we know that K_e is defined. With Lemma 4.1 it then follows that $f \in K_e$ for all these infinitely many f, which contradicts the fact that K_e is finite. Hence, $(K_e)_{e \in E(G)}$ is thin.

□

For an edge e to be of cycle- or of cut-type we have required that there is a *finite* $Z \in \mathcal{C}(G)$ with $Z + e \in \mathcal{C}^*(G)$. In light of Theorem 3.3, one could also very reasonably relax this requirement and say that an edge is of cycle- or of cut-type if there is any such Z, be it finite or infinite. A pedestrian graph, then, would be one without any *finite* bicycles, since this is precisely the case when all edges are of cycle- or of cut-type.

There are several problems with this definition. We have already seen (in Figures 1.2 and 3.1) that this would not give a proper tripartition. Furthermore, principal cuts in a pedestrian graph would not necessarily be unique and their family may not be thin. For instance, the cuts of the type indicated in the lower graph in Figure 3.1 would form a non-thin family of principal cuts.

The following corollary extends some basic properties of principal cycles and cuts. The proofs of these results for finite graphs use the finiteness substantially only in one point, namely that it is allowed to take arbitrary sums of principal cuts. While this is never an issue in finite graphs, such sums may be infinite in infinite graphs and then need to be thin in order to be well-defined. But with Proposition 4.2 this is the case, and we can extend these properties to locally finite graphs:

Corollary 4.3. [5]
Let G be a locally finite pedestrian graph. Then

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

(ii) *the union of all cut-type edges is a cut; and*

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

(iv) $(K_e)_{e \in E(G)}$ *generates the cut space.*

Proof. For (i) and (ii), we follow Godsil and Royle [16] closely.

For showing (i), we know that, since G is pedestrian, every edge e of G must be either of cycle- or of cut-type, and there exist unique $Z_e \in \mathcal{C}(G)$ and $K_e \in \mathcal{C}^*(G)$ such that $e = Z_e + K_e$.

Let $Z{:=}\sum_{e\in E(G)} Z_e$ and $K{:=}\sum_{e\in E(G)} K_e$ (which are well-defined by Proposition 4.2). Clearly, we know that $Z \in \mathcal{C}(G)$ and $K \in \mathcal{C}^*(G)$. Then it holds that
$E(G) = \sum_{e\in E(G)} e = Z + K$. We will use this fact in the following equation:

$$\begin{aligned} Z_e * \{e\} &= Z_e * \{Z_e + K_e\} = Z_e * Z_e + Z_e * K_e = Z_e * Z_e = Z_e * E(G) \\ &= Z_e * (Z + K) = Z_e * Z + Z_e * K = Z_e * Z = (e + K_e) * Z \\ &= \{e\} * Z + K_e * Z = \{e\} * Z, \end{aligned}$$

where $Z_e * K = 0$ follows from applying Lemma 2.3 to the finite cycle space element Z_e and $K \in \mathcal{C}^*(G)$. Similarly, $K_e * Z = 0$ holds by Theorem 2.2 for the finite cut K_e and $Z \in \mathcal{C}(G)$.

Since $Z_e * \{e\} = 1$ if and only if e is of cycle-type, it follows also that $\{e\} * Z = 1$ if and only if e is of cycle-type, which means that the elements of Z are exactly the edges of cycle-type, and hence their union is an element of the cycle space. We note that all of the above scalar products are well-defined, since the sets K_e and Z_e are finite.

Similarly, for (ii), we can deduce that

$$\begin{aligned} K_e * \{e\} &= K_e * \{Z_e + K_e\} = K_e * Z_e + K_e * K_e = K_e * K_e = K_e * E(G) \\ &= K_e * (Z + K) = K_e * Z + K_e * K = K_e * K = (e + Z_e) * K \\ &= \{e\} * K + Z_e * K = \{e\} * K, \end{aligned}$$

where again $K_e * Z = 0$ and $Z_e * K = 0$ by the above argument. Here too, $K_e * \{e\} = 1$ if and only if e is of cut-type, and hence $\{e\} * K = 1$ if and only if e is of cut-type, which means that the elements of K are exactly the edges of cut-type, and thus, their union is a cut.

The proof of the finite case for (iii) and (iv) can be found in Rosenstiehl and Read [24].

To see (iii), let $Z_A := \sum_{e\in A} Z_e$ and $K_A := \sum_{e\in A} K_e$. Then Z_A and K_A are well-defined since the families $(Z_e)_{e\in E(G)}$ and $(K_e)_{e\in E(G)}$ are thin. Thus,

$$Z_A + K_A := \sum_{e\in A} Z_e + \sum_{e\in A} K_e = \sum_{e\in A} (Z_e + K_e) = \sum_{e\in A} e = A$$

where all sums are thin and hence well-defined. From (i) and (ii) we know that $Z_A \in \mathcal{C}(G)$ and $K_A \in \mathcal{C}^*(G)$. Now, let any nonempty $C \in \mathcal{C}(G)$ be given. Then $C = Z_C + K_C$ which yields $C + Z_C = K_C$. Since C and Z_C both lie in $\mathcal{C}(G)$, and $\mathcal{C}(G)$ is closed under taking thin sums, we know that also $C + Z_C \in \mathcal{C}(G)$. Since, on the other hand, $K_C \in \mathcal{C}^*(G)$, from $\mathcal{B} = \{\emptyset\}$ it follows that $K_C = \emptyset$ and therefore $C = Z_C = \sum_{e\in C} Z_e$. Hence C is a sum of principal cycles Z_e, and thus $(Z_e)_{e\in E(G)}$ generates the cycle space.

In a similar way, to show (iv), let F be a nonempty cut in G. Then $F = Z_F + K_F$, and hence $F + K_F = Z_F$. Since F and K_F are both cuts in G, and $Z_F \in \mathcal{C}(G)$, we have that $Z_F = \emptyset$ and thus $F = K_F = \sum_{e\in F} K_e$. Therefore F is a sum of principal cuts K_e, and so $(K_e)_{e\in E(G)}$ generates the cut space.

$\square$

Chapter 5

Left-Right Tours

In Chapter 1 we already introduced left-right tours in finite graphs, albeit in an intuitive sense. Given a finite plane graph G, we start at any edge $e = uv$, traverse it in the direction of one of its endpoints, say v, 'turn left' at v (meaning we choose the leftmost edge at v), traverse along that edge, 'turn right' at the end of that edge and continue like this, alternately turning left and right. We stop when we reach an edge that we have already traversed and when we are about to traverse it again in the same direction and would make the same (left or right) turn again at its endpoint. That way, we obtain a closed walk in G.

How can we sensibly generalize this to infinite plane graphs? The challenge is that we have to avoid getting 'lost' in an end. In locally finite graphs, we do not necessarily return to an edge we have already traversed in the same direction and with the same left/right parity. We could traverse infinitely many edges, steering towards an end. This is why we will first introduce *left-right strings*, meaning we pick an arbitrary edge and traverse it in the familiar left-right fashion. We do the same for the other direction from our starting edge and maximally extend both sides. We will see that this gives either a closed and finite left-right tour, or a two-way infinite left-right string. Two examples are illustrated in Figure 5.2.

In general, the two ends of a left-right string will not be identical, and hence they do not form a closed walk. So in order to obtain something resembling a left-right tour in a locally finite graph, we will glue together several left-right strings to form a closed walk. This shall give us a topological tour in $|G|$ which is therefore closed, and locally it has the property of being 'left-right'. Unfortunately, we will generally have more than two left-right strings per end. So which ones should we pick to construct our infinite left-right tour?

Another way of viewing a left-right string is that after traversing an even number of edges from our starting point, we turn right, and if, on the other hand, we have traversed an odd number of edges, we will turn left. This kind of parity information gets lost in the ends. So are there certain left-right strings that fit together while others do not?

Now let us introduce these concepts formally. In order to define left-right strings, we need to describe what it means to do a 'left' turn followed by a 'right' turn. We will follow the treatment in [18] and [19].

Let G be a locally finite graph, and let $|G|$ be embedded in the sphere S. Recall that by Theorem 2.7, every locally finite planar graph has such an embedding. The interior of an edge of G is homeomorphic to the open unit interval $(0, 1)$. For each edge e, we fix a homeomorphism. Then, without loss of generality, let η_1 denote the image of the restriction of this homeomorphism to $(0, \frac{1}{2})$ and let η_2 be the image of the restriction to $(\frac{1}{2}, 1)$. Let us call η_1 and η_2 the *halves* of e. We will use the notation $\overline{\eta_1} = \eta_2$ and $\overline{\eta_2} = \eta_1$ to switch back and forth between the two halves of an edge. Furthermore, we fix for e two open, disjoint and connected subsets, σ_1 and σ_2, of $S \setminus |G|$ each of which has e in its boundary. We call these the *sides* of e, and as for the halves, we put $\overline{\sigma_1} = \sigma_2$ and $\overline{\sigma_2} = \sigma_1$. Now, a triple (e, η, σ), where $e \in E(G)$, η is a half of e, and σ is a side of e, is called a *corner* of $|G|$. We say that $c = (e, \eta, \sigma)$ is a *corner at* e, and it is a *corner at* $v \in V(G)$ if the boundary $\partial\eta$ contains v. Clearly, for each edge e there are four corners at e.

For every $v \in V(G)$ we choose an open disc D around v, so that each half of an edge at v intersects ∂D in exactly one point. Then ∂D defines in a natural way a rotation of the halves. We say that two corners (e, η, σ) and (e', η', σ') at v are *matched* if η and η' appear consecutively in the local rotation at v, and if the connected component K of $\sigma \cap D$ with $\eta \cap D \subseteq \partial K$ and the connected component K' of $\sigma' \cap D$ with $\eta' \cap D \subseteq \partial K'$ are contained in the same connected component of $D \setminus |G|$. We note that this definition is independent of the actual choice of D. Figure 5.1 gives an illustration.

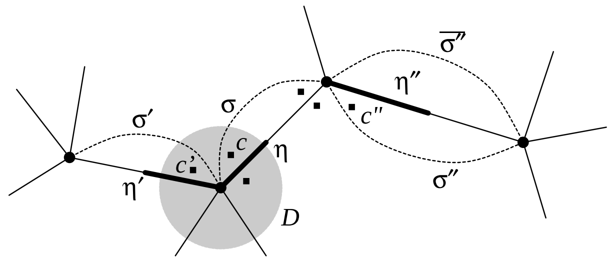

Figure 5.1: We think of a corner $c = (e, \eta, \sigma)$ at $v \in V(G)$ as a point close to v and η, and lying in σ. The corners c and $c' = (e', \eta', \sigma')$ are matched; the corners c and c'' describe a left-right step.

The reason for introducing corners is that they are very suitable when trying to describe a left-right step. So, now we are ready to define some objects with a left-right structure.

Let $W = \dots (e_{-1}, \eta_{-1}, \sigma_{-1}), (e_0, \eta_0, \sigma_0), (e_1, \eta_1, \sigma_1) \dots$ be a (finite, one-way infinite or two-way infinite) sequence of corners satisfying the following properties:

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

(ii) no corner appears twice in W.

Then such a sequence W is called a *left-right walk*, which is justified by the fact that the edges $\dots e_{-1} e_0 e_1 \dots$ do indeed form a walk. Moreover, we will sometimes pretend that a left-right walk is in fact a walk, ie. a sequence of vertices and edges, rather than a sequence of corners. As an example, the corners c and c'' in Figure 5.1 describe a left-right step as in (i).

We say that S is a *left-right string* (LRS for short) if it is a maximal left-right walk. It is not hard to check that if $S = \dots (e_{-1}, \eta_{-1}, \sigma_{-1}), (e_0, \eta_0, \sigma_0), (e_1, \eta_1, \sigma_1) \dots$ then $S' {:=} \dots (e_1, \overline{\eta_1}, \overline{\sigma_1}), (e_0, \overline{\eta_0}, \overline{\sigma_0}), (e_{-1}, \overline{\eta_{-1}}, \overline{\sigma_{-1}}) \dots$ is an LRS, too. In fact, the walks S and S' traverse the same edges, but in opposite directions. Although we will sometimes view S as an oriented walk, we will, in general, not distinguish between S and S' and consider them to be identical. This slight abuse of notation ensures that just as in the finite case, every edge is covered exactly twice by LRSs, as we will see in the next lemma. Figure 5.2 gives an example of two different LRSs in the double ladder.

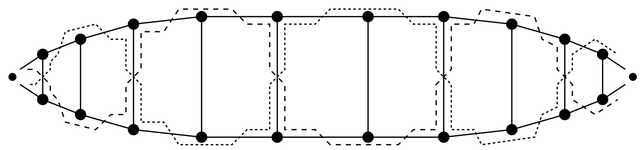

Figure 5.2: Two LRSs in the double ladder

Let us say that a set of walks $\mathcal{W}$ is a *double cover* of G if every edge $e \in E(G)$ is traversed exactly twice by walks in $\mathcal{W}$ (ie. either once in two walks or twice in one walk).

Lemma 5.1. [5]
For a locally finite graph G, let $|G|$ be embedded in the sphere. Then:

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

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

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

Proof. For (i), let us assume this is not the case and we have an LRS S with two corners $c_0 = (e_0, \eta_0, \sigma_0), d_0 = (f_0, \theta_0, \tau_0)$ such that c_0 and d_0 are matched. Then, by the definition of an LRS, there exists a corner $c_1 = (e_1, \eta_1, \sigma_1)$ of S such that $(e_0, \overline{\eta_0}, \overline{\sigma_0})$ and (e_1, η_1, σ_1) are matched, and similarly there is $d_{-1} = (f_{-1}, \theta_{-1}, \tau_{-1})$ of S such that $(f_0, \overline{\theta_0}, \overline{\tau_0})$ and $(f_{-1}, \theta_{-1}, \tau_{-1})$ are also matched. Inductively we obtain sequences of corners $\{c_i\}_i$ and $\{d_i\}_i$ with $i \in \mathbb{Z}$ that each describe left-right steps. Since c and d are both corners of S, then so are $\{c_i\}_i$ and $\{d_i\}_i$ for all i. Thus, since every successive pair or corners in an LRS describes a left-right step, there exist some i, j such that (without loss of generality) d_j is the successor of c_i. Hence, $(e_i, \overline{\eta_i}, \overline{\sigma_i})$ and $(f_j, \theta_j, \tau_j)$are matched. But $(e_i, \overline{\eta_i}, \overline{\sigma_i})$ is already matched with $(e_{i+1}, \eta_{i+1}, \sigma_{i+1})$. Thus, $c_i = d_j$ and inductively this holds for any appropriate c_k, d_l. It follows that $c_0 = d_0$ which contradicts the definition of left-right walk.

To see (ii), we assume that we have an LRS $S = \dots c_{-1}c_0c_1 \dots c_n$ that is not a two-way infinite walk. Then we have (at least) one 'last' corner $c_n = (e_n, \eta_n, \sigma_n) \in S$ (meaning that c_n has no successor). Since an LRS is maximal by definition, $(e_n, \overline{\eta_n}, \overline{\sigma_n})$ is either already matched with some $c_j \in S$, or $(e_n, \overline{\eta_n}, \overline{\sigma_n}) = c_k \in S$ for some k. Let us consider the first case, ie. the corners $(e_n, \overline{\eta_n}, \overline{\sigma_n})$ and c_j are matched. Then $c_n = c_{j-1}$ by definition, and it follows inductively that $c_{n-i} = c_{j-1-i}$. Therefore, S is closed and hence finite. Else, let $(e_n, \overline{\eta_n}, \overline{\sigma_n}) = c_k \in S$ for some k. If S is not closed, c_k has a successor c_{k+1} in S, and c_{k+1} and c_n are also matched, which are both corners of S. But by (i), no two corners in S can be matched, thus we obtain a contradiction.

To show (iii), we enumerate the edges of the graph G (which is locally finite and hence countable). To obtain a double cover, let us start with the first edge and start an LRS in any corner at this edge. As long as some edge of G is not covered twice yet, let us always pick the smallest such edge (in our enumeration), pick an uncovered corner at e and start a new LRS there. The set of LRSs we obtain this way clearly covers every edge at least twice. Every edge has four corners, two of which always induce the same LRS by definition (since if one of them is contained in an LRS, then the other one is matched to the predecessor or successor of it). Thus, every edge can occur in at most two LRSs. Altogether, we know that every edge occurs in exactly two LRSs (which may not necessarily be distinct). Therefore, the set of all LRSs forms a double cover of G.

$\square$

Let us note that since we used corners to define LRSs and LRTs, the statement (iii) also holds in degenerate cases. Consider for example a double ray. Then there exist exactly two (different) LRSs (forming an LRT), and they form a double cover of the edge set. We consider these LRSs distinct, since their corner sequences are distinct. But they both traverse the double ray from one end to the other and are indistinguishable as sequences of edges. Simply to avoid confusion, let us remark that this is not the same case as discussed earlier, where an LRS S' is considered identical to S. There, the corner sequences differed, but only because they described the same LRS being traversed in two different directions. Two of the four corners at an edge always induce the same LRS in different directions. In our example of the double ladder, however, we obtain two LRSs each of which is induced (independently) by two corners at an edge, but these corners differ for the two distinct LRSs.

Now that we have formalized left-right structures, let us consider tours in locally finite graphs. As mentioned in Chapter 2, in a locally finite graph G (not necessarily planar), a *tour* T in $|G|$ is a continuous map $T : S^1 \to |G|$ that is locally injective at every $x \in S^1$ for which $T(x)$ is an interior point of an edge. We remark that therefore, every edge with an interior point in the image of T, which we denote by rge T, is completely contained in rge T. We denote the set of all edges that lie in rge T by $E(T)$. This requirement is sensible and necessary because we now operate in the topological space $|G|$ as opposed to the graph G. Hence, we need to ensure that a tour cannot 'turn around' in the middle of an edge. The *residue* $\triangle T$ of a tour T is the symmetric difference of the edges it contains, and is therefore the set of those edges that are traversed exactly once by T.

At long last we are able to extend the definition of left-right tours to infinite graphs. Assume that $|G|$ is embedded in the sphere. Our aim is to define an LRT as a set of LRSs that are glued together at ends such that they constitute a tour in $|G|$. An example would be the two LRSs shown in Figure 5.2 together with the two ends of the double ladder.

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

Having defined LRTs, we are now able to extend the next theorem to locally finite graphs, ie. we will show that the residue of an LRT is always a bicycle. In finite graphs, this is due to Shank:

Theorem 5.2 (Shank [25]).
If G is a finite plane graph, then the residue of a left-right tour is a bicycle.

In the finite case, Theorem 5.2 is proven using plane duals. Unfortunately, a suitable theory of plane duality involving also infinite cycles has yet to be formulated. Therefore we will circumvent this obstacle by reducing the problem to finite graphs. The main construction for this lies in the proof of the following lemma.

The idea is that given a locally finite plane graph G with a set of LRTs L_i, and given a finite subgraph H of G, we will look at the pieces of those LRTs in H and try to connect these finite pieces using only finitely many vertices and edges (outside of H). We want to do this in such a way that the obtained finite LRTs L'_i simulate the infinite LRTs of G, meaning that locally, on H, the infinite L_i and the assembled finite pieces L'_i will behave in the same way.

Lemma 5.3. [5]
For a locally finite graph G, let $|G|$ be embedded in the sphere. Let $L_1, \dots, L_k$ be a set of LRTs of G so that no LRS of G lies in more than one L_i, and let H be a finite plane subgraph of G. Then there exist a finite plane supergraph H' of H and a set $L'_1, \dots, L'_k$ of LRTs of H', so that the LRT L_i traverses precisely the edges $e_1, \dots, e_n$ of H and in this order if and only if L'_i does, for all $i = 1, \dots, k$.

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

Since we only worry about the LRTs, let us denote by F those edges in the cut $E(H, G - H)$ that lie in some L_i, and find in the one face that contains

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

Now consider an LRT L, and let $\mathcal{S}$ be the set of those LRSs that lie in L (here, of the two orientations of an LRS $S \in \mathcal{S}$, we pick the one that is induced by L). In order to properly connect up the finite pieces of L that lie in H, we need to consider the respective corners. So let $\mathcal{K}_L$ be the set of corners of L, or, more precisely, $\mathcal{K}_L := \bigcup_{S \in \mathcal{S}} S$. We will also need to remember in which order the LRSs $S \in \mathcal{S}$ were traversed in L, but since L induces a cyclic ordering on its LRSs, it also does so on $\mathcal{K}_L$. Furthermore, we let $\mathcal{M}$ be those of the corners in $\bigcup_{i=1}^{k} \mathcal{K}_{L_i}$ that are corners at edges in F. Then each corner in $\mathcal{M}$, which is a corner in G, corresponds to a corner in H_0. For the sake of simplicity, we will not distinguish between these two and, depending on the context, view $\mathcal{M}$ as a set of corners either in G or in H_0. Corners in $\mathcal{M}$ come in two kinds: either they are *outgoing* corners, ie. corners at vertices in $V(H)$, or they are *ingoing* corners, ie. corners at vertices in X.

When we construct the finite subgraph that connects up the finite pieces of our LRTs L_i in H, we will need to know which piece to connect with which other one. Our aim is to connect those pieces that lie in the same LRT L_i in G, and to connect them in the same order as they appear in L_i. So let us construct a pairing of the corners in $\mathcal{M}$. For each i, we arbitrarily pick an outgoing corner c_1 in $\mathcal{M} \cap \mathcal{K}_{L_i}$. Then, let $c_1, \ldots, c_l$ be the corners in $\mathcal{M} \cap \mathcal{K}_{L_i}$ in the cyclic order of $\mathcal{K}_{L_i}$ (which is induced by L_i). We need to pair up consecutive corners, but we need to make sure that we do this in a way that will replace the *infinite* part of L_i with a finite subgraph, as opposed to replacing the part of L_i that is contained in H. We know that since L_i is a tour, l is even and for each odd j the corner c_j is outgoing, while the corner c_{j+1} is ingoing. Hence, we need to start with c_1 and then pair up consecutive corners: $\{c_1, c_2\}, \ldots, \{c_{l-1}, c_l\} \in \mathcal{P}$. For later use, we note that

$$\textit{if } \{c, c'\} \in \mathcal{P}\textit{, then one of } c, c' \textit{ is outgoing and one is ingoing.} \tag{5.1}$$

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

For every corner $c \in \mathcal{M}$ we will construct a sequence of left-right walks $K^j(c)$. To begin, let $K^0(c) := (c)$ for every c, so $K^0(c)$ is a walk of length one, which traverses an edge in F. To simplify the construction in the next steps we will, with the help of a suitable homeomorphism, identify D with $(0,3) \times (0,1) \subseteq \mathbb{R}^2$, where all the vertices in X are assumed to lie in the segment $\{0\} \times (0,1)$; see Figure 5.3.

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

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

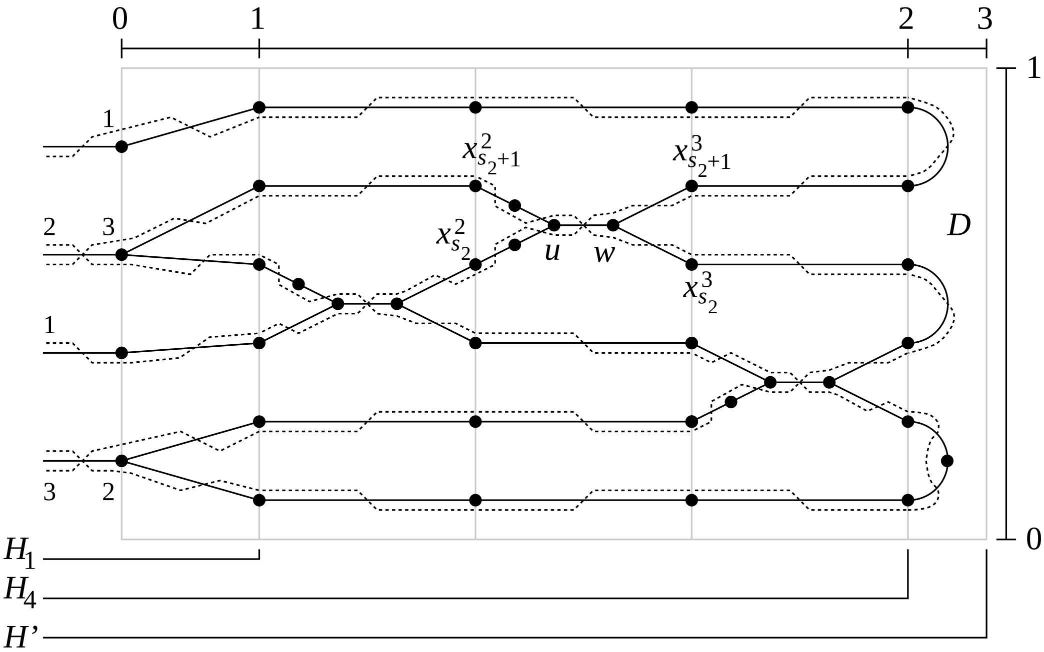

Figure 5.3: The construction of the H_i (not to scale). Corners with the same number are supposed to be paired.

The idea now is to proceed as follows: We would like to further extend the walks $K^1(c)$ in such a way that, eventually, they will connect corners c, c' that belong to the same pair in $\mathcal{P}$. Since there may be other corners between c and c' in $\{1\} \times (0,1)$, we will need to permute the vertices $x_1^1, \ldots, x_m^1$. We will achieve this with a sequence of transpositions, which we will call flips, such that after t steps, in $\{\tilde{t}\} \times (0,1)$ (for some $1 \leq \tilde{t} \leq 2$) every vertex x_k^{t+1} (which will be defined in a moment) will lie next to the vertex its paired corner belongs to. For this, we will construct a sequence of finite graphs $H_2, H_3, \ldots, H_{t+1}, H'$ that will allow us to make these flips and still remain planar and, even more, that will later allow us to connect the pieces of an LRT L_i in H such that the resulting walk is still left-right.

So, let us construct finitely many supergraphs H_i of H_1 with corresponding left-right walks $K^i(c) \supseteq K^1(c)$ for $c \in \mathcal{M}$. These supergraphs $H_1 \subset H_2 \subset \ldots \subset H_{t+1}$ will be nested and plane, and such that $H_i \setminus H_{i-1}$ is entirely drawn in $(a, b] \times (0,1)$ for some $1 \leq a < b < 3$ (we will determine the respective a and b in a moment). The intersection of H_i with $\{b\} \times (0,1)$ will consist of

m vertices. In the order we encounter them on $\{b\} \times (0,1)$ going from $(b,0)$ to $(b,1)$ we will denote these by $x^i_1, \dots, x^i_m$. For each $j = 1, \dots, m$ there will then be a unique corner $p^i_j \in \mathcal{M}$ so that the left-right walk $K^i(p^i_j)$ either starts or ends in x^i_j (and is otherwise disjoint from $x^i_1, \dots, x^i_m$).

Let $(p_1, \dots, p_m)$ be a permutation of $\mathcal{M}$. For the rest of the proof let us call a *flip* at $s \in \{1, \dots, m-1\}$ the operation that turns $(p_1, \dots, p_m)$ into $(p_1, \dots, p_{s-1}, p_{s+1}, p_s, p_{s+2}, \dots, p_m)$. Clearly, for some t there is a sequence of t flips at $s_1, \dots, s_t$ that turns $(p^1_1, \dots, p^1_m)$ into $(q_1, \dots, q_m)$ such that for each odd $j \in \{1, \dots, m\}$ it holds that $\{q_j, q_{j+1}\} \in \mathcal{P}$.

Our goal now is to define H_{i+1}, for $i \in \{1, \dots, t\}$, in such a way that $(p^{i+1}_1, \dots, p^{i+1}_m)$ is obtained from $(p^i_1, \dots, p^i_m)$ by performing a flip at s_i. Moreover, with the exception of the points $x^{i+1}_1, \dots, x^{i+1}_m$, we will draw $H_{i+1} \setminus H_i$ in $(1 + \frac{i-1}{t}, 1 + \frac{i}{t}) \times (0,1)$ for all $i > 1$. Let us assume $H_1, \dots, H_i$ to be already constructed. We put m distinct vertices $x^{i+1}_1, \dots, x^{i+1}_m$ (in this order) on the segment $\{1 + \frac{i}{t}\} \times (0,1)$. First, for each $j \in \{1, \dots, m\}$ with $j \neq s_i, s_i + 1$, draw a straight line between x^i_j and x^{i+1}_j. We thus extend $K^i(p^i_j)$ to a left-right walk $K^{i+1}(p^i_j)$ along the edge $x^i_j x^{i+1}_j$. Now, consider the vertices $x^i_{s_i}$ and $x^i_{s_i+1}$. We draw an edge uw in $(1 + \frac{i-1}{t}, 1 + \frac{i}{t}) \times (0,1)$ so that no crossing edges arise when we connect u to $x^i_{s_i}$ and $x^i_{s_i+1}$, and w to $x^{i+1}_{s_i}$ and $x^{i+1}_{s_i+1}$. If necessary (meaning, if the left-right walk would be on the 'wrong' side of $x^i_{s_i} u$ in order to traverse uw), we subdivide the edge $x^i_{s_i} u$ to guarantee the existence of a left-right walk from $x^i_{s_i}$ through uw to $x^{i+1}_{s_i+1}$ (that is disjoint from $x^i_{s_i+1}$). Now we can extend $K^i(p^i_{s_i})$ by this walk to a left-right walk $K^{i+1}(p^i_{s_i})$, and we proceed in an analogous way for $K^i(p^i_{s_i+1})$. This ensures that $(p^{i+1}_1, \dots, p^{i+1}_m)$ is indeed obtained from $(p^i_1, \dots, p^i_m)$ by performing a flip at s_i.

Finally, completing our construction, let all the H_i up to H_{t+1} be defined. Then every two vertices x^{t+1}_j and x^{t+1}_{j+1} for odd j belong to corners forming a pair in $\mathcal{P}$. So it only remains to connect them properly. For each odd j in $\{1, \dots, m\}$, we draw an edge in $(2,3) \times (0,1)$ that joins x^{t+1}_j to x^{t+1}_{j+1}. Subdividing $x^{t+1}_j x^{t+1}_{j+1}$ if necessary, we can join $K^{t+1}(p^{t+1}_j)$ by this (possibly subdivided) edge to $K^{t+1}(p^{t+1}_{j+1})$, so that the resulting walk $K^{t+1}(p^{t+1}_j) K^{t+1}(p^{t+1}_{j+1})$

is left-right (here, (5.1) ensures that the corner sequences fit with respect to orientation).

Now, let H' be the graph H_{t+1} joined with the possibly subdivided edges in $(2,3) \times (0,1)$. Then, by construction of the pairing $\mathcal{P}$, it is ensured that the resulting LRTs L'_i in the plane graph H' are indeed left-right tours, and behave on H in the same way as the L_i do. □

The following lemma lists some properties of our above construction. We will not need these properties until later, in Chapter 7, but since their proofs involve an understanding of the details of the proof of Lemma 5.3, we show them here.

Lemma 5.4.
For a locally finite graph G, let $|G|$ be embedded in the sphere. Let a set $L_1, \dots, L_k$ of LRTs of G, a finite plane subgraph H of G, a finite plane supergraph H' of H and a set $L'_1, \dots, L'_k$ of LRTs of H' be given as in Lemma 5.3. Then the following properties hold:

(i) *The set $\{L'_1, \dots, L'_k\}$ of LRTs of H' extends to a double cover $\{L'_1, \dots, L'_k, \dots, L'_l\}$ of LRTs of H' such that it still holds that the LRT L_i traverses precisely the edges $e_1, \dots, e_n$ of H and in this order if and only if L'_i does, for all $i = 1, \dots, k$.*

(ii) *If C is a cycle of G, and if $E(C) \subseteq E(H)$, then C is also a cycle of H'.*

(iii) *If F is a cut in G, and if $E(F) \subseteq E(H)$, then F is also a cut in H'.*

Proof. To see (i), let us construct a double cover of LRTs of H'. We start with the given LRTs $\mathcal{L}' := L'_1, \dots, L'_k$, and as long as any edge of H' is not covered twice by LRTs in $\mathcal{L}'$ yet, we start a new LRT on the uncovered side of that edge and add it to the set $\mathcal{L}'$. Since H' is finite, this is easily done, and the property that every LRS lies in at most one LRT is trivially preserved. Hence the set $\{L'_1, \dots, L'_k, \dots, L'_l\}$ of LRTs is a superset of the original LRTs $\mathcal{L}'$, and clearly it still holds that the LRT L_i traverses precisely the edges $e_1, \dots, e_n$ of H and in this order if and only if L'_i does, for all $i = 1, \dots, k$.

For (ii) it can easily be seen that if C is a cycle of G with $E(C) \subseteq E(H)$, then C is also a cycle of H. Since $H \subseteq H'$, it follows that C is a cycle in H' as well.

To show (iii), let a (necessarily finite) cut F in G be given such that $E(F) \subseteq E(H)$. Since F is a cut in G, it induces a partition $V(G) = A \mathbin{\dot\cup} B$ on its vertex set. Since H is a subgraph of G, it holds that F is a cut in H as well. For the purpose of contradiction, let us now assume that F is not a cut in H'. Thus, there exists a (necessarily finite) A–B path P in H' that does not meet F. Since F is also a cut in H, we know that P is not completely contained in H. Hence, there exists a subpath P' of P with $E(P') \subseteq E(H')$. Let P' be inclusion-maximal with this property. Now let a', b' be the first and last vertices of P', and let a and b be their neighboring vertices on P, ie. we have that $a, b \in V(H)$.

Since P' is completely contained in $H' \setminus H$, it arose from the construction in the proof of Lemma 5.3. Thus, there are corners c_a, c'_b at P such that c_a is a corner at the vertex a and c'_b is a corner at the vertex b' or vice versa, ie. c_a is a corner at a' and c'_b is a corner at b, and furthermore, the corners c_a, c'_b constitute a pair in the pairing $\mathcal{P}$ from the proof of Lemma 5.3. Without loss of generality, let c_a be outgoing, and c'_b be ingoing; thus $\{c_a, c'_b\} \in \mathcal{P}$. Let $c'_b = (e, \eta, \sigma)$ and let c_b be the corner at the vertex b that is matched with $(e, \overline{\eta}, \overline{\sigma})$. Then c_a and c_b are corners at some common LRT L in G, and between their occurrences, the LRT L is disjoint from H.

Thus, there exists some (topological) a–b path Q in G that does not meet H, except for its endpoints. Hence, $E(Q) \cap E(H) = \emptyset$. Since $F \subseteq E(H)$, we know that $E(Q) \cap F = \emptyset$ as well. Therefore, Q is an a–b path in G that does not meet the cut F, which is a contradiction. Hence, F is a cut in H', which completes the proof.

□

With the help of Lemma 5.3 we are now ready to prove Shank's theorem for locally finite graphs.

Theorem 5.5. [5]
For a locally finite graph G, let $|G|$ be embedded in the sphere. Then the residue of a left-right tour in $|G|$ is a bicycle.

Proof. Let L be an LRT in $|G|$. In order to prove that $\triangle L$ is a bicycle, we have to show that it lies in both the cycle space and the cut space of G. So, let us first show that $\triangle L \in \mathcal{C}(G)$. For this, let F be a finite cut in G. As a tour, L passes an even number of times through F. Therefore, $|\triangle L \cap F|$ is even and it follows by Theorem 2.2 that $\triangle L$ is an element of the cycle space.

To see that $\triangle L$ is also a cut, consider a finite cycle C in G. Lemma 5.3 (with $H = C$) then yields a finite plane supergraph H' of C and an LRT L' of H' so that $\triangle L \cap E(C) = \triangle L' \cap E(C)$. From Theorem 5.2 it follows that $\triangle L'$ is a cut in H', and since $C \subseteq H'$ is a cycle, we know that $|\triangle L' \cap E(C)|$ is even. Hence, $|\triangle L \cap E(C)| = |\triangle L' \cap E(C)|$ is also even, and it therefore follows from Lemma 2.3 that $\triangle L \in \mathcal{C}^*(G)$, and hence $\triangle L \in \mathcal{B}(G)$. □

Chapter 6

LRTs Generate the Bicycle Space

In Chapter 5 we saw an important connection between bicycles and left-right tours; the residue of an LRT is always a bicycle. But what about the other bicycles in a graph? Is there a way to characterize all of them? Indeed, the interaction between bicycles and LRTs continues. We will see that the LRTs determine *all* bicycles of a graph, and they do so in the anticipated way: Their residues generate the bicycle space. Thus, in this chapter we prove the analogue of the following theorem for locally finite graphs.

Theorem 6.1 (Shank [25]).
In a finite plane graph the residues of the left-right tours generate the bicycle space.

Let us consider a locally finite graph G for which $|G|$ is embedded in the plane. Now let B be any bicycle of G. Since B is, in particular, a cut in G, and since the cut space of G is generated by all cuts of the form $E(v)$ with $v \in V(G)$, there is a vertex set $X \subseteq V(G)$ such that $B = \sum_{x \in X} E(x)$. On the other hand, B is also an element of the cycle space. As for finite graphs, $\mathcal{C}(G)$ is generated by the residues of the face boundaries (this is shown in [6]). Thus, there is a set F of face boundaries such that $B = \sum_{f \in F} \triangle f$.

Now, for each bicycle B assume such a pair X, F to be fixed. Following Richter and Shank [22], we say that an LRS S is of *type I* if there is a corner $c = (e, \eta, \sigma)$ in S for which the following statements are either both true or both false:

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

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

Otherwise, if exactly one of the statements is true and one of them is false, we say that S is of *type II*.

We will show why this definition is sensible. Let S be an LRS of type I, and let $c = (e, \eta, \sigma)$ be the corresponding corner. Let $c' = (e', \eta', \sigma')$ be the successor of c, ie. $(e, \overline{\eta}, \overline{\sigma})$ and (e', η', σ') are matched. Now assume that both (i) and (ii) hold for c. Hence, c is a corner at some vertex $x \in X$ and lies in some face f with boundary in F. Let c' be a corner at some $y \in V(G)$ and which lies in some face g.

We first consider the case $e \in B$. Then, since $x \in X$, we know that $y \notin X$ since otherwise $\sum_{x \in X} E(x) = B$ would not contain e. Similarly, since f has a boundary in F, the boundary of the face g cannot be in F, since otherwise $\sum_{f \in F} \triangle f = B$ would not contain e. Thus, neither (i) nor (ii) hold for the corner c'. On the other hand, assume $e \notin B$. Then since $x \in X$ and $\sum_{x \in X} E(x) = B$, it follows that $y \in X$. And analogously, since f has its boundary in F, the boundary of the face g also has to be in F, since otherwise $\sum_{f \in F} \triangle f = B$ would contain e. Thus, in this case, both of (i) and (ii) hold for the corner c'.

Similarly, if neither (i) nor (ii) hold for c, then $x \notin X$ and the boundary of f is not in F. If $e \in B$, it follows that $y \in X$ and the boundary of g lies in F and therefore both (i) and (ii) hold for c'. Otherwise, if $e \notin B$, we have that $y \notin X$ and the boundary of g is not in F and hence, neither (i) nor (ii) hold for c'.

So inductively it follows that if both of (i) and (ii), or neither (i) nor (ii), hold for some corner in S, then this is true for every corner in S. Therefore if exactly one of (i) and (ii) hold for a corner in S, then the same is true for every corner in S. We have seen now that the classification of LRSs in type I and type II does not depend on the corner to which the statements (i) and (ii) refer (it does, however, depend on the bicycle). Compare also with Richter and Shank [22].

Lemma 6.2. [5]
Let G be a locally finite plane graph, and let B be a bicycle. Then an edge e of G lies in B if and only if it lies in exactly one LRS of type I and in one LRS of type II with respect to B.

Proof. The proof is the same as for the finite case, which is given in Richter and Shank [22].

First, let us assume that $e \in B$. Then let η be the half of e such that $\partial\eta$ contains a vertex in X, and let σ be the side of e that is incident with some face whose boundary lies in F. The two corners $c = (e, \eta, \sigma)$ and $c' = (e, \eta, \overline{\sigma})$ induce the two different LRSs containing e. Then c induces an LRS of type I, while c' induces an LRS of type II.

For the reverse implication, assume that $e \notin B$. Let η be any half of e, and let σ be any side of e. Then (i) holds for η if and only if it holds for $\overline{\eta}$ and analogously, (ii) holds for σ if and only if it holds for $\overline{\sigma}$. Hence for any two corners at e that induce the two different LRSs containing e, say $c = (e, \eta, \sigma)$ and $c' = (e, \eta, \overline{\sigma})$, it follows that c satisfies (i) if and only if c' satisfies (i), and similarly, c satisfies (ii) if and only if c' satisfies (ii). Hence, c and c' always induce LRSs of the same type, and thus e lies in two LRSs of type I or in two LRSs of type II. $\square$

In finite graphs, Lemma 6.2 is already enough to prove Theorem 6.1: we only need to sum up the residues of all the LRSs (which are identical to LRTs in finite graphs) of type I (or type II, for that matter). Then an edge e is in a bicycle if and only if it occurs in exactly one LRS of type I, which holds if and only if e is in this sum. Hence, for finite graphs, every bicycle is a sum of residues of LRTs.

In infinite graphs, however, things become a little more complicated. Since an LRT is formed by a set of LRSs, it is not clear what its type should be since the types of the contained LRSs could differ. If this is not the case, we have the following definition. An LRT L is called *B-uniform* if every two LRSs contained in L are of the same type. Even worse, it is not clear that in a locally finite graph there even exists a single B-uniform LRT, let alone a set of B-uniform LRTs with the properties as in the last lemma. Fortunately, though, we are able to show the existence of B-uniform LRTs:

Lemma 6.3. [5]
Let G be a locally finite graph, let $|G|$ be embedded in the sphere, and let B be a bicycle of G. Then there exists a set $\mathcal{L}$ of B-uniform LRTs so that each LRS of G is contained in exactly one $L \in \mathcal{L}$.

Proof. First, let us describe the main steps of this proof. Given a graph G with the set of all LRSs of G, we will construct a graph G' (which may not and need not be planar) by 'blowing up' the graph G; we will duplicate vertices and edges of G in order to disperse the LRSs S_i of G. Thereby we obtain vertex-disjoint walks S'_i in G' corresponding to the LRSs S_i in G. We will define a function ϕ that will map vertices and edges of G' to the corresponding ones in G. In particular, every walk S'_i gets mapped to the LRS S_i.

Next, we will classify the walks S'_i in G according to the type of their corresponding LRS S_i in G. Using this, we will show that for each type, the union of all $E(S'_i)$ in G' corresponding to LRSs of the same type forms an element of the cycle space of G'. Hence, this set decomposes into a set of edge-disjoint circuits in G'. We will see that each such circuit D contains a walk S'_i either entirely or not at all, and that, moreover, all these S'_i contained in D correspond to LRSs S_i that are all of the same type.

Since D is a circuit, there exists a homeomorphism $\sigma_D : S^1 \to |G'|$, and the composition of maps $\phi' \circ \sigma_D$ (where ϕ' is the extension of ϕ to the respective topological spaces) gives us a B-uniform LRT in $|G|$. Since this holds for every such D, we obtain a set of B-uniform LRTs in $|G|$.

We now begin with the rigorous proof. We may assume G to be connected. Hence, there is an enumeration $S_1, S_2, \ldots$ of all LRSs of G, since G is countable. To simplify matters, in this proof we view an LRS as its induced sequence of vertices and edges, rather than working with its corner sequence.

Now we construct from G another locally finite graph G' which, in all likelihood, will not be planar. Let p be a subwalk of the form $p = evf$ in some S_i in G, where $v \in V(G)$ and $e, f \in E(G)$. For each vertex v and each such subwalk p in each S_i, we call v_p a *clone* of v. Then let the vertex set of G' consist of all these clones. The edge set of G' is comprised of two disjoint sets, E' and F'. The set F' contains one edge between each pair of clones v_p and v_q of the same vertex $v \in V(G)$; ie. the clones of a vertex span a complete graph. Two clones u_p and v_q of distinct vertices $u, v \in V(G)$ are connected by an edge in E' if p and q are subwalks in the same LRS S_i and appear consecutively in S_i, ie. for $S_i = \ldots e_{-1}v_{-1}e_0v_0e_1v_1e_2\ldots$ we have $p = e_{j-1}v_{j-1}e_j$ and $q = e_jv_je_{j+1}$ (or vice versa) for some j. See Figure 6.1 for an illustration.

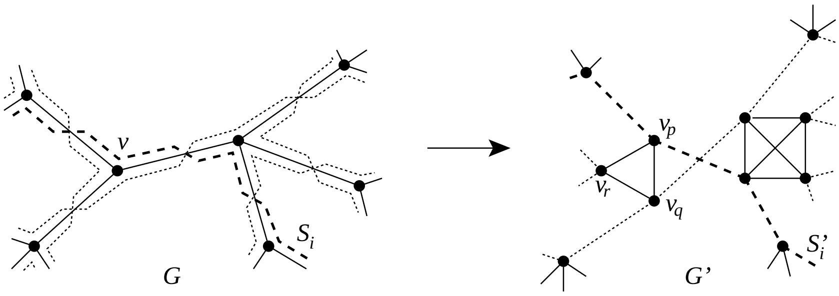

Figure 6.1: The construction of G' in the proof of Lemma 6.3; the edges in F' are solid, while the edges in E' are dotted.

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

We note, furthermore, that because the set of all LRSs of G forms a double cover (cf. Lemma 5.1 (iii)), it holds that

each $e \in E(G)$ has exactly two preimages under ϕ, and these are in E'. (6.1)

For each LRS $S_i = \ldots e_{-1}v_{-1}e_0v_0e_1v_1e_2\ldots$, the map ϕ defines a distinct walk in G'. Indeed, since there is a unique vertex v_{p_j} in G' for each subwalk $p_j{:=}e_jv_je_{j+1}$, and since each v_{p_j} is linked by a unique edge e'_{j+1} in E' to

$v_{p_{j+1}}$, the sequence $\ldots e'_{-1} v_{p_{-1}} e'_0 v_{p_0} e'_1 v_{p_1} e'_2 \ldots$ is a walk in G', which we will denote by S'_i. We claim that for all i it holds that

(i) if $S'_i = \ldots e'_{-1} v'_{-1} e'_0 v'_0 e'_1 v'_1 e'_2 \ldots$ with $e'_j \in E'$ for all j, then $S_i = \ldots \phi(e'_{-1})\phi(v'_{-1})\phi(e'_0)\phi(v'_0)\phi(e'_1)\phi(v'_1)\phi(e'_2)\ldots$; and

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

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

Claim (i) is clear by construction and our argument above. To see (ii) and (iii), note that in G', a clone v_p of a vertex $v \in V(G)$ is adjacent to exactly two vertices that are not clones of v. Hence, for (ii) we have that in S'_i every vertex has degree two. For (iii), on the other hand, we have different vertices v_p for different $v \in V(G)$, and every v_p is incident with exactly two edges from E', and both of these are its incident edges in S'_i.

For our next step, let us denote by $\mathcal{X}_I$ the set of all those S'_i for which S_i is of type I with respect to B, and let $\mathcal{X}_{II}$ be the set of the other S'_i, ie. those for which S_i is of type II. We will show that

$$\textit{both of } X_I := \bigcup_{S' \in \mathcal{X}_I} E(S') \textit{ and } X_{II} := \bigcup_{S' \in \mathcal{X}_{II}} E(S') \textit{ lie in } \mathcal{C}(G'). \tag{6.2}$$

To see that $X_I \in \mathcal{C}(G')$, consider a finite cut K' of G'. Then by Theorem 2.2, it suffices to prove that $|X_I \cap K'|$ is even.

We fix a vertex a' of G' and consider all finite cuts of the form $L = E_{G'}(A, B)$ in G' with $a' \in A$. For each such cut L, let $c(L)$ denote the number of vertices $w' \in B$ such that there exists a clone $u' \in A$ of the same vertex as w'. Since by definition each such w' is adjacent to a vertex in A, the number $c(L)$ is finite (since the cuts L are finite).

Now, among all cuts L for which $|L \cap X_I|$ has the same parity as $|K' \cap X_I|$, choose one, say K, such that $c(K)$ is minimal.

Suppose that $c(K) > 0$, and let $K = E_{G'}(A, B)$ with $a' \in A$. Since $c(K) > 0$ there exist $u' \in A$ and $w' \in B$ that are clones of the same vertex $v \in V(G)$.

As $w' = v_p$ for some subwalk p in some S_i, we obtain from (iii) that w' lies in exactly one S'_i, which implies that w' is incident with exactly zero or two edges in X_I, depending on whether S_i is of type II or of type I. Thus, the cut $\widetilde{K}{:=}K + E(w')$ meets X_I in an even number of edges if and only if $|K \cap X_I|$ is even. On the other hand, we have $\widetilde{K} = E_{G'}(A \cup \{w'\}, B \setminus \{w'\})$, which implies $c(\widetilde{K}) < c(K)$, but this contradicts the choice of K.

Therefore, it holds that $c(K) = 0$. Since all clones of a vertex are on the same side of K, it follows that $K \subseteq E'$ and that $\phi(K)$ is a finite cut in G. It also implies that for each $e \in \phi(K)$, both of the preimages of e under ϕ lie in K: if $e = xy$ for $x, y \in V(G)$, then all clones of x must be in one class of the partition induced by K, while all clones of y must lie in the other (since $c(K) = 0$). Hence, any other edge in G' between the two complete subgraphs induced by the clones of x and y, resp., has to lie in K as well.

Thus, if we can show that $\phi(K)$ is traversed an even number of times by LRSs of type I (with respect to B), then we know that also the cut K in G' is met by an even number of the corresponding walks S'_i. Hence $|X_I \cap K|$ is even, and because $|X_I \cap K| = |X_I \cap K'|$ by our choice of K, the cut $|X_I \cap K'|$ is also even.

Since the LRSs in G form a double cover (see Lemma 5.1 (iii)), any cut in G is met evenly by LRSs. However, from Lemma 6.2 we know that an edge e of G lies in B if and only if it lies in one LRS of type I and in one LRS of type II. Hence, every edge $e \notin B$ is traversed twice by LRSs, which are either both of type I or both of type II (in particular, e is traversed an even number of times by LRSs of the same type). Therefore, the set $\phi(K) \setminus B$ is traversed an even number of times by LRSs of type I. On the other hand, since B is an element of the cycle space, the set $B \cap \phi(K)$ has even cardinality by Theorem 2.2. Again applying Lemma 6.2, we know that any edge $e \in B$ is traversed exactly once by an LRS of type I. Thus, the even cut $B \cap \phi(K)$ is traversed an even number of times by LRSs of type I. So altogether we know that both $\phi(K) \setminus B$ and $B \cap \phi(K)$ are traversed an even number of times by LRSs of type I. Therefore, with (6.1), we obtain that $|X_I \cap K|$ is even. The proof for X_{II} is analogous.

Now that we have shown that both X_I and X_{II} are elements of $\mathcal{C}(G')$, it is clear that $X_I + X_{II}$ is as well. Hence we can use Theorem 2.4 to decompose $X_I + X_{II}$ into a set $\mathcal{D}$ of (edge-) disjoint circuits. Let us observe that

$$\text{for all } i \text{ and } D \in \mathcal{D} \text{ it holds that if } E(S_i') \cap D \neq \emptyset \text{ then } E(S_i') \subseteq D. \text{ Moreover, for each } D \in \mathcal{D}, \text{ all the } S_i \text{ with } E(S_i') \subseteq D \text{ are of the same type.} \tag{6.3}$$

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

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

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

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

To define ϕ' on ends, consider a ray R' in an end ω' of G'. Then $\phi(R')$ is a one-way infinite walk, and thus, by Lemma 2.5, it contains a ray in some end, say ω. We map ω' to ω.

It remains to check that ϕ' is continuous at ends. So let us consider an end ω' of G', and let a basic open neighborhood $C{:=}\hat{C}_G(U, \phi'(\omega'))$ of $\phi'(\omega')$ in $|G|$ be given (recall from Chapter 2 that U is a finite vertex set). By U' we denote the set of all clones of vertices in U. Then U' is also a finite vertex set in G', and we see that $C'{:=}\hat{C}_{G'}(U', \omega')$ is a basic open neighborhood of ω' in $|G'|$ with $\phi'(C') \subseteq C$. Therefore, ϕ' is continuous.

Finally, since each $D \in \mathcal{D}$ is a circuit, by definition there exists a homeomorphism $\sigma_D : S^1 \to |G'|$ with image $\overline{D}$. By the definition of $\mathcal{D}$ and by (b), the continuous mapping $\phi' \circ \sigma_D : S^1 \to |G|$ is locally injective at points $x \in S^1$ that are mapped to interior points of edges. Furthermore, (i) and (a) imply that each maximal subwalk in $\phi' \circ \sigma_D$ is an LRS, and that these are precisely those S_i for which $E(S_i') \subseteq D$. Therefore each $\phi' \circ \sigma_D$ describes an LRT in $|G|$. By (6.3), we know that each such LRT is B-uniform. Let us denote the set $\{\phi' \circ \sigma_D : D \in \mathcal{D}\}$ of LRTs by $\mathcal{L}$.

To conclude our proof, let us observe that since for every S_i the set $E(S_i')$ is contained in some $D \in \mathcal{D}$, every S_i occurs in one of the LRTs in $\mathcal{L}$. On the other hand, since all the $D \in \mathcal{D}$ are (edge-) disjoint, no S_i appears in two elements of $\mathcal{L}$. Hence, $\mathcal{L}$ is a set of B-uniform LRTs in $|G|$ with the desired property. □

We remark that the LRTs in $\mathcal{L}$ have an additional property, of which we will, however, make no use: each $L \in \mathcal{L}$ is *minimal* in the sense that, if L' is an LRT with $\emptyset \neq E(L') \subseteq E(L)$ then $E(L') = E(L)$. In order to briefly sketch the proof, let $D \in \mathcal{D}$ be the circuit in G' so that $\phi' \circ \sigma_D$ describes the LRT L. Let $\mathcal{Y}$ be the subset of LRSs contained in L that also lie in L'. Since L' is also an LRT, it is easy to check that $Y := \bigcup_{S \in \mathcal{Y}} E(S')$ is an element of the cycle space of G'. Since Y is nonempty and is a subset of the circuit D, it follows that $Y = D$, which implies $E(L) = E(L')$ as claimed.

With Lemma 6.3 we can extend Theorem 6.1 to locally finite graphs using arguments of Richter and Shank [22].

Theorem 6.4. [5]
Let G be a locally finite graph, and let $|G|$ be embedded in the sphere. Then the residues of the left-right tours in $|G|$ generate the bicycle space of G.

Proof. In Chapter 5 we have seen that, by Theorem 5.5, the residue of an LRT in $|G|$ is a bicycle. Since $\mathcal{B} = \mathcal{C}(G) \cap \mathcal{C}^*(G)$ is a subspace of $\mathcal{E}(G)$, it follows that all sums of residues of LRTs are also elements of the bicycle space. (Note that these sums are necessarily thin by Lemma 5.1 (iii).)

On the other hand, let a bicycle B in G be given. Then Lemma 6.2 tells us that an edge e is in B if and only if e occurs in exactly one LRS of type I. Hence, this holds if and only if $e \in \sum_{S \text{ is of type I}} \triangle S$.

Lemma 6.3 gives us a set $\mathcal{L}$ of B-uniform LRTs, so that every LRS appears in exactly one $L \in \mathcal{L}$. From this we obtain a set $\mathcal{M} \subseteq \mathcal{L}$ of B-uniform LRTs of type I (meaning that the contained LRSs are all of type I), so that every LRS of type I appears in exactly one $M \in \mathcal{M}$. Therefore e lies in B if and only if $e \in \sum_{M \in \mathcal{M}} \triangle M$. Thus $B = \sum_{M \in \mathcal{M}} \triangle M$, and we see that B is a sum of residues of LRTs in $|G|$.

□

Let us note one more difference between LRTs in finite graphs and in infinite graphs. We know that in a finite graph, the set of LRTs forms a double cover. For infinite graphs, however, let us consider the double ladder. There we have four different LRSs (which are shifts of the two LRSs we witness in Figure 5.2), and we construct different LRTs by gluing together any two of these four. Hence, the double ladder has a set of six LRTs, which clearly covers all edges more than twice. Moreover, while Lemma 6.3 asserts that there are double covers consisting of LRTs, none of these double covers is, for our purposes, sufficient and hence distinguished: In our example of the double ladder, none of the double covers suffices to generate the entire bicycle space of the graph.

Indeed, consider a double cover $\mathcal{L}$ of LRTs for the double ladder. Pick an LRT of the double cover and observe that it traverses some edge e twice (in Figure 5.2 this is the case for every second rung). Hence, e does not lie in the residue of any $L \in \mathcal{L}$, and therefore no bicycle containing e can be expressed as the sum of residues of $L \in \mathcal{L}$. It is easy to check that every edge in the double ladder lies in some bicycle, and since we can find such an edge e and a bicycle containing e for every double cover of LRTs of the double ladder, we do indeed need *all* of the LRTs of the double ladder in order to generate its bicycle space.

Chapter 7

The ABL Planarity Criterion

Planarity is a very important concept in graph theory, and therefore planarity criteria are especially valuable. The most famous planarity criterion is certainly the one known as Kuratowski's theorem (Theorem 2.6 from Chapter 2), which says that a finite graph is planar if and only if it contains neither K^5 nor $K_{3,3}$ as a minor. Another well-known planarity criterion is that of MacLane [21] which states that a finite graph is planar if and only if its cycle space has a basis such that every edge is contained in at most two members of it. Notice that MacLane's criterion characterizes planar graphs in terms of the cycle space. MacLane observed that in (finite) plane graphs, the set of facial walks is a double cover that generates the cycle space. Then he proved that, conversely, any double cover of closed walks with this property can be realized as a set of facial walks and is therefore a certificate for planarity.

The planarity criterion of Archdeacon, Bonnington and Little [1] works in a similar way. Instead of the facial walks leading to the planarity criterion though, they list the essential properties of left-right tours. These properties are rather more elaborate and necessitate a number of definitions, which we give below.

MacLane's criterion has been extended to locally finite graphs in [6], using the cycle space as defined by Diestel and Kühn. Similarly, in this chapter, we extend the planarity criterion of Archdeacon, Bonnington and Little to infinite, locally finite graphs.

Let us first start with the definitions. Consider a locally finite graph G, and let $\mathcal{W}$ be a double cover of tours in $|G|$, ie. every edge is traversed twice by $\mathcal{W}$. For any l, let $\mathcal{H}$ be a cyclic sequence $e = r_1, W_1, \ldots, r_l, W_l, r_{l+1} = e$ so that $W_i \in \mathcal{W}$, $r_i \in E(G)$, and W_i contains both of r_i and r_{i+1}, and $r_i \neq r_{i+1}$ for all i. We call such a sequence $\mathcal{H}$ a *ladder* (with respect to $\mathcal{W}$), and we say that the r_i are the *rungs* of $\mathcal{H}$. For a ladder of length one we require that W_1 traverses the rung r_1 twice.

Usually, the W_i will be distinct members of $\mathcal{W}$ and the r_j distinct edges of G. In this case, the sequence is called, more precisely, a *simple ladder*. With one exception in the proof of Theorem 7.2, every ladder we encounter will be simple, so for ease of terminology we will refer to them only as ladders.

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

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

We can now state the planarity criterion for the finite case:

Theorem 7.1 (Archdeacon, Bonnington & Little [1]).
A finite graph is planar if and only if it has a diagonal. In particular, the set of LRTs of a finite plane graph is a diagonal.

A simple proof of this criterion can be found in Keir and Richter [18]. So, let us extend Theorem 7.1 to locally finite graphs:

Theorem 7.2. [5]
A locally finite graph is planar if and only if it has a diagonal.

Proof. Let G be a locally finite graph. First, assume G to be planar. From Theorem 2.7 we know that $|G|$ has an embedding in the sphere, and thus Lemma 6.3 yields (with, for instance, $B = \emptyset$) a set $\mathcal{L}$ of LRTs so that every LRS of G lies in exactly one element of $\mathcal{L}$. Hence (with Lemma 5.1 (iii)), $\mathcal{L}$ is a double cover of G. Furthermore, Theorem 5.5 implies that since $\triangle L \in \mathcal{B}(G)$, in particular $\triangle L$ is a cut for each $L \in \mathcal{L}$.

For $\mathcal{L}$ to be a diagonal, it remains to show that for any side S of any ladder $\mathcal{H}$ (with respect to $\mathcal{L}$), the residue $\triangle S$ is a cut as well. We show that $\triangle S$ meets every finite cycle C in an even number of edges. With Lemma 2.3 it then follows that $\triangle S$ is a cut.

So, let $\mathcal{H} = r_1, L_1, \ldots, r_k, L_k, r_1$ be any given ladder with respect to $\mathcal{L}$. Let C be any given finite cycle in G, and let $R{:=}\{r_1, \ldots, r_k\}$ denote the set of rungs of $\mathcal{H}$. We define H to be the plane subgraph of G consisting of C and all the edges in R together with their incident vertices. Since now it is enough to consider a finite subgraph that behaves locally like G, we can apply Lemma 5.3 to H and the LRTs $\{L_1, \ldots, L_k\}$ in $\mathcal{H}$. This gives us a finite plane supergraph H' of H, and a set $\mathcal{L}'{:=}\{L'_1, \ldots, L'_k\}$ of LRTs of H' so that L_i and L'_i agree on H for all i.

By our construction, not every edge in $H' \setminus H$ is covered twice by LRTs. Hence, in order to obtain a properly defined ladder of H', let us apply Lemma 5.4 (i) which extends $\mathcal{L}'$ to a double cover $\mathcal{L}''{:=}\{L'_1, \ldots, L'_k, \ldots, L'_l\}$ of tours in H'. We now claim that $\mathcal{H}'{:=}r_1, L'_1, \ldots, r_k, L'_k, r_1$ is a ladder in H' with respect to $\mathcal{L}''$. Indeed, by our construction, $\mathcal{L}''$ is a double cover of tours in H', and $L'_i \in \mathcal{L}''$ for every $i = 1, \ldots, k$. Furthermore, since the tours L_i and L'_i agree on $H \supseteq R$, the L'_i remain distinct, and every L'_i contains the distinct rungs r_i and r_{i+1} for all i.

Altogether we now have the ladder $\mathcal{H}$ with respect to the LRTs $\{L_1, \ldots, L_k\}$ in G, and the ladder $\mathcal{H}'$ with respect to the LRTs $\{L'_1, \ldots, L'_k, \ldots, L'_l\}$ in H', such that on H, the tours L_i and L'_i agree for all i, and so do the rungs r_i. Thus, the ladder $\mathcal{H}'$ also has a side S' which agrees with the side S (of the ladder $\mathcal{H}$) on H. Indeed, since every tour $\vec{L}_i$ induces a direction on its corresponding finite counterpart $\vec{L}'_i$, we obtain paths P_i' which agree with P_i on H for all i. On the other hand, the set of inconsistent rungs R_{in} is a subset of R and hence is contained in $E(H)$. Therefore the side $S' := \bigcup_{i=1,\ldots,k} P_i' \cup R_{\text{in}}$ agrees with S on H. Hence, we have that $\triangle S' \cap E(H) = \triangle S \cap E(H)$. Now we can apply Theorem 7.1 to the finite side S' and obtain that $\triangle S'$ is a cut. Thus, the intersection $\triangle S' \cap E(C)$ is an even set by Lemma 2.3. Therefore, $|\triangle S \cap E(C)| = |\triangle S' \cap E(C)|$ is even, too (since $C \subseteq H$), which shows that $\triangle S$ is a cut as well. This proves $\mathcal{L}$ to be a diagonal.

For the reverse implication, let us first note that given a diagonal $\mathcal{D}$ of tours D_i in $|G|$, for any ladder $\mathcal{H}$ with a side S, the assertion that $\triangle S$ is a cut does not depend on the side of $\mathcal{H}$, meaning that it does not depend on the orientations $\vec{D}_i$ of the tours in $\mathcal{D}$. Indeed, let us change the orientation of one tour, say D_i, and let S' be the resulting side of $\mathcal{H}$. Then the rung r_i becomes consistent if it was inconsistent before, or vice versa, and the same holds for r_{i+1}. Hence, the number of consistent (resp. inconsistent) edges changes by 0 or 2. Therefore, $S' = S + P_i + \widetilde{P}_i \bmod 2$, and thus, $\triangle S' = \triangle S + \triangle D_i$. So we see that in general, the residues of two sides of a ladder only differ in the sum of residues of some D_i. Now, since $\mathcal{D}$ is a diagonal, we know that $\triangle D_i$ is a cut for every i. Hence, for two different sides S and T of a ladder it holds that $\triangle S$ is a cut if and only if $\triangle T$ is a cut.

Having said that, let us suppose that G has a diagonal $\mathcal{D}$ but also contains a subdivision X of $K_{3,3}$ or of K_5. We denote by H the (finite) induced subgraph of G on $V(X)$, and let $F := E(H, G-H)$, which is a finite cut. Our idea is to delete, one by one, all edges of F from G. Using arguments of Archdeacon, Bonnington, and Little [1], we will show that after each edge deletion, our graph still has a diagonal. Then, once we have deleted all of F, the diagonal will split into two parts: into the set $\mathcal{D}'$ of those tours that are completely contained in H, and into the tours that are disjoint from H. As $\mathcal{D}'$ remains (by what we will show) a diagonal of the finite non-planar graph H, we obtain a contradiction to Theorem 7.1.

We now show that deleting an edge still leaves a diagonal of the remaining graph. We will distinguish three cases, depending on whether the edge e

occurs in two different tours of $\mathcal{D}$, or whether it is traversed twice by the same tour D, in which case it is either consistent or inconsistent with respect to the orientation of D.

Case 1: e is traversed twice by the same tour D and is consistent with respect to $\vec{D}$.

Let Q_1 and Q_2 be the two topological subpaths in D between the two occurrences of e, with the orientation induced by $\vec{D}$. Hence, $\vec{D}{=}e\vec{Q}_1e\vec{Q}_2$. By Q_2^{-1} we denote the topological path Q_2 with the reverse direction of $\vec{Q}_2$. Then $\vec{D}'{:=}\vec{Q}_1Q_2^{-1}$ also describes a tour in $|G|$. Now, we claim that after deleting the edge e, the set $\mathcal{D}'{:=}(\mathcal{D} \setminus \{D\}) \cup \{D'\}$ is a diagonal of $G'{:=}G \setminus \{e\}$.

Clearly, $\mathcal{D}'$ still forms a double cover of the edges of G'. For any tour $D_i \neq D'$ its residue remains the same (since e was not contained in any other tour of the double cover), and hence still forms a cut in G'. On the other hand, for D' it holds that $\triangle D' = \triangle Q_1 + \triangle Q_2 = \triangle D$, since e occurred twice in D. Thus, $\triangle D'$ is also a cut in G'.

Now, let $\mathcal{H}'$ be any ladder in $\mathcal{D}'$. If $\mathcal{H}'$ does not contain D', clearly any side of $\mathcal{H}'$ is a side of the corresponding ladder in G, and hence their residues remain cuts (not containing e). So let us assume that D' appears in $\mathcal{H}'$. Let $\mathcal{H}$ be its corresponding ladder in G, obtained by replacing D' by D in $\mathcal{H}'$ (with the respective orientations), as described above. Let any side S of $\mathcal{H}$ be given. Then, by hypothesis, $\triangle S$ is a cut in G.

Let $\mathcal{H}' = r_1, D_1, \dots, r_l, D_l, r_1$. Then $D' = D_j$ for some j. Hence, D' contains the two rungs $r_j =: f$ and $r_{j+1} =: g$. Again, we need to consider different cases. First, let us assume that both f and g are contained in the same Q_i, say Q_1. Then let Q' be the topological subpath of Q_1 between f and g. Hence, the subpath in D corresponding to Q' does not contain e. Therefore $\mathcal{H}$ has a side S which does not contain e and which is identical to a side of $\mathcal{H}'$. Hence, its residue cannot contain e and remains a cut in G'.

So let us now consider the case that (without loss of generality) $f \in Q_1$ and $g \in Q_2$. Then the orientations of the ladders $\mathcal{H}$ and $\mathcal{H}'$ are the same for every D_i with $i \neq j$, and for the subpath Q_1. We would like to determine the difference between $\triangle S$ and $\triangle S'$ for the corresponding side S' of $\mathcal{H}'$. From the orientations we can conclude that all rungs $r_i \neq g$ remain inconsistent in G' if and only if they were inconsistent in G. For the rung g the consistency changes. For every tour $D_i \neq D$ the orientations remain the same, and hence so do their contributions to S'.

So let us consider the tour D. We denote the different subpaths of D between the edges e, f and g by different R_i; thus, let $\overrightarrow{D} = eR_1fR_2eR_3gR_4$. Then $\overrightarrow{D}' = R_1fR_2R_4^{-1}g^{-1}R_3^{-1}$ with the orientation of R_4gR_3 reversed. But this reversal does not influence the residue of a side, hence we can ignore it here. Let us denote by Y the change between the residues of the sides S and S', ie. let $Y = \triangle S + \triangle S'$. Then $Y = \{g\} + \triangle(R_2eR_3) + \triangle(R_2R_4) = \triangle(R_3gR_4e)$.

We will show that $\triangle(R_3gR_4)$ is a cut in G. Since $\mathcal{D}$ is, by hypothesis, a diagonal of G and contains the tour D, it also contains the ladder e, D, e of length one. We note that R_3gR_4 is a side of this ladder, and hence, $\triangle(R_3gR_4)$ is a cut in G. Thus, since $\triangle(R_3gR_4e) = Y = \triangle S + \triangle S'$, it follows that $Y + \{e\} + \triangle S = \triangle S' + \{e\}$, and therefore $\triangle(R_3gR_4) + \triangle S = \triangle S' + \{e\}$. Since both $\triangle(R_3gR_4)$ and $\triangle S$ are cuts in G, we obtain that $\triangle S' + \{e\}$ is a cut in G as well, and thus $\triangle S'$ is a cut in G'. This shows $\mathcal{D}'$ to be a diagonal of G'.

Case 2: e is traversed twice by the same tour D and is inconsistent with respect to $\overrightarrow{D}$.

Again, we denote by Q_1 and Q_2 the two topological subpaths in D between the two occurrences of e, with the orientation induced by $\overrightarrow{D}$. Then the Q_1, Q_2 form two closed walks in G not containing e, and hence they remain closed walks in $G \setminus \{e\}$. We set $\mathcal{D}' := (\mathcal{D} \setminus \{D\}) \cup \{Q_1, Q_2\}$. We claim that after deleting the edge e, the set $\mathcal{D}'$ is a diagonal of $G' := G \setminus \{e\}$.

Clearly $\mathcal{D}'$ is a double cover of G'. For any tour $D_i \neq Q_1, Q_2$ in $\mathcal{D}'$, we know that $\triangle D_i$ is a cut in G and does not contain e. Hence, it remains a cut in G'. So we still need to show that both $\triangle Q_1$ and $\triangle Q_2$ are cuts in G'. Let us consider the ladder e, D, e of length one in G. Since e is inconsistent, both $\{e\} \cup Q_1$ and $\{e\} \cup Q_2$ can occur as sides of this ladder. Hence, by hypothesis, $\triangle(eQ_1)$ and $\triangle(eQ_2)$ are both cuts in G. Therefore $\triangle(Q_1)$ and $\triangle(Q_2)$ are cuts in G'.

At this point, we need to refine our definition of a ladder (as mentioned before). For a double cover of tours W_i, a ladder remains a cyclic sequence $r_1, W_1, \ldots, r_l, W_l, r_{l+1} = r_1$ as described above. The difference now is that the tours W_i and the rungs r_i need not be distinct. More precisely, a *multi-ladder* is such a cyclic sequence where the W_i need not be distinct, and every W_i (still) contains the rungs r_i and r_{i+1} with $r_i \neq r_{i+1}$. So far, the W_i have always been distinct, in which case we speak of a *simple ladder*. Now, when we only refer to *ladders*, we mean both, multi- and simple ladders. We note that the assertion in the theorem refers to simple ladders, thus it suffices to consider simple ladders here.

Next, let $\mathcal{H}'$ be any simple ladder in $\mathcal{D}'$. Let $\mathcal{H}$ be the corresponding ladder in G which we obtain from replacing both Q_1 and Q_2 in $\mathcal{H}'$, each by D (with the respective orientations). We see that if $\mathcal{H}'$ contains neither Q_1 nor Q_2, then $\mathcal{H}$ is simple, and any side of $\mathcal{H}'$ is a side of $\mathcal{H}$ in G. Hence their residues remain cuts (not containing e) in G'. Now, let us assume that exactly one of the Q_i, say Q_1, appears in $\mathcal{H}'$. In this case, too, $\mathcal{H}$ is a simple ladder in G. Then let Q' be the topological subpath of Q_1 between its two rungs f and g, that does not contain an endvertex of e. Hence, the subpath in D corresponding to Q' does not contain e. Therefore $\mathcal{H}$ has a side S which does not contain e and which is identical to a side of $\mathcal{H}'$. It follows that its residue cannot contain e and remains a cut in G'.

So, let us assume that both Q_1 and Q_2 are tours in $\mathcal{H}'$. We distinguish two cases, depending on whether Q_1 and Q_2 appear consecutively or not. First, let us assume they are *not* consecutive in $\mathcal{H}'$. Hence, we may assume that $Q_1 = D_1$ and $Q_2 = D_j$ for some j. For simplicity, let us denote the rungs of the Q_i by a, b, c, d meaning $a = r_1$, $b = r_2$, $c = r_j$ and $d = r_{j+1}$. Let $\mathcal{H}$ denote the ladder (of length, say l) in G corresponding to $\mathcal{H}'$. Then $\mathcal{H}$ is not simple. Let us consider the subsequences $\mathcal{L}_1{:=}b, D_2, r_3, \ldots, r_{i-1}, D_{i-1}, c, D, b$ and $\mathcal{L}_2{:=}d, D_{i+1}, \ldots, r_l, D_l, a, D, d$ of $\mathcal{H}$. Then $\mathcal{L}_1$ and $\mathcal{L}_2$ are simple ladders in G since $\mathcal{H}'$ is simple. Now, orient the tours in $\mathcal{L}_1$ and $\mathcal{L}_2$ arbitrarily, and let every $D_i \neq Q_k$ have the same direction in $\mathcal{H}'$. Give Q_1 and Q_2 the orientation induced by D. Then this gives us a side S' of $\mathcal{H}'$ which induces a side S of $\mathcal{H}$. Let S_i denote the side of $\mathcal{L}_i$ corresponding to S for $i = 1, 2$.

As in Case 1, we ask ourselves how the sides S_1, S_2 and S' differ. The set of inconsistent edges remains the same. For all D_i except Q_1 and Q_2, the subpaths in S appear in either S_1 or S_2. Hence, the S_i differ from S only in the subpaths from D, Q_1, and Q_2. For simplicity, let us label them. We have two subcases.

First, let $\vec{D} = R_1aR_2bR_3eR_4cR_5dR_6e^{-1}$. Since the sequence between the rungs c and b is included in S_1, we have $R_5dR_6e^{-1}R_1aR_2 \in S_1$. Similarly, $R_2bR_3eR_4cR_5 \in S_2$, and $R_2, R_5 \in S'$. Now, let $Y{:=}\triangle S' + \triangle S_1 + \triangle S_2$. Then $\triangle Y = \triangle R_2 + \triangle R_5 + \triangle(R_5dR_6eR_1aR_2) + \triangle(R_2bR_3eR_4cR_5) = \triangle D$ which is a cut in G since $D \in \mathcal{D}$, and $\mathcal{D}$ is a diagonal in G by hypothesis. Since $e \notin \triangle D$, we have that $Y = \triangle D$ is also a cut in G'.

Since $\mathcal{L}_1$ and $\mathcal{L}_2$ are simple ladders in G, we know that $\triangle S_1$ and $\triangle S_2$ are both cuts in G, and they both contain the inconsistent edge e. Hence, $\triangle S_1 + \triangle S_2$ is also a cut in G which does not contain e. Therefore $\triangle S_1 + \triangle S_2$ is a cut in G' as well, and it follows that $\triangle S' = Y + \triangle S_1 + \triangle S_2$ is a cut in G'.

For the other subcase, in which the rungs c and d are reversed in Q_2, we let $\vec{D} = R_1aR_2bR_3eR_4dR_5cR_6e^{-1}$, and in a similar way we obtain that $\triangle S'$ is a cut in G'.

Now, to complete our (main) Case 2, assume that the tours Q_1 and Q_2 are consecutive in $\mathcal{H}'$. Let $\mathcal{H}' = r_1, Q_1, r_2, Q_2, r_3, D_3, \ldots, r_l, D_l$. For simplicity, we set $a{:=}r_1, b{:=}r_2$ and $c{:=}r_3$. Then b occurs once in Q_1 and once in Q_2. We denote these two occurrences by b_1 and b_2, respectively. Note that the ladder $\mathcal{H} = a, D, c, D_3 \ldots, r_l, D_l$ in G is simple. We direct D arbitrarily and let the Q_i have the directions induced by $\vec{D}$. Let every tour $D_i \neq D$ have the same directions in $\mathcal{H}'$ as in $\mathcal{H}$. As before, the consistency of the rungs does not change in G' and hence we only have to examine the subpaths of D, Q_1, and Q_2, which we will again label with some R_i. Again, we consider two subcases.

First, let $\vec{D} = R_1aR_2b_1R_3eR_4b_2R_5cR_6e^{-1}$. Then, as before, we know that in G, the side S of $\mathcal{H}$ that corresponds to the side S' of $\mathcal{H}'$ contains the topological subpath of $\vec{D}$ that lies between the two rungs a and c. Thus, we have $R_2b_1R_3eR_4b_2R_5 \in S$. Similarly, we obtain $R_2, R_5 \in S'$. As before, let $Y{:=}\triangle S + \triangle S'$. Then $Y = \triangle(R_2b_1R_3eR_4b_2R_5) + \triangle R_2 + \triangle R_5$.

If the 'new rung' b is consistent in $\vec{D}$, then it will not contribute to the side S'. Hence, we have $Y = \triangle(R_2b_1R_3eR_4b_2R_5)+\triangle R_2+\triangle R_5 = \triangle(R_3eR_4)$. Else, if b is inconsistent in $\vec{D}$, we have that $b \notin S$, but $b \in S'$ and thus, $b \in Y$. Therefore, $Y = \triangle(R_3eR_4) + b$. Note furthermore that b, D, b is a ladder of length one in G. Therefore, it has a side $T = R_3eR_4$ or $T = R_3eR_4b$, depending on whether b is consistent or not. In either case, Y is the residue of the corresponding side and is therefore a cut in G. Since $e \in \triangle S$ and $e \in Y$, we know that $\triangle S' = \triangle S + Y$ does not contain e. Since we have shown that both $\triangle S$ and Y are cuts in G, it follows that $\triangle S'$ is a cut in G not containing e, and hence it is also a cut in G'.

For the other subcase, we have $\vec{D} = R_1aR_2b_1R_3eR_4cR_5b_2R_6e^{-1}$. In a similar fashion we can show that in this case as well, for a side S' of $\mathcal{H}'$ it holds that $\triangle S'$ is a cut in G'. Thus we have shown that after deleting e, the set $\mathcal{D}'$ is a diagonal of G'.

Case 3: e is traversed by two different tours D_1 and D_2.

Let Q_1 and Q_2 be the two topological subpaths in D_1, resp. D_2 between e, with the orientation induced by $\vec{D}_1$, resp. $\vec{D}_2$; thus, $\vec{D}_i = e\vec{Q}_i$. By Q_2^{-1} we denote the topological path Q_2 with the reverse direction of $\vec{Q}_2$. If e is inconsistent with respect to D_1 and D_2, we let $\vec{D}'{:=}\vec{Q}_1\vec{Q}_2$. Otherwise, if e is consistent, let $\vec{D}'{:=}\vec{Q}_1Q_2^{-1}$. Then $\vec{D}'$ also describes a tour in $|G|$. Now, we claim that after deleting the edge e, the set $\mathcal{D}'{:=}(\mathcal{D} \setminus (\{D_1\} \cup \{D_2\})) \cup \{D'\}$ is a diagonal of $G'{:=}G \setminus \{e\}$.

It is clear that $\mathcal{D}'$ still forms a double cover of the edges of G'. For any tour D_i with $i \notin \{1, 2\}$, its residue remains the same (since e was not contained in any other tour of the double cover) and hence still forms a cut in G'. On the other hand, for D' it holds that $\triangle D' = \triangle D_1 + \triangle D_2$ where $\triangle D_1$ and $\triangle D_2$ are cuts in G containing e. Hence, $\triangle D'$ does not contain e and is a cut in G'.

Let $\mathcal{H}'$ be any simple ladder in $\mathcal{D}'$. If $\mathcal{H}'$ does not contain D', clearly any side of $\mathcal{H}'$ is a side of the corresponding ladder in G, and hence their residues remain cuts (not containing e). So let us assume that D' appears in $\mathcal{H}'$. We may assume that $\mathcal{H}' = r_1, D', r_2 \ldots, r_l, D_l, r_1$. Let $f{:=}r_1$ and $g{:=}r_2$. We need to consider two subcases.

First, let us assume that both f and g are contained in the same subpath Q_i of D', say Q_1. Then let $\mathcal{H}$ be the ladder in G which we obtain from replacing D' by D_1. We direct all tour in $\mathcal{H}$ arbitrarily, and let every D_i for $i \neq 1$ have the same orientation in $\mathcal{H}'$ as in $\mathcal{H}$. We orient D' according to the direction induced by D_1, and depending on whether e is consistent in $\mathcal{H}$ or not, as described above.

Now, let Q' be the topological subpath of Q_1 between f and g. Hence, the subpath in D_1 corresponding to Q' does not contain e. Therefore $\mathcal{H}$ has a side S which does not contain e and which is identical to a side of $\mathcal{H}'$ (remember here that since $D_2 \notin \mathcal{H}$, the edge e cannot be a rung). Hence, its residue cannot contain e and remains a cut in G'.

So let us now consider the case (without loss of generality) that $f \in Q_1$ and $g \in Q_2$. Then let $\mathcal{H}$ be the ladder of length $l+1$ in G which we obtain from replacing D' by D_1, e, D_2. Again, we direct all tours in $\mathcal{H}$ arbitrarily, and let the corresponding tours in $\mathcal{H}'$ inherit their orientations. First, assume that e is consistent in $\mathcal{H}$. Thus we direct D' as $\vec{D}' {:=} \vec{Q}_1 Q_2^{-1}$. Since e is a (consistent) rung of $\mathcal{H}$, it is not contained in any side S of $\mathcal{H}$, and moreover, there exists a side S' of $\mathcal{H}'$ with $S' = S$. Since $\triangle S$ is a cut in G by assumption, and it does not contain e, it follows that $\triangle S'$ is a cut in G'.

On the other hand, let us assume that e is inconsistent in $\mathcal{H}$. Then we direct D' as $\vec{D}' {:=} \vec{Q}_1 \vec{Q}_2$. Let S be a side of $\mathcal{H}$ that corresponds to a side S' of $\mathcal{H}'$. Then, since e is inconsistent, it follows that $S = S' \cup \{e\}$. Since $\triangle S$ is a cut in G, we have that $\triangle(S' \cup \{e\})$ is also a cut in G. Hence, $\triangle S'$ is a cut in $G' = G \setminus \{e\}$.

Furthermore we remark that since a diagonal forms a double cover, all considered sums are thin and therefore well-defined.

Altogether we have shown that if G contains a diagonal, then so does $G \setminus \{e\}$. The reverse implication of our theorem therefore follows. This completes the proof. $\square$

For pedestrian graphs, ie. those graphs G for which $\mathcal{B}(G) = \{\emptyset\}$, Rosenstiehl and Read [24] gave a slightly simpler planarity criterion (which led Archdeacon, Bonnington and Little to prove their extension in the form of Theorem 7.1). Let a tour W traverse an edge $e = uv$ twice. If e is consistent and is traversed from u to v, say, then W decomposes into four topological subpaths uv, H_1, uv and H_2. We call each of H_1 and H_2 a *half* of W (with respect to e). If e is inconsistent, then W is equally comprised of four topological subpaths: namely of uv, H_1', vu and H_2'. In this case we call the topological subpaths $uvH_1'{:=}H_1$ and $vuH_2'{:=}H_2$ *halves* of W.

We note two facts. First, e is contained in each half of W if and only if e is inconsistent in W. This results from the idea that a half should connect the two endvertices of e. Second, if e, W, e is seen as a ladder (of length one), then a half is simply a side of this ladder (more precisely, they have the same residues).

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

Theorem 7.3 (Rosenstiehl & Read [24])**.**
A finite connected pedestrian graph is planar if and only if it has an algebraic diagonal.

Let us extend this theorem to locally finite graphs.

Theorem 7.4. [5]
A locally finite connected pedestrian graph is planar if and only if it has an algebraic diagonal.

Proof. Let G be a locally finite connected pedestrian graph. If G is planar, then $|G|$ can be embedded in the sphere (recall Theorem 2.7), and by Lemma 6.3 there exists a family $\mathcal{L}$ of LRTs in $|G|$ that forms a double cover of $E(G)$.

We already know (from the proof of Theorem 7.2) that $\mathcal{L}$ is a diagonal of G. If $\mathcal{L}$ has only a single member D, then D is an algebraic diagonal of G. Indeed, let any edge e be given. Then (since e is inconsistent in D if and only if it is contained in each half of D) a half H of D with respect to e has the same residue as a side S of the (simple) ladder e, D, e. Since $\mathcal{L}$ is a diagonal of G, we know that $\triangle S$ is a cut, and hence it follows that $\triangle H = \triangle S$ is a cut, too.

So, assume that $\mathcal{L}$ has at least two members, and denote one of them by L. Since G is pedestrian and L is a left-right tour in $|G|$, Theorem 5.5 implies that $\triangle L = \emptyset$. Thus, since G is connected, there is a vertex $v \in V(G)$ which is incident with edges traversed by L and with edges not lying in L.

Consider an edge e incident with v that lies in L. Without loss of generality, let η be the half of e with $v \notin \partial\eta$, and let σ be a side of e. Since L traverses e twice (since $\triangle L = \emptyset$), we know that L (or, more precisely, the LRS lying in L that traverses e) contains one corner of the set $\{(e, \eta, \sigma), (e, \overline{\eta}, \overline{\sigma})\}$, and one corner of $\{(e, \eta, \overline{\sigma}), (e, \overline{\eta}, \sigma)\}$. Let (e_1, η_1, σ_1) be the corner that is matched with $(e, \overline{\eta}, \overline{\sigma})$, and let (e_2, η_2, σ_2) be the one matched with $(e, \overline{\eta}, \sigma)$. By definition, if L contains (e, η, σ), then it also contains (e_1, η_1, σ_1). If, on the other hand, $(e, \overline{\eta}, \overline{\sigma})$ lies in L, then $(e_1, \overline{\eta}_1, \overline{\sigma}_1)$ is also a corner of L. Hence, in any case, e_1 is traversed by L. Similarly, if $(e, \eta, \overline{\sigma})$ is contained in L, then so is (e_2, η_2, σ_2); and else, if $(e, \overline{\eta}, \sigma)$ is a corner of L, then $(e_2, \overline{\eta}_2, \overline{\sigma}_2)$ is also in L.

Hence, it follows that the predecessor and the successor of e in the local rotation of the edges at v both lie in L, and thus all of $E(v)$ is covered by L, which contradicts our assumption.

So, conversely, let us assume that G has an algebraic diagonal D. Then the set $\{D\}$ is a diagonal. By assumption, D is a double cover of G. For the only tour D in $\{D\}$ it holds that for any side S of D, we have $\triangle S = \triangle H$ for any half H of D, and $\triangle H$ is a cut by assumption. Last, we have to check that $\triangle D$ is also a cut. Let any edge e be given (consistent or inconsistent), and consider the two halves H_1, H_2 of D with respect to e. Then we obtain $\triangle D = \triangle H_1 + \triangle H_2$ which is a cut since $\triangle H_1$ and $\triangle H_2$ are cuts, since D is an algebraic diagonal.

Thus, $\{D\}$ is a diagonal of G which by Theorem 7.2 implies that G is planar. □

We come to the third and last result in this chapter that we will extend to locally finite graphs. Richter and Shank proved the following theorem for arbitrary surfaces:

Theorem 7.5 (Richter & Shank [22]).
Let G be a finite plane graph with a double cover consisting of LRTs. Orient each LRT arbitrarily. The edges that are traversed both times in the same direction form an element of the cycle space, while those traversed once in each direction form a cut.

We will show a similar result for locally finite plane graphs.

Theorem 7.6.
Let G be a locally finite graph, and let $|G|$ be embedded in the sphere. Let $L_1, \dots, L_k$ be a set of LRTs of G so that every LRS of G lies in exactly one L_i. Orient each LRT arbitrarily. Then the edges that are traversed both times in the same direction by some LRT form an element of the cycle space, while those traversed once in each direction by some LRT form a cut.

Proof. First, note that by Lemma 5.1 (iii) the set of all LRSs forms a double cover of G, and hence, by assumption, so does the set of LRTs. Thus, every edge is indeed traversed exactly twice. Let us denote by Y the set of edges that are traversed both times in the same direction by some LRT, and let K denote the set of edges that are traversed once in each direction.

First, we will show that $K \in \mathcal{C}^*(G)$. By Lemma 2.3, we know that K is a cut in G if and only if $|K \cap Z|$ is even for every finite circuit $Z \subseteq E(G)$. So let Z be given. Since Z is finite, there exists a finite subgraph H of G with $Z \subseteq E(H)$.

We will use Lemma 5.3 on G and H to obtain a finite plane supergraph H' of H and a set $L'_1, \ldots, L'_k$ of LRTs of H' such that L_i and L'_i agree on H for all i. Furthermore, Lemma 5.4 (i) yields a double cover $\mathcal{L}''$ of LRTs of H' that maintains this property.

Now, let K' be the set of edges in H' that are traversed once in each direction by the LRTs in $\mathcal{L}''$. Since H' is finite, we can apply Theorem 7.5 to it and see that K' is a cut in H'. From Lemma 5.4 (ii) we know that Z is also a circuit in H' and thus, with Lemma 2.3, we obtain that $|Z \cap K'|$ is even. But since $Z \subseteq E(H)$, we have that $|Z \cap K| = |(Z \cap E(H)) \cap K| = |Z \cap (E(H) \cap K)| = |Z \cap (E(H) \cap K')| = |(Z \cap E(H)) \cap K')| = |Z \cap K'|$, which is even, and thus $K \in \mathcal{C}^*(G)$.

Similarly, we will show that $Y \in \mathcal{C}(G)$. From Theorem 2.2 we know that $Y \in \mathcal{C}(G)$ if and only if $|Y \cap F|$ is even for every finite cut F in G. So let F be given. Then let H be a finite subgraph of G such that $F \subseteq E(H)$.

By using Lemma 5.3 on G and H, we obtain a finite plane supergraph H' of H and a set $L'_1, \ldots, L'_k$ of LRTs of H' such that L_i and L'_i agree on H. Furthermore, Lemma 5.4 (i) yields a double cover $\mathcal{L}''$ of LRTs of H' that maintains this property.

Let Y' denote those edges of the LRTs $L''_i \in \mathcal{L}''$ in H' that are traversed both times in the same direction by some L''_i. Since H' is finite, Theorem 7.5 yields that $Y' \in \mathcal{C}(H')$. Using Lemma 5.4 (iii), we obtain that F is also a cut in H' and thus, by Theorem 2.2, we see that $|F \cap Y'|$ is even. But since $F \subseteq E(H)$, we have $|F \cap Y| = |F \cap (E(H) \cap Y)| = |F \cap (E(H) \cap Y')| = |F \cap Y'|$, which is even. Therefore $Y \in \mathcal{C}(G)$, which concludes the proof.

□

As Richter and Shank already noted, their theorem gives another proof — for the case of plane graphs — of the well-known fact that the edges of *any* finite graph may be partitioned into a cycle space element and a cut. This follows from a theorem of Gallai (T. Gallai (unpublished), we refer the reader to [20]). The same is true for Theorem 7.6. Using the cycle space as defined by Diestel and Kühn, Gallai's theorem has been extended to the infinite case:

Theorem 7.7 (Bruhn, Diestel & Stein [4]).
For every locally finite graph G there is a partition of its vertex set into two (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.

The edges of a locally finite graph G thus may be partitioned into the sets $Z{:=}E(G[V_1])\cup E(G[V_2])$ and F, where F is the set of $G[V_1]-G[V_2]$ edges. By Theorem 7.7, we obtain that $Z \in \mathcal{C}(G)$, while by the definition of a cut, it holds that $F \in \mathcal{C}^*(G)$. Hence, the edge set of a locally finite graph may be partitioned into an element of the cycle space and a (possibly empty) cut, for which Theorem 7.6 provides an alternate proof in the case of plane graphs.

Chapter 8

Pedestrian Graphs

In the last chapter we have seen that things become a little easier if we consider the special case of pedestrian graphs. We have also seen examples of locally finite graphs where we lose a certain uniqueness that was given in the finite case:

Figures 1.2 and 3.1 show an edge e that is in some sense ambiguous — it is contained in an infinite bicycle, but there also exist infinite $Y, Z \in \mathcal{C}(G)$ such that $e \in Y$ with $Y + e \in \mathcal{C}^*(G)$, and also $Z \in \mathcal{C}(G)$ with $e \notin Z$ and $Z + e \in \mathcal{C}^*(G)$. Hence, the tripartition of edges ceases to be a partition if all of B, Y, Z are infinite.

Another example is given at the end of Chapter 6: As opposed to finite graphs, a locally finite graph can contain 'more' LRTs in the sense that they cover the entire edge set more than twice, yet there is still no distinguished double cover of LRTs, since in the example of the double ladder none of the double covers suffices to generate all of $\mathcal{B}$.

There is another theorem which, when extended to locally finite graphs, fails because a certain ambiguity arises. Recall from Chapter 7 that H_1, H_2 denote the two halves of a tour with respect to the edge e.

Theorem 8.1 (Richter & Shank [22]).
Let G be a finite plane graph, and let e be an edge of G.

(i) *If e lies in two distinct LRTs, say L and L', then e is a bicycle edge and both $\triangle L$ and $\triangle L'$ are bicycles.*

(ii) *If e is consistent in an LRT L, then e is of cycle-type, and $\triangle H_1$ and $\triangle H_2$ are each a principal cut of e.*

(iii) *If e is inconsistent in an LRT L, then e is of cut-type, and $\triangle H_1$ and $\triangle H_2$ are each a principal cut of e.*

In finite graphs, LRTs are also LRSs and hence form a double cover. Assuming that $\mathcal{L}$ is a *double cover* of tours in $|G|$, assertion (i) remains true in locally finite graphs. We have already shown in Theorem 5.5 that the residue of any LRT is a bicycle. If e is contained in two different LRTs $L, L' \in \mathcal{L}$ it can appear only once in each of them, and therefore, in particular, $e \in \triangle L \in \mathcal{B}$.

However, (ii) and (iii) may become false — even when only considering LRTs that form a double cover. Indeed, the edge e in Figure 8.1 lies twice in the drawn LRT, but it is neither of cycle- nor of cut-type but instead is a bicycle edge. Yet there is a double cover containing this LRT.

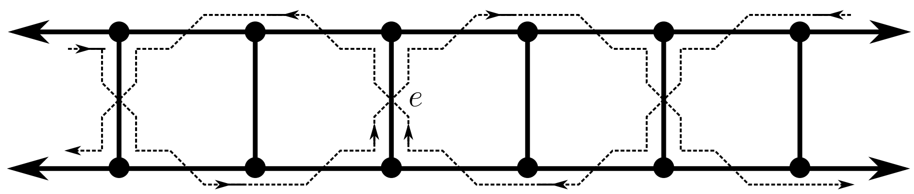

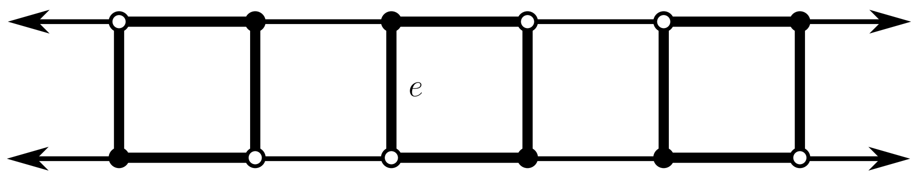

Figure 8.1: A counterexample to (ii) in Theorem 8.1

Furthermore, in Figure 8.2 the edge e is consistent in the drawn LRT L *and* is of cycle-type. Yet the residue of each half $\triangle H_1, \triangle H_2$ is an infinite cut, which is therefore not principal. In addition, any other LRS not contained in L is of the form S and is therefore finite. Thus, there exists no other LRT containing e only once.

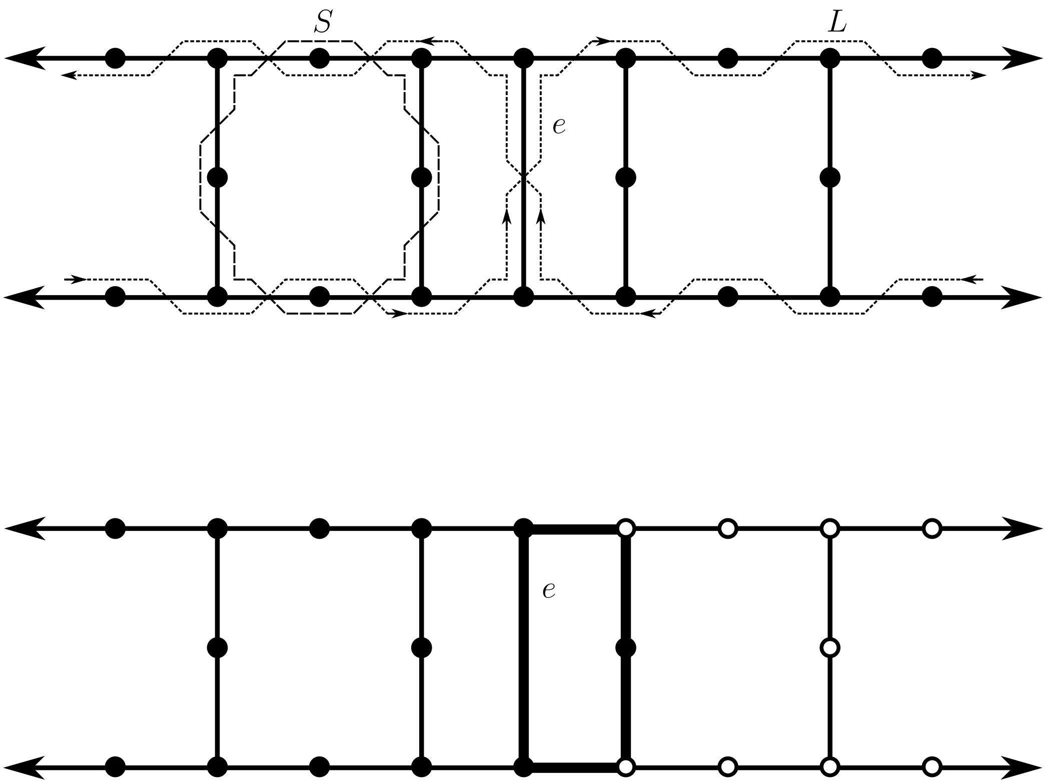

Figure 8.2: The edge e is of cycle-type, yet neither $\triangle H_1$ nor $\triangle H_2$ is a principal cut.

Let us note that given a plane graph G not containing any nonempty bicycle, we know by Theorem 5.5 that for any LRT L of $|G|$ its residue is empty, hence it traverses every edge of G exactly twice. Therefore, for any edge e in G it holds that $\triangle L = \triangle H_1 + \triangle H_2$ for the two halves of L with respect to e. Since $\mathcal{B}(G) = \{\emptyset\}$, we have that $\triangle L = \emptyset$ and obtain $\triangle H_1 = \triangle H_2$. In this chapter, we therefore use the notation $\triangle H_e := \triangle H_1 = \triangle H_2$.

So let us return to Theorem 8.1 that fails for locally finite graphs. We can show that it does not fail completely, but instead extends to locally finite graphs when considering pedestrian graphs.

Lemma 8.2.
Let G be a locally finite plane pedestrian graph, and let e be an edge of G. Then the following hold:

(i) *The edge e does not lie in two distinct LRTs.*

(ii) *If e is consistent in the LRT L, then e is of cycle-type, and $\triangle H_e$ is the principal cut of e.*

(iii) *If e is inconsistent in the LRT L, then e is of cut-type, and $\triangle H_e$ is the principal cut of e.*

Proof. For (i), we assume that e lies in two different LRTs, one of which we call L. Then, by Theorem 5.5, we have $e \in \triangle L \in \mathcal{B}$. Since G is pedestrian, we know that $\mathcal{B}$ is empty. Hence we obtain a contradiction.

To see (ii) and (iii), we first recall that in a pedestrian graph the principal cuts are unique. Since G is pedestrian, Theorem 5.5 tells us that the residue of any LRT in G must be empty. We may assume that G is connected since otherwise we could consider the connected component of G which contains e. Thus, all edges of G are traversed twice by the same LRT L. Using Lemma 5.1 (iii), we see that L forms a double cover of $E(G)$, and since G is plane, Theorem 7.4 then implies that this is an algebraic diagonal. Therefore, $\triangle H_e$ is a cut. Since $H_e + e$ is a closed walk in G (whether e is consistent in L or not), the sum of its edges modulo 2 is an element of the cycle space, and hence $\triangle H_e + e \in \mathcal{C}(G)$.

Now, suppose that $\triangle H_e$ is an infinite set. Since e is not a bicycle edge, by Theorem 3.2 there exist $Z \in \mathcal{C}_{\text{fin}}(G)$ and $K \in \mathcal{C}^*_{\text{fin}}(G)$ with $e = Z + K$. By assumption, $\triangle H_e$ is infinite and thus, in particular, $\triangle H_e \neq K$. Then, since $e+Z = K$, we obtain $\triangle H_e+e+Z = \triangle H_e+K$, or, $(\triangle H_e+e)+Z = \triangle H_e+K$. Since $\triangle H_e+e \in \mathcal{C}(G)$ and $Z \in \mathcal{C}(G)$ while $\triangle H_e \in \mathcal{C}^*(G)$ and $K \in \mathcal{C}^*(G)$, we see that $(\triangle H_e + e) + Z = \triangle H_e + K$ is a non-empty bicycle since $\triangle H_e \neq K$. Since G is pedestrian, this is a contradiction. Therefore, $|\triangle H_e|$ is finite, and because the principal cut in a pedestrian graph is unique, it follows that $\triangle H_e = K = e + Z$. Hence, $\triangle H_e$ is the principal cut of e.

Furthermore, if e is consistent in L, we know that $e \notin \triangle H_e$, and therefore $\triangle H_e+e$ is a principal cycle of e since $e \in (\triangle H_e+e) \in \mathcal{C}(G)$ and $\triangle H_e \in \mathcal{C}^*(G)$. Thus, e is of cycle-type. Similarly, if e is inconsistent in L, we have $e \in \triangle H_e$, and therefore $\triangle H_e$ is a principal cut of e since $e \in \triangle H_e \in \mathcal{C}^*(G)$ and $(\triangle H_e + e) \in \mathcal{C}(G)$. Thus, e is of cut-type. $\square$

These observations motivate us to further examine the class of locally finite pedestrian graphs. In finite graphs there is a simple characterization for pedestrian graphs:

Theorem 8.3 (Chen [9]).
A finite graph G is pedestrian if and only if the number of spanning forests of G is odd.

Since the number of spanning trees in a locally finite graph is usually infinite, this theorem certainly cannot hold for infinite graphs. In addition, the usual definition of a (finite) spanning tree fails to comprise the notion of spanning trees properly when used for infinite graphs: It is possible that an edge set T is a tree and covers all vertices of the graph, and yet $|T|$ contains an infinite cycle (cf. for example Figure 2 in [12]).

This motivates the definition of a *topological spanning tree* of a graph G: this is an arc-connected standard subspace of $|G|$ that contains every vertex and every end, but contains no circle. And indeed, fundamental properties such as the fact that the fundamental cycles with respect to a spanning tree generate the entire cycle space of the underlying graph hold in locally finite graphs precisely for the topological spanning trees.

Theorem 8.4 (Diestel & Kühn [12]).
Let G be a locally finite connected graph, and let T be any spanning tree of G. Then the closure of T in $|G|$ is a topological spanning tree of G if and only if the fundamental circuits of T generate $\mathcal{C}(G)$.

But the notion of a topological spanning tree still does nothing to fix the theorem for locally finite graphs. In general, the number of their topological spanning trees is infinite.

Consider for example the double ladder in Figure 8.3. Simply taking all edges of the form u_iu_{i+1} from the upper double ray and all edges of the form w_iw_{i+1} from the lower double ray, together with all 'half rungs' u_iv_i would be our approach to constructing a spanning tree with the notion of finite spanning trees in mind. In this case though, this edge set would contain an

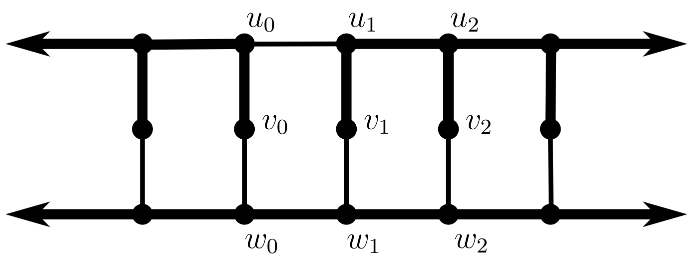

Figure 8.3: A graph with infinitely many topological spanning trees.

infinite cycle of its underlying graph, ie. the circle consisting of both double rays and both ends. Therefore, this edge set is not a topological spanning tree. But removing any edge from any of the two double rays fixes the problem: For example, in Figure 8.3 we removed the edge u_0u_1 from the edge set described above. This yields the topological spanning tree indicated in our figure. But we note that another topological spanning tree for the same graph is obtained by using the other 'half' of any rung — for example, deleting an edge u_iv_i from our edge set and instead adding the edge v_iw_i results in an edge set that still covers all vertices and all ends, and is indeed another topological spanning tree of the graph. Since we already have infinitely many choices for this operation (we can do this for any i), we see that the graph in Figure 8.3 has an infinite number of topological spanning trees.

Similar to the definition of topological spanning trees, a *topological spanning forest* is the closure in $|G|$ of a spanning forest of G that does not contain any (finite or infinite) cycles. We will show that if in a locally finite graph the number of topological spanning forests happens to be finite, then Theorem 8.3 still holds. To see this, let us cite two more results about spanning trees in infinite graphs. Recall that a *normal spanning tree* T of a graph G is, just as in finite graphs, a rooted spanning tree of G such that any two adjacent vertices of G are comparable in the tree-order of T.

Theorem 8.5 (Jung [17])**.**
Every countable connected graph has a normal spanning tree.

Lemma 8.6 (Diestel & Kühn [14])**.**
The closure of any normal spanning tree is a topological spanning tree.

Proofs of these statements can also be found in [11]. Now, let us return to pedestrian graphs.

Theorem 8.7.
Let G be a locally finite graph which has only a finite number of topological spanning forests. Then G is pedestrian if and only if the number of topological spanning forests of G is odd.

Proof. Since G has only finitely many topological spanning forests, every component K_i of G can have only finitely many topological spanning trees. So let us consider such a connected component K_i whose number $\tau(K_i)$ of topological spanning trees is finite. We will first show that there exist only finitely many finite cycles in K_i.

Since K_i is locally finite, it is in particular countable. Hence, by Theorem 8.5, there exists a normal spanning tree T of K_i. Applying Lemma 8.6, we know that its closure $\overline{T}$ in $|K_i|$ is a topological spanning tree. Let F consist of the edges of $E(K_i)$ that are not in $E(T)$, ie. $F{:=}\{e \mid e \in E(K_i) \backslash E(T)\}$. We now show that F is finite. For assume that this is not the case. Then for every edge $f \in F$ we may consider its fundamental circuit C_f with respect to $\overline{T}$. Thus there exists at least one other edge $f \neq f' \in C$ such that the closure of $T'{:=}(T \setminus \{f'\}) \cup \{f\}$ is also a topological spanning tree of K_i. If $|F|$ is infinite, then there exist infinitely many such T', similar to our argument above (see also Figure 8.3). But this contradicts the fact that $\tau(K_i)$ is finite. Since T is a tree, every finite cycle C in K_i contains an edge $e \in F$. Since $|F|$ is finite and K_i is locally finite, it follows that there exist only finitely many finite cycles in K_i.

Therefore, there exists a finite subgraph H of K_i such that $K_i \setminus H$ contains no finite cycle. We will show that K_i also contains no *infinite* cycle and therefore, $\mathcal{C}(K_i) = \mathcal{C}(H)$. The fundamental circuits of T and $\overline{T}$ are the same, and since T is connected, every fundamental circuit with respect to T is finite. From Theorem 8.4 we know that every circuit in K_i is a sum of fundamental circuits with respect to T. Since there are only finitely many finite cycles in K_i, in particular, there exist only finitely many fundamental circuits with respect to T in K_i. Thus, any sum of these circuits needs to be finite as well. Therefore K_i cannot contain an infinite cycle.

So $K_i \setminus H$ is a forest, and hence there exists exactly one set of edges, say E', such that E' covers all vertices of $K_i \setminus H$ (if $K_i \setminus H$ is connected then this is the edge set of the unique topological spanning tree, otherwise E' induces a forest). Let $\beta(H)$ be the number of topological spanning forests of H. Then every such forest corresponds to exactly one topological spanning tree of K_i since they all agree on $K_i \setminus H$. Hence, we have $\tau(K_i) = \beta(H)$ and therefore $\tau(K_i)$ is odd if and only if $\beta(H)$ is odd. Applying Theorem 8.3 to the finite graph H we know that $\beta(H)$ is odd if and only if H is pedestrian, which again holds if and only if K_i is pedestrian because $\mathcal{B}(H) = \mathcal{B}(K_i)$ since $\mathcal{C}(H) = \mathcal{C}(K_i)$ and every cut in H induces a cut in K_i. Thus, $\tau(K_i)$ is odd if and only if K_i is pedestrian.

Let $\beta(G)$ be the number of topological spanning forests of G. Then it holds that $\beta(G) = \prod_i \tau(K_i)$ and therefore we conclude that $\beta(G)$ is odd if and only if $\tau(K_i)$ is odd for every i. It remains to show that G is pedestrian if and only if every component K_i of G is pedestrian.

Since $\mathcal{C}(K_i) \subseteq \mathcal{C}(G)$ and $\mathcal{C}^*(K_i) \subseteq \mathcal{C}^*(G)$ for all i, we obtain that $\mathcal{B}(K_i) \subseteq \mathcal{B}(G)$. Hence, we have that every K_i is pedestrian if G is pedestrian. On the other hand, if G is not pedestrian, then there exists some bicycle B in G. By Theorem 2.4 it follows that B would also meet some $B_i \in \mathcal{C}(K_i)$, and since $B \in \mathcal{C}^*(G)$ and $B \cap \mathcal{C}^*(K_i) \neq \emptyset$, it holds that $B_i \in \mathcal{C}^*(K_i)$, too. Thus, $\mathcal{B}(K_i) \neq \emptyset$, and K_i is not pedestrian. Hence, K_i is pedestrian for every i if and only if G is pedestrian. Altogether we have thus shown that G is pedestrian if and only if K_i is pedestrian for every i which holds if and only if $\tau(K_i)$ is odd for every i (by our above argument), which is the case if and only if $\beta(G)$ is odd, meaning the number of topological spanning forests of G is odd.

□

Our aim is to obtain a better idea of what locally finite pedestrian graphs look like. Theorem 8.7 tells us that the number of their spanning forests must be odd in case it is finite. So if a locally finite graph contains only finitely many finite cycles, then we know that it is pedestrian if and only if the number of its spanning forests is odd.

There is another characterization of pedestrian graphs which can be found in [16].

Lemma 8.8.
A finite graph G is pedestrian if and only if for every subgraph H of G it holds that $E(H) = Z + K$ for some $Z \in \mathcal{C}(G)$ and some $K \in \mathcal{C}^*(G)$.

This extends to the locally finite case as follows:

Lemma 8.9.
A locally finite graph G is pedestrian if and only if for every finite subgraph H of G it holds that $E(H) = Z+K$ for some $Z \in \mathcal{C}_{\mathit{fin}}(G)$ and some $K \in \mathcal{C}^*_{\mathit{fin}}(G)$.

Proof. First, let us assume that G is pedestrian. Then by Theorem 3.2 we know that for every edge $e \in E(G)$ there exist $Z_e \in \mathcal{C}_{\text{fin}}(G)$ and $K_e \in \mathcal{C}^*_{\text{fin}}(G)$ such that $e = Z_e + K_e$. Now let H be some given finite subgraph of G, and let $F{:=}E(H)$. Then $F = \sum_{e\in F} e = \sum_{e\in F}(Z_e + K_e) = \sum_{e\in F} Z_e + \sum_{e\in F} K_e =: Z_F + K_F$, with $Z_F \in \mathcal{C}_{\text{fin}}(G)$ and $K_F \in \mathcal{C}^*_{\text{fin}}(G)$ since F is finite by assumption.

For the reverse implication, let us assume that G is not pedestrian. Hence, the graph G contains some bicycle B. Then, again by Theorem 3.2, we know that for every edge $e \in E(B)$ there exist no $Z \in \mathcal{C}_{\text{fin}}(G)$ and $K \in \mathcal{C}^*_{\text{fin}}(G)$ such that $e = Z + K$. Thus the set $\{e\}$ induces a finite subgraph H of G such that its edge set cannot be represented in the form $E(H) = Z + K$ with $Z \in \mathcal{C}_{\text{fin}}(G)$ and $K \in \mathcal{C}^*_{\text{fin}}(G)$.

We remark that the lemma holds not only for finite subgraphs H of G, but more generally for arbitrary finite subsets of $E(G)$, with the same proof. $\square$

We note that the restriction to *finite* subgraphs in Lemma 8.9 is indeed necessary. When also admitting infinite subgraphs, for example the graph G itself, we certainly need to allow infinite cuts and cycle space elements as well, since otherwise there is no hope of finding two edge sets constituting an infinite set of edges. As a counterexample, consider Figure 8.5. Choosing the graph G as our subgraph, we know that there exist no finite $Z \in \mathcal{C}_{\text{fin}}(G)$ and $K \in \mathcal{C}^*_{\text{fin}}(G)$ such that $E(G) = Z + K$, since $E(G)$ is an infinite set. Yet the graph G is pedestrian. There do, however, exist infinite elements of $\mathcal{C}(G)$ and $\mathcal{C}^*(G)$ which together constitute all edges of G: Let Z be the edges of the upper and lower rays together with the leftmost rung of G. Then this edge set forms an infinite cycle. On the other hand, the remaining edges (ie. all rungs except for the leftmost) form an infinite cut: let the middle vertices on these rungs form one partition class and let all other vertices constitute the other partition class. Thus, there exist infinite sets $Z \in \mathcal{C}(G)$ and $K \in \mathcal{C}^*(G)$ with $E(G) = Z + K$.

But the theorem will not extend to arbitrary subgraphs if we allow the infinite elements in $\mathcal{C}(G)$ and $\mathcal{C}^*(G)$ as summands. Figure 1.2 will serve as a counterexample. Here, the infinite graph G is not pedestrian, yet there exists an infinite element of the cycle space and an infinite cut which together constitute all edges of G: Let $Z \in \mathcal{C}(G)$ be the edges of the two double rays, which form a circle in the space $|G|$, and let $K \in \mathcal{C}^*(G)$ be the set of all rungs. Then K is an infinite cut since all vertices on the upper double ray provide one partition class, while the vertices on the lower double ray form the other partition class. Thus, there exist $Z \in \mathcal{C}(G)$ and $K \in \mathcal{C}^*(G)$ with $E(G) = Z + K$, yet the graph G contains a bicycle indicated by the bold edges in Figure 1.2.

For planar graphs, we can choose a different approach to characterize pedestrian graphs. From Theorem 5.5 we know that since plane pedestrian graphs do not contain any nonempty bicycles, the residue of every left-right tour must be empty. This means that the entire edge set of the graph is covered twice by only a single LRT. Intuitively one might think that this limits the number of ends a locally finite pedestrian graph can contain, and it seems even more likely that an end cannot have too large a degree.

Figure 8.4 gives an example of a locally finite pedestrian graph with a single end that has degree one; Figure 8.5 shows a locally finite pedestrian graph which is 2-connected and has a single end of degree two. The dotted lines indicate the single LRT in each graph.

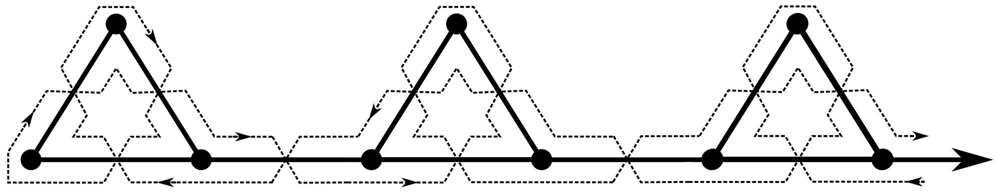

Figure 8.4: A pedestrian graph with an end of degree one.

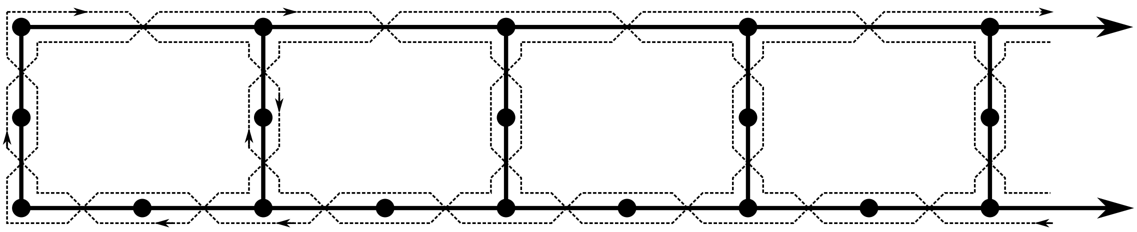

Figure 8.5: A pedestrian graph can have an end of degree two.

Surprisingly though, there is no limit to the degree of an end in a locally finite pedestrian graph. Figure 8.6 gives an example of a locally finite pedestrian graph with a (single) thick end, ie. an end that contains infinitely many disjoint rays. One can see that the subgraphs in this graph follow a certain structure and that with every vertical level towards the right in our figure, the number of these subgraphs increases by one. Hence for every $k \in \mathbb{N}$ there exists a 'next' level (further right) in which $k + 1$ disjoint rays belonging to ω begin.

By modifying the graph in Figure 8.6 a little we can stop the 'growth' that comes from adding subgraphs in every level and attain an end which contains any desired number of disjoint rays (at least three). Together with the examples from Figure 8.4 and Figure 8.5 this shows:

A locally finite pedestrian graph can contain an end of arbitrary (including infinite) degree. (8.1)

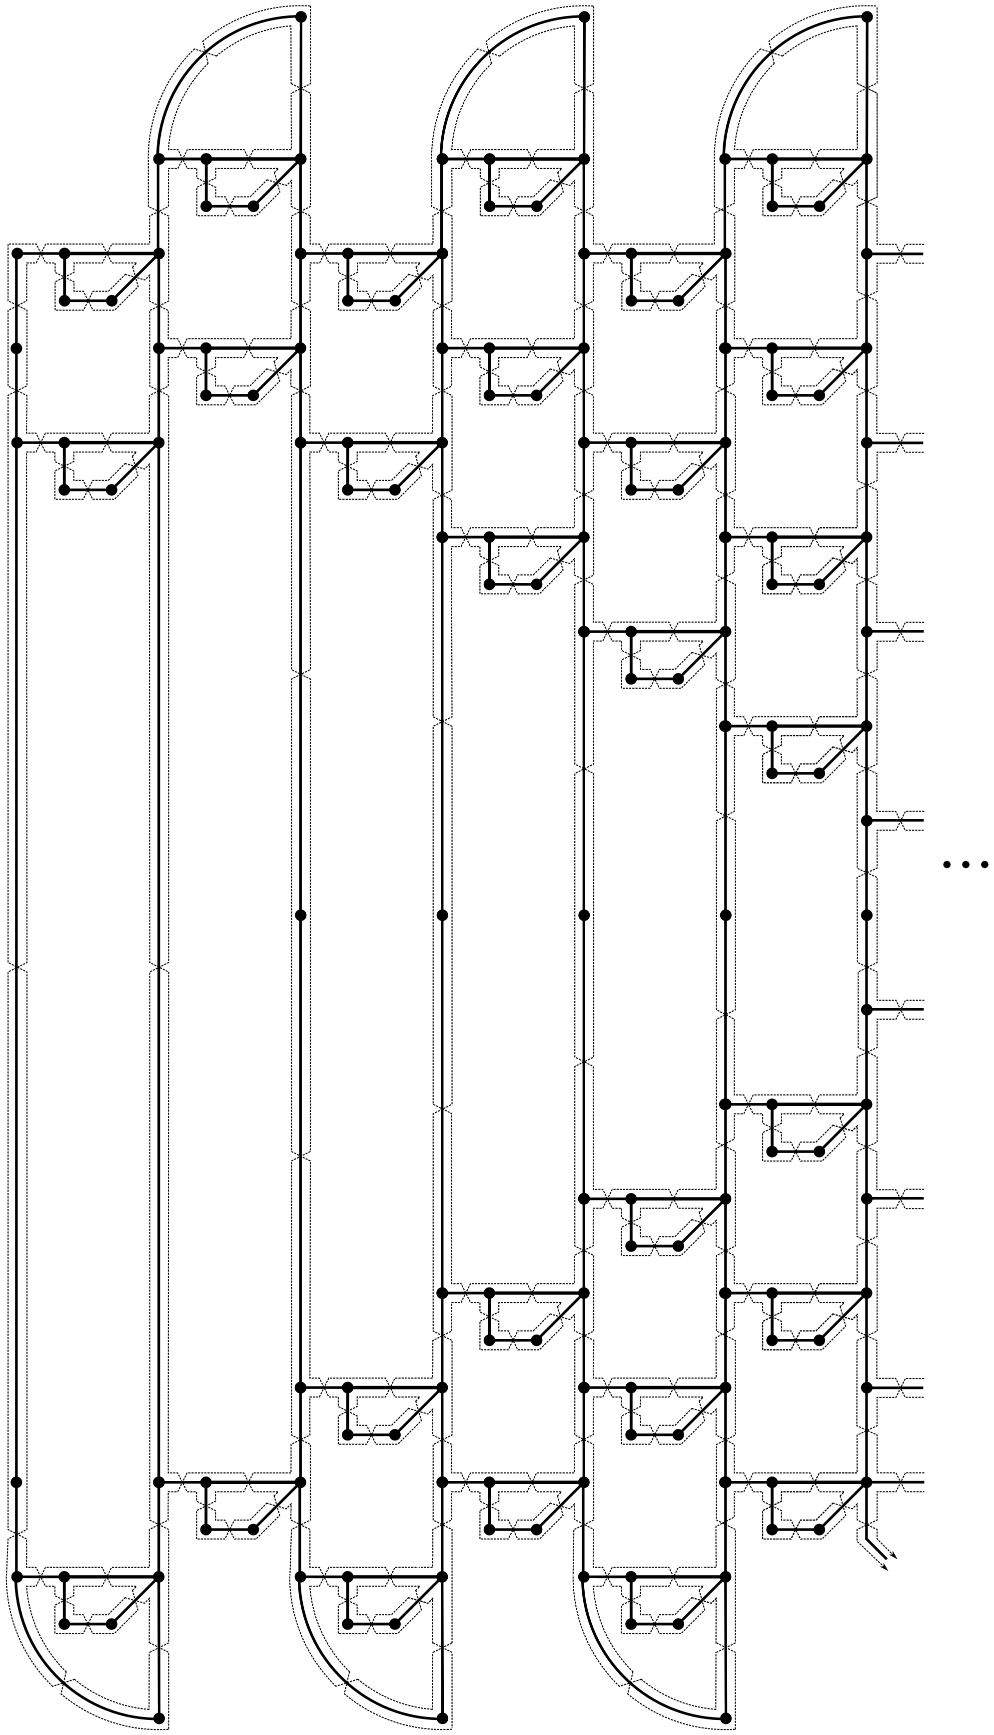

Figure 8.6: A locally finite pedestrian graph can have a thick end.

So is it possible, on the other hand, to control the *number* of ends in a locally finite pedestrian graph? Unfortunately, again, the answer is no. Figure 8.7 shows how we can pick any number of our illustrated locally finite pedestrian graphs with a certain end degree (in this case it is the graph from Figure 8.4 with degree one), connect them in the illustrated way with edges and obtain again a locally finite pedestrian graph. This construction also works for the other pedestrian graphs we have presented; the left-right strings 'connect up' properly as long as we add zero or an even number of edges to the new path that joins the different copies of the graphs. The new left-right strings thus obtained will then still form a single left-right tour.

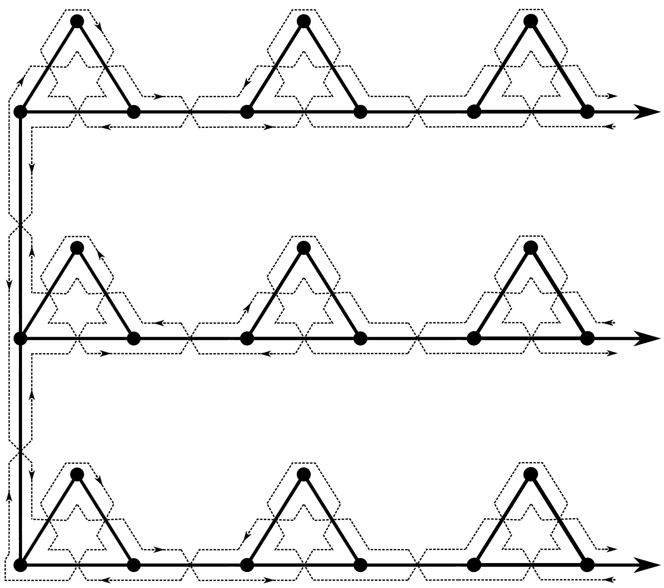

Figure 8.7: Constructing a pedestrian graph with any number of ends.

Thus, we see that:

> *A locally finite pedestrian graph can contain an arbitrary (including infinite) number of ends, each of which can be of arbitrary (including infinite) degree.* (8.2)

In Chapter 2 we have introduced the degree of an end as the maximum number of disjoint rays it contains. This is sometimes, more precisely, referred to as the *vertex-degree*, as opposed to the *edge-degree* of an end, which is the maximum number of *edge-disjoint* rays it contains. Our observations, however, hold for both the vertex- and the edge-degree since the number of edge-disjoint rays is the same as the number of (vertex-) disjoint rays in our examples.

In hope of finding some other characterization of locally finite pedestrian graphs, we have studied their subgraphs. Figure 8.7 suggests that some subgraphs of locally finite pedestrian graphs should again be pedestrian, but we have not succeeded at finding a good characterization of this type. Consider the example in Figure 8.8:

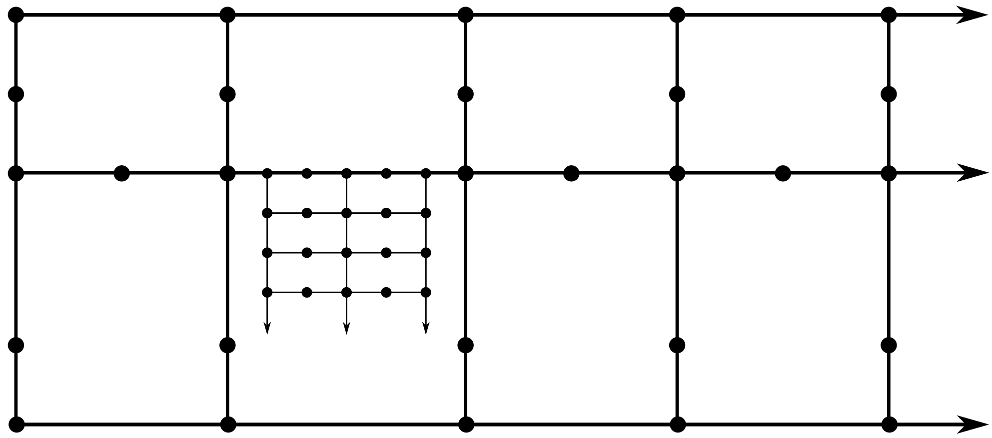

Figure 8.8: A pedestrian graph can have induced subgraphs that are not pedestrian.

The graph in Figure 8.8 is locally finite and pedestrian (in this case we do without the LRT since it would blur the structure of the illustrated graph). And yet, it has the following finite subgraph:

The graph in Figure 8.9 is an induced, 2-connected subgraph of the locally finite pedestrian graph from Figure 8.8. Furthermore, it is a subgraph whose deletion yields a collection of infinite components. Yet this graph is not pedestrian, as the bicycle indicated in Figure 8.9 shows.

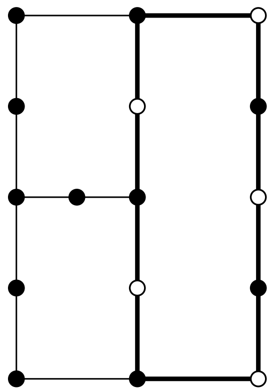

Figure 8.9: An induced subgraph of the graph in Figure 8.8 that contains a bicycle.

Bibliography

[1] D. Archdeacon, C.P. Bonnington, and C.H.C. Little, *An algebraic characterization of planar graphs*, J. Graph Theory **19** (1995), pp. 237–250.

[2] H. Bruhn, *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), pp. 235–256.

[3] H. Bruhn and R. Diestel, *Duality in infinite graphs*, Comb., Probab. Comput. **15** (2006), pp. 75–90.

[4] H. Bruhn, R. Diestel, and M. Stein, *Cycle-cocycle partitions and faithful cycle covers for locally finite graphs*, J. Graph Theory **50** (2005), pp. 150–161.

[5] H. Bruhn, S. Kosuch, and M. Win Myint, *Bicycles and left-right tours in locally finite graphs*, Europ. J. Comb. **30** (2009), pp. 356–371.

[6] H. Bruhn and M. Stein, *MacLane's planarity criterion for locally finite graphs*, J. Combin. Theory (Series B) **96** (2006), pp. 225–239.

[7] H. Bruhn and M. Stein, *On end degrees and infinite cycles in locally finite graphs*, Combinatorica **27 (3)** (2007), pp. 269–291.

[8] K. Casteels and R.B. Richter, *The bond and cycle spaces of an infinite graph*, J. Graph Theory **59** (2008), pp. 126–176.

[9] W.-K. Chen, *On vector spaces associated with a graph*, SIAM J. Appl. Math. **20** (1971), pp. 526–529.

[10] R. Diestel, *The cycle space of an infinite graph*, Comb., Probab. Comput. **14** (2005), pp. 59–79.

[11] R. Diestel, *Graph Theory* (3rd edition), Springer-Verlag, 2005.

[12] R. Diestel and D. Kühn, *On infinite cycles I*, Combinatorica **24** (2004), pp. 69–89.

[13] R. Diestel and D. Kühn, *On infinite cycles II*, Combinatorica **24** (2004), pp. 91–116.

[14] R. Diestel and D. Kühn, *Topological paths, cycles and spanning trees in infinite graphs*, Europ. J. Combinatorics **25** (2004), pp. 835–862.

[15] G.A. Dirac and S. Schuster, *A theorem of Kuratowski*, Nederl. Akad. Wetensch. Proc. Ser. A **57** (1954), pp. 343–348.

[16] C. Godsil and G. Royle, *Algebraic Graph Theory*, Springer-Verlag, 2001.

[17] H. Jung, *Wurzelbäume und unendliche Wege in Graphen*, Math. Nachr. **41** (1969), pp. 1–22.

[18] J. Keir and R.B. Richter, *Walks through every edge exactly twice II*, J. Graph Theory **21** (1996), pp. 301–309.

[19] S. Lins, R.B. Richter, and H. Shank, *The Gauss code problem off the plane*, Aequationes Math. **33** (1987), pp. 81–95.

[20] L. Lovász, *Combinatorial Problems and Exercises* (2nd edition), North-Holland, Amsterdam, 1993.

[21] S. MacLane, *A combinatorial condition for planar graphs*, Fund. Math. **28** (1937), pp. 22–32.

[22] R.B. Richter and H. Shank, *Left-right paths off the plane*, Congr. Numerantium **45** (1984), pp. 115–128.

[23] R.B. Richter and C. Thomassen, 3-*connected planar spaces uniquely embed in the sphere*, Trans. Am. Math. Soc. **354** (2002), pp. 4585–4595.

[24] P. Rosenstiehl and R.C. Read, *On the principal edge tripartition of a graph*, Ann. Disc. Math. **3** (1978), pp. 195–226.

[25] H. Shank, *The theory of left-right paths*, in Combinatorial Mathematics III, Lect. Notes Math. 452, Springer-Verlag, 1975, pp. 42–54.

[26] A. Vella and R.B. Richter, *Cycle spaces in topological spaces*, J. Graph Theory **59** (2008), pp. 115–144.