Bicycles and left-right tours in locally finite graphs

Henning Bruhn   Stefanie Kosuch   Melanie Win Myint

Abstract

We extend three results involving bicycles and left-right tours to infinite, locally finite graphs: Read and Rosenstiehl's tripartition theorem, Shank's theorem that the residues of left-right tours generate the bicycle space and the planarity criterion of Archdeacon, Bonnington and Little. In order to achieve this it is necessary to allow infinite cycles as defined by Diestel and Kühn.

1 Introduction

The set of edge sets in a graph G together with symmetric difference as addition forms a $\mathbb{Z}_2$-vector space, the *edge space* $\mathcal{E}(G)$. Two important subspaces of $\mathcal{E}(G)$ are the *cycle space* $\mathcal{C}(G)$, the set of all sums of (edge sets of) cycles, and the *cut space* $\mathcal{C}^*(G)$, the set of all cuts. Although cuts and cycle space elements are, in a sense, orthogonal to each other, it is possible for an edge set to be an element of both, $\mathcal{C}(G)$ and $\mathcal{C}^*(G)$. Such an edge set is called a *bicycle* and the space $\mathcal{B}(G) := \mathcal{C}(G) \cap \mathcal{C}^*(G)$ is the *bicycle space*. See Figure 1 for an example.

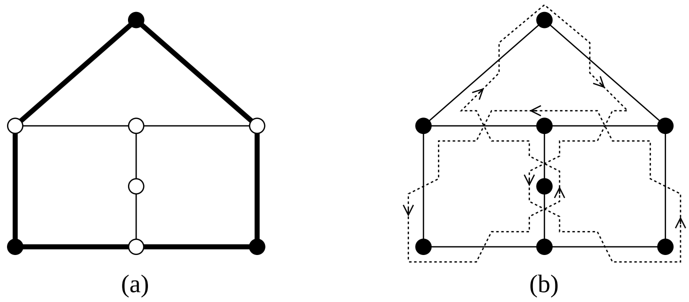

Figure 1: (a) a bicycle; (b) a left-right tour

In finite graphs, bicycles have been widely studied and a number of fundamental results involving bicycles are known. The aim of this work is to extend three of these to an important class of infinite graphs, namely to *locally finite* graphs, i.e. to graphs in which every vertex has finite degree.

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

Theorem 1 (Read and Rosenstiehl [15])**.** *Let e be an edge in a finite graph G. Then exactly one of the following holds:*

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

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

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

With the naive definition of $\mathcal{C}(G)$, in which every element of the cycle space is necessarily finite, Theorem 1 cannot be expected to carry over to locally finite graphs. The double ladder, depicted in Figure 2, constitutes an obvious counterexample: no finite bicycle contains the edge e, yet there is neither a finite Y nor a finite Z as in (ii) or (iii) of the theorem.

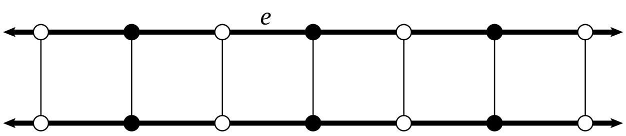

Figure 2: There is no finite B, Y or Z as in Theorem 1 for e

The almost self-evident solution is that infinite graphs demand infinite cycles. Indeed, lack of infinite cycles seems to be the reason why most properties of the cycle space fail in infinite graphs; see Diestel [5] for a number of examples. To remedy this, Diestel and Kühn [7, 8] 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. (Ends are equivalence classes of rays; formal definitions follow in the next section.) This definition has proved 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. We also refer to a more general approach pursued by Vella and Richter [17], that covers other compactifications of infinite graphs as well.

Coming back to our counterexample to Theorem 1, we see that the set of bold edges in the double ladder form an infinite cycle. (The two double rays together with the end to the left and the one to the right are homeomorphic to the unit circle.) Since this edge set is also a cut, we have found an infinite bicycle containing e, and thus the counterexample ceases to be one. More generally, we will prove in Sections 3 and 4 that the tripartition theorem becomes true for locally finite graphs once infinite cycles, as defined by Diestel and Kühn, are admitted.

In Sections 5 and 6 we will be concerned with plane graphs. In plane graphs, there is an easy way to find bicycles. Starting with any edge uv, we traverse uv from u to v, and then choose the leftmost edge at v, follow it along, then turn right, again turn left at the next vertex, and we continue alternating between left and right turns until we reach uv again. There we stop, provided we are about to traverse uv again from u to v and provided our turn at v would, again, be a left turn. The closed walk produced in this way is called a *left-right tour*. Its *residue*, the set of edges traversed exactly once, forms a bicycle; see Figure 1.

Shank [16] observed that left-right tours not only yield bicycles but that they, moreover, determine already all bicycles in the graph:

Theorem 2 (Shank [16]). *In a finite plane graph the residues of the left-right tours generate the bicycle space.*

This is the second of the theorems we shall extend to locally finite graphs. See also Richter and Shank [13] and Lins, Richter and Shank [11].

The third and final result we shall treat, in Section 7, is a planarity criterion that involves left-right tours and bicycles in a sophisticated way. For finite graphs this is due to Archdeacon, Bonnington and Little [1].

2 Definitions and preliminaries

All our graphs are simple and undirected, unless otherwise noted. In general, we follow the notation of [6], which also provides more background on the topological cycle space.

Let G be a locally finite graph. A *ray* is a one-way infinite path, and a *double ray* is a two-way infinite path. We say that two rays R, S are *equivalent* if there are infinitely many disjoint R–S paths. The equivalence classes are called the *ends of* G. As an example, the double ladder in Figure 3 has two ends, one to the left and one to the right. By contrast, the 3-regular tree has uncountably many ends.

We define on G, viewed as a 1-complex, together with its ends a topology, and denote the resulting topological space by $|G|$. The space $|G|$ is sometimes called the *Freudenthal compactification of* G. On G, the space carries the topology of a 1-complex, so every edge is homeomorphic to the unit interval, and a basic open neighbourhood of a vertex consists of half-open intervals, one for each edge incident with v. So, let us now define the basic open neighbourhoods for an end ω of G. Let S be a finite vertex set, and denote by $C(S,\omega)$ the component of $G-S$ that contains a ray in ω; then $C(S,\omega)$ contains a subray for every ray in ω. We define $\hat{C}(S,\omega)$ to be the union of $C(S,\omega)$ together with all interior points of edges between $C(S,\omega)$ and S, and all the ends that have a ray in $C(S,\omega)$. The sets $\hat{C}(S,\omega)$ for all finite sets $S\subseteq V(G)$ form a neighbourhood basis for ω. It can be shown that if G is connected, then $|G|$ is compact.

The image of a continuous mapping $[0,1]\to|G|$ is called a *topological path.* A *circle of* $|G|$ is a homeomorphic image C in $|G|$ of the unit circle; the subgraph $C\cap G$ is called a *cycle* and its edge set a *circuit.* A cycle may be finite or infinite; in the latter case it is the disjoint union of double rays.

The set of all subsets of $E(G)$ is the *edge space of* G and denoted by $\mathcal{E}(G)$. As noted in Section 1, together with the symmetric difference as addition, $\mathcal{E}(G)$ is a $\mathbb{Z}_2$-vector space. In order to define the topological cycle space of Diestel and Kühn we need to allow certain infinite sums as well. For this, we call a family $\mathcal{T}$ of edge sets *thin* if no edge appears in infinitely many of its members. The *sum* $\sum_{F\in\mathcal{F}}F$ is defined to be the set of edges that appear in exactly an odd number of members of $\mathcal{F}$. Whenever we take a sum over an (infinite) family it is tacitly assumed to be thin.

Now, we call the set of all (thin) sums of circuits the *topological cycle space* $\mathcal{C}(G)$ of G. If G is finite, it coincides with the usual cycle space. We will need two key properties of the topological cycle space:

Theorem 3 (Diestel and Kühn [7])**.** *Every element of the cycle space of a locally finite graph is the (edge-)disjoint union of circuits.*

An edge set F is called a *cut* if $F=\emptyset$ or if there is a set $U\subseteq V(G)$ so that each edge in F has precisely one endvertex in U and one outside U.

Theorem 4 (Diestel and Kühn [7])**.** *Let F be a set of edges in a locally finite graph G. Then F is an element of the cycle space if and only if it meets every finite cut in an even number of edges.*

In the *cut space* $\mathcal{C}^*(G)$, the set of all cuts, a result that is analogous to Theorem 4 holds; see the next lemma. A proof of this easy result can, for instance, be found in [2].

Lemma 5. *Let F be a set of edges in a graph G. Then F is a cut if and only if it meets every finite circuit in an even number of edges.*

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

In Sections 5 and 7 we will be concerned with infinite plane graphs. The usual drawings seem rather insufficient for infinite graphs. Indeed, several of the expected properties may fail. For instance, in a 2-connected graph the face boundaries do not need to be cycles. Moreover, they might even contain only half an edge (for instance, in the drawing there might be vertices converging against an interior point of an edge) or no edges at all. All these problems are overcome when, instead of G, the space $|G|$ is embedded in the sphere. Fortunately, this is not a restriction at all:

Theorem 6 (Richter and Thomassen [14])**.** *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.

3 The tripartition theorem

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

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

[1] There is a certain inconsistency here. Following Diestel [6], 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.

Proof of Theorem 1. Assume that there is no bicycle containing e. Thus

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

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

The first problem we encounter when we apply this proof to infinite graphs concerns the definition of the scalar product: What should the value of $X * Y$ be if the edge sets X, Y have infinite intersection? Fortunately, we will be able to circumvent this issue by only using the scalar product for $X, Y \in \mathcal{E}(G)$ with $|X \cap Y| < \infty$. A proper concept for orthogonal spaces appears to be more difficult, as however defined they seem to lose a number of their usual properties. For this reason, we will make do without them in infinite graphs. We remark that, these problems notwithstanding, Casteels and Richter [4] introduce orthogonal spaces in infinite graphs that still retain many of the usual properties.

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

Theorem 7. *Let e be an edge of a locally finite graph G. Then either*

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

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

but not both.

The reader will have noticed that the theorem only divides the edges into two classes rather than three. We will address this at the end of the section. The proof uses Kőnig's Infinity Lemma, a standard tool in infinite graph theory. For a proof we refer the reader to [6].

Lemma 8 (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 neighbour in W_{n-1}. Then H contains a ray $v_1 v_2 \ldots$ with $v_n \in W_n$ for all n.*

Proof of Theorem 7. We may assume G to be connected and therefore countable. For each $n \in \mathbb{N}$ denote by S_n the set of the first $n+1$ vertices in some fixed enumeration of the vertices of G that starts with the endvertices of e. Define G_n to be the graph $G[S_n]$ together with the edges in $E(S_n, V(G) \setminus S_n)$ and their incident vertices. Let $\widetilde{G}_n$ be the minor of G obtained by contracting the components of $G - S_n$ (where we keep parallel edges but delete loops). Note that $e \in E(G_n) = E(\widetilde{G}_n)$. Put $W_n := \{B \in \mathcal{C}^*(G_n) \cap \mathcal{C}(\widetilde{G}_n) : e \in B\}$.

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

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

and hence (ii) holds.

Second, assume $W_n \neq \emptyset$ for all n. It is not hard to check that for each $K \in \mathcal{C}^*(G_{n+1})$ it holds that $K \cap E(G_n) \in \mathcal{C}^*(G_n)$, and that for each $Z \in \mathcal{C}(\widetilde{G}_{n+1})$ the restriction $Z \cap E(\widetilde{G}_n)$ lies in $\mathcal{C}(\widetilde{G}_n)$. It follows that $B \in W_{n+1}$ implies $B \cap E(G_n) \in W_n$. We define a graph on $\bigcup_{n=1}^{\infty} W_n$ such that $B \in W_{n+1}$ is adjacent to $B' \in W_n$ if and only if $B \cap E(G_n) = B'$. Thus, the conditions for Lemma 8 are satisfied, and we obtain for each $n \in \mathbb{N}$ a $B_n \in W_n$ so that $B_{n+1} \cap E(G_n) = B_n$ for all n. Clearly, $B := \bigcup_{n \in \mathbb{N}} B_n$ contains e.

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

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

where the last equality follows since $B_N \in \mathcal{C}(\widetilde{G}_N)$. As F was arbitrary, Theorem 4 implies that $B \in \mathcal{C}(G)$. In a similar way, but using Lemma 5 in G_N instead of Theorem 4 in $\widetilde{G}_N$, we see that $B \in \mathcal{C}^*(G)$. Therefore, $B \in \mathcal{B}(G)$ and (i) holds.

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

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

which gives a contradiction. □

Casteels and Richter [4] independently proved a complementary result:

Theorem 9 (Casteels and Richter [4])**.** *Let e be an edge of a locally finite graph G. Then either*

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

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

but not both.

It should be noted that Casteels and Richter in fact prove a more general result of which Theorem 9 is but a consequence.

Theorems 7 and 9 look tantalisingly similar. The next lemma sheds some light on their relation.

Lemma 10. *Let G be a locally finite graph. If for an edge e of G two of the following conditions hold, then the third one is satisfied, too:*

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

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

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

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

The lemma is reminiscent of a theorem by Richter and Shank [13] 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 2 we have already seen that e lies in an infinite bicycle, while in Figure 3 we witness the other two cases.

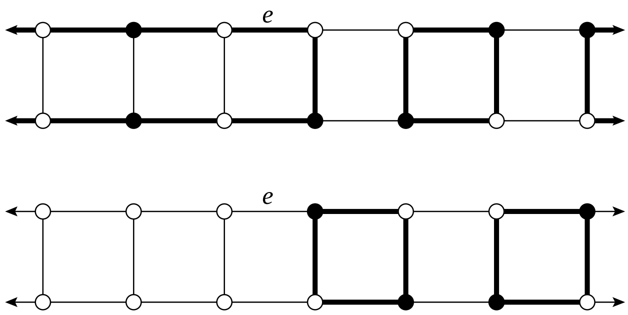

Figure 3: (i), (ii) in Lemma 10 hold for e

Proof of Lemma 10. First, assume (iii) and one of (i),(ii) to hold. Thus, there exist a $B \in \mathcal{B}(G)$ with $e \in B$ and an $X \in \mathcal{C}(G)$ so that $X + e \in \mathcal{C}^*(G)$. Since $X + B \in \mathcal{C}(G)$, we have that if X is as in (i), then $X + B$ satisfies (ii) and if, on the other hand, X is as in (ii), then $X + B$ satisfies (i).

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

For the second part of the lemma, assume that (i)-(iii) hold for e, and let $e \in B \in \mathcal{B}(G)$. By (the trivial part of) Theorem 9, it follows that B cannot be finite. On the other hand, Y and Z as in (i) resp. (ii) need to be infinite sets, too, since otherwise this would give a contradiction to Theorem 7. □

Read and Rosenstiehl's theorem partitions the edges of a finite graph into three classes. So far, our theorem yields only two classes. So, let us refine Theorem 7. For this, we say that an edge e in a locally finite graph G is of *cut-type* if there is a finite cut K containing e so that $K \setminus \{e\} \in \mathcal{C}(G)$. We say that e is of *flow-type* if there is a finite element Z of the cycle space with $e \in Z$ and $Z \setminus \{e\} \in \mathcal{C}^*(G)$. Then, the following immediate corollary of Lemma 10 turns Theorem 7 into a true tripartition theorem:

Corollary 11. *No edge in a locally finite graph can be of cut-type and of flow-type at the same time.*

We should point out that to denote by $\mathcal{C}^*(G)$ the set of all cuts is possibly a bit misleading as it might give the impression that it is the dual space of $\mathcal{C}(G)$. That, however, is not the case. Rather, Theorem 4 shows that, at least in some sense, $\mathcal{C}(G)$ and $\mathcal{C}^*_{\text{fin}}(G)$ are dual to each other. On the other hand, the dual space of $\mathcal{C}^*(G)$ is $\mathcal{C}_{\text{inf}}(G)$, see for instance [2].

In this respect, our bicycle space $\mathcal{B}(G)$ is situated between these two dualities. Examples as the graph in Figure 2 indicate that this is nevertheless justified since in order to make the tripartition theorem work in infinite graphs, whether it is in the form of Theorem 7 or in the form of Theorem 9, we need both spaces, $\mathcal{C}(G)$ and $\mathcal{C}^*(G)$.

4 Principal cuts

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

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

We need the following lemma, which (stated for finite graphs but with exactly the same proof) appears in Read and Rosenstiehl [15]. (We note that the lemma remains true in non-pedestrian graphs; Z_e (resp. Z_f) is then simply any principal flow through e (resp. f), as there is no longer a unique one. And similarly for K_e, K_f.)

Lemma 12. *Let e and f be edges in a locally finite pedestrian graph G. Then:*

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

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

Proof. To prove (i) consider

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

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

Proposition 13. *In a locally finite pedestrian graph G both of the families $(Z_e)_{e \in E(G)}$ and $(K_e)_{e \in E(G)}$ are thin.*

Proof. Suppose there is an edge e lying in infinitely many Z_f. Since G is a pedestrian graph, e is of flow- or cut-type and Z_e is therefore defined. Thus Lemma 12 implies that $f \in Z_e$ for all these infinitely many f, contradicting that Z_e is finite. Thus $(Z_e)_{e \in E(G)}$ is thin. The proof for the principal cuts is the same. □

For an edge e to be of flow- or of cut-type we have required that there is a *finite* $Z \in \mathcal{C}(G)$ with $Z + e \in \mathcal{C}^*(G)$. In the light of Theorem 9 one could also quite reasonably relax this, and say that an edge is of flow- or cut-type if there is any such Z, finite or infinite. A pedestrian graph, then, would be one without any *finite* bicycles, since in precisely this case all edges are of flow- or cut-type.

There are several problems with this definition. We have already seen (Figures 2 and 3) that this would not give a proper tripartition. Furthermore, principal cuts in a pedestrian graph would not necessarily be unique and their family may not be thin. For instance, the cuts in the lower graph in Figure 3 would form a non-thin family of principal cuts.

The following corollary lists verbatim extensions of some basic properties of principal flows and cuts. Their proofs for finite graphs (substantially) use the finiteness only in one point, namely that it is allowed to take arbitrary sums of principal cuts. While, clearly, this is never an issue in finite graphs, such sums may be infinite in infinite graphs and then need to be thin in order to be well-defined. But this is exactly what Proposition 13 asserts.

Corollary 14. *Let G be a locally finite pedestrian graph. Then*

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

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

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

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

Proof. (i) and (ii) can be found in Read and Rosenstiehl [15] and (iii) and (iv) in Godsil and Royle [9]. □

5 Left-right tours

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

The first requirement is fairly simple to guarantee. Just as with left-right tours in finite graphs, we start a walk at an arbitrary edge and then alternately turn left and right. If we reach our starting edge again in this way, we have found a finite left-right tour. Otherwise, we prolong our walk in the other direction from our starting edge, again taking left and right turns. The resulting walk, which we call a *left-right string*, will be two-way infinite; two examples can be seen in Figure 5. In general, the two ends of a left-right string will not be identical, and the walk will therefore not be closed. So to achieve that we do not get stuck in an end, it will be necessary to glue together several left-right strings at ends. In this way we shall obtain a topological tour in $|G|$.

Let us start with left-right strings. To define these properly we shall first need to describe what it means to do a left turn followed by a right turn. We follow the treatment of [10]. Let G be a locally finite graph, and let $|G|$ be embedded in the sphere S. Recall that, by Theorem 6, every locally finite planar graph has such an embedding. The interior of an edge of G is homeomorphic to the open unit interval $(0, 1)$. For each edge e, we fix a homeomorphism. If η_1 denotes the image of the restriction of this homeomorphism to $(0, \frac{1}{2})$ and η_2 is the image of the restriction to $(\frac{1}{2}, 1)$ then η_1, η_2 are the *halves of* e. We use the notation $\overline{\eta_1} = \eta_2$ and $\overline{\eta_2} = \eta_1$ to switch back and forth between the two halves of an edge. Furthermore, we fix for e two open, disjoint and connected subsets, σ_1 and σ_2, of $S \setminus |G|$ each of which has e in its boundary. These are the *sides of* e, and as for the halves, we put $\overline{\sigma_1} = \sigma_2$ and $\overline{\sigma_2} = \sigma_1$. A triple (e, η, σ), where $e \in E(G)$, η is a half of e and σ is a side of e, is called a *corner of* $|G|$. We say that $c = (e, \eta, \sigma)$ is a *corner at* e, and it is a *corner at* $v \in V(G)$ if the boundary $\partial\eta$ contains v. Clearly, for each edge e there are four corners at e.

For each $v \in V(G)$ choose an open disc D around v, so that each half of an edge at v intersects ∂D in exactly one point. Then ∂D defines in a natural way a rotation of the halves. We say that two corners (e, η, σ), (e', η', σ') at v are *matched* if η and η' appear consecutively in the local rotation at v, and if the connected component K of $\sigma \cap D$ with $\eta \cap D \subseteq \partial K$ and the connected component K' of $\sigma' \cap D$ with $\eta' \cap D \subseteq \partial K'$ are contained in the same connected component of $D \setminus |G|$. It can be seen that this definition is independent of the actual choice of D. See Figure 4 for an illustration.

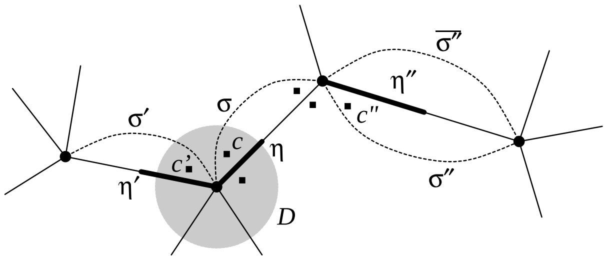

Figure 4: We think of a corner $c = (e, \eta, \sigma)$ at $v \in V(G)$ as a point close to v and η, and lying in σ. The corners c and $c' = (e', \eta', \sigma')$ are matched; the corners c and c'' describe a left-right step.

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

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

(ii) no corner appears twice in W.

We call such a sequence W a *left-right walk*, which is justified by the fact that the edges $\ldots e_{-1} e_0 e_1 \ldots$ do indeed form a walk. Moreover, we will sometimes pretend that a left-right walk is in fact a walk, i.e. a sequence of vertices and edges, rather than a sequence of corners. The corners c and c'' in Figure 4 describe a left-right step as in (i).

We say that S is a *left-right string* (LRS for short) if it is a maximal left-right walk. It is not hard to check that if $S = \ldots (e_{-1}, \eta_{-1}, \sigma_{-1}), (e_0, \eta_0, \sigma_0), (e_1, \eta_1, \sigma_1) \ldots$ then $S' := \ldots (e_1, \overline{\eta_1}, \overline{\sigma_1}), (e_0, \overline{\eta_0}, \overline{\sigma_0}), (e_{-1}, \overline{\eta_{-1}}, \overline{\sigma_{-1}}) \ldots$ is an LRS, too. Clearly, the walks S and S' traverse the same edges, but in opposite directions. Although we will sometimes view S as an oriented walk, we will, in general, not distinguish between S and S' and consider them to be identical. This slight abuse of notation ensures that every edge is covered exactly twice by LRS; see the next lemma. Figure 5 gives an example of two different LRS in the double ladder.

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

Lemma 15. *For a locally finite graph G, let $|G|$ be embedded in the sphere. Then:*

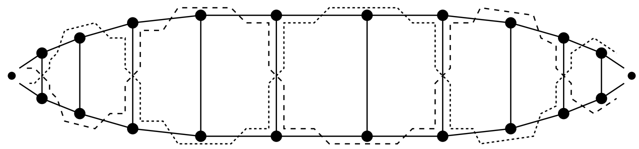

Figure 5: Two LRS in the double ladder

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

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

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

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

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

Now we can finally extend the definition of left-right tours to infinite graphs. Assume that $|G|$ is embedded in the sphere. Our aim is to give a definition so that an LRT consists of a number of LRS that are glued together at ends so as to constitute a tour in $|G|$. An example would be the two LRS shown in Figure 5 together with the two ends of the double ladder.

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

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

Lemma 16 (Shank [16])**.** *If G is a finite plane graph, then the residue of a left-right tour is a bicycle.*

Lemma 16 is proved with the help of plane dual graphs. While abstract dual graphs have been defined in [2], a suitable theory of plane dual graphs that involves infinite cycles has yet to be formulated. This is probably not overly difficult but checking the sometimes tedious geometrical details would take too much space and effort here. Rather, with the help of the next lemma, we will circumvent this obstacle by reducing the problem to finite graphs.

Lemma 17. *For a locally finite graph G, let $|G|$ be embedded in the sphere. Let H be a finite plane subgraph, and let $L_1, \ldots, L_k$ be a set of LRTs of G so that no*

LRS of G lies in more than one L_i. Then there exist a finite plane supergraph H' of H and a set $L'_1, \dots, L'_k$ of LRTs of H' so that for all $i = 1, \dots, k$, the LRT L_i traverses precisely the edges $e_1, \dots, e_n$ of H and in this order if and only if L'_i does.

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

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

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

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

$$\textit{if } \{c, c'\} \in \mathcal{P} \textit{ then one of } c, c' \textit{ is outgoing and one ingoing.} \tag{1}$$

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

Define for each $c \in \mathcal{M}$ a left-right walk $K^0(c) := (c)$, i.e. $K^0(c)$ is a walk of length 1, which traverses an edge in F. To simplify the construction in the next steps we will, with the help of a suitable homeomorphism, identify D with $(0, 3) \times (0, 1) \subseteq \mathbb{R}^2$, where all the vertices in X are assumed to lie in the open segment $\{0\} \times (0, 1)$; see Figure 6.

Next, we pick $m := |\mathcal{M}|$ distinct points $x^1_1, \dots, x^1_m$ in $\{1\} \times (0, 1)$, where we choose the labelling so that x^1_j has a smaller y-coordinate than x^1_{j+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^1_j to a vertex w in X so that w receives one edge if its incident edge in F is only traversed once by $L_1, \dots, L_k$; otherwise (when the edge is used twice) we make w adjacent to two of the x_j^1. Clearly, in the resulting plane supergraph H_1 of H_0 each vertex in $x_1^1, \dots, x_m^1$ has degree 1.

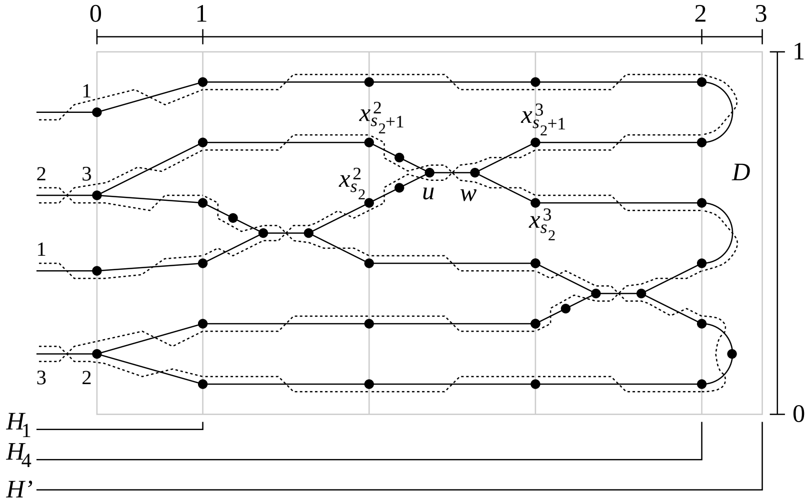

Figure 6: The construction of the H_i (not to scale)—corners with the same number are supposed to be paired.

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

We will construct supergraphs H_i of H_1 with corresponding left-right walks $K^i(c) \supseteq K^1(c)$, $c \in \mathcal{M}$. More precisely, we will construct finitely many nested plane supergraphs with $H_1 \subset H_2 \subset \dots \subset H_{t+1}$, where $H_i \setminus H_{i-1}$ is entirely drawn in $(a, b] \times (0, 1)$ for some $1 \le a < b < 3$ (we will determine the respective a and b in a moment). The intersection of H_i with $\{b\} \times (0, 1)$ will consist of m vertices; in the order we encounter them on $\{b\} \times (0, 1)$ going from $(b, 0)$ to $(b, 1)$ these will be denoted by $x_1^i, \dots, x_m^i$. For each $j = 1, \dots, m$ there will then be a unique corner $p_j^i \in \mathcal{M}$ so that the left-right walk $K^i(p_j^i)$ either starts or ends in x_j^i (and is otherwise disjoint from $x_1^i, \dots, x_m^i$).

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 $(p_1^1, \dots, p_m^1)$ into $(q_1, \dots, q_m)$ so that for each odd j in $\{1, \dots, m\}$ it holds that $\{q_j, q_{j+1}\} \in \mathcal{P}$.

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

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

Lemma 18. *For a locally finite graph G, let $|G|$ be embedded in the sphere. Then the residue of an LRT in G is an element of the bicycle space.*

Proof. Let F be a finite cut and L an LRT. As a tour, L passes an even number of times through F. Therefore, $|\bigtriangledown L \cap F|$ is even and it follows, by Theorem 4, that $\bigtriangledown L$ is an element of the cycle space.

To see that the residue $\bigtriangledown L$ is a cut, consider a finite cycle C. Lemma 17 (with $H = C$) yields a finite plane supergraph H' of C and an LRT L' of H' so that $\bigtriangledown L \cap E(C) = \bigtriangledown L' \cap E(C)$. As $\bigtriangledown L'$ is a cut in H' (by Lemma 16) and $C \subseteq H'$ a cycle, we have that $\bigtriangledown L' \cap E(C)$ is an even set. Since this implies that $\bigtriangledown L \cap E(C)$ is even, too, it follows from Lemma 5 that $\bigtriangledown L \in \mathcal{C}^*(G)$ and hence $\bigtriangledown L \in \mathcal{B}(G)$. □

6 LRTs generate the bicycle space

In this section we will prove the analogue of Theorem 2 for locally finite graphs.

Let G be a locally finite graph for which $|G|$ is embedded in the plane, and consider a bicycle B of G. Since the cuts of the form $E(v)$ generate the cut space, there is a vertex set X such that $B = \sum_{x \in X} E(x)$. On the other hand, B is also an element of the cycle space. As in finite graphs, $\mathcal{C}(G)$ is generated by the residues of the face boundaries (this is shown in [3]). Thus, there is a set F of face boundaries such that $B = \sum_{f \in F} \bigtriangledown f$. For each bicycle B assume such a pair X, F to be fixed. Following Richter and Shank [13], we say that an LRS S is of *type I* if there is a corner $c = (e, \eta, \sigma)$ in S for which the following statements are either both true or both false:

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

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

It is not hard to check that if for one corner in S either both of (i) and (ii) are true or are both false then this holds for every corner in S; see also Richter and Shank [13]. If S is not of type I, then S is of *type II*.

Lemma 19. *Let G be a locally finite plane graph, and let B be a bicycle. Then an edge e of G lies in B if and only if it lies in exactly one LRS of type I and in one LRS of type II with respect to B.*

Proof. The proof is identical to the one given for finite graphs in Richter and Shank [13]. □

An LRT L is called *B-uniform* if every two LRS contained in L are of the same type. In finite graphs, Lemma 19 is already enough to prove Theorem 2: we only need to sum up all LRS (which are identical to LRTs in finite graphs) of type I (or type II, for that matter). By contrast, in locally finite graphs, it is not even clear whether there is a single B-uniform LRT, let alone a set of B-uniform LRTs with the properties as in the last lemma. The next lemma asserts the existence of B-uniform LRTs.

Lemma 20. *Let G be a locally finite graph, let $|G|$ be embedded in the sphere, and let B be a bicycle of G. Then there exists a set $\mathcal{L}$ of B-uniform LRTs so that each LRS of G is contained in exactly one $L \in \mathcal{L}$.*

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

We construct from G another locally finite graph G' (which, in all likelihood, will not be planar). The vertex set of G' consists of vertices v_p, one for each vertex v of G and for each subwalk p of the form $p = evf$ in each S_i ($e, f \in E(G)$). Such a vertex $v_p \in V(G')$ is called a *clone of* v. The edge set of G' is comprised of two disjoint sets, E' and F'. The set F' contains one edge between each pair of clones v_p and v_q of the same vertex $v \in V(G)$; i.e. the clones of a vertex span a complete graph. Two clones u_p and v_q of distinct vertices $u, v \in V(G)$ are connected by an edge in E' if p and q are subwalks in the same LRS S_i and appear consecutively in S_i, i.e. if $S_i = \ldots e_{-1}v_{-1}e_0v_0e_1v_1e_2\ldots$ then $p = e_{j-1}v_{j-1}e_j$ and $q = e_jv_je_{j+1}$ (or the other way round) for some j. See Figure 7 for an illustration.

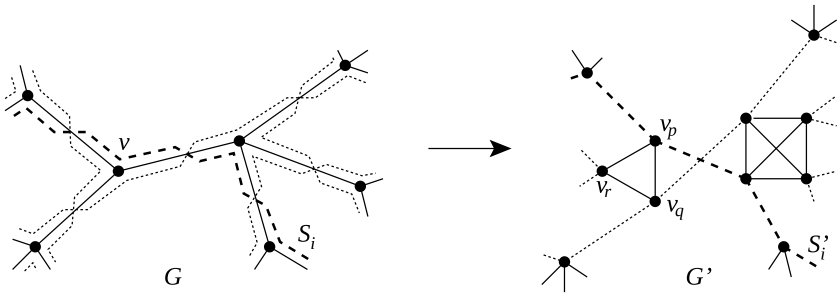

Figure 7: Construction of G' in the proof of Lemma 20; the edges in E' are dotted.

Let us define a mapping $\phi : V(G') \cup E(G') \to V(G) \cup E(G)$. For each $v \in V(G)$ we map all clones of v and all edges (in F') between two clones of v

to v. An edge $u_p v_q$ in E', where u_p is a clone of $u \in V(G)$ and v_q is a clone of $v \neq u$, is mapped to the edge uv of G. Clearly, this map is surjective.

We note, furthermore, that because of Lemma 15 (iii),

$$\textit{each } e \in E(G) \textit{ has exactly two preimages under } \phi\textit{, and these are in } E'. \quad (2)$$

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

(i) if $S'_i = \ldots e'_{-1}v'_{-1}e'_0v'_0e'_1v'_1e'_2\ldots$ 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 for (ii) and (iii) simply note that a clone v_p of a vertex $v \in V(G)$ is adjacent to exactly two vertices that are not clones of v.

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

$$\textit{both of } X_I := \bigcup_{S'\in\mathcal{X}_I} E(S') \textit{ and } X_{II} := \bigcup_{S'\in\mathcal{X}_{II}} E(S') \textit{ lie in } \mathcal{C}(G'). \quad (3)$$

To show that $X_I \in \mathcal{C}(G')$, consider a finite cut K' of G'; by Theorem 4, it suffices to prove that $|X_I \cap K'|$ is even.

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

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

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

Lemmas 15 (iii) and 19 imply that $\phi(K) \setminus B$ is traversed an even number of times by LRS of type I. Since B is an element of the cycle space, the set

$B \cap \phi(K)$ is even, by Theorem 4. Thus, Lemma 19 implies that also $B \cap \phi(K)$ is traversed an even number of times by LRS of type I. With (2) we get that $|X_I \cap K|$ is even. The proof for X_{II} is the same.

Next, we use Theorem 3 to decompose $X_I + X_{II}$ into a set $\mathcal{D}$ of (edge-)disjoint circuits. We observe that

for all i and $D \in \mathcal{D}$ it holds that if $E(S_i') \cap D \neq \emptyset$ then $E(S_i') \subseteq D$. Moreover, for each $D \in \mathcal{D}$, all the S_i with $E(S_i') \subseteq D$ are of the same type. (4)

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

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

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

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

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

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

Finally, since each $D \in \mathcal{D}$ is a circuit, by definition there exists a homeomorphism $\sigma_D : S^1 \to |G'|$ with image $\overline{D}$. By (b), the continuous mapping $\phi' \circ \sigma_D : S^1 \to |G|$ is locally injective at points $x \in S^1$ that are mapped to interior points of edges. Furthermore, (i) and (a) imply that each maximal subwalk in $\phi' \circ \sigma_D$ is an LRS, and that these are precisely those S_i for which $E(S_i') \subseteq D$. Therefore, each $\phi' \circ \sigma_D$ describes an LRT in $|G|$. By (4), each such LRT is B-uniform. We denote the set $\{\phi' \circ \sigma_D : D \in \mathcal{D}\}$ of LRTs by $\mathcal{L}$.

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

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

With Lemma 20 we can extend Theorem 2 to locally finite graphs using arguments of Richter and Shank [13]. Given a bicycle B, Lemma 20 yields a set $\mathcal{M}$ of LRTs, so that every LRS of type I appears in exactly one element of $\mathcal{M}$. Lemma 19 assures that $\sum_{M \in \mathcal{M}} \triangledown M = B$. On the other hand, Lemma 18 shows that all sums of residues of LRTs are elements of the bicycle space. In conclusion, we have proved:

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

In a finite graph, the set of LRTs is a double cover. In the double ladder, by contrast, we can construct LRTs by glueing together any two of the four LRS, which results in a set of six LRTs that cover all edges more than twice; see Figure 5. Moreover, while Lemma 20 asserts that there are double covers consisting of LRTs, in the case of the double ladder none of these are sufficient to generate the bicycle space. Indeed, consider a double cover $\mathcal{L}$ of LRTs for the double ladder. Pick an LRT of the double cover and observe that it traverses some edge e twice (in Figure 5 this is the case for every second rung). It is easy to check that every edge in the double ladder lies in a bicycle, and hence, no bicycle containing e can be expressed as the sum of residues of $L \in \mathcal{L}$.

7 The ABL planarity criterion

MacLane's well-known planarity criterion [12] characterises planar graphs in terms of the cycle space. MacLane observed that, in (finite) plane graphs, the set of facial walks is a double cover that generates the cycle space. Then he proved that, conversely, any double cover of closed walks with this property can be realised as a set of facial walks and is therefore a certificate for planarity.

The planarity criterion of Archdeacon, Bonnington and Little [1] works in a similar way with the difference that they list the essential properties of the left-right tours. These properties are rather more elaborate and necessitate a number of definitions, which we will give below. In this section it is our aim to show that the ABL criterion remains true in locally finite graphs.

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

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

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

We can now state the ABL criterion:

Theorem 22 (Archdeacon, Bonnington and 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 the ABL criterion can be found in Keir and Richter [10]. Theorem 22 extends to locally finite graphs:

Theorem 23. *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 6 we know that $|G|$ has an embedding in the sphere, and thus Lemma 20 yields (with, for instance, $B = \emptyset$) a set $\mathcal{L}$ of LRTs so that each LRS of G lies in exactly one element of $\mathcal{L}$. Hence, $\mathcal{L}$ is a double cover of G (Lemma 15 (iii)). Furthermore, Lemma 18 implies that $\bigtriangledown L$ 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 $\mathcal{L}$), $\bigtriangledown S$ is a cut as well. We show that $\bigtriangledown S$ meets every finite cycle C in an even number of edges, thereby proving $\bigtriangledown S$ to be a cut (Lemma 5).

If R is the set of rungs of $\mathcal{H}$, then we define H to be the plane subgraph of G consisting of C and all the edges in R together with their incident vertices. We apply Lemma 17 to H and the LRTs in $\mathcal{H}$, which yields a finite plane supergraph H' and a set $\mathcal{H}'$ of LRTs of H'. It is straightforward to see that $\mathcal{H}'$ is a ladder in H' with a side S' for which it holds that $\bigtriangledown S' \cap E(H) = \bigtriangledown S \cap E(H)$. Since $\bigtriangledown S'$ is a cut, by Theorem 22, the intersection $\bigtriangledown S' \cap E(C) = \bigtriangledown S \cap E(C)$ is even. This proves $\mathcal{L}$ to be a diagonal.

For the converse direction, let us now suppose that G has a diagonal $\mathcal{D}$ but also contains a subdivision X of $K_{3,3}$ or of K_5. Denote by H the (finite) induced subgraph of G on $V(X)$, and set $F := E(H, G - H)$, which is a finite cut. One by one, we delete the edges of F from G. We claim that after each edge deletion, the graph G still has a diagonal. For finite graphs, this is proved in Archdeacon, Bonnington and Little [1]. As their arguments remain still valid in locally finite graphs, we will not repeat them.

Once we have deleted all of F, the diagonal will split into two parts: into the set $\mathcal{D}'$ of those that are completely contained in H, and into those tours that are disjoint from H. Clearly, $\mathcal{D}'$ is then a diagonal of the finite non-planar graph H, which is impossible by Theorem 22. □

For pedestrian graphs, i.e. those graphs G for which $\mathcal{B}(G) = \{\emptyset\}$, Read and Rosenstiehl [15] gave a slightly simpler planarity criterion. Let a tour W traverse an edge $e = uv$ twice. If e is consistent, and traversed from u to v, say, then W decomposes into four topological subpaths uv, H_1, uv and H_2. We call each of H_1 and H_2 a *half of* W (with respect to e). If e is inconsistent, then W is equally comprised of four topological subpaths: namely of uv, H_1', vu and of H_2'. In this case we call the topological subpaths uvH_1' and vuH_2' *halves of* W.

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

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

Theorem 24 (Read and Rosenstiehl [15]). *A finite connected pedestrian graph is planar if and only if it has an algebraic diagonal.*

Theorem 25. *A locally finite connected pedestrian graph is planar if and only if it has an algebraic diagonal.*

Proof. Let G be a locally finite connected pedestrian graph. If G is planar, then $|G|$ can be embedded in the sphere (Theorem 6) and there is a family $\mathcal{L}$ of LRTs of G that forms a double cover (by Lemma 20). We already know (from the proof of Theorem 23) that $\mathcal{L}$ is a diagonal. If $\mathcal{L}$ has only a single member D, then D is an algebraic diagonal of G: since every half H of D is the side of a ladder, it follows that $\triangledown H$ is a cut.

So, assume that $\mathcal{L}$ has two members, and denote one of them by L. Since G is pedestrian, Lemma 18 implies $\triangledown L = \emptyset$. As G is connected there is therefore a vertex v which is met by L but also incident with edges not lying in L. Consider an edge e incident with v that lies in L. Let η be the half of e with $v \notin \partial\eta$, and let σ be a side of e. Since L traverses e twice (as $\triangledown L = \emptyset$), L (or, more precisely, the LRS lying in L) contains one corner of each of $\{(e,\eta,\sigma),(e,\overline{\eta},\overline{\sigma})\}$ and $\{(e,\eta,\overline{\sigma}),(e,\overline{\eta},\sigma)\}$. Let (e_1,η_1,σ_1) be the corner that is matched with $(e,\overline{\eta},\overline{\sigma})$, and let (e_2,η_2,σ_2) be the one matched with $(e,\overline{\eta},\sigma)$. By definition, if L contains (e,η,σ) then it also contains (e_1,η_1,σ_1). If, on the other hand, $(e,\overline{\eta},\overline{\sigma})$ lies in L, then $(e_1,\overline{\eta}_1,\overline{\sigma}_1)$ is a corner of L. In any case, e_1 is traversed by L. As, in a similar way, we see that e_2 lies in L as well, it follows that the predecessor and the successor of e in the local rotation at v both lie in L, and thus that all of $E(v)$ is covered by L, a contradiction to our assumption.

If, conversely, G has an algebraic diagonal D, then the set $\{D\}$ is a diagonal. Theorem 23 shows that G is planar. □

References

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

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

[3] H. Bruhn and M. Stein, *MacLane's planarity criterion for locally finite graphs*, To appear in *J. Combin. Theory (Series B)*.

[4] K. Casteels and B. Richter, *The cycle space of an infinite graph*, Preprint 2006.

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

[6] ———, *Graph theory* (3rd edition), Springer-Verlag, 2005.

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

[8] ———, *On infinite cycles II*, Combinatorica **24** (2004), 91–116.

[9] C. Godsil and G. Royle, *Algebraic graph theory*, Springer-Verlag, 2001.

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

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

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

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

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

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

[16] H. Shank, *The theory of left-right paths*, Combinatorial Mathematics III (A. Penfold Street and W.D. Wallis, eds.), Lecture notes in Math., vol. 452, Springer-Verlag, 1975, pp. 42–54.

[17] A. Vella and R.B. Richter, *Cycle spaces of topological spaces*, Preprint.

Version 02 Oct 2008

Henning Bruhn
`<hbruhn@gmx.net>`
Melanie Win Myint
`<mwinmyint@gmail.com>`
Mathematisches Seminar
Universität Hamburg
Bundesstraße 55
20146 Hamburg
Germany

Stefanie Kosuch
`<StefKosuch@web.de>`
Laboratoire de recherche en informatique
Université Paris-Sud
91405 Orsay Cedex
France